IBM System Storage DS Storage Manager Version 10

## *Command Line Interface and Script Commands Programming Guide*



GA32-0961-06

Note

Before using this information and the product it supports, read the information in "Notices" on page 667.

This edition applies to version 10 modification 35 of the IBM DS Storage Manager, and to all subsequent releases and modifications until otherwise indicated in new editions.

This edition replaces GA32-0961-05.

© Copyright IBM Corporation 2008, 2013. US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

## Contents

Figures.	•	•		•	•	•	•	•	•	•	•	•	. i	X
Tables .					•	•	•		•			•	. )	ci
About th	is d	loc	ume	ent									xi	ii
What's new DS3000, DS	7 in t 54000	:his ), D9	guid 55000	e ), aı	nd	the	Do	CS	Sei	ries	•	•	. X1	11
software ar	nd fii	mw	vare s	sup	por	ť							. xi	ii
Who shoul	d rea	nd tł	nis d	ocu	me	nt			•	•			. xi	V
Notices use	ed in	this	s doc	um	ent		•	•	•				. xi	V
Getting inf	orma	tior	ı, hel	p, a	nd	sei	rvi	ce	•	•	•	•	. xi	V
Before y	ou c	all .	•••	•	•	•	•	•	·	•	•	•	. xi	V
Using th	ie do	ocun	nenta	t101	n.		·	·	·	•		•	. x	V
Finding	Stora	age	Man	age	r so	oftv	var	e, o	con	itro	llei			
Escontia	e, an 1 mol	u re boite	aum		220		חס	2401	10	חס			. x	V
DS5020	and	the		Ser	500 Tipe	6111	D. nna	nt Srt	inf	orr	nat	ion	v	w
Software	o ser	vice	and	SUIT	000	ort	PP	511		011	iiai	1011	· ^	vi
Hardwa	re se	rvic	e an	d si	ipc ipc	ort	•						xv	ii
Taiwan	Cont	act	Infor	ma	tion	ι.							xv	ii
Fire sup	press	sion	syst	ems	5								xv	ii
1			5											
Chapter	1. A	bo	ut t	he	С	om	m	an	d	Lir	٦e			
Interface														1
How to Us	e the	e Co	mma	nd	Lir	ne I	nte	erfa	ce.					1
Usage N	lotes													2
CLI con	ıman	lds												2
Comma	nd L	ine	Term	ina	ls .									3
Formatt	ing c	ons	idera	tior	ns.									6
Detailed	Erro	or R	epor	ting	; ·					•			•	7
Exit stat	us.	÷		•	·	·	•	·	·	•		•	. 1	2
Usage e	xamp	oles		•	·	•	·	•	•	•	•	•	. 1	3
Chapter	2. A	bo	ut t	he	So	rip	ot	Сс	om	m	an	ds	1	5
Script Com	man	d St	ructi	ıre									. 1	6
Script Com	man	d Sy	ynop	sis		•							. 1	7
Recurring S	Synta	ax E	leme	nts	·	·	•	·	•	•		•	. 1	.9
Usage Guid	lelin	es	· ·		•		·	·	·	•	·	•	. 2	:5
Adding Co	mme	ents	to a	Scr	ıpt	File	e.	·	•	•	•	·	. 2	.5
Chapter	3. S	cri	pt C	or	nm	nar	nd	S	-				. 2	7
Naming co	nven	tior	ns.										. 2	7
Firmware (	Comp	patil	oility	Le	vel	s.	•	•					. 2	8
Hardware	Requ	iren	nents	s foi	r Fi	rm	wa	re	Co	mp	atil	oilit	y 2	<u>'</u> 9
Formatting	rule	s fo	r scri	ipt (	con	nma	anc	ls	•	•	·	·	. 2	9
DS3000, DS	\$4000	), DS	55000	), ai	nd :	the	D	CS	Sei	ries			~	0
Commands	with	n mi	nimi Clab	im	firr	nwa	are	le	vel	s .	·	·	. :	50 :1
Activate Er	inano	cea	GIOD	ar r	VIIT	rori	ng	~ Т	•	•	•	·	. 6	)1 :)
Activate Er	uiano ost P	ort	Kem	ote	1711	1101	un	g r	eat	ure	:.	•	. 6	i∠ 5
Activate iS	CSI I	niti	 ator	•	·	•	·	·	·	·	·	•	. 0	55
Activate St	orage	- S11	hsvs	tem	· Fi	۰ rmv	wa	re	•	·	·	•	. 6	55
Activate Er	hand	ced	Rem	ote	Mi	rro	rin	g.	•	•	•	•	. 6	6
Add logica	l dri	ve to	o Enl	nan	ced	l Gl	ob	al l	Miı	ror	G	rou	р 6	6

Add Member to Consistency Group	. 66
Add drives to Performance Read Cache	. 69
Autoconfigure Storage Subsystem	. 69
Autoconfigure Storage Subsystem Hot Spares	. 73
Cancel Enhanced Global Mirror Group role reversal	73
Change Performance Read Cache application type	74
Check Enhanced Global Mirror Group consistency	74
Check Logical Drive Parity	. 75
Check Remote Mirror Status.	. 76
Check Repository Consistency	. 76
Check storage subsystem connectivity	. 77
Check Enhanced Remote Mirroring status	. 78
Clear Enhanced Global Mirroring fault	. 78
Clear Drive Channel Statistics	. 79
Clear Logical Drive Reservations	. 79
Clear Logical Drive Unreadable Sectors	. 79
Clear Storage Subsystem Configuration	. 80
Clear storage subsystem Diagnostic Data Capture	81
Clear Storage Subsystem Diagnostic Data Capture	81
Clear Storage Subsystem Event Log	. 82
Clear Storage Subsystem Firmware Pending Area.	. 82
Clear Storage Subsystem Recovery Mode	. 82
Configure Automatic Support Bundle Collection .	. 83
Convert FlashCopy Logical Drive to Enhanced	
FlashCopy Group	. 83
Create subsystem	. 84
Create Enhanced Global Mirror Group	. 88
Create Consistency Group	. 88
Create Consistency Group Enhanced FlashCopy .	91
Create Consistency Group Enhanced FlashCopy	. /1
Create Consistency Group Enhanced FlashCopy Image	. 92
Create Consistency Group Enhanced FlashCopy Image	. 92
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118 119
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118 119
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118 119 123
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118 119 123 126
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118 119 123 126 131
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118 119 123 126 131
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118 119 123 126 131
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118 119 123 126 131 134 136
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118 119 123 126 131 134 136 137
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118 119 123 126 131 134 136 137 139
Create Consistency Group Enhanced FlashCopy Image	. 92 . 93 . 95 . 96 101 107 108 110 116 117 118 118 119 123 126 131 134 136 137 139 140

Create VolumeCopy	141
Deactivate Remote Mirror	143
Deactivate Enhanced Global Mirroring	143
Deactivate Enhanced Remote Mirroring	143
Delete subsystem	143
Delete Enhanced Global Mirror Group	144
Delete Consistency Group	144
Delete Consistency Group Enhanced FlashCopy	
Logical Drive	145
Delete Consistency Group Enhanced FlashCopy	
Image	145
Delete Disk Pool	146
Delete Enhanced FlashCopy Group	147
Delete Enhanced FlashCopy Image	148
Delete Enhanced FlashCopy Logical Drive.	150
Delete FlashCopy Logical Drive	151
Delete Host	151
Delete Host Group	152
Delete Host Group	152
Delete iSCSI Initiator	152
Delete Logical Drive	152
Delete VolumeGroup	154
Delete Logical Drive on a Disk Pool	154
Delete Porformance Read Cache	154
Diagnose Controller	155
Diagnose Controller :	150
Diagnose Controller 15C51 Host Cable	157
Diagnose Remote Mirror	150
Diagnose Ennanced Kemole Mirroring	150
Disable External Security Key Management	139
$\mathbf{D}^{\prime}$ 111 $\mathbf{F}$ 1. $\mathbf{C}$ .	1 50
Disable FlashCopy.	159
Disable FlashCopy.       .       .       .       .       .         Disable Storage Subsystem Feature       .       .       .       .       .         Disable Storage Subsystem Feature       .       .       .       .       .	159 160
Disable FlashCopy	159 160
Disable FlashCopy	159 160 162
Disable FlashCopy	159 160 162 162
Disable FlashCopy	159 160 162 162 162
Disable FlashCopy	159 160 162 162 162 164
Disable FlashCopy	159 160 162 162 162 164 164
Disable FlashCopy	159 160 162 162 162 164 164 164
Disable FlashCopy	159 160 162 162 162 164 164 165 166
Disable FlashCopy Disable Storage Subsystem Feature Display Automatic Support Bundle Collection Configuration	159 160 162 162 164 164 165 166 167
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167 168
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167 168 168
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167 168 168 168
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167 168 168 168 169 169
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167 168 168 169 169 170
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167 168 168 169 169 170
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167 168 168 169 169 170 170
Disable FlashCopy	159 160 162 162 164 164 165 166 167 168 168 169 169 170 170
Disable FlashCopy	159 160 162 162 164 164 165 166 167 168 168 169 169 170 170 171
Disable FlashCopy	159 160 162 162 164 164 165 166 167 168 168 169 169 170 170 171 172
Disable FlashCopy	159 160 162 162 164 164 165 166 167 168 168 169 169 170 170 170 171 172 173
Disable FlashCopy	159 160 162 162 164 164 165 166 167 168 168 169 169 170 170 170 171 172 173 174
Disable FlashCopy	159 160 162 162 164 164 164 165 166 167 168 168 169 169 170 170 171 172 173 174 174
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167 168 168 169 169 170 170 171 172 173 174 174 175
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167 168 168 169 169 170 170 171 172 173 174 174 175 176
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167 168 168 169 169 170 170 171 172 173 174 174 175 176 177
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167 168 169 169 170 170 171 172 173 174 174 175 176 177 178
Disable FlashCopy	159 160 162 162 164 164 165 166 167 167 168 169 169 170 170 171 172 173 174 174 175 176 177 178 179

Re-create Enhanced Remote Mirroring Repository	
Logical Drive	183
Re-create External Security Key	185
Re-create FlashCopy	186
Remove Logical Drive LUN Mapping	187
Remove drives from Performance Read Cache	188
Remove incomplete Enhanced Global Mirrored	
Pair from an Enhanced Global Mirror Group	189
Remove volume from Enhanced Global Mirror	
Group	190
Remove Logical Drive LUN Mapping	191
Remove Member Logical Drive from Consistency	
Group	192
Remove Remote Mirror	192
Remove Enhanced Remote Mirroring	192
Remove VolumeCopy	193
Rename Enhanced FlashCopy Logical Drive	193
Rename Performance Read Cache	194
Repair Logical Drive Parity.	194
Replace Disk Drive	195
Reset Enhanced Global Mirror Group statistics	196
Reset Controller	197
Reset iSCSI IP address	198
Reset Storage Subsystem Battery Install Date	199
Reset Storage Subsystem Diagnostic Data	199
Reset Storage Subsystem iSCSI Baseline	200
Reset Storage Subsystem Logical Drive Distribution	200
Reset Storage Subsystem RLS Baseline	201
Reset Storage Subsystem SAS PHY Baseline	201
Reset Storage Subsystem SOC Baseline	201
Resume Enhanced Global Mirror Group	001
	201
Resume Consistency Group Enhanced FlashCopy	201
Resume Consistency Group Enhanced FlashCopy Logical Drive	201
Resume Consistency Group Enhanced FlashCopy Logical Drive	201 202 203
Resume Enhanced FlashCopy Image Rollback Resume Enhanced FlashCopy Image Rollback	201 202 203 204
Resume Enhanced FlashCopy Index Copy Index Consistency Group Enhanced FlashCopy Logical Drive	201 202 203 204 206
Resume Consistency Group Enhanced FlashCopy Logical Drive	201 202 203 204 206 206
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache	201 202 203 204 206 206 207
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.	201 202 203 204 206 206 207 207
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.         Revive Disk Drive.	201 202 203 204 206 206 207 207 207
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.         Revive Disk Drive.         Revive Enhanced FlashCopy Group	201 202 203 204 206 206 207 207 207 208 208
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive	201 202 203 204 206 206 207 207 207 208 208 208 209
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive	201 202 203 204 206 206 207 207 208 208 208 209 209
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Logical Drive         Save Enhanced Global Mirror Group statistics	201 202 203 204 206 206 207 207 208 208 208 209 209 210
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Logical Drive         Save Enhanced Global Mirror Group statistics         Save Controller NVSRAM	201 202 203 204 206 206 207 207 208 209 209 209 210 212
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Logical Drive         Save Enhanced Global Mirror Group statistics         Save Controller NVSRAM         Save Disk Drive Channel Fault Isolation Diagnostic	201 202 203 204 206 206 207 207 207 208 209 209 209 210 212
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Logical Drive         Save Enhanced Global Mirror Group statistics         Save Controller NVSRAM         Save Disk Drive Channel Fault Isolation Diagnostic	201 202 203 204 206 206 207 207 208 208 209 209 210 212 212
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive         Save Enhanced Global Mirror Group statistics         Save Controller NVSRAM         Save Disk Drive Channel Fault Isolation Diagnostic         Status       .         Save Drive Log	201 202 203 204 206 207 207 208 209 209 209 209 210 212 212 212
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.         Revive array.         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Logical Drive         Save Enhanced Global Mirror Group statistics         Save Controller NVSRAM         Save Disk Drive Channel Fault Isolation Diagnostic         Status       .         Save Drive Log       .	201 202 203 204 206 207 207 208 209 209 209 210 212 212 213 213
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume FlashCopy Rollback         Resume Performance Read Cache         Revive array.         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive         Save Enhanced Global Mirror Group statistics         Save Controller NVSRAM         Save Disk Drive Channel Fault Isolation Diagnostic         Status       .         Save Enclosure Log       .         Save Storage Subsystem Configuration       .	201 202 203 204 206 207 207 208 208 209 209 210 212 212 212 213 213 214
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.         Revive array.         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Logical Drive         Save Enhanced Global Mirror Group statistics         Save Controller NVSRAM         Save Disk Drive Channel Fault Isolation Diagnostic         Status       .         Save Enclosure Log       .         Save Storage Subsystem Configuration       .	201 202 203 204 206 207 207 207 208 208 209 209 210 212 212 213 213 214 215
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.         Revive Disk Drive.         Revive Enhanced FlashCopy Group.         Revive Enhanced FlashCopy Logical Drive         Save Enhanced FlashCopy Logical Drive         Save Enhanced Global Mirror Group statistics         Save Controller NVSRAM         Save Disk Drive Channel Fault Isolation Diagnostic         Status         Save Enclosure Log         Save Storage Subsystem Configuration         Save Storage Subsystem DBM Database         Save Storage Subsystem DBM Validator	201 202 203 204 206 207 207 207 208 208 209 209 210 212 212 213 213 214 215
Resume Consistency Group Enhanced FlashCopy         Logical Drive         Resume Enhanced FlashCopy Image Rollback         Resume Enhanced FlashCopy Logical Drive         Resume Enhanced FlashCopy Logical Drive         Resume FlashCopy Rollback         Resume Remote Mirror         Resume Performance Read Cache         Revive array.         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Group         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Logical Drive         Revive Enhanced FlashCopy Logical Drive         Save Enhanced Global Mirror Group statistics         Save Controller NVSRAM         Save Disk Drive Channel Fault Isolation Diagnostic         Status         Save Enclosure Log         Save Storage Subsystem Configuration         Save Storage Subsystem DBM Database         Save Storage Subsystem DBM Validator         Information File	201 202 203 204 206 207 207 208 209 209 210 212 212 213 213 214 215
Resume Consistency Group Enhanced FlashCopyLogical DriveResume Enhanced FlashCopy Image RollbackResume Enhanced FlashCopy Logical DriveResume FlashCopy RollbackResume Remote MirrorResume Performance Read CacheRevive array.Revive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveSave Enhanced Global Mirror Group statisticsSave Controller NVSRAMSave Disk Drive Channel Fault Isolation DiagnosticStatusSave Enclosure LogSave Storage Subsystem DBM DatabaseSave Storage Subsystem DBM ValidatorInformation FileSave Storage Subsystem Diagnostic Data	201 202 203 204 206 207 207 208 209 209 210 212 213 213 213 214 215 215 216
Resume Consistency Group Enhanced FlashCopyLogical DriveResume Enhanced FlashCopy Image RollbackResume Enhanced FlashCopy Logical DriveResume FlashCopy RollbackResume Remote MirrorResume Performance Read CacheRevive array.Revive array.Revive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveSave Enhanced Global Mirror Group statisticsSave Controller NVSRAMSave Disk Drive Channel Fault Isolation DiagnosticStatusSave Enclosure LogSave Storage Subsystem DBM DatabaseSave Storage Subsystem DBM ValidatorInformation FileSave Storage Subsystem Diagnostic DataSave Storage Subsystem Diagnostic Data	201 202 203 204 206 207 207 208 209 209 210 212 213 213 213 214 215 216 217
Resume Consistency Group Enhanced FlashCopyLogical DriveResume Enhanced FlashCopy Image RollbackResume Enhanced FlashCopy Logical DriveResume FlashCopy RollbackResume Remote MirrorResume Performance Read CacheRevive array.Revive array.Revive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveSave Enhanced Global Mirror Group statisticsSave Controller NVSRAMSave Drive LogSave Drive LogSave Storage Subsystem ConfigurationSave Storage Subsystem DBM DatabaseSave Storage Subsystem Diagnostic DataSave Storage Subsystem Diagnostic DataSave Storage Subsystem Diagnostic DataSave Storage Subsystem Diagnostic Data CaptureSave Storage Subsystem Events	201 202 203 204 206 207 207 208 208 209 210 212 213 213 213 214 215 216 217 217
Resume Consistency Group Enhanced FlashCopyLogical DriveResume Enhanced FlashCopy Image RollbackResume Enhanced FlashCopy Logical DriveResume FlashCopy RollbackResume Remote MirrorResume Performance Read CacheRevive array.Revive array.Revive Disk Drive.Revive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveSave Enhanced Global Mirror Group statisticsSave Controller NVSRAMSave Disk Drive Channel Fault Isolation DiagnosticStatusSave Enclosure LogSave Storage Subsystem ConfigurationSave Storage Subsystem DBM ValidatorInformation FileSave Storage Subsystem Diagnostic DataSave Storage Subsystem Firmware Inventory	201 202 203 204 206 207 207 208 208 209 209 210 212 213 213 214 215 216 217 217 217 218
Resume Consistency Group Enhanced FlashCopyLogical DriveResume Enhanced FlashCopy Image RollbackResume Enhanced FlashCopy Logical DriveResume FlashCopy RollbackResume Remote MirrorResume Performance Read CacheRevive array.Revive array.Revive Disk Drive.Revive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveSave Enhanced Global Mirror Group statisticsSave Controller NVSRAMSave Disk Drive Channel Fault Isolation DiagnosticStatusSave Enclosure LogSave Storage Subsystem ConfigurationSave Storage Subsystem DBM DatabaseSave Storage Subsystem Diagnostic DataSave Storage Subsystem Firmware InventorySave storage Subsystem Firmware InventorySave storage Subsystem Firmware Inventory	201 202 203 204 206 207 207 208 209 209 210 212 213 213 214 215 216 217 217 217 218 218
Resume Consistency Group Enhanced FlashCopyLogical DriveResume Enhanced FlashCopy Image RollbackResume Enhanced FlashCopy Logical DriveResume FlashCopy RollbackResume Remote MirrorResume Performance Read CacheRevive array.Revive array.Revive Disk Drive.Revive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveRevive Enhanced Global Mirror Group statisticsSave Controller NVSRAMSave Disk Drive Channel Fault Isolation DiagnosticStatusSave Enclosure LogSave Storage Subsystem ConfigurationSave Storage Subsystem DBM DatabaseSave Storage Subsystem Diagnostic DataSave Storage Subsystem Firmware InventorySave Storage Subsystem Firmware InventorySave Storage Subsystem Firmware InventorySave Storage Subsystem infiniBand statisticsSave Storage Subsystem infiniBand statistics	201 202 203 204 206 207 207 208 209 209 210 212 213 213 214 215 216 217 217 218 218 219
Resume Consistency Group Enhanced FlashCopyLogical DriveResume Enhanced FlashCopy Image RollbackResume Enhanced FlashCopy Logical DriveResume FlashCopy RollbackResume Remote MirrorResume Performance Read CacheRevive array.Revive array.Revive Disk Drive.Revive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveRevive Enhanced Global Mirror Group statisticsSave Controller NVSRAMSave Drive LogSave Drive LogSave Storage Subsystem ConfigurationSave Storage Subsystem DBM DatabaseSave Storage Subsystem Diagnostic DataSave Storage Subsystem Firmware InventorySave Storage Subsystem Performance StatisticsSave Storage Subsystem Performance Statistics	201 202 203 204 206 207 207 208 209 209 210 212 213 213 214 215 216 217 217 218 218 219 220
Resume Consistency Group Enhanced FlashCopyLogical DriveResume Enhanced FlashCopy Image RollbackResume Enhanced FlashCopy Logical DriveResume FlashCopy RollbackResume Remote MirrorResume Performance Read CacheRevive array.Revive array.Revive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy GroupRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveRevive Enhanced FlashCopy Logical DriveRevive Enhanced Global Mirror Group statisticsSave Controller NVSRAMSave Disk Drive Channel Fault Isolation DiagnosticStatusSave Brive LogSave Storage Subsystem DBM DatabaseSave Storage Subsystem DBM ValidatorInformation FileSave Storage Subsystem Diagnostic DataSave Storage Subsystem Diagnostic DataSave Storage Subsystem Diagnostic DataSave Storage Subsystem Diagnostic DataSave Storage Subsystem Firmware InventorySave Storage Subsystem Performance StatisticsSave Storage Subsystem RLS Counts	201 202 203 204 206 207 207 208 209 210 212 213 213 214 215 216 217 217 218 218 219 220 220

Save Storage Subsystem SAS PHY Counts	221
Save Storage Subsystem SOC Counts	221
Save Storage Subsystem State Capture	221
Save Storage Subsystem Support Data	222
Schedule Support Bundle Collection Configuration	223
Set subsystem	224
Set subsystem Forced State	226
Set Enhanced Global Mirror Group	227
Set Consistency Group Attributes.	230
Set Consistency Group Enhanced FlashCopy	
Logical Drive	231
Set Controller	232
Set Controller Service Action Allowed Indicator	235
Set Disk Pool	236
Set Disk Pool (Modify Disk Pool).	238
Set Drawer Service Action Allowed Indicator.	239
Set Drive Channel Status	240
Set Disk Drive Hot Spare	241
Set Disk Drive Service Action Allowed Indicator	242
Set Drive State	243
Set Enclosure Alarm	243
Set Enclosure Identification	244
Set Enclosure Service Action Allowed Indicator	244
Set Enhanced FlashCopy Group Attributes	246
Set Enhanced FlashCopy Group Media Scan	247
Set Enhanced FlashCopy Group Schedule	247
Set Enhanced FlashCopy Logical Drive Media Scan	251
Set Enhanced FlashCopy Logical Drive Repository	201
Logical Drive Canacity	252
Set Enhanced ElashCony Croup Repository Logical	202
Set Enhanced Thasheopy Gloup Repository Logical	
Drive Capacity	252
Drive Capacity	252 253
Drive Capacity	252 253 254
Drive Capacity	252 253 254 255
Drive Capacity	252 253 254 255 256
Drive Capacity	<ul> <li>252</li> <li>253</li> <li>254</li> <li>255</li> <li>256</li> <li>257</li> </ul>
Drive Capacity	<ul> <li>252</li> <li>253</li> <li>254</li> <li>255</li> <li>256</li> <li>257</li> <li>258</li> </ul>
Drive Capacity	252 253 254 255 256 257 258 258
Drive Capacity	<ul> <li>252</li> <li>253</li> <li>254</li> <li>255</li> <li>256</li> <li>257</li> <li>258</li> <li>258</li> <li>259</li> </ul>
Drive Capacity	<ul> <li>252</li> <li>253</li> <li>254</li> <li>255</li> <li>256</li> <li>257</li> <li>258</li> <li>258</li> <li>259</li> <li>260</li> </ul>
Drive Capacity	252 253 254 255 256 257 258 258 258 259 260 265
Drive Capacity	252 253 254 255 256 257 258 258 258 258 259 260 265 271
Drive Capacity	252 253 254 255 256 257 258 258 258 259 260 265 271 276
Drive Capacity	252 253 254 255 256 257 258 258 258 258 259 260 265 271 276
Drive Capacity	252 253 254 255 256 257 258 258 258 259 260 265 271 276
Drive Capacity	252 253 254 255 256 257 258 258 259 260 265 271 276 278 278
Drive Capacity	252 253 254 255 256 257 258 258 259 260 265 271 276 278 278 278 280
Drive Capacity	252 253 254 255 256 257 258 258 259 260 265 271 276 278 278 278 278 280 281
Drive Capacity	252 253 254 255 256 257 258 258 259 260 265 271 276 278 278 278 280 281 284
Drive Capacity	252 253 254 255 256 257 258 259 260 265 271 276 278 278 278 280 281 284 285
Drive Capacity	252 253 254 255 256 257 258 259 260 265 271 276 278 278 278 278 278 280 281 284 285 285
Drive Capacity	252 253 254 255 256 257 258 259 260 265 271 276 278 278 278 278 280 281 284 285 285 286
Drive Capacity	252 253 254 255 256 257 258 259 260 265 271 276 278 278 278 278 280 281 284 285 285 285 286
Drive Capacity	252 253 254 255 256 257 258 259 260 265 271 276 278 278 278 280 281 284 285 285 286 286 287
Drive Capacity	252 253 254 255 256 257 258 259 260 265 271 276 278 278 280 281 284 285 285 286 285 286 287 287
Drive Capacity	252 253 254 255 256 257 258 259 260 265 271 276 278 278 280 281 284 285 285 285 286 285 286 287 287
Drive Capacity	252 253 254 255 256 257 258 259 260 265 271 276 265 271 276 278 280 281 284 285 285 286 285 286 286 287 287 288 289
Drive Capacity	252 253 254 255 256 257 258 259 260 265 271 276 278 280 281 284 285 285 285 286 285 286 287 288 288 289
Drive Capacity	252 253 254 255 256 257 258 259 260 265 271 276 278 280 281 284 285 285 285 286 285 286 287 288 288 289 280
Drive Capacity	252 253 254 255 256 257 258 259 260 265 271 276 278 278 280 281 284 285 285 286 285 286 287 288 288 288 289 289 201
Drive Capacity	252 253 254 255 256 257 258 259 260 265 271 276 278 278 280 281 284 285 285 286 285 286 285 286 287 288 288 289 289 291 202

Set VolumeCopy	. 293
Show subsystem	. 293
Show subsystem Export Dependencies	. 294
Show subsystem Import Dependencies	. 294
Show Enhanced Global Mirror Group	295
Show Enhanced Global Mirror Group	>0
synchronization progress	296
Show Cache Backup Device Diagnostic Status	· 207
Show Cache Dackup Device Diagnostic Status	. 207
Show Cache Memory Diagnostic Status	. 297
Show consistency group.	. 297
Snow Consistency Group Enhanced FlashCopy	••••
Image	. 298
Show Controller	. 300
Show Controller Diagnostic Status	. 301
Show Controller NVSRAM	. 301
Show Current iSCSI Sessions	. 302
Show Disk Pool	. 302
Show Disk Drive	. 303
Show Drive Channel Statistics.	. 305
Show Drive Download Progress	. 305
Show Enhanced FlashCopy Group	. 306
Show Enhanced FlashCopy Image	. 307
Show Enhanced FlashCopy Logical Drives	309
Show Enhanced Remote Mirroring Logical Drive	
Candidates	310
Show Enhanced Romete Mirroring Logical Drive	. 510
Silow Enhanced Remote Millioning Logical Drive	210
Synchronization Progress	. 510
Show Host Interface Card Diagnostic Status	. 311
Show Host Ports	. 311
Show Logical Drive	. 312
Show Logical Drive Action Progress	. 313
Show Performance Read Cache	. 313
Show Array	. 314
Show array export dependencies	. 315
Show Array import dependencies	. 315
Show Logical Drive Performance Statistics	. 316
Show Logical Drive Reservations	. 316
Show Storage Subsystem	. 317
Show Storage Subsystem Auto Configure	. 320
Show storage subsystem Diagnostic Data Capture	323
Show Storage Subsystem DBM Database	. 323
Show Storage Subsystem Diagnostic Data Capture	324
Show Storage Subsystem Features	. 324
Show storage subsystem unconfigured iSCSI	
Initiators	. 324
Show Storage Subsystem Host Topology	325
Show Storage Subsystem Profile	325
Show Storage Subsystem LUN Mappings	326
Show Storage Subsystem Lord Mappings	226
Show Storage Subsystem Unreadable Sectors	. 320
Show Storage Subsystem Unreadable Sectors.	. 320
Show Enhanced Remote Mirroring volume	007
Snow Enhanced Remote Mirroring volume	. 327
synchronization progress	. 327
Show String	. 327
	. 327 . 327 . 328
Show Thin Logical Drive	. 327 . 327 . 328 . 328
Show Thin Logical Drive	. 327 . 327 . 328 . 328
Show Thin Logical Drive	. 327 . 327 . 328 . 328 . 328
Show Thin Logical Drive	. 327 . 327 . 328 . 328 . 328 . 329 . 330
Show Thin Logical Drive	. 327 . 327 . 328 . 328 . 328 . 329 . 330 . 330

Start subsystem Defragment	331
Start subsystem Export	331
Start array import	332
Start subsystem Locate	333
Start Enhanced Global Mirroring synchronization	333
Start Cache Backup Device Diagnostic	334
Start Cache Memory Diagnostic	335
Start Disk Pool Locate	337
Start Consistency Group Enhanced FlashCopy	
Rollback	337
Start Controller Diagnostic	338
Start Controller Trace	340
Start Drive Channel Fault Isolation Diagnostics	341
Start Drive Channel Locate	342
Start Disk Drive Initialize	342
Start Drive Locate	3/3
Start Dick Drive Reconstruction	242
Start Engloguro Logato	244
Start Enclosure Locate	244
Start Enhanced FlashCopy Image Rollback	344
Start Ennanced Remote Mirroring Synchronization	346
	346
Start Host Interface Card Diagnostic.	346
Start ISCSI DHCP Retresh	349
Start Array defragment	349
Start Array export.	350
Start Array import.	350
Start Array locate	351
Start Logical Drive Initialization	351
Start Secure Disk Drive Erase	352
Start Performance Read Cache performance	
modeling.	352
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate	352 353
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database	352 353
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic	352 353 353
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture	352 353 353 355
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem iSNS Server Refresh.	352 353 353 355 355
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem iSNS Server Refresh.         Stop array locate	352 353 353 355 355 355
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem iSNS Server Refresh.         Stop array locate         Stop Cache Backup Device Diagnostic	352 353 353 355 355 355 355
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh.         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic	352 353 355 355 355 355 356 356
Start Performance Kead Cache performancemodeling.Start Storage Subsystem LocateStart Storage Subsystem Configuration DatabaseDiagnosticStart Storage Subsystem Diagnostic Data CaptureStart Storage Subsystem iSNS Server Refresh.Stop array locateStop Cache Backup Device DiagnosticStop Storage Subsystem Configuration Database	352 353 353 355 355 355 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Diagnostic	352 353 353 355 355 355 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy	352 353 355 355 355 356 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem iSNS Server Refresh.         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback	352 353 355 355 355 356 356 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem iSNS Server Refresh         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback       .         Stop Consistency Group Enhanced FlashCopy	352 353 355 355 355 356 356 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem iSNS Server Refresh.         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback       Stop Consistency Group Enhanced FlashCopy         Logical Drive	352 353 355 355 355 356 356 356 356 357
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh.         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback         Stop Consistency Group Enhanced FlashCopy         Logical Drive         Stop Controller Diagnostic	352 353 355 355 355 356 356 356 356 356 357 358 358
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem iSNS Server Refresh         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback       .         Stop Consistency Group Enhanced FlashCopy         Logical Drive       .         Stop Controller Diagnostic       .         Stop Controller Diagnostic       .         Stop Controller Diagnostic       .         Stop Controller Diagnostic       .	352 353 355 355 355 356 356 356 356 356 357 358 358 359
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem ISNS Server Refresh         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Cache Memory Diagnostic         Stop Cache Memory Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback       .         Stop Consistency Group Enhanced FlashCopy         Logical Drive       .         Stop Controller Diagnostic       .         Stop Disk Pool Locate       .         Stop Drive Channel Fault Isolation Diagnostics       .	352 353 355 355 355 356 356 356 356 356 357 358 358 359 359
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback       .         Stop Consistency Group Enhanced FlashCopy         Logical Drive       .         Stop Controller Diagnostic       .         Stop Disk Pool Locate       .         Stop Disk Pool Locate       .         Stop Drive Channel Fault Isolation Diagnostics       .	352 353 355 355 355 356 356 356 356 356 357 358 358 359 359
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback       .         Stop Consistency Group Enhanced FlashCopy         Logical Drive       .         Stop Controller Diagnostic       .         Stop Disk Pool Locate       .         Stop Drive Channel Fault Isolation Diagnostics       .         Stop Drive Channel Locate       .	352 353 355 355 355 356 356 356 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem ISNS Server Refresh         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback         Stop Consistency Group Enhanced FlashCopy         Logical Drive         Stop Controller Diagnostic         Stop Drive Channel Fault Isolation Diagnostics         Stop Drive Channel Locate         Stop Drive Channel Locate	352 353 355 355 355 356 356 356 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback         Stop Consistency Group Enhanced FlashCopy         Logical Drive         Stop Controller Diagnostic         Stop Controller Diagnostic         Stop Disk Pool Locate         Stop Drive Channel Fault Isolation Diagnostics         Stop Drive Channel Locate         Stop Drive Channel Locate         Stop Enclosure Locate         Stop Enclosure Locate	352 353 355 355 356 356 356 356 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback       .         Stop Consistency Group Enhanced FlashCopy         Logical Drive       .         Stop Controller Diagnostic       .         Stop Drive Channel Fault Isolation Diagnostics       .         Stop Drive Channel Locate       .         Stop Enclosure Locate       .	352 353 355 355 355 356 356 356 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback         Stop Consistency Group Enhanced FlashCopy         Logical Drive         Stop Controller Diagnostic         Stop Drive Channel Fault Isolation Diagnostics         Stop Drive Channel Locate         Stop Drive Channel Locate         Stop Drive Channel Locate         Stop Drive Channel Locate         Stop Enclosure Locate         Stop Enhanced FlashCopy Group Pending         Enhanced FlashCopy Images	352 353 355 355 355 356 356 356 356 356 357 358 358 359 359 359 359 360 360 261
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback         Stop Consistency Group Enhanced FlashCopy         Logical Drive         Stop Controller Diagnostic         Stop Drive Channel Fault Isolation Diagnostics         Stop Drive Channel Locate         Stop Enclosure Locate         Stop Enclosure Locate         Stop Enhanced FlashCopy Images         Stop Enhanced FlashCopy Images	352 353 355 355 355 356 356 356 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh.         Stop array locate         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback         Stop Consistency Group Enhanced FlashCopy         Logical Drive         Stop Disk Pool Locate         Stop Drive Channel Fault Isolation Diagnostics         Stop Drive Channel Fault Isolation Diagnostics         Stop Enclosure Locate         Stop Enhanced FlashCopy Images         Stop Enhanced FlashCopy Image	352 353 355 355 355 356 356 356 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh.         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback         Stop Consistency Group Enhanced FlashCopy         Logical Drive         Stop Controller Diagnostic         Stop Drive Channel Fault Isolation Diagnostics         Stop Drive Channel Fault Isolation Diagnostics         Stop Enclosure Locate         Stop Enhanced FlashCopy Images         Stop Enhanced FlashCopy Image         Stop Enhanced FlashCopy Image	352 353 355 355 355 356 356 356 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh.         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback         Stop Consistency Group Enhanced FlashCopy         Logical Drive         Stop Controller Diagnostic         Stop Drive Channel Fault Isolation Diagnostics         Stop Drive Channel Fault Isolation Diagnostics         Stop Drive Channel Locate         Stop Enhanced FlashCopy Images         Stop Enhanced FlashCopy Image         Stop Enhanced FlashCopy Image Rollback.	352 353 355 355 355 356 356 356 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh.         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback         Stop Consistency Group Enhanced FlashCopy         Logical Drive         Stop Controller Diagnostic         Stop Drive Channel Fault Isolation Diagnostics         Stop Drive Channel Fault Isolation Diagnostics         Stop Enclosure Locate         Stop Enhanced FlashCopy Images         Stop Enhanced FlashCopy Image         Stop Enhanced FlashCopy	352 353 355 355 355 356 356 356 356 356 356
Start Performance Read Cache performance         modeling.         Start Storage Subsystem Locate         Start Storage Subsystem Configuration Database         Diagnostic         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem Diagnostic Data Capture         Start Storage Subsystem SNS Server Refresh.         Stop Cache Backup Device Diagnostic         Stop Cache Memory Diagnostic         Stop Storage Subsystem Configuration Database         Diagnostic         Stop Consistency Group Enhanced FlashCopy         Rollback         Stop Consistency Group Enhanced FlashCopy         Logical Drive         Stop Disk Pool Locate         Stop Drive Channel Fault Isolation Diagnostics         Stop Drive Channel Fault Isolation Diagnostics         Stop Enclosure Locate         Stop Enhanced FlashCopy Images         Stop Enhanced FlashCopy Image         Stop Drive Channel Fault Isolation Diagnostics         Stop Drive Channel Fault Isolation Diagnostics         Stop Enclosure Locate         Stop Enhanced FlashCopy Images         Stop Enhanced FlashCopy Image         Stop Enhanced FlashCopy Image         Stop Enhanced FlashCopy Image         Stop Enhanced FlashCopy Image Rollback         Stop Enhanc	352 353 355 355 355 356 356 356 356 356 356

Stop Performance Read Cache performance	
modeling	55
Stop VolumeCopy	6
Stop Array locate	57
Stop Pending Enhanced FlashCopy Images on	
Consistency Group	57
Stop Storage Subsystem Drive Firmware Download 36	57
Stop Storage Subsystem Locate	68
Stop VolumeCopy	68
Suspend Enhanced Global Mirror Group 36	68
Suspend Remote Mirror	<u>,</u> 9
Suspend Performance Read Cache	'0
Suspend Enhanced Global Mirror Group 37	0'
Test Enhanced Global Mirror Group connectivity 37	'1
Validate Storage Subsystem Security Key 37	'2
Script Commands Listed by Function	'3
subsystem Commands	'3
Enhanced Global Mirroring commands 37	'4
Controller Commands	'4
Database Commands	'6
Diagnostic Data Capture Commands 37	'6
Disk Pool Commands	'6
Disk Drive Commands	'6
Enclosure Commands	7
Enhanced FlashCopy Commands	7
Enhanced Remote Mirroring Commands 37	'9
Feature Management Commands	30
FlashCopy Commands	30
Host Topology Commands	30
iSCSI Commands	31
Logical Drive Commands (For Standard and	
Thin Logical Drives and Logical Drives in a	
Disk pool)	32
Repository Logical Drive Analysis Commands 38	32
Session Command	33
Performance Read Cache commands 38	33
Storage Subsystem Commands	33
String Commands	35
Support Bundle Collection Commands 38	66
VolumeCopy Commands	66
Chapter 4. Configuring a Storage	
Subsystem	7
	-

Subsystem	•	387
Configuration Concepts		. 388
Controllers		. 388
Disk Drives		. 395
Hot Spares		. 396
Subsystems		. 397
Disk Pools		. 398
Logical Drives		. 399
RAID Levels		. 401
Hosts		. 403
Host Groups		. 404
Host Bus Adapter Host Ports		. 404
Logical Unit Numbers		. 405
Configuring a Storage Subsystem		. 405
Determining What Is on Your Storage		
Subsystem		. 405
Clearing the Configuration		. 420
Using the Auto Configure Command		. 421
Using the Create LogicalDrive Command .		. 424

Configuring a storage subsystem with disk p	000	ls	428
Modifying Your Configuration			. 435
Setting the Controller Clocks			. 435
Setting the Storage Subsystem Password			. 435
Setting the Storage Subsystem Host Type			. 436
Setting the Storage Subsystem Cache .			. 437
Setting the Modification Priority			. 440
Assigning Global Hot Spares			. 441
Saving a Configuration to a File			. 441

## Chapter 5. Using the FlashCopy

Premium Feature	143
How FlashCopy Works	443
Creating a FlashCopy Logical Drive	444
Creating a FlashCopy Logical Drive with	
User-Assigned Disk Drives	445
Creating a FlashCopy Logical Drive with	
Software-Assigned Disk Drives	446
Creating a FlashCopy Logical Drive by	
Specifying a Number of Disk Drives	446
User-Defined Parameters	447
FlashCopy Logical Drive Names and FlashCopy	
Repository Logical Drive Names	449
Changing FlashCopy Logical Drive Settings	449
Stopping, Restarting, and Deleting a FlashCopy	
Logical Drive	450

#### Chapter 6 Using the Enhanced

FlashCopy Images premium feature 453
About Enhanced FlashCopy Images Premium
Feature
Working With Enhanced FlashCopy Images
Premium Feature
Creating an Enhanced FlashCopy Image Group 460
Deleting an Enhanced FlashCopy Image Group 461
Creating an Enhanced FlashCopy Image 461
Canceling a Pending Enhanced FlashCopy
Image
Creating an Enhanced FlashCopy Image
Schedule
Deleting an Enhanced FlashCopy Image Group
(with options)
Creating a Consistency Group Enhanced
FlashCopy logical drives
Deleting a Consistency Group Enhanced
FlashCopy Image
Creating an Enhanced FlashCopy Image Logical
drive
Restarting a Consistency Group Enhanced
FlashCopy Image Logical drive
Deleting an Enhanced FlashCopy Image Logical
drive
Changing the Size of a Repository Logical drive 470
Starting, Stopping, and Resuming an Enhanced
FlashCopy Image Rollback 472
Chapter 7. About Enhanced Global

Mirroring premium feature	477
Creating an Enhanced Global Mirrored Pair .	. 482
Changing Enhanced Global Mirroring settings	. 487

Suspending and resuming the Enhanced Global		
Mirror Group		488
Manually resynchronizing logical drives in an		
Enhanced Global Mirror Group		489
Changing Enhanced Global Mirroring roles		489
Canceling a pending Enhanced Global Mirror		
Group role change		491
Resolving role conflicts		491
Removing logical drives from the Enhanced Globa	1	
Mirror group		491
Deleting an Enhanced Global Mirror Group		493

### Chapter 8. About Enhanced Remote

Mirroring premium feature	495
How Enhanced Remote Mirroring Works	495
Mirror Repository Logical Drives	496
Mirror Relationships	497
Data Replication	497
Link Interruptions or Secondary Logical Drive	
Errors	498
Resynchronization.	499
Creating a Remote Mirror Pair	499
Performance Considerations	500
Enabling the Enhanced Remote Mirroring	
Feature	500
Activating the Enhanced Remote Mirroring	
Feature	501
Determining Candidates for a Remote-Mirror	
Pair	503
Creating a Remote Mirror Pair	503
Changing Enhanced Remote Mirroring Settings	504
Suspending and Resuming a Mirror Relationship	505
Removing a Mirror Relationship	506
Deleting a Primary Logical Drive or a Secondary	
Logical Drive	506
Disabling the Enhanced Remote Mirroring Feature	506
Deactivating the Enhanced Remote Mirroring	
Feature	507
Interaction with Other Features	507
Storage Partitioning	507
FlashCopy Logical Drives	507
VolumeCopy	508
Dynamic Logical Drive Expansion (DVE)	508
Enhanced Remote Mirroring utility	508
Description of the Enhanced Remote Mirroring	
Utility	508
Operation of the Enhanced Remote Mirroring	
Utility	509
Running the Enhanced Remote Mirroring Utility	510
Configuration Utility	510
-	

## Chapter 9. Using the VolumeCopy Premium Feature

onapter 5. Osing the volumeot	- Py		
Premium Feature			513
How VolumeCopy Works			. 513
Source Logical Drive			. 513
Target Logical Drive			. 514
VolumeCopy and Persistent Reservat	ions		. 515
Storage Subsystem Performance			. 515
Restrictions			. 515
VolumeCopy Commands			. 516

Creating a logical VolumeCopy		516
Enabling the logical VolumeCopy Feature .		517
Determining logical VolumeCopy Candidate	s	517
Creating a logical VolumeCopy		517
Viewing logical VolumeCopy Properties		518
Changing logical VolumeCopy Settings		519
Recopying a Logical Drive		520
Stopping a logical VolumeCopy		521
Removing Copy Pairs		521
Interaction with Other Features		522
Storage Partitioning		522
FlashCopy Logical Drives		522
Enhanced Remote Mirroring		523

#### Chapter 10. Using the Performance

Read Cache premium feature	525
Creating the Performance Read Cache, adding	
volumes, and removing volumes	. 526
Performance Read Cache performance modeling	527
Performance Read Cache management tasks	. 530

#### Chapter 11. Maintaining a storage

subsystem
Routine Maintenance
Running a Media Scan
Running a Redundancy Check
Resetting a Controller
Enabling a Controller Data Transfer
Resetting the Battery Age
Removing Persistent Reservations
Synchronizing the Controller Clocks 535
Locating Disk Drives
Relocating an subsystem
Performance Tuning
Monitoring the Performance
Changing the RAID Levels
Changing the Segment Size
Changing the Cache Parameters
Defragmenting an subsystem
Troubleshooting and Diagnostics
Collecting All Support Data
Collecting Drive Data
Diagnosing a Controller
Running Read Link Status Diagnostics 543
Collecting Switch-on-a-Chip Error Statistics 545
Recovery Operations
Setting the Controller Operational Mode 546

Changing the Controller Ownership.		. 547
Initializing a Drive		. 547
Reconstructing a Drive		. 548
Initializing a Logical Drive		. 548
Redistributing Logical Drives		. 548
Replacing CRUs		. 549

## Chapter 12. Examples of Information

Returned by the Show	V C	om	۱m	an	ds	5.	551
Show Storage Subsystem Co	mm	and	d.				. 551
Show Controller NVSRAM .							. 591
Show Drive							. 595
Show Drive Channel Status.							. 599
Show Logical Drives							. 604
	<b>D</b>				-		044

Chapter 13. Example Script	1 1	-ile	S.		. 6	11
Configuration Script Example 1 .					. (	611
Configuration Script Example 2 .					. (	513

#### Chapter 14. Deprecated Commands

and Parameters .					615
Deprecated Commands					. 615
Deprecated Parameters					. 615

Appendix A. Examples of information returned by the show commands	617
Appendix B. Example script files	649
Appendix C. Global Copy mirror utility	653
Appendix D. Simplex-to-Duplex Conversion	657
Appendix E. Additional Documentation	663
Notices	<b>667</b> . 668 . 668 . 669
Index	671

## Figures

1.	Host and Storage Relationship
2.	DS3200 Host Ports
3.	DS3400 Host Ports
4.	DS3500 Series Host Ports
5.	DCS3700 storage subsystem host ports 391
6.	DCS3700 storage subsystem with
	Performance Module Controllers series Host
	Ports with FC HIC installed

7.	DCS3860 storage subsystem with SAS host				
	ports installed				. 392
8.	DS4200 and DS4700 Host Ports				. 393
9.	DS4800 Host Ports				. 393
10.	DS3950 and DS5020 Host Ports				. 394
11.	DS5300 / DS5100 Host Ports .				. 395

X IBM System Storage DS Storage Manager Version 10: Command Line Interface and Script Commands Programming Guide

## Tables

1.	Command Name Syntax Conventions 2	15.
2.	Command Line Terminals	
3.	Valid Trace Constants	16.
4.	Exit Status	
5.	Configuration and Management Operations 15	17.
6.	Object Types and Identifiers	
7.	General Form of the Script Commands 17	18.
8.	Recurring Syntax Elements	19.
9.	Host Ports and the Type of Host Interfaces on	20.
	Controller Modules	21.
10.	Maximum Number of Disk Drives	22.
11.	Maximum Number of Disk Drives in a	23.
	Subsystem Based on Capacity	24.
12.	Logical Drive Specifications by Supported	25.
	Controller Model (For DS3000, DS4000.	
	DS5000)	26.
13.	Logical Drive Specifications by Supported	27.
	Controller Model (For DCS Series) 401	28.
14.	RAID Level Configurations	29.
	-	

15.	Maximum Number of HBA Host Ports per	
	Controller	404
16.	Default Values for Segment Size and Cache	
	Read Prefetch	423
17.	Drive Counts and the Number of Disk Drives	
	Reserved for Reconstruction	428
18.	Values for the <b>cacheFlushModifier</b> Parameter	438
19.	Components of a FlashCopy Logical Drive	443
20.	FlashCopy Logical Drive Commands	444
21.	FlashCopy Logical Drive Parameters	447
22.	FlashCopy Rollback Operation Status	475
23.	Maximum Number of Defined Mirrors	495
24.	VolumeCopy Commands	516
25.	Information About Storage Subsystem	
	Performance	538
26.	Support Data for the Storage Subsystem	540
27.	RLS Baseline Data for the Storage Subsystem	543
28.	Deprecated Commands	615
29.	Deprecated Parameters	615

## About this document

This document is a programming guide that describes the command line interface and the script commands for the IBM System Storage DS Storage Manager software to configure, operate, and maintain a storage subsystem. This document describes all of the script commands, explains the purpose of the commands, shows the complete syntax of the commands, and defines the parameters for the commands. The command line interface code is automatically installed as part of the IBM DS Storage Manager client installation.

For instructions to install the DS Storage Manager V10.77 or earlier, see the *IBM System Storage DS Storage Manager Version 10 Installation and Host Support Guide.* 

For instructions to install the DS Storage Manager V10.83 or later, see the *IBM System Storage DS Storage Manager Version 10.8 Installation and Host Support Guide.* 

This document does not cover hardware installation or integration. For information about these topics, see the Installation and Support Guide that is appropriate for your storage subsystem. See Appendix E, "Additional Documentation," on page 663 for a list of these documents.

Check the DS3000, DS4000, DS5000, and the DCS Series readme files for the most up-to-date information regarding hardware, software, or firmware products that might not be described in this document.

### What's new in this guide

This guide contains references to the new DCS3860 storage subsystem.

## DS3000, DS4000, DS5000, and the DCS Series software and firmware support

The command line interface and the script commands described in this document are supported with the following software and firmware levels:

- **DS3000 and DCS series:** DS3000 Storage Manager Version 10.xx software in conjunction with firmware 07.xx.xx and later.
- **DS4000**<sup>®</sup>: DS4000 Storage Manager Version 10.xx software in conjunction with the DS4000 controller firmware version 07.xx.xx and later.

**Note:** Some of the script commands are also supported in earlier versions of DS4000 controller firmware. Please see "DS3000, DS4000, DS5000, and the DCS Series commands with minimum firmware levels" on page 30 for the minimum controller firmware that is required to support each script command. This table also shows which commands are supported with the DS3000, DS4000, DS5000, and the DCS Series products only, two of the products, or all the products.

• **DS5000:** DS5000 Storage Manager Version 10.50 and later software in conjunction with the DS5000 controller firmware version 07.50 and later.

## Who should read this document

This document assumes that the user has a knowledge of basic storage area network (SAN) hardware and installation skills.

This document is intended for system operators, system administrators, and service personnel who are responsible for operating, maintaining, troubleshooting, and servicing a DS3000, DS4000, DS5000, or the DCS Series storage subsystem. Users must be familiar with computer system operation, maintenance, and repair. In addition, they should understand disk storage technology, Redundant subsystem of Independent Disks (RAID) concepts, networking, and Fibre Channel and iSCSI technologies. The reader must have a basic knowledge of SAN hardware functionality (controllers, drives, and hosts) and SAN cabling.

#### Notices used in this document

This document contains notices that are designed to highlight the following key information:

**Note:** These notices provide tips, guidance, or advice.

#### Important:

These notices provide information or advice that might help you avoid inconvenient or problem situations.

#### Attention:

These notices indicate possible damage to programs, devices or data. An attention notice is placed just before the instruction or situation in which damage could occur.

#### Getting information, help, and service

If you need help, service, or technical assistance or just want more information about IBM<sup>®</sup> products, you will find a wide variety of sources available from IBM to assist you. This section contains information about where to go for additional information about IBM and IBM products, what to do if you experience a problem with your system, and whom to call for service, if it is necessary.

#### Before you call

Before you call, take these steps to try to solve the problem yourself:

- · Check all cables to make sure that they are connected.
- Check the power switches to make sure that the system is turned on.
- Use the troubleshooting information in your system documentation, and use the diagnostic tools that come with your system.
- Check for technical information, hints, tips, and new device drivers at the IBM support website pages that are listed in this section.
- Use an IBM discussion forum on the IBM Web site to ask questions.

You can solve many problems without outside assistance by following the troubleshooting procedures that IBM provides in the Storage Manager online help or in the documents that are provided with your system and software. The information that comes with your system also describes the diagnostic tests that you can perform. Most subsystems, operating systems, and programs come with

information that contains troubleshooting procedures and explanations of error messages and error codes. If you suspect a software problem, see the information for the operating system or program.

## Using the documentation

Information about your IBM system and preinstalled software, if any, is available in the documents that come with your system. This includes printed books, online documents, readme files, and help files. See the troubleshooting information in your system documentation for instructions for using the diagnostic programs. The troubleshooting information or the diagnostic programs might tell you that you need additional or updated device drivers or other software.

## Finding Storage Manager software, controller firmware, and readme files

Storage Manager software and controller firmware are available on the product DVD and can also be downloaded from the Web. Storage Manager readme files are also found on the Web.

**Important:** Before you install Storage Manager software, consult the readme file for your host operating system. Updated readme files contain the latest device driver versions, firmware levels, limitations, and other information not found in this document.

1. Go to the following website:

www.ibm.com/servers/storage/support/disk/

- 2. Click the link for your storage subsystem (for example, DS3400 or DS4800).
- **3**. When the support page for your storage subsystem opens, click the **Download** link in the Support & Downloads box. The Software and device drivers page opens.
- 4. In the Storage Manager section of the table, locate your operating system and version level, and click the version link in the right hand column. The specific page for your storage subsystem and operating system version opens.
- 5. In the table under File details, click the **\*.txt** file link, and the README will open in your Web browser.

# Essential websites for DS3000, DS4000, DS5000, DS5020 and the DCS Series support information

The most up-to-date information about DS3000, DS4000, DS5000, and the DCS Series storage subsystems and Storage Manager, including documentation and the most recent software, firmware, and NVSRAM downloads, can be found at the following websites.

#### IBM System Storage<sup>®</sup> Disk Storage Systems

Find links to software and firmware downloads, readme files, and support pages for all IBM System Storage disk storage systems, including DS3000, DS4000, DS5000, DS5020 and the DCS Series:

www.ibm.com/systems/support/storage/disk

#### **IBM System Storage Interoperation Center (SSIC)**

Find technical support information for your specific DS3000, DS4000, DS5000, DS5020 and the DCS Series storage subsystem/host configuration, including the latest recommended firmware versions for your system, by using this interactive Web-based utility:

www.ibm.com/systems/support/storage/config/ssic/index.jsp

**IBM DS3000, DS4000, DS5000 and the DCS Series Premium Feature Activation** Activate a DS3000, DS4000, DS5000 and the DCS Series premium feature by using this Web-based utility:

https://www-912.ibm.com/PremiumFeatures/jsp/keyPrereq.jsp

#### IBM System Storage DS3000, DS4000, DS5000 and the DCS Series Interoperability Support

Find the latest information about operating system and HBA support, clustering support, storage area network (SAN) fabric support, and Storage Manager feature support:

www.ibm.com/systems/support/storage/config/ssic/index.jsp

#### Storage Area Network (SAN) Support

Find information about using SAN switches, including links to SAN user guides and other documents:

www.ibm.com/systems/support/storage/san

#### Support for IBM System p<sup>®</sup>, AIX 5L<sup>™</sup>, and Linux servers

Find the latest support information for System p, AIX<sup>®</sup>, Linux, BladeCenter<sup>®</sup>, and i5/OS<sup>™</sup> servers:

www.ibm.com/systems/support/supportsite.wss/ brandmain?brandind=5000025

#### Support for IBM System x<sup>®</sup> servers

Find the latest support information for System x Intel- and AMD-based servers:

www.ibm.com/systems/support/supportsite.wss/ brandmain?brandind=5000008

#### eServer<sup>™</sup> System p and AIX Information Center

Find everything you need to know about using AIX with System p and POWER<sup>®</sup> servers:

publib.boulder.ibm.com/infocenter/pseries/index.jsp?

#### Fix Central

Find fixes and updates for your system's software, hardware, and host operating system:

www.ibm.com/eserver/support/fixes

#### **IBM System Storage products**

Find information about all IBM System Storage products:

www.storage.ibm.com

#### **IBM Publications Center**

Find IBM publications:

www.ibm.com/shop/publications/order/

#### Software service and support

Through IBM Support Line, for a fee you can get telephone assistance with usage, configuration, and software problems. For information about which products are supported by Support Line in your country or region, go to the following website:

www.ibm.com/services/sl/products

For more information about the IBM Support Line and other IBM services, go to the following websites:

- www.ibm.com/services
- www.ibm.com/planetwide

#### Hardware service and support

You can receive hardware service through IBM Integrated Technology Services or through your IBM reseller, if your reseller is authorized by IBM to provide warranty service. Visit this website for support telephone numbers:

www.ibm.com/planetwide

In the U.S. and Canada, hardware service and support is available 24 hours a day, 7 days a week. In the U.K., these services are available Monday through Friday, from 9 a.m. to 6 p.m.

## **Taiwan Contact Information**

This topic contains the product service contact information for Taiwan.

IBM Taiwan Product Service Contact Information: IBM Taiwan Corporation 3F, No 7, Song Ren Rd., Taipei Taiwan Tel: 0800-016-888

2c00790

台灣IBM產品服務聯絡方式: 台灣國際商業機器股份有限公司 台北市松仁路7號3樓 電話:0800-016-888

#### Fire suppression systems

A fire suppression system is the responsibility of the customer. The customer's own insurance underwriter, local fire marshal, or a local building inspector, or both, should be consulted in selecting a fire suppression system that provides the correct level of coverage and protection. IBM designs and manufactures equipment to internal and external standards that require certain environments for reliable operation. Because IBM does not test any equipment for compatibility with fire suppression systems, IBM does not make compatibility claims of any kind nor does IBM provide recommendations on fire suppression systems.

xviii IBM System Storage DS Storage Manager Version 10: Command Line Interface and Script Commands Programming Guide

## **Chapter 1. About the Command Line Interface**

Attention: IBM recommends using the Storage Manager client GUI to manage your storage subsystems. – The command-line interface does not have any mechanisms to prevent you from inadvertently making unwanted changes to the storage subsystem. Because the script commands are capable of damaging a configuration and causing loss of data access if not used correctly, IBM recommends using the Storage Manager client GUI to manage your storage subsystem configurations.

The command-line interface (CLI) is a software tool that lets storage subsystem installers, developers, and engineers configure and monitor storage subsystems using script commands. Using the CLI, you can run commands from an operating system prompt, such as the Windows command prompt, a Linux operating system console, or a Solaris operating system console. You have to install the IBM DS Storage Manager client in order to run the script commands either through the script window, which is invoked from the IBM DS Storage Manager client Enterprise window, or through the command line interface using the SMcli program. The script command engine is automatically installed as part of the IBM DS Storage Manager client installation.

Each command performs a specific action for managing a storage subsystem or returning information about the status of a storage subsystem. You can enter individual commands, or you can run script files when you need to perform operations more than once. For example, you can run script files when you want to install the same configuration on several storage subsystems. The CLI lets you load a script file from a disk and run the script file. The CLI provides a way to run storage management commands on more than one network storage subsystem. You can use the CLI both in installation sites and in development environments.

### How to Use the Command Line Interface

The commands that you run on the CLI provide access to the script engine, specify the storage subsystem to receive the script commands, and set operation environment parameters.

A CLI command consists of these elements:

- The term SMcli
- The storage subsystem identifier
- Parameters
- Script commands

A CLI command takes this form:

SMcli storageSubsystem parameters script-commands;

• SMcli invokes the command line interface.

**Note:** If you issue the command from the directory or folder that does not contain the SMcli.exe program, you need to include the appropriate path: directoryName/SMcli.

• *storageSubsystem* is the name or the IP address of the storage subsystem.

- *parameters* are CLI parameters that define the environment and the purpose for the command.
- *script-commands* are one or more script commands or the name of a script file that contains script commands. (The script commands are the storage subsystem configuration commands.)

#### **Usage Notes**

If you enter **SMcli** and a storage subsystem name but do not specify CLI parameters, script commands, or a script file, the command line interface runs in interactive mode. Interactive mode lets you run individual commands without prefixing the commands with **SMcli**. In interactive mode, you can enter a single command, view the results, and enter the next command without typing the complete **SMcli** string. Interactive mode is useful for determining configuration errors and quickly testing configuration changes.

To end an interactive mode session, type the operating system-specific command for terminating a program (such as **Control-C** on the UNIX operating system or the Windows operating system). Typing the termination command (**Control-C**) while in interactive mode turns off interactive mode and returns operation of the command prompt to an input mode that requires you to type the complete **SMcli** string.

If you enter an incomplete or inaccurate **SMcli** string that does not have the correct syntax, parameter names, options, or terminals, the script engine returns usage information.

## **CLI commands**

Conventions used in the command statement syntax to manage the storage subsystem are listed in the following table:

Convention	Definition		
(a   b   c)	Alternative ("a" or "b" or "c")		
italicized-words	A terminal that needs user input to fulfill a parameter (a response to a variable)		
[ ] (square brackets)	Zero or one occurrence (square brackets are also used as a delimiter for some command parameters)		
{ } (curly braces)	Zero or more occurrences		
bold	A terminal that needs a command parameter entered to start an action		

Table 1. Command Name Syntax Conventions

General forms of the CLI commands follow, which show the terminals and the parameters that are used in each command.

```
SMcli DNS-network-name-or-IP-address [DNS-network-name-or-IP-address]
[-c "command; {command2};"]
[-n storage-system-name | -w wwID]
[-o outputfile] [-p password] [-e] [-S] [-quick]
SMcli DNS-network-name-or-IP-address[DNS-network-name-or-IP-address]
[-f scriptfile]
[-n storage-system-name | -w wwID]
[-o outputfile] [-p password] [-e] [-S] [-quick]
```

```
SMcli (-n storage-system-name | -w wwID)
[-c "command; {command2};"]
[-o outputfile] [-p password] [-e] [-S] [-quick]
SMcli (-n storage-system-name -w wwID)
[-f scriptfile]
[-o outputfile] [-p password] [-e] [-S] [-quick]
SMcli -a email: email-address [host-name-or-IP-address1
[host-name-or-IP-address2]]
[-n storage-system-name ] -w wwID | -h host-name | -r (host sa | direct sa)]
[-I information-to-include] [-q frequency] [-S]
SMcli -x email: email-address [host-name-or-IP-address1
[host-name-or-IP-address2]]
[-n storage-system-name | -w wwID | -h host-name | -r (host_sa | direct_sa)] [-S]
SMcli (-a | -x) trap: community, host-name-or-IP-address
[host-name-or-IP-address1 [host-name-or-IP-address2]]
[-n storage-system-name | -w wwID | -h host-name | -r (host_sa | direct sa)] [-S]
SMcli -d [-w] [-i] [-s] [-v] [-S]
SMcli -m host-name-or-IP-address -F email-address
[-g contactInfoFile] [-S]
SMcli -A [host-name-or-IP-address [host-name-or-IP-address]] [-S]
SMcli -X (-n storage-system-name | -w wwID | -h host-name)
SMcli -?
```

## **Command Line Terminals**

Table 2. Command Line Terminals

Terminal	Definition		
host-name-or-IP-address	<ul> <li>Specifies either the host name or the Internet Protocol (IP) address (<i>xxx.xxx.xxx</i>) of an in-band managed storage subsystem or an out-of-band managed storage subsystem.</li> <li>If you are managing a storage subsystem by using a host through in-band storage management, you must use the -n terminal or the -w terminal if more than one storage subsystem is connected to the host</li> </ul>		
	<ul> <li>If you are managing a storage subsystem by using out-of-band storage management through the Ethernet connection on each controller, you must specify the <i>host-name-or-IP-address</i> of the controllers.</li> </ul>		
	<ul> <li>If you have previously configured a storage subsystem in the Enterprise Management Window, you can specify the storage subsystem by its user-supplied name by using the -n terminal.</li> </ul>		
	• If you have previously configured a storage subsystem in the Enterprise Management Window, you can specify the storage subsystem by its World-Wide Identifier (WWID) by using the <b>-w</b> terminal.		
-A	Adds a storage subsystem to the configuration file. If you do not follow the <b>-A</b> terminal with a <i>host-name-or-IP-address</i> , auto-discovery scans the local subnet for storage subsystems.		

Table 2. Command Line Terminals (continued)

Terminal	Definition
-a	Adds a Simple Network Management Protocol (SNMP) trap destination or an email address alert destination.
	<ul> <li>When you add an SNMP trap destination, the SNMP community is automatically defined as the community name for the trap, and the host is the IP address or Domain Name Server (DNS) host name of the system to which the trap should be sent.</li> <li>When you add an email address for an alert destination, the email-address is the email address to which you want the alert message to be sent.</li> </ul>
-c	Indicates that you are entering one or more script commands to run on the specified storage subsystem. End each command with a semicolon (;). You cannot place more than one -c terminal on the same command line. You can include more than one script command after the -c terminal.
-d	Shows the contents of the script configuration file. The file content takes this format:
	storage-system-name host-name1 host-name2
-е	Runs the commands without performing a syntax check first.
-F (uppercase)	Specifies the email address from which all alerts will be sent.
-T (lowercase)	Specifies a file name that contains script commands that you want to run on the specified storage subsystem. (This terminal is similar to the <b>-c</b> terminal in that both terminals are intended for running script commands. The <b>-c</b> terminal runs individual script commands. The <b>-f</b> terminal runs a file of script commands.)
	script commands in a file are ignored, and the file continues to run. To override this behavior, use the <b>set session</b> <b>errorAction=stop</b> command in the script file.
-g	Specifies an ASCII file that contains email sender contact information that will be included in all email alert notifications. The CLI assumes that the ASCII file is text only, without delimiters or any expected format. Do not use this terminal if a userdata.txt file exists.
-h	Specifies the host name that is running the SNMP agent to which the storage subsystem is connected. Use this terminal with the <b>-a</b> terminal and the <b>-x</b> terminal.
-I	<pre>Specifies the type of information to be included in the email alert notifications. You can select these values: eventOnly profile supportBundle</pre>
-i	Shows the IP address of the known storage subsystems. Use this terminal with the <b>-d</b> terminal. The file contents takes this format: <i>storage-system-name IP-address1 IPaddress2</i>

Terminal	Definition
-m	Specifies the host name or the IP address of the email server from which email alert notifications are sent.
-n	Specifies the name of the storage subsystem on which you want to run the script commands. This name is optional when you use a <i>host-name-or-IP-address</i> . If you are using the in-band method for managing the storage subsystem, you must use the <b>-n</b> terminal if more than one storage subsystem is connected to the host at the specified address. The storage subsystem name is required when the <i>host-name-or-IP-address</i> is not used. The name of the storage subsystem that is configured for use in the Enterprise Management Window (that is, the name is listed in the configuration file) must not be a duplicate name of any other configured storage subsystem.
-0	Specifies a file name for all output text that is a result of running the script commands when you use this terminal with the <b>-c</b> terminal or the <b>-f</b> terminal. If you do not specify an output file, the output text goes to standard output (stdout). All output from commands that are not script commands is sent to stdout, regardless of whether this terminal is set.
-p	Specifies the password for the storage subsystem on which you want to run commands. A password is not necessary under these conditions:
	• A password has not been set on the storage subsystem.
	• The password is specified in a script file that you are running.
	• You specify the password by using the <b>-c</b> terminal and the <b>set session password</b> = <i>password</i> command.
-q	Specifies the frequency that you want to include additional profile or support bundle information in the email alert notifications. An email alert notification containing at least the basic event information is always generated for every critical event. If you set the <b>-I</b> terminal to event0n1y, the only valid value for the <b>-q</b> terminal is everyEvent. If you set the <b>-I</b> terminal to either the profile value or the supportBundle value, this information is included with the emails with the frequency specified by the <b>-q</b> terminal. These values are valid frequency values:
	<ul> <li>everyEvent – Information is returned with every email alert notification.</li> </ul>
	• 2 – Information is returned no more than once every two hours.
	• 4 – Information is returned no more than once every four hours.
	• 8 – Information is returned no more than once every eight hours.
	• 12 – Information is returned no more than once every 12 hours.
	• 24 – Information is returned no more than once every 24 hours.

Table 2. Command Line Terminals (continued)

Table 2. Command Line	Terminals	(continued)
-----------------------	-----------	-------------

Terminal	Definition
-quick	Reduces the amount of time that is required to run a single-line operation. An example of a single-line operation is the <b>re-create FlashCopy LogicalDrive</b> command. This terminal reduces time by not running background processes for the duration of the command.
	Do not use this terminal for operations that involve more than one single-line operation. Extensive use of this command can overrun the controller with more commands than the controller can process, which causes operational failure. Also, status updates and configuration updates that are collected usually from background processes will not be available to the CLI. This terminal causes operations that depend on background information to fail.
-r	Sets or changes the alert notifications for all of the storage subsystems that are managed by a specific storage management station. Use the <b>-r</b> terminal with the <b>-a</b> terminal or the <b>-x</b> terminal. For storage subsystems that are managed out-of-band, use the direct_sa value. For storage subsystems that are managed in-band, use the host_sa value.
-S (uppercase)	Suppresses informational messages describing the command progress that appear when you run script commands. (Suppressing informational messages is also called silent mode.) This terminal suppresses these messages: • Performance syntax check
	<ul> <li>Syntax check complete</li> <li>Executing script</li> <li>Script execution complete</li> <li>SMcli completed successfully</li> </ul>
-s (lowercase)	Shows the alert settings in the configuration file when used with the <b>-d</b> terminal.
-v	Show the current global status of the known devices in a configuration file when used with the <b>-d</b> terminal.
-w	Specifies the WWID of the storage subsystem. This terminal is an alternate to the <b>-n</b> terminal. Use the <b>-w</b> terminal with the <b>-d</b> terminal to show the WWIDs of the known storage subsystems. The file content take this format: <i>storage-system-name world-wide-ID IPaddress1 IP-address2</i> .
-X (uppercase)	Deletes a storage subsystem from a configuration.
-x (lowercase)	Removes an SNMP trap destination or an email address alert destination. The <i>community</i> is the SNMP community name for the trap, and the <i>host</i> is the IP address or DNS host name of the system to which you want the trap sent.
-?	Shows usage information about the CLI commands.

## Formatting considerations

Double quotation marks ("") that are used as part of a name or label require special consideration when you run the CLI commands and the script commands on a Microsoft Windows operating system. This section describes how to use

double quotation marks in names while running CLI commands and script commands on a Windows operating system.

When double quotation marks ("") are part of a name or value, you must insert a backslash (\) before each double quotation mark character. For example: -c set storageSubsystem userLabel=\"Engineering\";"

where "Engineering" is the storage subsystem name. A second example is: -n  $\My\$ 

where "My"\_StorageSubsystem is the name of the storage subsystem.

You cannot use double quotation marks ("") as part of a character string (also called *string literal*) within a script command. For example, you cannot enter the string below in order to set the storage subsystem name to "Finance Subsystem":

-c "set storageSubsystem userLabel=\"\"Finance\"Subsystem\";"

In the Linux operating system and the Solaris operating system, the delimiters around names or labels are single quotation marks (' '). The UNIX versions of the previous examples are as follows:

-c 'set storageSubsystem userLabel="Engineering";'

-n "My"\_StorageSubsystem

In a Windows operating system, if you do not use double quotation marks ("") around a name, you must insert a caret (  $^{\land}$  ) before each special script character. Special characters are  $^{\land}$ , |, <, and >.

Insert a caret before each special script character when used with the terminals **-n**, **-o**, **-f**, and **-p**. For example, to specify storage subsystem CLI>CLIENT, enter this string:

-n CLI^>CLIENT

Insert one caret (^) before each special script character when used within a string literal in a script command. For example, to change the name of a storage subsystem to FINANCE\_|\_PAYROLL, enter the following string:

-c set storageSubsystem userLabel=\"FINANCE\_^|\_PAYROLL\";"

#### **Detailed Error Reporting**

Error data collected from an error encountered by the CLI is written to a file. Detailed error reporting under the CLI works as follows:

- If the CLI must abnormally end running CLI commands and script commands, error data is collected and saved before the CLI finishes.
- The CLI saves the error data by writing the data to a standard file name.
- The CLI automatically saves the data to a file. Special command line options are not required to save the error data.
- You are not required to perform any action to save the error data to a file.
- The CLI does not have any provision to avoid over-writing an existing version of the file that contains error data.

For error processing, errors appear as two types:

- Terminal errors or syntax errors that you might enter
- Exceptions that occur as a result of an operational error

When the CLI encounters either type of error, the CLI writes information that describes the error directly to the command line and sets a return code. Depending on the return code, the CLI also might write additional information about which terminal caused the error. The CLI also writes information about what it was expecting in the command syntax to help you identify any syntax errors that you might have entered.

When an exception occurs while a command is running, the CLI captures the error. At the end of processing the command (after the command processing information has been written to the command line), the CLI automatically saves the error information to a file.

The name of the file to which error information is saved is excprpt.txt. The CLI tries to place the excprpt.txt file in the directory that is specified by the system property devmgr.datadir. If for any reason the CLI cannot place the file in the directory specified by devmgr.datadir, the CLI saves the excprpt.txt file in the same directory from which the CLI is running. You cannot change the file name or the location. The excprpt.txt file is overwritten every time that an exception occurs. If you want to save the information in the excprpt.txt file, you must copy the information to a new file or a new directory.

Note: In Windows, this is specified by the registry key:

HKLM\SOFTWARE\storage\SMclient\SunJVMOption1

In a typical installation with no changes in the default installation directory, the location is:

- Windows operating systems: c:\Program Files\IBM\_DS....\client\data
- UNIX-based operating systems: /var/opt/SM

#### Trace facility

The SMcli program has a trace facility. The purpose of the trace facility is to collect the SMcli information during the command execution to assist IBM Technical Support in troubleshooting SMcli problems. The trace facility is enabled by entering **-trace** on the command line followed by a list of trace options separated by colons, and optionally a **-tracefile** command, followed by a filename you designate for the trace file. To get the list of options for a release, you can use the command **SMcli -trace**, as noted in the example section.

The most common options for debugging interaction with the storage subsystem are:

- 0 JAL
- 1 JAL Detail
- 79 All Errors and Exceptions

**Note:** Any time tracing is enabled, 79 - All Errors and Exceptions is also enabled.

If the storage subsystem is being managed through the Storage Manager host agent, these options would be helpful to use:

- 2 RPC-UTM Agent
- 3 RPC-UTM Agent Detail

#### Example

The examples listed below show the profile using SMcli with tracing enabled. The tracing information is being written to a local file for a storage subsystem named subsystemName. The storage subsystem was previously added to the Enterprise Management window of your Storage Manager client program.

```
SMcli -n subsystemName -c "show storageSubsystem profile;"
-trace 0:1:79 -tracefile trace_file.txt
SMcli -n subsystemName -c " clear storageSubsystem configuration;"
-trace 0:1:2:3:79 -tracefile trace file.txt
```

#### Trace options available with SMcli with Storage Manager version 10.77

A list of trace options that are available with SMcli in Storage Manager version 10.77. The available trace options can also be displayed by running the **SMcli** command with the **-trace** parameter are shown in the following table:

C:\Program Files\IBM\_DS\client>SMcli -trace is an example of a missing or an invalid trace constant.

Trace Options	Trace Constant	Description
0	JAL	JAL
1	JAL	JAL Detail
2	Agent	RPC-UTM Agent
3	Agent	RPC-UTM Agent Detail
4	JRPC	JRPC
5	JRPC-In	JRPC Raw Data In
6	JRPC-Out	JRPC Raw Data Out
7	DMV Errors	DMV Errors and Exceptions
8	SAM Main	SAM Main
9	SAM Physical	SAM Main Physical View
10	CreateVol	Create Logical drive
11	EventViewer	Event Viewer
12	PerfMon	Performance Monitor
13	PerfMon	Performance Monitor Detail
14	RecGuru	Recovery Guru
15	ScriptEngine	Script Engine
16	Download	Download
17	StoragePools	Storage Pools
18	StoragePools	Storage Pools Detail
19	NodeProvider	DMV Internal: General
20	NP-LM	DMV Internal: Link Monitor
21	NP-Device	DMV Internal: Device Interaction
22	DataStore	DMV Internal: Data Store
23	SADeviceCoupler	SADeviceCoupler

Table 3. Valid Trace Constants

Trace Options	Trace Constant	Description
24	SAStatusMonitor	SAStatusMonitor
25	DMV Node State	DMV Events: Node State
26	DMV Edge State	DMV Events: Edge State
27	DMV Domain Change	DMV Events: Domain Changes
28	DMV Progress Evt	DMV Events: Progress
29	DMV Alert Evt	DMV Events: Alerts
30	DMV Menus	DMV Display: Toolbar and Menus
31	DMV Tree	DMV Display: Tree
32	DMV Table	DMV Display: Table
33	DMV Health	DMV Display: Overall Health
34	DMV Progress	DMV Display: Progress Label/Bar
35	DMV Add	DMV Op: Add Device
36	DMV Remove	DMV Op: Remove Device
37	DMV Comment	DMV Op: Edit Comment
38	DMV AutoDisc	DMV op: Automatic Discovery
39	DMV Rescan	DMV Op: Rescan Server
40	DMV Manage	DMV Op: Manage Device
41	DMV Script	DMV Op: Script Launch
42	DMV Help	DMV Op: Online Help
43	DMV Alert	DMV Op: Alert Setting
44	DMV Focus	DMV: Focus Request Attempts
45	Tools Broker	Tools Broker
46	DMV Persistent Monitor	DMV Persistent Monitor
47	System Logger Errors	System Logger Errors and Exceptions
48	OVW Integration	OpenView Integration
49	Agent	RPC-UTM Agent Full Detail
50	Unreadable Sector	Unreadable Sector
51	Parallel Drive Firmware Download	Parallel Drive Firmware Download
52	RVM Ping Test	Mirror ping test
53	TALP	Task Assistant Launch Pad
54	Framework	Framework
55	Framework Engine	Framework Engine
56	Profile	Profile
57	НСА	Host Context Agent Detail

Table 3. Valid Trace Constants (continued)

Trace Options	Trace Constant	Description
58	Features	Features and Feature Bundles
59	State Capture	State Capture
60	Reset Cntrl	Reset Controller
61	Increase Capacity	Increase Capacity
62	RLS	RLS
63	FlashCopy	Create FlashCopy
64	Remote Logical drive Mirroring	Remote Logical drive Mirroring
65	Configuration Replication	Configuration Replication
66	Gather Support Data	Gather Support Data
67	ESM Autosync	Persistent Monitor - ESM Auto sync Manager
68	JSNMP	JSNMP General
69	JSNMP_IN	JSNMP Raw In Message
70	JSNMP_OUT	JSNMP Raw Out Message
71	Logical drive Copy Wizard	Logical drive Copy Wizard
72	Logical drive Copy Manager	Logical drive Copy Manager
73	Launch In Context	Launch In Context
74	CSD POLL AND TRAP	~ CSD Polling and Trap deprecated
75	CSD MAIN	~ CSD Main Window deprecated
76	CSD ALERT	~ CSD Alert deprecated
77	RECOVERY_PROFILE	Recovery Profile
78	RECOVERY_PROFILE_DIFFS	Recovery Profile Diffs
79	ALL ERRORS	All Errors and Exceptions
80	Memory Usage	Memory Usage
81	Menu Errors	Menu System Errors
82	Auto Config Wizard	Auto Config Wizard
83	Engine General	Engine General
84	Engine Discovery	Engine Discovery
85	Engine Bundle Aggregation	Engine Bundle Aggregation
86	Engine Change State	Engine Change State
87	Engine Backend Communication	Engine Backend Communication
88	Diagnostics Report	Diagnostics Report
89	Wizard	Wizard
90	DDC	Diagnostic Data Capture
91	Set enclosure Tags	Set enclosure Tags
92	Drive Channel Diagnostics	Drive Channel Diagnostics
93	View ISCSI Sessions Connections	View ISCSI Sessions Connections

Table 3. Valid Trace Constants (continued)

Trace Options	Trace Constant	Description
94	View ISCSI Statistics	View ISCSI Statistics
95	Import subsystem	Import subsystem
96	Export subsystem	Export subsystem
97	Factory	Factory
98	Dynamic Capacity Expansion	Dynamic Capacity Expansion
99	Replace Drive	Replace Drive
100	Rename Logical drive And subsystem	Rename Logical drive And subsystem
101	Delete Logical drive And subsystem	Delete Logical drive And subsystem
102	Pre-Read Redundancy Check	Pre-Read Redundancy Check
103	Dynamic RAID Migration	Dynamic RAID Migration
104	Lock Key Management	Lock Key Management
105	Drive Icon Factory	Drive Icon Factory
106	Media Scan	Media Scan
107	Trace errors	Trace errors and exceptions
108	Task	Task
109	Fru Diagnostics	Fru Diagnostics
110	Logical drive Operations	Logical drive Operations
111	Phonehome	Phonehome
112	DQ Trace Retrieval	DQ Trace Retrieval
113	DBM Validation	DBM Validation
114	DBM Data Collection	DBM Data Collection
115	Mapping	Mapping
116	HistoricalPerfMonitor	HistoricalPerfMonitor SMcli failed.

Table 3. Valid Trace Constants (continued)

### **Exit status**

The exit statuses that might be returned, and the meaning of each status are listed in the following table:

Status Value	Meaning
0	The command terminated without an error.
1	The command terminated with an error. Information about the error also appears.
2	The script file does not exist.
3	An error occurred while opening an output file.
4	A storage subsystem was not at the specified address.
5	Addresses specify different storage subsystems.

Table 4. Exit Status

Table 4. Exit Status (continued)

Status Value	Meaning
6	A storage subsystem name does not exist for the host agent that is connected.
7	The storage subsystem name was not at the specified address.
8	The storage subsystem name was not in the configuration file.
10	A management class does not exist for the storage subsystem.
11	A storage subsystem was not found in the configuration file.
12	An internal error occurred.
13	Invalid script syntax was found.
14	The controller was unable to communicate with the storage subsystem.
15	A duplicate argument was entered.
16	An execution error occurred.
17	A host was not at the specified address.
18	The WWID was not in the configuration file.
19	The WWID was not at the address.
20	An unknown IP address was specified.
21	The Event Monitor configuration file was corrupted.
22	The storage subsystem was unable to communicate with Event Monitor.
23	The controller was unable to write alert settings.
24	The wrong organizer node was specified.
25	The command was not available.
26	The device was not in the configuration file.
27	An error occurred while updating the configuration file.
28	An unknown host error occurred.
29	The sender contact information file was not found.
30	The sender contact information file could not be read.
31	The userdata.txt file exists.
32	An invalid <b>-I</b> value in the email alert notification was specified.
33	An invalid <b>-f</b> value in the email alert notification was specified.

## Usage examples

Here are some examples that show how CLI commands are entered on a command line. The examples show the syntax, the form, and, in some examples, script commands. Examples are shown for both the Windows operating system and the UNIX operating system. Note that the usage for the **-c** terminal varies depending on your operating system. On Windows operating systems, enclose the script command following the **-c** terminal in double quotation marks (" "). On UNIX operating systems, enclose the script command following the script command following the **-c** terminal in single quotation marks (' '). (For descriptions of the script commands used in these examples, see Chapter 3, "Script Commands," on page 27.

This example shows how to change the name of a storage subsystem. The original name of the storage subsystem is Payroll\_Subsystem. The new name is Finance\_Subsystem. The storage subsystem is managed in-band through the host name ICTSANT.

• Windows operating system:

```
SMcli ICTSANT -n "Payroll_Subsystem"
  -c "set storageSubsystem userLabel=\"Finance Subsystem\";"
```

• UNIX operating system:

```
SMcli ICTSANT -n 'Payroll_Subsystem'
  -c 'set storageSubsystem userLabel="Finance Subsystem";'
```

This example shows how to delete an existing logical drive and create a new logical drive on a storage subsystem. The existing logical drive name is Stocks\_<\_Bonds. The new logical drive name is Finance. The controller host names are finance1 and finance2. The storage subsystem is protected, requiring the password TestSubsystem.

• Windows operating system:

```
SMcli finance1 finance2
   -c 'set session password=\"TestSubsystem\"; delete logicalDrive
   [\"Stocks_<Bonds\"];
   create logicalDrive driveCount[3] RAIDLEVEL=3 capacity=10GB userLabel=\
   "Finance\";
   show storageSubsystem healthStatus;"</pre>
```

#### UNIX operating system:

```
SMcli finance1 finance2 -c 'set session password="TestSubsystem";
delete logicalDrive ["Stocks_<Bonds"];
    create logicalDrive driveCount[3] RAIDLEVEL=3 capacity=10GB userLabel=
"Finance";
    show storageSubsystem healthStatus;'
```

This example shows how to run commands in a script file named scriptfile.scr on a storage subsystem named Example. The **-e** terminal causes the file to run without checking the syntax. Running a script file without checking the syntax lets the file run more quickly; however, the file might not run correctly because the syntax for a command might be incorrect.

SMcli -n Example -f scriptfile.scr -e

This example shows how to run commands in a script file named scriptfile.scr on a storage subsystem named Example. In this example, the storage subsystem is protected by the password MySubsystem. Output, as a result of commands in the script file, goes to file output.txt.

• Windows operating system:

```
SMcli -n Example -f scriptfile.scr -p "My_Subsystem" -o output.txt
```

UNIX operating system:

```
SMcli -n Example -f scriptfile.scr -p 'My_Subsystem' -o output.txt
```

This example shows how to show all of the storage subsystems in the current configuration. The command in this example returns the host name of each storage subsystem.

SMcli -d

If you want to know the IP address of each storage subsystem in the configuration, add the **-i** terminal to the command.

SMcli -d -i

## **Chapter 2. About the Script Commands**

You can use the script commands to configure and manage a storage subsystem. You can enter individual script commands, or you can run a file of script commands.

To invoke the Storage Manager Command-Line Interface (SMcli), use the SMcli command.

**Note:** When you enter an individual script command, you include it as part of a SMcli command. When you run a file of script commands, you include the file name as part of a SMcli command.

The script commands are processed by a script engine that performs the following functions:

- Verifies the command syntax
- Interprets the commands
- · Converts the commands to the appropriate protocol-compliant commands
- Passes the commands to the storage subsystem

At the storage subsystem level, the storage subsystem controllers run script commands.

The script engine and the script commands support the storage subsystem configuration and management operations that are listed in the following table:

Table 5. Configuration and Management Operations

Operation	Activities
Cache configuration	Controlling all cache parameters, both at the storage subsystem level and the individual logical drive level
Logical drive configuration, subsystem and disk pool configuration	Creating, deleting, and setting the reconstruction priority control; labeling; setting drive composition when creating logical drives; setting the segment size; and setting the media scan control
Disk drive configuration	Assigning hot spares
Controller configuration	Defining logical drive ownership, changing mode settings, defining network settings, and setting host channel IDs
General storage subsystem configuration	Resetting a configuration to defaults, labeling, checking the health status, setting the time of day, clearing the Event Log, and setting the media scan rate
NVSRAM configuration	Downloading and modifying the user configuration region at the bit level and the byte level, showing nonvolatile static random access memory (NVSRAM) values
Product identification	Retrieving the enclosure profile display data
Battery management	Setting the battery installation date

Table 5. Configuration and Management Operations (continued)

Operation	Activities
Firmware management	Downloading controller firmware, the environmental services monitor (ESM) firmware, and the disk drive firmware
Copy services	Configuring, changing settings and monitoring copy services functions such as VolumeCopy, Enhanced FlashCopy, and Enhanced Remote Mirroring.

#### Script Command Structure

All script commands have the following structure: *command operand-data* (*statement-data*)

where *command* identifies the action to be performed, *operand-data* represents the objects associated with a storage subsystem that you want to configure or manage, and *statement-data* provides the information needed to perform the command.

The syntax for *operand-data* is as follows:

```
(object-type | allobject-types | [qualifier] (object-type [identifier]
{object-type [identifier]} |
object-types [identifier-list]))
```

An object can be identified four ways: object type, **all** parameter prefix, square brackets, or a list of identifiers. Use an object type when the command is not referencing a specific object. The **all** parameter prefix means all objects of the specified type in the storage subsystem (for example, **allLogicalDrives**). To perform a command on a specific object, use square brackets to identify the object (for example, **logicalDrive [engineering]**). Specify a subset of objects with a list of identifiers in square brackets (for example, **logicalDrives [sales engineering marketing]**). A qualifier is required if you want to include additional information to describe the objects.

The object type and the identifiers that are associated with each object type:

Identifier
a or b
Enclosure ID and slot ID only, or enclosure ID, drawer ID, and slot ID, depending on the type of storage enclosure
Enclosure ID and slot ID only, or enclosure ID, drawer ID, and slot ID, depending on the type of storage enclosure
Drive channel identifier
User label
Host channel identifier
User label
User label

Table 6. Object Types and Identifiers

Object Type	Identifier	
iscsiInitiator	Logical drive user label or iSCSI Qualified Name (IQN)	
iscsiTarget	Logical drive user label or IQN	
remote mirror	Primary logical drive user label	
flashcopy	Logical drive user label	
storagesubsystem	Not applicable	
enclosure	Enclosure ID	
logicalDrive	Logical drive user label or logical drive World-Wide Identifier (WWID) ( <b>set</b> command only)	
VolumeCopy	Target logical drive user label and, optionally, the source logical drive user label	
subsystem	subsystem number or user label <b>Note:</b> A user label can have the following characters: alphanumeric characters, a hyphen, a pound sign, and an underscore character.	
disk pool	User label <b>Note:</b> A user label can have the following characters: alphanumeric characters, a hyphen, a pound sign, and an underscore.	

Table 6. Object Types and Identifiers (continued)

A user-defined entry (such as a user label) is called a variable. In the syntax, it is shown in *italic* (such as *enclosureD* or *subsystemName*).

Statement data is in the form of:

- *Parameter=value* (such as **raidLevel=5**)
- *Parameter-name* (such as **batteryInstallDate**)
- *Operation-name* (such as **redundancyCheck**)

## **Script Command Synopsis**

Because you can use the script commands to define and manage the different aspects of a storage subsystem (such as host topology, disk drive configuration, controller configuration, logical drive definitions, and subsystem definitions), the actual number of commands is extensive. The commands, however, fall into general categories that are reused when you apply the commands to the different aspects of a storage subsystem. The general form of the script commands and provides a definition of each command are listed in the following table:

Syntax	Description
activate object {statement-data}	Sets up the environment so that an operation can take place or performs the operation if the environment is already set up correctly.
autoConfigure storageSubsystem { <i>statement-data</i> }	Automatically creates a configuration that is based on the parameters that are specified in the command.

Syntax	Description
check object {statement-data}	Starts an operation to report on errors in the object, which is a synchronous operation.
clear object {statement-data}	Discards the contents of some attribute of an object. This operation is destructive and cannot be reversed.
<pre>create object {statement-data}</pre>	Creates an object of the specified type.
deactivate <i>object</i> { <i>statement-data</i> }	Removes the environment for an operation.
delete <i>object</i>	Deletes a previously created object.
diagnose <i>object</i> { <i>statement-data</i> }	Runs a test and shows the results.
disable object {statement-data}	Prevents a feature from operating.
download <i>object</i> { <i>statement-data</i> }	Transfers data to the storage subsystem or to the hardware that is associated with the storage subsystem.
<pre>enable object {statement-data}</pre>	Sets a feature operate.
recopy object {statement-data}	Restarts a VolumeCopy operation by using an existing VolumeCopy pair. You can change the parameters before the operation is restarted.
recover object {statement-data}	Re-creates an object from saved configuration data and the statement parameters. (This command is similar to the <b>create</b> command.)
re-create <i>object</i> { <i>statement-data</i> }	Restarts a FlashCopy operation by using an existing FlashCopy logical drive. You can change the parameters before the operation is restarted.
<pre>remove object {statement-data}</pre>	Removes a relationship from between objects.
repair object {statement-data}	Repairs errors found by the <b>check</b> command.
reset object {statement-data}	Returns the hardware or an object to an initial state.
resume <i>object</i>	Starts a suspended operation. The operation starts where it left off when it was suspended.
revive <i>object</i>	Forces the object from the Failed state to the Optimal state. Use this command only as part of an error recovery procedure.
save object {statement-data}	Writes information about the object to a file.
<pre>set object {statement-data}</pre>	Changes object attributes. All changes are completed when the command returns.
<pre>show object {statement-data}</pre>	Shows information about the object.

Table 7. General Form of the Script Commands (continued)
Syntax	Description
start object {statement-data}	Starts an asynchronous operation. You can stop some operations after they have started. You can query the progress of some operations.
<pre>stop object {statement-data}</pre>	Stops an asynchronous operation.
suspend object {statement-data}	Stops an operation. You can then restart the suspended operation, and it continues from the point where it was suspended.

Table 7. General Form of the Script Commands (continued)

## **Recurring Syntax Elements**

Recurring syntax elements are a general category of variables and options that you can use in one or more script commands. The recurring syntax is used in the general definitions of the script commands that are listed in Chapter 3, "Script Commands," on page 27

The recurring syntax and the syntax values that you can use with the syntax are listed in the following table:

**Note:** A command statement appears in a monospace font. The statement variables are in a monospace italic font. Options are described separately, and usually the choices are shown as TRUE or FALSE, as a letter (a), a number (42), or a number-range (0-99) selection.

Recurring Syntax	Syntax Value
raid-level	(0   1   3   5   6)
repository-raid-level	(1   3   5   6)
capacity-spec	integer-literal[KB   MB   GB   TB   Bytes]
segment-size-spec	integer-literal
boolean	(TRUE   FALSE)
	string-literal
user-label	Valid characters are alphanumeric, the dash, and the underscore.
user-label-list	user-label {user-label}
create-raid-logicalDrive-attr-value- list	create-raid-logicalDrive-attribute-value- pair {create-raid-logicalDrive-attribute- value-pair}

Table 8. Recurring Syntax Elements

Recurring Syntax	Syntax Value
create-raid-logicalDrive-attribute- value-pair	<pre>capacity=capacity-spec   owner=(a   b)   cacheReadPrefetch=(TRUE   FALSE)   segmentSize=integer-literal usageHint=usage-hint-spec</pre>
noncontroller-enclosureID	(0-99)
slotID	(1-32)
portID	(0-127)
	enclosureID,slotID or enclosureID, drawerID,slotID, depending on the type of storage enclosure
drive-spec	A drive is defined as two or three integer literal values separated by a comma. Low-density enclosures require two values. High-density enclosures, those enclosures that have disk drive drawers, require three values.
drive-spec-list	drive-spec {drive-spec}
enclosureID-list	enclosureID {enclosureID}
esm-spec-list	esm-spec {esm-spec}
esm-spec	enclosureID, (left right)
hex-literal	Øxhexadecimal-literal
subsystem-number	integer-literal
filename	string-literal
error-action	(stop   continue)
	(1   2   3   4   5   6  7   8)
drive-channel-identifier	<b>Note:</b> You might not be able to use all the supported values, depending on your storage subsystem model.
drive-channel-identifier-list	drive-channel-identifier {drive-channel-identifier}
	(a1   a2   a3   a4   a5   a6   a7   a8  b1   b2   b3   b4   b5   b6   b7   b8)
host-channel-identifier	<b>Note:</b> You might not be able to use all the supported values, depending on your storage subsystem model.
drive-type	(fibre   SATA   SAS)

 Table 8. Recurring Syntax Elements (continued)

Recurring Syntax	Syntax Value
	(HDD   SSD   unknown  allMedia)
drive-media-type	<i>HDD</i> means hard disk drive. <i>SSD</i> means solid state disk.
feature-identifier <sup>8</sup>	VolumeCopy   mixedDriveTypes   highPerformanceTier   SSDSupport   storagePartition   flashcopy   enhancedFlashCopyImage   driveSlotLimit   FullDiskEncryption   ExternalKeyManager   T10PI)
	To use the High Performance Tier premium feature, you must configure a storage subsystem as one of these: • SHIPPED_ENABLED • SHIPPED_ENABLED=FALSE; KEY_ENABLED=TRUE
repository-spec	instance-based-repository-spec   count-based-repository-spec
instance-based-repository-spec	<pre>repositoryRAIDLevel=repository-raid-level repositoryDrives=(drive-spec-list) [enclosureLossProtect=boolean<sup>1</sup>]</pre>
	<pre>[drawerLossProtect=(TRUE   FALSE)<sup>2</sup>])   repositorysubsystem=subsystem-number freeCapacityArea=integer-literal<sup>3</sup></pre>
	Specify <b>repositoryRAIDLevel</b> with <b>repositoryDrives</b> . Do not specify the RAID level or the disk drives with the subsystem. Do not set <b>enclosureLossProtect</b> when you specify a subsystem.
count-based-repository-spec	<pre>repositoryRAIDLevel=repository-raid-level repositoryDriveCount=integer-literal [repositorysubsystemUserLabel=user-label] [driveType=drive-type<sup>4</sup>] [enclosureLossProtect=boolean<sup>2</sup>]</pre>
	[drawerLossProtect=(TRUE   FALSE) <sup>2</sup> ]   [T10PI=(none   enabled) <sup>5</sup> ]
wwID	World-Wide Identifier
gid	Group Identifier
host-type	string-literal   integer-literal
host-card-identifier	

Table 8. Recurring Syntax Elements (continued)

Recurring Syntax	Syntax Value
backup-device-identifier	(1   n   a]])
	n is a specific slot number.
	Specifying all includes all of the cache backup devices available to the entire storage subsystem.
nvsram-offset	hex-literal
nvsram-byte-setting	nvsram-value-> 0xhexadecimal integer-literal
nvsram-bit-setting	nvsram-mask, nvsram-value-> 0xhexadecimal, 0xhexadecimal   integer-literal
inn/-address	Four groups of four numeric characters (0-9) separated by period "."
	(0255).(0255).(0255).
inthe address	Eight groups of four hexadecimal characters (0-9 and A-F) separated by semicolon ":"
tpv6-address	(02AF):(02AF):(02AF):(02AF): (02AF):(02AF):(02AF):(02AF)
autoconfigure-logicaldrive-attr- value-list	autoconfigure-logicaldrive-attr-value-pair {autoconfigure-logicaldrive-attr-value-pair}
autoconfigure-logicaldrive-attr- value-pair	<pre>driveType=drive-type   raidLevel=raid-level   subsystemWidth=integer-literal   subsystemCount=integer-literal   logicalDrivesPersubsystemCount=integer- literal <sup>6</sup>   hotSpareCount=integer-literal   segmentSize=segment-size-spec   cacheReadPrefetch=(TRUE   FALSE)   securityType=(none   capable   enabled)<sup>7</sup>   T10PI=(none   enabled)<sup>5</sup></pre>
create-logicalDrive-copy-attr-value- list	create-logicalDrive-copy-attr-value-pair {create-logicalDrive-copy-attr-value-pair}
create-logicalDrive-copy-attr-value- pair	<pre>copyPriority=(highest   high   medium   low   lowest)   targetReadOnlyEnabled=(TRUE   FALSE)   copyType=(offline   online)   repositoryPercentOfBase=(20   40   60   120   default)   repositoryGroupPreference=(sameAsSource   otherThanSource   default)</pre>
recover-raid-logicalDrive-attr- value-list	recover-raid-logicalDrive-attr-value-pair {recover-raid-logicalDrive-attr-value-pair}
recover-raid-logicalDrive-attr- value-pair	owner=(a   b)   cacheReadPrefetch=(TRUE   FALSE)   T10PI=(none   enabled)

Table 8. Recurring Syntax Elements (continued)

Recurring Syntax	Syntax Value
cache-flush-modifier-setting	immediate, 0, .25, .5, .75, 1, 1.5, 2, 5, 10, 20, 60, 120, 300, 1200, 3600, infinite
serial-number	string-literal
usage-hint-spec	usageHint=(multiMedia   database   fileSystem)
	[session-identifier]
iscsiSession	The host port number might be 2, 3, or 4 depending on the type of controller you are using.
iscsi-host-port	(1   2   3   4)
ethernet-port-options	<pre>enableIPv4=(TRUE   FALSE)   enableIPv6=(TRUE   FALSE)  IPv6LocalAddress=ipv6-address   IPv6RoutableAddress=ipv6-address   IPv6RouterAddress=ipv6-address   IPv4Address=ip-address   IPv4ConfigurationMethod=(static   dhcp)   IPv4GatewayIP=ip-address   IPv4SubnetMask=ip-address   duplexMode=(TRUE   FALSE)   portSpeed=(autoNegotiate   10   100   1000)</pre>
iscsi-host-port-options	IPv4Address= <i>ip-address</i>   IPv6RoutableAddress= <i>ipv6-address</i>   IPv6RoutableAddress= <i>ipv6-address</i>   enableIPv4= (TRUE   FALSE)   enableIPv6= (TRUE   FALSE)   enableIPv4Priority= (TRUE   FALSE)   enableIPv6Priority= (TRUE   FALSE)   IPv4ConfigurationMethod= (static   dhcp)   IPv6ConfigurationMethod= (static   auto)   IPv6ConfigurationMethod= (static   auto)   IPv6AdtewayIP= <i>ip-address</i>   IPv6NdDetectDuplicateAddress= <i>integer</i>   IPv6NdReachableTime= <i>time-interval</i>   IPv6NdRetransmitTime= <i>time-interval</i>   IPv4SubnetMask= <i>ip-address</i>   IPv4VlanId= <i>integer</i>   IPv6VlanId= <i>integer</i>   maxFramePayload= <i>integer</i>   10)
test-devices-list	controller=(a   b)   esms=(esm-speclist)   drives=(drive-spec-list)
test-devices	<pre>controller=(a   b)esms=(esm-spec- list)drives=(drive-spec-list)</pre>

Table 8. Recurring Syntax Elements (continued)

Recurring Syntax	Syntax Value
flashcopy-schedule-attribute-value- list	flashcopy-schedule-attribute-valuepair {flashcopy-schedule-attribute-valuepair}
time-zone-spec	(GMT+HH:MM   GMTHH:MM) [dayLightSaving=HH:MM]
flashcopy-schedule-attribute-value- pair	<pre>startDate=MM:DD:YY scheduleDay=(dayOfWeek all)startTime=HH:MM scheduleInterval=integer endDate=(MM:DD:YY noEndDate) timesPerDay=integer</pre>

Table 8. Recurring Syntax Elements (continued)

<sup>1</sup>For enclosure loss protection to work, each disk drive in an subsystem must be in a separate enclosure. If you set the **enclosureLossProtect** parameter to TRUE and you have selected more than one disk drive from any one enclosure, the storage subsystem returns an error. If you set **enclosureLossProtect** parameter to FALSE, the storage subsystem performs operations, but the subsystem that you create might not have enclosure loss protection.

If you set the **enclosureLossProtect** parameter to TRUE, the storage subsystem returns an error if the controller firmware cannot find drives that will enable the new subsystem to have enclosure loss protection. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs the operation even if it means that the subsystem might not have enclosure loss protection.

<sup>2</sup> In the enclosures that have drawers for holding the drives, such as EXP5060, DCS3700, DCS3700 storage subsystem with Performance Module Controllers, DCS3860, drawer loss protection determines whether data on a logical drive is accessible or inaccessible if a drawer fails. To help make sure that your data is accessible, set the **drawerLossProtect** parameter to TRUE. For drawer loss protection to work, each drive in an subsystem must be in separate drawers. If you have a storage subsystem configuration in which an subsystem spans several enclosures, you must make sure that the setting for drawer loss protection works with the setting for enclosure loss protection. If you set the **enclosureLossProtect** parameter to TRUE, you must set the **drawerLossProtect** parameter to FALSE, the storage subsystem returns an error message and a storage subsystem configuration will not be created.

<sup>3</sup>To determine if a free capacity area exists, run the **show subsystem** command.

<sup>4</sup>The default disk drive (drive type) for DS3000 is SAS. The default disk drive (drive type) for DS4000/DS5000 is fibre (Fibre Channel).

If you set the **enclosureLossProtect** parameter to TRUE, the storage subsystem returns an error if the controller firmware cannot find disk drives that will enable the new subsystem to have enclosure loss protection. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs the operation even if it means that the subsystem might not have enclosure loss protection.

<sup>5</sup>The **T10PI** parameter applies to the drives in an subsystem. Using the **T10PI** parameter, you can specify that protected drives must be selected for an

subsystem. If you want to set the **T10PI** parameter to **enabled**, all of the drives in the subsystem must be capable of T10 PI (Protection Information). You cannot have a mix of drives that are capable of T10 PI (Protection Information) and drives that are not capable of T10 PI (Protection Information) in the subsystem.

<sup>6</sup>The **logicalDrivesPersubsystemCount** parameter is the number of equal-capacity logical drives per subsystem.

<sup>7</sup>The **securityType** parameter enables you to specify the security setting for an subsystem that you are creating. All of the subsystems are also set to the security setting that you choose. Available options for setting the security setting include:

- **none** The subsystem is not secure.
- **capable** The subsystem is security capable, but security has not been enabled.
- **enabled** The subsystem is security enabled.

**Note:** A storage subsystem security key must already be created for the storage subsystem if you want to set **securityType=enabled**. (To create a storage subsystem security key, use the create **storageSubsystemsecurityKey** command).

<sup>8</sup>In certain versions of Storage Manager, **FullDiskEncryption** was referred to as **safestoreSecurity** and **ExternalKeyManager** was referred to as **safeStoreExtenalKeyMgr**. Depending on the version of the controller firmware, you might have to use the following terms instead of the ones listed here

```
(storagePartition2 | storagePartition4 | storagePartition8 |
storagePartition16 | storagePartition64 | storagePartition96 |
storagePartition128 | storagePartition256 | storagePartition512 |
storagePartitionMax | flashcopy | flashcopy2 | flashcopy4 | flashcopy8 |
flashcopy16 | remoteMirror8 | remoteMirror16 | remoteMirror32 |
remoteMirror64 | remoteMirror128 |.
```

## Usage Guidelines

This list provides guidelines for writing script commands on the command line:

- You must end all commands with a semicolon (;).
- You can enter more than one command on a line, but you must separate each command with a semicolon (;).
- You must separate each base command and its associated primary parameters and secondary parameters with a space.
- The script engine is not case sensitive. You can enter commands by using uppercase letters, lowercase letters, or mixed-case letters.
- Add comments to your scripts to make it easier for you and future users to understand the purpose of the script commands. (For information about how to add comments, see ""Adding Comments to a Script File.")

**Note:** While the CLI commands and the script commands are not case sensitive, user labels (such as for logical drives, hosts, or host ports) are case sensitive. If you try to map to an object that is identified by a user label, you must enter the user label exactly as it is defined, or the CLI commands and the script commands will fail.

## Adding Comments to a Script File

The script engine looks for certain characters or a command to show comments. You can add comments to a script file in three ways: 1. Add text after two forward slashes (//) as a comment until an end-of-line character is reached. If the script engine does not find an end-of-line character in the script after processing a comment, an error message appears, and the script operation is terminated. This error usually occurs when a comment is placed at the end of a script and you have forgotten to press the **Enter** key.

// Deletes the existing configuration.
set storageSubsystem resetConfiguration=true;

2. Add text between / \* and \* / as a comment. If the script engine does not find both a starting comment notation and an ending comment notation, an error message appears, and the script operation is terminated.

/\* Deletes the existing configuration \*/
set storageSubsystem resetConfiguration=true;

**3**. Use the **show** statement to embed comments in a script file that you want to appear while the script file is running. Enclose the text that you want to appear by using double quotation marks ("").

show "Deletes the existing configuration"; set storageSubsystem resetConfiguration=true;

## **Chapter 3. Script Commands**

Attention: The script commands are capable of damaging a configuration and causing loss of data access if not used correctly – Command operations are performed as soon as you run the commands. Some commands can immediately delete configurations or data. Before using the script commands, make sure that you have backed up all data, and have saved the current configuration so that you can reinstall it if the changes you make do not work.

**IBM recommends using the Storage Manager client GUI to manage your storage subsystems.** – The command line interface does not have any mechanisms to prevent you from inadvertently making unwanted changes to the storage subsystem; therefore, IBM recommends using the Storage Manager client GUI to manage your storage subsystem configurations.

This chapter has five sections to help you use script commands:

- "Naming conventions" lists the general formatting rules for entering the names of storage subsystem entities, such as logical drives or disk drives, with the script commands.
- "Firmware Compatibility Levels" on page 28 describes how to interpret the firmware level information.
- "Formatting rules for script commands" on page 29 lists the general formatting rules that apply to the script command syntax.
- "DS3000, DS4000, DS5000, and the DCS Series commands with minimum firmware levels" on page 30 indicates which commands you can use for DS3000, DS4000, DS5000, or the DCS Series with their minimum controller firmware requirements.
- "Script Commands Listed by Function" on page 373 lists the script commands organized into groups related to the physical features, the logical features, and the operational features of the storage subsystem.
- Finally, the script commands are listed alphabetically with detailed information including the command name, syntax, and parameters.

## Naming conventions

- Names can have a maximum of 30 characters.
- You can use any combination of alphanumeric characters, hyphens, and underscores for the names of the following components:
  - Storage subsystems
  - Host groups
  - Hosts
  - subsystems
  - Logical drives
  - HBA host ports
- You must use unique names. If you do not use unique names, the controller firmware returns an error.
- If the name contains more than one word, hyphens, or underscores, enclose the name in double quotation marks (""). In some usages, you must also surround

the name with square brackets ([]). The description of each parameter indicates whether you need to enclose a parameter in double quotation marks, square brackets, or both.

- The name character string cannot contain a new line.
- On Windows operating systems, you must enclose the name between two back slashes (\) in addition to other delimiters. For example, the following name is used in a command that runs under a Windows operating system: [\"Engineering\"]
- For a UNIX operating system and, when used in a script file, the name appears as the following example:

["Engineering"]

• When you enter a World-Wide Identifier (WWID) of an HBA host port, some usages require that you surround the WWID with double quotation marks. In other uses, you must surround the WWID with angle brackets (<>). The description of the WWID parameter indicates whether you need to enclose the WWID in double quotation marks ("") or angle brackets.

## Entering numerical names

When the storage management software automatically configures a storage subsystem, the storage management software assigns names that consist of numerical characters. Names that consist only of numerical characters are valid names. Numerical character names, however, must be treated differently than names that start with alphabetic characters.

When you enter a script command that requires a name, the script engine looks for a name that starts with an alphabetic character. The script engine might not recognize the following names:

- Names that are only numbers, such as 1 or 2
- Names that start with a number, such as 1Disk or 32logical drive

To enter a name that consists only of numerical characters so that the script engine will recognize the name, use a combination of back slashes and double quotation marks. Here are a few examples of how you can enter names that consist only of numerical characters or start with numerical characters:

- [\"1\"]
- [\"1Disk\"]

## Firmware Compatibility Levels

Some of the script commands and the command parameters are not supported with every version of the controller firmware. Minimum controller firmware levels are listed in "DS3000, DS4000, DS5000, and the DCS Series commands with minimum firmware levels" on page 30.

**Note:** For many commands, you can enable additional parameters by upgrading to a higher firmware level than the minimum.

## Hardware Requirements for Firmware Compatibility

Even though a command is supported by a storage subsystem, the feature must be installed to use the command. For example, the Full Disk Encryption (FDE) commands are supported for all DS5000 models, but you cannot use the command unless the FDE premium feature is installed in the storage subsystem.

Not all models in the DS3000, DS4000, DS5000, or the DCS Series product lines support the same set of functionality. For example, the DS3500 series of the DS3000 product line supports Enhanced Remote Mirroring premium features, but the DS3200, DS3300, and DS3400 models of the DS3000 product line do not. You can only run the commands for the functions supported by your storage subsystem model.

## Formatting rules for script commands

Syntax unique to a specific script command is explained in the Notes section at the end of each script command description.

**Case sensitivity** – The script commands are not case sensitive. You can type the script commands in lowercase, uppercase, or mixed case. (mixed case is used as an aid to reading the following command names and understanding the purpose of the command).

**Spaces** – You must enter spaces in the script commands as they are shown in the command descriptions.

Square brackets – Square brackets are used in two ways:

- As part of the command syntax
- To indicate that the parameters are optional. The description of each parameter tells you if you need to enclose a parameter value in square brackets.

**Parentheses** – Parentheses shown in the command syntax enclose specific choices for a parameter. That is, if you want to use the parameter, you must enter one of the values enclosed in parentheses. Generally, you do not include parentheses in a script command; however, in some instances, when you enter lists, you must enclose the list in parentheses. Such a list might be a list of enclosure ID values and slot ID values. The description of each parameter tells you if you need to enclose a parameter value in parentheses.

**Vertical bars** – Vertical bars in a script command indicate "or" and separate the valid values for the parameter. For example, the syntax for the **raidLevel** parameter in the command description appears as follows: raidLevel=(0 | 1 | 3 | 5 | 6)

To use the **raidLevel** parameter to set a RAID level of 5, enter this value: raidLevel=5

**Drive locations** – When you specify disk drive locations by using enclosure ID values and slot ID values, separate the ID values with a comma. If you enter more than one set of ID values, separate each set of values with a space. Enclose the set of values in parentheses. For example:

(1,1 1,2 1,3 1,4 2,1 2,2 2,3 2,4)

For storage enclosures with multiple hard drive drawers, you specify drive locations by using the combination of enclosure ID, drawer ID, and slot ID values, which are separated by commas. For example:

(1,1,1 1,2,1 1,3,1 1,4,1 1,5,1 1,2,12 2,1,12 2,2,10)

**Italicized terms** – Italicized terms in the command indicate a value or information that you need to provide. For example, when you encounter the italicized term: *numberOfDrives* 

replace the italicized term with a value for the number of disk drives that you want to include with the script command.

**Semicolon** – Script commands must end with a semicolon (;). You can enter more than one script command on the command line each time you enter a CLI command.

# DS3000, DS4000, DS5000, and the DCS Series commands with minimum firmware levels

#### Alphabetic command list

The DS3000, DS4000, DS5000, and the DCS Series products share many common Storage Manager script commands. Some commands are specific to one of the products, two of the products, or all the products and require a particular minimum level of controller firmware.

DCS Series (DCS3700 and DCS3700 storage subsystem with Performance Module Controllers, and the DCS3860 storage subsystem) are part of the DS3000 series products. Not all commands that are supported by DS3000 can be used in the DS3200, DS3300, and DS3400 storage subsystems. Some commands are supported only with the DS3500, DCS3700, DCS3700 storage subsystem with Performance Module Controllers, and the DCS3860 storage subsystem.

Commands with minimum controller firmware requirements are alphabetically listed below:

#### Activate Enhanced Remote Mirroring Feature

DS3000: 7.7x.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Note:

- 7.10 adds RAID Level 6 capability
- 7.60 adds the drawerID user input, the **driveMediaType** parameter, and the **drawerLossProtect** parameter.
- For version 7.60.xx.xx, you can change the **repositoryDrive** parameter to accommodate the enclosures with internal disk drawers.
- 7.77 adds the T10 PI parameter.

#### Activate Host Port

DS3000: 7.35.xx.xx

DS4000: N/A

DS5000: 7.6x.xx.xx

### Activate iSCSI Initiator

DS3000: 7.35.xx.xx

DS4000: N/A

DS5000: 7.6x.xx.xx

## Activate Storage Subsystem Firmware

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Add Member to Consistency Group

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Autoconfigure Storage Subsystem

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Note:

- 7.50.xx.xx adds securityType parameter.
- 7.77 adds the **T10 PI** parameter.

## Autoconfigure Storage Subsystem Hot Spares

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

Check Logical Drive Parity DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### **Check Remote Mirror Status**

DS3000: 7.7x.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

Check Repository Consistency DS3000: 7.83.xx.xx

DS4000: N/A

DS5000: N/A

## **Clear Drive Channel Statistics**

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

Note: 6.14.xx.xx adds discreteLines parameter on DS4800.

#### **Clear Logical Drive Reservations**

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

#### Clear Logical Drive Unreadable Sectors DS3000: 6.17.xx.xx

-----

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

## **Clear Storage Subsystem Configuration**

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

Note: 7.1x.xx.xx and higher adds the all and subsystems parameters.

#### Clear Storage Subsystem Diagnostic Data Capture

DS3000: 7.35.xx.xx

DS4000: 6.6x.xx.xx

DS5000: 7.3x.xx.xx

## Clear Storage Subsystem Event Log

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Clear Storage Subsystem Firmware Pending Area

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Clear Storage Subsystem Recovery Mode

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Configure Automatic Support Bundle Collection

DS3000: 7.xx.xx.xx

DS4000: 7.xx.xx.xx

DS5000: 7.3x.xx.xx

**Note:** This command is only compatible with Storage Manager version 10.83 and later.

**Convert FlashCopy**<sup>®</sup> **Logical Drive to Enhanced FlashCopy Group** DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Create subsystem

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Note:

- For version 7.60.xx.xx, you can change the drive parameter to accommodate the enclosures with internal disk drawers.
- 7.77 adds the T10 PI parameter.

### **Create Consistency Group**

DS3000: 7.8x.xx.xx

DS4000: N/A

DS5000: N/A

#### Create Consistency Group Enhanced FlashCopy Image DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

Create Consistency Group Enhanced FlashCopy Logical Drive DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

### Create Consistency Group Enhanced FlashCopy Logical Drive Mapping DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### **Create Disk Pool**

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Create Enhanced FlashCopy Image

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Create Enhanced FlashCopy Group DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Create FlashCopy Logical Drive

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

#### Note:

- 7.1x.xx.xx and higher adds RAID 6 parameter.
- For version 7.60.xx.xx, you can change the repository drive parameter to accommodate the enclosures with internal disk drawers.

#### Create Enhanced FlashCopy Logical Drive

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### **Create Host**

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

#### Create Host Group

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

#### Create Host Port

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

#### Create iSCSI Initiator

DS3000: 6.50.xx.xx

DS4000: Not supported

DS5000: 7.6x.xx.xx

**Note:** 7.10.xx.xx deprecates the **hostType** parameter.

#### Create RAID Logical Drive (Automatic Drive Select)

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

#### Note:

- 7.1x.xx.xx and higher adds **RAID 6** parameter.
- 7.50.xx.xx adds **securityType** parameter.
- 7.60.xx.xx adds DriveMediaType parameter.

## Create RAID Logical Drive (Free Capacity-based Select)

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

#### Note:

- 7.1x.xx.xx and higher adds **RAID 6** parameter.
- 7.50.xx.xx adds securityType parameter.

• 7.77.xx.xx and higher adds **T10PI** parameter.

#### Create RAID Logical Drive (Manual Drive Select)

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

## Note:

- 7.1x.xx.xx and higher adds RAID 6 parameter
- 7.50.xx.xx adds the **securityType** parameter.
- For version 7.60.xx.xx, you can change the Drives parameter to accommodate the enclosures with internal disk drawers.
- 7.77 adds the T10 PI parameter.

## Create Logical Drive on Disk Pool

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Create Read-Only Enhanced FlashCopy Logical Drive

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## **Create Remote Mirror**

DS3000: 7.7x.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

Note: 7.1x.xx.xx and higher adds **RAID 6** parameter.

### Create Storage Subsystem Security Key

DS3000: 7.7x.xx.xx

DS4000: Not supported

DS5000: 7.3x.xx.xx

#### Create VolumeCopy

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

Note: 7.1x.xx.xx and higher adds RAID 6 parameter.

#### Deactivate Remote Mirror DS3000: 7.7x.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

Note: 7.1x.xx.xx and higher adds RAID 6 parameter.

#### Delete subsystem

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Delete Consistency Group

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Delete Consistency Group Enhanced FlashCopy Image DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

### Delete Consistency Group Enhanced FlashCopy Logical Drive DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## **Delete Disk Pool**

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Delete Enhanced FlashCopy Group DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Delete Enhanced FlashCopy Image DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Delete FlashCopy Logical Drive DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Delete Enhanced FlashCopy Logical Drive DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### **Delete Host**

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

## Delete Host Group

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

Note: 7.1x.xx.xx and higher adds RAID 6 parameter.

## **Delete Host Port**

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

Note: 7.1x.xx.xx and higher adds RAID 6 parameter.

#### Delete iSCSI Initiator

DS3000: 6.5x.xx.xx

DS4000: Not supported

DS5000: 7.6x.xx.xx

#### Delete Logical Drive

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

Note: 7.1x.xx.xx adds the removesubsystem parameter

### Delete Logical Drive on a Disk Pool

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## **Diagnose Controller**

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

## Note:

- 6.1x.xx.xx adds parameters:
  - readTest
  - writeTest
  - dataLoopBackTest
- 6.14.xx.xx adds the discreteLinesDiagnosticTest parameter.

#### **Diagnose Remote Mirror**

DS3000: 7.7x.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Diagnose Controller iSCSI Host Cable

DS3000: 7.77.xx.xx

DS4000: N/a

DS5000: 7.77.xx.xx

### Disable External Security Key Management

DS3000: N/A

DS4000: N/A

DS5000: 7.70.xx.xx

#### Disable FlashCopy

DS3000: 7.83.xx.xx

DS4000: N/A

DS5000: N/A

## Disable Storage Subsystem Feature

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Display Automatic Support Bundle Collection Configuration

DS3000: 7.xx.xx.xx

DS4000: 7.xx.xx.xx

DS5000: 7.3x.xx.xx

**Note:** This command is only compatible with Storage Manager version 10.83 and later.

## **Display Support Bundle Collection Schedule**

DS3000: 7.xx.xx.xx

DS4000: 7.xx.xx.xx

DS5000: 7.3x.xx.xx

**Note:** This command is only compatible with Storage Manager version 10.83 and later.

#### **Disable Storage Subsystem Remote Status Notification**

DS3000: 7.77.xx.xx

DS4000: N/A

DS5000: 7.77.xx.xx

## Download Drive Firmware

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

**Note:** For version 7.60.xx.xx, you can change the drive parameter to accommodate the enclosures with internal disk drawers.

### **Download Enclosure Configuration Settings**

DS3000: 7.70.xx.xx

DS4000: 7.60.xx.xx

DS5000: 7.60.xx.xx

### Download Environmental Card Firmware DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

## Download Power Supply Firmware

DS3000: 7.77.xx.xx

DS4000: N/A

DS5000: 7.77.xx.xx

#### Download Storage Subsystem Drive Firmware DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

## Download Storage Subsystem Firmware

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

#### Download Storage Subsystem NVSRAM DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

## Enable subsystem Security

DS3000: Not supported

DS4000: Not supported

DS5000: 7.5x.xx.xx

## Enable Controller Data Transfer DS3000: 7.83.xx.xx

DS4000: N/A

DS5000: N/A

## Enable Disk Pool Security

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## **Enable Controller**

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Enable External Security Key Management

DS3000: N/A

DS4000: N/A

DS5000: 7.70.xx.xx

## Enable Storage Subsystem Feature Key DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

Note:

- 6.5x.xx.xx adds the featurePack parameter.
- 7.5x.xx.xx adds the highPerformanceTier parameter.
- Enable Storage Subsystem Remote Status Notification DS3000: 7.77.xx.xx

DS4000: N/A

DS5000: 7.77.xx.xx

Export Storage Subsystem Security Key DS3000: 7.7x.xx.xx

DS4000: Not supported

DS5000: 7.5x.xx.xx

Import Storage Subsystem Security Key DS3000: 7.7x.xx.xx

DS4000: Not supported

DS5000: 7.5x.xx.xx

Increase Logical Drive Capacity on a Disk Pool DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

Initialize Thin Logical Drive DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Load Storage Subsytem DBM Database DS3000: 7.77.xx.xx

D00000. 7.77.XX.

DS4000: N/A

DS5000: 7.77.xx.xx

## Recopy VolumeCopy

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Recover RAID Logical Drive DS3000: 6.17.xx.xx

-----

DS4000: 5.43.xx.xx

DS5000: 7.3x.xx.xx

#### Note:

- For version 7.60.xx.xx, you can change the drive/drives parameter to accommodate the enclosures with internal disk drawers.
- 7.77 adds the T10 PI parameter.

## Re-create Enhanced Remote Mirroring Repository Logical Drive

DS3000: 7.7x.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Note:

- For version 7.60.xx.xx, you can change the repositoryDrive parameter to accommodate the enclosures with internal disk drawers.
- 7.77 adds the T10 PI parameter.

#### **Re-create External Security Key**

DS3000: N/A

DS4000: N/A

DS5000: 7.70.xx.xx

#### Re-create FlashCopy

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

## **Remove Logical Drive LUN Mapping**

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Remove Member Logical Drive from Consistency Group DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### **Remove Remote Mirror**

DS3000: 7.7x.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Remove VolumeCopy

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

#### Rename Enhanced FlashCopy Logical Drive DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Repair Logical Drive Parity DS3000: 6.17.xx.xx

D33000. 0.17.XX.XX

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### **Replace** Drive

DS3000: 6.17.xx.xx

DS4000: 7.1x.xx.xx

DS5000: 7.3x.xx.xx

**Note:** For version 7.60.xx.xx, you can change the drive parameter to accommodate the enclosures with internal disk drawers

#### **Reset Controller**

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

#### Reset Storage Subsystem SAS PHY Baseline

DS3000: 7.77.xx.xx

DS4000: N/A

DS5000: 7.77.xx.xx

### Reset Storage Subsystem Battery Install Date

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

**Note:** 7.3x.xx.xx and higher adds the ability to reset the battery installation dates on specific batteries that are inside the enclosure.

#### Reset Storage Subsystem Diagnostic Data

DS3000: 6.17.xx.xx

DS4000: 6.16.xx.xx

DS5000: 7.3x.xx.xx

#### Reset Storage Subsystem iSCSI Baseline

DS3000: 6.50.xx.xx

DS4000: Not supported

DS5000: 7.6x.xx.xx

#### Reset Storage Subsystem Logical Drive Distribution

DS4000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

#### DS5000: 7.3x.xx.xx

## Reset Storage Subsystem RLS Baseline

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

#### DS5000: 7.3x.xx.xx

#### Reset Storage Subsystem SOC Baseline DS3000: Not supported

#### DS4000: 6.16.xx.xx

DS5000: 7.3x.xx.xx

#### Resume Consistency Group Enhanced FlashCopy Logical Drive DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## **Resume Remote Mirror**

DS3000: 7.7x.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Resume Enhanced FlashCopy Image Rollback DS3000: 7.83.xx.xx and later

DS4000: 7.83.xx.xx and later

DS5000: 7.83.xx.xx and later

## Resume FlashCopy Rollback

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Resume Enhanced FlashCopy Logical Drive DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

### Revive subsystem

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### **Revive Drive**

DS3000: 6.17.xx.xx

DS4000: 5.43.xx.xx

DS5000: 7.3x.xx.xx

**Note:** For version 7.60.xx.xx, you can change the drive parameter to accommodate the enclosures with internal disk drawers.

## Revive Enhanced FlashCopy Group

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Revive Enhanced FlashCopy Logical Drive DS3000: 7.83.xx.xx and later

DS4000: N/A

. . . . . ,

## DS5000: N/A

## Save Controller NVSRAM

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Save Drive Channel Fault Isolation Diagnostic Status

DS3000: Not supported

DS4000: Not supported

DS5000: 7.3x.xx.xx

#### Save Drive Log

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Save Enclosure Log

DS3000: 7.35.xx.xx

DS4000: 6.60.xx.xx

#### DS5000: 7.xx.xx.xx

## Save Storage Subsystem Configuration

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Save Storage Subsystem Diagnostic Data Capture DS3000: 7.35.xx.xx

DS4000: 6.6x.xx.xx

DS5000: 7.3x.xx.xx

## Save Storage Subsystem DBM Database

DS3000: 7.77.xx.xx

#### DS4000: N/A

DS5000: 7.77.xx.xx

### Save Storage Subsystem Diagnostic Data

DS3000: 6.17.xx.xx

DS4000: 6.16.xx.xx

#### DS5000: 7.3x.xx.xx

## Save Storage Subsystem Events

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

#### DS5000: 7.3x.xx.xx

#### Save Storage Subsystem Firmware Inventory DS3000: 7.70.xx.xx

D00000. 7.7 0.7.7.

## DS4000: N/A

#### DS5000: 7.70.xx.xx

#### Save Storage Subsystem iSCSI Statistics DS3000: 6.50.xx.xx

#### DS4000: Not supported

DS5000: 7.6x.xx.xx

## Save Storage Subsystem Performance Statistics

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Save Storage Subsystem RLS Counts

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Save Storage Subsystem SAS PHYCounts DS3000: 7.77.xx.xx

DS4000: N/A

DS5000: 7.77.xx.xx

## Save Storage Subsystem SOC Counts

DS3000: Not supported

DS4000: 6.16.xx.xx

DS5000: 7.3x.xx.xx

#### Save Storage Subsystem State Capture DS3000: 6.17.xx.xx

\_\_\_\_

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Save Storage Subsystem Support Data

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Save Storage Subsystem Validator

DS3000: 7.77.xx.xx

DS4000: N/A

DS5000: 7.77.xx.xx

#### Schedule Support Bundle Collection Configuration

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

**Note:** This command is only compatible with Storage Manager version 10.83 and later.

#### Set subsystem

DS3000: 6.17.xx.xx

DS4000: 6.10.xx.xx

DS5000: Not supported

## Note:

• RAID 6 is supported with 7.xx.xx.and higher.

- 7.30.xx.xx removes the availability parameter.
- For version 7.60.xx.xx, you can change the **addDrive** parameter to accommodate the enclosures with internal disk drawers.

## Set subsystem Forced State

DS3000: 6.17.xx.xx

DS4000: 7.1x.xx.xx

DS5000: 7.3x.xx.xx

## Set Consistency Group Attributes

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Set Consistency Group Enhanced FlashCopy Logical Drive

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Set Controller

DS3000: 6.17.xx.xx

DS4000: 6.14.xx.xx

DS5000: 7.3x.xx.xx

#### Note:

- 7.30.xx.xx removed the **bootp** parameter and added the new Ethernet port options and the new iSCSI host-port options.
- 7.50.xx.xx moved the **IPV4Gateway** parameter and the **IPV6RouterAddress** parameter from the iSCSI host port options, to the command.

#### Set Controller Service Action Allowed Indicator

DS3000: 6.17.xx.xx

DS4000: 6.16.xx.xx

DS5000: 7.3x.xx.xx

#### Set Disk Pool

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Set Disk Pool (Modify Disk Pool)

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Set Drawer Service Action Allow Indicator

DS3000: N/A

DS4000: N/A

DS5000: 7.60.xx.xx

## Set Drive Channel Status

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

## Set Drive Hot Spare

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

### Set Disk Drive Service Action Allow Indicator

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

#### Set Drive State

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

**Note:** For version 7.60.xx.xx, you can change the drive parameter to accommodate the enclosures with internal disk drawers.

## Set Enclosure Alarm

DS3000: 6.17.xx.xx

DS4000: 6.16.xx.xx

DS5000: 7.3x.xx.xx

## Set Enclosure Identification

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Note:

- 6.14.xx.xx adds support for DS4800.
- 6.16.xx.xx adds support for subsystems and enclosures that set enclosure IDs through the controller firmware.

## Set Enclosure Service Action Allowed Indicator

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Note:

- 6.14.xx.xx adds the **powerfan** and **interconnect** parameters.
- 6.16.xx.xx adds the enclosure and esm parameters.

### Set Enhanced FlashCopy Group Attributes

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Set Enhanced FlashCopy Group Media Scan

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Set Enhanced FlashCopy Group Repository Logical Drive Capacity DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Set Enhanced FlashCopy Group Schedule

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Set FlashCopy Logical Drive

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Set Enhanced FlashCopy Logical Drive Media Scan

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

# Set Enhanced FlashCopy Logical Drive Repository Logical Drive Capacity DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Set Foreign Drive to Native

DS3000: 7.xx.xx.xx

DS4000: 7.xx.xx.xx

DS5000: 7.3x.xx.xx

**Note:** For version 7.60.xx.xx, you can change the drive parameter to accommodate the enclosures with internal disk drawers.

#### Set Host

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Set Host Channel

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Set Host Group

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Set Host Port

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Set iSCSI Initiator

DS3000: 6.50.xx.xx

DS4000: Not supported

DS5000: 7.6x.xx.xx

## Set iSCSI Target Properties

DS3000: 6.50.xx.xx

DS4000: Not supported

DS5000: 7.6x.xx.xx

### Set Logical Drive or Set Logical Drive Attributes

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

#### Note:

- 7.1.xx.xx adds the **preReadRedundancyCheck** parameter.
- For version 7.60.xx.xx, you can change the **addDrive** parameter to accommodate the enclosures with internal disk drawers.

## Set Logical Drive Attributes for a Disk Pool

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Set Logical Drive Mapping

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

## Set read-only to a Read/Write Enhanced FlashCopy Logical Drive

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Set Remote Mirror

DS3000: 7.7x.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Set Session

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

#### Set Storage Subsystem

DS3000: 6.17.xx.xx

DS4000: 5.2x.xx.xx

DS5000: 7.3x.xx.xx

## Note:

- 5.4x.xx.xx adds the failoverAlertDelay parameter.
- 6.1x.xx.xx adds the **alarm** parameter.

#### Set Storage Subsystem Enclosure Positions

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

Note: For 6.14.xx.xx and 6.16.xx.xx, controller is not a valid value.

## Set Storage Subsystem ICMP Response

DS3000: 6.50.xx.xx (DS3300 only)

DS4000: 7.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Set Storage Subsystem iSNS Server IPv4 Address DS3000: 6.50.xx.xx

DS4000: N/A

DS5000: 7.6x.xx.xx

#### Set Storage Subsystem iSNS Server IPv6 Address DS3000: 6.50.xx.xx

DS4000: N/A

#### DS5000: 7.6x.xx.xx

Set Storage Subsystem iSNS Server Listening Port DS3000: 6.50.xx.xx

DS4000: N/A

DS5000: 7.6x.xx.xx

#### Set Storage Subsystem iSNS Server Refresh DS3000: 6.50.xx.xx

DS4000: N/A

#### DS5000: 7.6x.xx.xx

**Note:** This command is replaced by Start Storage Subsystem iSNS Server Refresh in controller.

## Set Storage Subsystem Learn Cycle

DS3000: 7.35.xx.xx and higher

DS4000: N/A

DS5000: 7.6x.xx.xx

## Set Storage Subsystem Redundancy Mode

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Set Storage Subsystem Remote Status Notification Time DS3000: 7.77.xx.xx

D00000. 7.77.XX.X/

DS4000: N/A

DS5000: 7.7x.xx.xx

## Set Storage Subsystem Security Key

DS3000: 7.7x.xx.xx

DS4000: Not supported

DS5000: 7.5x.xx.xx

#### Set Storage Subsystem Time

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

## Set Storage Subsystem Unnamed Discovery Session

DS3000: 6.50.xx.xx

DS4000: 7.1x.xx.xx

DS5000: 7.3x.xx.xx

## Set Thin Logical Drive Attributes

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Set VolumeCopy

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

#### Show subsystem

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Show subsystem Export Dependencies DS3000: 7.35.xx.xx and higher

DS4000: 7.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Show subsystem Import Dependencies DS3000: 7.35.xx.xx and higher

DS4000: 7.1x.xx.xx

DS5000: 7.3x.xx.xx

## Show Cache Backup Device Diagnostic Status

DS3000: N/A

DS4000: N/A

DS5000: 7.70.xx.xx (DS5000 only)

Show Cache Memory Diagnostic Status DS3000: N/A

DS4000: N/A

DS5000: 7.70.xx.xx (DS5000 only)

#### Show Consistency Group Enhanced FlashCopy Image DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Show Controller

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

Note: 5.43.x.x.xx adds the summary parameter.

#### Show Controller Diagnostic Status

DS3000: 7.70.xx.xx

DS4000: N/A

DS5000: 7.70.xx.xx

## Show Controller NVSRAM

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Show Current iSCSI Sessions

DS3000: 6.50.xx.xx

DS4000: Not supported

DS5000: 7.6x.xx.xx

#### Show Disk Pool

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### **Show Drive**

DS3000: 6.17.xx.xx

DS4000: 5.43.xx.xx

#### DS5000: 7.3x.xx.xx

#### Show Drive Channel Stats DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Show Drive Download Progress

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Show Enhanced Remote Mirroring Logical Drive Candidates

DS3000: 7.7x.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

# Show Enhanced Remote Mirroring Logical Drive Synchronization Progress DS3000: 7.7x.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

### Show Enhanced FlashCopy Group

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

## Show Enhanced FlashCopy Image

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Show Enhanced FlashCopy Logical Drives DS3000: 7.83.xx.xx and later

DS4000: N/A

#### DS5000: N/A

#### Show Host Interface Card Diagnostic Status DS3000: 7.70.xx.xx

DS4000: N/A

DS5000: 7.70.xx.xx

#### Show Host Ports

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

#### Show Logical Drive

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

#### Note: 5.43.xx.xx adds the **summary** parameter.

## Show Logical Drive Action Progress DS3000: 6.17.xx.xx

DS4000: 5.43.xx.xx

DS5000: 7.3x.xx.xx

#### Show Logical Drive Performance Statistics DS3000: 6.17.xx.xx

D33000. 0.17.XX.XX

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

## Show Logical Drive Reservations

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

#### Show Storage Subsystem

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

**Note:** 5.43.xx.xx adds the **summary** parameter. 6.14.xx.xx adds the **connections** parameter.

### Show Storage Subsystem Auto Configure

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Show Storage Subsystem Diagnostic Data Capture DS3000: 7.35.xx.xx

DS4000: 6.6x.xx.xx

DS5000: 7.3x.xx.xx

#### Show Storage Subsystem DBM Database

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

#### Show Storage Subsystem Features DS3000: 6.17.xx.xx

D00000. 0.17.XX.XX

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

## Show Storage Subsystem Host Topology

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

#### DS5000: 7.3x.xx.xx

#### Show Storage Subsystem LUN Mappings DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx
### Show Storage Subsystem Negotiation Defaults

DS3000: 6.17.xx.xx

DS4000: 7.1x.xx.xx

DS5000: 7.3x.xx.xx

### Show Storage Subsystem Profile

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

#### Show Storage Subsystem Remote Status Notification DS3000: 7.77.xx.xx

DS4000: N/A

DS5000: 7.77.xx.xx

### Show Storage Subsystem Unreadable Sectors

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### **Show String**

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Show Thin Logical Drive

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Show Storage subsystem Unconfigured iSCSI Initiators DS3000: 6.50.xx.xx

DS4000: N/A

DS5000: 7.6x.xx.xx

### Show VolumeCopy

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

# Show VolumeCopy Source Candidates

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Show VolumeCopy Target Candidates

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

### Start subsystem Defragment

DS3000: 6.17.xx.xx

DS4000: 7.1x.xx.xx

DS5000: 7.3x.xx.xx

Start subsystem Export DS3000: 7.35.xx.xx and higher

DS4000: 7.1x.xx.xx

DS5000: 7.3x.xx.xx

### Start subsystem Import DS3000: 7.35.xx.xx and higher

DS4000: 7.1x.xx.xx

DS5000: 7.3x.xx.xx

### Start subsystem Locate

DS3000: 6.17.xx.xx

DS4000: 6.16.xx.xx

DS5000: 7.3x.xx.xx

# Start Cache Backup Device Diagnostic

DS3000: N/A

DS4000: N/A

DS5000: 7.70.xx.xx

### Start Cache Memory Diagnostic

DS3000: N/A

# DS4000: N/A

DS5000: 7.70.xx.xx

#### Start Consistency Group Enhanced FlashCopy Rollback DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

# Start Configuration Database Diagnostic

DS3000: 7.77.xx.xx

# DS4000: N/A

DS5000: 7.77.xx.xx

#### Start Controller Diagnostic DS3000: 7.70.xx.xx

DS4000: N/A

DS5000: 7.70.xx.xx

#### Start Controller Trace DS3000: 7.77.xx.xx

DS4000: N/A

D34000. IN/A

DS5000: 7.77.xx.xx

### Start Disk Pool Locate

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

### Start Drive Channel Fault Isolation Diagnostics

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

**Note:** For version 7.60.xx.xx, you can change the drive parameter to accommodate the enclosures with internal disk drawers.

### Start Drive Channel Locate

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Start Drive Initialize

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

**Note:** For version 7.60.xx.xx, you can change the drive parameter to accommodate the enclosures with internal disk drawers.

#### Start Drive Locate

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

**Note:** For version 7.60.xx.xx, you can change the drive parameter to accommodate the enclosures with internal disk drawers.

#### Start Drive Reconstruction

DS3000: 6.17.xx.xx

DS4000: 5.43.xx.xx

DS5000: 7.3x.xx.xx

**Note:** For version 7.60.xx.xx, you can change the drive parameter to accommodate the enclosures with internal disk drawers.

### Start Enclosure Locate

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Start Enhanced Remote Mirroring Synchronization

DS3000: 7.7x.xx.xx

DS4000: 6.1x.xx.xx

#### Start Enhanced FlashCopy Image Rollback

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

# Start Host Interface Card Diagnostic

DS3000: 7.70.xx.xx

DS4000: N/A.xx.xx

DS5000: 7.70.xx.xx

### Start iSCSI DHCP Refresh

DS3000: 6.50.xx.xx

DS4000: N/A

DS5000: 7.6x.xx.xx

#### Start Logical Drive Initialization

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

# Start Secure Drive Erase

DS3000: 7.7x.xx.xx

DS4000: Not supported

DS5000: 7.5x.xx.xx

**Note:** For version 7.60.xx.xx, you can change the drive parameter to accommodate the enclosures with internal disk drawers.

# Start Storage Subsystem Configuration Database Diagnostic

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Start Storage Subsystem Diagnostic Data Capture DS3000: 7.35.xx.xx

DS4000: 6.6x.xx.xx

DS5000: 7.3x.xx.xx

# Start Storage Subsystem iSNS Server Refresh

DS3000: 6.50.xx.xx

DS4000: N/A

DS5000: 7.6x.xx.xx

**Note:** This command replaces **Set Storage Subsystem iSNS Server Refresh** starting with Storage Manager version 10.83.xx.xx

#### Start Storage Subsystem Locate

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

#### Stop subsystem Locate

DS3000: 6.17.xx.xx

DS4000: 6.16.xx.xx

DS5000: 7.3x.xx.xx

# Stop Cache Backup Device Diagnostic

DS3000: N/A

DS4000: N/A

DS5000: 7.70.xx.xx (DS5000 only)

#### Stop Configuration Database Diagnostic DS3000: 7.77.xx.xx

-----

DS4000: N/A

DS5000: 7.77.xx.xx

### Stop Cache Memory Diagnostic

DS3000: N/A

DS4000: N/A

DS5000: 7.70.xx.xx (DS5000 only)

Stop Consistency Group Enhanced FlashCopy Logical Drive DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

Stop Consistency Group Enhanced FlashCopy Rollback DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

#### Stop Controller Diagnostic DS3000: 7.70.xx.xx

DS4000: N/A

DS5000: 7.70.xx.xx

# Stop Disk Pool Locate

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

### Stop Drive Channel Fault Isolation Diagnostics DS3000: Not supported

DS4000: 7.3x.xx.xx

DS5000: 7.3x.xx.xx

### Stop Drive Channel Locate

DS3000: 6.17.xx.xx

### DS4000: 6.1x.xx.xx

#### Stop Drive Locate

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

# Stop Enclosure Locate

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Stop Enhanced FlashCopy Group Pending Enhanced FlashCopy Images DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

### Stop Enhanced FlashCopy Image

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

# Stop Enhanced FlashCopy Image Rollback

DS3000: 7.83.xx.xx and later

DS4000: N/A

DS5000: N/A

### Stop FlashCopy

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Stop FlashCopy Rollback

DS3000: 7.83.xx.xx

DS4000: N/A

DS5000: 7.83.xx.xx

# Stop Enhanced FlashCopy Logical Drive

DS3000: 7.83.xx.xx and later

DS4000: N/A

### DS5000: N/A

# Stop Host Interface Card Diagnostic

DS3000: 7.70.xx.xx

DS4000: N/A

DS5000: 7.70.xx.xx

#### Stop Pending Enhanced FlashCopy Images on Consistency Group DS3000: 7.83.xx.xx and later

DS4000: N/A

#### DS5000: N/A

#### Stop Storage Subsystem Configuration Database Diagnostic DS3000: 7.77.xx.xx and later

DS4000: N/A

DS5000: 7.77.xx.xx and later

# Stop Storage Subsystem Drive Firmware Download

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Stop Storage Subsystem iSCSI Session DS3000: 6.50.xx.xx

DS4000: Not supported

DS5000: 7.6x.xx.xx

### Stop Storage Subsystem Drive Firmware Download

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Stop Storage Subsystem Locate

DS3000: 6.17.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

### Stop VolumeCopy

DS3000: 6.17.xx.xx

DS4000: 5.4x.xx.xx

DS5000: 7.3x.xx.xx

# Suspend Remote Mirror

DS3000: 7.7x.xx.xx

DS4000: 6.1x.xx.xx

DS5000: 7.3x.xx.xx

#### Validate Storage Subsystem Security Key

DS3000: N/A

DS4000: N/A

DS5000: 7.70.xx.xx

# **Activate Enhanced Global Mirroring**

This command activates the Enhanced Global Mirroring premium feature. After you activate the Enhanced Global Mirroring premium feature, you must set up an Enhanced Global Mirror Group and an Enhanced Global Mirrored Pair.

### Syntax

activate storagesubsystem feature=enhancedGlobalMirror

### Parameters

None

# Activate Enhanced Remote Mirroring Feature

This command creates the mirror repository logical drive and activates the Enhanced Remote Mirroring feature. When you use this command, you can define the mirror repository logical drive in one of three ways:

- User-defined disk drives
- User-defined subsystem
- User-defined number of disk drives

If you choose to define a number of drives, the controller firmware chooses which disk drives to use for the mirror repository logical drive.

### Syntax (User-Defined Disk Drives)

```
activate storageSubsystem feature=remoteMirror
repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDrives=(enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn)
repositorysubsystemUserLabel=[subsystemName]
[enclosureLossProtect=(TRUE | FALSE)]
[drawerLossProtect=(TRUE | FALSE)];
T10PI=(none | enabled)]
```

## Syntax (User-Defined subsystem)

```
activate storagesubsystem feature=remoteMirror
repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDrives=(enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn)
repositorysubsystemUserLabel=[subsystemName]
driveMediaType=(HDD | SSD | unknown | allMedia)
driveType=(fibre | SATA | SAS)
[enclosureLossProtect=(TRUE | FALSE)
drawerLossProtect=(TRUE | FALSE)
T10 PI=(none | enabled)]
```

## Syntax (User-Defined Number of Disk Drives)

```
activate storagesubsystem feature=remoteMirror
repositorysubsystem=subsystemName
[freeCapacityArea=freeCapacityIndexNumber]
activate storagesubsystem feature=remoteMirror
repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDriveCount=numberOfDrives
repositorysubsystemUserLabel=[subsystemName]
driveMediaType=(HDD | SSD | unknown | allMedia)
driveType=(fibre | SATA | SAS)]
[enclosureLossProtect=(TRUE | FALSE)
drawerLossProtect=(TRUE | FALSE)
T10 PI=(none | enabled)]
```

### **Parameters**

Parameter	Description
repositoryRAIDLevel	The RAID level for the mirror repository logical drive. Valid values are 1, 3, 5, or 6.

Parameter	Description
repositoryDrives	The disk drives for the mirror repository logical drive. EnclosureID1,drawerID1, slotID1 enclosureIDn,drawerIDn,slotIDn is for enclosures with multiple disk drive drawers. For enclosures without disk drive drawers, use enclosureID1,slotID1 enclosureIDn,slotIDn. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID values, the drawer ID values, and the slot ID values in parentheses.
repositorysubsystemUserLabel	The alphanumeric identifier (including - and _) that you want to give the new subsystem in which the mirror repository logical drive will be located. Enclose the subsystem identifier in square brackets ([]).
repositorysubsystem	The name of the mirror repository subsystem, where the mirror repository logical drive is located. (To determine the names of the subsystems in your storage subsystem, run the <b>show storageSubsystem</b> <b>profile</b> command.)
driveType	<ul> <li>The type of drive that you want to use in the mirror logical drive. You cannot mix drive types.</li> <li>You must use this parameter when you have more than one type of drive in your storage subsystem.</li> <li>Valid drive types are : <ul> <li>fibre</li> <li>SATA</li> <li>SAS</li> </ul> </li> <li>If you do not specify a drive type, the command defaults to fibre.</li> <li>Use this parameter when you use the repositoryDriveCount parameter.</li> </ul>
driveMediaType	<ul> <li>The type of drive media that you want to use for the mirror repository subsystem. Valid drive media are these:</li> <li>HDD – Use this option when you have hard drives in the drive expansion enclosure.</li> <li>SSD – Use this option when you have solid state drives in the drive expansion enclosure.</li> <li>unknown – Use if you are not sure what types of drive media are in the drive expansion enclosure.</li> <li>allMedia – Use this option when you want to use all types of drive media that are in the drive expansion enclosure.</li> <li>Use this parameter when you use the repositoryDriveCount parameter.</li> <li>You must use this parameter when you have more than one type of drive media in your storage subsystem.</li> </ul>

Parameter	Description
freeCapacityArea	The index number of the free space in an existing subsystem that you want to use to create the mirror repository logical drive. Free capacity is defined as the free capacity between existing logical drives in an subsystem. For example, an subsystem might have the following areas: logicalDrive 1, free capacity, logicalDrive 2, free capacity, logicalDrive 3, free capacity. To use the free capacity following logical drive 2, you would specify:
	freeCapacityArea=2
	Run the <b>show subsystem</b> command to determine if a free capacity area exists.
repositoryDriveCount	The number of unassigned disk drives that you want to use for the mirror repository logical drive.
enclosureLossProtect	The setting to enforce enclosure loss protection when you create the mirror repository logical drive. To enforce loss protection, set this parameter to <b>TRUE</b> . The default value is <b>FALSE</b> .
drawerLossProtect	The setting to enforce drawer loss protection when you create the mirror repository logical drive. To enforce drawer loss protection, set this parameter to TRUE. The default value is FALSE.
T10PI	The setting to specify that an subsystem, and the logical drives within the subsystem, has T10 PI (Protection Information) to help you ensure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the subsystem. These settings are valid: <ul> <li>none—The subsystem does not have T10 PI (Protection Information)</li> </ul>
	<ul> <li>(Protection Information).</li> <li>enabled—The subsystem has T10 PI (Protection Information). The subsystem support is protected information and is formatted with protection information enabled.</li> </ul>

# Notes

If the disk drives that you select for the **repositoryDrives** parameter are not compatible with other parameters (such as the **repositoryRAIDLevel** parameter), the script command returns an error, and Enhanced Remote Mirroring is not activated. The error returns the amount of space that is needed for the mirror repository logical drive. You can then re-enter the command, and specify the appropriate amount of space.

If you enter a value for the repository storage space that is too small for the mirror repository logical drives, the controller firmware returns an error message that provides the amount of space that is needed for the mirror repository logical drives. The command does not try to activate Enhanced Remote Mirroring. You can re-enter the command by using the value from the error message for the repository storage space value.

When you assign the disk drives, if you set the **enclosureLossProtect** parameter to TRUE and have selected more than one disk drive from any one enclosure, the

storage subsystem returns an error. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs operations, but the subsystem that you create might not have enclosure loss protection.

When the controller firmware assigns the disk drives, if you set the **enclosureLossProtect** parameter to TRUE, the storage subsystem returns an error if the controller firmware cannot provide disk drives that result in the new subsystem having enclosure loss protection. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs the operation even if it means that the subsystem might not have enclosure loss protection.

# **Activate Host Port**

This command activates an inactive host port that was created when the Host Context Agent (HCA) registered the host port to a host.

### **Syntax**

activate hostPort "user-label"

### Parameters

Parameter	Description
userLabel	The name of the HCA host port. Enclose the host port name in double quotation marks (" ").

# Activate iSCSI Initiator

This command activates an inactive iSCSI initiator that was created when the Host Context Agent (HCA) registered the iSCSI initiator to a host.

## **Syntax**

activate iscsiInitiator "iSCSI-ID"

## Parameters

Parameter	Description
iscsiInitiator	The name of the iSCSI initiator. Enclose the name in double quotation marks (" ").

# Activate Storage Subsystem Firmware

This command activates firmware that you have previously downloaded to the pending configuration area on the controllers in the storage subsystem.

## **Syntax**

activate storageSubsystem firmware

## **Parameters**

None.

# Activate Enhanced Remote Mirroring

This command creates the mirror repository logical drive and activates the Enhanced Remote Mirroring premium feature. When you use this command, you can define the mirror repository logical drive in one of three ways:

### **Syntax**

- User-defined drives
- User-defined Volume Group
- User-defined number of drives

If you choose to define a number of drives, the controller firmware chooses which drives to use for the mirror repository logical drive.

**Note:** In previous versions of this command the feature identifier was remoteMirror. This feature identifier is no longer valid and is replaced by syncMirror.

activate storagesubsystem firmware

### **Parameters**

None

# Add logical drive to Enhanced Global Mirror Group

This command adds a primary logical drive to an Enhanced Global Mirror group. This command is valid only on the local storage subsystem that contains the Enhanced Global Mirror Group to which you want to add the primary logical drive. An Enhanced Global Mirror Group has a repository logical drive that is used to save data for all of the point-in-time images that are part of the Enhanced Global Mirror Group. Each primary logical drive in the Enhanced Global Mirror Group has a corresponding mirror logical drive on a remote storage subsystem.

### Syntax

```
add logicalDrive="logicalDriveName"
enhancedGlobalMirrorGroup="enhancedGlobalMirrorGroupName" △
remotePassword="password" △
(repositorylogicalDrive="repos_xxxx" | △
repositorylogicalDrive=(logicalDriveGroupName
[capacity=capacityValue]) △
repositorylogicalDrive=(diskPoolName
[capacity=capacityValue]))
```

### Parameters

None

# Add Member to Consistency Group

This command adds a base logical drive as a member to an existing consistency group. You can specify an existing repository logical drive for the new consistency group member, or create a repository logical drive. When you create a repository logical drive, you identify an existing subsystem or an existing disk pool where you want the repository logical drive.

# Syntax for Use With an Existing Repository Logical Drive

set consistencyGroup ["consistencyGroupName"]
addCGMemberLogicalDrive="baseLogicalDriveName"
repositoryLogicalDrive="repos\_XXXX"

# Syntax for Use When Creating a New Repository Logical Drive in an subsystem

set consistencyGroup ["consistencyGroupName"]
addCGMemberLogicalDrive="baseLogicalDriveName"
repositoryLogicalDrive=("subsystemName" capacity=capacityValue(KB|MB|GB|TB|bytes))

# Syntax for Use When Creating a New Repository Logical Drive in a Disk Pool

set consistencyGroup ["consistencyGroupName"]
addCGMemberLogicalDrive="baseLogicalDriveName"
repositoryLogicalDrive=("diskpoolName" capacity=capacityValue(KB|MB|GB|TB|bytes))

#### Parameter Description The name of the consistency group to which you want consistencyGroup to add a member logical drive. The new member logical drive is the base logical drive for Enhanced FlashCopy operations. Enclose the consistency group name in double quotation marks (" ") inside square brackets ([ ]). The name of a base logical drive that you want to add. addCGMemberLogicalDrive Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the member name in double quotation marks (" "). If the specified logical drive is an existing repository logical drive or an existing Enhanced FlashCopy logical drive, the command fails. repositoryLogicalDrive This parameter performs 2 functions: In an existing consistency group that has a repository logical drive, this parameter identifies the repository logical drive. When creating a repository logical drive this parameter identifies either an subsystem or disk pool in which to create the repository logical drive. The size of a new repository logical drive in either an capacity subsystem or a disk pool. Size is defined in units of bytes, KB, MB, GB, or TB.

### **Parameters**

#### Notes

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

A consistency group is a collection of base logical drives that are the source logical drives for Enhanced FlashCopies. You collect the base logical drives in a consistency group so that you can perform the same Enhanced FlashCopy operations on each of the base logical drives. In the context of this command, the term *member* means a base logical drive for Enhanced FlashCopy operations. You

can manipulate Enhanced FlashCopy images associated with members of a consistency group through batch-style operations, while maintaining consistency across the Enhanced FlashCopy images.

Each member logical drive must have a corresponding repository logical drive. You can define the relationship between the member logical drive and the repository logical drive using the **repositoryLogicalDrive** parameter. The **repositoryLogicalDrive** parameter can perform one of these actions:

- Identify an existing repository logical drive that is connected to the consistency group.
- Identify either an subsystem or a disk pool in which you want to create a repository logical drive.

Adding a member to a consistency group with an existing repository has two purposes:

• You can create an entirely new repository logical drive by running the command without the repositoryLogicalDrive parameter. When you run the command without the **repositoryLogicalDrive** parameter the command creates a repository logical drive in the subsystem or disk pool in which all the other repository logical drives are stored. An example of this command usage is: set consistencyGroup ["First\_Images"] addCGMemberLogicalDrive="Data\_020212"

You can reuse an existing repository logical drive if:

- that repository logical drive is empty
- · that repository logical drive is not related to another member logical drive
- you want to maintain a particular sequence or relationship for the repository logical drives

To reuse an existing, empty repository logical drive you need to know the name of the repository logical drive. To determine the name of the repository logical drive use the show allLogicalDrives summary command. All repository logical drive names have the form **repos\_XXXX**, where XXXX is a unique identifier generated by the storage management software. An example of this command usage is: set consistencyGroup ["*First\_Images*"] addCGMemberLogicalDrive="*Data\_020212*" repositoryLogicalDrive="*repos\_0011*"

You can place the repository logical drive in an subsystem or a disk pool of your choosing. You are not required to have the repository logical drive in the same location as other repository logical drives. To place the repository logical drive in an subsystem or a disk pool of your choice, use the **repositoryLogicalDrive** parameter and identify the subsystem or the disk pool and a size for the repository logical drive. An example of this command usage is:

```
set consistencyGroup ["First_Images"] addCGMemberLogicalDrive="Data_020212"
repositoryLogicalDrive="12" capacity=2 GB)
```

In this example, "12" is the name of an existing subsystem or an existing disk pool. The capacity parameter defines the size that you want for the repository subsystem.

When you create a repository logical drive in either an subsystem or a disk pool, you must include parenthesis around the subsystem name and capacity, or the disk pool name and capacity.

# Add drives to Performance Read Cache

Use this command to increase the capacity of an existing Performance Read Cache by adding additional solid state disks (SSDs).

### **Synatx**

```
set performanceReadCache [performanceReadCacheName]
addDrives=(enclosureID1,drawerID1,slotID1 ...
enclosureIDn,drawerIDn,slotIDn)
```

# **Parameters**

Parameter	Description
performanceReadCache	The alphanumeric identifier (including - and _) of the Performance Read Cache to which you want to add SSDs. Enclose the identifier in square brackets ([]). If the Performance Read Cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks ("") inside square brackets.
addDrives	The drives that you want to add to the Performance Read Cache. For high-capacity drive enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for each SSD that you want to add. For low-capacity drive enclosures, specify the enclosure ID value and the slot ID value for each SSD that you want to add.
	<ul> <li>Enclosure ID values are 0 to 99</li> <li>Drawer ID values are 1 to 5</li> <li>Slot ID values are 1 to 32</li> <li>Enclose the enclosure ID values, the drawer ID values, and the slot ID values in parentheses.</li> </ul>

**Note:** The Performance Read Cache can contain any number of SSDs. The maximum size of the Performance Read Cache is 5 TB, but might be less depending on the size of the controller's primary cache.

# **Minimum Firmware Level**

7.84

# Autoconfigure Storage Subsystem

This command automatically configures a storage subsystem. Before you enter the autoConfigure storageSubsystem command, run the show storageSubsystem autoConfiguration command. The show storageSubsystem autoConfiguration command returns configuration information in the form of a list of valid disk drive types, RAID levels, logical drive information, and hot spare information. (This list corresponds to the parameters for the autoConfigure storageSubsystem command.) The controllers audit the storage subsystem and then determine the highest RAID level that the storage subsystem can support and the most efficient logical drive

definition for the RAID level. If the configuration that is described by the returned list is acceptable, you can enter the **autoConfigure storageSubsystem** command without any parameters. If you want to modify the configuration, you can change the parameters to meet your configuration requirements. You can change a single parameter or all of the parameters. After you enter the **autoConfigure storageSubsystem** command, the controllers set up the storage subsystem by using either the default parameters or those you selected.

# **Syntax**

```
autoConfigure storageSubsystem
[driveType=(fibre | SATA | SAS |)
raidLevel=(0 | 1 | 3 | 5 | 6)
subsystemWidth=numberOfDrives
subsystemCount=numberOfsubsystems
logicalDrivesPersubsystemCount=numberOfLogicalDrivesPersubsystem
hotSpareCount=numberOfHotSpares
segmentSize=segmentSizeValue
cacheReadPrefetch=(TRUE | FALSE)
securityType=(none | capable | enabled)]
T10PI=(none | enabled)]
```

# Parameters

Parameter	Description
driveType	The type of disk drives that you want to use for the storage subsystem. You must use this parameter when you have more than one type of drive in your storage subsystem. Valid disk drive types are fibre, SATA, and SAS. If you do not specify a drive type, the command defaults to fibre.
raidLevel	The RAID level of the subsystem that contains the disk drives in the storage subsystem. Valid RAID levels are $0$ , 1, 3, 5, or 6.
subsystemWidth	The number of disk drives in an subsystem in the storage subsystem.
subsystemCount	The number of subsystems in the storage subsystem. Use integer values.
logicalDrivesPerGroupCount	The number of equal-capacity logical drives per subsystem. Use integer values.
hotSpareCount	The number of hot spares that you want in the storage subsystem. Use integer values.
segmentSize	The amount of data (in KB) that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.

Parameter	Description
securityType	The setting to specify the security level when creating the subsystem and all associated logical drives. These settings are valid:
	• <b>none</b> —The subsystem and logical drives are not secure.
	• <b>capable</b> —The subsystem and logical drives are capable of having security set, but security has not been enabled.
	• <b>enabled</b> —The subsystem and logical drives have security enabled.
T10PI	The setting to specify that an subsystem, and the logical drives within the subsystem, has T10 PI (Protection Information) to help you ensure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the subsystem. These settings are valid:
	• <b>none</b> —The subsystem does not have T10 PI (Protection Information).
	• <b>enabled</b> —The subsystem has T10 PI (Protection Information). The subsystem support is protected information and is formatted with protection information enabled.

# Notes

# **Drives and subsystem**

An subsystem is a set of disk drives that are logically grouped together by the controllers in the storage subsystem. The number of disk drives in an subsystem is a limitation of the RAID level and the controller firmware. Follow these guidelines when you create an subsystem:

- Beginning with firmware version 7.10, you can create an empty subsystem to reserve the capacity for later use.
- You cannot mix drive types, such as SAS, SATA, and Fibre Channel, within a single subsystem.
- The maximum number of drives in an subsystem depends on these conditions:
  - The type of controller
  - The RAID level
- RAID levels include: 0, 1, 10, 3, 5, and 6.
  - In a DS4700 Model 70 or a DS4700 Model 72 storage subsystem, an subsystem with RAID level 0 and an subsystem with RAID level 10 can have a maximum of 112 drives.
  - In a DS4800 (Models 82, 84, 88) storage subsystem, an subsystem with RAID level 0 and an subsystem with RAID level 10 can have a maximum of 224 drives.
  - An subsystem with RAID level 3, RAID level 5, or RAID level 6 cannot have more than 30 drives.
  - An subsystem with RAID level 6 must have a minimum of five drives.
  - If an subsystem with RAID level 1 has four or more drives, the storage management software automatically converts the subsystem to a RAID level 10, which is RAID level 1 + RAID level 0.

For information about the maximum number of disk drives in a RAID subsystem or disk pool, see Table 10 on page 395.

- If an subsystem contains disk drives that have different capacities, the overall capacity of the subsystem is based on the smallest capacity drive.
- To enable enclosure loss protection, you must create an subsystem that uses disk drives located in at least three drive expansion enclosures.

Disk drives do not report their exact capacity (for example, a 73 GB drive does not report 73 GB as its capacity).

### Hot Spares

Hot spare disk drives can replace any failed disk drive in the storage subsystem. The hot spare must be the same type of disk drive as the disk drive that failed (that is, a SATA hot spare cannot replace a Fibre Channel disk drive). A hot spare must have capacity greater than or equal to any disk drive that can fail. If a hot spare is smaller than a failed disk drive, you cannot use the hot spare to rebuild the data from the failed disk drive. Hot spares are available only for RAID levels 1, 3, 5, or 6.

### Segment Size

The size of a segment determines how many data blocks that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single disk drive for a single request leaves other disk drives available to simultaneously service other requests.

If the logical drive is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of disk drives in the subsystem that are used for data transfers.) In this case, multiple disk drives are used for the same request, but each disk drive is accessed only once. For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of disk drives that are required to satisfy a data transfer request.

### **Cache Read Prefetch**

Cache read prefetch lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the disk drive into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage subsystem that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the **cacheReadPrefetch** parameter are TRUE or FALSE.

# Security Type

The **securityType** parameter is valid for drives that are capable of full disk encryption (FDE). With FDE, the controller firmware can create a key and activate the Full Disk Encryption feature. The Full Disk Encryption feature encrypts data as the data is written to the drive and decrypts the data as the data is read from the drive. Without the key created by the controller, the data written to the drive is inaccessible

Before you can set the **securityType** parameter to *capable* or *enabled*, you must create a storage subsystem security key. Use the create storageSubsystem securityKey command to create a storage subsystem security key. These commands are related to the security key:

- create storageSubsystem securityKey
- set storageSubsystem securityKey
- import storageSubsystem securityKey
- export storageSubsystem securityKey
- start secureErase (drive | drives)
- enable subsystem [subsystemName] security

# Autoconfigure Storage Subsystem Hot Spares

This command automatically defines and configures the hot spares in a storage subsystem. You can run this command at any time. This command provides the best hot spare coverage for a storage subsystem.

### Syntax

autoConfigure storageSubsystem hotSpares

### **Parameters**

None.

# Notes

When you run the autoconfigure storageSubsystem hotSpares command, the controller firmware determines the number of hot spares to create based on the total number and type of disk drives in the storage subsystem. For Fibre Channel disk drives, SATA disk drives, and SAS disk drives, the controller firmware creates one hot spare for the storage subsystem and one additional hot spare for every 60 disk drives in the storage subsystem.

# **Cancel Enhanced Global Mirror Group role reversal**

This command automatically defines and configures the hot spares in a storage subsystem. You can run this command at any time. This command provides the best hot spare coverage for a storage subsystem.

### **Syntax**

autoConfigure storagesubsystem hotSpares

# Parameters

None

# Change Performance Read Cache application type

This command changes the application type associated with the Performance Read Cache. The application type can be web server, database, or multimedia. Changing the application type changes the block size, subblock size, populate on read threshold, and populate on write threshold for the logical drives underlying the Performance Read Cache.

# **Syntax**

set performanceReadCache [performanceReadCacheName]
usageHint=(webServer|dataBase|fileSystem)

### **Parameters**

Parameter	Description
performanceReadCache	The alphanumeric identifier (including - and _) of the Performance Read Cache for which you want to change the application type. Enclose the identifier in square brackets ([]). If the Performance Read Cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.
usageHint	The values are based on the typical I/O usage pattern of the application that is using the Performance Read Cache. Valid values are webServer, dataBase, or fileSystem.

# **Minimum Firmware Level**

7.84

# **Check Enhanced Global Mirror Group consistency**

This command changes the application type associated with the Performance Read Cache. The application type can be web server, database, or multimedia. Changing the application type changes the block size, subblock size, populate on read threshold, and populate on write threshold for the logical drives underlying the Performance Read Cache.

## **Syntax**

set performanceReadCache [performanceReadCacheName]
usageHint=(webServer|dataBase|fileSystem)

## **Parameters**

None

# **Check Logical Drive Parity**

This command checks a logical drive for parity and media errors and writes the results of the check to a file.

# **Syntax**

```
check logicalDrive [logicalDriveName]
parity [parityErrorFile=filename]
[mediaErrorFile=filename]
[priority=(highest | high | medium | low | lowest)]
[startingLBA=LBAvalue] [endingLBA=LBAvalue]
[verbose=(TRUE | FALSE)]
```

# **Parameters**

Parameter	Description
logicalDrive	The name of the specific logical drive for which you want to check parity. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks ("").
parityErrorFile	The file path and the file name to which you want to save the parity error information. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\parerr.txt"
	This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.
mediaErrorFile	The file path and the file name to which you want to save the media error information. Enclose the file name in double quotation marks (" "). For example:
	file="C:\Program Files\CLI\logs\mederr.txt"
	This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.
priority	The priority that the parity check has relative to host I/O activity. Valid values are highest, high, medium, low, or lowest.
startingLBA	The starting logical block address. Use integer values.
endingLBA	The ending logical block address. Use integer values.
verbose	The setting to capture progress details, such as percent complete, and to show the information as the logical drive parity is being repaired. To capture progress details, set this parameter to TRUE. To prevent capturing progress details, set this parameter to FALSE.

# Notes

The starting logical block address and the ending logical block address are useful for very large single-logical drive LUNs. Running a logical drive parity check on a very large single logical drive LUN can take a long time. By defining the beginning address and ending address of the data blocks, you can reduce the time that a logical drive parity check takes to complete.

# **Related Topic**

Naming Conventions

# **Check Remote Mirror Status**

This command returns the status of a remote-mirror logical drive. Use this command to determine when the status of the remote-mirror logical drive becomes Optimal.

### **Syntax**

check remoteMirror localLogicalDrive [logicalDriveName] optimalStatus
timeout=timeoutValue

### **Parameters**

Parameter	Description
localLogicalDrive	The name of any remote-mirror logical drive. The remote-mirror logical drive can be the primary logical drive or the secondary logical drive of a remote-mirror pair. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks ("").
timeout	The time interval within which the software can return the remote-mirror logical drive status. The timeout value is in minutes.

### Notes

This command waits until the status becomes Optimal or the timeout interval expires. Use this command when you run the Enhanced Remote Mirroring utility.

### **Related Topics**

Naming Conventions

# **Check Repository Consistency**

This command applies to a number of objects that have underlying repository logical drives. The command produces a report based on analysis of the data in the repository.

### **Syntax**

```
check [enhancedFlashCopyGroup[enhancedFlashCopyGroupName] repositoryConsistency |
enhancedFlashCopyLogicalDrive[enhancedFlashCopyLogicalDriveName]
repositoryConsistency |
logicaldrive[logicaldriveName] repositoryConsistency |
logicalVolumeCopy target[targetName] repositoryConsistency |
localLogicalDrive="localLogicalDriveName"]
file="filePath"
```

# **Parameters**

Parameter	Description
repositoryConsistency	Specifies that the Enhanced FlashCopy group, Enhanced FlashCopy logical drive, or logical drive (depending on which is specified by the corresponding parameter) is checked for consistency.
enhancedF1ashCopyGroup	The alphanumeric identifier (including - and _) of the Enhanced FlashCopy group on which to run a consistency check. Enclose the Enhanced FlashCopy group identifier in square brackets ([]).
enhancedFlashCopyLogicalDrive	The alphanumeric identifier (including - and _) of the Enhanced FlashCopy logical drive on which to run a consistency check. Enclose the Enhanced FlashCopy logical drive identifier in square brackets ([]).
logicaldrive	The alphanumeric identifier (including - and _) of the thin logical drive on which to run a consistency check. Enclose the logical drive identifier in square brackets ([]).
VolumeCopy	Specifies that the target of a VolumeCopy relationship is checked for consistency.
target	Use only in conjunction with the logical VolumeCopy parameter to specify the alphanumeric identifier of the logical drive that is the target of the logical drive copy relationship. Enclose the logical drive identifier in double quotation marks (" ").
file	The file path and the file name to which you want to save the report that results from the consistency check. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\ repoConsistency.txt" This parameter must appear last, after any of the optional parameters.

# Check storage subsystem connectivity

This command verifies that the local storage subsystem and the remote storage subsystem have a communication path and displays the connection details between the local and remote storage subsystem. Before creating an Enhanced Global Mirror Group, you should check whether the local storage subsystem and the remote storage subsystem can communicate with each other. When you execute this command, the system queries for the remote storage subsystem on all eligible host ports to determine what ports are connected to the remote storage subsystem. The result of the test is a list of all ports on the local storage subsystem along with a list of the remote storage subsystem port addresses accessible through that port.

**Note:** Connectivity is tested using all possible channels, and if it is a dual controller configuration, connectivity is checked from each controller. It might take up to 20 minutes to check connectivity between two storage subsystems.

### Syntax

```
check storagesubsystem connectivity
(remoteStoragesubsystemName="storagesubsystemName" |
remoteStoragesubsystemWwid=<wwID>
```

### **Parameters**

None

# Check Enhanced Remote Mirroring status

This command returns the status of a remote-mirror logical drive. Use this command to determine when the status of the remote-mirror logical drive status changes to Optimal.

**Note:** In previous versions of this command the feature identifier was remoteMirror. This feature identifier is no longer valid and is replaced by syncMirror.

#### Syntax

check syncMirror locallogicalDrive [logicalDriveName]
optimalStatus timeout=timeoutValue

### Parameters

None

# Clear Enhanced Global Mirroring fault

This command clears an Enhanced Global Mirroring "sticky" fault from one or more Enhanced Global Mirror Groups and one or more Enhanced Global Mirror Group member logical drives. An Enhanced Global Mirror Group and its member logical drives can encounter Enhanced Global Mirroring "sticky" faults, which occur at a single point-in-time but do not impact the functionality of the mirrors. These type of faults must be reviewed, but may not require any changes to the configuration. An Enhanced Global Mirror Group and its member logical drives might or can have more than one associated sticky fault. This command clears all of the faults associated with the Enhanced Global Mirror Group and its member logical drive. However, if an Enhanced Global Mirror Group has a fault and one of its member logical drives has a fault, clearing the fault on the Enhanced Global Mirror Group does not clear the fault on its member logical drive.

### Syntax

```
clear enhancedGlobalMirrorFault(all | enhancedGlobalMirrorGroup
["enhancedGlobalMirrorGroupName"] | △
enhancedGlobalMirrorGroups ["enhancedGlobalMirrorGroupName1" ...
"enhancedGlobalMirrorGroupNameN"] | △
logicalDrive ["logicalDriveName"] | △
logicalDrives ["logicalDriveName1" ... "logicalDriveNameN"])
```

### Parameters

None

# **Clear Drive Channel Statistics**

This command resets the statistics for all of the disk drive channels.

### **Syntax**

clear allDriveChannels stats

### **Parameters**

None.

# **Clear Logical Drive Reservations**

This command clears persistent logical drive reservations.

### Syntax

```
clear (allLogicalDrives | logicaldrive [logicalDriveName] |
logicalDrives [logicalDriveName1 ... logicalDriveNameN]) reservations
```

### **Parameters**

Parameter	Description
allLogicalDrives	The setting to clear persistent logical drive reservations on all of the logical drives in the storage subsystem.
logicalDrive or logicalDrives	The name of the specific logical drive for which you want to clear persistent logical drive reservations. You can enter more than one logical drive name. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks ("").

# Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

### **Related topics**

Naming Conventions

# **Clear Logical Drive Unreadable Sectors**

This command clears unreadable sector information from one or more logical drives.

### Syntax

clear (allLogicalDrives | logicaldrive [logicalDriveName] |
logicalDrives [logicalDriveName1 ... logicalDriveNameN]) unreadableSectors

# **Parameters**

Parameter	Description
allLogicalDrives	The setting to clear unreadable sector information from all of the logical drives in the storage subsystem.
logicalDrive or logicalDrives	The name of the specific logical drive for which you want to clear unreadable sector information. You can enter more than one logical drive name. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks ("").

## Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

# **Related Topic**

Naming Conventions

# **Clear Storage Subsystem Configuration**

Use this command to perform one of these operations:

- Clear the entire storage subsystem configuration, and return it back to the initial installation state
- Clear the configuration except for security information and identification information
- Clear subsystem configuration information and logical drive configuration information only

**Attention: Possible damage to the storage subsystem configuration** – When you run this command, the existing storage subsystem configuration is deleted.

## Syntax

clear storageSubsystem configuration [all | subsystems]

## **Parameters**

Parameter	Description
None	If you do not enter a parameter, this command removes all configuration information for the storage subsystem, except for information related to security and identification.
all	The setting to remove the entire configuration of the storage subsystem, including security information and identification information. Removing all configuration information returns the storage subsystem to its initial state.
subsystems	The setting to remove the logical drive configuration and the subsystem configuration. The rest of the configuration stays intact.

### Notes

When you run this command, the storage subsystem becomes unresponsive, and all script processing is canceled. You must remove and re-add the storage subsystem to resume communication with the host. To remove an unresponsive storage subsystem, access the Enterprise Management Window, and select Edit **> Remove**. To re-add the storage subsystem, access the Enterprise Management Window, select Edit **>** Add Storage Subsystem, and enter the appropriate IP addresses.

# Clear storage subsystem Diagnostic Data Capture

This command sets a flag on a controller to allow a new Diagnostic Data Capture to overwrite an existing Diagnostic Data Capture.

#### Syntax

set storagesubsystem coreDumpAllowOverWrite

### Parameters

None

# Clear Storage Subsystem Diagnostic Data Capture

This command sets a flag on a controller to allow a new diagnostic data capture to overwrite an existing diagnostic data capture.

#### Syntax

set storageSubsystem ddcAllowOverWrite

#### Parameters

None.

# Notes

When you retrieve a diagnostic data capture from the controller cache to a host, a flag is set on the controller to indicate that the diagnostic data capture does not need to be retrieved. This setting persists for 48 hours. If a new diagnostic data capture occurs during that period the new diagnostic data capture is saved to the controller cache and overwrites any previous diagnostic data capture data in cache.

You can use the set storageSubsystem ddcAllowOverWrite command to set the controller flag so that a new diagnostic data capture will overwrite any previous diagnostic data capture. Without retrieving a diagnostic data capture, this command sets the flag as if you had.

# Clear Storage Subsystem Event Log

This command clears the Event Log in the storage subsystem by deleting the data in the Event Log buffer.

#### Attention:

**Possible damage to the storage subsystem configuration** – As soon as you run this command, the existing Event Log in the storage subsystem is deleted.

#### Syntax

clear storageSubsystem eventLog

### **Parameters**

None.

# Clear Storage Subsystem Firmware Pending Area

This command deletes a firmware image or NVSRAM values that you have previously downloaded from the pending area buffer.

#### Attention:

**Possible damage to the storage subsystem configuration** – As soon as you run this command, the contents of the existing pending area in the storage subsystem are deleted.

### Syntax

clear storageSubsystem firmwarePendingArea

#### **Parameters**

None.

# **Clear Storage Subsystem Recovery Mode**

This command forces a storage subsystem to exit recovery mode.

#### Syntax

clear storageSubsystem recoveryMode

### Notes

Recovery mode is entered during start-of-day operations when the system configuration is cleared and a valid on board backup is available. This mode is exited by either restoring the system configuration from a backup location, or by clearing the existing on board backups. While recovery mode is in force, a needs attention condition is raised and the Recovery Guru will be available from the user interface. However, the system configuration will be empty during recovery mode.

# **Configure Automatic Support Bundle Collection**

This command enables or disables the automatic collection of support bundles on one or more storage subsystems when a critical MEL event is detected.

# **Syntax**

**Attention:** The SM CLI support bundle commands must be run from the SM CLI command line interface only. They cannot be run from the storage subsystem SM CLI script window. In addition, the storage subsystem IP addresses are not required. Also, the **-c** parameter is not required to precede the SM CLI support bundle command.

<install Path>SMcli -supportBundle auto (enable|disable)
(all | storageSubsystemName) [data=pathName]

Parameter	Description
enable	Enable automatic collection of support bundles when a critical MEL event is detected.
disable	Disable automatic collection of support bundles where the feature was previously enabled.
all	Apply the command to all storage subsystems detected from the host.
storageSubsystemName	Apply the command to the named storage subsystem that is in the Storage Manager Enterprise management window domain. The names of the storage subsystems can be displayed using the SMcli command with the -d parameter.
data	Store the data generated by the command at the location specified by the <i>pathName</i> .

## **Parameters**

# Convert FlashCopy Logical Drive to Enhanced FlashCopy Group

This command migrates from a FlashCopy logical drive to an Enhanced FlashCopy group, and converts from an existing FlashCopy repository logical drive to an Enhanced FlashCopy image repository logical drive. A maximum of four FlashCopies can be converted to an Enhanced FlashCopy group. The FlashCopies must be in a Disabled state and must not be involved in a VolumeCopy operation. The base logical drive can have only four FlashCopies . If the base logical drive has more than four FlashCopies, you must delete the extra FlashCopies before running this command.

**Note:** You must disable any FlashCopies before attempting to convert from a FlashCopy logical drive to an Enhanced FlashCopy group. Use the disableFlashcopy command to disable a FlashCopy .

# **Syntax**

convert flashcopyLogicalDrive baseLogicalDrive="baseLogicalDriveName"

# **Parameters**

Parameter	Description
baseLogicalDrive	The alphanumeric identifier (including - and _) of the base logical drive that has the FlashCopies that you want to convert into an Enhanced FlashCopy group. Enclose the base logical drive name in double quotation marks (" ").

# **Create subsystem**

This command creates either a free-capacity subsystem or one logical drive on a set of unassigned disk drives.

### Syntax

```
create subsystem [userLabel]
Drives=(enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn)
raidLevel=(0 | 1 | 3 | 5 | 6)
userLabel=subsystemName
DriveMediaType=(HDD | SSD | unknown | allMedia)
enclosureLossProtect=(TRUE | FALSE)
drawerLossProtect=(TRUE | FALSE)
securityType=(none | capable | enabled)]
T10PI=(none | enabled)]
```

# **Parameters**

Parameter	Description
drives	The disk drives that you want to assign to the subsystem to be created. For high-capacity disk drive expansion enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for each disk drive that you assign to the subsystem. For low-capacity drive expansion enclosures, specify the enclosure ID value and the slot ID value for each disk drive that you assign to the subsystem. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID values, the drawer ID values, and the slot ID values in parentheses.
raidLevel	The RAID level of the subsystem that contains the logical drive. Valid values are $0$ , 1, 3, 5, or 6.
userLabel	The alphanumeric identifier (including - and _) that you want to give the new subsystem. Enclose the subsystem identifier in double quotation marks (" ").

Parameter	Description
DriveMediaType	The type of disk drive medium that you want to use for the create RAID logical drive. Values are
	<ul> <li>HDD – Use this option when you have hard disk drives in the expansion drawer.</li> </ul>
	• SSD – Use this option when you have solid state disk drives in the expansion drawer.
	• unknown – Use if you are not sure what types of disk drive media are in the expansion drawer.
	• allMedia – Use this option when you want to use all types of disk drive media that are in the expansion drawer.
	Use this parameter only when you use the DriveCount parameter. If you have only one type of disk drive media, you do not need to use this parameter. You must use this parameter when you have more than one type of disk drive media in your storage subsystem.
driveType	The type of drive that you want to use in the subsystem. You cannot mix drive types.
	You must use this parameter when you have more than one type of drive in your storage subsystem.
	Valid drive types are :
	• fibre
	• SAS
	If you do not specify a drive type, the command defaults to fibre.
enclosureLossProtect	The setting to enforce enclosure loss protection when you create the subsystem. To enforce enclosure loss protection, set this parameter to TRUE. The default value is FALSE.
drawerLossProtect	The setting to enforce drawer loss protection when you create the mirror repository logical drive. To enforce drawer loss protection, set this parameter to TRUE. The default value is FALSE.
securityType	The setting to specify the security level when creating the subsystems and all associated logical drives. These settings are valid:
	<ul> <li><i>none</i> – The subsystem and logical drives are not secure.</li> </ul>
	<ul> <li><i>capable</i> – The subsystem and logical drives are capable of having security set, but security has not been enabled.</li> </ul>
	<ul> <li><i>enabled</i> – The subsystem and logical drives have security enabled.</li> </ul>

Parameter	Description
T10PI	The setting to specify that an subsystem, and the logical drives within the subsystem, has T10 PI (Protection Information) to help you ensure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the subsystem. These settings are valid:
	<ul> <li>none—The subsystem does not have T10 PI (Protection Information).</li> </ul>
	• <b>enabled</b> —The subsystem has T10 PI (Protection Information). The subsystem support is protected information and is formatted with protection information enabled.

### Notes

The **drives** parameter lets you choose the number of disk drives that you want to use in the subsystem. If you choose this option, you do not need to specify the disk drives by enclosure ID and slot ID. The controllers choose the specific disk drives to use for the subsystem.

If you do not specify a capacity by using the **capacity** parameter, all of the disk drive capacity that is available in the subsystem is used. If you do not specify capacity units, bytes is used as the default value.

## **Cache Read Prefetch**

The **cacheReadPrefetch** command lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the disk drives into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage subsystem that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the **cacheReadPrefetch** parameter are TRUE or FALSE.

You do not need to enter a value for the **cacheReadPrefetch** parameter or the **segmentSize** parameter. If you do not enter a value, the controller firmware uses the **usageHint** parameter with fileSystem as the default value. Entering a value for the **usageHint** parameter and a value for the **cacheReadPrefetch** parameter or a value for the **segmentSize** parameter does not cause an error. The value that you enter for the **cacheReadPrefetch** parameter takes priority over the value for the **usageHint** parameter.

## Segment Size

The size of a segment determines how many data blocks that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8 KB segment holds 16 data blocks. A 64 KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that

you entered is not valid, the controller returns a list of valid values. Using a single disk drive for a single request leaves other disk drives available to simultaneously service other requests.

If the logical drive is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of disk drives in the subsystem that are used for data transfers.) In this case, multiple disk drives are used for the same request, but each disk drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of disk drives that are required to satisfy a data transfer request.

# Security Type

The **securityType** parameter is valid for drives that are capable of full disk encryption (FDE). With FDE, the controller firmware can create a key and activate the Full Disk Encryption feature. The Full Disk Encryption feature encrypts data as the data is written to the drive and decrypts the data as the data is read from the drive. Without the key created by the controller, the data written to the drive is inaccessible.

Before you can set the **securityType** parameter to *capable* or *enabled*, you must create a storage subsystem security key. Use the create storageSubsystem securityKey command to create a storage subsystem security key. These commands are related to the security key:

- create storageSubsystem securityKey
- enable subsystem [subsystemName] security
- export storageSubsystem securityKey
- import storageSubsystem securityKey
- set storageSubsystem securityKey
- start secureErase (drive | drives)

# **Enclosure Loss Protection and Drawer Loss Protection**

For enclosure loss protection to work, each disk drive in an subsystem must be in a separate enclosure. If you set the **enclosureLossProtect** parameter to TRUE and have selected more than one disk drive from any one enclosure, the storage subsystem returns an error. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs operations, but the subsystem that you create might not have enclosure loss protection.

Enclosure loss protection is not valid when you create logical drives on existing subsystems.

The **drawerLossProtect** parameter defines if data on a logical drive is accessible even if a drawer fails. When you assign the disk drives, if you set the **drawerLossProtect** parameter to TRUE and select more than one disk drive from any one drawer, the storage subsystem returns an error. If you set the **drawerLossProtect** parameter to FALSE, the storage subsystem performs operations, but the subsystem that you create might not have drawer loss protection. You must set the same value for the **enclosureLossProtect** and **drawerLossProtect** parameters. Both of the parameters must be either TRUE or FALSE. If the **enclosureLossProtect** parameter and the **drawerLossProtect** parameter are set to different values, the storage subsystem returns an error.

# Create Enhanced Global Mirror Group

This command migrates from a FlashCopy logical drive to a Enhanced FlashCopy group, and converts from an existing FlashCopy repository logical drive to a Enhanced FlashCopy image repository logical drive. A maximum of four Enhanced FlashCopies can be converted to a Enhanced FlashCopy group. The Enhanced FlashCopies must be in a Disabled state and must not be involved in a VolumeCopy operation. The base logical drive can have only four Enhanced FlashCopies (legacy). If the base logical drive has more than four Enhanced FlashCopies (legacy), you must delete the extra Enhanced FlashCopies before running this command.

**Note:** Users must disable any Enhanced FlashCopies before attempting to convert from a FlashCopy logical drive to a Enhanced FlashCopy group. Use the disableEnhanced FlashCopy command to disable a FlashCopy.

### **Syntax**

convert Enhanced FlashCopylogicalDrive baselogicalDrive="baselogicalDriveName"

### **Parameters**

None

# Create Consistency Group

Note: This command does not apply to the FlashCopy commands.

This command creates a new, empty consistency group that can contains Enhanced FlashCopy groups. You must add the Enhanced FlashCopy groups using the set consistencyGroup addCGMember command.

## **Syntax**

```
create consistencyGroup userLabel="consistencyGroupName"
[repositoryFullPolicy=(failBaseWrites | purgeEnhancedFlashcopyImages) |
repositoryFullLimit=percentValue |
autoDeleteLimit=numberOfEnhancedFlashcopyImages |
enableSchedule=(TRUE | FALSE) |
schedule (immediate | EnhancedFlashcopySchedule) |
rollbackPriority=(lowest | low | medium | high | highest)]
```

### **Parameters**

Parameter	Description
userLabel	The name of the new consistency group that you want to create. Enclose the new consistency group name in double quotation marks (" ").

Parameter	Description
repositoryFullPolicy	How you want Enhanced FlashCopy processing to continue if the Enhanced FlashCopy repository logical drives are full. You can choose to fail writes to the base logical drive (failBaseWrites) or delete (purge) the Enhanced FlashCopy images ( purgeEnhancedFlashcopyImages). The default action is purgeEnhancedFlashcopyImages.
repositoryFullLimit	The percentage of repository capacity at which you receive a warning that the Enhanced FlashCopy repository logical drive is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.
autoDeleteLimit	The maximum number of Enhanced FlashCopy images that you want to automatically delete if you have selected to purge the Enhanced FlashCopy images for a repository full policy. The default value is 32.
enableSchedule	Whether the ability to schedule an Enhanced FlashCopy operation is turned on or turned off. To turn on Enhanced FlashCopy scheduling, set this parameter to TRUE. To turn off Enhanced FlashCopy scheduling, set this parameter to FALSE.
schedule	Use this parameter to schedule an Enhanced FlashCopy operation. You can use one of these options for setting a schedule for an Enhanced FlashCopy operation: • <i>immediate</i> • <i>startDate</i> • <i>startDate</i> • <i>scheduleDay</i> • <i>startTime</i> • <i>scheduleInterval</i> • <i>endDate</i> • <i>timesPerDay</i> • <i>timeZone</i> See the "Notes" section for information explaining how to use these options.
rollBackPriority	Determines whether system resources should be allocated to the rollback operation at the expense of system performance. A value of high indicates that the rollback operation is prioritized over all other host I/O. A value of low indicates that the rollback operation should be performed with minimal impact to host I/O.

**Note:** A consistency group is a logical entity that enables you to manage in batch form all of the Enhanced FlashCopy images that you add to the collection. The consistency group is a collection of Enhanced FlashCopy groups that have mutual consistency requirements or dependencies for their Enhanced FlashCopy images.

Any Enhanced FlashCopy images that you create and use for this collection must be managed in accordance with the consistency dependencies.

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

The Enhanced FlashCopy images in a consistency group can be deduced based on the existence of an Enhanced FlashCopy image within a consistency group. All Enhanced FlashCopy images that reside within a consistency group share a common time stamp and sequence number.

An operation on an Enhanced FlashCopy image consistency group is treated as a single request, and causes all pending I/O operations to the associated base logical drive of each member to be drained and suspended before creating the Enhanced FlashCopy images. If creation of the Enhanced FlashCopy images cannot be completed successfully for all of the consistency group members, the operation fails and has no affect (that is, new Enhanced FlashCopy images are not created).

Based on this behavior all members for a consistency group usually have the same number of Enhanced FlashCopy images. However, when a new member is added to a consistency group, that new member lacks the Enhanced FlashCopy images that were previously created on the established members of the consistency group. This is not considered an error condition. Ensuing requests for deletion or rollback of Enhanced FlashCopy images that only exist on a subset of the consistency group members will only affect the members for which the specified Enhanced FlashCopy images actually exists

## Scheduling Enhanced FlashCopy Images in a Consistency Group

The enableSchedule parameter and the schedule parameter provide a way for you to schedule automatic FlashCopies . Using these parameters, you can schedule Enhanced FlashCopies daily, weekly, or monthly (by day or by date). The enableSchedule parameter turns on or turns off the ability to schedule Enhanced FlashCopies. When you enable scheduling, you use the schedule parameter to define when you want the Enhanced FlashCopies to occur.

This list explains how to use the options for the schedule parameter:

- immediate As soon as you enter the command, an Enhanced FlashCopy logical drive is created and a copy-on-write operation begins.
- startDate A specific date on which you want to create an Enhanced FlashCopy logical drive and perform a copy-on-write operation. The format for entering the date is MM:DD:YY. If you do not provide a start date, the current date is used. An example of this option is startDate=06:27:11.
- scheduleDay A day of the week on which you want to create an Enhanced FlashCopy logical drive and perform a copy-on-write operation. The values that you can enter are: monday, tuesday, wednesday, thursday, friday, saturday, sunday. An example of this option is scheduleDay=wednesday.
- startTime The time of a day that you want to create an Enhanced FlashCopy logical drive and start performing a copy-on-write operation. The format for entering the time is HH:MM, where HH is the hour and MM is the minute past the hour. Use a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is startTime=14:27
- scheduleInterval An amount of time, in minutes, that you want to have as a minimum between copy-on-write operation. It is possible for you to create a schedule in which you have overlapping copy-on-write operations because of
the duration a copy operation. You can make sure that you have time between copy-on-write operations by using this option. The maximum value for the scheduleInterval option is 1440 minutes. An example of this option is scheduleInterval=180.

- endDate A specific date on which you want to stop creating an Enhanced FlashCopy logical drive and end the copy-on-write operations. The format for entering the date is MM:DD:YY. An example of this option is endDate=11:26:11.
- timesPerDay The number of times that you want the schedule to run in a day. An example of this option is timesPerDay=4.
- timeZone Use this parameter to define the time zone in which the storage subsystem is operating. You can define the time zone in one of two ways:
  - GMT±HH:MM The time zone offset from GMT. Enter the offset in hours and minutes. For example GMT-06:00 is the central time zone in the United States.
  - Text string Standard time zone text strings. For example: "USA/Chicago" or "Germany/Berlin". Time zone text strings s are case sensitive. If you enter an incorrect text string, GMT time is used. Enclose the text string in double quotation marks.

The code string for defining a schedule is similar to these examples:

```
enableSchedule=true schedule startTime=14:27
enableSchedule=true schedule scheduleInterval=180
enableSchedule=true schedule timeZone=GMT-06:00
enableSchedule=true schedule timeZone="USA/Chicago"
```

If you also use the scheduleInterval option, the firmware will choose between the timesPerDay option and the scheduleInterval option by selecting the lowest value of the two options. The firmware calculates an integer value for the scheduleInterval option by dividing 1440 by a the scheduleInterval option value that you set. For example, 1440/180 = 8. The firmware then compares the timesPerDay integer value with the calculated scheduleInterval integer value and uses the smaller value.

To remove a schedule, use the delete logicaldrive command with the schedule parameter. The delete logicaldrive command with the schedule parameter deletes only the schedule, not the Enhanced FlashCopy logical drive.

When performing a rollback in a consistency group, the default operation is to rollback all members of the consistency group. If a rollback cannot be started successfully for all of the members in the consistency group, the rollback fails and has no effect. The Enhanced FlashCopy image is not rolled back.

# Create Consistency Group Enhanced FlashCopy

This command creates a logical mapping from a consistency group Enhanced FlashCopy logical drive to a host or a host group.

#### **Syntax**

create mapping cgEnhancedFlashCopyLogicalDrive="enhancedFlashCopyLogicalDriveName"
(host="hostName" | hostGroup=("hostGroupName" | defaultGroup)

# Parameters

Parameter	Description
cgEnhancedFlashCopyLogicalDrive	The name of the consistency group Enhanced FlashCopy logical drive for which you want to create a logical mapping. Enclose the consistency group Enhanced FlashCopy logical drive name in double quotation marks (" ").
host	The name of a host to which you want to create a logical mapping. Enclose the host name in double quotation marks (" ").
hostGroup	The name of a host group to which you want to create a logical mapping. Enclose the host group name in double quotation marks (" "). If you use the defaultGroup keyword, do not enclose it in quotation marks.

#### Notes

A host is a computer that is attached to the storage subsystem and accesses the logical drives on the storage subsystem through the host ports. You can define specific mappings to an individual host. You also can assign the host to a host group that shares access to one or more logical drives.

A host group is an optional topological element that you can define if you want to designate a collection of hosts that share access to the same logical drives. The host group is a logical entity. Define a host group only if you have two or more hosts that share access to the same logical drives.

# Create Consistency Group Enhanced FlashCopy Image

This command creates a new Enhanced FlashCopy image for each base logical drive that is a member of an Enhanced FlashCopy consistency group.

#### **Syntax**

create cgEnhancedFlashCopyImage consistencyGroup="consistencyGroupName"

#### **Parameters**

Parameter	Description
consistencyGroup	The name of the consistency group for which you are creating Enhanced FlashCopy images. Enclose the consistency group name in double quotation marks (" ").

#### Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters. The command causes all pending I/O operations to each base logical drive that is a member of the consistency group to be drained and suspended before creating the Enhanced FlashCopy images. If the creation of all of the Enhanced FlashCopy images cannot be completed successfully for all of the consistency group members, the operation fails and new Enhanced FlashCopy images are not created.

Normally, all members of an Enhanced FlashCopy consistency group will have the same number of Enhanced FlashCopy images. When you add a new member to an Enhanced FlashCopy consistency group, that new member lacks the Enhanced FlashCopy images that were previously created on the established members of the Enhanced FlashCopy consistency group. This is not an error condition. Requests for deletion or rollback of Enhanced FlashCopy images that exist on only a subset of the Enhanced FlashCopy consistency group members affects only those members for which the specified Enhanced FlashCopy image actually exists.

# Create Consistency Group Enhanced FlashCopy Logical Drive

This command creates an Enhanced FlashCopy logical drive of specific images in the base logical drives in a consistency group. You can select one base logical drive or more than one base logical drives from the consistency group to include in the Enhanced FlashCopy logical drive. When you create an Enhanced FlashCopy logical drive of a consistency group you are creating a logical drive with contents that you can view.

# Syntax With User Specified Base Logical Drives

create cgEnhancedFlashCopyLogicalDrive userLabel="cgLogicalDriveName"
cgEnhancedFlashCopyImageID="enhancedFlashCopyCGID:imageID"
members=(baseVol1:repos\_XXXX ... baseVoln:repos\_YYYY)

# Syntax When Setting the Consistency Group Enhanced FlashCopy Logical Drive to read-only

create cgEnhancedFlashCopyLogicalDrive userLabel="cgLogicalDriveName"
cgEnhancedFlashCopyImageID="enhancedFlashCopyCGID:imageID"
readOnly

# Syntax When Setting the Repository Full Limit

create cgEnhancedFlashCopyLogicalDrive userLabel="cgLogicalDriveName"
cgEnhancedFlashCopyImageID="enhancedFlashCopyCGID:imageID"
repositoryFullLimit=percentValue

Parameter	Description
userLabel	The name that you want to give the consistency group Enhanced FlashCopy logical drive that you are creating. Enclose the consistency group Enhanced FlashCopy logical drive name in double quotation marks (" ").
cgEnhancedFlashCopyImageID	The identifier of the Enhanced FlashCopy image from the Enhanced FlashCopy consistency group that you want to include in the consistency group Enhanced FlashCopy logical drive. Enclose the Enhanced FlashCopyImage identifier in double quotation marks (" ").

Parameter	Description
members	The identifier of one base logical drive or more than one base logical drives that you want to add. The members identifier is comprised of the base logical drive name concatenated with the repository logical drive name. You must use a colon (:) between the two names. Enclose all of the member identifiers in parenthesizes. If you enter more than one member separated the members with a space. When you do not use the members parameter, all of the members are automatically added to the new consistency group Enhanced ElsebCony logical drive
repositoryFullLimit	The percentage of repository capacity at which the consistency group Enhanced FlashCopy repository logical drive is nearing full. Use integer values. For example, a value of 70 means 70 percent.
readOnly	The setting to determine whether you can write to the Enhanced FlashCopy logical drive or only read from the Enhanced FlashCopy logical drive. To write to the Enhanced FlashCopy logical drive, do not include this parameter. To prevent writing to the Enhanced FlashCopy logical drive, include this parameter.

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

If you do not specify the repositoryVolType or readOnly parameters, the storage management software selects the repositories for the consistency group Enhanced FlashCopy logical drive. If the subsystem or disk pool where the base logical drive resides does not have enough space, this command fails.

The create cgEnhancedFlashCopyLogicalDrive command has unique forms that are explained by these examples:

• Creating a read/write consistency group Enhanced FlashCopy logical drive on an Enhanced FlashCopy consistency group named "enhancedFlashCopyCG1" that has three memebers cgm1, cgm2, and cgm3. The repository logical drives already exist and are selected by the user in this command.

create cgEnhancedFlashCopyLogicalDrive userLabel="cgEnhancedFlashCopyVol1"
cgEnhancedFlashCopyImageID="enhancedFlashCopyCG1:oldest"
members=(cgm1:repos1\_0010 cgm2:repos2\_0011 cgm3:repos3\_0007);

Note the use of the colon (:) in the name of the Enhanced FlashCopy image to be included in the consistency group Enhanced FlashCopy logical drive. The colon is a delimiter that separates the name of the Enhanced FlashCopy logical drive from a particular Enhanced FlashCopy image that you might want to use. You can use one of these options following the colon:

- An integer value that is the actual sequence number of the Enhanced FlashCopy image.
- newest Use this option when you want to show the latest consitency group Enhanced FlashCopy image.
- oldest Use this option when you want to show the earliest Enhanced FlashCopy image created.

The use of the colon following the names of the members of the Enhanced FlashCopy consistency group acts define the mapping between the member and a repository logical drive. For example, in cgm1:repos1\_10, member cgm1 maps to repository logical drive repos1\_0010.

• Creating a read/write consistency group Enhanced FlashCopy logical drive on an Enhanced FlashCopy consistency group named "enhancedFlashCopyCG1" of only members cgm1 and cgm2:

create cgEnhancedFlashCopyLogicalDrive userLabel="cgEnhancedFlashCopyVol2"
cgEnhancedFlashCopyImageID="enhancedFlashCopyCG1:14214"
members=(cgm1:repos\_1000,cgm2:repos\_1001);

• Creating a read-only consistency group Enhanced FlashCopy logical drive on an Enhanced FlashCopy consistency group named enhancedFlashCopyCG1 that has three members: cgm1, cgm2, and cgm3:

create cgEnhancedFlashCopyLogicalDrive userLabel="cgEnhancedFlashCopyVol3"
cgEnhancedFlashCopyImageID="enhancedFlashCopyCG1:oldest" readOnly;

• Creating a consistency group Enhanced FlashCopy logical drive that has a repository full limit set to 60 percent on an Enhanced FlashCopy consistency group named enhancedFlashCopyCG1 that has three members cgm1, cgm2, and cgm3:

create cgEnhancedFlashCopyLogicalDrive userLabel="cgEnhancedFlashCopyVol3"
cgEnhancedFlashCopyImageID="enhancedFlashCopyCG1:oldest"
repositoryFullLimit=60;

• Creating a read/write consistency group Enhanced FlashCopy logical drive with automatic repository selection on an Enhanced FlashCopy consistency group named enhancedFlashCopyCG1 that has three members cgm1, cgm2, and cgm3:

createc gEnhancedFlashCopyLogicalDriveu serLabel="cgEnhancedFlashCopyVol4"c
gEnhancedFlashCopyImageID="enhancedFlashCopyCG1:oldest";

# Create Consistency Group Enhanced FlashCopy Logical Drive Mapping

This command creates a logical mapping from a consistency group Enhanced FlashCopy logical drive to a host or a host group.

# **Syntax**

create mapping cgEnhancedFlashCopyLogicalDrive="enhancedFlashCopyLogicalDriveName"
(host="hostName" | hostGroup=("hostGroupName" | defaultGroup

### Parameters

Parameter	Description
cgEnhancedFlashCopyLogicalDrive	The name of the consistency group Enhanced FlashCopy logical drive for which you want to create a logical mapping. Enclose the consistency group Enhanced FlashCopy logical drive name in double quotation marks (" ").
host	The name of a host to which you want to create a logical mapping. Enclose the host name in double quotation marks (" ").
hostGroup	The name of a host group to which you want to create a logical mapping. Enclose the host group name in double quotation marks (" "). If you use the <b>defaultGroup</b> keyword, do not enclose it in quotation marks.

#### Notes

A host is a computer that is attached to the storage subsystem and accesses the logical drives on the storage subsystem through the host ports. You can define specific mappings to an individual host. You also can assign the host to a host group that shares access to one or more logical drives.

A host group is an optional topological element that you can define if you want to designate a collection of hosts that share access to the same logical drives. The host group is a logical entity. Define a host group only if you have two or more hosts that share access to the same logical drives.

# **Create Disk Pool**

This command creates a new disk pool based on the specified parameters. You can create the disk pool by entering either a list of drives or a type of drive that you want to use for the disk pool.

#### Syntax

```
create diskPool
(drives=(enclosureID1,drawerID1,slotID1 ... enclosureIDN,drawerIDN,slotIDN) |
driveType=(fibre|sas|sata))
userLabel="diskPoolName"
[driveCount=driveCountValue |
warningThreshold=(warningThresholdValue|default) |
criticalThreshold=(criticalThresholdValue|default) |
criticalPriority=(highest|high|medium|low|lowest) |
backgroundPriority=(highest|high|medium|low|lowest) |
degradedPriority=(highest|high|medium|low|lowest) |
securityType=(none|capable|enabled) |
driveMediaType=(hdd | ssd | allMedia | unknown) |
T10PI=(none|enabled)]
```

Parameter	Description
drives	The drives that you want to assign to the disk pool that you want to create. For high capacity drive enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for each drive that you assign to the disk pool. For low-capacity drive expansion enclosures, specify the enclosure ID value and the slot ID value for each drive that you assign to the disk pool. For low-capacity drive expansion enclosures, specify the enclosure ID value and the slot ID value for each drive that you assign to the disk pool. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID values, the drawer ID values, and the slot ID values in parentheses.
driveType	The type of drive that you want to use in the disk pool. You cannot mix drive types. You must use this parameter when you have more than one type of drive in your storage subsystem. Valid drive types are: • <i>fibre</i> • <i>SATA</i>
	• <i>SAS</i> If you do not specify a drive type, the command defaults to <i>fibre</i> .
userLabel	The name that you want to give the new disk pool. Enclose the disk pool name in double quotation marks (" ").
driveCount	The <b>driveCount</b> parameter limits the disk pool candidates to the given number. When using this parameter, the minimum value that you can enter is 11. Use this parameter only when the <b>driveType</b> parameter is specified.
warningThreshold	The percentage of storage capacity at which you receive a warning alert that the disk pool is nearing full. Use integer values. For example, a value of 60 means 60 percent. The default value is 50. Setting this parameter to 0 disables warning alerts.
	If you set this to default, the warning alert threshold value is determined by the controller firmware.

Parameter	Description
criticalThreshold	The percentage of storage capacity at which you receive a critical alert that the disk pool is nearing full. Use integer values. For example, a value of 70 means 70 percent.
	The <i>default</i> value is 85.
	Setting this parameter to $0$ disables both warning alerts and critical alerts.
	If you set this to <i>default</i> , the critical alert threshold value is determined by the controller firmware.
backgroundPriority	The priority for background operations on the disk pool.
	Valid values are <i>highest, high, medium, low,</i> and <i>lowest</i> . The default value is <i>low</i> .
degradedPriority	The priority for degraded activities on the disk pool. For example, disk pool reconstruction after one drive failures.
	Valid values are <i>highest, high, medium, low,</i> and <i>lowest</i> . The default value is <i>high</i> .
securityType	The setting to specify the security level when creating the disk pool. All logical drive candidates for the disk pool will have the specified security type. These settings are valid:
	• <i>none</i> - The logical drive candidates are not secure.
	<ul> <li><i>capable</i> - The logical drive candidates are capable of having security set, but security has not been enabled.</li> </ul>
	• <i>enabled</i> - The logical drive candidates have security enabled.
	The default value is <i>none</i>

Parameter	Description
driveMediaType	The type of drive media that you want to use for the disk pool
	You must use this parameter when you have more than one type of drive media in your storage subsystem.
	Valid drive media are:
	• <i>hdd</i> - Use this option when you have hard drives.
	• <i>ssd</i> - Use this option when you have solid state disks.
	• <i>unknown</i> - Use if you are not sure what types of drive media are in the drive expansion enclosure.
	• <i>allMedia</i> - Use this option when you want to use all types of drive media that are in the drive expansion enclosure.
	The default value is <i>hdd</i> .
T10PI	The setting to specify that a disk pool has T10 PI (Protection Information) protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the disk pool. These settings are valid:
	• <i>none</i> - The disk pool does not have T10 PI (Protection Information) protection.
	• <i>enabled</i> - The disk pool has T10 PI (Protection Information) protection. The disk pool supports protected information and is formatted with protection information enabled.
	The default value is <i>none</i>

Each disk pool name must be unique. You can use any combination of alphanumeric characters, underscore (\_), hyphen(-), and pound (#) for the user label. User labels can have a maximum of 30 characters.

If the parameters you specify cannot be satisfied by any of the available candidate drives, the command fails. Normally, all drives that match the quality of service attributes will be returned as the top candidates. However, if you specifying a drive list, some of the available drives returned as candidates might not match the quality of service attributes.

If you do not specify a value for an optional parameter, a default value is assigned.

# Drives

When you use the **driveType** parameter, all of the unassigned drives that are of that drive type are used to create the disk pool. If you want to limit the number of drives found by the **driveType** parameter in the disk pool, you can specify the

number of drives using the **driveCount** parameter. You can use the **driveCount** parameter only when you use the **driveType** parameter.

The **drives** parameter supports both high-capacity drive expansion enclosures and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need only specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

If you enter specifications for a high capacity drive expansion enclosure, but a drive expansion enclosure is not available, the storage management software returns an error message.

#### **Disk Pool Alert Thresholds**

Each disk pool has two progressively severe levels of alerts to inform users when the storage capacity of the disk pool is approaching full. The threshold for an alert is a percent of the used capacity to the total usable capacity in the disk pool. The alerts are:

- Warning This is the first level of alert that the used capacity in a disk pool is approaching full. When the threshold for the warning alert is reached a Needs Attention condition is generated and an event is posted to the storage management software. The warning threshold is superseded by the critical threshold. The default warning threshold is 50 percent.
- Critical This is the most severe level of alert that the used capacity in a disk pool is approaching full. When the threshold for the critical alert is reached a Needs Attention condition is generated and an event is posted to the storage management software. The warning threshold is superseded by the critical threshold. The default critical warning threshold is 85 percent.

To be effective, the value for a warning alert must always be less than the value for a critical alert. If the value for the warning alert is the same as the value for a critical alert, only the critical alert is sent.

#### **Disk Pool Background Operations**

Disk pools support these background operations:

- Reconstruction
- Instant Availability Format (IAF)
- Format
- Dynamic Capacity Expansion (DCE)
- Dynamic Volume Expansion (DVE) (For disk pools, DVE is actually not a background operation, but DVE is supported as a synchronous operation).

Disk pools do not queue background commands. You can start several background commands sequentially, but starting more than one background operation at a time

delays the completion of commands that you started previously. The relative priority levels for the supported background operations are:

- 1. Reconstruction
- 2. Format
- **3**. IAF
- 4. DCE

# Security Type

The **securityType** parameter is valid for drives that are capable of full disk encryption (FDE). With FDE, the controller firmware can create a key and activate the Full Disk Encryption feature. The Full Disk Encryption feature encrypts data as the data is written to the drive and decrypts the data as the data is read from the drive. Without the key created by the controller, the data written to the drive is inaccessible.

Before you can set the **securityType** parameter to **capable** or **enabled**, you must create a storage subsystem security key. Use the create **storageSubsystem securityKey** command to create a storage subsystem security key. These commands are related to the security key:

- create storageSubsystem securityKey
- enable subsystem [subsystemName] security
- export storageSubsystem securityKey
- import storageSubsystem securityKey
- set storageSubsystem securityKey
- start secureErase (drive | drives)

# Create Enhanced FlashCopy Group

This command creates a new Enhanced FlashCopy group and the associated repository logical drive. An Enhanced FlashCopy group contains a sequence of Enhanced FlashCopy images of an associated base logical drive. An Enhanced FlashCopy group has a repository logical drive that is used to save data for all of the Enhanced FlashCopy images that are part of the Enhanced FlashCopy group.

# **Syntax**

```
create enhancedFlashCopyGroup userLabel="enhancedFlashCopyGroupName"
sourceLogicalDrive="logicaldriveName"
(repositoryLogicalDrive="repos_xxxx" |
repositoryLogicalDrive=(subsystemName [capacity=capacityValue])
repositoryLogicalDrive=(diskPoolName [capacity=capacityValue]))
[repositoryFullPolicy=(failBaseWrites | purgeEnhancedFlashCopyImages) |
rollbackPriority=(highest | high | medium | low | lowest) |
repositoryFullLimit=percentValue |
autoDeleteLimit=numberOfEnhancedFlashCopyImages |
enableSchedule=(TRUE | FALSE)
schedule=(immediate | flashcopySchedule)]
```

#### DITA

DITA

Parameter	Description
userLabel	The name that you want to give the new Enhanced FlashCopy group. Enclose the Enhanced FlashCopy group identifier in double quotation marks (" ").
sourceLogicalDrive	The name of the logical drive that you want to use as the source for your Enhanced FlashCopy images. Enclose the source logical drive name in double quotation marks (" ").
repositoryLogicalDrive	The name of the repository logical drive that will contain the Enhanced FlashCopy group.
	You have two options for defining the name of a repository logical drive:
	<ul><li>Use an existing repository logical drive: name</li><li>Create a new repository logical drive when you run this command</li></ul>
	The name of an existing repository logical drive is comprised of two parts: • The term <i>repos</i>
	• A four digit numerical identifier that you assign to the repository logical drive name
	Enclose the name of the existing repository logical drive in double quotation marks (" ").
	If you want to create a new repository logical drive when you run this command you must enter the name of either a an subsystem or a disk pool in which you want the repository logical drive. Optionally, you can also define the capacity of the repository logical drive. If you want to define the capacity you can use these values:
	<ul> <li>An integer value that represents a percentage of the base logical drive capacity</li> </ul>
	• An decimal fraction value that represents a percentage of the base logical drive capacity
	• A specific size for the repository logical drive. Size is defined in units of <i>bytes, KB, MB, GB,</i> or <i>TB</i> .
	If you do not use the capacity option, the storage management software sets the capacity to 20 percent of the base logical drive capacity.
	Enclose the name of the new repository logical drive in parenthesis ( ).
repositoryFullPolicy	Defines how Enhanced FlashCopy image processing continues if the Enhanced FlashCopy group repository logical drive is full. You can choose to fail I/O writes to the base logical drive ( <i>failBaseWrites</i> ) or delete (purge) the Enhanced FlashCopy images ( <i>purgeEnhancedFlashCopyImages</i> ) in the repository logical drive. The <i>purgeEnhancedFlashCopyImages</i> option deletes the oldest Enhanced FlashCopy images to free up space. The default action is <i>purgeEnhancedFlashCopyImages</i> .

Description
Determines whether system resources should be allocated to the rollback operation at the expense of system performance. A value of high indicates that the rollback operation is prioritized over all other host I/O. A value of low indicates that the rollback operation should be performed with minimal impact to host I/O. The default value is medium.
The percentage of repository capacity at which you receive a warning that the Enhanced FlashCopy group repository logical drive is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.
The maximum number of Enhanced FlashCopy images that you want to keep if you have selected to purge the Enhanced FlashCopy images for the repository full policy. The storage management software deletes the oldest Enhanced FlashCopy images first and continues until the auto delete limit is reached. The default value is 32.
Use this parameter to turn on or to turn off the ability to schedule an Enhanced FlashCopy operation. To turn on Enhanced FlashCopy scheduling, set this parameter to TRUE. To turn off Enhanced FlashCopy scheduling, set this parameter to FALSE.
Use this parameter to schedule an Enhanced FlashCopy group operation. You can use one of these options for setting a schedule for an Enhanced FlashCopy group operation: • immediate • startDate • scheduleDay • startTime • scheduleInterval • endDate • noEndDate • timesPerDay • timeZone See the Notes section for information explaining how to use

Each Enhanced FlashCopy group name must be unique. You can use any combination of alphanumeric characters, underscore (\_), hyphen(-), and pound (#) for the user label. User labels can have a maximum of 30 characters.

To create an Enhanced FlashCopy group, you must have an associated repository logical drive in which you store the Enhanced FlashCopy images. You can either use an existing repository logical drive or create a new repository logical drive. You can create the repository logical drive when you create the Enhanced FlashCopy group. An Enhanced FlashCopy group repository logical drive is an expandable logical drive that is structured as a concatenated collection of up to 16 standard logical drive entities. Initially, an expandable repository logical drive has only a single element. The capacity of the expandable repository logical drive is exactly that of the single element. You can increase the capacity of an expandable repository logical drive by attaching additional standard logical drives to it. The composite expandable repository logical drive capacity then becomes the sum of the capacities of all of the concatenated standard logical drives.

An Enhanced FlashCopy group has a strict ordering of Enhanced FlashCopy images based on the time that each Enhanced FlashCopy image is created. An Enhanced FlashCopy image that is created after another Enhanced FlashCopy image is a successor relative to that other Enhanced FlashCopy image. An Enhanced FlashCopy image that is created before another Enhanced FlashCopy image is a predecessor relative to that other one.

An Enhanced FlashCopy group repository logical drive must satisfy a minimum capacity requirement that is the sum of the following:

- 32 MB to support fixed overhead for the Enhanced FlashCopy group and for copy-on-write processing.
- Capacity for rollback processing, which is 1/5000th of the capacity of the base logical drive.

The minimum capacity is enforcement by the controller firmware and the storage management software.

When a you first create an Enhanced FlashCopy group, it does not contains any Enhanced FlashCopy images. When you create Enhanced FlashCopy images, you add the Enhanced FlashCopy images to an Enhanced FlashCopy group. Use the create enhancedFlashCopyImage command to create Enhanced FlashCopy images and add the Enhanced FlashCopy images to an Enhanced FlashCopy group.

An Enhanced FlashCopy group can have one of these states:

- Optimal The Enhanced FlashCopy group is operating normally.
- Full The Enhanced FlashCopy group repository is full. Additional copy-on-write operations can not be performed. This state is possible only for Enhanced FlashCopy groups that have the Repository Full policy set to Fail Base Writes. Any Enhanced FlashCopy group in a Full state causes a Needs-Attention condition to be posted for the storage subsystem.
- Over Threshold The Enhanced FlashCopy group repository logical drive usage is at or beyond its alert threshold. Any Enhanced FlashCopy group in this state causes a Needs-Attention condition to be posted for the storage subsystem.
- Failed The Enhanced FlashCopy group has encountered a problem that has made all Enhanced FlashCopy images in the Enhanced FlashCopy group unusable. For example, certain types of repository logical drive failures can cause a Failed state. To recover from a Failed state use the revive enhancedFlashCopyGroup command.

#### Automatic Enhanced FlashCopy Image Deletion

You can configure each Enhanced FlashCopy group to automatically delete Enhanced FlashCopy images by using the **autoDeleteLimit** parameter. Automatically deleting the Enhanced FlashCopy images enables you to avoid having to routinely, manually delete the images that you do not want and that might prevent the creation of future Enhanced FlashCopy images because the repository logical drive is full. When you use the **autoDeleteLimit** parameter it causes the storage management software to automatically delete Enhanced FlashCopy images. The storage management software deletes Enhanced FlashCopy images until it reaches a number of Enhanced FlashCopy images that is equal to the number that you enter with **autoDeleteLimit** parameter. When new Enhanced FlashCopy images are added to the repository logical drive, the storage management software deletes the oldest Enhanced FlashCopy images until the **autoDeleteLimit** parameter number is reached.

# Scheduling Enhanced FlashCopies

The **enableSchedule** parameter and the schedule parameter provide a way for you to schedule creating Enhanced FlashCopy images for an Enhanced FlashCopy group. Using these parameters, you can schedule Enhanced FlashCopies daily, weekly, or monthly (by day or by date). The **enableSchedule** parameter turns on or turns off the ability to schedule Enhanced FlashCopies. When you enable scheduling, you use the **schedule** parameter to define when you want the Enhanced FlashCopies to occur.

This list explains how to use the options for the **schedule** parameter:

- immediate As soon as you enter the command, an Enhanced FlashCopy image is created, and a copy-on-write operation begins.
- startDate A specific date on which you want to create an Enhanced FlashCopy image and perform a copy-on-write operation. The format for entering the date is MM:DD:YY. If you do not provide a start date, the current date is used. An example of this option is startDate=06:27:11
- scheduleDay A day of the week on which you want to create an Enhanced FlashCopy image and perform a copy-on-write operation. You can enter these values: monday, tuesday, wednesday, thursday, friday, saturday, and sunday. An example of this option is scheduleDay=wednesday.
- startTime The time of a day that you want to create an Enhanced FlashCopy image and start performing a copy-on-write operation. The format for entering the time is HH:MM, where HH is the hour and MM is the minute past the hour. Use a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is startTime=14:27.
- scheduleInterval An amount of time, in minutes, that you want to have as a minimum between copy-on-write operations. You can possibly create a schedule in which you have overlapping copy-on-write operations because of the duration a copy operation. You can make sure that you have time between copy-on-write operations by using this option. The maximum value for the scheduleInterval option is 1440 minutes. An example of this option is scheduleInterval=180.
- endDate A specific date on which you want to stop creating an Enhanced FlashCopy image and end the copy-on-write operation. The format for entering the date is MM:DD:YY. An example of this option is endDate=11:26:11.
- noEndDate Use this option if you do not want your scheduled copy-on-write operation to end. If you later decide to end the copy-on-write operations you must re-enter the set enhancedFlashCopyGroup command and specify an end date.
- timesPerDay The number of times that you want the schedule to run in a day. An example of this option is timesPerDay=4.
- timeZone Use this parameter to define the time zone in which the storage subsystem is operating. You can define the time zone in one of two ways:
  - GMT±HH:MM The time zone offset from GMT. Enter the offset in hours and minutes. For example GMT-06:00 is the central time zone in the United States.

 Text string - Standard time zone text strings. For example: "USA/Chicago" or "Germany/Berlin". Time zone text strings s are case sensitive. If you enter an incorrect text string, GMT time is used. Enclose the text string in double quotation marks.

The code string for defining a schedule is similar to these examples:

enableSchedule=true schedule startTime=14:27
enableSchedule=true schedule scheduleInterval=180
enableSchedule=true schedule timeZone=GMT-06:00
enableSchedule=true schedule timeZone="USA/Chicago"

If you also use the scheduleInterval option, the firmware chooses between the timesPerDay option and the scheduleInterval option by selecting the lowest value of the two options. The firmware calculates an integer value for the scheduleInterval option by dividing 1440 by the scheduleInterval option value that you set. For example, 1440/180 = 8. The firmware then compares the timesPerDay integer value with the calculated scheduleInterval integer value and uses the smaller value.

To remove a schedule, use the delete logicaldrive command with the **schedule** parameter. The delete logicaldrive command with the **schedule** parameter deletes only the schedule, not the Enhanced FlashCopy logical drive.

The **enableSchedule** parameter and the **schedule** parameter provide a way for you to schedule automatic Enhanced FlashCopy. Using these parameters, you can schedule Enhanced FlashCopy daily, weekly, or monthly (by day or by date). The **enableSchedule** parameter turns on or turns off the ability to schedule Enhanced FlashCopy. When you enable scheduling, you use the **schedule** parameter to define when you want the Enhanced FlashCopy to occur.

This list explains how to use the options for the schedule parameter:

- immediate As soon as you enter the command, an Enhanced FlashCopy logical drive is created, and a copy-on-write operation begins.
- startDate A specific date on which you want to create an Enhanced FlashCopy logical drive and perform a copy-on-write operation. The format for entering the date is MM:DD:YY. If you do not provide a start date, the current date is used. An example of this option is startDate=06:27:11.
- scheduleDay A day of the week on which you want to create an Enhanced FlashCopy logical drive and perform a copy-on-write operation. You can enter these values: monday, tuesday, wednesday, thursday, friday, saturday, and sunday. An example of this option is scheduleDay=wednesday
- startTime The time of a day that you want to create an Enhanced FlashCopy logical drive and start performing a copy-on-write operation. The format for entering the time is HH:MM, where HH is the hour and MM is the minute past the hour. Use a 24- hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is startTime=14:27.
- scheduleInterval An amount of time, in minutes, that you want to have as a minimum between copy-on-write operation. You can possibly create a schedule in which you have overlapping copy-on-write operations because of the duration a copy operation. You can make sure that you have time between copy-on-write operations by using this option. The maximum value for the scheduleInterval option is 1440 minutes. An example of this option is scheduleInterval=180.

- endDate A specific date on which you want to stop creating an Enhanced FlashCopy logical drive and end the copy-on-write operation. The format for entering the date is MM:DD:YY. An example of this option is endDate=11:26:11
- noEndDate Use this option if you do not want your scheduled copy-on-write operation to end. If you later decide to end the copy-on-write operations you must re-enter the set enhancedFlashCopyGroup command and specify an end date.
- timesPerDay The number of times that you want the schedule to run in a day. An example of this option is timesPerDay=4.
- timeZone Use this parameter to define the time zone in which the storage subsystem is operating. You can define the time zone in one of two ways:
  - GMT±HH:MM The time zone offset from GMT. Enter the offset in hours and minutes. For example GMT-06:00 is the central time zone in the United States.
  - Text string Standard time zone text strings. For example: "USA/Chicago" or "Germany/Berlin". Time zone text strings are case sensitive. If you enter an incorrect text string, GMT time is used. Enclose the text string in double quotation marks.

The code string for defining a schedule is similar to these examples:

enableSchedule=true schedule startTime=14:27
enableSchedule=true schedule scheduleInterval=180
enableSchedule=true schedule timeZone=GMT-06:00
enableSchedule=true schedule timeZone="USA/Chicago"

If you also use the scheduleInterval option, the firmware chooses between the timesPerDay option and the scheduleInterval option by selecting the lowest value of the two options. The firmware calculates an integer value for the scheduleInterval option by dividing 1440 by the scheduleInterval option value that you set. For example, 1440/180 = 8. The firmware then compares the *timesPerDay* integer value with the calculated *scheduleInterval integer* value and uses the smaller value.

To remove a schedule, use the command set enhancedFlashCopyGroup ["enhancedFlashCopyGroupName"] enableSchedule=FALSE. This command deletes only the schedule, not the Enhanced FlashCopy group.

# Create Enhanced FlashCopy Image

This command creates a new Enhanced FlashCopy image in one or more existing Enhanced FlashCopy groups. Before you can create an Enhanced FlashCopy image, you must first have at least one Enhanced FlashCopy group into which you can place the Enhanced FlashCopy image. To create an Enhanced FlashCopy group use the create enhancedFlashCopyGroup command.

#### Syntax

create enhancedFlashCopyImage (enhancedFlashCopyGroup="enhancedFlashCopyGroupName" |
enhancedFlashCopyGroupS=("enhancedFlashCopyGroupName1"... "enhancedFlashCopyGroupNamen"))

# Parameters

Parameter	Description
enhancedFlashCopyGroup or enhancedFlashCopyGroups	The name of the Enhanced FlashCopy group that will contain the Enhanced FlashCopy image. Enclose the name of the Enhanced FlashCopy group in double quotation marks (" ").
	If you enter more than one Enhanced FlashCopy group name, enclose each of the Enhanced FlashCopy group names in double quotation marks (" "). Enclose all of the Enhanced FlashCopy group names in parenthesizes.

#### Notes

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

An Enhanced FlashCopy image is a logical point-in-time image of the contents of an associated base logical drive. The Enhanced FlashCopy image is created instantly and records the state of the base logical drive at that moment. Every Enhanced FlashCopy image is created in the context of exactly one Enhanced FlashCopy group. An Enhanced FlashCopy group is a sequence of Enhanced FlashCopy images of the associated base logical drive. An Enhanced FlashCopy group has one repository logical drive used to save all of the data from the Enhanced FlashCopy images. The Enhanced FlashCopy images in an Enhanced FlashCopy group have a specific order. The specific order of the Enhanced FlashCopy images enables you to manage the Enhanced FlashCopy images, such as restoring a specific Enhanced FlashCopy image to the base logical drive or deleting the Enhanced FlashCopy images that you no longer need.

The result of creating an Enhanced FlashCopy image of a consistency group is an Enhanced FlashCopy image of every member logical drive of the consistency group.

# Create Enhanced FlashCopy Logical Drive

This command creates an Enhanced FlashCopy logical drive with read-write capabilities for Enhanced FlashCopy images of a base logical drive. You can map the Enhanced FlashCopy logical drive to a host and all of the host writes reside in the repository logical drive associated with the Enhanced FlashCopy logical drive. You can assign the new Enhanced FlashCopy logical drive to an existing repository logical drive, or you can create a new repository logical drive in an subsystem or disk pool.

**Note:** You cannot use this command for an Enhanced FlashCopy image that is used in online logical drive copy.

#### Syntax

```
create enhancedFlashCopyLogicalDrive userLabel="enhancedFlashCopyLogicalDriveName"
enhancedFlashCopyImageID="enhancedFlashCopyImageID"
[repositoryLogicalDrive="repos_xxxx" |
repositoryFullLimit=percentValue]
```

# Parameter

Parameter	Description
userLabel	The name that you want to give to an Enhanced FlashCopy logical drive. Enclose the Enhanced FlashCopy logical drive name in double quotation marks (" ").
enhancedFlashCopyImageID	<ul> <li>The alphanumeric identifier of an Enhanced FlashCopy image that you want to add to the new Enhanced FlashCopy logical drive.</li> <li>The identifier of an Enhanced FlashCopy image is comprised of two parts:</li> <li>The name of the Enhanced FlashCopy group</li> <li>An identifier for the Enhanced FlashCopy image in the Enhanced FlashCopy group.</li> </ul>
	<ul> <li>The identifier for the Enhanced FlashCopy image can be one of these:</li> <li>An integer value that is the sequence number of the Enhanced FlashCopy group.</li> <li>NEWEST - Use this option when you want to show the latest Enhanced FlashCopy image created in the Enhanced FlashCopy group.</li> <li>OLDEST - Use this option when you want to show the earliest Enhanced FlashCopy image created in the Enhanced FlashCopy group.</li> <li>DLDEST - Use this option when you want to show the earliest Enhanced FlashCopy image created in the Enhanced FlashCopy image</li></ul>
	name in double quotation marks (" ").
repositoryLogicalDrive	The alphanumeric identifier of the repository logical drive has the Enhanced FlashCopy image. All repository identifiers have this form: repos_xxxx
	where xxxx represents a four digit numerical value. Enclose the Enhanced FlashCopy logical drive name in double quotation marks (" ").
repositoryFullLimit	The percentage of repository capacity at which you receive a warning that the Enhanced FlashCopy repository logical drive is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.

# Notes

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

The identifier of an Enhanced FlashCopy image has two parts separated by a colon (:):

- The name of the Enhanced FlashCopy group
- The identifier of the Enhanced FlashCopy image

For example, if you want to create an Enhanced FlashCopy logical drive named *enhancedFlashCopyData1* using the most recent Enhanced FlashCopy image in an Enhanced FlashCopy group that has the name **enhancedFlashCopyGroup1** with a maximum fill limit of 80 percent for the repository logical drive, you would use this command:

create enhancedFlashCopyLogicalDrive userLabel="enhancedFlashCopyDatal"
enhancedFlashCopyImageID="enhancedFlashCopyGroup1:newest"
repositoryLogicalDrive="repos\_1234" repositoryFullLimit=80;

The repository logical drive identifier is automatically created by the storage management software and the firmware when you create a new Enhanced FlashCopy group. You cannot rename the repository logical drive because renaming the repository logical drive will break the linkage with the Enhanced FlashCopy images.

# Create FlashCopy Logical Drive

This command creates a FlashCopy logical drive of a base logical drive. You can also use this command to create a new repository subsystem if one does not already exist, or if you would prefer a different repository subsystem. This command defines three ways to create a FlashCopy logical drive:

- · In a new repository subsystem created from user-defined drives
- In a new repository subsystem created from a user-defined number of drives
- In an existing repository subsystem

If you choose to define a number of disk drives, the controller firmware chooses which disk drives to use for the FlashCopy logical drive.

#### Syntax (User-Defined Number of Drives)

```
create flashcopyLogicalDrive baseLogicalDrive="baseLogicalDriveName"
(repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDrives=(enclosureID1,drawerID1,slotID1 ...
enclosureIDn,drawerIDn,slotIDn))
[repositorysubsystemUserLabel="repositorysubsystemName"
enclosureLossProtect=(TRUE | FALSE)
drawerLossProtect=(TRUE | FALSE)
freeCapacityArea=freeCapacityIndexNumber
userLabel="flashcopyLogicalDriveName"
warningThresholdPercent=percentValue
repositoryUserLabel="repositoryName"
repositoryUserLabel="repositoryName"
repositoryFullPolicy=(failBaseWrites | failFlashcopy)
enableSchedule=(TRUE | FALSE)
schedule=(immediate | flashcopySchedule)]
```

#### Syntax (Existing Repository subsystem)

create flashcopyLogicalDrive baseLogicalDrive="baseLogicalDriveName"
[repositorysubsystem="repositorysubsystemName"
repositoryUserLabel="repositoryName"
freeCapacityArea=freeCapacityIndexNumber
userLabel="flashcopyLogicalDriveName"
warningThresholdPercent=percentValue

repositoryPercentOfBase=percentValue
repositoryFullPolicy=(failBaseWrites | failFlashcopy)
enableSchedule=(TRUE | FALSE)
schedule=(immediate | flashcopySchedule)]

Parameter	Description
baseLogicalDrive	The name of the base logical drive from which you want to take a FlashCopy. Enclose the base logical drive name in double quotation marks ("").
repositoryRAIDLevel	Use this parameter when you create a new subsystem.
	The RAID level for the FlashCopy repository logical drive. Valid values are 1, 3, 5, or 6.
repositoryDrives	Use this parameter when you create a new subsystem.
	The disk drives that you want to assign to the FlashCopy repository subsystem.For high-capacity disk drive expansion enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for the disk drive. For low-capacity drive expansion enclosures, specify the enclosure ID value and the slot ID value for the disk drive. Enclosure ID values are $\theta$ to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID value, the drawer ID value, and the slot ID value in parentheses.
repositoryDriveCount	Use this parameter when you create a new subsystem.
	The number of unassigned disk drives that you want to use for the FlashCopy repository logical drive.
repositorysubsystemUserLabel	The alphanumeric identifer (including - and _) of the subsystem where the FlashCopy repository logical drive is located.
repositorysubsystem	The name of an existing subsystem where you want to place the repository logical drive. Use this parameter if you do not want to put the repository logical drive in the same subsystem as the base logical drive. The default is to use the same subsystem for both the base logical drive and the repository logical drive. Enclose the name of the repository subsystem in double quotation marks (" ").
UserLabel	The name that you want to give the FlashCopy logical drive. Enclose the FlashCopy logical drive name in double quotation marks ("").

Parameter	Description
enclosureLossProtect	The setting to enforce enclosure loss protection when you create the FlashCopy repository logical drive. To enforce enclosure loss protection, set this parameter to TRUE. The default value is FALSE.
drawerLossProtect	The setting to enforce drawer loss protection when you create the mirror repository logical drive. To enforce drawer loss protection, set this parameter to TRUE.The default value is FALSE.
driveMediaType	The type of drive medium that you want to use for the mirror repository logical drive. Valid drive media are these:
	<ul> <li>HDD – Use this option when you have hard drives in the drive expansion enclosure.</li> <li>SSD – Use this option when you have</li> </ul>
	solid state drives in the drive expansion enclosure.
	• <b>unknown</b> – Use if you are not sure what types of drive media are in the drive expansion enclosure.
	• <b>allMedia</b> – Use this option when you want to use all types of drive media that are in the drive expansion enclosure.
	Use this parameter when you use the <b>repositoryDriveCount</b> parameter.
	You must use this parameter when you have more than one type of drive media in your storage subsystem.
DriveType	The type of disk drives that you want to use for the FlashCopy repository logical drive. Valid disk drive types are SCSI, fibre, or SAS.
freeCapacityArea	The index number of the free space in an existing subsystem that you want to use to create the FlashCopy repository logical drive. Free capacity is defined as the free capacity between existing logical drives in an subsystem. For example, an subsystem might have the following areas: logical drive 1, free capacity, logical drive 2, free capacity, logical drive 3, free capacity. To use the free capacity following logical drive 2, you would specify: freeCapacityArea=2
	Run the <b>show subsystem</b> command to determine if a free capacity area exists.
warningThresholdPercent	The percentage of repository capacity at which you receive a warning that the FlashCopy repository logical drive is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 50.

Parameter	Description
repositoryPercentOfBase	The size of the FlashCopy repository logical drive as a percentage of the base logical drive. Use integer values. For example, a value of 40 means 40 percent. The default value is 20.
repositoryUserLabel	The name that you want to give to the FlashCopy repository logical drive. Enclose the FlashCopy repository logical drive name in double quotation marks ("").
repositoryFullPolicy	How you want FlashCopy processing to continue if the FlashCopy repository logical drive is full. You can choose to fail writes to the base logical drive (failBaseWrites) or fail the FlashCopy logical drive (failFlashCopy). The default value is failFlashCopy.
enableSchedule	Use this parameter to turn on or to turn off the ability to schedule a FlashCopy operation. To turn on FlashCopy scheduling, set this parameter to TRUE. To turn off FlashCopy scheduling, set this parameter to FALSE.
schedule	Use this parameter to schedule a FlashCopy operation. You can use one of these options for setting a schedule for a FlashCopy operation: immediate startDate scheduleDay startTime scheduleInterval noEndDate timesPerDay timeZone See the "Notes" section for information explaining how to use these options.

The logical drive that you are taking a FlashCopy of must be a standard logical drive in the storage subsystem. The maximum number of FlashCopy logical drives that you can create is one-half of the total number of logical drives that are supported by a controller.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

One technique for naming the FlashCopy logical drive and the FlashCopy repository logical drive is to add a hyphenated suffix to the original base logical drive name. The suffix distinguishes between the FlashCopy logical drive and the FlashCopy repository logical drive. For example, if you have a base logical drive

with a name of Engineering Data, the FlashCopy logical drive can have a name of Engineering Data-S1, and the FlashCopy repository logical drive can have a name of EngineeringData-R1.

If you do not choose a name for either the FlashCopy logical drive or the FlashCopy repository logical drive, the storage management software creates a default name by using the base logical drive name. An example of the FlashCopy logical drive name that the controllers might create is, if the base logical drive name is AAA and does not have a FlashCopy logical drive, the default FlashCopy logical drive name is AAA-1. If the base logical drive already has *n*-1 number of FlashCopy logical drive name that the controller might create is, if the base logical drive repository logical drives, the default name is AAA-*n*. An example of the FlashCopy repository logical drive name that the controller might create is, if the base logical drive name is AAA and does not have a FlashCopy repository logical drive, the default FlashCopy repository logical drive name is AAA and does not have a FlashCopy repository logical drive, the default FlashCopy repository logical drive name is AAA and does not have a FlashCopy repository logical drive, the default FlashCopy repository logical drive name is AAA and does not have a FlashCopy repository logical drive, the default FlashCopy repository logical drive name is AAA-R1. If the base logical drive already has *n*-1 number of FlashCopy repository logical drives, the default name is AAA-R1.

If you do not specify the unconfigured space or free space, the FlashCopy repository logical drive is placed in the same subsystem as the base logical drive. If the subsystem where the base logical drive resides does not have enough space, this command fails.

When you assign the disk drives, if you set the **enclosureLossProtect** parameter to TRUE and have selected more than one disk drive from any one enclosure, the storage subsystem returns an error. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs operations, but the subsystem that you create might not have enclosure loss protection.

When the controller firmware assigns the disk drives, if you set the **enclosureLossProtect** parameter to TRUE, the storage subsystem returns an error if the controller firmware cannot provide disk drives that result in the new subsystem having enclosure loss protection. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs the operation even if it means the subsystem might not have enclosure loss protection.

# **Enclosure Loss Protection and Drawer Loss Protection**

When you assign the drives, if you set the **enclosureLossProtect** parameter to TRUE and have selected more than one drive from any one enclosure, the storage subsystem returns an error. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs operations, but the subsystem that you create might not have enclosure loss protection.

When the controller firmware assigns the drives, if you set the **enclosureLossProtect** parameter to TRUE, the storage subsystem returns an error if the controller firmware cannot provide drives that result in the new subsystem having enclosure loss protection. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs the operation even if it means the subsystem might not have enclosure loss protection.

The **drawerLossProtect** parameter defines if data on a logical drive is accessible if a drawer fails. When you assign the drives, if you set the **drawerLossProtect** parameter to TRUE and select more than one drive from any one drawer, the storage subsystem returns an error. If you set the **drawerLossProtect** parameter to FALSE, the storage subsystem performs operations, but the subsystem that you create might not have drawer loss protection. If you have a storage configuration that includes a drive expansion enclosure that has drawers to hold the drives, follow these guidelines when configuring enclosure loss protection:

- If you set **enclosureLossProtect** to TRUE, then you must set **drawerLossProtect** to TRUE.
- If you set **enclosureLossProtect** to FALSE, then you can set **drawerLossProtect** to either TRUE or FALSE.

If you set **enclosureLossProtect** to TRUE and **drawerLossProtect** to FALSE, the storage subsystem returns an error.

# Scheduling FlashCopies

The **enableSchedule** parameter and the **schedule** parameter provide a way for you to schedule automatic FlashCopies . Using these parameters, you can schedule FlashCopies daily, weekly, or monthly (by day or by date). The **enableSchedule** parameter turns on or turns off the ability to schedule FlashCopies . When you enable scheduling, you use the schedule parameter to define when you want the FlashCopies to occur.

This list explains how to use the options for the schedule parameter:

- immediate As soon as you enter the command, a FlashCopy logical drive is created and a copy-on-write operation begins.
- startDate A specific date on which you want to create a FlashCopy logical drive and perform a copy-on-write operation. The format for entering the date is MM:DD:YY. If you do not provide a start date, the current date is used. An example of this option is startDate=06:27:11.
- scheduleDay A day of the week on which you want to create a FlashCopy logical drive and perform a copy-on-write operation. The values that you can enter are: monday, tuesday, wednesday, thursday, friday, saturday, and sunday. An example of this option is scheduleDay=wednesday.
- startTime The time of a day that you want to create a FlashCopy logical drive and start performing a copy-on-write operation. The format for entering the time is HH:MM, where HH is the hour and MM is the minute past the hour. Use a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is startTime=14:27.
- scheduleInterval An amount of time, in minutes, that you want to have as a
  minimum between copy-on-write operation. It is possible for you to create a
  schedule in which you have overlapping copy-on-write operations because of
  the duration a copy operation. You can make sure that you have time between
  copy-on-write operations by using this option. The maximum value for the
  scheduleInterval option is 1440 minutes. An example of this option is
  scheduleInterval=180.
- endDate A specific date on which you want to stop creating a FlashCopy logical drive and end the copy-on-write operations. The format for entering the date is MM:DD:YY. An example of this option is endDate=11:26:11.
- timesPerDay The number of times that you want the schedule to run in a day. An example of this option is timesPerDay=4.
- timeZone Use this parameter to define the time zone in which the storage subsystem is operating. You can define the time zone in one of two ways:
  - GMT±HH:MM The time zone offset from GMT. Enter the offset in hours and minutes. For example GMT-06:00 is the central time zone in the United States.

 Text string – Standard time zone text strings. For example: "USA/Chicago" or "Germany/Berlin". Time zone text strings s are case sensitive. If you enter an incorrect text string, GMT time is used. Enclose the text string in double quotation marks.

The code string for defining a schedule is similar to these examples:

enableSchedule=true schedule startTime=14:27
enableSchedule=true schedule scheduleInterval=180
enableSchedule=true schedule timeZone=GMT-06:00
enableSchedule=true schedule timeZone="USA/Chicago"

If you also use the **scheduleInterval** option, the firmware will choose between the **timesPerDay** option and the **scheduleInterval** option by selecting the lowest value of the two options. The firmware calculates an integer value for the **scheduleInterval** option by dividing 1440 by a the **scheduleInterval** option value that you set. For example, 1440/180 = 8. The firmware then compares the **timesPerDay** integer value with the calculated **scheduleInterval** integer value and uses the smaller value.

To remove a schedule, use the delete logicaldrive command with the **schedule** parameter. The delete logicaldrive command with the **schedule** parameter deletes only the schedule, not the FlashCopy logical drive.

#### **Related Topic**

Naming Conventions

# **Create Host**

This command creates a new host. If you do not specify a host group in which to create the new host, the new host is created in the Default Group.

#### Syntax

```
create host userLabel="hostName"
[hostGroup=("hostGroupName" | defaultGroup)]
[hostType=(hostTypeIndexLabel | hostTypeIndexNumber)]
```

Parameter	Description
userLabel	The name that you want to give the host that you are creating. Enclose the host name in double quotation marks ("").
hostGroup	The name of the host group in which you want to create a new host. Enclose the host group name in double quotation marks (""). (If a host group does not exist, you can create a new host group by using the <b>create hostGroup</b> command.) The defaultGroup option is the host group that contains the host to which the logical drive is mapped.
hostType	The index label or the index number that identifies the host type. Use the <b>show storageSubsystem hostTypeTable</b> command to generate a list of available host type identifiers. If the host type has special characters, enclose the host type in double quotation marks ("").

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

A host is a computer that is attached to the storage subsystem and accesses the logical drives on the storage subsystem through its HBA host ports. You can define specific logical drive-to-LUN mappings to an individual host. You also can assign the host to a host group that shares access to one or more logical drives.

A host group is an optional topological element that you can define if you want to designate a collection of hosts that share access to the same logical drives. The host group is a logical entity. Define a host group only if you have two or more hosts that share access to the same logical drives.

If you do not specify a host group in which to place the host that you are creating, the newly defined host belongs to the default host group.

#### **Related Topic**

Naming Conventions

# **Create Host Group**

This command creates a new host group.

#### **Syntax**

create hostGroup userLabel="hostGroupName"

#### Parameter

Parameter	Description
userLabel	The name that you want to give the host group that you are creating. Enclose the host name in double quotation marks ("").

#### **Notes**

A host group is an optional topological element that you can define if you want to designate a collection of hosts that share access to the same logical drives. The host group is a logical entity. Define a host group only if you have two or more hosts that can share access to the same logical drives.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

#### **Related Topic**

Naming Conventions

# **Create Host Port**

This command creates a new host port identification on a host bus adapter (HBA) or on a host channel adapter (HCA). The identification is a software value that represents the physical HBA or HCA host port to the controller. Without the correct host port identification, the controller cannot receive instructions or data from the host port.

# **Syntax**

```
create hostPort identifier=(wwID | gid)
userLabel=portLabel
host=hostName interfaceType=(FC | SAS | IB)
```

# Parameters

Parameter	Description
identifier	The 8-byte World Wide Identifier (WWID) or the 16-byte group identifier (GID) of the HBA or HCA host port. Enclose the unique ID in double quotation marks ("").
userLabel	The name that you want to give to the new HBA or HCA host port. Enclose the HBA host port label in double quotation marks ("").
host	The name of the host for which you are defining an HBA or HCA host port. Enclose the host name in double quotation marks ("").
interfaceType	<ul> <li>The identifier of the type of interface for the host port. The choices are:</li> <li>FC – Fibre Channel</li> <li>SAS – Serial-Attached SCSI</li> <li>IB – Infiniband</li> <li>An FC or a SAS selection requires an 8-byte WWID. An IB selection requires a 16-byte group identifier (gid).</li> <li>FC is the default interface for the host port.</li> </ul>

#### **Notes**

An HBA or HCA host port is a physical connection on a host adapter that resides within a host computer. An HBA or HCA host port provides access to the logical drives in a storage subsystem. If the host bus adapter has only one physical connection (one host port), the terms host port and host bus adapter are synonymous.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Host names can have a maximum of 30 characters.

# **Related Topic**

Naming Conventions

# **Create iSCSI Initiator**

This command creates a new iSCSI initiator object.

# Syntax

create iscsiInitiator iscsiName=iSCSI-ID
userLabel="name"
host=host-name
[chapSecret=securityKey]

#### **Parameters**

Parameters	Description
iscsiName	The default identifier of the iSCSI initiator. Enclose the name in double quotation marks (" ").
userLabel	The name that you want to use for the iSCSI initiator. Enclose the name in double quotation marks ("").
host	The name of the host in which the iSCSI initiator is installed. Enclose the name in double quotation marks ("").
chapSecret	The security key that you want to use to authenticate a peer connection. Enclose the name in double quotation marks ("").

#### Notes

Challenge Handshake Authentication Protocol (CHAP) is a protocol that authenticates the peer of a connection. CHAP is based upon the peers sharing a "secret." A secret is a security key that is similar to a password.

Use the **chapSecret** parameter to set up the security keys for initiators that require a mutual authentication.

#### **Related Topic**

Naming Conventions

# Create RAID Logical Drive (Automatic Drive Select)

This command creates an subsystem across the disk drives in the storage subsystem, and a new logical drive in the subsystem. The storage subsystem controllers choose the disk drives to be included in the logical drive.

# **Syntax**

**Important:** If you have disk drives with different capacities, you cannot automatically create logical drives by specifying the driveCount parameter. If you want to create logical drives with disk drives of different capacities, see "Create RAID Logical Drive (Manual Drive Select)" on page 126.

```
create logicalDrive driveCount=numberOfDrives
subsystemUserLabel=subsystemName
raidLevel=(0 | 1 | 3 | 5 | 6)
userLabel=logicalDriveName
DriveType=(fibre | SATA | SAS)
DriveMediaType=(HDD | SSD | unknown | allMedia)
capacity=logicalDriveCapacity
owner=(a | b)
cacheReadPrefetch=(TRUE | FALSE)
segmentSize=segmentSizeValue
usageHint=(fileSystem | dataBase | multiMedia)
enclosureLossProtect=(TRUE | FALSE)
enclosureLossProtect=(TRUE | FALSE)
```

drawerLossProtect=(TRUE | FALSE) dssPreAllocate=(TRUE | FALSE)] securityType=(none | capable | enabled)] T10PI=(none | enabled)]

Parameter	Description
driveCount	The number of unassigned disk drives that you want to use in the subsystem.
subsystemUserLabel	The alphanumeric identifier (including – and _) that you want to give the new subsystem. Enclose the new subsystem name in double quotation marks (" ").
raidLevel	The RAID level of the subsystem that contains the logical drive. Valid values are $0$ , 1, 3, 5, or 6.
userLabel	The name that you want to give to the new logical drive• . Enclose the new logical drive name in double quotation marks ("").
driveType	The type of disk drive that you want to use in the logical drive. You cannot mix disk drive types. Valid disk drive types are fibre, SATA, or SAS.
DriveMediaType	The type of disk drive medium that you want to use for the create RAID logical drive. Values are
	• HDD – Use this option when you have hard disk drives in the expansion drawer.
	• SSD – Use this option when you have solid state disk drives in the expansion drawer.
	<ul> <li>unknown – Use if you are not sure what types of disk drive media are in the expansion drawer.</li> </ul>
	<ul> <li>allMedia – Use this option when you want to use all types of disk drive media that are in the expansion drawer.</li> </ul>
	Use this parameter only when you use the DriveCount parameter. If you have only one type of disk drive media, you do not need to use this parameter. You must use this parameter when you have more than one type of disk drive media in your storage subsystem.
driveType	The type of drive that you want to use in the logical drive. You cannot mix drive types.
	You must use this parameter when you have more than one type of drive in your storage subsystem.
	Valid drive types are :
	• fibre
	• SATA
	• SAS
	If you do not specify a drive type, the command defaults to fibre.
capacity	The size of the logical drive that you are adding to the storage subsystem. Size is defined in units of bytes, KB, MB, GB, or TB.
owner	The controller that owns the logical drive. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. If you do not specify an owner, the controller firmware determines the owner.

Parameter	Description
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.
segmentSize	The amount of data (in KB) that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
usageHint	The setting for both <b>cacheReadPrefetch</b> parameter and the <b>segmentSize</b> parameter to be default values. The default values are based on the typical I/O usage pattern of the application that is using the logical drive. Valid values are fileSystem, dataBase, or multiMedia.
enclosureLossProtect	The setting to enforce enclosure loss protection when you create the subsystem. To enforce enclosure loss protection, set this parameter to TRUE. The default setting is FALSE.
drawerLossProtect	The setting to enforce drawer loss protection when you create the mirror repository subsystem. To enforce drawer loss protection, set this parameter to TRUE. The default value is FALSE.
dssPreAllocate	The setting to make sure that reserve capacity is allocated for future segment size increases. The default value is TRUE.
securityType	The setting to specify the security level when creating the subsystems and all associated logical drives. These settings are valid:
	• <b>none</b> —The subsystem and logical drives are not secure.
	• <b>capable</b> —The subsystems and logical drives are capable of having security set, but security has not been enabled.
	• <b>enabled</b> —The subsystem and logical drives have security enabled.
T10PI	The setting to specify that an subsystem, and the logical drives within the subsystem, has T10 PI (Protection Information) protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the subsystem. These settings are valid:
	<ul> <li>none – The subsystem does not have T10 PI (Protection Information) protection.</li> </ul>
	• enabled – The subsystem has T10 PI (Protection Information) protection. The subsystem supports protected information and is formatted with protection information enabled.

The **driveCount** parameter lets you choose the number of disk drives that you want to use in the subsystem. You do not need to specify the disk drives by enclosure ID and slot ID. The controllers choose the specific disk drives to use for the subsystem.

The **owner** parameter defines which controller owns the logical drive.

If you do not specify a capacity using the **capacity** parameter, all of the disk drive capacity that is available in the subsystem is used. If you do not specify capacity units, bytes is used as the default value.

# **Cache Read Prefetch**

Cache read prefetch lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the disk drives into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage subsystem that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the **cacheReadPrefetch** parameter are TRUE or FALSE.

#### Segment Size

The size of a segment determines how many data blocks that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single disk drive for a single request leaves other disk drives available to simultaneously service other requests.

If the logical drive is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of disk drives in the subsystem that are used for data transfers.) In this case, multiple disk drives are used for the same request, but each disk drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of disk drives that are required to satisfy a data transfer request.

You do not need to enter a value for the **cacheReadPrefetch** parameter or the **segmentSize** parameter. If you do not enter a value, the controller firmware uses the **usageHint** parameter with fileSystem as the default value. Entering a value for the **usageHint** parameter and a value for the **cacheReadPrefetch** parameter or a value for the **segmentSize** parameter does not cause an error. The value that you enter for the **cacheReadPrefetch** parameter takes priority over the value for the **usageHint** parameter.

# **Enclosure Loss Protection and Drawer Loss Protection**

For enclosure loss protection to work, each disk drive in an subsystem must be on a separate enclosure. If you set the **enclosureLossProtect** parameter to TRUE and have selected more than one disk drive from any one enclosure, the storage subsystem returns an error. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs operations, but the subsystem that you create might not have enclosure loss protection.

Enclosure loss protection is not valid when you create logical drives on existing subsystems.

The **drawerLossProtect** parameter defines if data on a logical drive is accessible if a drawer fails. When you assign the drives, if you set the **drawerLossProtect** parameter to TRUE and select more than one drive from any one drawer, the storage subsystem returns an error. If you set the **drawerLossProtect** parameter to FALSE, the storage subsystem performs operations, but the subsystem that you create might not have drawer loss protection.

You must set the **enclosureLossProtect** parameter and the **drawerLossProtect** parameter to the same value. Both of the parameters must be either TRUE or FALSE. If the **enclosureLossProtect** parameter and the **drawerLossProtect** parameter are set to different values, the storage subsystem returns an error.

# **Security Type**

See "Security Type" on page 423 for an explanation of this new parameter.

# **Related Topic**

Naming Conventions

# Create RAID Logical Drive (Free Extent Based Select)

This command creates a logical drive in the free space of an subsystem.

#### **Syntax**

```
create logicalDrive subsystem="subsystemNumber"
userLabel="logicalDriveName"
[freeCapacityArea=freeCapacityIndexNumber
capacity=logicalDriveCapacity
owner=(a | b)
cacheReadPrefetch=(TRUE | FALSE)
segmentSize=segmentSizeValue
usageHint=(fileSystem | dataBase | multiMedia)
[dssPreAllocate=(TRUE | FALSE)
securityType=(none | capable | enabled)
T10PI=(none | enabled)]
```

Parameter	Description
subsystem	The alphanumeric identifier (including - and _) for a specific subsystem in your storage subsystem.
userLabel	The name that you want to give the new logical drive. Enclose the new logical drive name in double quotation marks ("").

Parameter	Description
freeCapacityArea	The index number of the free space in an existing subsystem that you want to use to create the new logical drive. Free capacity is defined as the free capacity between existing logical drives in an subsystem. For example, an subsystem might have the following areas: logical drive 1, free capacity, logical drive 2, free capacity, logical drive 3, free capacity. To use the free capacity following logical drive 2, you would specify: <b>freeCapacityArea=2</b>
	Run the <b>show subsystem</b> command to determine if the free capacity area exists.
capacity	The size of the logical drive that you are adding to the storage subsystem. Size is defined in units of bytes, KB, MB, GB, or TB.
owner	The controller that owns the logical drive. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. If you do not specify an owner, the controller firmware determines the owner.
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.
segmentSize	The amount of data (in KB) that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
usageHint	The settings for both the <b>cacheReadPrefetch</b> parameter and the <b>segmentSize</b> parameter to be default values. The default values are based on the typical I/O usage pattern of the application that is using the logical drive. Valid values are fileSystem, dataBase, or multiMedia.
dssPreAllocate	The setting to make sure that reserve capacity is allocated for future segment size increases. The default value is TRUE.
securityType	The setting to specify the security level when creating the subsystem and all associated logical drives. These settings are valid:
	• <b>none</b> —The subsystem and logical drives are not secure.
	• <b>capable</b> —The subsystem and logical drives are capable of having security set, but security has not been enabled.
	• <b>enabled</b> —The subsystem and logical drives have security enabled.

Parameter	Description
T10PI	The setting to specify that an subsystem, and the logical drives within the subsystem, has T10 PI (Protection Information) protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the subsystem. These settings are valid:
	• <b>none</b> — The subsystem does not have T10 PI (Protection Information) protection.
	• <b>enabled</b> — The subsystem has T10 PI (Protection Information) protection. The subsystem supports protected information and is formatted with protection information enabled.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

The **owner** parameter defines which controller owns the logical drive. The preferred controller ownership of a logical drive is the controller that currently owns the subsystem.

If you do not specify a capacity using the **capacity** parameter, all of the available capacity in the free capacity area of the subsystem is used. If you do not specify capacity units, bytes is used as the default value.

# **Segment Size**

The size of a segment determines how many data blocks that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single disk drive for a single request leaves other disk drives available to simultaneously service other requests.

If the logical drive is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of disk drives in the subsystem that are used for data transfers.) In this case, multiple disk drives are used for the same request, but each disk drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of disk drives that are required to satisfy a data transfer request.

# **Cache Read Prefetch**

Cache read prefetch lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the disk drives into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage subsystem that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the **cacheReadPrefetch** parameter are TRUE or FALSE. You do not need to enter a value for the **cacheReadPrefetch** parameter or the **segmentSize** parameter. If you do not enter a value, the controller firmware uses the **usageHint** parameter with fileSystem as the default value.

Entering a value for the **usageHint** parameter and a value for the **cacheReadPrefetch** parameter or a value for the **segmentSize** parameter does not cause an error. The value that you enter for the **cacheReadPrefetch** parameter or the **segmentSize** parameter takes priority over the value for the **usageHint** parameter.

# Security Type

The **securityType** parameter is valid for drives that are capable of full disk encryption (FDE). With FDE, the controller firmware can create a key and activate the Full Disk Encryption feature. The Full Disk Encryption feature encrypts data as the data is written to the drive and decrypts the data as the data is read from the drive. Without the key created by the controller, the data written to the drive is inaccessible.

Before you can set the **securityType** parameter to capable or enabled, you must create a storage subsystem security key. Use the create storageSubsystem securityKey command to create a storage subsystem security key. These commands are related to the security key:

- create storageSubsystem securityKey
- set storageSubsystem securityKey
- import storageSubsystem securityKey
- export storageSubsystem securityKey
- start secureErase (drive | drives)
- enable subsystem [subsystemName] security
- create hostPort identifier

#### **Related Topic**

Naming Conventions

# Create RAID Logical Drive (Manual Drive Select)

This command creates a new subsystem and logical drive and lets you specify the disk drives for the logical drive.

**Important:** You cannot use mixed disk drive types in the same subsystem and logical drive. This command fails if you specify different types of disk drives for the RAID logical drive.
# Syntax

```
create logicalDrive drives=(enclosureID1,drawerID1,slotID1 ... enclosureIDn,
drawerIDn,slotIDn)
[subsystem=[subsystemName]]
raidLevel=(0 | 1 | 3 | 5 | 6)
userLabel="logicalDriveName"
owner=(a | b)
[capacity=logicalDriveCapacity
cacheReadPrefetch=(TRUE | FALSE)
segmentSize=segmentSizeValue
usageHint=(fileSystem | dataBase | multiMedia)
enclosureLossProtect=(TRUE | FALSE)
drawerLossProtect=(TRUE | FALSE)
dssPreAllocate=(TRUE | FALSE)
securityType=(none | capable | enabled)]
T10PI=(none | enabled)]
```

## **Parameters**

Parameter	Description
drives	The disk drives that you want to assign to the logicalDrive. EnclosureID1, drawerID1, slotID1 enclosureIDn, drawerIDn, slotIDn is for enclosures with multiple disk drive drawers. For enclosures without disk drive drawers, use enclosureID1, slotID1 enclosureIDn, slotIDn instead. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the values in parentheses.
subsystemUserLabel	The alphanumeric identifier (including – and _) that you want to give the new subsystem. Enclose the subsystem identifier in double quotation marks (" ").
raidLevel	The RAID level of the subsystem that contains the logical drive. Valid values are $0$ , 1, 3, 5, or 6.
userLabel	The name that you want to give the new logical drive. Enclose the new logical drive name in double quotation marks ("?).
capacity	The size of the logical drive that you are adding to the storage subsystem. Size is defined in units of bytes, KB, MB, GB, or TB.
owner	The controller that owns the logical drive. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. If you do not specify an owner, the controller firmware determines the owner.
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.
segmentSize	The amount of data (in KB) that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
usageHint	The settings for both the <b>cachReadPrefetch</b> parameter and the <b>segmentSize</b> parameter to be default values. The default values are based on the typical I/O usage pattern of the application that is using the logical drive. Valid values are fileSystem, dataBase, or multiMedia.

Parameter	Description
enclosureLossProtect	The setting to enforce enclosure loss protection when you create the repository. To enforce enclosure loss protection, set this parameter to TRUE. The default value is FALSE.
drawerLossProtect	The setting to enforce drawer loss protection when you create the mirrored repository logical drive. To enforce drawer loss protection, set this parameter to TRUE. The default value is FALSE.
dssPreAllocate	The setting to make sure that reserve capacity is allocated for future segment size increases. This default value is TRUE.
securityType	The setting to specify the security level when creating the subsystem and all associated logical drives. These settings are valid:
	• <b>none</b> —The subsystem and logical drives are not secure.
	• <b>capable</b> —The subsystem and logical drives are capable of having security set, but security has not been enabled.
	• <b>enabled</b> —The subsystem and logical drives have security enabled.
T10PI	The setting to specify that an subsystem, and the logical drives within the subsystem, has T10 PI (Protection Information) to help you ensure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the subsystem. These settings are valid:
	• <b>none</b> —The subsystem does not have T10 PI (Protection Information).
	• <b>enabled</b> —The subsystem has T10 PI (Protection Information). The subsystem support is protected information and is formatted with protection information enabled.

The **drives** parameter supports both high-capacity drive expansion enclosures and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive resides. For a low-capacity drive expansion enclosure, you need only specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

If you set the **raidLevel** parameter to RAID 1, the controller firmware takes the list of drives and pairs them by using this algorithm:

Data drive = XParity drive = N/2 + X

In this algorithm X is 1 to N/2, and N is the number of drives in the list. For example, if you have six drives, the mirror pairs are as follows:

Data	Parity
1	N/2 + 1 = 4
2	N/2 + 2 = 5
3	N/2 + 3 = 6

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

The **owner** parameter defines which controller owns the logical drive. The preferred controller ownership of a logical drive is the controller that currently owns the subsystem.

If you do not specify a capacity using the **capacity** parameter, all of the disk drive capacity that is available in the subsystem is used. If you do not specify capacity units, bytes is used as the default value.

## Segment Size

The size of a segment determines how many data blocks that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single disk drive for a single request leaves other disk drives available to simultaneously service other requests.

If the logical drive is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of disk drives in the subsystem that are used for data transfers.) In this case, multiple disk drives are used for the same request, but each disk drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of disk drives that are required to satisfy a data transfer request.

## Cache Read Prefetch

Cache read prefetch lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the disk drive into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage subsystem that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the **cacheReadPrefetch** parameter are TRUE or FALSE.

You do not need to enter a value for the **cacheReadPrefetch** parameter or the **segmentSize** parameter. If you do not enter a value, the controller firmware uses the **usageHint** parameter with fileSystem as the default value. Entering a value for the **usageHint** parameter and a value for the **cacheReadPrefetch** parameter or a value for the **segmentSize** parameter does not cause an error. The value that you enter for the **cacheReadPrefetch** parameter takes priority over the value for the **usageHint** parameter.

For enclosure loss protection to work, each disk drive in an subsystem must be on a separate enclosure. If you set the **enclosureLossProtect** parameter to TRUE and have selected more than one disk drive from any one enclosure, the storage subsystem returns an error. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs operations, but the subsystem that you create might not have enclosure loss protection. enclosure loss protection is not valid when you create logical drives on existing subsystems.

### **Enclosure Loss Protection and Drawer Loss Protection**

For enclosure loss protection to work, each drive in an subsystem must be on a separate enclosure. If you set the **enclosureLossProtect** parameter to TRUE and have selected more than one drive from any one enclosure, the storage subsystem returns an error. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs operations, but the subsystem that you create might not have enclosure loss protection.

Enclosure loss protection is not valid when you create logical drives on existing subsystems.

The **drawerLossProtect** parameter defines if data on a logical drive is accessible if a drawer fails. When you assign the drives, if you set the **drawerLossProtect** parameter to TRUE and select more than one drive from any one drawer, the storage subsystem returns an error. If you set the **drawerLossProtect** parameter to FALSE, the storage subsystem performs operations, but the subsystem that you create might not have drawer loss protection

You must set the **enclosureLossProtect** parameter and the **drawerLossProtect** parameter to the same value. Both of the parameters must be either TRUE or FALSE. If the **enclosureLossProtect** parameter and the **drawerLossProtect** parameter are set to different values, the storage subsystem returns an error.

### Security Type

The **securityType** parameter is valid for drives that are capable of full disk encryption (FDE). With FDE, the controller firmware can create a key and activate the Full Disk Encryption feature. The Full Disk Encryption feature encrypts data as the data is written to the drive and decrypts the data as the data is read from the drive. Without the key created by the controller, the data written to the drive is inaccessible Before you can set the securityType parameter to capable or enabled, you must create a storage subsystem security key. Use the create storageSubsystem securityKey command to create a storage subsystem security key. These commands are related to the security key:

- create storageSubsystem securityKey
- enable subsystem [subsystemName] security
- export storageSubsystem securityKey
- import storageSubsystem securityKey
- set storageSubsystem securityKey
- start secureErase (drive | drives)

## **Related Topic**

Naming Conventions

## Create Logical Drive on Disk Pool

This command creates a new standard RAID logical drive or a thin logical drive on an existing disk pool. Note that some of the traditional logical drive creation parameters are left off because they have no applicability to disk pools. Some of the parameters included in **create-raid-vol-attr-value-list** must be ignored, or should cause an error if specified by the user.

**Note:** Some parameters for creating a standard RAID logical drive on an subsystem are not compatible for creating logical drives of any type on a disk pool. When using older scripts to create logical drives on disk pools, make sure that all of the parameters are valid for disk pools. Invalid parameters prevent the scripts from running correctly and cause an error to be posted.

## Syntax for Creating a Standard Logical Drive

create logicaldrive diskPool="diskPoolName"
userLabel="logicalDriveName"
capacity=logicaldriveCapacity
[thinProvisioned=(TRUE | FALSE) |
owner=(a|b) |
mapping=(none|default) |
T10PI=(none|enabled) |
cacheReadPrefetch=(TRUE | FALSE)]

## Syntax for Creating a Thin Provision Logical Drive

```
create logicaldrive diskPool="diskPoolName"
userLabel="logicalDriveName"
capacity=logicaldriveCapacity
[thinProvisioned=(TRUE | FALSE) |
owner=(a|b) |
mapping=(none|default) |
T10PI=(none|enabled) |
(existingRepositoryLabel=existingRepositoryName |
newRepositoryCapacity=newRepositoryCapacityValue [KB | MB | GB | TB | Bytes]) |
repositoryMaxCapacity=repositoryMaxCapacityValue[KB|MB|GB|TB|Bytes] |
warningThresholdPercent=warningThresholdPercentValue |
repositoryExpansionPolicy=(automatic|manual) |
cacheReadPrefetch=(TRUE | FALSE)]
```

# Parameter

Parameter	Description
diskPool	The name of the disk pool on which to create the new logical drive. Enclose the disk pool name in double quotation marks (" ").
userLabel	The name that you want to give the new logical drive. Enclose the logical drive name in double quotation marks (" ").
capacity	The size of the logical drive that you are creating.
	For a standard RAID logical drive, the capacity that will be allocated for the logical drive.
	For a thin logical drive, the virtual capacity value that will be exposed by the thin logical drive.
	Size is defined in units of bytes, KB, MB, GB, or TB. The following are examples of the syntax:
	<ul> <li>capacity=500MB</li> <li>capacity=2GB</li> </ul>
thinProvisioned	This parameter enables thin provisioning for the new logical drive. To use thin provisioning, set this parameter to TRUE. If you do not want thin provisioning, set this parameter to FALSE.
owner	The controller that owns the logical drive. Valid controller identifiers are <b>a</b> or <b>b</b> , where <b>a</b> is the controller in slot A, and <b>b</b> is the controller in slot B. If you do not specify an owner, the controller firmware determines the owner.
mapping	This parameter enables you to map the logical drive to a host. If you want to map later, set this parameter to none. If you want to map now, set this parameter to default. The logical drive is mapped to all hosts that have access to the storage pool.
	The default value is none.

Parameter	Description
T10PI	<ul> <li>The setting to specify that an subsystem, and the logical drives within the subsystem, has T10 PI (Protection Information) protection to make sure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the subsystem. These settings are valid:</li> <li>none - The subsystem does not have T10 PI (Protection Information) protection.</li> <li>enabled - The subsystem has T10 PI (Protection Information) protection. The subsystem supports protected information and is formatted with protection information enabled.</li> </ul>
	The default value is <i>enabled</i> if the disk pool is T10 PI (Protection Information) capable.
existingRepositoryLabel	This parameter applies only for thin provisioning. If you use the <b>existingRepositoryLabel</b> parameter, you must not use the <b>newRepositoryCapacity</b> parameter.
newRepositoryCapacity	The size of the logical drive that you are creating. Use this parameter only if you see the value of the <b>thinProvisioned</b> parameter to TRUE.
	<pre>Size is defined in units of MB, GB, or TB. The following are examples of the syntax: capacity=500MB capacity=2GB</pre>
	The default value is 50% of the virtual capacity.
warningThreshold	Setting this value to <i>zero</i> disables warnings. The default value is determined by the
	controller firmware.
repositoryExpansionPolicy	<ul><li> automatic</li><li> manual</li></ul>
cacheReadPrefetch	• TRUE • FALSE

Each logical drive name must be unique. You can use any combination of alphanumeric characters, underscore (\_), hyphen(-), and pound (#) for the user label. User labels can have a maximum of 30 characters.

For thin logical drives, the *capacity* value specifies the virtual capacity of the logical drive, and the *repositoryCapacity* specifies the capacity of the logical drive created as the repository logical drive. Use the *existingRepositoryLabel* to specific an existing unused repository logical drive instead of creating a new logical drive.

For best results when creating a thin logical drive, the repository logical drive must already exist or must be created in an already existing disk pool. If you do not specify some of the optional parameters when creating thin logical drives the storage management software will attempt to create the repository logical drive. The most desirable candidate logical drive is a repository logical drive that already exists and that is within the size requirements. The next most desirable candidate logical drive is a new repository logical drive that is created on the disk pool free extent.

Repository logical drives for thin logical drives cannot be created on subsystems.

# Create Read-Only Enhanced FlashCopy Logical Drive

This command creates a read-only Enhanced FlashCopy logical drive for the Enhanced FlashCopy images of a base logical drive. To change a read-only Enhanced FlashCopy logical drive to a read/write logical drive, use the set enhancedFlashCopyLogicalDrive convertToReadWrite command.

**Note:** You cannot use this command for an Enhanced FlashCopy image that is used in online logical drive copy.

### **Syntax**

create enhancedFlashCopyLogicalDrive userLabel="enhancedFlashCopyLogicalDriveName"
enhancedFlashCopyImageID="enhancedFlashCopyImageID"
readOnly

### **Parameter**

Parameter	Description
userLabel	The name that you want to give to an Enhanced FlashCopy logical drive. Enclose
	in double quotation marks (" ").

Parameter	Description
enhancedFlashCopyImageID	<ul> <li>The alphanumeric identifier of an Enhanced FlashCopy image that you want to add to the new Enhanced FlashCopy logical drive.</li> <li>The identifier of an Enhanced FlashCopy image is comprised of two parts:</li> <li>The name of the Enhanced FlashCopy group</li> <li>An identifier for the Enhanced FlashCopy</li> </ul>
	image in the Enhanced FlashCopy group.
	The identifier for the Enhanced FlashCopy image can be one of these:
	• An integer value that is the sequence number of the Enhanced FlashCopy in the Enhanced FlashCopy group.
	• <i>NEWEST</i> - Use this option when you want to show the latest Enhanced FlashCopy image created in the Enhanced FlashCopy group.
	• <i>OLDEST</i> - Use this option when you want to show the earliest Enhanced FlashCopy image created in the Enhanced FlashCopy group.
	Enclose the Enhanced FlashCopy image name in double quotation marks (" ").
readOnly	This parameter sets the Enhanced FlashCopy logical drive to read-only. This parameter is actually a Boolean; however, in the context of this command, the Boolean value is always TRUE. <b>Note:</b> With the <b>read0n1v</b> parameter.
	Enhanced FlashCopy logical drives are not created.

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

The identifier of an Enhanced FlashCopy image has two parts separated by a colon (:):

- The name of the Enhanced FlashCopy group
- The identifier of the Enhanced FlashCopy image

For example, if you want to create a read-only logical drive named **engData1** using the most recent Enhanced FlashCopy image in an Enhanced FlashCopy group that has the name *enhancedFlashCopyGroup1*, you would use this command:

create enhancedFlashCopyLogicalDrive userLabel="engData1"
enhancedFlashCopyImageID="enhancedFlashCopyGroup1:newest" readOnly;

# **Create Remote Mirror**

This command creates both the primary logical drive and the secondary logical drive for a remote-mirror pair. This command also sets the write mode (Synchronous or Asynchronous) and the synchronization priority.

## **Syntax**

```
create remoteMirror primary="primaryLogicalDriveName"
secondary="secondaryLogicalDriveName"
(remoteStorageSubsystemName="storageSubsystemName" |
remoteStorageSubsystemWwn="wwID")
[remotePassword="password"
syncPriority=(highest | high | medium | low | lowest)
autoResync=(enabled | disabled)
writeOrder=(preserved | notPreserved)
writeMode=(synchronous | asynchronous)]
```

Parameter	Description
primary	The name of an existing logical drive on the local storage subsystem that you want to use for the primary logical drive. Enclose the primary logical drive name in double quotation marks ("").
secondary	The name of an existing logical drive on the remote storage subsystem that you want to use for the secondary logical drive. Enclose the secondary logical drive name in double quotation marks ("").
remoteStorageSubsystemName	The name of the remote storage subsystem. Enclose the remote storage subsystem name in double quotation marks ("").
remoteStorageSubsystemWwn	The WWID of the remote storage subsystem. Enclose the WWID in double quotation marks ("").
remotePassword	The password for the remote storage subsystem. Use this parameter when the remote storage subsystem is password protected. Enclose the password in double quotation marks ("").
syncPriority	The priority that full synchronization has relative to host I/O activity. Valid values are highest, high, medium, low, or lowest.
autoResync	The settings for automatic resynchronization between the primary logical drives and the secondary logical drives of a remote-mirror pair. This parameter has the following values:
	<ul> <li>enabled – Automatic resynchronization is turned on. You do not need to do anything further to resynchronize the primary logical drive and the secondary logical drive.</li> </ul>
	<ul> <li>disabled – Automatic resynchronization is turned off. To resynchronize the primary logical drives and the secondary logical drive, you must run the resume remoteMirror command.</li> </ul>
writeOrder	The write order for data transmission between the primary logical drive and the secondary logical drive. Valid values are preserved or not Preserved

### **Parameters**

Parameter	Description
writeMode	How the primary logical drive writes to the secondary logical drive. Valid values are synchronous or asynchronous.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

When you choose the primary logical drive and the secondary logical drive, the secondary logical drive must be of equal or greater size than the primary logical drive. The RAID level of the secondary logical drive does not have to be the same as the primary logical drive.

Product shipments using the DS4800 (Models 82, 84, 88) or DS5100, DS5300 controller define a maximum of 128 remote mirrors. The DS4700 Model 70 and DS4700 Model 72 controllers can define a maximum of 64 remote mirrors.

Passwords are stored on each storage subsystem in a management domain. If a password was not previously set, you do not need a password. The password can be any combination of a alphanumeric characters with a maximum of 30 characters. (You can define a storage subsystem password by using the **set storageSubsystem** command).

Synchronization priority defines the amount of system resources that are used to synchronize the data between the primary logical drive and the secondary logical drive of a mirror relationship. If you select the highest priority level, the data synchronization uses the most system resources to perform the full synchronization, which decreases performance for host data transfers.

The **writeOrder** parameter applies only to Enhanced Global Mirrors and makes them become part of a consistency group. Setting the **writeOrder** parameter to preserved causes the remote-mirror pair to transmit data from the primary logical drive to the secondary logical drive in the same order as the host writes to the primary logical drive. In the event of a transmission link failure, the data is buffered until a full synchronization can occur. This action can require additional system overhead to maintain the buffered data, which slows operations. Setting the **writeOrder** parameter to notPreserved frees the system from having to maintain data in a buffer, but it requires forcing a full synchronization to make sure that the secondary logical drive has the same data as the primary logical drive.

## **Create SSD cache**

This command creates a read cache for a storage subsystem using Solid State Disks (SSDs). Using high performance SSDs to cache read data improves the application I/O performance and response times, and delivers sustained performance improvement across different workloads, especially for high-IOP workloads. Performance Read Cache maximizes the use of expensive fast SSDs.

Performance Read Cache works in addition to the primary cache in the controller DRAM. With controller cache, the data is stored in DRAM after a host read. With Performance Read Cache, the data is copied from user-specified base logical drives, and then cached on SSDs.

# **Syntax**

create performanceReadCache userLabel="performanceReadCacheName"
drives=(enclosureID1,drawerID1,slotID1 ...
enclosureIDn,drawerIDn,slotIDn)
[updateExistinglogicalDrives=(TRUE|FALSE)]

## **Parameters**

Parameter	Description
userLabel	The alphanumeric identifier (including - and _) that you want to name the new Performance Read Cache. Enclose the identifier in double quotation marks (" "). You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the identifier. Identifiers can have a maximum of 30 characters.
drives	<ul> <li>The drives that you want to use to create the Performance Read Cache. For high-capacity drive enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for each drive that you want to assign to the Performance Read Cache. For low-capacity drive enclosures, specify the enclosure ID value and the slot ID value for each drive that you assign to the Performance Read Cache.</li> <li>Enclosure ID values are 0 to 99</li> <li>Drawer ID values are 1 to 5</li> <li>Slot ID values are 1 to 32</li> <li>Enclose the enclosure ID values, the drawer ID values, and the slot ID values in parentheses.</li> </ul>
updateExistinglogicalDrives	This optional parameter specifies whether the Performance Read Cache should be enabled for all existing logical drives in the storage subsystem. To enable the Performance Read Cache for all existing logical drives, set this parameter to TRUE. If you want to enable the Performance Read Cache for individual logical drives later, set this parameter toFALSE. The default value is TRUE.

## Notes

A storage subsystem can have only one Performance Read Cache.

Only logical drives created using hard disk drives can use the Performance Read Cache. You cannot enable Performance Read Cache on FlashCopy logical drives or Enhanced FlashCopy images.

SSDs do not currently support full disk encryption (FDE). Therefore, logical drives that have FDE enabled cannot use the Performance Read Cache.

If all of the SSDs in the Performance Read Cache are T10PI-capable and the T10PI premium feature is enabled, T10PI is automatically enabled for the Performance Read Cache and cannot be disabled. In addition, you cannot add non-T10PI capable SSDs to a T10PI-enabled Performance Read Cache.

## **Minimum Firmware Level**

7.84

# Create Storage Subsystem Security Key

## Syntax

This command creates a new security key for a storage subsystem that has Full Disk Encryption drives. This command also sets the security definitions and sets the state to Security Enabled.

```
create storageSubsystem securityKey
[keyIdentifier=keyIdentifierString] |
passPhrase=passPhraseString |
file=fileName |
commitSecurityKey=(TRUE | FALSE)
```

### Parameters

Parameter	Description
keyIdentifier	The character string that combines the storage subsystem ID and a randomly generated string to produce the security key identifier.
passPhrase	The character string that wraps the security key identifier with a pass phrase that is 8 to 32 characters in length. You must use at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character (a space is not permitted).
file	The name of the file to which you save the security key identifier. Enclose the file name in double quotes (" "). <b>Note:</b> Add a file extension of .slk to the end of the file name.
commitSecurityKey	This parameter commits the security key identifier to the storage subsystem for all FDE drives as well as the controllers. After the security key identifier is committed, a key is required to read data or write data. The data can only be read or changed by using a key, and the drive can never be used in a non-secure mode without rendering the data useless or totally erasing the drive.

## Notes

Use this command for local key management only.

The controller firmware creates a lock that restricts access to the FDE drives. FDE drives have a state called Security Capable. When you create a security key, the state is set to Security Enabled, which restricts access to all FDE drives that exist within the storage subsystem.

You can have a storage subsystem configuration with more than one set of encrypted subsystems or disk pools. Each subsystem or disk pool can have a unique security key. The character string generated by the **keyIdentifier** parameter is a string that you can read and that enables you to identify the security key that you need. You can create a **keyIdentifer** by using one of these methods:

You can enter up to 189 alphanumeric characters for a key identifier. The key identifier cannot have these characters:

- White spaces
- Punctuation
- Symbols

If you do not enter the **keyIdentifer** parameter, the controller automatically generates the **keyIdentifer** parameter.

Additional characters are automatically generated and appended to the end of the string that you enter for the key identifier. If you do not enter any string for the **keyIdentifier** parameter, the key identifier consists of only the characters that are automatically generated.

Your pass phrase must meet these criteria:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, <> @ +.

**Note:** If your pass phrase does not meet these criteria, you will receive an error message and will be asked to retry the command.

## Create Enhanced Remote Mirroring

This command creates both the primary logical drive and the secondary logical drive for a Synchronous Remote Mirrored pair. This command also sets the write mode (Metro Mirror or Global Copy) and the synchronization priority.

### Syntax

**Note:** In previous versions of this command the feature identifier was remoteMirror. This feature identifier is no longer valid and is replaced by syncMirror.

```
create syncMirror primary="primarylogicalDriveName" fn
secondary="secondarylogicalDriveName" fn
(remoteStoragesubsystemName="storagesubsystemName" | fn
remoteStoragesubsystemWwn="wwID") fn
[remotePassword="password" fn
syncPriority=(highest | high | medium | low | lowest) fn
autoResync=(enabled | disabled) fn
writeOrder=(preserved | notPreserved) fn
writeMode=(synchronous | asynchronous)]
```

## Parameters

None

## Create VolumeCopy

This command creates a VolumeCopy and starts the VolumeCopy operation.

**Attention:** Starting a VolumeCopy operation overwrites all existing data on the target logical drive, makes the target logical drive read-only to hosts, and fails all FlashCopy logical drives associated with the target logical drive, if any exist. If you have used the target logical drive as a copy before, be sure you no longer need the data or have it backed up.

This command creates VolumeCopies in two ways:

- VolumeCopy without FlashCopy also called offline VolumeCopy
- VolumeCopy with FlashCopy, also called online VolumeCopy

If you use VolumeCopy without FlashCopy you cannot write to the source logical drive until the copy operation is complete. If you want to be able to write to the source logical drive before the copy operation is complete, use VolumeCopy with FlashCopy . You can select VolumeCopy with FlashCopy through the optional parameters in the command syntax.

After completion of the VolumeCopy with FlashCopy operation, the FlashCopy is disabled.

**Note:** You can have a maximum of eight VolumeCopies in progress at one time. If you try to create more than eight VolumeCopies at one time, the controllers return a status of **Pending** until one of the logical driveCopies that is in progress finishes and returns a status of Complete

## **Syntax**

```
create VolumeCopy source="sourceName"
target="targetName"
[copyPriority=(highest | high | medium | low | lowest)
targetReadOnlyEnabled=(TRUE | FALSE)
copyType=(offline | online)
repositoryPercentOfBase=(20 | 40 | 60 | 120 | default) |
repositoryGroupPreference=(sameAsSource | otherThanSource | default)]
```

## **Parameters**

Parameter	Description
source	The name of an existing logical drive that you want to use as the source logical drive. Enclose the source logical drive name in double quotation marks ("").
target	The name of an existing logical drive that you want to use as the target logical drive. Enclose the target logical drive name in double quotation marks ("").
copyPriority	The priority that VolumeCopy has relative to host I/O activity. Valid values are highest, high, medium, low, or lowest.

Parameter	Description
targetRead0n1yEnab1ed	The setting so that you can write to the target logical drive or only read from the target logical drive. To write to the target logical drive, set this parameter to FALSE. To prevent writing to the target logical drive, set this parameter to TRUE.
соруТуре	Use this parameter to create a VolumeCopy with a FlashCopy . Creating a VolumeCopy with a FlashCopy enables you to continue to write to the source logical drive while creating the VolumeCopy. To create a VolumeCopy with a FlashCopy, set this parameter to online. To create a VolumeCopy without a FlashCopy, set this parameter to offline.
	If you do not use this parameter, the VolumeCopy is created without a FlashCopy .
repositoryPercentOfBase	This parameter determines the size of the repository logical drive for the FlashCopy when you are creating a VolumeCopy with a FlashCopy . The size of the repository logical drive is expressed as a percentage of the source logical drive, which is also called the base logical drive. Valid values for this parameter are 20, 40, 60, 120, and default. The default value is 20. If you do not use this parameter, the firmware uses a value of 20 percent.
repositoryGroupPreference	This parameter determines to which subsystem the FlashCopy repository logical drive is written. You have these choices:
	• sameAsSource – The FlashCopy repository logical drive is written to the same subsystem as the source logical drive if space is available.
	• otherThanSource – The FlashCopy repository logical drive is written to a different subsystem. Firmware determines which subsystem based on available space on the subsystems.
	<ul> <li>default – The FlashCopy repository logical drive is written to any subsystem that has space.</li> </ul>
	For best performance, use the sameAsSource option.
	You must use the <b>copyType</b> parameter with the <b>repositoryGroupPreference</b> parameter.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

Copy priority defines the amount of system resources that are used to copy the data between the source logical drive and the target logical drive of a VolumeCopy pair. If you select the highest priority level, the VolumeCopy uses the most system resources to perform VolumeCopy, which decreases performance for host data transfers.

# **Deactivate Remote Mirror**

This command deactivates the Enhanced Remote Mirroring feature, disassembles the mirror repository logical drive, and releases the controller-owner of the secondary logical drive. The controller host port that is dedicated to the secondary logical drive is available for host data transfers.

### **Syntax**

deactivate storageSubsystem feature=remoteMirror

### **Parameters**

None.

# **Deactivate Enhanced Global Mirroring**

This command deactivates the Enhanced Global Mirroring premium feature.

### **Syntax**

**Note:** All existing Enhanced Global Mirror Groups or Enhanced Global Mirrored pairs must be deleted from the local storage subsystem and the remote storage subsystem before the Enhanced Global Mirroring feature can be deactivated. deactivate storageArray feature=enhancedGlobalMirror

### **Parameters**

None

## **Deactivate Enhanced Remote Mirroring**

This command deactivates the Enhanced Remote Mirroring premium feature, disassembles the mirror repository logical drive, and releases the controller owner of the secondary logical drive. The controller host port that is dedicated to the secondary logical drive is available for host data transfers.

### Syntax

deactivate storagesubsystem feature=syncMirror

**Note:** In previous versions of this command the feature identifier was remoteMirror. This feature identifier is no longer valid and is replaced by syncMirror.

## **Parameters**

None

## Delete subsystem

This command deletes an entire subsystem and its associated logical drives.

Attention: Possible damage to the storage subsystem configuration – All of the data in the subsystem is lost as soon as you run this command.

## **Syntax**

delete subsystem [subsystemName]

## Parameter

Parameter	Description
subsystem	The alphanumeric identifier (including - and _) of the subsystem that you want to delete. Enclose the subsystem identifier in square brackets ([]).

## **Delete Enhanced Global Mirror Group**

This command deletes one or more Enhanced Global Mirror Groups from the local storage subsystem and the remote storage subsystem.

### **Syntax**

**Note:** The Enhanced Global Mirror Group must be empty before it can be successfully deleted. You must remove all Enhanced Global Mirrored pairs from the Enhanced Global Mirror Group before using this command.

```
delete enhancedGlobalMirrorGroup △
(allenhancedGlobalMirrorGroups |
enhancedGlobalMirrorGroup["enhancedGlobalMirrorGroupName"] | △
enhancedGlobalMirrorGroups
["enhancedGlobalMirrorGroupName_01""enhancedGlobalMirrorGroupName 02"])
```

## **Parameters**

None

## **Delete Consistency Group**

This command deletes an Enhanced FlashCopy consistency group. This command works in two ways:

- You can delete both the consistency group and the repository logical drives contained by the consistency group
- You can delete only the consistency group and leave the repository logical drives that are contained by the consistency group intact

## **Syntax**

delete consistencyGroup [enhancedFlashCopyConsistencyGroupName
[deleteRepositoryMembers=(TRUE | FALSE)]]

### **Parameter**

Parameter	Description
enhancedF1ashCopyConsistencyGroup	The name of the Enhanced FlashCopy consistency group that you want to delete. Enclose the Enhanced FlashCopy consistency group name in square brackets ([]). If the Enhanced FlashCopy consistency group name has special characters, you also must enclose the Enhanced FlashCopy consistency group name in double quotation marks (" ").

Parameter	Description
deleteRepositoryMembers	The setting to delete on or retain the repository logical drives. To delete the repository logical drives, set this parameter to TRUE. To retain the repository logical drives, set this parameter to FALSE. The default setting is FALSE.

# Delete Consistency Group Enhanced FlashCopy Logical Drive

This command deletes the Enhanced FlashCopy logical drive of a consistency group. Optionally, you can also delete the repository members.

## **Syntax**

delete cgEnhancedFlashCopyLogicalDrive ["enhancedFlashCopyLogicalDriveName"]
[deleteRepositoryMembers=(TRUE | FALSE)]

## **Parameter**

Parameter	Description
cgEnhancedFlashCopyLogicalDrive	The name of the consistency group Enhanced FlashCopy logical drive that you want to delete. Enclose the name of the consistency group Enhanced FlashCopy logical drive in double quotation marks (" ") inside square brackets ([ ]).
deleteRepositoryMembers	The parameter to save or delete the member logical drives. To save the member logical drives, set this parameter to TRUE. To delete the member logical drives, set this parameter to FALSE. The default value is TRUE.

# Delete Consistency Group Enhanced FlashCopy Image

This command deletes the Enhanced FlashCopy images in a consistency group.

## **Syntax**

delete cgEnhancedFlashCopyImage consistencyGroup="consistencyGroupName"
[(deleteCount=numberOfEnhancedFlashCopyImages |
retainCount=numberOfEnhancedFlashCopyImages) |
ignoreEnhancedFlashCopyLogicalDrive=(TRUE | FALSE)

## **Parameters**

Parameter	Description
consistencyGroup	The name of the consistency group from which you want to delete the Enhanced FlashCopy images. Enclose the consistency group name in double quotation marks (" ").

Parameter	Description
deleteCount	The number of Enhanced FlashCopy images that you want to delete from the consistency group. Use integer values.
	This parameter deletes the oldest Enhanced FlashCopy image first and continues to delete the oldest Enhanced FlashCopy images until reaching the number that you enter.
retainCount	The number of Enhanced FlashCopy images that you want to keep in the consistency group. Use integer values.
	This parameter keeps the most recent Enhanced FlashCopy images in the consistency group.
ignoreEnhancedFlashCopyLogicalDrive	Use this parameter to make sure that you do not delete an Enhanced FlashCopy image that has an Enhanced FlashCopy logical drive associated with the Enhanced FlashCopy image. You can use one of these values:
	• <b>TRUE</b> - Use this value to delete the Enhanced FlashCopy image even though the Enhanced FlashCopy image has an associated Enhanced FlashCopy logical drive.
	• <b>FALSE</b> - This value will cause the command to terminate when the Enhanced FlashCopy image has an associated Enhanced FlashCopy logical drive.
	The default value is <b>TRUE</b> .

If the Enhanced FlashCopy images cannot be deleted for all of the relevant consistency groups, the operation fails and none of the Enhanced FlashCopy images are deleted.

When you delete a consistency group Enhanced FlashCopy image that is associated with a consistency group Enhanced FlashCopy logical drive, the corresponding Enhanced FlashCopy logical drive member in the consistency group Enhanced FlashCopy logical drive are transitioned to the Stopped state. An Enhanced FlashCopy logical drive member in the Stopped state no longer has a relationship to the Enhanced FlashCopy group of the deleted Enhanced FlashCopy image. However, an Enhanced FlashCopy logical drive member in the Stopped state keeps its relationship to its consistency group Enhanced FlashCopy logical drive.

## **Delete Disk Pool**

This command deletes a disk pool.

## Syntax

delete diskPool [diskPoolName] [force=(true | false))]

## Parameter

Parameter	Description
diskPoolName	The name of the disk pool that you want to delete. Enclose the disk pool name in square brackets ([]). If the disk pool name has special characters, you also must enclose the host group name in double quotation marks ("").
force	Under normal operation, deleting a disk pool will automatically delete the pool whether there are logical drives present or not. The force parameter has special meaning in the following cases: If the GlobalFeatureState <b>DELETE_LOGICALDRIVE_RESTRICTION</b> feature is present, and the Preference Store key <b>DELETE_LOGICALDRIVE_RESTRICTIONS_KEY</b> are both true, then restrictions for deleting disk pools with logical drives are in effect. If both conditions are true, and the user has specified "force=true", then the disk pool will be deleted even if it has logical drives. If both flags are true, the disk pool has logical drives and the force flag has not been specified, the operation will fail.

## Notes

Each disk pool name must be unique. You can use any combination of alphanumeric characters, underscore (\_), hyphen(-), and pound (#) for the user label. User labels can have a maximum of 30 characters.

If the disk pool contains logical drives, the logical drives will be deleted along with the disk pool.

# Delete Enhanced FlashCopy Group

This command deletes an entire Enhanced FlashCopy group and optionally the associated member logical drives.

**Attention:** Possible damage to the storage subsystem configuration All of the data in the Enhanced FlashCopy group is lost as soon as you run this command.

This command deletes an entire Enhanced FlashCopy group and optionally the associated member logical drives.

### Syntax

delete enhancedFlashCopyGroup ["enhancedFlashCopyGroupName"]
[deleteRepositoryMembers=(TRUE | FALSE)

## Parameters

Parameter	Description
enhancedF1ashCopyGroup	The alphanumeric identifier (including -and _) of the Enhanced FlashCopy group that you want to delete. Enclose the Enhanced FlashCopy group identifier in double quotation marks (" ") inside square brackets ([]).
deleteRepositoryMembers	The parameter to save or delete the member logical drives. To save the member logical drives, set this parameter to FALSE. To delete the member logical drives, set this parameter to TRUE. The default value is FALSE.

### Notes

You can delete an Enhanced FlashCopy group if it is empty or if it contains Enhanced FlashCopy images. All of the Enhanced FlashCopy images in the Enhanced FlashCopy group are deleted along with the Enhanced FlashCopy group. If any existing Enhanced FlashCopy image within the Enhanced FlashCopy group has an associated Enhanced FlashCopy logical drive, each Enhanced FlashCopy logical drive is stopped and detached from the Enhanced FlashCopy image. When you delete an Enhanced FlashCopy group the associated repository logical drive is also deleted. By default, all member logical drives in the repository logical drive are retained as unused, unmapped standard logical drives. To remove the member logical drives set the **deleteRepositoryMembers** parameter to TRUE, or do not use this parameter. To keep the member logical drives, set the **deleteRepositoryMembers** parameter to FALSE.

# Delete Enhanced FlashCopy Image

This command deletes one or more Enhanced FlashCopy images from an Enhanced FlashCopy group.

### Syntax

```
delete enhancedFlashCopyImage (enhancedFlashCopyGroup="enhancedFlashCopyGroupName" |
enhancedFlashCopyGroups=("enhancedFlashCopyGroupName1"... "enhancedFlashCopyGroupNamen"))
(deleteCount=numberOfEnhancedFlashCopyImages |
retainCount=numberOfEnhancedFlashCopyImages |
ignoreEnhancedFlashCopyLogicalDrive=(TRUE | FALSE) |
enhancedFlashCopyImageID=0LDEST)
```

# Parameters

Parameter	Description
enhancedFlashCopyGroup or enhancedFlashCopyGroups	The name of the Enhanced FlashCopy group that has the Enhanced FlashCopy image that you want to delete. Enclose the name of the Enhanced FlashCopy group in double quotation marks (" ").
	If you enter more than one Enhanced FlashCopy group name, enclose each of the Enhanced FlashCopy group names in double quotation marks (" "). Enclose all of the Enhanced FlashCopy group names in parenthesizes.
	If you do not use any other parameters with the enhancedFlashCopyGroup parameter or enhancedFlashCopyGroups parameter, then by default the oldest Enhanced FlashCopy image is deleted.
deleteCount	The number of Enhanced FlashCopy images that you want to delete from the Enhanced FlashCopy group. Use integer values.
	This parameter deletes the oldest Enhanced FlashCopy image first and continues to delete the oldest Enhanced FlashCopy images until reaching the number that you enter.
	If the number that you enter is greater than the actual number of all of the Enhanced FlashCopy images in the Enhanced FlashCopy group, all of the Enhanced FlashCopy images will be deleted. The Enhanced FlashCopy group is left empty.
retainCount	The number of Enhanced FlashCopy images that you want to keep in the Enhanced FlashCopy group. Use integer values.
	This parameter keeps the most recent Enhanced FlashCopy images in the Enhanced FlashCopy group and deletes older Enhanced FlashCopy images.
	If the number of existing Enhanced FlashCopy images in the Enhanced FlashCopy group is less than the number that you enter, none of the Enhanced FlashCopy images are deleted.

Parameter	Description
ignoreEnhancedFlashCopyLogicalDrive	Use this parameter to make sure that you do not delete an Enhanced FlashCopy image that has an Enhanced FlashCopy logical drive associated with the Enhanced FlashCopy image. You can use one of these values:
	• TRUE - Use this value to delete the Enhanced FlashCopy image even though the Enhanced FlashCopy image has an associated Enhanced FlashCopy logical drive.
	• FALSE - Use this value to stop the Enhanced FlashCopy image deletion when the Enhanced FlashCopy image has an associated Enhanced FlashCopy logical drive.
	The default value is TRUE.
enhancedFlashCopyImageID	The <b>enhancedFlashCopyImageID</b> parameter accepts only the 0LDEST option. This parameter deletes the earliest Enhanced FlashCopy image created.

You can delete the oldest Enhanced FlashCopy image from an Enhanced FlashCopy group repository logical drive. The definition of an Enhanced FlashCopy image that you delete is removed from the system. The space occupied by the Enhanced FlashCopy image that you delete from the Enhanced FlashCopy group repository logical drive is released and made available for reuse within the Enhanced FlashCopy group.

Any Enhanced FlashCopy logical drives that exist for an Enhanced FlashCopy image migrate to the Stopped state when the Enhanced FlashCopy image is deleted.

This command will not run when the controller is in Lockdown mode.

# Delete Enhanced FlashCopy Logical Drive

This command deletes an Enhanced FlashCopy logical drive and optionally the associated Enhanced FlashCopy repository members.

**Note:** You cannot use this command for Enhanced FlashCopy images involved in online VolumeCopy.

### Syntax

delete enhancedFlashCopyLogicalDrive ["enhancedFlashCopyLogicalDriveName"]
[deleteRepositoryMembers=(TRUE | FALSE)]

## **Parameters**

Parameter	Description
enhancedFlashCopyLogicalDrive	The name of the Enhanced FlashCopy logical drive that you want to delete. Enclose the Enhanced FlashCopy logical drive identifier in double quotation marks (" ") inside square brackets ([]).
deleteRepositoryMembers	The parameter to save or delete the repository members. To save the repository members, set this parameter to FALSE. To delete the repository members, set this parameter to TRUE. The default value is TRUE. If you do not use this parameter the repository members are automatically deleted.

# **Delete FlashCopy Logical Drive**

This command deletes one or more FlashCopy logical drives or FlashCopy repository logical drives. You can also use this command to remove schedules for creating FlashCopies .

**Attention:** Possible damage to the storage subsystem configuration All of the data in the logical drive is lost as soon as you run this command.

## **Syntax**

delete (logicaldrive [logicaldriveName] |
logicaldrives [logicaldriveName1 ... logicaldriveNameN])
[schedule]

## **Parameters**

Parameter	Description
logicaldrive or logicaldrives	The name of the FlashCopy logical drive that you want to delete. You can enter more than one FlashCopy logical drive name. Enclose the FlashCopy logical drive name in square brackets ([]). If the FlashCopy logical drive name has special characters, you also must enclose the FlashCopy logical drive name in double quotation marks (" ").
schedule	This parameter deletes the schedule for a specific FlashCopy logical drive. Only the schedule is deleted, the FlashCopy logical drive remains.

## **Delete Host**

This command deletes a host.

## **Syntax**

delete host [hostName]

# **Delete Host Group**

This command deletes a host group.

**Attention: Possible damage to the storage subsystem configuration** – This command deletes all of the host definitions in the host group.

### Syntax

delete hostGroup [hostGroupName]

### Parameter

Parameter	Description
hostGroup	The name of the host group that you want to delete. Enclose the host group name in square brackets ([]). If the host group name has special characters, you also must enclose the host group name in double quotation marks ("").

### Notes

A host group is an optional topological element that is a collection of hosts that share access to the same logical drives. The host group is a logical entity.

## **Delete Host Port**

This command deletes an HBA host port identification. The identification is a software value that represents the physical HBA host port to the controller. By deleting the identification, the controller no longer recognizes instructions and data from the HBA host port.

### Syntax

delete hostPort [hostPortName]

### Parameter

Parameter	Description
hostPort	The name of the HBA host port that you want to delete. Enclose the name of the HBA host port in square brackets ([]).

### Notes

An HBA host port is a physical connection on a host bus adapter that resides within a host computer. An HBA host port provides a host access to the logical drives in a storage subsystem. If the host bus adapter has only one physical connection (one host port), the terms HBA host port and host bus adapter are synonymous.

## **Delete iSCSI Initiator**

This command deletes a specific iSCSI initiator object.

## Syntax

delete iscsiInitiator (iSCSI-ID | name)

## Parameters

Parameter	Description
iSCSI-ID	The identifier of the iSCSI initiator that you want to delete. Enclose the name in double quotation marks ("").
name	The name of the iSCSI initiator that you want to delete. Enclose the name in double quotation marks ("").

# **Delete Logical Drive**

This command deletes one or more standard logical drives, FlashCopy logical drives, or FlashCopy repository logical drives.

**Attention:** All of the data in the logical drive is lost as soon as you run this command.

## **Syntax**

```
delete (allLogicalDrives | logicalDrive [logicalDriveName] |
logicalDrives [logicalDriveName1 ... logicalDriveNameN])
removesubsystem=(TRUE | FALSE)
```

## Parameters

Parameter	Description
allLogicalDrives	This parameter deletes all of the logical drives in a storage subsystem.
logicalDrive or logicalDrives	The name of the logical drive that you want to delete. You can enter more than one logical drive name separated by spaces. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks ("").
removesubsystem	Deleting the last logical drive in an subsystem does not delete the subsystem. You can have a stand-alone subsystem (minus any logical drives). To remove the stand-alone subsystem, set this parameter to TRUE. To keep stand-alone subsystems intact, set this parameter to FALSE.

## Notes

When you use the **allLogicalDrives** parameter, this command deletes logical drives until all of the logical drives are removed or until an error is encountered. If an error is encountered, this command does not try to delete the remaining logical drives. Deleting logical drives from different subsystems is possible. All of the subsystems that become empty are deleted if you set the **removesubsystem** parameter to TRUE.

If you want to delete an entire subsystem, you can also use the **delete subsystem** command.

# **Delete VolumeGroup**

This command deletes one or more standard logical drives, Enhanced FlashCopy logical drives, or FlashCopy repository logical drives.

### Syntax

**Attention:** All of the data in the logical drive group is lost as soon as you run this command.

This command deletes an entire VolumeGroup and its associated logical drives. delete VolumeGroup [VolumeGroupName]

#### **Parameters**

None

## **Delete Logical Drive on a Disk Pool**

This command deletes either normal or thin logical drives on a disk pool.

**Attention:** Possible damage to the storage subsystem configuration All of the data in the logical drive is lost as soon as you run this command.

#### Syntax

```
delete (allLogicalDrives |
logicaldrive [logicaldriveName] |
logicaldrives [logicaldriveName1 ... logicaldriveNameN] | alllogicalDrives)
[removesubsystem=(TRUE | FALSE)
force=(TRUE | FALSE)
schedule
retainRepositoryMembers=(TRUE | FALSE)]
```

### **Parameters**

Parameter	Description
allLogicalDrives	This parameter deletes all of the logical drives on a disk pool.
logicaldrive or logicaldrives	The name of the logical drive that you want to delete. You can enter more than one logical drive name. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks (" ").
removesubsystem	Deleting the last logical drive in an subsystem does not delete the subsystem. You can have a standalone subsystem (minus any logical drives). To remove the standalone subsystem, set this parameter to TRUE. To keep standalone subsystems intact, set this parameter to FALSE.

Parameter	Description
force	Use this parameter to force the immediate deletion of a logical drive even if the controllers are performing other operations. To immediately force the deletion of a logical drive, set this parameter to TRUE. To wait until the controllers have finished performing other operations, do not use this parameter or set this parameter to FALSE.
schedule	
retainRepositoryMembers	When you delete a thin logical drive, the associated repository logical drive is deleted by default. However, when the <b>retainRepositoryMembers</b> is set to TRUE, the repository logical drive is retained. For normal logical drives, this parameter has no effect.

When you use the **allLogicalDrives** parameter, this command deletes logical drives until all of the logical drives are removed or until an error is encountered. If an error is encountered, this command does not try to delete the remaining logical drives. Deleting logical drives from different subsystems is possible. All of the subsystems that become empty are deleted if you set the **removesubsystem** parameter to TRUE.

# **Delete Performance Read Cache**

This command deletes the Performance Read Cache. All data in the SSD cache is purged.

## **Syntax**

delete performanceReadCache [performanceReadCacheName]

## **Parameters**

Header	Header
performanceReadCache	The alphanumeric identifier (including - and _) of the Performance Read Cache that you want to delete. Enclose the identifier in square brackets ([]). If the Performance Read Cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.

## **Minimum Firmware Level**

7.84

# **Diagnose Controller**

This command runs diagnostic tests on the controller. The diagnostic tests consist of loopback tests in which data is written to the disk drives and read from the disk drives.

## **Syntax**

```
diagnose controller [(a | b)]
loopbackDriveChannel=(allchannels | (1 | 2 | 3 | 4 | 5 | 6 | 7 | 8))
testID=(1 | 2 | 3 | discreteLines)
[patternFile="filename"]
```

## **Parameters**

Parameter	Description
controller	The controller on which you want to run the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
loopbackDriveChannel	The disk drive channels on which you want to run the diagnostic tests. You can either choose to run the diagnostics on all channels or select a specific channel on which to run diagnostics. If you select a specific channel, valid values for the disk drive channels are 1, 2, 3, or 4.
testID	<ul> <li>The identifier for the diagnostic test you want to run. The identifier and corresponding tests are as follows:</li> <li>1 - Read test</li> <li>2 - Write test</li> <li>3 - Data loop-back test</li> <li>discreteLines - Discrete lines diagnostic test</li> </ul>
patternFile	The name of a file that contains a data pattern that you want to use as test data. Enclose the file name of the data pattern in double quotation marks ("").

## Notes

When you run a data loop-back test, you can optionally specify a file that contains a data pattern. If you do not specify a file, the controller firmware provides a default pattern.

Discrete lines are control lines and status lines that are connected between two controllers in a controller module. The discrete lines diagnostic test lets each controller check that control signal migrations can be observed at the control inputs of the alternate controller. The discrete lines diagnostic test automatically runs after each power-cycle or each controller-reset. You can run the discrete lines diagnostic test after you have replaced a component that failed the initial discrete lines diagnostic test. This test applies only to the DS5100, DS5300, and DS4800 (models 82, 84, and 88) series controller modules. The discrete lines diagnostic test returns one of the following messages:

- When the discrete lines diagnostic test runs successfully, this message appears: The controller discrete lines successfully passed the diagnostic test. No failures were detected.
- If the discrete lines diagnostic test fails, this message appears:

One or more controller discrete lines failed the diagnostic test.

- If the CLI cannot run the discrete lines diagnostic test, the CLI returns Error 270, which means that the discrete lines diagnostic test could not start nor complete.
- The discrete lines diagnostic test is not functioning in SM release version 10.60.x5.11. Check the SM release readme files for information.

# **Diagnose Controller iSCSI Host Cable**

This command runs diagnostic tests on the copper cables between iSCSI Host interface cards and a controller. You can run diagnostics on a selected port or all ports. The ports must be able to support the cable diagnostics. If the ports do not support cable diagnostics an error is returned.

## Syntax

diagnose controller [(a | b)] iscsiHostPorts=(all | ("wwID" | "gID") testID=cableDiagnostics

## **Parameters**

Parameter	Description
controller	The controller on which you want to run the cable diagnostic test. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
iscsiHostPorts	The 8-byte World Wide Identifier (WWID) or the 16-byte group identifier (GID) of the HBA or HCA host port. Enclose the WWID or the GID in double quotation marks (" ").
testID	The identifier for the diagnostic test that you want to run. For this diagnostic test, the only choice is <b>cableDiagnostics</b> .

## Notes

When you run the cable diagnostic test, the firmware returns the following information:

- Host Port: The port on which the diagnostic test was run.
- HIC: The host interface card associated with this port.
- The date and time the test was run.
- Status:
  - OK: All of the pairs of cables are good and do not have any faults.
  - Open: One or more of the four pairs of cables are open.
  - Short: One or more of the four pairs of cables are shorted.
  - Incomplete: One or more of the four pairs returned incomplete or invalid test results.
- Length The Length of the cables are listed in meters and the following information about the cables is returned:

- When the cable status is OK, the approximate lengths of the cable pairs are returned. The lengths of the cable pair is are shown as a range (L1-L2), which are the shortest and the longest lengths of the cable pairs.
- If the cable status is Open or Short, the approximate distance to the failure in the cable pairs. If there is one failure, the length is reported for that cable pair. If there is more than one failure, the information returned is both the shortest and longest lengths to the failures. The lengths are listed as a range (L1-L2) where L1<L2.</li>
- If the cable status is Incomplete, the information returned are the lengths for the shortest and longest cable pairs that the firmware can successfully test. The lengths are listed for the valid cable pairs as a range (L1-L2) where L1<L2.</li>
- Register values for the cable diagnostic registers. The values are in a hexadecimal format:
  - Two bytes show the combined cable status (four bits per port).
  - Four two-byte numbers show the length of each channel.

## **Diagnose Remote Mirror**

This command tests the connection between the specified primary logical drives and the mirror logical drives on a storage subsystem with the Enhanced Remote Mirroring feature enabled.

### **Syntax**

```
diagnose remoteMirror (primary [primarylogicalDriveName] |
primaries [primaryLogicalDriveName1 ... primarylogicalDriveNameN])
testID=connectivity
```

### **Parameters**

Parameter	Description
primary or primaries	The name of the primary logical drive of the remote mirror pair that you want to test. You can enter more than one primary logical drive name. Enclose the primary logical drive names in square brackets ([]). If the primary logical drive name has special characters, you also must enclose the primary logical drive name in double quotation marks ("").

# **Diagnose Enhanced Remote Mirroring**

This command tests the connection between the specified primary logical drives and the mirror logical drives on a storage subsystem with the Enhanced Remote Mirroring premium feature enabled.

### Syntax

```
diagnose syncMirror (primary [primarylogicalDriveName] | △
primaries [primarylogicalDriveName1 ... primarylogicalDriveNameN]) △
testID=connectivity
```

**Note:** In previous versions of this command the feature identifier was remoteMirror. This feature identifier is no longer valid and is replaced by syncMirror.

## **Parameters**

None

# **Disable External Security Key Management**

This command disables external security key management for a storage subsystem that has full disk encryption drives.

## **Syntax**

disable storageSubsystem externalKeyManagement
file="fileName"
passPhrase="passPhraseString"

## Parameters

Parameter	Description
file	The file path and the file name that has the security key. For example: file="C:\Program Files\CLI\sup\seckey.slk" <b>Important:</b> The file name must have an extension of <i>.slk</i> .
passPhrase	A character string that encrypts the security key so that you can store the security key in an external file.

## Notes

Your pass phrase must meet these criteria:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, <> @ +.

**Note:** If your pass phrase does not meet these criteria, you will receive an error message.

# **Disable FlashCopy**

This command stops a copy-on-write operation. This command performs the same action as the legacy stop flashcopy command.

## **Syntax**

disableFlashcopy (logicaldrive [logicaldriveName] |
logicaldrives [logicaldriveName1 ... logicaldriveNameN])

Parameter	Description
logicaldrive or logicaldrives	The name of the specific logical drive for which you want to stop a copy-on-write operation. You can enter more than one logical drive name.
	Enclose the logical drive names using one of these forms:
	<ul> <li>On a Windows command line: \" logicaldriveName\"</li> </ul>
	<ul> <li>In a Windows script engine window: [" logicaldriveName"]</li> </ul>
	<ul> <li>On a Linux command line: \" logicaldriveName\"</li> </ul>
	<ul> <li>In a Linux script engine window: [\" logicaldriveName\"]</li> </ul>

## **Parameters**

### Notes

Names can be any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#). Names can have a maximum of 30 characters.

One technique for naming the FlashCopy logical drive and the FlashCopy repository logical drive is to add a hyphenated suffix to the original base logical drive name. The suffix distinguishes between the FlashCopy logical drive and the FlashCopy repository logical drive. For example, if you have a base logical drive with a name of Engineering Data, the FlashCopy logical drive can have a name of Engineering Data-S1, and the FlashCopy repository logical drive can have a name of EngineeringData-R1.

If you do not choose a name for either the FlashCopy logical drive or the FlashCopy repository logical drive, the storage management software creates a default name by using the base logical drive name. An example of the FlashCopy logical drive name that the controllers might create is, if the base logical drive name is aaa and does not have a FlashCopy logical drive, the default FlashCopy logical drive name is *aaa-1*. If the base logical drive already has n-1 number of FlashCopy logical drive name is aaa and does not have a FlashCopy repository logical drive, the base logical drive name is *aaa-n*. An example of the FlashCopy repository logical drive name is aaa and does not have a FlashCopy repository logical drive, the default flashCopy repository logical drive name is aaa and does not have a FlashCopy repository logical drive, the default FlashCopy repository logical drive name is *aaa-R1*. If the base logical drive name is *aaa-R1*. If the base logical drive name is *aaa-R1*. If the base logical drive name is *aaa-R1*.

## **Disable Storage Subsystem Feature**

This command disables a storage subsystem premium feature. Run the **show storagesubsystem** command to show a list of all enabled premium features in the storage subsystem.

## **Syntax**

disable storageSubsystem [featurePack |

The **featureAttributeList** can be one or more of these attribute values. If you enter more than one attribute value, separate the values with a white space.

- logical VolumeCopy
- Enhanced FlashCopy
- remoteMirror
- mixedDriveTypes
- Full Disk Encryption (safeStoreSecurity)
- storagePartition or storagePartition
   [2|4|8|16|32|64|96|128|192|256|512|Max]

In certain versions of the controller firmware, the form for this attribute value is a combination of the alphabetical term merged with the numerical value, as shown by this example: **storagePartition256** 

```
    driveSlotLimit or
driveSlotLimit[16|24|32|48|60|64|72|96|112|10|128|136|144|180|192|256|
272|300|360|384|448|480|Max]
```

In certain versions of the controller firmware, the form for this attribute value is a combination of the alphabetical term merged with the numerical value, as shown by this example: driveSlotLimit360

- ExternalKeyMgr (safeStoreExternalKeyMgr)
- T10PI
- SSDSupport
- highPerformanceTier
- raid6
- FlashCopy or FlashCopy [2|4|8|16]

The form for this attribute value is a combination of the alphabetical term merged with the numerical value, as shown by this example: **flashcopy16** 

Note: This attribute is for enabling the FlashCopy feature only.

```
    remoteMirror or remoteMirror [8 | 16 | 32 | 64 | 128]
```

In certain versions of the controller firmware, the form for this attribute value is a combination of the alphabetical term merged with the numerical value, as shown by this example: **remoteMirror128** 

## Parameter

None

### Notes

If you specify the **remoteMirror** parameter, this command disables the Enhanced Remote Mirroring premium feature and takes away the structure of the mirror repository logical drive.

If you disable the High Performance Tier feature, all of the included features are disabled.

# **Display Automatic Support Bundle Collection Configuration**

This command displays the automatic support bundle collection settings.

### Syntax

**Attention:** The SM CLI support bundle commands must be run from the SM CLI command line interface only. They cannot be run from the storage subsystem SM CLI script window. In addition, the storage subsystem IP addresses are not required. Also, the **-c** parameter is not required to precede the SM CLI support bundle command.

-supportBundle auto show program path SMcliadd program path SMcli

#### **Parameters**

None

## **Display Support Bundle Collection Schedule**

This command displays the schedule for collecting support bundles for all storage subsystems.

### Syntax

**Attention:** The SM CLI support bundle commands must be run from the SM CLI command line interface only. They cannot be run from the storage subsystem SM CLI script window. In addition, the storage subsystem IP addresses are not required. Also, the **-c** parameter is not required to precede the SM CLI support bundle command.

-supportBundle schedule show

### **Parameters**

None.

# **Download Drive Firmware**

### Syntax

**Attention:** Possible damage to the storage subsystem configuration – downloading drive firmware incorrectly can result in damage to the drives or a loss of data access.

This command is intended for downloading a firmware image to only one drive at a time. If you use this command in a script, make sure that you use this command only once. If you use this command more than once, the operation can fail. You can download firmware images to all of the drives in a storage subsystem at onetime by using the download **storageSubsystem driveFirmware** command. download drive [enclosureID,drawerID,slotID] firmware file="filename"
# Parameters

Parameter	Description
drive	The drive to which you want to download the firmware image. For high-capacity drive enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for each drive to which you want to download firmware. For low-capacity drive enclosures, specify the enclosure ID value and the slot ID value for each drive to which you want to download firmware. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID values, the drawer ID values, and the slot ID values in brackets ([]).
file	The file path and the file name of the file that contains the firmware image. Enclose the file path and the file name of the firmware image in double quotation marks (" "). For example: file="C:\Program Files\CLI\dnld\drvfrm.dlp" Valid file names have a <i>.dlp</i> extension.

## Notes

Before trying to download drive firmware, take these precautions:

- Stop all I/O activity to the storage subsystem before you download the firmware image. The download drive command blocks all I/O activity until the download finishes or fails; however, as a precaution, make sure that all I/O activity that might affect the drive is stopped.
- Ensure that the firmware image file is compatible with the drive enclosure. If you download a firmware image file that is not compatible with the drive enclosure that you have selected, the drive enclosure might become unusable.
- Do not make any configuration changes to the storage subsystem while you download drive firmware. Trying to make a configuration change can cause the firmware download to fail and make the selected drives unusable.

When you download the firmware to the drives, you must provide the full path and file name to the firmware image that is stored on your system.

You can use download drive command to test the firmware on one drive before you install the firmware on all of the drives in a storage subsystem. The download returns one of these statuses:

- Successful
- Unsuccessful With Reason
- Never Attempted With Reason

The **drive** parameter supports both high-capacity drive enclosures and low-capacity drive enclosures. A high-capacity drive enclosure has drawers that hold the drives. The drawers slide out of the drive enclosure to provide access to the drives. A low-capacity drive enclosure does not have drawers. For a high-capacity drive enclosure, you must specify the identifier (ID) of the drive enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive enclosure, you need only specify the ID of the drive enclosure, an alternative method for identifying a location for a drive is to specify the ID of the slot in which a drive resides.

# **Download Enclosure Configuration Settings**

This command downloads the factory default settings to all of the drive enclosures in a storage subsystem or to a specific drive enclosure in a storage subsystem.

#### **Syntax**

download (allEnclosures | enclosure [enclosureID]) configurationSettings
firmware file="filename"

#### **Parameters**

Parameter	Description
allEnclosure	This parameter downloads new firmware to all of the enclosures in the storage subsystem.
enclosure	The drive enclosure that contains the ESM card to which you want to load new firmware. Enclosure ID values are $\theta$ to 99. Enclose the enclosure ID value in square brackets ([]).
file	The file path and the file name of the file that contains the firmware image. Enclose the file path and the file name of the firmware image in double quotation marks (" "). For example: file="C:\Program Files\CLI\dnld\ enclosureset.dlp". Valid file names have a .dlp extension.

#### Notes

The enclosure parameter downloads the factory default configuration settings to a specific drive expansion enclosure. If you need to download the factory default configuration settings to more than one drive expansion enclosure, but not all drive expansion enclosures, you must enter this command for each drive expansion enclosure.

# **Download Environmental Card Firmware**

This command downloads environmental services module (ESM) firmware.

#### Syntax

```
download (allEnclosures | enclosure [enclosureID])
firmware file="filename"
```

Parameter	Description
allEnclosure	This parameter downloads new firmware to all of the enclosures in the storage subsystem.
enclosure	The drive expansion enclosure that contains the ESM card to which you want to load new firmware. Enclosure ID values are $0$ to 99. Enclose the enclosure ID value in square brackets ([ ]).

Parameter	Description
file	The file path and the file name of the file that contains the firmware image. Enclose the file path and the file name of the firmware image in double quotation marks (""). For example:
file="C:\Program Files\CLI\dnld\esmfrm.esm" Valid file names have a .esm extension.	file="C:\Program Files\CLI\dnld\esmfrm.esm"
	Valid file names have a .esm extension.

The **enclosure** parameter downloads new firmware to a specific drive expansion enclosure. If you need to download new firmware to more than one drive expansion enclosure, but not all drive expansion enclosures, you must enter this command for each drive expansion enclosure.

# **Download Power Supply Firmware**

This command downloads firmware updates to the power supplies. You can schedule simultaneous firmware updates for several power supplies, and the power supplies can be in different enclosures. A single firmware file can contain updates for several different power supplies. Matching firmware updates are automatically chosen for the power supplies. Firmware download occurs only if the new firmware version is not the same as the version of the power supplies on the enclosure. A download succeeds only if the power supply is in an Optimal state and there is a redundant power supply that is in an Optimal state.

To bypass these checks 'forceUpdate' can be used.

## **Syntax**

```
download (allEnclosures |
...enclosure [enclosureID1]...[enclosureIDn] |
...enclosure [enclosureID])
powerSupplyUpdate file="filename"
powerSupplyUnit [(left | right) | (top | bottom)] |
[forceUpdate]
```

Parameter	Description
allenclosures	This parameter downloads new power supplyfirmware to all of the enclosures in the storage subsystem.
enclosure or enclosures	The enclosure that contains the power supply to which you want to download new firmware. Enclosure ID values are $\theta$ to 99. Enclose the enclosure ID value in square brackets ([]).
powerSupplyUpdate file	The file path and the file name of the file that contains the firmware image. Enclose the file path and the file name of the firmware image in double quotation marks (" "). For example: file="C:\Program Files\CLI\dnld\esmfrm.esm". Valid file names have an <i>.esm</i> extension.
powerSupplyUnit	The power supply to which you want to download new firmware. Valid power supply identifiers are left, right, top, or bottom. Enclose the power-fan CRU/FRU identifier in square brackets ([])

Parameter	Description	
forceUpdate	This parameter bypasses these checks:	
	• To determine if the new firmware version is the same as the existing firmware version.	
	• To determine if the power supply is in an Optimal state.	
To determ     in an Opt	• To determine if there is a redundant power supply that is in an Optimal state.	

# **Download Storage Subsystem Drive Firmware**

This command downloads firmware images to all of the disk drives in the storage subsystem.

#### Syntax

```
download storageSubsystem driveFirmware file="filename"
[file="filename2"... file="filenameN"]
```

#### Parameter

Parameter	Description
file	The file path and the file name of the file that contains the firmware image. Enclose the file path and the file name of the firmware image in double quotation marks ("").

#### **Notes**

When you run this command, you can download more than one firmware image file to the disk drives in a storage subsystem. The number of firmware image files that you can download depends on the storage subsystem. The storage management software returns an error if you try to download more firmware image files than the storage subsystem can accept.

You can schedule downloads for multiple disk drives at the same time, including multiple disk drives in a redundant subsystem. Each firmware image file contains information about the disk drive types on which the firmware image runs. The specified firmware images can be downloaded only to a compatible disk drive. Use the **download drive firmware** command to download a firmware image to a specific disk drive.

The download storageSubsystem driveFirmware command blocks all I/O activity until a download try has been made for each candidate disk drive or you run the stop storageSubsystem downloadDriveFirmware command. When the download storageSubsystem driveFirmware command finishes downloading the firmware image, each candidate disk drive shows the download status for each disk drive. One of these statuses is returned:

- Successful
- Unsuccessful With Reason
- Never Attempted With Reason

# **Download Storage Subsystem Firmware/NVSRAM**

This command downloads firmware and, optionally, NVSRAM values for the storage subsystem controller. If you want to download only NVSRAM values, use the **downLoad storageSubsystem NVSRAM** command.

## **Syntax**

```
download storageSubsystem firmware [, NVSRAM ]
file= "filename" [, "NVSRAM-filename"]
[downgrade=(TRUE | FALSE)]
[activateNow=(TRUE | FALSE)]
```

## **Parameters**

Parameter	Description	
NVSRAM	The setting to download a file with NVSRAM values when you download a firmware file. Do not include square brackets with this parameter. Include a comma after the <b>firmware</b> parameter.	
file	The file path and the name of the file that contains the firmware. Valid file names must end with a .dlp extension. Enclose the file path and the file name in double quotation marks ("").	
NVSRAM-filename	The file path and the name of the file that contains the NVSRAM values. Valid file names must end with a .dlp extension. Enclose the NVSRAM file name in double quotation marks (""). Include a comma after the file name.	
downgrade	Attention:	
	• Downgrading the controller firmware is not a supported function. This option should be used <b>only</b> under the direction of IBM Support.	
	See "Software service and support" on page xvi for more information.	
	• Downgrading from 07.xx to 06.xx firmware levels is not supported and will return an error if attempted.	
	The <b>downgrade</b> parameter is the setting used to load firmware from a previous version. The default value is FALSE. Set the <b>downgrade</b> parameter to TRUE if you want to download an earlier version of firmware.	
activateNow	The setting to activate the firmware image and the NVSRAM image. The default value is TRUE. If you set the <b>activateNow</b> parameter to FALSE, you must run the <b>activate storageSubsystem firmware</b> command to activate the firmware values and the NVSRAM values at a later time.	

# **Download Storage Subsystem NVSRAM**

This command downloads NVSRAM values for the storage subsystem controller.

## Syntax

download storageSubsystem NVSRAM file="filename"

## Parameter

file	The file path and the file name that contains the NVSRAM values. Enclose the NVSRAM file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\dnld\afrm.dlp" Valid file names have a .dlp extension.

# **Enable subsystem Security**

This command converts a non-secure subsystem to a secure subsystem.

#### **Syntax**

enable subsystem [subsystemName] security

#### Parameters

Parameter	Description	
subsystem	The alphanumeric identifier (including - and _) of the subsystem that you want to place in the Security Enabled state. Enclose the subsystem identifier in square brackets.	

#### Notes

The following conditions must be met to successfully run this command.

- All drives in the subsystem must be FDE drives.
- The Full Disk Encryption premium feature must be enabled.
- The storage subsystem security key has to be set.
- The subsystem is Optimal, and it does not have FlashCopy logical drives or repository logical drives.

Once the subsystem is security enable, all of the drives in the subsystem will be in "secured" state, which will be in locked state upon power on until the controllers unlock them with the appropriate key. In addition, the data will be encrypted or decrypted by the drive before it is written to or read from to the disk surface. Once the FDE drive is placed into "secured state", it will stay in the secure state until one re-provision the FDE drive.

# **Enable Controller Data Transfer**

This command revives a controller that has become quiesced while running diagnostics.

#### Syntax

enable controller [(a | b)] dataTransfer

# Parameter

Parameter	Description
controller	The controller that you want to revive. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

# **Enable Disk Pool Security**

This command converts a non-secure disk pool to a secure disk pool.

Note: All of the drives that comprise the disk pool must be security capable.

#### **Syntax**

enable diskPool [diskPoolName] security

#### Parameters

Parameter	Description
diskPool	The name of the disk pool that you want to place in the Security Enabled state. Enclose the disk pool identifier in square brackets ([ ]).

#### **Notes**

Each disk pool name must be unique. You can use any combination of alphanumeric characters, underscore (\_), hyphen(-), and pound (#) for the user label. User labels can have a maximum of 30 characters.

# **Enable External Security Key Management**

This command enables external security key management for a storage subsystem that has full disk encryption drives.

#### Syntax

enable storageSubsystem externalKeyManagement
file="fileName" |
passPhrase="passPhraseString"

Parameter	Description
file	The file path and the file name that has the security key. Enclose the file path and the file name that has the security key in double quotation marks (" "). For example: file="C:\Program Files\CLI\sup\seckey.slk" Important: The file name must have an extension of <i>.slk</i> .
passPhrase	A character string that encrypts the security key so that you can store the security key in an external file. Enclose the pass phrase character string in double quotation marks (" ").

Your pass phrase must meet these criteria:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, <> @ +.

**Note:** If your pass phrase does not meet these criteria, you will receive an error message.

## Enable array security

This command converts a non-secure array to a secure array.

#### Syntax

enable volumeGroup [volumeGroupName] security

#### Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier (including - and _) of the array that you want to place in the Security Enabled state. Enclose the array identifier in square brackets ([]).

#### Notes

The following conditions must be met to successfully run this command:

- All drives in the storage subsystem must be full disk encryption drives.
- The Full Disk Encryption premium feature must be enabled.
- The storage subsystem security key has to be set.
- The storage subsystem is Optimal, and it does not have FlashCopy volumes or repository volumes.

The controller firmware creates a lock that restricts access to the FDE drives. FDE drives have a state called Security Capable. When you create a security key, the state is set to Security Enabled, which restricts access to all FDE drives that exist within the storage subsystem.

#### Minimum Firmware Level

7.40

# Enable or disable Performance Read Cache for a volume

This command turns on or off caching using the Performance Read Cache feature for a specific volume. The volume can be either a standard volume, a Enhanced FlashCopy volume, or a consistency group Enhanced FlashCopy volume.

# Syntax Applicable to a Standard Volume

set volume ["volumeName"] performanceReadCacheEnabled=(TRUE |
FALSE)

## Syntax Applicable to a Enhanced FlashCopy Volume

set enhancedFlashCopyLogicalDrive ["enhancedFlashCopyLogicalDriveName"]
performanceReadCacheEnabled=(TRUE | FALSE)

# Syntax Applicable to a Consistency Group Enhanced FlashCopy Volume

set cgenhancedFlashCopyLogicalDrive ["cgenhancedFlashCopyLogicalDriveName"]
performanceReadCacheEnabled=(TRUE | FALSE)

#### **Parameters**

Parameter	Description
<pre>volume, enhancedFlashCopyLogicalDrive, or cgenhancedFlashCopyLogicalDrive</pre>	The name of the specific volume for which you want to turn on or off the Performance Read Cache. Enclose the volume name in double quotation marks (" ") inside of square brackets ([]).
performanceReadCacheEnabled	To turn on Performance Read Cache, set this parameter to TRUE. To turn off Performance Read Cache, set this parameter to FALSE.

#### Notes

You can turn the Performance Read Cache on or off for only one volume at a time.

When you turn off Performance Read Cache for a volume, the Performance Read Cache for that volume is purged.

#### **Minimum Firmware Level**

7.84

# Enable Enhanced FlashCopy volume Performance Read Cache

This command turns on or turns off Performance Read Cache for Enhanced FlashCopy volume.

#### **Syntax**

set enhancedFlashCopyLogicalDrive ["enhancedFlashCopyLogicalDriveName"]
performanceReadCacheEnabled=(TRUE | FALSE)
performanceReadCacheName="performanceReadCacheName"]

Parameter	Description
enhancedFlashCopyLogicalDrive	The alphanumeric identifier (including - and _) of the Enhanced FlashCopy volume for which you are setting properties. Enclose the Enhanced FlashCopy volume identifier in double quotation marks (" ") inside of square brackets ([]).

Parameter	Description
performanceReadCacheEnabled	The setting to turn on or turn off Performance Read Cache for the Enhanced FlashCopy volume. To turn on Performance Read Cache, set this parameter to TRUE. To turn off Performance Read Cache, set this parameter to FALSE
performanceReadCacheName	The alphanumeric identifier (including - and _) that you want to give to the Performance Read Cache. Enclose the Performance Read Cache identifier in double quotation marks (" ").

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

#### **Minimum Firmware Level**

7.83

# Enable Storage Subsystem Feature Key

This command enables a premium feature by using a feature key file.

#### Syntax

**Attention:** When enabling the featurePack, this command must be performed when the subsystem is off-line, without IO activities from the host servers, because both storage subsystem controllers will be rebooted during the disabling process. Schedule storage subsystem down time before running this command.

enable storageSubsystem [featurePack |feature]
file="filename"

#### Parameters

Parameter	Description
file	The file path and the file name of a valid feature key file. Enclose the file path and the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\dnld\ftrkey.key" Valid file names for feature key files end with a <i>.key</i> extension.

## Notes

A premium feature is an additional application to enhance the capabilities of a storage subsystem. The premium features are these:

- Storage partitioning
- FlashCopies
- Enhanced Remote Mirroring
- Mixed drive types
- High performance tier

- SSD support
- Full Disk Encryption

A feature pack is a predefined set of enhanced storage subsystem capabilities and/or premium features, such as Storage Partitioning and Enhanced Remote Mirroring. These premium features are combined for the convenience of the users.

# **Establish Enhanced Global Mirrored Pair**

Use this command to complete an Enhanced Global Mirrored Pair on the remote storage subsystem by adding a secondary logical drive to an existing Enhanced Global Mirror Group. Before you run this command, the Enhanced Global Mirror Group must exist and the primary logical drive must exist in the Enhanced Global Mirror Group. After this command successfully completes, Enhanced Global Mirroring starts between the primary logical drive and the secondary logical drive.

#### Syntax

establish enhancedGlobalMirror logicalDrive="secondarylogicalDriveName" enhancedGlobalMirrorGroup="enhancedGlobalMirrorGroupName" primarylogicalDrive="primaylogicalDriveName"

## **Parameters**

Parameter	Description
logicalDrive	The name of an existing logical drive on the remote storage subsystem that you want to use for the secondary logical drive. Enclose the logical drive name in double quotation marks (" ").
enhancedGlobalMirrorGroup	The name of an existing Enhanced Global Mirror Group that you want to use to contain the Enhanced Global Mirrored Pair. Enclose the Enhanced Global Mirror Group name in double quotation marks (" ").
primarylogicalDrive	The name of an existing logical drive on the local storage subsystem that you want to use for the primary logical drive. Enclose the logical drive name in double quotation marks (" ").

## Notes

An Enhanced Global Mirrored Pair comprises of two logical drives, one primary and the other secondary logical drive, that contain identical copies of the same data. The mirrored pair is a part of an Enhanced Global Mirror Group, which allows the mirrored pair to synchronize at the same time as any other mirrored pairs within the Enhanced Global Mirror Group .

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

When you choose the primary logical drive and the secondary logical drive, the secondary logical drive must be of equal or greater size than the primary logical drive. The RAID level of the secondary logical drive need not be the same as the primary logical drive.

## **Minimum Firmware Level**

#### 7.84

# Export Storage Subsystem Security Key

This command saves a full disk encryption (FDE) security key to a file. You can transfer the file from one storage subsystem to another storage subsystem. The file enables you to move FDE drives between storage subsystems.

#### Syntax

export storageSubsystem securityKey
passPhrase=passPhraseString |
file=fileName

#### **Parameters**

Parameter	Description
passPhrase	The character string that provides authentication for the security key. The pass phrase is 8 to 32 characters in length. You must use at least one number, one lowercase letter, one uppercase letter, and one nonalphanumeric character in the pass phrase. A space is not permitted.
file	The name of the file to which you want to save the security key. You must add a file extension of .slk to the end of the file name.

#### Notes

For more information about exporting the storage subsystem security key, see the *DS Storage Manager Installation and User's Guide* version 10.

The pass phrase must:

- Be between eight and 32 characters long.
- · Contain at least one uppercase letter.
- Contain at least one lowercase letter.
- Contain at least one number.
- Contain at least one non-alphanumeric character, for example, < > @ +.

If your pass phrase does not meet these criteria, you receive an error message and are prompted to retry the command.

## Import Storage Subsystem Security Key

This command unlocks one or more Full Disk Encryption (FDE) drives that you have imported from one storage subsystem to another storage subsystem. Only the FDE drives with the matching security key from the imported storage subsystem are unlocked. After they are unlocked, the security key for the new storage subsystem is applied.

#### Syntax

```
import storageSubsystem securityKey file=fileName |
passPhrase=passPhraseString
```

# Parameters

Parameter	Description
file	The file that has the original security key of the imported FDE drives. The file has an extension of .slk.
passPhrase	The character string that provides authentication for the security key. The pass phrase is 8 to 32 characters in length. You must use at least one number, one lowercase letter, one uppercase letter, and one non-alphanumeric character in the pass phrase. A space is not permitted.

## Notes

For more information about exporting the storage subsystem security key, see the DS Storage Manager version 10 Installation and User's Guide.

You must use the same pass phrase that you use when exporting the storage subsystem security key. If your pass phrase does not match, you receive an error message and the FDE drives remain locked. For more information about exporting the storage subsystem security key, see the *DS Storage Manager Installation and User's Guide* version 10.

# **Increase Logical Drive Capacity on a Disk Pool**

Short description

Note: You cannot use this command to increase the capacity of a thin logical drive.

#### **Syntax**

```
start increaseVolCapacity logicaldrive="logicaldriveName"
incrementalCapacity=logicaldriveCapacity
[addDrives=(enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn)]
```

Parameter	Description
logicaldrive	The name of the disk pool logical drive for which you want to increase capacity. Enclose the logical drive name in double quotation marks (" ").
incrementalCapacity	The setting to increase the storage size (capacity) for the logical drive. Size is defined in units of <i>bytes</i> , <i>KB</i> , <i>MB</i> , <i>GB</i> , or <i>TB</i> . The default value is <i>bytes</i> .
addDrives	The setting to add new drives to the logical drive. For high-capacity drive expansion enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive expansion enclosures, specify the enclosure ID value and the slot ID value for the drive. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID value, drawer ID value, and the slot ID value, drawer ID value, and the slot ID value in parentheses.

In some cases a drive parameter might appear as valid input for the command syntax. However, you cannot use the drive parameter with this command.

Setting the **incrementalCapacity** parameter, starts a long-running operation that you cannot stop. Long-running operations are performed in the background and do not prevent you from running other commands. To show the progress of long-running operations, use the show logicaldrive actionProgress command.

The **addDrives** parameter supports both high-capacity drive expansion enclosures and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need to only specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

# Initialize Thin Logical Drive

This command initializes or re-initializes a thin logical drive.

**Note:** Initializing a thin logical drive starts a long-running operation that you cannot stop.

#### Syntax

```
start logicaldrive [logicaldriveName] initialize
[existingRepositoryLabel=existingRepositoryName |
diskPool="diskPoolName" capacity=capacityValue|
retainRepositoryMembers=[TRUE|FALSE]]
```

Parameter	Description
logicaldrive	The name of the logical drive for which you are starting the formatting. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks (" ").

Parameter	Description
existingRepositoryLabel	This parameter replaces the repository logical drive by the candidate logical drive specified:
	• The value specified is an existing repository logical drive user label. The logical drive specified must be an unused repository logical drive with the name in the proper form.
	• If the newly specified repository logical drive is on a different disk pool, the thin logical drive will change ownership to that pool.
	• The old repository logical drive will be deleted by default.
diskPool	This parameter identifies the disk pool in which you want to create a new repository logical drive with the specified capacity.
	You must use this parameter with the <b>capacity</b> parameter to create a new repository logical drive with the specified capacity.
capacity	The size that you want to set for the repository logical drive that you are creating. Size is defined in units of bytes, KB, MB, GB, or TB.
	You must use this parameter with the <b>diskPool</b> parameter to create a new repository logical drive with the specified capacity.
retainRepositoryMembers	If this parameter is set to TRUE, the old repository is retained. By default, the old repository is deleted. This parameter is ignored if the existing repository is reused.

# Load Storage Subsystem DBM Database

This command uploads a Database Management (DBM) database image from a file or from cache. This command restores a storage subsystem DBM database to the exact configuration that existed when the database image was captured using the save storageSubsystem dbmDatabase command. The data in a file can be just RAID configuration data or all data, including data for RAID configuration, subsystems and disk pools. The data in a cache location always includes all data.

Before using this command with the file option, you must first obtain a validator string (a security code) from your IBM Technical Support representative. To obtain a validator, use the save storageSubsystem dbmValidator command to generate an XML file that contains validator information. Your IBM Technical Support representative uses the XML file to generate the validator string required for this command.

## Syntax

load storageSubsystem dbmDatabase
((file="filename" validator="validatorValue") | sourceLocation=onboard)
[controller [(a|b)]]

#### **Parameters**

Parameter	Description
file	The file path and the file name of the DBM database you want to upload. Enclose the file name in double quotation marks (" "). For example: file="C:\Subsystem Backups\DBMbackup_03302010.dbm"
	This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.
validator	The alphanumeric security code required to restore a storage subsystem to an existing configuration. Use the <b>save</b> <b>storageSubsystemdbmValidator</b> command to generate the required validation information XML file. After the validation information XML file is available, contact your IBM Technical Support representative to obtain the Validator.
sourceLocation	This parameter specifies the location from which to retrieve backup database information. The parameter must be included for consistency, but the only allowed value is onboard
controller	This parameter specifies the controller from which data will be exclusively retrieved, if sourceLocation is set to cache. Enclose the controller identifier in squarebrackets ([]). If the controller parameter is not specified, data might be retrieved from either controller.

## Notes

Depending on the size of the database image, restoring the database might take up as much as 30 minutes. The host software will not show the controllers in an Optimal state until after all actions for loading the database image are completed on the controllers.

# Locate Performance Read Cache

The start locate command identifies the Solid State Disks (SSDs) that are being used in the Performance Read Cache by turning on the indicator lights for the SSDs. The stop locate command turns off the indicator lights on the SSDs.

# Syntax Applicable to Starting a Locate Operation

start performanceReadCache [performanceReadCacheName] locate

## Syntax Applicable to Stopping a Locate Operation

stop performanceReadCache locate

#### **Parameters**

None

# **Recopy VolumeCopy**

This command re-initiates a VolumeCopy operation by using an existing VolumeCopy pair.

Attention: Starting a VolumeCopy operation overwrites all existing data on the target logical drive, makes the target logical drive read-only to hosts, and fails all FlashCopy logical drives associated with the target logical drive, if any exist. If you have used the target logical drive as a copy before, be sure you no longer need the data or have it backed up.

This command works with logical VolumeCopy pairs that you created with a FlashCopy logical drive or without a FlashCopy logical drive.

#### Syntax

```
recopy VolumeCopy target [
targetName]
[source [sourceName]]
[copyPriority=(highest | high | medium | low | lowest)
targetReadOnlyEnabled=(TRUE | FALSE)
copyType=(online | offline)]
```

Parameter	Description
target	The name of the target logical drive for which you want to reinitiate a VolumeCopy operation. Enclose the target logical drive name in square brackets ([]). If the target logical drive name has special characters, you also must enclose the target logical drive name in double quotation marks ("").
source	The name of the source logical drive for which you want to reinitiate a VolumeCopy operation. Enclose the source logical drive name in square brackets ([]). If the source logical drive name has special characters, you also must enclose the source logical drive name in double quotation marks ("").
copyPriority	The priority that the VolumeCopy has relative to host I/O activity. Valid values are highest, high, medium, low, or lowest.
targetReadOnlyEnabled	The setting so that you can write to the target logical drive or only read from the target logical drive. To write to the target logical drive, set this parameter to FALSE. To prevent writing to the target logical drive, set this parameter to TRUE.
соруТуре	Use this parameter to create a VolumeCopy with a FlashCopy . Creating a VolumeCopy with a FlashCopy enables you to continue to write to the source logical drive while creating the VolumeCopy. To reinitiate a VolumeCopy with a FlashCopy, set this parameter to online. To reinitiate a VolumeCopy without a FlashCopy, set this parameter to offline.

Copy priority defines the amount of system resources that are used to copy the data between the source logical drive and the target logical drive of a VolumeCopy pair. If you select the highest priority level, the VolumeCopy uses the most system resources to perform the VolumeCopy, which decreases performance for host data transfers.

# **Recover RAID Logical Drive**

This command creates a RAID logical drive with the given properties without initializing any of the user data areas on the disk drives. Parameter values are derived from the Recovery Profile data file for the storage subsystem. You can create the recover logical drive in an existing subsystem or create a new subsystem by using this command.

**Note:** You can run this command only from a command line. You cannot run this command from the GUI script editor. You cannot use the DS Storage Management GUI to recover a logical drive.

#### Syntax

```
recover logicalDrive (drive=(enclosureID,drawerID,slotID) |
Drives=(enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn) |
subsystem=subsystemName)
[newsubsystem=subsystemName]
userLabel=("logicalDriveName"
logicaldriveWWN="logicaldriveWWN")
capacity=logicalDriveCapacity
offset=offsetValue
raidLevel=(0 | 1 | 3 | 5 | 6)
segmentSize=segmentSizeValue
dssPreallocate=(TRUE | FALSE)
SSID=subsystemLogicalDriveID
[owner=(a | b)
cacheReadPrefetch=(TRUE | FALSE)]
T10PI=(none | enabled)]
```

Parameter	Description
<b>drive</b> or <b>drives</b>	The disk drives that you want to assign to the logical drive that is to be recovered. For high-capacity drive expansion enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for each disk drive that you assign to the logical drive. For low-capacity drive expansion enclosures, specify the enclosure ID value and the slot ID value for each disk drive that you assign to the logical drive. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID, drawer ID, and slot ID values in square parentheses ([]).
subsystem	The name of the subsystem in which you want to create the logical drive. (To determine the names of the subsystems in your storage subsystem, run the <b>show storageSubsystem profile</b> command.)

Parameter	Description
newsubsystem	The name that you want to give to a new subsystem. Enclose the new subsystem name in double quotation marks (" ").
userLabel	The name of the logical drive that you want to recover. Enclose the logical drive name in double quotation marks ("").
logicaldriveWWN	The world wide name of the logical drive that you want to recover. The name is a 16-byte identifier, for example, 60080E500017B432000000049887D77. Enclose the identifier in double quotation marks ("").
capacity	The size of the logical drive that you are adding to the storage subsystem. Size is defined in units of bytes, KB, MB, GB, or TB.
offset	The number of blocks from the start of the subsystem to the start of the referenced logical drive.
raidLevel	The RAID level of the subsystem that contains the disk drives. Valid values are $0$ , 1, 3, 5, or 6.
segmentSize	The amount of data (in KB) that the controller writes on a single disk drive in an subsystem before writing data on the next disk drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
dssPreallocate	The setting to turn on or turn off allocating logical drive storage capacity for future segment size changes. To turn on allocation, set this parameter to TRUE. To turn off allocation, set this parameter to FALSE.
SSID	The storage subsystem identifier of a logical drive. Use the <b>show logical drive</b> command to determine the storage subsystem identifier.
owner	The controller that owns the logical drive. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. If you do not specify an owner, the controller firmware determines the owner.
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.
T10PI	<ul> <li>The setting to specify that an subsystem, and the logical drives within the subsystem, has T10 PI (Protection Information) to help you ensure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the subsystem. These settings are valid:</li> <li>none—The subsystem does not have T10 PI (Protection Information).</li> <li>enabled—The subsystem has T10 PI</li> </ul>
	(Protection Information). The subsystem support is protected information and is formatted with protection information enabled.

The storage management software collects recovery profiles of the monitored storage subsystems and saves the profiles on a storage management station.

The **drive** parameter supports both high-capacity and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need to specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

If you attempt to recover a logical drive using the **drive** parameter or the **drives** parameter and the drives are in an unassigned state, the controller automatically creates a new subsystem. Use the **newLogicalDriveGroup** parameter to specify a name for the new subsystem.

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

The **owner** parameter defines which controller owns the logical drive. The preferred controller ownership of a logical drive is the controller that currently owns the subsystem.

#### Preallocating Storage Capacity

With the **dssPreallocate** parameter, you can assign space in a logical drive for storing information that is used to rebuild a logical drive. When you set the **dssPreallocate** parameter to TRUE, the storage space allocation logic in the controller firmware preallocates space for future segment size changes. The preallocated space is the maximum allowable segment size. The **dssPreallocate** parameter is necessary for recovering logical drive configurations that are not retrievable from the controller database. To turn off the preallocation capability, set **dssPreallocate** to FALSE.

#### Segment Size

The size of a segment determines how many data blocks that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single disk drive for a single request leaves other disk drives available to simultaneously service other requests.

If the logical drive is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of disk drives in the subsystem that are used for data transfers.) In this case, multiple disk drives are used for the same request, but each disk drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of disk drives that are required to satisfy a data transfer request.

## **Cache Read Prefetch**

Cache read prefetch lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from disk into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage subsystem that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the **cacheReadPrefetch** parameter are TRUE or FALSE.

# **Re-create Enhanced Remote Mirroring Repository Logical Drive**

**Note:** With firmware version 7.80, the **re-create storageSubsystem mirrorRepository** command is deprecated. This command is no longer supported in either the GUI or the CLI. If you attempt to run this command, you get an error message that this functionality is no longer supported and that no changes will be made to the specified remote mirror repositories.

This command creates a new Enhanced Remote Mirroring repository logical drive (also called a mirror repository logical drive) by using the parameters defined for a previous mirror repository logical drive. The underlying requirement is that you have previously created a mirror repository logical drive. When you use this command, you can define the mirror repository logical drive in one of three ways: user-defined disk drives, user-defined subsystem, or user-defined number of disk drives for the mirror repository logical drive. If you choose to define a number of disk drives, the controller firmware chooses which disk drives to use for the mirror repository logical drive.

## Syntax (User-Defined Disk Drives)

```
re-create storageSubsystem mirrorRepository
repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDrives=(enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn)
[enclosureLossProtect=(TRUE | FALSE)]
T10=(none | enabled)]
```

# Syntax (User-Defined subsystem)

re-create storageSubsystem mirrorRepository
repositorysubsystem=subsystemNumber[freeCapacityArea=freeCapacityIndexNumber]

# Syntax (User-Defined Number of Disk Drives)

re-create storageSubsystem mirrorRepository
repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDriveCount=numberOfDrives
[driveType=(fibre | SATA | SAS)]

DriveMediaType=(HDD | SSD | unknown | allMedia) [repositorysubsystemUserLabel="repositorysubsystemName" enclosureLossProtect=(TRUE | FALSE) drawerLossProtect=(TRUE | FALSE)] T10PI=(none | enabled)]

Parameter	Description
repositoryRAIDLevel	The RAID level for the mirror repository logical drive. Valid values are 1, 3, 5, or 6.
repositoryDrives	The disk drives for the mirror repository logical drive. EnclosureID1,drawerID1,slotID1 enclosureIDn,drawerIDn,slotIDn is for enclosures with multiple disk drive drawers. For enclosures without disk drive drawers, use enclosureID1,slotID1 enclosureIDn,slotID1 instead. Enclosure ID values are $\theta$ to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID values and slot ID values in parentheses.
repositorysubsystem	The name of the subsystem where the mirror repository logical drive is located.
freeCapacityArea	The index number of the free space in an existing subsystem that you want to use to re-create the mirror repository logical drive. Free capacity is defined as the free capacity between existing logical drives in an subsystem. For example, an subsystem might have the following areas: logical drive 1, free capacity, logical drive 2, free capacity, logical drive 3, free capacity. To use the free capacity following logical drive 2, you would specify: freeCapacityArea=2 Run the <b>show subsystem</b> command to determine if a free capacity area exists.
repositoryDriveCount	The number of unassigned disk drives that you want to use for the mirror repository logical drive.
driveType	The type of disk drives that you want to use for the mirror repository logical drive. You cannot mix drive types. You must use this parameter when you have more than one type of drive in your storage subsystem. Valid disk drive types are fibre, SATA, and SAS. If you do not specify a drive type, the command defaults to fibre.
enclosureLossProtect	The setting to enforce enclosure loss protection when you create the mirror repository logical drive. To enforce enclosure loss protection, set this parameter to TRUE. The default value is FALSE.

Parameter	Description
T10PI	The setting to specify that an subsystem, and the logical drives within the subsystem, has T10 PI (Protection Information) to help you ensure that the data maintains its integrity. When you use this parameter, only protected drives can be used for the subsystem. These settings are valid:
	• <b>none</b> —The subsystem does not have T10 PI (Protection Information).
	• <b>enabled</b> —The subsystem has T10 PI (Protection Information). The subsystem support is protected information and is formatted with protection information enabled.

If you enter a value for the storage space of the mirror repository logical drive that is too small, the controller firmware returns an error message, which states the amount of space that is needed for the mirror repository logical drive. The command does not try to change the mirror repository logical drive. You can re-enter the command by using the value from the error message for the storage space value of the mirror repository logical drive.

When you assign the disk drives, if you set the **enclosureLossProtect** parameter to TRUE and have selected more than one disk drive from any one enclosure, the storage subsystem returns an error. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs operations, but the mirror repository logical drive that you create might not have enclosure loss protection.

When the controller firmware assigns the disk drives, if you set the **enclosureLossProtect** parameter to TRUE, the storage subsystem returns an error if the controller firmware cannot provide disk drives that result in the new mirror repository logical drive having enclosure loss protection. If you set the **enclosureLossProtect** parameter to FALSE, the storage subsystem performs the operation even if it means that the mirror repository logical drive might not have enclosure loss protection.

# **Re-create External Security Key**

This command regenerates a storage subsystem security key for use with the external security key management feature.

#### **Syntax**

```
re-create storageSubsystem securityKey
passPhrase="passPhraseString"
file="fileName"
```

Parameter	Description
passPhrase	A character string that encrypts the security key so that you can store the security key in an external file.

Parameter	Description
file	The file path and the file name that has the security key. For example: file="C:\Program Files\CLI\sup\seckey.slk" <b>Important:</b> The file name must have an extension of <i>.slk</i> .

Your pass phrase must meet these criteria:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, <> @ +.

**Note:** If your pass phrase does not meet these criteria, you will receive an error message.

# **Re-create FlashCopy**

This command starts a fresh copy-on-write operation by using an existing FlashCopy logical drive.

#### Syntax

```
re-create FlashCopy (logicalDrive
[logicalDriveName] | logicalDrives [logicalDriveName1 ... logicalDriveNameN)
[userLabel="flashcopyLogicalDriveName"
[warningThresholdPercent=percentValue
repositoryFullPolicy (failBaseWrites | failFlashCopy)]
```

Parameter	Description
logicalDrive or logicalDrives	The name of the specific logical drive. You can enter more than one logical drive name. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you must also enclose the logical drive name in double quotation marks ("").
userLabel	The name of the FlashCopy logical drive. Enclose the FlashCopy logical drive name in double quotation marks (""). If you enter more than one FlashCopy logical drive name, this command fails.
warningThresholdPercent	The percentage of repository capacity at which you receive a warning that the FlashCopy repository logical drive is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 50.

Parameter	Description
repositoryFullPolicy	The type of processing that you want to continue if the FlashCopy repository logical drive is full. You can choose to fail writes to the base logical drive (failBaseWrites) or fail writes to the FlashCopy logical drive (failFlashCopy). The default value is failFlashCopy.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

If you do not specify a value for the **warningThresholdPercent** parameter or the **repositoryFullPolicy** parameter, the previously set value is used.

#### Recreating Multiple FlashCopy Logical Drives without Optional Parameters

- If you list multiple FlashCopy logical drives to be re-created but do not specify any of the optional parameters, the re-create operation processes the FlashCopy logical drives as a "batch" process.
- Before restarting any FlashCopy, validation checks for the necessary restarted preconditions of the FlashCopy are performed. If any of the listed FlashCopy logical drives fail the validation, the entire command fails and the FlashCopy logical drives are not re-created. If the validation is successful for all of the FlashCopy logical drives in the list, but one or more of the legacy FlashCopies in the list fail to restart, the entire command fails and none of the FlashCopies are re-created.
- During FlashCopy recreation, all affected logical drives (FlashCopy, base, and repository) are appropriately quiesced and I/O operations are resumed to all affected logical drives after all FlashCopy are re-created.

# **Remove Logical Drive LUN Mapping**

This command removes the logical unit number (LUN) mapping.

#### Syntax

remove (allLogicalDrives | logicalDrive [logicalDriveName] |
logicalDrives [logicalDriveName1 ... logicalDriveNameN] | accesslogicalDrive)
lunMapping (host="hostName" |
hostGroup=("hostGroupName" | defaultGroup))

Parameter	Description
allLogicalDrives	This parameter removes the LUN mapping from all of the logical drives.
logicalDrive or logicalDrives	The name of the specific logical drive that you want to remove from the LUN mapping. You can enter more than one logical drive name. Enclose the logical drive name in double quotation marks ("") inside of square brackets ([]).
accessLogicalDrive	This parameter removes the access logical drive.

Parameter	Description
host	The name of the host to which the logical drive is mapped. Enclose the host name in double quotation marks ("").
hostGroup	The name of the host group that contains the host to which the logical drive is mapped. Enclose the host group name in double quotation marks (""). The <b>defaultGroup</b> value is the host group that contains the host to which the logical drive is mapped. The HostGroup and Host parameters can not be specifed together. You can specify either parameter but not both at the same time.

The access logical drive is the logicalDrive in a SAN environment that is used for communication between the storage management software and the storage subsystem controller. The access logical drive uses a LUN address and consumes 20 MB of storage space that is not available for application data storage. An access logical drive is required only for in-band managed storage subsystems.

Attention: Removing an access logical drive can damage your configuration – The agent uses the access logical drives to communicate with a storage subsystem. If you remove an access logical drive mapping for a storage subsystem from a host that has an agent running on it, the storage management software is no longer able to manage the storage subsystem through the agent.

You must use the **host** parameter and the **hostGroup** parameter when you specify a non-access logical drive or an access logical drive. The script engine ignores the **host** parameter or the **hostGroup** parameter when you use the **allLogicalDrives** parameter or the **logicalDrives** parameter.

# **Remove drives from Performance Read Cache**

This command reduces the capacity of the disk pool by removing the drives from the pool.

#### **Syntax**

set performanceReadCache [performanceReadCacheName]
removeDrives=(enclosureID1,drawerID1,slotID1 ...
enclosureIDn,drawerIDn,slotIDn)

Parameter	Description
performanceReadCache	The alphanumeric identifier (including - and _) of the Performance Read Cache from which you want to remove SSDs. Enclose the identifier in square brackets ([]). If the Performance Read Cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.

Parameter	Description
removeDrives	<ul> <li>The drives that you want to remove from the Performance Read Cache. For high-capacity drive enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for each SSD that you want to remove. For low-capacity drive enclosures, specify the enclosure ID value and the slot ID value for each SSD that you want to remove:</li> <li>Enclosure ID values are 0 to 99</li> <li>Drawer ID values are 1 to 5</li> <li>Slot ID values are 1 to 32</li> </ul>
	ID values, and the slot ID values in
	parentheses.

You cannot remove all of the SSDs from the Performance Read Cache using this command; at least one SSD must remain in the Performance Read Cache. If you want to completely remove the Performance Read Cache, use the delete performanceReadCache command instead.

## **Minimum Firmware Level**

7.84

# Remove incomplete Enhanced Global Mirrored Pair from an Enhanced Global Mirror Group

This command removes an orphaned mirrored pair volume on the storage subsystem. An orphaned mirrored pair volume exists when a member volume in an Enhanced Global Mirror Group has been removed on one side of the Enhanced Global Mirror Group (either the primary side or secondary side) but not on the other side.

Orphan mirrored pair volumes are detected when inter-controller communication is restored and the two sides of the mirror configuration reconcile mirror parameters.

Use this command when the mirror relationship is successfully removed on the local or remote storage subsystem, but cannot be removed on the corresponding storage subsystem because of a communication problem.

#### **Syntax**

remove enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"]
incompleteMirror volume="volumeName"

## Parameters

Parameter	Description
enhancedGlobalMirrorGroup	The name of the Enhanced Global Mirror Group that contains the orphaned volume that you want to remove. Enclose the Enhanced Global Mirror Group name in double quotation marks (" "). If the Enhanced Global Mirror Group name has special characters, you also must enclose the Enhanced Global Mirror Group name in double quotation marks (" ") inside square brackets.
volume	The name of the orphaned volume that you want to remove from the Enhanced Global Mirror Group. Enclose the volume name in double quotation marks (" ").

## **Minimum Firmware Level**

7.84

# **Remove volume from Enhanced Global Mirror Group**

This command removes a member volume from an existing Enhanced Global Mirror Group. Optionally, you can delete the repository volume members from the Enhanced Global Mirror Group.

This command is valid only on the local storage subsystem that contains the Enhanced Global Mirror Group whose member volume that you want to remove.

## **Syntax**

```
remove volume ["volumeName"]
enhancedGlobalMirrorGroup="enhancedGlobalMirrorGroupName"
[deleteRepositoryMembers=(TRUE | FALSE)]
```

Parameter	Description
volume	The name of the specific volume that you want to remove from the Enhanced Global Mirror Group. Enclose the volume name in double quotation marks (" ") inside square brackets ([ ]).
enhancedGlobalMirrorGroup	The name of the Enhanced Global Mirror Group that contains the member volume that you want to remove. Enclose the enhanced global mirror group name in double quotation marks (" ").
deleteRepositoryMembers	Determines whether to remove all of the repository members from the Enhanced Global Mirror Group.

## **Minimum Firmware Level**

7.84

# **Remove Logical Drive LUN Mapping**

This command removes the logical unit number (LUN) mapping.

#### **Syntax**

remove (allLogicalDrives | logicalDrive [logicalDriveName] |
logicalDrives [logicalDriveName1 ... logicalDriveNameN] | accesslogicalDrive)
lunMapping (host="hostName" |
hostGroup=("hostGroupName" | defaultGroup))

#### **Parameters**

Parameter	Description
allLogicalDrives	This parameter removes the LUN mapping from all of the logical drives.
logicalDrive or logicalDrives	The name of the specific logical drive that you want to remove from the LUN mapping. You can enter more than one logical drive name. Enclose the logical drive name in double quotation marks ("") inside of square brackets ([]).
accessLogicalDrive	This parameter removes the access logical drive.
host	The name of the host to which the logical drive is mapped. Enclose the host name in double quotation marks ("").
hostGroup	The name of the host group that contains the host to which the logical drive is mapped. Enclose the host group name in double quotation marks (""). The <b>defaultGroup</b> value is the host group that contains the host to which the logical drive is mapped. The HostGroup and Host parameters can not be specifed together. You can specify either parameter but not both at the same time.

#### Notes

The access logical drive is the logicalDrive in a SAN environment that is used for communication between the storage management software and the storage subsystem controller. The access logical drive uses a LUN address and consumes 20 MB of storage space that is not available for application data storage. An access logical drive is required only for in-band managed storage subsystems.

Attention: Removing an access logical drive can damage your configuration – The agent uses the access logical drives to communicate with a storage subsystem. If you remove an access logical drive mapping for a storage subsystem from a host that has an agent running on it, the storage management software is no longer able to manage the storage subsystem through the agent.

You must use the **host** parameter and the **hostGroup** parameter when you specify a non-access logical drive or an access logical drive. The script engine ignores the **host** parameter or the **hostGroup** parameter when you use the **allLogicalDrives** parameter or the **logicalDrives** parameter.

# **Remove Member Logical Drive from Consistency Group**

This command removes a member logical drive from a an existing Enhanced FlashCopy consistency group. Optionally, you can delete the repository logical drive members from the consistency group.

#### **Syntax**

set consistencyGroup ["consistencyGroupName"] removeCGMemberLogicalDrive=memberName
[deleteRepositoryMembers=(TRUE | FALSE)]

#### **Parameters**

Parameter	Description
consistencyGroupName	The name of the consistency group from which you want to remove a member. Enclose the consistency group name in double quotation marks (" ") inside square brackets ([ ]).
memberName	The name of the member logical drive that you want to remove.
deleteRepositoryMembers	Determines whether to remove all of the repository members from the consistency group.

# **Remove Remote Mirror**

This command removes the mirror relationship between the primary logical drive and the secondary logical drive in a remote-mirror pair.

#### **Syntax**

```
remove remoteMirror (localLogicalDrive [logicalDriveName] |
localLogicalDrives [logicalDriveName1 ... logicalDriveNameN])
```

#### Parameters

Parameter	Description
localLogicalDrive or localLogicalDrives	The name of the primary logical drive (the logical drive on the local storage subsystem) that you want to remove. You can enter more than one logical drive name. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks ("").

# **Remove Enhanced Remote Mirroring**

This command removes the mirror relationship between the primary volume and the secondary volume in a remote-mirrored pair.

**Note:** In previous versions of this command the feature identifier was remoteMirror. This feature identifier is no longer valid and is replaced by syncMirror.

## Syntax

remove syncMirror (localVolume [volumeName] |
localVolumes [volumeName1 ... volumeNameN])

#### **Parameters**

Parameter	Description
localVolume or localVolumes	The name of the primary volume (the volume on the local storage subsystem) that you want to remove. You can enter more than one volume name. Enclose the volume name in square brackets ([]). If the volume name has special characters, you also must enclose the volume name in double quotation marks (" ").

## **Minimum Firmware Level**

6.10

# **Remove VolumeCopy**

This command removes a VolumeCopy pair.

#### **Syntax**

remove VolumeCopy target [
targetName] [source [sourceName]
copyType=(online | offline)]

#### **Parameters**

Parameter	Description
target	The name of the target logical drive that you want to remove. Enclose the target logical drive name in square brackets ([]). If the target logical drive name has special characters, you also must enclose the target logical drive name in double quotation marks ("").
source	The name of the source logical drive that you want to remove. Enclose the source logical drive name in square brackets ([]). If the source logical drive name has special characters, you also must enclose the source logical drive name in double quotation marks ("").
соруТуре	Use this parameter to identify that a VolumeCopy has a FlashCopy . If the VolumeCopy has a FlashCopy , set this parameter to online. If the VolumeCopy does not have a FlashCopy, set this parameter to offline.

# **Rename Enhanced FlashCopy Logical Drive**

This command renames an existing Enhanced FlashCopy logical drive.

## **Syntax**

set enhancedFlashCopyLogicalDrive ["enhancedFlashCopyLogicalDriveName"]
userLabel="enhancedFlashCopyImageLogicalDriveName

## Parameters

Parameter	Description
enhancedFlashCopyLogicalDrive	The name of the Enhanced FlashCopy logical drive that you want to rename. Enclose the Enhanced FlashCopy logical drive name in double quotation marks (" ") inside of square brackets ([ ]).
userLabel	A new name that you want to give to the Enhanced FlashCopy logical drive. Enclose the new Enhanced FlashCopy logical drive name in double quotation marks (" ").

#### Notes

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

# **Rename Performance Read Cache**

This command changes the name of the Performance Read Cache.

#### **Syntax**

set performanceReadCache [old\_performanceReadCacheName]
userLabel="new\_performanceReadCacheName"

#### **Parameters**

Parameter	Description
performanceReadCache	The alphanumeric identifier (including - and _) of the Performance Read Cache that you want to rename. Enclose the identifier in square brackets ([]). If the Performance Read Cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.
userLabel	The new name for the Performance Read Cache. Enclose the name in double quotation marks (" ").You can use any combination of alphanumeric characters, underscore (_), hyphen (-), and pound (#) for the identifier. Identifiers can have a maximum of 30 characters.

## **Minimum Firmware Level**

7.84

# **Repair Logical Drive Parity**

This command repairs the parity errors on a logical drive.

# Syntax

repair logicaldrive [logicalDriveName] parity
parityErrorFile="filename"
[verbose=(TRUE | FALSE)]

## **Parameters**

Parameter	Description
logicalDrive	The name of the specific logical drive for which you want to repair parity. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks ("").
parityErrorFile	The name of the file that contains the parity error information that you use to repair the errors. Enclose the file name in double quotation marks ("").
verbose	The setting to capture progress details, such as percent complete, and to show the information as the logical drive parity is being repaired. To capture progress details, set this parameter to TRUE. To prevent capturing progress details, set this parameter to FALSE.

# **Replace Disk Drive**

This command redefines the composition of an subsystem. You can use this command to replace a disk drive with either an unassigned disk drive or a fully integrated hot spare.

## **Syntax**

replace Drive [enclosureID,drawerID,slotID] |
<"wwID"> replacementDrive=enclosureID,drawerID,slotID

Parameter	Description
drive	The location of the disk drive that you want to replace. For high-capacity expansion drawers, specify the enclosure ID, the drawer ID, and the slot ID. For low-capacity expansion drawers, specify the enclosure ID and the slot ID. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID value, the drawer ID value, and the slot ID value in square brackets ([]). You can also use the World Wide Identifier (WWID) of the disk drive. Enclose the WWID in double quotation marks <" ">
replacementDrive	The location of the replacement disk drive. For high-capacity expansion drawers, specify the enclosure ID, the drawer ID, and the slot ID for the disk drive. For low-capacity expansion drawers, specify the enclosure ID and the slot ID. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32.

The **diskDrive** parameter supports both high-capacity expansion drawers and low-capacity expansion drawers. A high-capacity expansion drawer has drawers that hold the disk drives. The drawers slide out of the expansion drawer to provide access to the disk drives. A low-capacity expansion drawer does not have drawers. For a high-capacity expansion drawer, you must specify the identifier (ID) of the expansion drawer, the ID of the drawer, and the ID of the slot in which a disk drive resides. For a low-capacity expansion drawer, you need only specify the ID of the expansion drawer and the ID of the slot in which a disk drive resides. For a low-capacity expansion drawer, an alternative method for identifying a location for a disk drive is to specify the ID of the expansion drawer, set the ID of the drawer to 0, and specify the ID of the slot in which a disk drive resides.

# **Reset Enhanced Global Mirror Group statistics**

This command resets the synchronization statistics for one or more member volumes in an Enhanced Global Mirror Group to a relative 0.

#### Syntax

```
reset storageArray arvmStats enhancedGlobalMirrorGroup
["enhancedGlobalMirrorGroupName"]
volume="volumeName" sampleType=(all | mostRecent |
longestSyncTime | errors)
```

Parameter	Description
enhancedGlobalMirrorGroup	The name of the Enhanced Global Mirror Group for which you are resetting the synchronization statistics. Enclose the Enhanced Global Mirror Group name in square brackets ([]). If the enhanced global mirror group name has special characters, you also must enclose the Enhanced Global Mirror Group name in double quotation marks (" ") inside square brackets. If the Enhanced Global Mirror Group name consists only of numbers, such as "1002," you also must enclose the enhanced global mirror group name in double quotation marks (" ") inside square brackets.
volume	This parameter is optional. The name of the specific member volume in the Enhanced Global Mirror Group for which you are resetting the synchronization statistics. If no volume is specified, the statistics for every member volume in the Enhanced Global Mirror Group mirror group are reset. Enclose the volume name in in double quotation marks (" ").

Parameter	Description
sampleType	This parameter is optional. The default value for sampleType is all
	• all— Data for all three sample types are reset.
	• mostRecent — Statistics are reset for the most recent 50 resynchronization samples.
	<ul> <li>longestSyncTime — Statistics are reset for the most recent 20 longest resynchronization samples.</li> </ul>
	• errors— Statistics are reset for the most recent 20 failed resynchronization samples.

Statistics are reset for mirrored volumes in the Primary role. The statistics that are reset include the following data:

- Synchronization start time
- Synchronization type (manual or periodic)
- Synchronization duration
- Number of bytes sent
- Maximum and minimum write time (for a single write)
- · Maximum and minimum synchronization data rate
- Total write time
- Repository utilization (%)
- Recovery point age

## Minimum Firmware Level

7.84

## **Reset Controller**

This command resets a controller, and it is disruptive to I/O operations.

**Important:** When you reset a controller, the controller is removed from the data path and is not available for I/O operations until the reset operation is complete. If a host is using logical drives that are owned by the controller being reset, the I/O directed to the controller is rejected. Before resetting the controller, either make sure that the logical drives that are owned by the controller are not in use or make sure that a multi-path driver is installed on all of the hosts that use these logical drives.

#### Syntax

reset controller [(a | b)]

## Parameter

Parameter	Description
controller	The controller that you want to reset. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller owner, the controller firmware returns a syntax error.

#### Notes

The controller that receives the reset controller command resets the controller specified. For example, if the reset controller command is sent to controller A to request a reset of controller A, then controller A reboots itself by doing a soft reboot. If the reset controller command is sent to controller A to request a reset of controller A holds controller B in reset and then releases controller B from reset, which is a hard reboot. A soft reboot in some products only resets the IOC chip. A hard reboot resets both the IOC and the expander chips in the controller.

# **Reset iSCSI IP address**

This command resets the IP address for the remote storage subsystem to re-establish connection with the local storage subsystem. You can use this command to notify the local storage subsystem that the iSCSI IP address of the remote storage subsystem has changed and needs to be updated.

When you establish an Enhanced Global Mirroring relationship with an iSCSI connection, both storage subsystems store the IP addresses of the remote storage subsystem. If the IP address of an iSCSI port changes, the remote storage subsystem that is attempting to use that port encounters a communication error.

The storage subsystem with the changed IP address sends a message to each remote storage subsystem associated with the Enhanced Global Mirror Groups that are configured to mirror over an iSCSI connection. Storage subsystems that receive this message automatically update their remote-target IP address.

If the storage subsystem with the changed IP address is unable to send its inter-controller message to a remote storage subsystem, the system sends you an alert of the connectivity issue. Use the reset command to re-establish connection with the local storage subsystem.

#### Syntax

```
reset (remoteStorageArrayName="storageArrayName" |
remoteStorageArrayWwid=<wwID>)
iscsiIpAddress
```

Parameter	Description
remoteStorageArrayName	The name for the remote storage subsystem for which you are resetting the iSCSI IP address. Enclose the storage subsystem name in double quotation marks (" ").
Parameter	Description
----------------------------	--
remoteStorageArrayNameWwid	The World Wide Identifier (WWID) of the storage subsystem for which you are resetting the iSCSI IP address. You can use the WWID instead of the storage subsystem name to identify the storage subsystem. Enclose the WWID in angle brackets (< >).

### **Minimum Firmware Level**

7.84

# **Reset Storage Subsystem Battery Install Date**

This command resets the age of the batteries in a storage subsystem to zero days. You can reset the batteries for an entire storage subsystem or just the battery for a specific controller or in a specific battery pack.

#### **Syntax**

reset storageSubsystem batteryInstallDate [controller=(a | b)
| batteryPack=(left | right)]

### Parameter

Parameter	Description
controller	The controller that contains the battery for which you want to reset the age. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B.
batteryPack	The battery pack contains both a left battery and a right battery. Valid identifiers are left or right, where left is the battery that supports the controller in slot A, and right is the battery that supports the controller in slot B. Use the batteryPack parameter only for controllers with battery packs. The controller and batterypack parameters are mutual exclusive. You can specify either parameter but not both at the same time. <b>Note:</b> The DS3000 does not have battery packs.

#### Notes

A controller might have a battery associated with it, so the controller is identified as either a or b. With the release of the DS5100, DS5300 controller enclosure, battery packs inside the interconnect-battery CRU/FRU are identified as either left or right. If the command statement uses the wrong parameter, an error appears.

# **Reset Storage Subsystem Diagnostic Data**

This command resets the NVSRAM that contains the diagnostic data for the storage subsystem. This command does not delete the diagnostic data. This command replaces the Needs Attention status with the Diagnostic Data Available status. The old diagnostic data is written automatically over when new data is captured. The memory that contains the diagnostic data is also cleared when the

controllers reboot. Before resetting the diagnostic data, use the **save storageSubsystem diagnosticData** command to save the diagnostic data to a file.

**Attention:** Run this command only with the assistance of an IBM Technical Support representative.

#### Syntax

reset storageSubsystem diagnosticData

#### **Parameters**

None.

## Reset Storage Subsystem iSCSI Baseline

This command resets the iSCI baseline for the storage subsystem to 0.

#### Syntax

reset storageSubsystem iscsiStatsBaseline

#### **Parameters**

None.

#### Notes

This command resets the baseline to 0 for both controllers in the storage subsystem. The purpose of resetting both of the controller baselines is to help make sure that the controller counts are synchronized between the controllers. If one controller resets but the second controller does not reset, the host is informed that the controllers are out of synchronization. The host is informed by the time stamps that are reported with the statistics.

### Reset Storage Subsystem Logical Drive Distribution

This command reassigns (moves) all of the logical drives to their preferred controller.

#### Syntax

reset storageSubsystem logicalDriveDistribution

#### **Parameters**

None.

#### Notes

If you use this command on a host without a multi-path driver, to prevent application errors, you must stop I/O operations to the logical drives until this command has completed.

Under certain host operating system environments, you might be required to reconfigure the multi-path host driver. You might also need to make operating system modifications to recognize the new I/O path to the logical drives.

# **Reset Storage Subsystem RLS Baseline**

This command resets the Read Link Status (RLS) baseline for all devices by setting all of the RLS counts to 0.

#### **Syntax**

reset storageSubsystem RLSBaseline

#### **Parameters**

None.

### Reset Storage Subsystem SAS PHY Baseline

This command resets the SAS physical layer (SAS PHY) baseline for all devices, and removes the list of errors from the .csv file. The .csv file is generated when you run the save storageSubsystem SASPHYCounts command.

**Note:** The previous release of the reset storageSubsystem SASPHYBaseline command cleared error counts for all devices except the drives. The reset storageSubsystem SASPHYBaseline command now resets the SAS PHY baseline for the drives as well as the other devices. All errors are deleted from the .csv file.

#### Syntax

reset storageSubsystem SASPHYBaseline

#### **Parameters**

None.

### Reset Storage Subsystem SOC Baseline

This command resets the baseline for all switch-on-a-chip (SOC) devices that are accessed through the controllers. This command resets the baseline by setting all of the SOC counts to 0. This command is valid only for Fibre Channel devices in an Arbitrated Loop topology.

#### Syntax

reset storageSubsystem SOCBaseline

#### Parameters

None.

# **Resume Enhanced Global Mirror Group**

This command resumes data transfer between all mirrored pairs in an Enhanced Global Mirror Group. Data written to the primary volumes while the Enhanced Global Mirror Group was suspended is written to the secondary volumes immediately. Periodic synchronization resumes if an automatic synchronization interval has been set.

#### Syntax

resume enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"]

### Parameter

Parameter	Description
enhancedGlobalMirrorGroup	The name of the Enhanced Global Mirror Group that you want to resume. Enclose the Enhanced Global Mirror Group name in square brackets ([]). If the Enhanced Global Mirror Group name has special characters, you also must enclose the Enhanced Global Mirror Group name in double quotation marks (" ") inside square brackets. If the enhanced global mirror group name consists only of numbers, such as "1002," you also must enclose the Enhanced Global Mirror Group name in double quotation marks (" ") inside square brackets.

### **Minimum Firmware Level**

7.84

# **Resume Consistency Group Enhanced FlashCopy Logical Drive**

This command restarts a copy-on-write operation for creating a consistency group Enhanced FlashCopy logical drive that you stopped using the stop cgEnhancedFlashCopyLogicalDrive command.

#### **Syntax**

resume cgEnhancedFlashCopyLogicalDrive ["enhancedFlashCopyLogicalDriveName"]
cgEnhancedFlashCopyImage="enhancedFlashCopyImageName"

Parameter	Description
cgEnhancedFlashCopyLogicalDrive	The name of the consistency group Enhanced FlashCopy logical drive for which you want to restart a rollback operation. Enclose the name of the consistency group Enhanced FlashCopy logical drive in double quotation marks (" ") inside square brackets ([]).

Parameter	Description
cgEnhancedFlashCopyImage	The name of the Enhanced FlashCopy image in a consistency group that you are restarting. The name of an Enhanced FlashCopy image is comprised of two parts:
	• The name of the consistency group
	• An identifier for the Enhanced FlashCopy image in the consistency group.
	The identifier for the Enhanced FlashCopy image can be one of these:
	• An integer value that is the sequence number of the Enhanced FlashCopy in the consistency group.
	• NEWEST - Use this option when you want to show the latest Enhanced FlashCopy image created in the consistency group.
	• OLDEST - Use this option when you want to show the earliest Enhanced FlashCopy image created in the consistency group.
	Enclose the Enhanced FlashCopy image name in double quotation marks (" ").

The name of an Enhanced FlashCopy image has two parts separated by a colon (:):

- The identifier of the Enhanced FlashCopy group
- The identifier of the Enhanced FlashCopy image

For example, if you want to restart a copy-on-write operation for Enhanced FlashCopy image **12345** which is in Enhanced FlashCopy consistency group enhancedFlashCopygroup1 in a consistency group Enhanced FlashCopy logical drive that has the name *enhancedFlashCopyVol1*, you would use this command:

resume cgEnhancedFlashCopyLogicalDrive ["enhancedFlashCopyVol1"]
cgEnhancedFlashCopyImage=["enhancedFlashCopygroup1:12345"]

# **Resume Enhanced FlashCopy Image Rollback**

This command resumes a rollback operation that has entered a paused state. A rollback operation can enter a paused state due to processing errors, which causes a **Needs Attention** condition for the storage subsystem.

If the rollback operation cannot be resumed, the selected Enhanced FlashCopy logical drive reverts to a paused state, and the **Needs Attention** condition is displayed.

**Note:** You cannot use this command for Enhanced FlashCopy images involved in online logical VolumeCopy.

#### **Syntax**

resume enhancedFlashCopyImage [enhancedFlashCopyImageName] rollback

Parameter	Description
enhancedF1ashCopyImage	The name of the Enhanced FlashCopy image for which you want to restart a rollback operation. The name of an Enhanced FlashCopy image is comprised of two parts:
	• The name of the Enhanced FlashCopy group
	• An identifier for the Enhanced FlashCopy image in the Enhanced FlashCopy group.
	The identifier for the Enhanced FlashCopy image can be one of these:
	• An integer value that is the sequence number of the Enhanced FlashCopy in the Enhanced FlashCopy group.
	• NEWEST - Use this option when you want to show the latest Enhanced FlashCopy image created in the Enhanced FlashCopy group.
	• OLDEST - Use this option when you want to show the earliest Enhanced FlashCopy image created in the Enhanced FlashCopy group.
	Enclose the Enhanced FlashCopy image name in double quotation marks (" ") inside square brackets ([ ]).

### **Parameters**

#### Notes

The name of an Enhanced FlashCopy image has two parts separated by a colon (:):

- The identifier of the Enhanced FlashCopy group
- The identifier of the Enhanced FlashCopy image

For example, if you want to restart a rollback operation for Enhanced FlashCopy image 12345 in an Enhanced FlashCopy group that has the name enhancedFlashCopygroup1, you would use this command:

resume enhancedFlashCopyImage ["enhancedFlashCopygroup1:12345"] rollback;

# **Resume Enhanced FlashCopy Logical Drive**

This command resumes an Enhanced FlashCopy logical drive operation that was stopped.

#### Syntax

resume enhancedFlashCopyLogicalDrive [enhancedFlashCopyLogicalDriveName] enhancedFlashCopyImage="enhancedFlashCopyImageName"

Parameters	;
------------	---

Parameter	Description
enhancedFlashCopyLogicalDrive	The name of the specific Enhanced FlashCopy logical drive for which you want to resume a rollback operation. Enclose the Enhanced FlashCopy logical drive identifier in square brackets ([]). If the Enhanced FlashCopy logical drive identifier has special characters, you must also enclose the Enhanced FlashCopy logical drive identifier in double quotation marks (" ").
enhancedFlashCopyImage	<ul> <li>The name of the Enhanced FlashCopy image for which you want to resume Enhanced FlashCopy logical drive operations. The name of an Enhanced FlashCopy image is comprised of two parts:</li> <li>The name of the Enhanced FlashCopy group</li> <li>An identifier for the Enhanced FlashCopy group.</li> <li>The identifier for the Enhanced FlashCopy group.</li> <li>The identifier for the Enhanced FlashCopy image can be one of these:</li> <li>An integer value that is the sequence number of the Enhanced FlashCopy in the Enhanced FlashCopy group.</li> <li>NEWEST - Use this option when you want to show the latest Enhanced FlashCopy group.</li> <li>OLDEST - Use this option when you want to show the earliest Enhanced FlashCopy group.</li> <li>Enclose the Enhanced FlashCopy image</li> </ul>
	name in double quotation marks (" ") inside square brackets ([ ]).

The name of an Enhanced FlashCopy image has two parts separated by a colon (:):

- The identifier of the Enhanced FlashCopy group
- The identifier of the Enhanced FlashCopy image

For example, if you want to resume Enhanced FlashCopy logical drive operations for Enhanced FlashCopy image **12345** in an Enhanced FlashCopy group that has the name *enhancedFlashCopyGroup1*, you would use this command: resume enhancedFlashCopyLogicalDrive [enhancedFlashCopyVol1] enhancedFlashCopyImage=["enhancedFlashCopyGroup1:12345"];

To resume Enhanced FlashCopy logical drive operations for the most recent Enhanced FlashCopy image in an Enhanced FlashCopy group that has the name *enhancedFlashCopyGroup1*, you would use this command:

# **Resume FlashCopy Rollback**

This command resumes a rollback operation that has entered a paused state. A rollback operation can enter a paused state due to processing errors, which will trigger a **Needs Attention** condition for the storage subsystem.

If the rollback operation cannot be resumed, the selected FlashCopy logical drive reverts to a paused state, and the **Needs Attention** condition is displayed.

### **Syntax**

resume rollback logicaldrive [flashcopyLogicalDriveName]

### **Parameters**

Parameter	Description
logicaldrive	The name of the specific FlashCopy logical drive for which you want to resume a rollback operation. Enclose the FlashCopy logical drive name in square brackets ([]). If the FlashCopy logical drive name has special characters, you must also enclose the FlashCopy logical drive name in double quotation marks (" ").

## **Resume Remote Mirror**

This command resumes a suspended Enhanced Remote Mirroring operation.

### **Syntax**

```
resume remoteMirror (primary [logicalDriveName] |
primaries [logicalDriveName1 ... logicalDriveNameN])
[writeConsistency=(TRUE | FALSE)}
```

Parameter	Description
primary or primaries	The name of the primary logical drive for which you want to resume operation. You can enter more than one primary logical drive name. Enclose the primary logical drive name in square brackets ([]). If the primary logical drive name has special characters, you also must enclose the primary logical drive name in double quotation marks ("").
writeConsistency	The setting to identify the logical drives in this command that are in a write-consistency group or are separate. For the logical drives to be in the same write-consistency group, set this parameter to TRUE. For the logical drives to be separate, set this parameter to FALSE.

If you set the **writeConsistency** parameter to TRUE, the logical drives must be in a write-consistency group (or groups). This command resumes all write-consistency groups that contain the logical drives. For example, if logical drives A, B, and C are in a write-consistency group and they have remote counterparts A', B', and C', the **resume remoteMirror logicalDrive ["A"] writeConsistency=TRUE** command resumes A-A', B-B', and C-C'.

## **Resume Performance Read Cache**

This command restarts the caching for all of the volumes using the Performance Read Cache that was temporarily stopped with the suspend performanceReadCache command.

#### Syntax

resume performanceReadCache [performanceReadCacheName]

#### Parameter

Parameter	Description
performanceReadCache	The alphanumeric identifier (including - and _) of the Performance Read Cache that you want to resume. Enclose the identifier in square brackets ([]). If the Performance Read Cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.

#### **Minimum Firmware Level**

7.84

# **Revive array**

This command forces the specified array and its associated failed disk drives to the Optimal state.

**Attention: Possible loss of data access or data integrity issues** – Correct use of this command depends on the data configuration on all of the disk drives in the array. Never try to revive a disk drive unless you are supervised by an IBM Technical Support representative.

#### Syntax

revive array [arrayName]

Parameter	Description
subsystem	The alphanumeric identifier (including - and _) of the array to be set to the Optimal state. Enclose the array identifier in square brackets ([]).

# **Revive Disk Drive**

This command forces the specified disk drive to the Optimal state.

Attention: Possible loss of data access or possible data integrity issues – Correct use of this command depends on the data configuration on all of the disk drives in the subsystem. Never try to revive a disk drive unless you are supervised by an IBM Technical Support representative.

#### Syntax

revive drive [enclosureID,drawerID,slotID]

#### Parameter

Parameter	Description
drive	The location of the drive that you want to revive. For high-capacity disk drive expansion enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for the disk drive. For low-capacity drive expansion enclosures, specify the enclosure ID value and the slot ID value for the disk drive. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID value, the drawer ID value, and the slot ID value in square brackets ([ ]).

### Notes

The **drive** parameter supports both high-capacity drive expansion enclosures and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need only specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

## **Revive Enhanced FlashCopy Group**

This command forces the specified Enhanced FlashCopy group to the Optimal state. If the Enhanced FlashCopy group is not in a **Failed** state, the firmware displays an error message and does not run this command.

#### **Syntax**

revive enhancedFlashCopyGroup [enhancedFlashCopyGroupName]

### Parameters

Parameter	Description
enhancedF1ashCopyGroup	The alphanumeric identifier (including - and _) of the Enhanced FlashCopy group that you want to set to the <i>Optimal</i> state. Enclose the Enhanced FlashCopy group identifier in square brackets ([]).

# **Revive Enhanced FlashCopy Logical Drive**

This command forces a specified Enhanced FlashCopy logical drive to the *Optimal* state. The Enhanced FlashCopy logical drive can be one of these:

- A standalone Enhanced FlashCopy logical drive
- An Enhanced FlashCopy logical drive that is a member of a consistency group

If the Enhanced FlashCopy logical drive is not in a **Failed** state, the firmware displays an error message and does not run this command.

**Note:** You cannot use this command for an Enhanced FlashCopy logical drive that is used in online logical drive copy.

#### Syntax

revive enhancedFlashCopyLogicalDrive [enhancedFlashCopyLogicalDriveName]

### **Parameters**

Parameter	Description
enhancedFlashCopyLogicalDrive	The name of the Enhanced FlashCopy logical drive that you want to set to the <i>Optimal</i> state. Enclose the Enhanced FlashCopy logical drive name in square brackets ([]).

#### Notes

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

# **Revive Array**

This command forces the specified Array and its associated failed drives to the Optimal state.

**ATTENTION Possible loss of data access** – Correct use of this command depends on the data configuration on all of the drives in the Array. Never try to revive a drive unless you are supervised by your Technical Support representative.

#### **Syntax**

revive volumeGroup [volumeGroupName]

### Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier (including - and
	_) of the Array to be set to the Optimal state.
	Enclose the Array identifier in square
	brackets ([]).

#### **Minimum Firmware Level**

6.10

# Save Enhanced Global Mirror Group statistics

This command saves to a file the synchronization statistics for one or more member volumes in an Enhanced Global Mirror Group. The statistics collected are available only for Enhanced Global Mirror Group member volumes in a primary role on the local storage subsystem.

A set of data is collected during the synchronization process that can be used to evaluate how well the mirror configuration is functioning. The data is collected as a set of samples. A sample is created at the beginning of a synchronization process and updated regularly while the synchronization process proceeds.

A sample collects data until the synchronization process completes or until a disruption to the synchronization process occurs such as a volume ownership transfer or a read-write error. When a synchronization process disruption is resolved (for example, the volume is transferred to the alternate controller), a new sample is created and updated as the synchronization process continues.

### **Syntax**

```
save storageArray arvmStats enhancedGlobalMirrorGroup
["enhancedGlobalMirrorGroupName"]
arvmStats file="filename"
[volume="volumeName"
sampleType=(all | mostRecent | longestSyncTime |
errors) △
recordLimit=(1-90)]
```

Parameter	Description
enhancedG1oba1MirrorGroup	The name of the Enhanced Global Mirror Group for which you are retrieving synchronization statistics. Enclose the enhanced global mirror group name in square brackets ([]). If the Enhanced Global Mirror Group name has special characters, you also must enclose the enhanced global mirror group name in double quotation marks ("") inside square brackets. If the Enhanced Global Mirror Group name consists only of numbers, such as "1002," you also must enclose the Enhanced Global Mirror Group name in double quotation marks ("") inside square brackets.

Parameter	Description
file	The file path and the file name to which you want to save the synchronization statistics. Enclose the file name in double quotation marks (" "). For example:
	file="C:\Program Files\CLI\logs\ iscsistat.csv"
	This command does not automatically append a file extension to the saved file. You can use any file name, but you must use the .csv extension.
volume	This parameter is optional. The name of the specific member volume in the Enhanced Global Mirror Group for which you are retrieving synchronization statistics. If no volume is specified, the statistics for every member volume in the Enhanced Global Mirror Group are saved.
	If statistics are collected for more than one volume in an Enhanced Global Mirror Group, all the data will be written to the same file.
	Enclose the volume name in double quotation marks (" ").
sampleType	This parameter is optional. The default value for sampleType is all.
	• all— Data for all three sample types are collected and written to the same file.
	<ul> <li>mostRecent — Statistics are recorded for the most recent 50 resynchronization samples.</li> </ul>
	<ul> <li>longestSyncTime — Statistics are collected for the most recent 20 longest resynchronization samples.</li> </ul>
	• errors— Statistics are recorded for the most recent 20 failed resynchronization samples. These samples include a failure code.
recordLimit	This parameter is optional. The default value for record limit is no limit. The recordLimit must be greater than 0 and less than or equal to 90.

Statistics are captured for mirrored volumes in the Primary role. The statistics collected include the following data:

- Synchronization start time
- Synchronization type (manual or periodic)
- Synchronization duration
- Number of bytes sent

- Maximum write time (for a single write)
- Minimum write time (for a single write)
- Minimum synchronization data rate
- · Maximum synchronization data rate
- Total write time
- Repository utilization (%)
- Recovery point age

During initial synchronization, statistics samples are captured approximately every 15 minutes.

The synchronization statistics are included in the support bundle.

#### **Minimum Firmware Level**

7.84

## Save Controller NVSRAM

This command saves a copy of the controller NVSRAM values to a file. This command saves all of the regions.

#### **Syntax**

save controller [(a | b)]
NVSRAM file="filename"

#### **Parameters**

Parameter	Description
controller	The controller with the NVSRAM values that you want to save. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]).
file	The file path and the file name to which you want to save the NVSRAM values. Enclose the NVSRAM file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\nvsramb.txt" This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

# Save Disk Drive Channel Fault Isolation Diagnostic Status

This command saves the drive channel fault isolation diagnostic data that is returned from the **start diskDriveChannel faultDiagnostics** command. You can save the diagnostic data to a file as standard text or as XML.

See the "Start Drive Channel Fault Isolation Diagnostics" on page 341 command for more information.

### Syntax

save DriveChannel [drive-channel-Identifier] faultDiagnostics file="filename"

### Parameter

Parameter	Description
file	The file path and the file name to which you want to save the results of the fault isolation diagnostics test on the disk drive channel. Enclose the file name in double quotation marks (""). For example: file="C:\Program Files\CLI\sup\fltdiag.txt" This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

#### Notes

A file extension is not automatically appended to the saved file. You must specify the applicable format suffix for the file. If you specify a file extension of .txt, then the output will be in a text file format. If you specify a file extension of .xml, then the output will be in an XML file format.

# Save Drive Log

This command saves the log sense data to a file. Log sense data is maintained by the storage subsystem for each disk drive.

#### **Syntax**

save allDrives logFile="filename"

### Parameter

Parameter	Description
logFile	The file path and the file name to which you want to save the log sense data. Enclose the file name in double quotation marks (" "). For example:
	file="C:\Program Files\CLI\logs\lgsendat.txt"
	This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

# Save Enclosure Log

This command saves the log sense data to a file. Log sense data is maintained by the environmental cards for each enclosure. Not all of the environmental cards contain log sense data.

#### **Syntax**

save allEnclosures logFile="filename"

## Parameter

Parameter	Description
logFile	The file path and the file name to which you want to save the log sense data. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\enclosurelogdat.txt" This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name
	when entering the file name.

# Save Storage Subsystem Configuration

This command creates a script file that you can use to create the current storage subsystem logical drive configuration.

### **Syntax**

```
save storageSubsystem configuration file="filename"
[(allconfig | globalSettings=(TRUE | FALSE)
logicalDriveConfigAndSettings=(TRUE | FALSE)
hostTopology=(TRUE | FALSE)
lunMappings=(TRUE | FALSE))]
```

Parameter	Description
file	The name of the file that contains the configuration values. Enclose the file name in double quotation marks ("").
allConfig	The setting to save all of the configuration values to the file. (If you choose this parameter, all of the configuration parameters are set to TRUE.)
globalSettings	The setting to save the global settings to the file. To save the global settings, set this parameter to TRUE. To prevent saving the global settings, set this parameter to FALSE. The default value is TRUE.
logicalDriveConfigAndSettings	The setting to save the logical drive configuration settings and all of the global settings to the file. To save the logical drive configuration settings and global settings, set this parameter to TRUE. To prevent saving the logical drive configuration settings and global settings, set this parameter to FALSE. The default value is TRUE.
hostTopology	The setting to save the host topology to the file. To save the host topology, set this parameter to TRUE. To prevent saving the host topology, set this parameter to FALSE. The default value is FALSE.
lunMappings	The setting to save the LUN mapping to the file. To save the LUN mapping, set this parameter to TRUE. To prevent saving the LUN mapping, set this parameter to FALSE. The default value is FALSE.

When you use this command, you can specify any combination of the parameters for the global setting, the logical drive configuration setting, the host topology, or the LUN mapping. If you want to enter all settings, use the **allConfig** parameter. The parameters are all optional.

### Save Storage Subsystem DBM Database

This command saves the current state of the storage subsystem's Database Management (DBM) database into a local file. The output file that is produced can be used as the input file for the **save storageSubsystemdbmValidator** and the **load storageSubsystem dbmDatabase** commands.

#### **Syntax**

save storageSubsystem dbmDatabase file="filename"

#### Parameter

Parameter	Description
file	The file path and the file name of the DBM database you want to save. Enclose the file name in double quotation marks (" "). For example: file="C:\Subsystem Backups\DBMbackup_03302010.dbm"
	This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

# Save Storage Subsystem DBM Validator Information File

This command saves a storage subsystem's Database Management (DBM) validation information in an XML file, which can be used by an IBM Technical Support representative to generate a security code or Validator. The Validator must be included in the **load storageSubsystem dbmDatabase** command when restoring a storage subsystem back to a pre-existing configuration.

#### Syntax

save storageSubsystem dbmValidatorInfo file="filename" dbmDatabase="filename"

Parameter	Description
file	The file path and the file name of the DBM Validator required for IBM Technical Support representative. Enclose the file name in double quotation marks (" "). For example: file="C:\Subsystem Backups \DBMvalidator.xml" This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

Parameter	Description
dmbDatabase	The file path and the file name of the DBM database you want to use to restore a storage subsystem. Enclose the file name in double quotation marks (" "). For example: dbmDatabase="C:\Subsystem Backups \DBMbackup_03302010.dbm" This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

If you are generating the XML information file in order to obtain a validator string, you must run this command while you are connected to the controller where you intend to restore the database. The following example show the format of the XML file:

```
<?xml version="1.0" encoding="utf-8"?>
<DbmImageMetadata>
<Controllers>
<A>1ITOnnnnnnnABCD</A>
<B>1TOnnnnnnnnABCD</B>
</Controllers>
<Header>
<ImageVersion>1</ImageVersion>
<TimeStamp>1269388283</TimeStamp>
</Header>
<Trailer>
<CRC>nnnnnnnnn</CRC><
/Trailer>
</DbmImageMetadata>
```

# Save Storage Subsystem Diagnostic Data

This command saves the storage subsystem diagnostic data from either the controllers or the environmental services monitors (ESMs) to a file. You can review the file contents at a later time. You can also send the file to an IBM Technical Support representative for further review.

After you saving the diagnostic data, you can reset the NVSRAM registers that contain the diagnostic data so that the old data can be overwritten. Use the **reset storageSubsystem diagnosticData** command to reset the diagnostic data registers.

**Important:** Run this command only with the assistance of an IBM Technical Support representative.

#### **Syntax**

```
save storageSubsystem diagnosticData
[(controller | esm)]
file="filename"
```

### **Parameters**

Parameter	Description
diagnosticData	This parameter allows you to downloads the diagnostic data from either the controllers or the ESMs.
file	The name of the file to which you want to save the diagnostic data. Enclose the file name in double quotation marks ("").

# Save Storage Subsystem Diagnostic Data Capture

This command saves a diagnostic data capture to a specified file on a host.

### **Syntax**

save storageSubsystem ddc file="filename"

### **Parameters**

Parameter	Description
file	The file path and the file name to which you save the diagnostic data capture. Enclose the file name in double quotation marks (" "). For example: file="C:\subsystem
	Backups\DBMbackup_03302010.core"

# Save Storage Subsystem Events

This command saves events from the Major Event Log to a file. You can save all of the events or only the critical events.

### **Syntax**

```
save storageSubsystem (allEvents | criticalEvents)
file="filename"
[count=numberOfEvents]
```

### Parameters

Parameter	Description
file	The name of the file to which you want to save the events. Enclose the file name in double quotation marks ("").
count	The number of events or critical events that you want to save to a file. If you do not enter a value for the count, all events or all critical events are saved to the file. If you enter a value for the count, only that number of events or critical events (starting with the last event entered) are saved to the file. Use integer values.

### Notes

You have the option to save all events (allEvents) or only the critical events (criticalEvents).

## Save Storage Subsystem Firmware Inventory

This command saves a report to a file of all of the firmware currently running on the storage subsystem. The report lists the firmware for these components:

- Controllers
- Drives
- Drawers (if applicable)
- Environmental services monitors (ESMs)

You can use the information to help identify out-of-date firmware or firmware that does not match the other firmware in your storage subsystem. You can also send the report to your IBM Technical Support representative for further review.

#### **Syntax**

save storageSubsystem firmwareInventory file="filename"

#### Parameter

Parameter	Description
file	The file path and the file name to which you want to save the firmware inventory. Enclose the file name in double quotation marks (" "). For example: file="C:\Subsystem Backups\DBMbackup_03302010.dbm"C:\Program Files\CLI\logs\fwinvent.txt"
	This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

# Save storage subsystem infiniBand statistics

This command saves the InfiniBand performance statistics of the storage subsystem to a file.

#### **Syntax**

save storageArray ibStats [raw | baseline]
file="filename"

Parameter	Description
rāw	The statistics that are collected are all statistics from the controller start-of-day. Enclose the parameter in square brackets ([ ]).
baseline	The statistics that are collected are all statistics from the time the controllers were reset to zero using the reset storageArray ibStatsBaseline command. Enclose the parameter in square brackets ([]).

Parameter	Description
file	The file path and the file name to which you want to save the performance statistics. Enclose the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\sup\ ibctat txt"
	This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.

If you have not reset the InfiniBand baseline statistics since the controller start-of-day, the time at the start-of-day is the default baseline time.

### **Minimum Firmware Level**

7.32

# Save Storage Subsystem iSCSI Statistics

This command saves the iSCSI performance statistics of the storage subsystem to a file.

#### **Syntax**

```
save storageSubsystem iscsiStatistics [raw | baseline]
file="filename"
```

#### **Parameters**

Parameter	Description
raw	The statistics collected are all statistics from the controller start-of-day. Enclose the parameter in square brackets ([ ]).
baseline	The statistics that are collected are all statistics from the time the controllers were reset to zero using the <b>reset storageSubsystem iscsiStatsBaseline</b> command. Enclose the parameter in square brackets ([]).
file	The name of the file to which you want to save the performance statistics. Enclose the file name in double quotation marks ("").

### Notes

If you have not reset the iSCSI baseline statistics since the controller start–of–day, the time at the start–of–day is the default baseline time.

# Save Storage Subsystem Performance Statistics

This command saves the performance statistics to a file. Before you use this command, run the **set session performanceMonitorInterval** command and the **set session performanceMonitorIterations** command to specify how often statistics are collected.

### **Syntax**

save storageSubsystem performanceStats file="filename"

#### **Parameter**

Parameter	Description
file	The name of the file to which you want to save the performance statistics. Enclose the file name in double quotation marks (""). For example:
	<pre>save storageSubsystem performanceStats file="C:\Program Files\CLI\reports\stats.csv"</pre>
	This command does not automatically append a file extension to the saved file. You can use any file name, but you must specify the <b>.csv</b> extension.

# Save Storage Subsystem RLS Counts

This command saves the read link status (RLS) counters to a file.

#### Syntax

save storageSubsystem RLSCounts file="filename"

#### Parameter

Parameter
file

### Notes

To more effectively save RLS counters to a file, perform these steps:

- 1. Run the **reset storageSubsystem RLSBaseline** command to set all of the RLS counters to 0.
- 2. Run the storage subsystem for a predetermined amount of time (for instance, two hours).
- 3. Run the save storageSubsystem RLSCounts file="filename" command.

# Save Storage Subsystem SAS PHY Counts

This command saves the SAS physical layer (SAS PHY) counters to a file. To reset the SAS PHY counters, run the **reset storageSubsystem SASPHYBaseline** command.

#### Syntax

save storageSubsystem SASPHYCounts file="filename"

### Parameter

Parameter	Description
file	The file path and the file name to which you want to save the SYSPHY counters. Enclose the file path and the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\logs\sasphy.csv"
	This command does not automatically append a file extension to the saved file. You can use any file name but you must use the <i>.csv</i> extension.

# Save Storage Subsystem SOC Counts

This command saves the SOC error statistics to a file. This command is valid only for Fibre Channel devices in an Arbitrated Loop topology.

### **Syntax**

save storageSubsystem SOCCounts file="filename"

#### Parameter

Parameter	Description
file	The name of the file to which you want to save the SOC error statistics. Enclose the file name in double quotation marks ("").

#### **Notes**

To more effectively save SOC error statistics to a file, perform these steps:

- 1. Run the **reset storageSubsystem SOCBaseline** command to set all of the SOC counters to 0.
- 2. Run the storage subsystem for a predetermined amount of time (for example, two hours).
- 3. Run the save storageSubsystem SOCCounts file="filename" command.

# Save Storage Subsystem State Capture

This command saves the state capture of a storage subsystem to a file.

#### **Syntax**

save storageSubsystem stateCapture file="filename"

### Parameter

Parameter	Description
file	The name of the file to which you want to save the state capture. Enclose the file name in double quotation marks (""). For example:
	file="C:\Program Files\CLI\logs\state.zip"
	This command does not automatically append a file extension to the saved file. You must specify a file extension when entering the file name.
	CAUTION: Only do this when Technical Support requests it.

# Save Storage Subsystem Support Data

This command saves the support-related information of the storage subsystem to a file. Support-related information includes these items:

- The storage subsystem profile
- The Major Event Log information
- The read ink status (RLS) data
- The NVSRAM data
- Current problems and associated recovery information
- The performance statistics for the entire storage subsystem
- · The persistent registration information and the persistent reservation information
- Detailed information about the current status of the storage subsystem
- The diagnostic data for the disk drive
- A recovery profile for the storage subsystem
- · The unreadable sectors that are detected on the storage subsystem
- The state capture data
- An inventory of the versions of the firmware running on the controllers, the drives, the drawers, and the environmental services modules (ESMs)

### Syntax

```
save storageSubsystem supportData file="
filename" [force=(TRUE | FALSE)]
```

Parameter	Description
file	The name of the file to which you want to save the support-related data for the storage subsystem. Enclose the file name in double quotation marks ("").
force	This parameter forces the collection of support data if there are any failures in securing a lock on controller operations. To force the collection of support data, set this parameter to TRUE. The default value is FALSE.

# Schedule Support Bundle Collection Configuration

This command specifies a schedule for collecting support bundle data.

### Syntax

**Attention:** The SM CLI support bundle commands must be run from the SM CLI command line interface only. They cannot be run from the storage subsystem SM CLI script window. In addition, the storage subsystem IP addresses are not required. Also, the **-c** parameter is not required to precede the SM CLI support bundle command.

```
<installed-SM-Directory>SMcli -supportBundle schedule (enable|disable)
(all|storageSubsystemName)
[data=pathName]
[startTime=HH:MM]
[startDate=MM:DD:YY]
[endDate=MM:DD:YY]
(daysOfWeek={} | months={}
(onDays={} |
weekNumber=(First|Second|Third|Fourth|Lasdta)y
OfWeek=(Sunday|Monday|Tuesday|Wednesday|Thursday|Friday|Saturday))
```

### **Parameters**

Parameter	Description
enable	Enable automatic collection of support bundles when a critical MEL event is detected.
disable	Disable automatic collection of support bundles where the feature was previously enabled.
a11	Apply the command to all storage subsystems detected from the host. The possible names can be displayed by using the SMcli command with the -d parameter
storageSubsystemName	Apply the command to the named storage subsystem.
data	Store the data generated by the command at the location specified by the pathname.

**Note:** The defined schedule is stored in the management station that the Storage Manager host software has been installed. This is where the schedule is initially defined and enabled for the managed storage subsystems. For example, you have installed the Storage Manager host software on two different management stations, and defined the schedule in only one subsystem. The schedule show command in the management station where the schedule is not defined will display a no schedule found message. In addition, if multiple management stations have been enabled for the schedule to automatically collect the support data from a single storage subsystem, this subsystem will have support data collected multiple times, depending on the schedule defined within the management stations. IBM recommends that you define support data collection in only one workstation, if there are multiple management stations that are available to manage a given storage subsystem.

An example of the SMcli command to disable automatic support bundle in a Microsoft Windows management station, with default Storage Manager software installed path, is as follows:

```
c:\Program Files (x86)\IBM_DS\client>SMcli -supportBundle auto disable all
SMcli completed successfully.
```

Another example of the SMcli command that displays the automatically support bundle collection schedule in a Microsoft Windows management station is as follows:

```
c:\Program Files (x86)\IBM_DS\client>SMcli -supportBundle schedule show
subsystem Name Existing Schedule Starting Ending
DS3524-DTL139140 Daily / Weekly - Mon - Fri 05/23/2012 02:00AM 05/30/2012
DS3524-DTL Daily / Weekly - Mon - Fri 05/23/2012 02:00AM 05/30/2012
DS3400A None
DS5100-DTL None
```

SMcli completed successfully.

**Note:** In Microsoft Windows OS, the command prompt window must have administrator authority or the command will fail with an error message as shown below:

C:\Program Files (x86)\IBM DS\client>SMcli -supportBundle auto disable all

Incorrect Usage

. . . . . . . . . .

Do not modify any parameters to the command. Some examples are as follows:

Incorrectly modified parameters	Advisory
SMcli 10.12.128.1 10.12.128.2 -supportBundle schedule show	Do not specify IP addresses
SMcli 10.12.128.1 10.12.128.2 -c "-supportBundle schedule show"	Do not specify IP addresses and <b>-c</b> parameter
SMcli -supportbundle schedule show	Note that <b>supportBundle</b> parameter is case sensitive and must be typed accordingly.

### Set subsystem

This command defines the properties for an subsystem.

#### Syntax

```
set subsystem [subsystemNumber]
addDrives=(enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn)
raidLevel=(0 | 1 | 3 | 5 | 6)
owner=(a | b)
```

Parameter	Description
subsystem	The alphanumeric identifier (including - and _) of the subsystem for which you want to set properties. Enclose the subsystem identifier in square brackets ([]).

Parameter	Description
addDrives	The location of the disk drive that you want to add to the subsystem. For high-capacity disk drive expansion enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for the disk drive. For low-capacity drive expansion enclosures, specify the enclosure ID value and the slot ID value for the disk drive. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID value, the drawer ID value, and the slot ID value in parentheses.
raidLevel	The RAID level for the subsystem. Valid values are $0, 1, 3, 5$ , or $6$ .
owner	The controller that owns the subsystem. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Use this parameter only if you want to change the subsystem owner.

Host I/O errors might result in subsystems with more than 32 logical drives. This operation also might result in internal controller reboots because the timeout period ends before the subsystem definition is set. If you experience this issue, quiesce the host I/O operations, and try the command again.

When you use this command, you can specify one or more of the parameters.

**Important:** Specifying the **addDrives** parameter or the **raidLevel** parameter starts a long-running operation that you cannot stop.

The **addDrives** parameter supports both high-capacity and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need to specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the slot in which a drive resides.

## Troubleshooting

An attempt to expand large subsystems by adding drives, also called Dynamic Capacity Expansion (DCE), might fail with the following message:

Return code: Error 26 - The modification operation cannot complete because of the number of drives in the subsystem and the segment size of the associated logical drives. Reduce the segment size of all logical drives in the subsystem to 128 KB or below using the Change Segment Size option. Then, retry the operation.

Systems running 7.35.xx.xx firmware might fail with the following message instead of the message given earlier:

Return code: Error 462 - A SYMbol procedure could not be carried out because the firmware could not allocate sufficient cache memory. Operation when error occurred: PROC\_startVolum

In addition to the above messages, a Major Event Log (MEL) event that indicates unavailability of sufficient cache to complete the DCE operation might occur.

Any controller running 7.xx firmware can encounter this problem. DCE requires enough memory to buffer the data read from the original logical drive and the data to be written to the expanded logical drive. Some combination of number of drives in the expansion operation, stripe size, and whether mirroring is enabled may result in not enough memory being available to complete the DCE operation.

If such a situation occurs, possible workarounds are as follows:

- Create the desired size subsystem using other unassigned drives.
- Delete the current subsystem and then re-create the subsystem with the desired number of drives.
- Reduce the segment size being used and then retry the operation.
- If possible, add additional memory to the controller and then retry the operation.

# Set subsystem Forced State

This command moves an subsystem into a Forced state. Use this command if the **start subsystem import** command does not move the subsystem to an Imported state or if the import operation does not work because of hardware errors. In a Forced state, the subsystem can be imported, and you can then identify the hardware errors.

#### Syntax

set subsystem [subsystemName] forcedState

#### Parameter

Parameter	Description
subsystem	The alphanumeric identifier (including - and
	_) of the subsystem that you want to place
	in a Forced state. Enclose the subsystem
	identifier in square brackets ([ ]).

#### Notes

You can move the disk drives that comprise an subsystem from one storage subsystem to another storage subsystem. The CLI provides three commands that let you move the disk drives. The commands are **start subsystem export**, **start subsystem import**, and **set subsystem forcedState**.

In the Forced state, you can perform an import operation on the subsystem.

# Set Enhanced Global Mirror Group

Use this command to change the synchronization settings and warning thresholds that the controller owner of the primary side of the Enhanced Global Mirror Group uses when it performs an initial synchronization or resynchronization. Changing the synchronization settings affects the synchronization operations of all mirrored pairs within the Enhanced Global Mirror Group.

### Syntax

```
set enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"]
[syncInterval=integer (minutes | hours | days)
warningSyncThreshold=integer (minutes | hours | days)
warningRecoveryThreshold=integer (minutes | hours |
days)
warningThresholdPercent=percentValue
userLabel="New_enhancedGlobalMirrorGroupName"
autoResync=(TRUE | FALSE)
volume="repos_xxxx" increaseRepositoryCapacity
(repositoryVolumes=("repos_xxxx") |
repositoryVolumes=("repos_xxxx" "repos_xxxx"))
role=(primary | secondary)
(force=TRUE | FALSE | noSync=TRUE | FALSE)
```

Parameter	Description
enhancedGlobalMirrorGroup	The name of the Enhanced Global Mirror Group that you want to modify. Enclose the Enhanced Global Mirror Group name in square brackets ([]). If the Enhanced Global Mirror Group name has special characters, you also must enclose the name in double quotation marks (" ") inside square brackets. If the Enhanced Global Mirror Group name consists only of numbers, such as "1002," you also must enclose the name in double quotation marks (" ") inside square brackets.
syncInterval	Specify the length of time between automatically sending updates of modified data from the local storage subsystem to the remote storage subsystem. You can specify the length of time in minutes, hours, or days.
warningSyncThreshold	Specify the length of time to wait until a warning is triggered for cases in which the synchronization of all of the volumes within the Enhanced Global Mirror Group takes longer than the defined time. You can specify the length of time in minutes, hours, or days.

Parameter	Description
warningRecoveryThreshold	Specify the length of time to wait until a warning is triggered when the automatic data update for the point-in-time image on the remote storage subsystem is older than the defined time. Define the threshold from the end of the previous update. You can specify the length of time in minutes, hours, or days. <b>NOTE</b> You must set the Recovery Point Threshold to be twice as large as the synchronization interval threshold.
warningThresholdPercent	Specify the percent value that determines when a warning is triggered when the capacity of a mirror repository volume reaches the defined percentage. Define the threshold by percentage (%) of the capacity remaining.
userLabel	The new name that you want to give to the Enhanced Global Mirror Group. Use this parameter when you want to rename the Enhanced Global Mirror Group. Enclose the new Enhanced Global Mirror Group name in double quotation marks (" ").
autoResync	<ul> <li>The settings for automatic resynchronization between the primary volumes and the secondary volumes of an Enhanced Global Mirrored pair within an Enhanced Global Mirror Group. This parameter has these values:</li> <li>enabled– Automatic resynchronization is turned on. You do not need to do anything further to resynchronize the primary volume and the secondary volume.</li> <li>disabled– Automatic resynchronization is turned off. To resynchronize the primary volume and the secondary volume, you must run the resume enhancedGlobalMirrorGroup command.</li> </ul>
volume	<ul> <li>The name of an Enhanced Global Mirror Group repository volume for which you want to increase the capacity.</li> <li>The name of a repository volume is comprised of two parts: <ul> <li>The term repos</li> <li>A four digit numerical identifier that the storage management software assigns to the repository volume name</li> </ul> </li> <li>Enclose the name of the repository volume in double quotation marks (" ").</li> </ul>

Parameter	Description
repositoryVolume or repositoryVolumes	The name of an unused repository volume that you want to use to increase the capacity of another repository volume. An available standard volume is added to the repository volume to increase the capacity of the repository volume.
	The name of the unused repository volume is comprised of two parts:
	<ul> <li>The term repos</li> <li>A four digit numerical identifier that the storage management software assigns to the repository volume name</li> </ul>
	Enclose the name of the existing repository volume in double quotation marks (" ").
	You can enter more than one repository volume name. Enclose each repository volume name in double quotation marks (" ") inside parenthesizes (()). Separate each repository volume name with a white space.
role	Use this parameter to promote the Enhanced Global Mirror Group to a primary role or demote the Enhanced Global Mirror Group to a secondary role. To define the Enhanced Global Mirror Group as the primary role, set this parameter to primary. To define the Enhanced Global Mirror Group as the secondary role, set this parameter to secondary.
force	The role reversal is forced if the communications link between the storage subsystems is down and promotion or demotion on the local side results in a dual-primary condition or a dual-secondary condition. To force a role reversal, set this parameter to TRUE. The default value is FALSE.
nosync	This parameter defines whether to perform an initial synchronization between the primary storage subsystem and the secondary storage subsystem before a role reversal operation is performed. To perform an initial synchronization, set this parameter to TRUE. The default value is FALSE.

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

When you use this command, you can specify one or more of the parameters. You do not, however, need to use all of the parameters.

An Enhanced Global Mirror Group repository volume is an expandable volume that is structured as a concatenated collection of up to 16 standard volume entities. Initially, an expandable repository volume has only a single volume. The capacity of the expandable repository volume is exactly that of the single volume. You can increase the capacity of an expandable repository volume by attaching additional unused repository volumes to it. The composite expandable repository volume capacity then becomes the sum of the capacities of all of the concatenated standard volumes.

To increase the capacity of a repository volume use the set enhancedGlobalMirrorGroup command as shown in this example: set enhancedGlobalMirrorGroup ["amg\_001"] volume="repos\_006" increaseRepositoryCapacity repositoryVolumes=("repos\_0021" "repos\_0021");

An Enhanced Global Mirror Group repository volume must satisfy a minimum capacity requirement that is the sum of the following:

- 32 MB to support fixed overhead for the Enhanced Global Mirror Group and for copy-on-write processing. [comment: should the text in these two bullets be in blue?]
- Capacity for rollback processing, which is 1/5000th of the capacity of the base volume.

The minimum capacity is enforced by the controller firmware and the storage management software.

#### **Minimum Firmware Level**

7.84

## Set Consistency Group Attributes

This command defines the properties for an Enhanced FlashCopy consistency group.

#### Syntax

```
set consistencyGroup ["consistencyGroupName"]
[userLabel="consistencyGroupName" |
repositoryFullPolicy=(failBaseWrites | purgeEnhancedFlashCopyImages)|
repositoryFullLimit=percentValue|
autoDeleteLimit=numberOfEnhancedFlashCopyImages|
rollbackPriority=(lowest | low | medium | high | highest)]
```

Parameter	Description
consistencyGroupName	The name of the consistency group for which you are setting properties. Enclose the consistency group name in double quotation marks (" ") inside square brackets ([ ]).
userLabel	The new name that you want to give to the Enhanced FlashCopy consistency group. Enclose the new consistency group name in double quotation marks (" ").

Parameter	Description
repositoryFullPolicy	How you want Enhanced FlashCopy processing to continue if the Enhanced FlashCopy repository logical drives are full. You can choose to fail writing to the base logical drive (failBaseWrites) or delete (purge) the Enhanced FlashCopy images (purgeEnhancedFlashCopyImages). The default action is purgeEnhancedFlashCopyImages.
repositoryFullLimit	The percentage of repository capacity at which you receive a warning that the Enhanced FlashCopy repository logical drive is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.
autoDeleteLimit	The maximum number of Enhanced FlashCopy images that you want to automatically delete if you have selected to purge the Enhanced FlashCopy images for a repository full policy. The default value is 32.
rollBackPriority	The priority for rollback operations for a consistency group while the storage subsystem is operational. Valid values are highest, high, medium, low, or lowest.

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

When you use this command, you can specify one or more of the parameters. You do not, however, need to use all of the parameters.

The rollback priority defines the amount of system resources that should be allocated to the rollback operation at the expense of system performance. A value of high indicates that the rollback operation is prioritized over all other host I/O. A value of low indicates that the rollback operation should be performed with minimal impact to host I/O.

# Set Consistency Group Enhanced FlashCopy Logical Drive

This command creates a unique name for an Enhanced FlashCopy logical drive of a consistency group.

#### Syntax

set cgEnhancedFlashCopyLogicalDrive [consistencyGroupName]
userLabel="enhancedFlashCopylogicalDriveName"

### Parameters

Parameter	Description
cgEnhancedFlashCopyLogicalDrive	The name of the consistency group logical drive that you want to rename. Enclose the name of the consistency group Enhanced FlashCopy logical drive in square brackets ([ ]).
userLabel	The new name that you want to give to the Enhanced FlashCopy logical drive in the consistency group. Enclose the new Enhanced FlashCopy logical drive name in double quotation marks (" ").

### Notes

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

# Set Controller

This command defines the attributes for the controllers.

#### Syntax

```
set controller [(a | b)]
availability=(online | offline | serviceMode) |
ethernetPort [(1 | 2)] ethernetPortOptions |
globalNVSRAMByte [nvsramOffset]=(nvsramByteSetting | nvsramBitSetting) |
hostNVSRAMByte [hostType, nvsramOffset]=(nvsramByteSetting | nvsramBitSetting) |
IPv4GatewayIP=ipAddress | IPv6RouterAddress=ipv6Address |
iscsiHostPort [(1 | 2 | 3 | 4)] iscsiHostPortOptions |
rloginEnabled=(TRUE | FALSE) |
serviceAllowedIndicator=(on | off
```

Parameter	Description
controller	The controller for which you want to define properties. Valid identifiers for the controller are a or b, where a is the controller in slot A and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the firmware for the controller returns a syntax error.
availability	The mode for the controller. The controller modes are online, offline, or serviceMode.
ethernetPort	The attributes (options) for the management Ethernet ports. The entries to support this parameter are listed in the "Syntax Element Statement Data" table that follows. Many settings are possible, including setting the IP address, the gateway address, and the subnet mask address.
globalNVSRAMByte	A portion of the controller NVSRAM. Specify the region that you want to modify by using the starting byte offset within the region, and the size and value of the new data to be stored into NVSRAM.

Parameter	Description
hostNVSRAMByte	The NVSRAM for the host-specific region. The setting specifies the host index for the specific host, the starting offset within the region, the number of bytes, and the values to be written.
hostNVSRAMByte	The NVSRAM for the host-specific region. The setting specifies the host index for the specific host, the starting offset within the region, the number of bytes, and the the byte value or bit value of the new data to be stored into the NVSRAM.
IPv4GatewayIP	The IP address of the node that provides the interface to the network. The address format for the IPv4 gateway is (0–255).(0–255).(0–255).(0–255)
IPv6RouterAddress	The IP address of IPv6 router that connects two or more logical subnets. The address format for the IPv6 router is (0–FFFF):(0–FFFF):(0–FFFF):(0–FFFF): (0–FFFF):(0–FFFF):(0–FFFF).
iscsiHostPort	The entries to support this parameter are listed in the "Syntax Element Statement Data" table that follows. Many settings are possible, including setting the IP address, the gateway address, the subnet mask address, the IPv4 priority, and the IPv6 priority.
rloginEnabled	The setting for whether the remote login feature is turned on or turned off. To turn on the remote login feature, set this parameter to TRUE. To turn off the remote login feature, set this parameter to FALSE.
serviceAllowedIndicator	The setting for whether the Service Action Allowed indicator light is turned on or turned off. To turn on the Service Action Allowed indicator light, set this parameter to on. To turn off the Service Action Allowed indicator light, set this parameter to off.

## Syntax Element Statement Data

```
The following options are available for the ethernetPort parameter:
enableIPv4=(TRUE | FALSE) |
enableIPv6=(TRUE | FALSE) |
IPv6LocalAddress=(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):
(0-FFFF):(0-FFFF):(0-FFFF)|
IPv6RoutableAddress=(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):
(0-FFFF):(0-FFFF):(0-FFFF) |
IPv6RouterAddress=(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):
(0-FFFF):(0-FFFF):(0-FFFF) |
IPv4Address=(0-255).(0-255).(0-255) |
IPv4ConfigurationMethod=[(static | dhcp)] |
IPv4GatewayIP=(0-255).(0-255).(0-255).(0-255) |
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255) |
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255).(0-255) |
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255).(0-255) |
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255).(0-255) |
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255).(0-255) |
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255).(0-255).(0-255) |
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255).(0-255).(0-255) |
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255).(0-255) |
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255).(0-255).(0-255) |
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-255).(0-
```

The following options are available for the **iscsiHostPort** parameter: IPv4Address=(0-255).(0-255).(0-255).(0-255).

```
IPv6LocalAddress=(0-FFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):
(0-FFFF):(0-FFFF):(0-FFFF)
IPv6RoutableAddress=(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):
(0-FFFF):(0-FFFF):(0-FFFF)
IPv6RouterAddress=(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):(0-FFFF):
(0-FFFF):(0-FFFF):(0-FFFF) |
enableIPv4=(TRUE | FALSE) |
enableIPv6=(TRUE | FALSE) |
enableIPv4Vlan=(TRUE | FALSE)
enableIPv6Vlan=(TRUE | FALSE) |
enableIPv4Priority=(TRUE | FALSE) |
enableIPv6Priority=(TRUE | FALSE) |
IPv4ConfigurationMethod=(static | dhcp) |
IPv6ConfigurationMethod=(static | auto) |
IPv4GatewayIP=(TRUE | FALSE) |
IPv6HopLimit=[0-255]
IPv6NdDetectDuplicateAddress=[0-256]
IPv6NdReachableTime=[0-65535]
IPv6NdRetransmitTime=[0-65535]
IPv6NdTimeOut=[0-65535]
IPv4Priority=[0-7]
IPv6Priority=[0-7]
IPv4SubnetMask=(0-255).(0-255).(0-255).(0-255)
IPv4VlanID=[1-4094] |
IPv6VlanID=[1-4094] |
maxFramePayload=[framesize] |
tcpListeningPort=[3260, 49152-65536]
portSpeed=[(autoNegotiate | 1 | 10)]
```

Before firmware version 7.75, the set controller command supported an NVSRAMByte parameter. The NVSRAMByte parameter is deprecated and must be replaced with either the hostNVSRAMByte parameter or the globalNVSRAMByte parameter.

When you use this command, you can specify one or more of the parameters. You do not need to use all of the parameters.

Setting the **availability** parameter to serviceMode causes the alternate controller to take ownership of all of the virtual disks. The specified controller no longer has any logical drives and refuses to take ownership of any more logical drives. Service mode is persistent across reset cycles and power cycles until the **availability** parameter is set to online.

Use the **show controller NVSRAM** command to show the NVSRAM.

When the **duplexMode** option is set to TRUE, the selected Ethernet port is set to full duplex. The default value is half duplex (the **duplexMode** parameter is set to FALSE).

To make sure that the IPv4 settings or the IPv6 settings are applied, you must set these **iscsiHostPort** options:
- enableIPV4= TRUE
- enableIPV6= TRUE

The IPv6 address space is 128 bits. It is represented by eight 16-bit hexadecimal blocks separated by colons.

The **maxFramePayload** parameter is shared between IPv4 and IPv6. The payload portion of a standard Ethernet frame is set at 1500, and a jumbo Ethernet frame is set at 9000. When using jumbo frames, all of the devices that are in the network path should be capable of handling the larger frame size.

The portSpeed option is expressed as megabits per second (Mb/s).

Values for the portSpeed option of the iscsiHostPort parameter are in megabits per second (Mb/s).

The following values are the default values for the iscsiHostOptions:

- The IPv6HopLimit option is 64.
- The IPv6NdReachableTime option is 30000 milliseconds.
- The IPv6NdRetransmitTime option is 1000 milliseconds.
- The IPv6NdTimeOut option is 30000 milliseconds.
- The tcpListeningPort option is 3260.

# Set Controller Service Action Allowed Indicator

This command turns off or on the Service Action Allowed indicator light of a controller in a controller module. If the storage subsystem does not support the Service Action Allowed indicator light feature, this command returns an error. If the storage subsystem supports the command but is unable to turn on or turn off the indicator light, this command returns an error. (To turn on or turn off the Service Action Allowed indicator light on the power-fan CRU or the interconnect-battery CRU, use the **set enclosure serviceAllowedIndicator** command.)

### **Syntax**

set controller=[(a | b)] serviceAllowedIndicator=(on | off

Parameter	Description
controller	The controller that has the Service Action Allowed indicator light that you want to turn on or turn off. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the controller firmware returns a syntax error.
serviceAllowedIndicator	The setting to turn on or turn off the Service Action Allowed indicator light. To turn on the Service Action Allowed indicator light, set this parameter to on. To turn off the Service Action Allowed indicator light, set this parameter to off.

This command was originally defined for use with the DS4800 (models 82, 84, and 88) controller module. This command is not supported by controller modules that were shipped before the introduction of the DS4800 (models 82, 84, and 88) controller module. The DS3500 series, DS3950, DS4200, DS4700, DS5020, DS5100, and DS5300 controllers also support this command.

# Set Disk Pool

This command sets the attributes associated with a disk pool based on the specified parameters.

#### Syntax

```
set diskPool (diskPool=[
diskPoolName] | diskPools=[diskPoolName1... diskPoolNamen] | allDiskPools)
[reserveDriveCount=reserveDriveCountValue |
warningThreshold=(warningThresholdValue | default) |
criticalThreshold=(criticalThresholdValue | default) |
criticalPriority=(highest|high|medium|low|lowest) |
degradedPriority=(highest|high|medium|low|lowest)
backgroundPriority=(highest|high|medium|low|lowest) | userLabel=diskPoolName]
```

Parameter	Description
diskPool or diskPools	The name of the disk pool for which you are setting attributes. Enclose the disk pool name in square brackets ([]).
	You can enter more than one disk pool name if you use the <b>diskPools</b> parameter. Enclose all of the disk pool names in a single set of square brackets ([]). Separate each disk pool name with a space. If you have special characters in a disk pool name, you must also enclose that name in double quotation marks (" ").
allDiskPools	This parameter sets attributes for all of the disk pools in the storage subsystem
reserveDriveCount	This parameter sets the number of drives to be reserved in the disk pool for drive failures or drive removals. Setting this value to $\theta$ results in a possible loss of data if a drive failure occurs in a disk pool.
warningThreshold	The percentage of disk pool capacity at which you receive a warning alert that the disk pool is nearing full. Use integer values. For example, a value of 70 means 70 percent. For best operation, the value for this parameter must be less than the value for the <b>criticalThreshold</b> parameter. Setting this parameter to 0 disables warning alerts.

Parameter	Description
criticalThreshold	The percentage of disk pool capacity at which you receive a critical alert that the disk pool is nearing full. Use integer values. For example, a value of 70 means 70 percent. For best operation, the value for this parameter must be greater than the value for the <b>warningThreshold</b> parameter. Setting this parameter to 0 disables both warning alerts and critical alerts.
degradedPriority	The priority for reconstruction operations for degraded events on the disk pool. For example, disk pool reconstruction after at one drive failure. Valid values are highest, high, medium, low, and lowest. The default value is high.
backgroundPriority	The priority for background operations on the disk pool. Valid values are highest, high, medium, low, and lowest. The default value is low.
userLabel	The new alphanumeric identifier (including - and _) that you want to give the disk pool. Enclose the disk pool identifier in double quotation marks (" ").

Each disk pool name must be unique. You can use any combination of alphanumeric characters, underscore (\_), hyphen(-), and pound (#) for the user label. User labels can have a maximum of 30 characters.

You can specify an arbitrary set of disk pools. If you select multiple disk pools, setting a value for the **userLabel** parameter causes an error.

If you do not specify a value for an optional parameter, a default value is assigned.

# **Disk Pool Alert Thresholds**

Each disk pool has two progressively severe levels of alerts to inform users when the storage capacity of the disk pool is approaching full. The threshold for an alert is a percent of the used capacity to the total usable capacity in the disk pool. The alerts are:

- **Warning** This is the first level of alert that the used capacity in a disk pool is approaching full. When the threshold for the warning alert is reached a Needs Attention condition is generated and an event is posted to the storage management software. The warning threshold is superseded by the critical threshold. The default warning threshold is 50 percent.
- **Critical** This is the most severe level of alert that the used capacity in a disk pool is approaching full. When the threshold for the **critical** alert is reached a **Needs Attention** condition is generated, and an event is posted to the storage management software. The **warning** threshold is superseded by the **critical** threshold. The default **critical** warning threshold is 85 percent.

To be effective, the value for a warning alert must always be less than the value for a critical alert. If the value for the warning alert is the same as the value for a critical alert, only the critical alert is sent.

### **Disk Pool Background Operations**

Disk pools support these background operations:

- Reconstruction
- Instant Availability Format (IAF)
- Format
- Dynamic Capacity Expansion (DCE)
- Dynamic Volume Expansion (DVE) (For disk pools, DVE is actually not a background operation, but DVE is supported as a synchronous operation.)

Disk pools do not queue background commands. You can start several background commands sequentially, but starting more than one background operation at a time delays the completion of commands that you started previously. The relative priority levels for the supported background operations are:

- 1. Reconstruction
- 2. Format
- **3**. IAF
- 4. DCE

# Set Disk Pool (Modify Disk Pool)

This command adds capacity to a disk pool or changes the controller ownership for the entire disk pool. These two operations are mutually exclusive.

### **Syntax**

```
set diskPool [diskPoolName]
((addDrives=[enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn] |
addCapacity=(diskPoolCapacity)) | owner=(a | b))
```

Parameter	Description
diskPool	The name of the disk pool that you want to modify. Enclose the disk pool name in square brackets ([ ]).
addDrives	The drive s that you want to add to the disk pool. For high-capacity drive enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value of the drive that you want to add. For low-capacity drive enclosures, specify the enclosure ID value and the slot ID value of the drive that you want to add. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID value, the drawer ID value, and the slot ID value in square brackets ([]).

Parameter	Description
addCapacity	The amount of additional storage capacity that you want to add to the disk pool. This parameter automatically selects the drives to meet the capacity that you want to add. The capacity is defined in units of bytes, KB, MB, GB, or TB.
owner	The controller that owns the disk pool. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. If you do not specify an owner, the controller firmware determines the owner.

Logical Drives already in the disk pool remain on line and available for I/O operations while you add new drives. The disk pool must be in the Complete state before you add capacity.

To add capacity, specify individual drives with the **addDrives** parameter, or an amount of drive capacity with the **addCapacity** parameter. If you use **addDrives**, the host must validate the drive set before allowing the operation to run. If you use the **addCapacity** parameter, the capacity you specify is taken as the minimum capacity to be added. The candidate drives with the best match for quality of service and a capacity greater than or equal to what you specified are used. If no candidate is available with a minimum match, or the drive list specified is not available or attribute mismatches are detected, the operation fails.

You can also use this command to change ownership of a disk pool from one controller in the storage subsystem to the other. Using this command to change ownership is mutually exclusive with using the command to add drives or to add capacity.

# Set Drawer Service Action Allowed Indicator

This command turns on or turns off the Service Action Allowed indicator light on a drawer that holds disk drives. Drawers are used in high-capacity expansion drawers. The drawers slide out of the expansion drawer to provide access to the disk drives. Use this command only for expansion enclosures that use drawers. If the storage subsystem does not support the Service Action Allowed indicator light feature, this command returns an error. If the storage subsystem supports the command but is unable to turn on or turn off the indicator light, this command returns an error.

### **Syntax**

set enclosure [enclosureID]
drawer [drawerID]
serviceAllowedIndicator=(on | off | forceOnWarning)

# **Parameters**

Parameter	Description
enclosure	The enclosure where the drawer resides. Enclosure ID values are 0 to 99. Enclose the enclosure ID value in square brackets ([]). If you do not enter an enclosure ID value, the enclosure ID of the controller module is the default value.
drawer	The location of the drawer for which you want to turn on or turn off the service action allowed indicator. Drawer ID values are 1 to 5. Enclose the value in square brackets ([]).
serviceAllowedIndicator	The setting to turn on or turn off the Service Action Allowed indicator light. To turn on the Service Action Allowed indicator light, set this parameter to on. To turn off the Service Action Allowed indicator light, set this parameter to off.

### Notes

Before you can enter this command, the expansion drawer must meet these conditions:

- The expansion drawer cannot be over temperature.
- The fans must have a status of Optimal.
- All expansion drawer components must be in place. The logical drives in the disk drive drawer cannot be in a Degraded state. If you remove disk drives from the disk drive drawer and a logical drive is already in a Degraded state, the logical drive can fail.

### Attention

Read the following before you issue this command.

- All logical drives with disk drives in the affected disk drive drawer are checked to make sure that the logical drives have drawer loss protection before the command is sent. If the logical drives have drawer loss protection, the Set Service Action Allowed command proceeds without stopping I/O activity to the logical drive. If any logical drives in the affected drawer do not have drawer loss protection, you must stop I/O activity to those logical drives. A warning appears, which indicates that this command should not be completed.
- If you are preparing a component for removal and want to override the warning that the logical drives do not have drawer loss protection, enter this parameter: *serviceAllowedIndicator=forceOnWarning*. This parameter sends the request to prepare to remove a component to the controller firmware, and forces the set drawer serviceAllowedIndicator command to proceed.
- To turn on or turn off the Service Action Allowed indicator light for the entire high-capacity expansion drawer, use the set enclosure serviceAllowedIndicator command.

# Set Drive Channel Status

This command defines how the disk drive channel performs.

### Syntax

set driveChannel [(1 | 2 | 3 | 4 | 5 | 6 | 7 | 8)]
status=(optimal | degraded)

#### **Parameters**

Parameter	Description
driveChannel	The identifier number of the disk drive channel for which you want to set the status. Valid disk drive channel values are 1, 2, 3, or 4. Enclose the disk drive channel number in square brackets ([]).
status	The condition of the disk drive channel. You can set the disk drive channel status to optimal or degraded.

#### Notes

Use the optimal value to move a degraded disk drive channel back to the Optimal state. Use the degraded value when the disk drive channel is experiencing problems, and the storage subsystem requires additional time for data transfers.

# Set Disk Drive Hot Spare

This command assigns or unassigns one or more disk drives as a hot spare.

#### Syntax

```
set (Drive [enclosureID,drawerID1,slotID] |
Drives (enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn)
hotSpare=(TRUE | FALSE)
```

#### **Parameters**

Parameter	Description	
drive or drives	EnclosureID1,drawerID1,slotID1 enclosureIDn,drawerIDn,slotIDn is for enclosures with multiple disk drive drawers. For enclosures without disk drive drawers, use enclosureID1,slotID1 enclosureIDn,slotIDn instead. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the values in square brackets ([ ]).	
hotSpare	The setting to assign the disk drive as the hot spare. To assign the disk drive as the hot spare, set this parameter to TRUE. To remove a hot spare assignment from a disk drive, set this parameter to FALSE.	

#### Notes

The **drive** parameter supports both high-capacity drive expansion enclosures and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need only specify the ID of the drive

expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

# Set Disk Drive Service Action Allowed Indicator

This command turns on or turns off the Service Action Allowed indicator light on a disk drive in drive expansion enclosures that support the Service Action Allowed indicator light feature. If the storage subsystem does not support the Service Action Allowed indicator light feature, this command returns an error. If the storage subsystem supports the command but is unable to turn on or turn off the indicator light, this command returns an error.

### **Syntax**

set (Drive [enclosureID,drawerID1,slotID] |
Drives (enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn)
serviceAllowedIndicator=(on | off)

### **Parameters**

Parameter	Description
drive or drives	The location of the disk drive for which you want to turn on or turn off the Service Action Allowed indicator light. For high-capacity drive expansion enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive expansion enclosures, specify the enclosure ID value and the slot ID value for the drive. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID value, the drawer ID value, and the slot ID value in square brackets ([]).
serviceAllowedIndicator	The setting to turn on or turn off the Service Action Allowed indicator light. To turn on the Service Action Allowed indicator light, set this parameter to on. To turn off the Service Action Allowed indicator light, set this parameter to off.

### Notes

The **drive** parameter supports both high-capacity and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need to specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

# Set Drive State

This command sets a disk drive to the Failed state. (To return a disk drive to the Optimal state, use the revive drive command.)

### **Syntax**

set Drive [enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn]
operationalState=failed

### Parameter

Parameter	Description
drive	The enclosure, drawer, and the slot where the disk drive resides. EnclosureID1,drawerID1,slotID1 enclosureIDn,drawerIDn,slotIDn is for enclosures with multiple disk drive drawers. For enclosures without disk drive drawers, use enclosureID1,slotID1 enclosureIDn,slotIDn instead. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the values in square brackets ([ ]).

### Notes

The **drive** parameter supports both high-capacity drive expansion enclosures and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need only specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

# Set Enclosure Alarm

This command turns on, turns off, or mutes the audible alarm for a specific enclosure or all of the enclosures in a storage subsystem.

### **Syntax**

set (allenclosures | enclosure [enclosureID]
alarm=(enable | disable | mute))

Parameter	Description
allEnclosures	The setting to select all of the enclosures in a storage subsystem that have audible alarms that you want to turn on, turn off, or mute.
enclosure	The specific enclosure that has the audible alarm that you want to turn on, turn off, or mute. Enclosure ID values are 0 to 99. Enclose the enclosure ID value in square brackets ([]). You can specify only one enclosure ID at a time

Parameter	Description
alarm	The setting for the audible alarm. This alarm has these values:
	<ul> <li>enable – The audible alarm is turned on and sounds if a fault occurs.</li> </ul>
	<ul> <li>disable – The audible alarm is turned off and does not sound if a fault occurs.</li> </ul>
	• mute – The audible alarm is turned off if it is sounding.
	(If another fault occurs after you set the audible alarm to mute, the audible alarm sounds again.)

# Set Enclosure Identification

This command sets the enclosure ID of a controller module or an expansion drawer in a storage subsystem. This command is valid only for controller modules or expansion drawers that have enclosure IDs that you can set through the controller firmware. You cannot use this command for controller modules or expansion drawers that have an enclosure ID that you set with a switch.

### **Syntax**

set enclosure ["serial-number"] id=enclosureID

### **Parameters**

Parameter	Description
enclosure	The serial number of the controller module, controller module, or the expansion drawer for which you are setting the enclosure ID. Serial numbers can be any combination of alphanumeric characters and any length. Enclose the serial number in double quotation marks ("").
id	The value for the controller module enclosure ID, controller module enclosure ID, or the expansion drawer enclosure ID. Enclosure ID values are 00 through 99. You do not need to enclose the enclosure ID value in parentheses.

### **Notes**

This command originally supported the DS4800 (models 82, 84, and 88) controller module. The DS4200–, DS4700–, and DS4800– (models 82, 84, and 88) series controller modules can connect to a variety of expansion drawers, including those whose enclosure IDs are set by switches. When connecting a controller module to expansion drawers whose enclosure IDs are set by switches, recommended values for enclosure IDs for the controller module are 80 through 99. This range avoids conflicts with enclosure IDs that are used for attached expansion drawers.

# Set Enclosure Service Action Allowed Indicator

This command turns on or turns off the Service Action Allowed indicator light on a power-fan CRU, an interconnect-battery CRU, or an environmental services monitor (ESM) CRU. If the storage subsystem does not support the Service Action Allowed indicator light feature, this command returns an error. If the storage subsystem supports the command but is unable to turn on or turn off the indicator light, this command returns an error. To turn on or turn off the Service Action Allowed indicator light on the controller CRU, use the set controller serviceAllowedIndicator command.

### **Syntax**

```
set enclosure [enclosureID]
(powerFan [(left | right)] |
interconnect |
battery [(left | right)] |
esm [(left | right)])
serviceAllowedIndicator=(on | off)
```

## **Parameters**

Parameter	Description
enclosure	The enclosure where the power-fan CRU, the interconnect-battery CRU, or the ESM CRU resides. Enclosure ID values are $\theta$ to 99. Enclose the enclosure ID value in square brackets ([]). If you do not enter a enclosure ID value, the enclosure ID of the controller module is the default value.
powerFan	The Service Action Allowed indicator light on the power-fan CRU that you want to turn on or turn off. Valid power-fan CRU identifiers are left or right. The values left and right are with respect to the front of the DS5300–, DS5100–, and DS4800–series controller modules. Enclose the power-fan CRU identifier in square brackets ([]).
interconnect	The Service Action Allowed indicator light for the interconnect-battery CRU DS5300–, DS5100–, and DS4800–series controller modules only.
battery	The Service Action Allowed indicator light for a battery. Valid battery identifiers are left or right. The values left and right are with respect to the front of the DS5300–, DS5100–, and DS4800–series controller modules.
esm	The Service Action Allowed indicator light for an ESM CRU. Valid ESM CRU identifiers are left or right. (The values left and right are with respect to the rear of the drive expansion enclosure.)
serviceAllowedIndicator	The setting to turn on or turn off the Service Action Allowed indicator light. To turn on the Service Action Allowed indicator light, set this parameter to on. To turn off the Service Action Allowed indicator light, set this parameter to off.

### Example

This command turns on the Service Action Allowed indicator light for the left ESM on enclosure 5 with the IP address of 155.155.155.155.

SMcli 155.155.155.155 -c "set enclosure [5] ESM [left] serviceAllowedIndicator=on";

### Notes

This command was originally defined for use with the DS4800 (Models 82, 84, 88) controller enclosure. This command is not supported by controller enclosures that

were shipped before the introduction of the DS4800 (Models 82, 84, 88) controller enclosure.

# Set Enhanced FlashCopy Group Attributes

This command defines the properties for an Enhanced FlashCopy group.

### **Syntax**

```
set enhancedFlashCopyGroup ["
enhancedFlashCopyGroupName"]
[userLabel="enhancedFlashCopyGroupName" |
repositoryFullPolicy=(failBaseWrites |
purgeEnhancedFlashCopyImages) |
repositoryFullLimit=percentValue |
autoDeleteLimit=numberOfEnhancedFlashCopyImages |
rollbackPriority=(lowest | low | medium | high | highest)
```

Parameter	Description
enhancedF1ashCopyGroup	The alphanumeric identifier (including - and _) of the Enhanced FlashCopy group for which you are setting properties. Enclose the Enhanced FlashCopy group name in double quotation marks (" ") inside of square brackets ([ ]).
userLabel	The new alphanumeric identifier (including - and _) that you want to give to the Enhanced FlashCopy group. Use this parameter when you want to rename the Enhanced FlashCopy group. Enclose the new Enhanced FlashCopy group name in double quotation marks (" ").
repositoryFullPolicy	How you want Enhanced FlashCopy image processing to continue if the repository logical drives for the Enhanced FlashCopy group are full. You can choose to <i>fail writes</i> to the base logical drive ( <b>failBaseWrites</b> ) or <i>delete</i> (purge) the Enhanced FlashCopy images ( <b>purgeEnhancedFlashCopyImages</b> ). The default action is <b>purgeEnhancedFlashCopyImages</b>
repositoryFullLimit	The percentage of repository logical drive capacity at which you receive a warning that the Enhanced FlashCopy image repository logical drive is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.
autoDeleteLimit	The maximum number of Enhanced FlashCopy images that you want to automatically delete if you have selected to purge the Enhanced FlashCopy images for a repository full policy. Use integer values. The default value is 32.

Parameter	Description
rollBackPriority	Use this parameter to determine whether system resources should be allocated to the rollback operation at the expense of system performance. Valid values are highest, high, medium, low, or lowest. A value of <i>high</i> indicates that the rollback operation is prioritized over all other host I/O. A value of <i>lowest</i> indicates that the rollback operation should be performed with minimal impact to host I/O.

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

When you use this command, you can specify one or more of the parameters. You do not, however, need to use all of the parameters.

# Set Enhanced FlashCopy Group Media Scan

This command runs a media scan on an Enhanced FlashCopy group.

### Syntax

```
set enhancedFlashCopyGroup ["
enhancedFlashCopyGroupName"] mediaScanEnabled=(TRUE | FALSE)
redundancyCheckEnabled=(TRUE | FALSE)
```

### Parameters

Parameter	Description
enhancedF1ashCopyGroup	The alphanumeric identifier (including - and _) of the Enhanced FlashCopy group on which you want to run a media scan. Enclose the Enhanced FlashCopy group name in double quotation marks (" ") inside of square brackets ([ ]).
mediaScanEnabled	The setting to turn <i>on</i> or turn <i>off</i> media scan for the logical drive. To turn on media scan, set this parameter to TRUE. To turn off media scan, set this parameter to FALSE. (If media scan is disabled at the storage subsystem level, this parameter has no effect.)
redundancyCheckEnabled	The setting to turn <i>on</i> or turn <i>off</i> redundancy checking during a media scan. To turn on redundancy checking, set this parameter to TRUE. To turn off redundancy checking, set this parameter to FALSE.

# Set Enhanced FlashCopy Group Schedule

This command defines the schedule for taking Enhanced FlashCopy images for an Enhanced FlashCopy group.

# Syntax

```
set enhancedFlashCopyGroup ["
enhancedFlashCopyGroupName"] enableSchedule=(TRUE | FALSE)
schedule=(immediate | flashcopySchedule)
```

### **Parameters**

Parameter	Description
enhancedF1ashCopyGroup	The name of the Enhanced FlashCopy group for which you are setting properties. Enclose the Enhanced FlashCopy group name in double quotation marks (" ") inside of square brackets ([ ]).
enableSchedule	Use this parameter to turn <i>on</i> or to turn <i>off</i> the ability to schedule an Enhanced FlashCopy operation. To turn <i>on</i> Enhanced FlashCopy scheduling, set this parameter to TRUE. To turn <i>off</i> Enhanced FlashCopy scheduling, set this parameter to FALSE.
schedule	Use this parameter to schedule an Enhanced FlashCopy group operation. You can use one of these options for setting a schedule for an Enhanced FlashCopy group operation: • immediate
	• startDate
	<ul><li>scheduleDay</li><li>startTime</li></ul>
	• scheduleInterval
	• endDate
	• noEndDate
	• timeZone
	See the <b>Notes</b> section for information explaining how to use these options.

### Notes

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

# Scheduling Enhanced FlashCopies

The **enableSchedule** parameter and the **schedule** parameter provide a way for you to schedule creating Enhanced FlashCopy images for an Enhanced FlashCopy group. Using these parameters, you can schedule Enhanced FlashCopies daily, weekly, or monthly (by day or by date). The **enableSchedule** parameter turns *on* or turns *off* the ability to schedule Enhanced FlashCopies. When you enable scheduling, you use the **schedule** parameter to define when you want the Enhanced FlashCopies to occur.

This list explains how to use the options for the schedule parameter:

- **immediate** As soon as you enter the command, an Enhanced FlashCopy image is created, and a copy-on-write operation begins.
- **startDate** A specific date on which you want to create an Enhanced FlashCopy image and perform a copy-on-write operation. The format for entering the date is MM:DD:YY. If you do not provide a start date, the current date is used. An example of this option is startDate=06:27:11.
- **scheduleDay** A day of the week on which you want to create an Enhanced FlashCopy image and perform a copy-on-write operation. You can enter these values: monday, tuesday, wednesday, thursday, friday, saturday and sunday. An example of this option is scheduleDay=wednesday.
- startTime The time of a day that you want to create an Enhanced FlashCopy image and start performing a copy-on-write operation. The format for entering the time is HH:MM, where HH is the hour and MM is the minute past the hour. Use a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is startTime=14:27.
- scheduleInterval An amount of time, in minutes, that you want to have as a minimum between copy-on-write operations. You can possibly create a schedule in which you have overlapping copy-on-write operations because of the duration a copy operation. You can ensure that you have time between copy-on-write operations by using this option. The maximum value for the scheduleInterval option is 1440 minutes. An example of this option is scheduleInterval=180.
- endDate A specific date on which you want to stop creating an Enhanced FlashCopy image and end the copy-on-write operation. The format for entering the date is MM:DD:YY. An example of this option is endDate=11:26:11.
   noEndDate is a value that can be specified for the endDate parameter.
- **timesPerDay** The number of times that you want the schedule to run in a day. An example of this option is timesPerDay=4.
- **timeZone** Use this parameter to define the time zone in which the storage subsystem is operating. You can define the time zone in one of two ways:
  - GMT±HH:MM The time zone offset from GMT. Enter the offset in hours and minutes. For example GMT-06:00 is the central time zone in the United States.
  - Text string Standard time zone text strings. For example: "USA/Chicago" or "Germany/Berlin". Time zone text strings s are case sensitive. If you enter an incorrect text string, GMT time is used. Enclose the text string in double quotation marks.

The code string for defining a schedule is similar to these examples: enableSchedule=true schedule startTime=14:27 enableSchedule=true schedule scheduleInterval=180 enableSchedule=true schedule scheduleInterval=180 enableSchedule=true schedule timeZone="USA/Chicago"

If you also use the **scheduleInterval** option, the firmware chooses between the **timesPerDay** option and the **scheduleInterval** option by selecting the lowest value of the two options. The firmware calculates an integer value for the **scheduleInterval** option by dividing 1440 by the **scheduleInterval** option value that you set. For example, 1440/180 = 8. The firmware then compares the **timesPerDay** integer value with the calculated **scheduleInterval** integer value and uses the smaller value.

To remove a schedule, use the delete logicaldrive command with the schedule parameter. The delete logicaldrive command with the schedule parameter deletes only the schedule, not the Enhanced FlashCopy logical drive.

The **enableSchedule** parameter and the **schedule** parameter provide a way for you to schedule automatic Enhanced FlashCopy. Using these parameters, you can schedule Enhanced FlashCopy daily, weekly, or monthly (by day or by date). The **enableSchedule** parameter turns *on* or turns *off* the ability to schedule Enhanced FlashCopy. When you enable scheduling, you use the schedule parameter to define when you want the Enhanced FlashCopy to occur.

This list explains how to use the options for the schedule parameter:

- **immediate** As soon as you enter the command, an Enhanced FlashCopy logical drive is created, and a copy-on-write operation begins.
- **startDate** A specific date on which you want to create an Enhanced FlashCopy logical drive and perform a copy-on-write operation. The format for entering the date is MM:DD:YY. If you do not provide a start date, the current date is used. An example of this option is startDate=06:27:11.
- **scheduleDay** A day of the week on which you want to create an Enhanced FlashCopy logical drive and perform a copy-on-write operation. You can enter these values: monday, tuesday, wednesday, thursday, friday, saturday, and sunday. An example of this option is scheduleDay=wednesday.
- startTime The time of a day that you want to create an Enhanced FlashCopy logical drive and start performing a copy-on-write operation. The format for entering the time is HH:MM, where HH is the hour and MM is the minute past the hour. Use a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is startTime=14:27
- scheduleInterval An amount of time, in minutes, that you want to have as a minimum between copy-on-write operation. You can possibly create a schedule in which you have overlapping copy-on-write operations because of the duration a copy operation. You can make sure that you have time between copy-on-write operations by using this option. The maximum value for the scheduleInterval option is 1440 minutes. An example of this option is scheduleInterval=180.
- **endDate** A specific date on which you want to stop creating an Enhanced FlashCopy logical drive and end the copy-on-write operation. The format for entering the date is MM:DD:YY. An example of this option is endDate=11:26:11.
- noEndDate Use this option if you do not want your scheduled copy-on-write operation to end. If you later decide to end the copy-on-write operations you must re-enter the set enhancedFlashCopyGroup command and specify an end date.
- **timesPerDay** The number of times that you want the schedule to run in a day. An example of this option is timesPerDay=4.
- **timeZone** Use this parameter to define the time zone in which the storage subsystem is operating. You can define the time zone in one of two ways:
  - GMT±HH:MM The time zone offset from GMT. Enter the offset in hours and minutes. For example GMT-06:00 is the central time zone in the United States.
  - Text string Standard time zone text strings. For example: "USA/Chicago" or "Germany/Berlin". Timezone text strings are case sensitive. If you enter an incorrect text string, GMT time is used. Enclose the text string in double quotation marks.

The code string for defining a schedule is similar to these examples:

enableSchedule=true schedule startTime=14:27
enableSchedule=true schedule scheduleInterval=180
enableSchedule=true schedule timeZone=GMT-06:00
enableSchedule=true schedule timeZone="USA/Chicago"

If you also use the **scheduleInterval** option, the firmware chooses between the **timesPerDay** option and the **scheduleInterval** option by selecting the lowest value of the two options. The firmware calculates an integer value for the **scheduleInterval** option by dividing 1440 by the **scheduleInterval** option value that you set. For example, 1440/180 = 8. The firmware then compares the **timesPerDay** integer value with the calculated **scheduleInterval** integer value and uses the smaller value.

To remove a schedule, use the command set enhancedFlashCopyGroup ["enhancedFlashCopyGroupName"] enableSchedule=FALSE. This command deletes only the schedule, not the Enhanced FlashCopy group.

# Set Enhanced FlashCopy Logical Drive Media Scan

This command runs a media scan on the drives used for an Enhanced FlashCopy logical drive. Optionally, you also can perform a redundancy check on the data.

### **Syntax**

```
set enhancedFlashCopyLogicalDrive ["
enhancedFlashCopyLogicalDriveName"] mediaScanEnabled=(TRUE | FALSE)
[redundancyCheckEnabled=(TRUE | FALSE)]
```

Parameter	Description
enhancedFlashCopyLogicalDrive	The name of the Enhanced FlashCopy logical drive for which you are setting properties. Enclose the Enhanced FlashCopy logical drive identifier in double quotation marks (" ") inside of square brackets ([ ]).
mediaScanEnabled	The setting to turn <i>on</i> or turn <i>off</i> media scan for the Enhanced FlashCopy logical drive. To turn <i>on</i> media scan, set this parameter to TRUE. To turn <i>off</i> media scan, set this parameter to FALSE. (If media scan is disabled at the storage subsystem level, this parameter has no effect.)
redundancyCheckEnabled	The setting to turn <i>on</i> or turn <i>off</i> redundancy checking during a media scan. To turn <i>on</i> redundancy checking, set this parameter to TRUE. To turn <i>off</i> redundancy checking, set this parameter to FALSE.

#### Parameters

### Notes

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

# Set Enhanced FlashCopy Logical Drive Repository Logical Drive Capacity

This command defines the capacity that you can set for an Enhanced FlashCopy logical drive repository logical drive.

### **Syntax**

```
set enhancedFlashCopyLogicalDrive ["enhancedFlashCopyLogicalDriveName"]
(repositoryFullLimit=percentValue |
increaseRepositoryCapacity (repositoryLogicalDrive=(repos_xxxx) |
decreaseRepositoryCapacity count=numberOfLogicalDrives)
```

### **Parameters**

Parameter	Description
enhancedFlashCopyLogicalDrive	The name of the Enhanced FlashCopy logical drive for which you are setting properties. Enclose the Enhanced FlashCopy logical drive identifier in double quotation marks (" ") inside of square brackets ([ ]).
repositoryFullLimit	The percentage of repository logical drive capacity at which you receive a warning that the Enhanced FlashCopy repository logical drive is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.
repositoryLogicalDrive	The alphanumeric identifier of the repository logical drive has the Enhanced FlashCopy image. All repository identifiers have this form: repos_xxxx where xxxx represents a four digit numerical value. Enclose the Enhanced FlashCopy logical drive name in double quotation marks (" ").
count	The number of repository logical drives that you want to remove from the Enhanced FlashCopy group. Use integer values.

### Notes

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

When you use this command, you can specify only one of the parameters. You cannot enter more than one parameter when you run this command.

# Set Enhanced FlashCopy Group Repository Logical Drive Capacity

This command increases or decreases the capacity of an Enhanced FlashCopy group repository logical drive.

# Syntax for Increasing Capacity

set enhancedFlashCopyGroup ["
enhancedFlashCopyGroupName"]
increaseRepositoryCapacity repositoryLogicalDrives=(repos\_xxxx)

# Syntax for Decreasing Capacity

set enhancedFlashCopyGroup ["
enhancedFlashCopyGroupName"]
decreaseRepositoryCapacity count=numberOfLogicalDrives

# **Parameters**

Parameter	Description
enhancedF1ashCopyGroup	The alphanumeric identifier (including - and _) of the Enhanced FlashCopy group on which you want to run a media scan. Enclose the Enhanced FlashCopy group name in double quotation marks (" ") inside of square brackets ([ ]).
repositoryLogicalDrive	<ul> <li>The name of the repository logical drive for which you want to increase capacity. An available standard logical drive is added to the repository logical drive to increase the capacity of the repository logical drive.</li> <li>The repository logical drive name is comprised of two parts:</li> <li>The term <b>repos</b></li> <li>A numerical identifier generated by the storage management software and the firmware</li> </ul>
count	The number of repository logical drives that you want to remove from the Enhanced FlashCopy group. Use integer values.

# Set FlashCopy Logical Drive

This command defines the properties for a FlashCopy logical drive and lets you rename a FlashCopy logical drive.

## **Syntax**

```
set (logicalDrive [logicalDriveName] |
logicalDrives [logicalDriveName1 ... logicalDriveNameN])
warningThresholdPercent=percentValue
repositoryFullPolicy=(failBaseWrites | failFlashCopy)
```

# Parameters

Parameter	Description
logicalDrive or logicalDrives	The name of the specific FlashCopy logical drive for which you want to define properties. (You can enter more than one logical drive name if you use the <b>logicalDrives</b> parameter). Enclose the FlashCopy logical drive name in double quotation marks ("") inside of square brackets ([]).
userLabel	A new name that you want to give to a FlashCopy logical drive. Enclose the new FlashCopy logical drive name in double quotation marks ("").
warningThresholdPercent	The percentage of repository capacity at which you receive a warning that the FlashCopy repository is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 50.
repositoryFullPolicy	How you want FlashCopy processing to continue if the FlashCopy repository logical drive is full. You can choose to fail writes to the base logical drive (failBaseWrites) or fail writes to the FlashCopy logical drive (failFlashCopy). The default value is failFlashCopy.

### Notes

When you use this command, you can specify one or more of the optional parameters.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

You can set the **warningThresholdPercent** parameter and the **repositoryFullPolicy** parameter for both the FlashCopy repository logical drive or the FlashCopy logical drive.

# Set Foreign Disk Drive to Native

A drive is considered to be native when it is a part of the storage subsystem's subsystem configuration. A drive is considered to be foreign when it does not belong to the storage subsystem's subsystem configuration or when it fails to be imported with the drives of an subsystem that is transferred to a new storage subsystem. The latter failure creates an incomplete subsystem on the new storage subsystem.

Run this command to add the missing (foreign) drives back into their original subsystem and to make them part of the new storage subsystem subsystem configuration of the new storage subsystem.

Use this operation for emergency recovery only: when one or more drives need to be changed from a foreign drive status and returned to a native status within their original subsystem. **Attention: Possible data corruption or data loss**—Using this command for reasons other than what is stated previously might result in data loss without notification.

#### Syntax

set (drive [enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn] |

allDrives) nativeState

#### **Parameters**

Parameter	Description
drive	EnclosureID1,drawerID1,slotID1 enclosureIDn,drawerIDn,slotIDn is for enclosures with multiple disk drive drawers. For enclosures without disk drive drawers, use enclosureID1,slotID1 enclosureIDn,slotIDn instead. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the values in square brackets ([]).
allDrives	The setting to select all of the drives.

### Notes

The **drive** parameter supports both high-capacity drive expansion enclosures and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive resides. For a low-capacity drive expansion enclosure, you need only specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to  $\theta$ , and specify the ID of the slot in which a drive resides.

# Set Host

This command assigns a host to a host group or moves a host to a different host group. You can also create a new host group and assign the host to the new host group with this command. The actions performed by this command depend on whether the host has individual Volume-to-LUN mappings or does not have individual Volume-to-LUN mappings.

### **Syntax**

```
set host [hostName]
hostGroup=("hostGroupName" | none | defaultGroup)
userLabel="newHostName"
hostType=(hostTypeIndexLabel | hostTypeIndexNumber)
```

# **Parameters**

Parameter	Description
host	The name of the host that you want to assign to a host group. Enclose the host name in square brackets ([]). If the host name has special characters, you also must enclose the host name in double quotation marks ("").
hostGroup	The name of the host group to which you want to assign the host. (The following table defines how the command runs if the host does or does not have individual logical drive-to-LUN mappings.) Enclose the host group name in double quotation marks (""). The defaultGroup option is the host group that contains the host to which the logical drive is mapped.
userLabel	The new host name. Enclose the host name in double quotation marks ("").
hostType	The index label or number of the host type for the HBA host port. Use the <b>show storageSubsystem hostTypeTable</b> command to generate a list of available host type identifiers. If the host type has special characters, enclose the host type in double quotation marks ("").

Host Group Parameter	Host Has Individual Volume-to-LUN mappings	Host Does Not Have Individual Logical drive-to-LUN Mappings
hostGroupName	The host is removed from the present host group and is placed under the new host group defined by <i>hostGroupName</i> .	The host is removed from the present host group and is placed under the new host group defined by <i>hostGroupName</i> .
none	The host is removed from the host group as an independent partition and is placed under the root node.	The host is removed from the present host group and is placed under the default group.
defaultGroup	The command fails.	The host is removed from the present host group and is placed under the default group.

# Notes

When you use this command, you can specify one or more of the optional parameters.

For the names, you can use any combination of alphanumeric characters, hyphens, and underscores. Names can have a maximum of 30 characters.

# **Related Topic**

Naming Conventions

# **Set Host Channel**

This command defines the loop ID for the host channel.

# Syntax

set hostChannel [hostChannelNumber]
preferredID=portID

### **Parameters**

Parameter	Description
hostChannel	The identifier number of the host channel for which you want to set the loop ID. Enclose the host channel identifier number in square brackets ([ ]).
	Valid host channel values are a1, a2, a3, a4, a5, a6, a7, or a8 for controller A host ports, and b1, b2, b3, b4, b5, b6, b7, or b8 for controller B host ports. Depending on the number of host channels that are currently installed in your storage subsystems, you can specify two, four or up to 8 host channels per controller.
	For example the DS5100 and DS5300 series controller modules, valid host channel values are a1, a2, a3, a4, a5, a6, a7, a8, b1, b2, b3, b4, b5, b6, b7, or b8.
preferredID	The new port ID for the specified host channel. Port ID values are $\theta$ to 127.

# Set Host Group

This command renames a host group.

### **Syntax**

set hostGroup [hostGroupName]
userLabel="newHostGroupName"

### **Parameters**

Parameter	Description
hostGroup	The name of the host group that you want to rename. Enclose the host group name in square brackets ([]). If the host group name has special characters, you also must enclose the host group name in double quotation marks ("").
userLabel	The new name for the host group. Enclose the new host group name in double quotation marks ("").

### Notes

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

### **Related Topic**

Naming Conventions

# Set Host Port

This command changes the host type for an HBA host port. You can also change an HBA host port label with this command.

#### **Syntax**

set hostPort [portLabel] host="hostName" userLabel="newPortLabel"

#### **Parameters**

Parameter	Description
hostPort	The name of the HBA host port for which you want to change the host type, or for which you want to create a new name. Enclose the HBA host port name in square brackets ([]). If the HBA host port label has special characters, enclose the HBA host port label in double quotation marks ("").
host	The name of the host to which the HBA host port is connected. Enclose the host name in double quotation marks ("").
userLabel	The new name that you want to give to the HBA host port. Enclose the new name of the HBA host port in double quotation marks ("").

### Notes

When you use this command, you can specify one or more of the optional parameters.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

### **Related Topic**

Naming Conventions

# Set iSCSI Initiator

This command sets the attributes for an iSCSI initiator.

#### Syntax

```
set iscsiInitiator iscsiName] |
userLabel=newName |
host=newHostName |
chapSecret=newSecurityKey
```

Parameter	Description
iscsiName	The name of the iSCSI initiator for which you want to set attributes. Enclose the value in double quotation marks ("").

Parameter	Description
userLabel	The new name that you want to use for the iSCSI initiator. Enclose the value in double quotation marks ("").
host	The name of the new host to which the HBA host port is connected. Enclose the value in double quotation marks ("").
chapSecret	The security key that you want to use to authenticate a peer connection. Enclose the value in double quotation marks ("").

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

Challenge Handshake Authentication Protocol (CHAP) is a protocol that authenticates the peer of a connection. CHAP is based upon the peers sharing a *secret*. A secret is a security key that is similar to a password.

Use the **chapSecret** parameter to set up the security keys for initiators that require a mutual authentication.

# **Related Topic**

Naming Conventions

# **Set iSCSI Target Properties**

This command defines properties for an iSCSI target.

### **Syntax**

```
set iscsiTarget ["iSCSI-ID"] authenticationMethod=chap chapSecret=securityKey
[targetAlias=user-label]
Or
set iscsiTarget ["iSCSI-ID"] authenticationMethod= none
[targetAlias=user-label]
```

### **Parameters**

Parameter	Description
iscsiTarget	The iSCSI target for which you want to set properties. Enclose the userLabel in double quotation marks (" "). You must also enclose the userLabel in either square brackets ([ ]) or angle brackets (< >).
authenticationMethod	The means of authenticating your iSCSI session.
chapSecret	The security key that you want to use to authenticate a peer connection.
targetAlias	The name that you want to use for the target.

There is no isnsRegistration parameter in this command.

Challenge Handshake Authentication Protocol (CHAP) is a protocol that authenticates the peer of a connection. CHAP is based upon the peers sharing a *secret*. A secret is a security key that is similar to a password.

Use the **chapSecret** parameter to set up the security keys for initiators that require a mutual authentication.

# Set Logical Drive Attributes

This command defines the properties for a logical drive. You can use most parameters to define properties for one or more logical drives. You also can use some parameters to define properties for only one logical drive. The syntax definitions are separated to show which parameters apply to several logical drives and which apply to only one logical drive. Also, the syntax for logical drive mapping is listed separately.

**Note:** In configurations with subsystems that have more than 32 logical drives, the timeout period might expire before an operation completes. This causes host I/O errors or internal controller reboots. If such errors or reboots occur, quiesce the host I/O and retry the operation.

#### Syntax Applicable to One or More Logical Drives

```
set (allLogicalDrives | logicalDrive ["logicalDriveName"] |
logicalDrives ["logicalDriveName1" ... "logicalDriveNameN"] | logicalDrive <wwID>)
cacheFlushModifier=cacheFlushModifierValue
cacheWithoutBatteryEnabled=(TRUE | FALSE)
mediaScanEnabled=(TRUE | FALSE)
mirrorCacheEnabled=(TRUE | FALSE)
modificationPriority=(highest | high | medium | low | lowest)
owner=(a | b)
preReadRedundancyCheck=(TRUE | FALSE)
readCacheEnabled=(TRUE | FALSE)
writeCacheEnabled=(TRUE | FALSE)
cacheReadPrefetch=(TRUE | FALSE)
T10PIDisabled=(TRUE | FALSE)
```

### Syntax Applicable to Only One Logical Drive

```
set (logicalDrive ["logicalDriveName"] | logicalDrive <wwID>)
addCapacity=logicalDriveCapacity
[addDrives=(enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn))]
redundancyCheckEnabled=(TRUE | FALSE)
segmentSize=segmentSizeValue
userLabel=logicalDriveName
preReadRedundancyCheck=(TRUE | FALSE)
```

Parameter	Description
alllogicalDrives	The properties for all logical drives in the storage subsystem.
logicalDrive or logicalDrives	The name of the specific logical drive for which you want to define properties. (You can enter more than one logical drive name if you use the <b>logicalDrives</b> parameter). Enclose the logical drive name in double quotation marks ("") inside of square brackets ([]).

Parameter	Description
logicalDrive	The World Wide Identifier (WWID) of the logical drive for which you are setting properties. You can use the WWID instead of the logical drive name to identify the logical drive. Enclose the WWID in angle brackets (< >).
cacheFlushModifier	The maximum amount of time that data for the logical drive stays in cache before the data is flushed to physical storage. Valid values are listed in the "Notes" section.
cacheWithoutBatteryEnabled	The setting to turn on or turn off caching without batteries. To turn on caching without batteries, set this parameter to TRUE. To turn off caching without batteries, set this parameter to FALSE.
mediaScanEnabled	The setting to turn on or turn off media scan for the logical drive. To turn on media scan, set this parameter to TRUE. To turn off media scan, set this parameter to FALSE. (If media scan is disabled at the storage subsystem level, this parameter has no effect.)
mirrorCacheEnabled	The setting to turn on or turn off the mirror cache. To turn on the mirror cache, set this parameter to TRUE. To turn off the mirror cache, set this parameter to FALSE.
modificationPriority	The priority for logical drive modifications while the storage subsystem is operational. Valid values are highest, high, medium, low, or lowest.
owner	The controller that owns the logical drive. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Use this parameter only if you want to change the logical drive owner.
preReadRedundancyCheck	The setting to turn on or turn off preread redundancy checking. Turning on preread redundancy checking verifies the consistency of RAID redundancy data for the stripes containing the read data. Preread redundancy checking is performed on read operations only. To turn on preread redundancy checking, set this parameter to TRUE. To turn off preread redundancy checking, set this parameter to FALSE. <b>Note:</b> Do not use this parameter on non-redundant logical drives, such as RAID 0 logical drives.
readCacheEnabled	The setting to turn on or turn off the read cache. To turn on the read cache, set this parameter to TRUE. To turn off the read cache, set this parameter to FALSE.
writeCacheEnabled	The setting to turn on or turn off the write cache. To turn on the write cache, set this parameter to TRUE. To turn off the write cache, set this parameter to FALSE.

Parameter	Description
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.
T10PIDisabled	The setting to turn on or turn off T10 PI (Protection Information) for a specific logical drive. For this parameter to have meaning, your logical drive must be capable of T10 PI (Protection Information). This parameter changes a logical drive from one that supports T10 PI (Protection Information) to a logical drive that cannot support T10 PI (Protection Information).
	To remove T10 PI (Protection Information) from a logical drive that supports T10 PI (Protection Information), set this parameter to TRUE. To return a logical drive to supporting T10 PI (Protection Information), set this parameter to FALSE. <b>Note:</b> If you remove T10 PI (Protection Information) from a logical drive, you cannot reset T10 PI (Protection Information) for that logical drive.
	<ul> <li>To reset T10 PI (Protection Information) for the data on a logical drive, from which you removed T10 PI (Protection Information):</li> <li>1. Remove the data from the logical drive.</li> <li>2. Delete the logical drive.</li> <li>3. Re-create a new logical drive with the properties of the deleted logical drive.</li> <li>4. Set T10 PI (Protection Information) for the new logical drive.</li> <li>5. Move the data to the new logical drive.</li> </ul>
addCapacity	The setting to increase the storage size (capacity) of the logical drive for which you are defining properties. Size is defined in units of bytes, KB, MB, GB, or TB. The default value is bytes.
addDrives	The setting to add new disk drives to the logical drive. For high-capacity disk drive expansion enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for the disk drive. For low-capacity drive expansion enclosures, specify the enclosure ID value and the slot ID value for the disk drive. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID value in parentheses. Use this parameter with the <b>addCapacity</b> parameter if you need to specify additional disk drives to accommodate the new size.

Parameter	Description
redundancyCheckEnabled	The setting to turn on or turn off redundancy checking during a media scan. To turn on redundancy checking, set this parameter to TRUE. To turn off redundancy checking, set this parameter to FALSE.
segmentSize	The amount of data (in KB) that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
userLabel	The new name that you want to give an existing logical drive. Enclose the new logical drive name in double quotation marks ("").
preReadRedundancyCheck	The setting to check the consistency of RAID redundancy data on the stripes during read operations. Do not use this operation for non-redundant logical drives, for example RAID Level 0. To check redundancy consistency, set this parameter to TRUE. For no stripe checking, set this parameter to FALSE.
accesslogicalDrive	The logical unit number for the access logical drive. The logical unit number is the only property that you can set for the access logical drive.
logicalUnitNumber	Specific Volume-to-LUN mappings to an individual host. This parameter also assigns the host to a host group.
host	The name of the host to which the logical drive is mapped. Enclose the host name in double quotation marks ("").
hostGroup	The name of the host group to which the logical drive is mapped. Enclose the host group name in double quotation marks (""). defaultGroup is the host group that contains the host to which the logical drive is mapped.

In configurations with subsystems that have more than 32 logical drives, the timeout period might expire before an operation completes. This causes host I/O errors or internal controller reboots. If such errors or reboots occur, quiesce the host I/O and retry the operation.

When you use this command, you can specify one or more of the optional parameters.

You can apply these parameters to only one logical drive at a time:

- addCapacity
- segmentSize
- userLabel
- logicalUnitNumber

# Add Capacity, Add Drives, and Segment Size

Setting the **addCapacity** parameter, the **addDrives** parameter, or the **segmentSize** parameter starts a long-running operation that you cannot stop. These long-running operations are performed in the background and do not prevent you from running other commands. To show the progress of long-running operations, use the **show logicalDrive actionProgress** command.

The **addDrives** parameter supports both high-capacity and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need to specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the slot in which a drive expansion enclosure, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

### **Cache Flush Modifier**

Value	Description
Immediate	Data is flushed as soon as it is placed into the cache.
250	Data is flushed after 250 ms.
500	Data is flushed after 500 ms.
750	Data is flushed after 750 ms.
1	Data is flushed after 1 s.
1500	Data is flushed after 1500 ms.
2	Data is flushed after 2 s.
5	Data is flushed after 5 s.
10	Data is flushed after 10 s.
20	Data is flushed after 20 s.
60	Data is flushed after 60 s (1 min.).
120	Data is flushed after 120 s (2 min.).
300	Data is flushed after 300 s (5 min.).
1200	Data is flushed after 1200 s (20 min.).
3600	Data is flushed after 3600 s (1 hr).
Infinite	Data in cache is not subject to any age or time constraints. The data is flushed based on other criteria that are managed by the controller.

Valid values for the cache flush modifier are listed in this table.

### **Cache Without Battery Enabled**

Write caching without batteries enables write caching to continue if the controller batteries are completely discharged, not fully charged, or not present. If you set this parameter to TRUE without an uninterruptible power supply (UPS) or other

backup power source, you can lose data if the power to the storage subsystem fails. This parameter has no effect if write caching is disabled.

### **Modification Priority**

Modification priority defines the amount of system resources that are used when modifying logical drive properties. If you select the highest priority level, the logical drive modification uses the most system resources, which decreases the performance for host data transfers.

### **Cache Read Prefetch**

The **cacheReadPrefetch** parameter enables the controller to copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the disk drive into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage subsystem that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the **cacheReadPrefetch** parameter are TRUE or FALSE.

### Segment Size

The size of a segment determines how many data blocks that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single disk drive for a single request leaves other disk drives available to simultaneously service other requests.

If the logical drive is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of disk drives in the subsystem that are used for data transfers.) In this case, multiple disk drives are used for the same request, but each disk drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of disk drives that are required to satisfy a data transfer request.

### Set Logical Drive Attributes

This command defines the properties for a logical drive. You can use most parameters to define properties for one or more logical drives. You also can use some parameters to define properties for only one logical drive. The syntax definitions are separated to show which parameters apply to several logical drives and which apply to only one logical drive. **Note:** In configurations where subsystems consist of more than 32 logical drives, the operation can result in host I/O errors or internal controller reboots due to the expiration of the timeout period before the operation completes. If you experience host I/O errors or internal controller reboots, quiesce the host I/O and try the operation again.

# Syntax Applicable to One or More Logical Drives

```
set (allLogicalDrives | logicaldrive ["
logicaldriveName"] | logicaldrives ["logicaldriveName1" ... "logicaldriveNameN"] |
logicaldrive <wwID>)
cacheFlushModifier=cacheFlushModifierValue
cacheWithoutBatteryEnabled=(TRUE | FALSE)
mediaScanEnabled=(TRUE | FALSE)
mirrorCacheEnabled=(TRUE | FALSE)
modificationPriority=(highest | high | medium | low | lowest)
owner=(a | b)
preReadRedundancyCheck=(TRUE | FALSE)
readCacheEnabled=(TRUE | FALSE)
writeCacheEnabled=(TRUE | FALSE)
cacheReadPrefetch=(TRUE | FALSE)
T10PIDisabled=(TRUE | FALSE)
```

# Syntax Applicable to Only One Logical Drive

```
set (logicaldrive ["logicaldriveName"] | logicaldrive <wwID>)
addCapacity=logicaldriveCapacity
[addDrives=(enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn)]
redundancyCheckEnabled=(TRUE | FALSE)
segmentSize=segmentSizeValue
userLabel=logicaldriveName
preReadRedundancyCheck=(TRUE | FALSE)
```

Parameter	Description
allLogicalDrives	The properties for all logical drives in the storage subsystem.
logicaldrive or logicaldrives	The name of the specific logical drive for which you want to define properties. (You can enter more than one logical drive name if you use the <b>logicaldrives</b> parameter.) Enclose the logical drive name in double quotation marks (" ") inside of square brackets ([ ]).
logicaldrive	The World Wide Identifier (WWID) of the logical drive for which you are setting properties. You can use the WWID instead of the logical drive name to identify the logical drive. Enclose the WWID in angle brackets (< >).
cacheFlushModifier	The maximum amount of time that data for the logical drive stays in cache before the data is flushed to physical storage. Valid values are listed in the <b>Notes</b> section.
cacheWithoutBatteryEnabled	The setting to turn <i>on</i> or turn <i>off</i> caching without batteries. To turn <i>on</i> caching without batteries, set this parameter to TRUE. To turn <i>off</i> caching without batteries, set this parameter to FALSE.

Parameter	Description
mediaScanEnabled	The setting to turn <i>on</i> or turn <i>off</i> media scan for the logical drive. To turn <i>on</i> media scan, set this parameter to TRUE. To turn <i>off</i> media scan, set this parameter to FALSE. (If media scan is disabled at the storage subsystem level, this parameter has no effect.)
mirrorCacheEnabled	The setting to turn <i>on</i> or turn <i>off</i> the mirror cache. To turn <i>on</i> the mirror cache, set this parameter to TRUE. To turn <i>off</i> the mirror cache, set this parameter to FALSE.
modificationPriority	The priority for logical drive modifications while the storage subsystem is operational. Valid values are highest, high, medium, low, or lowest.
owner	The controller that owns the logical drive. Valid controller identifiers are a or b, where $a$ is the controller in slot A, and $b$ is the controller in slot B. Use this parameter only if you want to change the logical drive owner.
preReadRedundancyCheck	The setting to turn <i>on</i> or turn <i>off</i> pre-read redundancy checking. Turning on pre-read redundancy checking verifies the consistency of RAID redundancy data for the stripes containing the read data. Pre-read redundancy checking is performed on read operations only. To turn <i>on</i> pre-read redundancy checking, set this parameter to TRUE. To turn <i>off</i> pre-read redundancy checking, set this parameter to FALSE. <b>Note:</b> Do not use this parameter on non-redundant logical drives, such as RAID 0 logical drives.
readCacheEnabled	The setting to turn <i>on</i> or turn <i>off</i> the read cache. To turn <i>on</i> the read cache, set this parameter to TRUE. To turn <i>off</i> the read cache, set this parameter to FALSE.
writeCacheEnabled	The setting to turn <i>on</i> or turn <i>off</i> the write cache. To turn <i>on</i> the write cache, set this parameter to TRUE. To turn <i>off</i> the write cache, set this parameter to FALSE.
cacheReadPrefetch	The setting to turn <i>on</i> or turn <i>off</i> cache read prefetch. To turn <i>off</i> cache read prefetch, set this parameter to FALSE. To turn <i>on</i> cache read prefetch, set this parameter to TRUE.

Parameter	Description
T10PIDisabled	The setting to turn <i>on</i> or turn <i>off</i> T10 PI (Protection Information) for a specific logical drive.
	For this parameter to have meaning, your logical drive must be capable of T10 PI (Protection Information). This parameter changes a logical drive from one that supports T10 PI (Protection Information) to a logical drive that cannot support T10 PI (Protection Information).
	To remove T10 PI (Protection Information) from a logical drive that supports T10 PI (Protection Information), set this parameter to TRUE. To return a logical drive to supporting T10 PI (Protection Information), set this parameter to FALSE. <b>Note:</b> If you remove T10PI (Protection Information) from a logical drive, you cannot reset T10 PI (Protection Information) for that logical drive.
	To reset T10 PI (Protection Information) for the data on a logical drive, from which you removed T10 PI (Protection Information), perform these steps:
	1. Remove the data from the logical drive.
	<ol> <li>Delete the logical drive.</li> <li>Re-create a new logical drive with the</li> </ol>
	properties of the deleted logical drive.
	4. Set T10 PI (Protection Information) for the new logical drive.
	5. Move the data to the new logical drive
addCapacity	The setting to increase the storage size (capacity) of the logical drive for which you are defining properties. Size is defined in units of bytes, KB, MB, GB, or TB. The default value is bytes.
addDrives	The setting to add new drives to the logical drive. For high-capacity drive expansion enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive expansion enclosures, specify the enclosure ID value and the slot ID value for the drive. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID value in parentheses. Use this parameter with the <b>addCapacity</b> parameter if you need to specify additional drives to accommodate the new size.

Parameter	Description
redundancyCheckEnabled	The setting to turn <i>on</i> or turn <i>off</i> redundancy checking during a media scan. To turn <i>on</i> redundancy checking, set this parameter to TRUE. To turn <i>off</i> redundancy checking, set this parameter to FALSE.
segmentSize	The amount of data (in KB) that the controller writes on a single drive in a logical drive before writing data on the next drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
userLabel	The new name that you want to give an existing logical drive. Enclose the new logical drive name in double quotation marks (" ").
preReadRedundancyCheck	The setting to check the consistency of RAID redundancy data on the stripes during read operations. Do not use this operation for non-redundant logical drives, for example RAID Level 0. To check redundancy consistency, set this parameter to TRUE. For no stripe checking, set this parameter to FALSE.

Host I/O errors might result in subsystems with more than 32 logical drives. This operation might also result in internal controller reboots due to the expiration of the timeout period before the operation completes. If you experience this issue, quiesce host I/O, and try the operation again.

When you use this command, you can specify one or more of the optional parameters.

You can apply these parameters to only one logical drive at a time:

- addCapacity
- segmentSize
- userLabel
- logicalUnitNumber

# Add Capacity, Add Drives, and Segment Size

Setting the **addCapacity** parameter, the **addDrives** parameter, or the **segmentSize** parameter starts a long-running operation that you cannot stop. These long-running operations are performed in the background and do not prevent you from running other commands. To show the progress of long-running operations, use the show logicaldrive actionProgress command.

The **addDrives** parameter supports both high-capacity drive expansion enclosures and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need only specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

# **Cache Flush Modifier**

Valid values for the cache flush modifier are listed in this table.

Value	Description
Immediate	Data is flushed as soon as it is placed into the cache.
250	Data is flushed after 250 ms.
500	Data is flushed after 500 ms.
750	Data is flushed after 750 ms.
1	Data is flushed after 1 s.
1500	Data is flushed after 1500 ms.
2	Data is flushed after 2 s.
5	Data is flushed after 5 s.
10	Data is flushed after 10 s.
20	Data is flushed after 20 s.
60	Data is flushed after 60 s (1 min.).
120	Data is flushed after 120 s (2 min.).
300	Data is flushed after 300 s (5 min.).
1200	Data is flushed after 1200 s (20 min.).
3600	Data is flushed after 3600 s (1 hr).
Infinite	Data in cache is not subject to any age or time constraints. The data is flushed based on other criteria that are managed by the controller.

### **Cache Without Battery Enabled**

Write caching without batteries enables write caching to continue if the controller batteries are completely discharged, not fully charged, or not present. If you set this parameter to TRUE without an uninterruptible power supply (UPS) or other backup power source, you can lose data if the power to the storage subsystem fails. This parameter has no effect if write caching is disabled.

### **Modification Priority**

Modification priority defines the amount of system resources that are used when modifying logical drive properties. If you select the highest priority level, the logical drive modification uses the most system resources, which decreases the performance for host data transfers.
# **Cache Read Prefetch**

The **cacheReadPrefetch** parameter enables the controller to copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the drive into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage subsystem that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the **cacheReadPrefetch** parameter are TRUE or FALSE.

# Segment Size

The size of a segment determines how many data blocks that the controller writes on a single drive in a logical drive before writing data on the next drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single drive for a single request leaves other drives available to simultaneously service other requests.

If the logical drive is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of drives in the subsystem that are used for data transfers.) In this case, multiple drives are used for the same request, but each drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of drives that are required to satisfy a data transfer request.

### **Related Topic**

Naming Conventions

# Set Logical Drive Attributes for a Disk Pool

This command defines the properties for a logical drive. You can use most parameters to define properties for one or more logical drives. You also can use some parameters to define properties for only one logical drive. The syntax definitions are separated to show which parameters apply to several logical drives and which apply to only one logical drive. Also, the syntax for logical drive mapping is listed separately.

**Note:** In configurations where subsystems consist of more than 32 logical drives, the operation can result in host I/O errors or internal controller reboots due to the expiration of the timeout period before the operation completes. If you experience host I/O errors or internal controller reboots, bring the host to a quiescent state and try the operation again.

```
set (allLogicalDrives | logicaldrive ["
logicaldriveName"] | logicaldrives ["logicaldriveName1" ... "logicaldriveNameN"]
 logicaldrive <wwID>)
[(addCapacity = capacityValue[KB|MB|GB|TB|Bytes] |
[addDrives = (drive-spec-list)])
cacheFlushModifier=cacheFlushModifierValue
cacheReadPrefetch = (TRUE | FALSE)
cacheWithoutBatteryEnabled=(TRUE | FALSE) |
mediaScanEnabled=(TRUE | FALSE) |
mirrorCacheEnabled=(TRUE | FALSE) |
modificationPriority=(highest | high | medium | low | lowest) | owner=(a | b) |
preReadRedundancyCheck = (TRUE | FALSE) |
readCacheEnabled=(TRUE | FALSE) |
redundancyCheckEnabled = (TRUE | FALSE) |
segmentSize = segmentSizeValue
userLabel = userlabelValue
writeCacheEnabled=(TRUE | FALSE) |
T10PIDisabled=(TRUE | FALSE)]
```

Parameter	Description
allLogicalDrives	The properties for all logical drives in the storage subsystem.
logicaldrive or logicaldrives	The name of the specific logical drive for which you want to define properties. (You can enter more than one logical drive name if you use the <b>logicaldrives</b> parameter.) Enclose the logical drive name in double quotation marks (" ") inside of square brackets ([ ]).
logicaldrive	The World Wide Identifier (WWID) of the logical drive for which you are setting properties. You can use the WWID instead of the logical drive name to identify the logical drive. Enclose the WWID in angle brackets (< >).
addCapacity	The setting to increase the storage size (capacity) of the logical drive for which you are defining properties. Size is defined in units of bytes, KB, MB, GB, or TB. The default value is bytes. This parameter is not valid for thin logical drives.

Parameter	Description
addDrives	The setting to add new drives to the logical drive. For high-capacity drive expansion enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for the drive. For low-capacity drive expansion enclosures, specify the enclosure ID value and the slot ID value for the drive. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID value in parentheses. Use this parameter with the <b>addCapacity</b> parameter if you need to specify additional drives to accommodate the new size.
	logical drives or for thin logical drives.
cacheFlushModifier	The maximum amount of time that data for the logical drive stays in cache before the data is flushed to physical storage. Valid values are listed in the <b>Notes</b> section.
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.
cacheWithoutBatteryEnabled	The setting to turn on or turn off caching without batteries. To turn on caching without batteries, set this parameter to TRUE. To turn off caching without batteries, set this parameter to FALSE.
mediaScanEnabled	The setting to turn on or turn off media scan for the logical drive. To turn on media scan, set this parameter to TRUE. To turn off media scan, set this parameter to FALSE. (If media scan is disabled at the storage subsystem level, this parameter has no effect.)
mirrorCacheEnabled	The setting to turn on or turn off the mirror cache. To turn on the mirror cache, set this parameter to TRUE. To turn off the mirror cache, set this parameter to FALSE.
modificationPriority	The priority for logical drive modifications while the storage subsystem is operational. Valid values are highest, high, medium, low, or lowest.
owner	The controller that owns the logical drive. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Use this parameter only if you want to change the logical drive owner.

Parameter	Description
preReadRedundancyCheck	The setting to turn on or turn off pre-read redundancy checking. Turning on pre-read redundancy checking verifies the consistency of RAID redundancy data for the stripes containing the read data. Pre-read redundancy checking is performed on read operations only. To turn on pre-read redundancy checking, set this parameter to TRUE. To turn off pre-read redundancy checking, set this parameter to FALSE. <b>Note:</b> Do not use this parameter on non-redundant logical drives, such as RAID 0 logical drives.
readCacheEnabled	The setting to turn on or turn off the read cache. To turn on the read cache, set this parameter to TRUE. To turn off the read cache, set this parameter to FALSE.
redundancyCheckEnabled	The setting to turn on or turn off the write cache. To turn on the write cache, set this parameter to TRUE. To turn off the write cache, set this parameter to FALSE.
segmentSize	The amount of data (in KB) that the controller writes on a single drive in a logical drive before writing data on the next drive. Valid values are 8, 16, 32, 64, 128,256, or 512
userLabel	The new name that you want to give an existing logical drive. Enclose the new logical drive name in double quotation marks (" ").
writeCacheEnabled	The setting to turn on write cache capability.

When you use this command, you can specify one or more of the optional parameters.

You can apply these parameters to only one logical drive at a time:

- addCapacity
- segmentSize
- userLabel

# Add Capacity, Add Drives, and Segment Size

Setting the **addCapacity** parameter, the **addDrives** parameter, or the **segmentSize** parameter starts a long-running operation that you cannot stop. These long-running operations are performed in the background and do not prevent you from running other commands. To show the progress of long-running operations, use the show logicaldrive actionProgress command.

The **addDrives** parameter supports both high-capacity drive expansion enclosures and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive

expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need only specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

# **Cache Flush Modifier**

Value	Description
Immediate	Data is flushed as soon as it is placed into the cache.
250	Data is flushed after 250 ms.
500	Data is flushed after 500 ms.
750	Data is flushed after 750 ms.
1	Data is flushed after 1 s.
1500	Data is flushed after 1500 ms.
2	Data is flushed after 2 s.
5	Data is flushed after 5 s.
10	Data is flushed after 10 s.
20	Data is flushed after 20 s.
60	Data is flushed after 60 s (1 min.).
120	Data is flushed after 120 s (2 min.).
300	Data is flushed after 300 s (5 min.).
1200	Data is flushed after 1200 s (20 min.).
3600	Data is flushed after 3600 s (1 hr).
Infinite	Data in cache is not subject to any age or time constraints. The data is flushed based on other criteria that are managed by the controller.

Valid values for the cache flush modifier are listed in this table.

# **Cache Without Battery Enabled**

Write caching without batteries enables write caching to continue if the controller batteries are completely discharged, not fully charged, or not present. If you set this parameter to TRUE without an uninterruptible power supply (UPS) or other backup power source, you can lose data if the power to the storage subsystem fails. This parameter has no effect if write caching is disabled.

# **Modification Priority**

Modification priority defines the amount of system resources that are used when modifying logical drive properties. If you select the highest priority level, the logical drive modification uses the most system resources, which decreases the performance for host data transfers.

# **Cache Read Prefetch**

The **cacheReadPrefetch** parameter enables the controller to copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the drive into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage subsystem that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the **cacheReadPrefetch** parameter are TRUE or FALSE.

### Segment Size

The size of a segment determines how many data blocks that the controller writes on a single drive in a logical drive before writing data on the next drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single drive for a single request leaves other drives available to simultaneously service other requests.

If the logical drive is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of drives in the subsystem that are used for data transfers.) In this case, multiple drives are used for the same request, but each drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of drives that are required to satisfy a data transfer request.

# Set Logical Drive Mapping

This command defines the logical unit number (LUN) mapping between a logical drive a host or host group. This command is applicable to logical drives in either an subsystem or disk pools.

**Note:** You cannot use this command for an Enhanced FlashCopy logical drive that is used in online logical drive copy.

#### Syntax

```
set (logicaldrive ["logicaldriveName"] | logicaldrive <wwID> | accessLogicalDrive)
logicalUnitNumber=LUN
(host="hostName" |
hostGroup=("hostGroupName" | defaultGroup)
```

# **Parameters**

Parameter	Description
logicaldrive	The name of the specific logical drive for which you want to define mapping. Enclose the logical drive name in double quotation marks (" ") inside of square brackets ([ ]). The World Wide Identifier (WWID) of the logical drive for which you are setting properties. You can use the WWID instead of the logical drive name to identify the logical drive. Enclose the WWID in angle brackets (< >).
accessLogicalDrive	The logical unit number for the access logical drive. The logical unit number is the only property that you can set for the access logical drive.
logicalUnitNumber	The logical unit number that you want to use to map to a specific host. This parameter also assigns the host to a host group.
host	The name of the host to which the logical drive is mapped. Enclose the host name in double quotation marks (" ").
hostGroup	The name of the host group to which the logical drive is mapped. Enclose the host group name in double quotation marks (" "). defaultGroup is the host group that contains the host to which the logical drive is mapped.

### Notes

A host group is an optional topological element that you can define if you want to designate a collection of hosts that share access to the same logical drives. The host group is a logical entity. Define a host group only if you have two or more hosts that can share access to the same logical drives.

You can use any combination of alphanumeric characters, hyphens, and underscores for the names. Names can have a maximum of 30 characters.

The access logical drive is the logical drive in a SAN environment that is used for in-band communication between the storage management software and the storage subsystem controller. This logical drive uses a LUN address and consumes 20 MB of storage space that is not available for application data storage. An access logical drive is required only for in-band managed storage subsystems. If you specify the **accessLogicalDrive** parameter, the only property you can set is the **logicalUnitNumber** parameter.

# Set read-only Enhanced FlashCopy Logical Drive to a Read/Write Logical Drive

This command changes an Enhanced FlashCopy logical drive that is a read<sup>only</sup> logical drive to an Enhanced FlashCopy logical drive that is read/write logical drive. You can also use this command to define the repository for the read/write logical drive.

### **Syntax**

```
set enhancedFlashCopyLogicalDrive ["
enhancedFlashCopyImageLogicalDriveName"] convertReadOnlyToReadWrite
[repositoryLogicalDrive="repos_xxxx" |
repositoryFullLimit=percentValue]
```

# **Parameters**

Parameter	Description
enhancedFlashCopyLogicalDrive	The name of the Enhanced FlashCopy logical drive that you want to change from read-only to read/write. Enclose the Enhanced FlashCopy logical drive identifier in double quotation marks (" ") inside of square brackets ([ ]).
repositoryLogicalDrive	The name of the repository logical drive has the Enhanced FlashCopy image. All repository identifiers have this form: repos_xxxx where xxxx represents a four digit numerical value. Enclose the Enhanced FlashCopy logical drive name in double quotation marks (" ").
repositoryFullLimit	The percentage of repository logical drive capacity at which you receive a warning that the Enhanced FlashCopy image repository logical drive is nearing full. Use integer values. For example, a value of 70 means 70 percent. The default value is 75.

### Notes

The repository logical drive name is automatically created by the storage management software and the firmware when you create a new Enhanced FlashCopy group. You cannot rename the repository logical drive because renaming the repository logical drive will break the linkage with the Enhanced FlashCopy images.

# Set Remote Mirror

This command defines the properties for a remote-mirror pair.

# **Syntax**

```
set remoteMirror (localLogicalDrive [logicalDrive] |
localLogicalDrives [logicalDriveName1 ... logicalDriveNameN])
role=(primary | secondary)
```

```
[force=(TRUE | FALSE)]
syncPriority=(highest | high | medium | low | lowest)
autoResync=(enabled | disabled)
writeOrder=(preserved | notPreserved)
writeMode=(synchronous | asynchronous)
```

#### **Parameters**

Parameter	Description
localLogicalDrive or localLogicalDrives	The name of the primary logical drive for which you want to define properties. You can enter more than one primary logical drive name. Enclose the primary logical drive name in square brackets ([]). If the primary logical drive name has special characters, you also must enclose the primary logical drive name in double quotation marks ("").
role	The setting for the logical drive to act as the primary logical drive or the secondary logical drive. To define the logical drive as the primary logical drive, set this parameter to primary. To define the logical drive as the secondary logical drive, set this parameter to secondary. This parameter applies only when the logical drive is part of a mirror relationship.
force	The role reversal is forced if the communications link between the storage subsystems is down and promotion or demotion on the local side results in a dual-primary condition or a dual-secondary condition. To force a role reversal, set this parameter to TRUE. The default value is FALSE.
syncPriority	The priority that full synchronization has relative to host I/O activity. Valid values are highest, high, medium, low, or lowest.
autoResync	<ul> <li>The settings for automatic resynchronization between the primary logical drives and the secondary logical drives of a remote-mirror pair. This parameter has these values:</li> <li>enabled – Automatic resynchronization is turned on. You do not need to do anything further to resynchronize the primary logical drive and the secondary logical drive.</li> <li>disabled – Automatic resynchronization is turned off. To resynchronize the primary logical drives and the secondary logical drive, you must run the <b>resume</b></li> </ul>
	remoteMirror command.
writeUrder	This parameter defines write order for data transmission between the primary logical drive and the secondary logical drive. Valid values are preserved or notPreserved.
writeMode	This parameter defines how the primary logical drive writes to the secondary logical drive. Valid values are synchronous or asynchronous.

# Notes

When you use this command, you can specify one or more of the optional parameters.

Synchronization priority defines the amount of system resources that are used to synchronize the data between the primary logical drives and the secondary logical drives of a mirror relationship. If you select the highest priority level, the data synchronization uses the most system resources to perform the full synchronization, which decreases the performance for host data transfers.

The **writeOrder** parameter applies only to Enhanced Global Mirrors and makes them become part of a consistency group. Setting the **writeOrder** parameter to preserved causes the remote-mirror pair to transmit data from the primary logical drive to the secondary logical drive in the same order as the host writes to the primary logical drive. In the event of a transmission link failure, the data is buffered until a full synchronization can occur. This action can require additional system overhead to maintain the buffered data, which slows operations. Setting the **writeOrder** parameter to notPreserved frees the system from having to maintain data in a buffer, but it requires forcing a full synchronization to make sure that the secondary logical drive has the same data as the primary logical drive.

# Set Session

This command defines how you want the current script engine session to run. For the purpose of this command a session is the duration for the running of the commands. This command does not permanently set the parameters for the storage subsystem.

#### Syntax

set session errorAction=(stop | continue)
password="storageSubsystemPassword"
performanceMonitorInterval=intervalValue
performanceMonitorIterations=iterationValue

Parameter	Description
errorAction	How the session responds if an error is encountered during processing. You can choose to stop the session if an error is encountered, or you can continue the session after encountering an error. The default value is stop. (This parameter defines the action for execution errors, not syntax errors. Some error conditions might override the continue value.)
password	The password for the storage subsystem. Enclose the password in double quotation marks ("").
performanceMonitorInterval	The frequency of gathering performance data. Enter an integer value for the polling interval, in seconds, for which you want to capture data. The range of values is 3 to 3600 seconds. The default value is 5 seconds.
performanceMonitorIterations	The number of samples to capture. Enter an integer value. The range of values for samples captured is 1 to 3600. The default value is 5.

When you use this command, you can specify one or more of the optional parameters.

Passwords are stored on each storage subsystem in a management domain. If a password was not previously set, you do not need a password. The password can be any combination of alphanumeric characters with a maximum of 30 characters. (You can define a storage subsystem password by using the **set storageSubsystem** command.)

The polling interval and the number of iterations that you specify remain in effect until you end the session. After you end the session, the polling interval and the number of iterations return to their default values.

# Set Storage Subsystem

This command defines the properties of the storage subsystem.

# **Syntax**

```
set storageSubsystem alarm=(enable | disable | mute)|
{autoSupportConfig (enable | disable) |
cacheBlockSize=cacheBlockSizeValue
cacheFlushStart=cacheFlushStartSize
cacheFlushStop=cacheFlushStopSize
defaultHostType=("hostTypeName" | hostTypeIdentifier)
failoverAlertDelay=delayValue |
mediaScanRate=(disabled | 1-30) |
password="password" |
userLabel="storageSubsystemName"
isnsRegistration=(TRUE | FALSE))
```

Parameter	Description
alarm	The setting for the audible alarm. This parameter has these values:
	• enable – The audible alarm is turned on and sounds if a fault occurs.
	<ul> <li>disable – The audible alarm is turned off and does not sound if a fault occurs.</li> </ul>
	<ul> <li>mute – The audible alarm is turned off if it is sounding.</li> </ul>
	If another fault occurs after you set the audible alarm to mute, the audible alarm sounds again.
autoSupportConfig	The setting for automatically collecting support data. This parameter has these values:
	• enable – Turns on the collection of support data
	<ul> <li>disable – Turns off the collection of support data</li> </ul>
cacheBlockSize	The cache block size that is used by the controller for managing the cache. Valid values are 4 (4 KB), 8 (8 KB), 16 (16 KB), or 32 (32 KB).
cacheFlushStart	The percentage of unwritten data in the cache that causes a cache flush. Use integer values from 0 to 100 to define the percentage. The default value is 80.

Parameter	Description
cacheFlushStop	The percentage of free data area in the cache that stops a cache flush in progress. Use integer values from 0 to 100 to define the percentage. This value must be less than the value of the <b>cacheFlushStart</b> parameter. <b>Note:</b> This is not the percentage of unwritten data in the cache. Instead, it is the percentage of the unused area in the cache.
defaultHostType	The default host type of any unconfigured HBA host port to which the controllers are connected. To generate a list of valid host types for the storage subsystem, run the <b>show storageSubsystem hostTypeTable</b> command. Host types are identified by a name or a numerical index. Enclose the host type name in double quotation marks (""). Do not enclose the host type numerical identifier in double quotation marks.
failoverAlertDelay	The failover alert delay time in minutes. The valid values for the delay time are 0 to 60 minutes. The default value is 5.
mediaScanRate	The number of days over which the media scan runs. Valid values are disabled, which turns off the media scan, or 1 day to 30 days, where 1 day is the fastest scan rate, and 30 days is the slowest scan rate. A value other than disabled or 1 to 30 does not allow the media scan to function.
password	The password for the storage subsystem. Enclose the password in double quotation marks ("").
userLabel	The name for the storage subsystem. Enclose the storage subsystem name in double quotation marks ("").
isnsRegistration	The means of listing the iSCSI target on the iSNS server. Set the parameter to TRUE to list it.

When you use this command, you can specify one or more of the optional parameters.

# Auto Support Data

When enabled, the **set storageSubsystem autoSupportConfig** command causes all configuration and state information for the storage subsystem to be returned each time a critical Major Event Log (MEL) event is detected. The configuration and state information is returned in the form of an object graph. The object graph contains all relevant logical and physical objects and their associated state information for the storage subsystem.

The **set storageSubsystem autoSupportConfig** command collects configuration and state information in this way:

• Automatic collection of the configuration and state information occurs every 72 hours. The configuration and state information is saved to the storage subsystem zip archive file. The archive file has a time stamp that is used to manage the archive files.

- Two storage subsystem zip archive files are maintained for each storage subsystem. The zip archive files are kept on a disk drive. After the 72-hour time period is exceeded, the oldest archive file is always overwritten during the new cycle.
- After you enable automatic collection of the configuration and state information using this command, an initial collection of information starts. Collecting information after the you issue the command makes sure that one archive file is available and starts the time stamp cycle.

You can run the **set storageSubsystem autoSupportConfig** command on more than one storage subsystem.

# **Cache Block Size**

When you define cache block sizes, use the 4-KB cache block size for storage subsystems that require I/O streams that are typically small and random. Use the 8-KB cache block size when the majority of your I/O streams are larger than 4 KB but smaller than 8 KB. Use the 16-KB cache block size or the 32-KB cache block size for storage subsystems that require large data transfer, sequential, or high-bandwidth applications.

This parameter defines the supported cache block size for all of the logical drives in the storage subsystem. Not all controller types support all cache block sizes. For redundant configurations, this parameter includes all of the logical drives that are owned by both controllers within the storage subsystem.

# **Cache Flush Start and Cache Flush Stop**

When you define values to start a cache flush, a value that is too low increases the chance that data needed for a host read is not in the cache. A low value also increases the number of disk drive writes that are necessary to maintain the cache level, which increases system overhead and decreases performance.

When setting storage subsystem cache settings, the value of **cacheFlushStart** must be greater than or equal to the value of **cacheFlushStop**. For example, if the value of **cacheFlushStart** is set to 80, set the value of **cacheFlushStop** within the range of 0 to 80.

When you define values to stop cache flush, the higher the value, the higher the chance that the data for a host read requires a drive read, rather than a reading from the cache.

# **Default Host Type**

When you define host types, if Storage Partitioning is enabled, the default host type affects only those logical drives that are mapped in the default group. If Storage Partitioning is not enabled, all of the hosts that are attached to the storage subsystem must run the same operating system and be compatible with the default host type.

# Media Scan Rate

Media scan runs on all of the logical drives in the storage subsystem that have Optimal status, do not have modification operations in progress, and have the **mediaScanRate** parameter enabled. Use the **set logicalDrive** command to enable or disable the **mediaScanRate** parameter.

# Password

Passwords are stored on each storage subsystem. For best protection, the password must:

- Be between eight and 32 characters long.
- Contain at least one uppercase letter.
- · Contain at least one lowercase letter.
- Contain at least one number.
- Contain at least one non-alphanumeric character, for example, <> @ +.

**Note:** If you are using Full Disk Encryption (FDE) disk drives in your storage subsystem, you must use these criteria for your storage subsystem password.

If you are using full disk encryption drives in your storage subsystem, you must use these criteria for your storage subsystem password.

**Important:** You must set a password for your storage subsystem before you can create a security key for encrypted Full Disk Encryption (FDE) disk drives.

# Set Storage Subsystem Enclosure Positions

This command defines the position of the enclosures in a storage subsystem. You must include all of the enclosures in the storage subsystem when you enter this command.

#### Syntax

```
set storageSubsystem enclosurePositions=(controller |
enclosureID ... enclosureIDn)
```

#### Parameter

Parameter	Description
enclosurePositions	A list of all of the enclosure IDs. The sequence of the enclosure IDs in the list defines the positions for the controller module and the expansion drawers in a storage subsystem. Valid values are 0 to 99. Enter the enclosure ID values separated by a space. Enclose the list of enclosure ID values in parentheses. For storage subsystems where the controller module has a predefined identifier that is not in the range of valid enclosure position values, use the controller value.

#### Notes

This command defines the position of a enclosure in a storage subsystem by the position of the enclosure ID in the **enclosurePositions** list. For example, if you have a controller module with an ID set to 84 and expansion drawers with IDs set to 1, 12, and 50, the **enclosurePositions** sequence (84 1 12 50) places the controller module in the first position, expansion drawer 1 in the second position, expansion drawer 12 in the third position, and expansion drawer 50 in the fourth position. The **enclosurePositions** sequence (1 84 50 12) places the controller module in the second position, expansion drawer 1 in the first position, expansion drawer 50 in the fourth position, expansion drawer 1 in the first position, expansion drawer 50 in the third position, and expansion drawer 12 in the fourth position.

You must include all of the enclosures in the storage subsystem in the list defined by the **enclosurePositions** parameter. If the number of enclosures in the list does not match the total number of enclosures in the storage subsystem, an error message appears.

# Set Storage Subsystem ICMP Response

This command returns the default values for negotiable settings for sessions and connections, which represent the starting point for the storage subsystem for negotiations.

### **Syntax**

set storageSubsystem icmpPingResponse=(TRUE | FALSE)

### Parameter

Parameter	Description
icmpPingResponse	This parameter turns on or turns off Echo Request messages. Set the parameter to TRUE to turn on Echo Request messages. Set the parameter to FALSE to turn off Echo Request messages.

### Notes

The Internet Control Message Protocol (ICMP) is used by operating systems in a network to send error messages, test packets, and informational messages related to the IP, such as a requested service is not available or that a host or router could not be reached. The ICMP response command sends ICMP Echo Request messages and receives Echo Response messages to determine if a host is reachable and how long packets take to get to and from that host.

# Set Storage Subsystem iSNS Server IPv4 Address

This command sets the configuration method and address for an IPv4 Internet Storage Name Service (iSNS).

### **Syntax**

set storageSubsystem isnsIPV4ConfigurationMethod=[static | dhcp]
isnsIPV4Address=ip-address

Parameters	Description
isnsIPv4ConfigurationMethod	The method that you want to use to define the iSNS server configuration. You can enter the IP address for the IPv4 iSNS servers by selecting static. For IPv4, you can choose to have a Dynamic Host Configuration Protocol (DHCP) server select the iSNS server IP address by entering dhcp. To enable DCHP, you must set the <b>isnsIPv4Address</b> parameter to 0.0.0.

Parameters	Description
isnsIPv4Address	The IP address that you want to use for the iSNS server. Use this parameter with the static value for IPv4 configurations. If you choose to have a DHCP server set the IP address for an IPv4 Internet iSNS server, you must set the <b>isnsIPv4Address</b> parameter to 0.0.0.0.

The iSNS protocol facilitates the automated discovery, management, and configuration of iSCSI devices and Fibre Channel devices on a TCP/IP network. iSNS provides intelligent storage discovery and management services comparable to those found in Fibre Channel networks, which allow a commodity IP network to function in a similar capacity as a storage area network. iSNS also facilitates a seamless integration of IP networks and Fibre Channel networks, due to its ability to emulate Fibre Channel fabric services and manage both iSCSI devices and Fibre Channel devices.

The DHCP server passes configuration parameters, such as network addresses, to IP nodes. DHCP enables a client to acquire all of the IP configuration parameters that it needs to operate. DHCP lets you automatically allocate reusable network addresses.

# Set Storage Subsystem iSNS Server IPv6 Address

This command sets the IPv6 address for the iSNS server.

#### **Syntax**

set storageSubsystem isnsIPV6Address=ip-address

#### **Parameters**

Parameters	Description
isnsIPv6Address	The IPv6 address that you want to use for the iSNS server.

#### **Notes**

The iSNS protocol facilitates the automated discovery, management, and configuration of iSCSI and Fibre Channel devices on a TCP/IP network. iSNS provides intelligent storage discovery and management services comparable to those found in Fibre Channel networks, permitting a commodity IP network to function in a similar capacity as a storage area network. iSNS also facilitates a seamless integration of IP networks and Fibre Channel networks, due to its ability to emulate Fibre Channel fabric services, and manage both iSCSI devices and Fibre Channel devices. iSNS thereby provides value in any storage network that has iSCSI devices, Fibre Channel devices, or any combination thereof.

# Set Storage Subsystem iSNS Server Listening Port

This command sets the iSNS server listening port.

set storageSubsystem isnsListeningPort=listeningPortIPAddress

### Parameter

Parameter	Description
isnsListeningPort	The IP address that you want to use for the iSNS server listening port. The range of values for the listening port is 49152 to 65535. The default value is 3205.

#### Notes

A listening port resides on the database server and is responsible for these activities:

- · Listening (monitoring) for incoming client connection requests
- Managing the traffic to the server

When a client requests a network session with a server, a listener receives the actual request. If the client information matches the listener information, then the listener grants a connection to the database server.

# Set Storage Subsystem iSNS Server Refresh

This command refreshes the network address information for the iSNS server. This command is valid for only IPv4.

**Note:** With controller firmware 10.83.xx.xx, use the command **Storage Subsystem iSNS Server Refresh** instead.

#### Syntax

set storageSubsystem isnsServerRefresh

#### **Parameters**

None.

#### Notes

If the DHCP server is not operating at full capability, or if the DHCP server is unresponsive, the refresh operation can take between two and three minutes to complete.

The **set storageSubsystem isnsServerRefresh** command returns an error if you did not set the configuration method to DHCP. To set the configuration method to DHCP, use the **set storageSubsystem isnsIPV4ConfigurationMethod** command.

# Set Storage Subsystem Learn Cycle

This command sets the learn cycle for the battery backup unit, which enables the storage management software to predict the remaining battery life. Learn cycles run at set intervals, and they store the results for software analysis.

set storageSubsystem learnCycleDate
(daysToNextLearnCycle=numberOfDays |
day=dayOfTheWeek) time=HH:MM

### **Parameters**

Parameter	Description
daysToNextLearnCycle	Valid values are θ through 7, where 0 is immediately and 7 is in seven days. The <b>daysToNextLearnCycle</b> parameter takes place up to seven days after the next scheduled learn cycle.
day	Valid values for the day parameter include the days of the week (Sunday, Monday, Tuesday, Wednesday, Thursday, Friday, and Saturday). Setting the day causes the next learn cycle to be scheduled on the specified day, after the currently scheduled learn cycle.
time	The time in 24-hour format; for example 8:00 a.m. is entered as 08:00. Nine o'clock p.m. is entered as 21:00, and 9:30 p.m. is entered as 21:30.

### Notes

You can set the learn cycle to occur only once during a seven-day period.

The **time** parameter selects a specific time that you want to run the learn cycle. If a value is not entered, the command uses a default value of 00:00 (midnight).

If the day and time specified are in the past, the next learn cycle takes place on the next possible day specified.

# Set Storage Subsystem Redundancy Mode

This command sets the redundancy mode of the storage subsystem to either simplex or duplex.

### **Syntax**

set storageSubsystem redundancyMode=(simplex | duplex)

### Parameter

Parameter	Description
redundancyMode	Use simplex mode when you have a single controller. Use duplex mode when you have two controllers.

# Set Storage Subsystem Security Key

Use this command to set the security key that is used throughout the storage subsystem to implement the Full Disk Encryption (FDE) premium feature. When any security-capable drive in the storage subsystem is assigned to a secured subsystem, that drive is security-enabled using the security key. Before you can set the security key, you must use the create **storageSubsystem securityKey** command to create the security key.

set storageSubsystem securityKey

# **Parameters**

None.

### Notes

Security-capable disk drives have hardware to accelerate cryptographic processing and each has a unique disk drive key. A security-capable disk drive behaves like any other disk drive until it is added to a secured subsystem. When a security-capable disk drive is added to a secured subsystem, it becomes security-enabled.

Whenever a security-enabled disk drive is powered on, it requires the correct security key from the controller before it can read or write data. So, a security-enabled drive uses two keys: the disk drive key that encrypts and decrypts the data and the security key that authorizes the encryption and decryption processes. The **set storageSubsystem securityKey** command commits the security key to all of the controllers and security-enabled disk drives in the storage subsystem. The Full Disk Encryption (FDE) feature ensures that if a security-enabled disk drive is physically removed from a storage subsystem, its data cannot be read by any other device unless the security key is known.

# Set Storage Subsystem Time

This command sets the clocks on both controllers in a storage subsystem by synchronizing the controller clocks with the clock of the host from which you run this command.

### **Syntax**

set storageSubsystem time

### Parameters

None.

# Set Enhanced Remote Mirroring

This command defines the properties for a remote-mirrored pair.

**Note:** In previous versions of this command the feature identifier was remoteMirror. This feature identifier is no longer valid and is replaced by syncMirror.

### **Syntax**

```
set syncMirror (localVolume [volumeName] |
localVolumes [volumeName1 ... volumeNameN])
role=(primary | secondary)
[force=(TRUE | FALSE)]
syncPriority=(highest | high | medium | low | lowest)
autoResync=(enabled | disabled)
writeOrder=(preserved | notPreserved)
writeMode=(synchronous | asynchronous)
```

Parameter	Description
localVolume or localVolumes	The name of the primary volume for which you want to define properties. You can enter more than one primary volume name. Enclose the primary volume name in square brackets ([]). If the primary volume name has special characters, you also must enclose the primary volume name in double quotation marks (" ").
role	The setting for the volume to act as the primary volume or the secondary volume. To define the volume as the primary volume, set this parameter to primary. To define the volume as the secondary volume, set this parameter to secondary. This parameter applies only when the volume is part of a mirror relationship.
force	The role reversal is forced if the communications link between the storage subsystems is down and promotion or demotion on the local side results in a dual-primary condition or a dual-secondary condition. To force a role reversal, set this parameter to TRUE. The default value is FALSE.
syncPriority	The priority that full synchronization has relative to host I/O activity. Valid values are highest, high, medium, low, or lowest.
autoResync	<ul> <li>The settings for automatic resynchronization between the primary volumes and the secondary volumes of a remote-mirrored pair. This parameter has these values:</li> <li>enabled – Automatic resynchronization is turned on. You do not need to do anything further to resynchronize the primary volume and the secondary volume.</li> <li>disabled – Automatic resynchronize the primary volumes and the secondary volume, you must run the resume syncMirror command.</li> </ul>
writeOrder	This parameter defines write order for data transmission between the primary volume and the secondary volume. Valid values are preserved or notPreserved.
writeMode	This parameter defines how the primary volume writes to the secondary volume. Valid values are synchronous or asynchronous.

When you use this command, you can specify one or more of the optional parameters. Synchronization priority defines the amount of system resources that are used to synchronize the data between the primary volumes and the secondary volumes of a mirror relationship. If you select the highest priority level, the data synchronization uses the most system resources to perform the full synchronization, which decreases the performance for host data transfers.

The writeOrderparameter applies only to Enhanced Global Mirrors and makes them become part of a consistency group. Setting the writeOrder parameter to preserved causes the remote-mirrored pair to transmit data from the primary volume to the secondary volume in the same order as the host writes to the primary volume. In the event of a transmission link failure, the data is buffered until a full synchronization can occur. This action can require additional system overhead to maintain the buffered data, which slows operations. Setting the writeOrderparameter to notPreserved frees the system from having to maintain data in a buffer, but it requires forcing a full synchronization to make sure that the secondary volume has the same data as the primary volume.

# **Minimum Firmware Level**

```
6.10
```

# Set Thin Logical Drive Attributes

This command defines the properties for a thin logical drive. You can use the parameters to define properties for one or more thin logical drives.

# Syntax

```
set (logicaldrive ["
logicaldriveName"] | logicaldrives ["logicaldriveName1" ...
"logicaldriveNameN"] | logicaldrive <wwID>)
[newCapacity=capacity-spec |repositoryMaxCapacity=capacity-spec |
repositoryExpansionPolicy=(automatic|manual) |
warningThresholdPercent=integer-literal |
addRepositoryCapacity=capacity-spec]
```

Parameter	Description
logicaldrive or logicaldrives	The name of the specific logical drive for which you want to define properties. (You can enter more than one logical drive name if you use the <b>logicaldrives</b> parameter.) Enclose the logical drive name in double quotation marks (" ") inside of square brackets ([ ]).
logicaldrive	The World Wide Identifier (WWID) of the logical drive for which you are setting properties. You can use the WWID instead of the logical drive name to identify the logical drive. Enclose the WWID in angle brackets (< >).

Parameter	Description
newCapacity	Increase the virtual capacity of the thin logical drive. This is the value that the logical drive will report to a host that is mapped to the logical drive. Values smaller or equal to the existing capacity will cause an error.
repositoryMaxCapacity	Set the maximum capacity of the repository logical drive. The value must not be smaller that the physical capacity of the repository logical drive. If the new value results in a reduction in capacity to a level below the warning threshold, the command will produce an error.
repositoryExpansionPolicy	Set the expansion policy to automatic or manual. When you change the policy from automatic to manual, the maximum capacity value (quota) changes to the physical capacity of the repository logical drive.
warningThresholdPercent	Adjust the warning threshold. If a change to this value would cause a warning condition, the command will fail and produce an error.
addRepositoryCapacity	Allocate capacity from the free extent of the disk pool. If insufficient space is available the command fails.

When you use this command, you can specify one or more of the optional parameters.

# Set Storage Subsystem Unnamed Discovery Session

This command enables the storage subsystem to participate in unnamed discovery sessions.

#### **Syntax**

set storageSubsystem unnamedDiscoverySession=(TRUE | FALSE)

### Parameter

Parameter	Description
unnamedDiscoverySession	This parameter turns on or turns off unnamed discovery sessions. Set the parameter to TRUE to turn on unnamed discovery sessions. Set the parameter to FALSE to turn off unnamed discovery sessions.

### Notes

Discovery is the process where initiators determine the targets that are available. Discovery occurs at power-on/initialization and also if the bus topology changes, for example, if an extra device is added.

An unnamed discovery session is established without specifying a target ID in the login request. For unnamed discovery sessions, neither the target ID nor the target portal group ID are available to the targets.

# Set VolumeCopy

This command defines the properties for a VolumeCopy pair.

### **Syntax**

```
set VolumeCopy target [targetName]
[source [sourceName]]
copyPriority=(highest | high | medium | low | lowest)
targetReadOnlyEnabled=(TRUE | FALSE)
copyType=(online | offline)
```

### **Parameters**

Parameter	Description
target	The name of the target logical drive for which you want to define properties. Enclose the target logical drive name in square brackets ([]). If the target logical drive name has special characters, you also must enclose the target logical drive name in double quotation marks ("").
source	The name of the source logical drive for which you want to define properties. Enclose the source logical drive name in square brackets ([]). If the source logical drive name has special characters, you also must enclose the source logical drive name in double quotation marks ("").
copyPriority	The priority that the VolumeCopy has relative to host I/O activity. Valid values are highest, high, medium, low, or lowest.
targetRead0n1yEnab1ed	The setting so that you can write to the target logical drive or only read from the target logical drive. To write to the target logical drive, set this parameter to FALSE. To prevent writing to the target logical drive, set this parameter to TRUE.
соруТуре	Use this parameter to identify that a VolumeCopy has a FlashCopy . If the VolumeCopy has a FlashCopy , set this parameter to online. If the VolumeCopy does not have a FlashCopy, set this parameter to offline.

### Notes

When you use this command, you can specify one or more of the optional parameters.

# Show subsystem

This command returns this information about an subsystem:

- The status (online or offline)
- The disk drive type (Fibre Channel, SATA, or SAS)
- Enclosure loss protection (yes or no)

- The current owner (the controller in slot A or the controller in slot B)
- The associated logical drives and free capacity
- The associated disk drives

show subsystem [subsystemName]

#### **Parameter**

Parameter	Description
subsystem	The alphanumeric identifier of the subsystem (including - and _) for which you want to show information. Enclose the subsystem identifier in square brackets ([]).

# Show subsystem Export Dependencies

This command shows a list of dependencies for the disk drives in an subsystem that you want to move from one storage subsystem to a second storage subsystem.

#### Syntax

show subsystem [subsystemName] exportDependencies

#### Parameter

Parameter	Description
subsystem	The alphanumeric identifier (including - and _) of the subsystem for which you want to show export dependencies. Enclose the subsystem identifier in square brackets ([]).

#### Notes

This command spins up the disk drives in an subsystem, reads the DACstore, and shows a list of import dependencies for the subsystem. The subsystem must be in an Exported state or a Forced state.

# Show subsystem Import Dependencies

This command shows a list of dependencies for the disk drives in an subsystem that you want to move from one storage subsystem to a second storage subsystem.

#### Syntax

show subsystem [subsystemName] importDependencies
[cancelImport=(TRUE | FALSE)]

Parameter	Description
subsystem	The alphanumeric identifier (including - and _) of the subsystem for which you want to show import dependencies. Enclose the subsystem identifier in square brackets ([]).

Parameter	Description
cancelImport	The setting to spin the disk drives back down after the subsystem dependencies have been read. To spin down the disk drives, set this parameter to TRUE. To let the disk drives stay spinning, set this parameter to FALSE.

This command returns the dependencies of a specific subsystem, which must be in an Exported state or a Forced state. If a decision is made to retain the listed dependencies, then the **cancelImport** parameter can be enforced to spin the disk drives back down.

The **show subsystem importDependencies** command must be run before the **start subsystem import** command.

# Show Enhanced Global Mirror Group

This command displays configuration information for one or more Enhanced Global Mirror Groups. This command also displays the Enhanced Global Mirrored Pairs associated with each Enhanced Global Mirror Group, including incomplete Enhanced Global Mirrored Pairs.

You also can use this command to show the progress of periodic data synchronization on all of the mirrored pairs within the Enhanced Global Mirror Group.

#### Syntax

```
show (allenhancedGlobalMirrorGroups | enhancedGlobalMirrorGroup
["enhancedGlobalMirrorGroupName"])
[summary]
```

Parameter	Description
allenhancedGlobalMirrorGroups	Use this parameter if you want to display the properties for all Enhanced Global Mirror Groups.
enhancedGlobalMirrorGroup	Use this parameter to display the properties for an Enhanced Global Mirror Group. Enclose the Enhanced Global Mirror Group name in square brackets ([]). If the Enhanced Global Mirror Group name has special characters, you also must enclose the Enhanced Global Mirror Group name in double quotation marks ("") inside square brackets. If the Enhanced Global Mirror Group name consists only of numbers, such as "1002," you also must enclose the Enhanced Global Mirror Group name in double quotation marks ("") inside square brackets.

Parameter	Description
summary	Use this optional parameter to show a concise list of information about the synchronization progress of one or more Enhanced Global Mirror Groups.

### **Minimum Firmware Level**

7.840

# Show Enhanced Global Mirror Group synchronization progress

This command displays the progress of periodic synchronization of the Enhanced Global Mirror Group between the local and remote storage array. This command returns the progress of data synchronization on all of the mirrored pairs within the Enhanced Global Mirror Group. This command shows the progress as a percentage of data synchronization that has been completed.

**Note:** There are two types of synchronization: initial synchronization and periodic synchronization. Initial Enhanced Global Mirror Group synchronization progress is displayed in the **Long Running Operations** dialog and by running the show storageArray longRunningOperations command.

#### Syntax

show enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"]
[synchronizationProgress]

Parameter	Description
enhancedGlobalMirrorGroup	The name of the Enhanced Global Mirror Group for which you want to check the synchronization progress. Enclose the Enhanced Global Mirror Group name in square brackets ([]). If the enhanced global mirror group name has special characters, you also must enclose the Enhanced Global Mirror Group name in double quotation marks (" ") inside square brackets. If the Enhanced Global Mirror Group name consists only of numbers, such as "1002," you also must enclose the enhanced global mirror group name in double quotation marks (" ") inside square brackets.
synchronizationProgress	Use this optional parameter to display the periodic synchronization progress of the Enhanced Global Mirror Group.

# Parameters

# **Minimum Firmware Level**

7.84

# Show Cache Backup Device Diagnostic Status

This command returns the status of backup device diagnostic tests started by the **startcacheBackupDevice diagnostic** command. If the diagnostics have finished, all of the results of the diagnostic tests are shown. If the diagnostics have not finished, only the results of the diagnostic tests that finished are shown. The results of the test are shown on the terminal, or you can write the results to a file.

### **Syntax**

show cacheBackupDevice controller [(a | b)] diagnosticStatus [file="fileName"]

# **Parameters**

Parameter	Description
controller	The controller that has the cache backup device on which you are running the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
file	The name of the file that contains the result of the diagnostic tests. Enclose the file name in double quotation marks (" "). This command does not automatically append a file extension to the file name. You must add an extension when you enter the file name.

# Show Cache Memory Diagnostic Status

This command returns the status of cache memory diagnostics started by the **start controllerdiagnostic** command. If the diagnostics have finished, all of the results of the diagnostic tests are shown. If all of the diagnostics have not finished, only the results of the diagnostic tests that finished are shown.

# **Syntax**

show cacheMemory controller [(a | b)] diagnosticStatus file="fileName"

### Parameters

Parameter	Description
controller	The controller that has the cache memory on which you are running the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]).
file	The name of the file that contains the result of the diagnostic tests. Enclose the file name in double quotation marks (" "). This command does not automatically append a file extension to the file name. You must add an extension when you enter the file name.

# Show consistency group

This command returns information about one or more consistency groups.

show (consistencyGroup [consistencyGroupName] |
consistencyGroups [consistencyGroupName1 ...
consistencyGroupNameN] |
allConsistencyGroups)
[(summary | schedule)]

### **Parameters**

Parameter	Description
consistencyGroup or consistencyGroups	The name of the specific consistency group for which you are retrieving information. You can enter more than one consistency group name. Enclose the consistency group name in square brackets ([]). You can enter more than one Enhanced FlashCopy image name or sequence number. Enclose all of the Enhanced FlashCopy image names in one set of double quotation marks (" ") inside square brackets ([]). Separate each Enhanced FlashCopy image name with a white space.
allConsistencyGroups	This setting returns information about all of the consistency groups in the storage subsystem.
summary	This setting returns a concise list of information about the consistency groups.
schedul	This parameter returns information about any schedules for a consistency group.

# **Minimum Firmware Level**

7.83

# Show Consistency Group Enhanced FlashCopy Image

This command shows one Enhanced FlashCopy image or several Enhanced FlashCopy images that are in one or more Enhanced FlashCopy consistency groups.

# **Syntax**

```
show ((cgEnhancedFlashCopyImage [(cgEnhancedFlashCopyImageName |
cgEnhancedFlashCopyImageSequenceNumber)]) |
(cgEnhancedFlashCopyImages [(cgEnhancedFlashCopyImageNumbern |
cgEnhancedFlashCopyImageSequenceNumber1 ...
cgEnhancedFlashCopyImageSequenceNumber1 )) |
allCGEnhancedFlashCopyImages
[summary]
```

# Parameters

Parameter	Description
cgEnhancedFlashCopyImage or cgEnhancedFlashCopyImages	The name of the Enhanced FlashCopy image in a consistency group. The name of an Enhanced FlashCopy image is comprised of two parts:
	• The name of the consistency group
	• An identifier for the Enhanced FlashCopy image in the consistency group.
	The identifier for the Enhanced FlashCopy image can be one of these:
	• An integer value that is the sequence number of the Enhanced FlashCopy in the consistency group.
	• <b>NEWEST</b> - Use this option when you want to show the latest Enhanced FlashCopy image created in the consistency group.
	• <b>OLDEST</b> - Use this option when you want to show the earliest Enhanced FlashCopy image created in the consistency group.
	Enclose the Enhanced FlashCopy image name in double quotation marks (" ") inside square brackets ([ ]).
	You can enter more than one Enhanced FlashCopy name or sequence number. Enclose all of the Enhanced FlashCopy names in one set of double quotation marks (" ") inside square brackets ([ ]). Separate each Enhanced FlashCopy image name with a white space.
allCGEnhancedFlashCopyImages	The setting to return all of the Enhanced FlashCopy images from the consistency groups.
summary	The setting to return a concise list of information about all of the Enhanced FlashCopy images in the storage subsystem.

### Notes

The name of an Enhanced FlashCopy image has two parts separated by a colon (:):

- The identifier of the Enhanced FlashCopy consistency group
- The identifier of the Enhanced FlashCopy image

For example, if you want to show Enhanced FlashCopy image 12345 in an Enhanced FlashCopy consistency group that has the name enhancedFlashCopyCGroup1, you would use this command: show enhancedFlashCopyImage ["enhancedFlashCopyCGroup1:12345"];

To show the most recent Enhanced FlashCopy image in an Enhanced FlashCopy consistency group that has the name **enhancedFlashCopyCGroup1**, you would use this command:

show enhancedFlashCopyImage ["enhancedFlashCopyCGroup1:newest"];

To show the Enhanced FlashCopy images in several Enhanced FlashCopy consistency groups that have the names **enhancedFlashCopyCGroup1**, **enhancedFlashCopyCGroup2**, and **enhancedFlashCopyCGroup3**, you would use this command:

show enhancedFlashCopyImages ["enhancedFlashCopyCGroup1:12345 enhancedFlashCopyCGroup2:newest enhancedFlashCopyCGroup3:oldest"];

**Note:** In these examples, the Enhanced FlashCopy consistency group name is separated from the Enhanced FlashCopy image identifier by a colon (:).

# **Show Controller**

For each controller in a storage subsystem, this command returns the following information:

- The status (Online or Offline)
- The current firmware and NVSRAM configuration
- The pending firmware configuration and NVSRAM configuration (if any)
- The board ID
- The product ID
- The product revision
- The serial number
- The date of manufacture
- The cache size or the processor size
- The date and the time to which the controller is set
- The associated logical drives (including the preferred owner)
- · The Ethernet port
- The physical disk interface
- The host interface, which applies only to Fibre Channel host interfaces

#### Syntax

show (allControllers | controller [(a | b)]) [summary]

Parameter	Description
allControllers	The setting to return information about both controllers in the storage subsystem.
controller	The setting to return information about a specific controller in the storage subsystem. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]).
summary	The setting to return a concise list of information about both controllers in the storage subsystem.

# **Show Controller Diagnostic Status**

This command returns the status of controller diagnostics started by the **start controller diagnostic** command. If the diagnostics have finished, the entire results of the diagnostic tests are shown. If the diagnostic tests have not finished, only the results of the of the tests that are finished are shown. The results of the test are shown on the terminal, or you can write the results to a file.

### **Syntax**

show controller [(a | b)] diagnosticStatus [file=filename]

# **Parameters**

Parameter	Description
controller	The setting to return information about a specific controller in the storage subsystem. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]).
file	The name of the file that contains the results of the diagnostic tests. This command does not automatically append a file extension to the file name. You must add an extension when you enter the file name.

# Show Controller NVSRAM

This command returns a list of the NVSRAM byte values for the specified host type. If you do not enter the optional parameters, this command returns a list of all of the NVSRAM byte values.

#### Syntax

show (allControllers | controller [(a | b)]) NVSRAM [hostType=hostTypeIndexLabel
| host="hostName"]

### Parameters

Parameter	Description
allControllers	The setting to return information about both controllers in the storage subsystem.
controller	The setting to return information about a specific controller in the storage subsystem. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]).
hostType	The index label or number of the host type. Use the <b>show storageSubsystem hostTypeTable</b> command to generate a list of available host type identifiers.
host	The name of the host that is connected to the controllers. Enclose the host name in double quotation marks ("").

#### Notes

Use the show controller NVSRAM command to show parts of or all of the NVSRAM before using the set controller command to change the NVSRAM values. Before

making any changes to the NVSRAM, contact your IBM Technical Support representative to learn what regions of the NVSRAM you can modify.

# Show Current iSCSI Sessions

This command returns information about an iSCSI session.

#### **Syntax**

show (iscsiInitiator | iscsiTarget) [iSCSI-ID] iscsiSessions

#### Parameters

Parameter	Description
iSCSI-ID	The name of the iSCSI initiator or target for which you want to obtain session information. Enclose the iSCSI initiator name or target in square brackets ([]).

#### Notes

If you enter this command without defining any arguments, this command returns information about all of the iSCSI sessions that are currently running. To limit the information that is returned, enter a specific iSCSI initiator or a specific iSCSI target. This command then returns information about the session for only the iSCSI initiator or the iSCSI target that you named.

# **Show Disk Pool**

This command returns this information about a disk pool:

- The status (such as Optimal, online, offline)
- The total capacity
- The current owner (the controller in slot A or the controller in slot B)
- The RAID level
- The drive media (Fibre Channel, SATA, or SAS)
- The drive interface (Fibre Channel, SATA, or SAS)
- Enclosure loss protection (yes or no)
- · The associated logical drives and free capacity
- The associated drives

#### Syntax

show diskPool [diskPoolName]

Parameter	Description
diskPool	The name of the disk pool for which you are retrieving information. Enclose the disk pool name in square brackets ([]).

Use this command to show the disk pool content of the storage subsystem profile.

For disk pools, the show diskPool command returns information in this form: DETAILS

Name: 1 Status: Optimal Capacity: 67.860 GB Current owner: Controller in slot 0,1 Quality of Service (QoS) Attributes RAID level: 0 Drive media type: Disk Drive Drive interface type: Fibre Channel enclosure loss protection: No Secure Capable: No Secure: No 7 Total logical drives: 3 Standard logical drives: Repository logical drives: 4 Free Capacity: 58.645 GB Associated drives - present (in piece order) Total drives present: 1 Enclosure Slot 0 0

# **Show Disk Drive**

For each disk drive in the storage subsystem, this command returns the following information:

- The total number of disk drives
- The type of disk drive (Fibre Channel, SATA, or SAS)
- Information about the basic disk drive:
  - The enclosure location and the slot location
  - The status
  - The capacity
  - The data transfer rate
  - The product ID
  - The firmware level
- Information about the disk drive channel:
  - The enclosure location and the slot location
  - The preferred channel
  - The redundant channel
- Hot spare coverage
- Details for each disk drive

Depending on the size of your storage subsystem, this information can be several pages long. In addition, the disk drive information is returned for the **show storageSubsystem profile** command.

#### Syntax

```
show (allDrives
[DriveMediaType=(HDD | SSD | unknown | allMedia)] |
[DriveType=(fibre | SATA | SAS | PATA)]) |
Drive [enclosureID,drawerID,slotID] |
Drives [enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn])
summary
```

### **Parameters**

Parameter	Description
allDrives	The setting to return information about all of the disk drives in the storage subsystem.
DriveMediaType	The type of disk drive media for which you want to retrieve information. Valid are:
	<ul> <li>HDD – Use this option when you have hard disk drives in the expansion drawer.</li> </ul>
	• SSD – Use this option when you have solid state disk drives in the expansion drawer.
	<ul> <li>unknown – Use if you are not sure what types of disk drive media are in the expansion drawer.</li> </ul>
	<ul> <li>allMedia – Use this option when you want to use all types of disk drive media that are in the expansion drawer.</li> </ul>
DriveType	The type of disk drive for which you want to retrieve information. Valid are fibre, SATA, SAS, or PATA.
Drive or Drives	EnclosureID1,drawerID1,slotID1 enclosureIDn,drawerIDn,slotIDn is for enclosures with multiple disk drive drawers. For enclosures without disk drive drawers, use enclosureID1,slotID1 enclosureIDn,slotIDn instead. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the values in parentheses.
summary	The setting to return the status, the capacity, the data transfer rate, the product ID, and the firmware version for the specified disk drives.

#### Notes

To determine information about the type and location of all of the disk drives in the storage subsystem, use the **allDrives** parameter.

To determine the information about the Fibre Channel, SATA, SAS, PATA disk drives in the storage subsystem, use the **DriveType** parameter.

To determine the type of disk drive in a specific location, use the **Drive** parameter, and enter the enclosure ID and the slot ID for the disk drive.

The **drive** parameter supports both high-capacity drive expansion enclosures and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive

expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need only specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to 0, and specify the ID of the slot in which a drive resides.

# **Show Drive Channel Statistics**

This command shows the cumulative data transfer for the disk drive channel and error information. If the controller has automatically degraded a disk drive channel, this command also shows interval statistics. When you use this command, you can show information about one specific disk drive channel, several disk drive channels, or all disk drive channels.

### **Syntax**

show (driveChannel [(1 | 2 | 3 | 4 | 5 | 6 | 7 | 8)] | driveChannels [1 2 3 4 5 6 7 8] | allDriveChannels) stats

#### **Parameters**

Parameter	Description
driveChannel	The identifier number of the disk drive channel for which you want to show information. Valid disk drive channel values are 1, 2, 3, 4, 5, 6, 7, or 8. Enclose the disk drive channel in square brackets ([]).
driveChannels	The identifier number of the disk drive channels for which you want to show information. Valid disk drive channel values are 1, 2, 3, 4, 5, 6, 7, or 8. Enclose the disk drive channels in square brackets ([]) with the disk drive channel value separated by a space.
allDriveChannels	The identifier that selects all of the disk drive channels.

# Show Drive Download Progress

This command returns the status of firmware downloads for the disk drives that are targeted by the **download drive firmware** command or the **download storageSubsystem driveFirmware** command.

### **Syntax**

show allDrives downloadProgress

### **Parameters**

None.

When all of the firmware downloads have successfully completed, this command returns good status. If any firmware downloads fail, this command shows the firmware download status of each disk drive that was targeted. This command returns the statuses shown in this table.

Status	Definition
Successful	The downloads completed without errors.
Not Attempted	The downloads did not start.
Partial Download	The download is in progress.
Failed	The downloads completed with errors.

# Show Enhanced FlashCopy Group

This command returns this information about one or more Enhanced FlashCopy groups.

#### Syntax

```
show (allEnhancedFlashCopyGroups | enhancedFlashCopyGroup [
enhancedFlashCopyGroupName] | enhancedFlashCopyGroupS
[enhancedFlashCopyGroupName1 ... enhancedFlashCopyGroupNamen])
[summary | schedule]
```

#### **Parameters**

Parameter	Description
allEnhancedFlashCopyGroups	The parameter to return information about all of the Enhanced FlashCopy groups in the storage subsystem.
enhancedFlashCopyGroup or enhancedFlashCopyGroups	The name of the specific Enhanced FlashCopy group for which you are retrieving information. Enclose the Enhanced FlashCopy group name in square brackets ([ ]).
	If you enter more than one Enhanced FlashCopy group name, separate each name with a space. Enclose all of the Enhanced FlashCopy group names in only one set of double quotation marks (" ").
summary	The parameter to return a concise list of information about the Enhanced FlashCopy groups.
schedule	The parameter to return a concise list of information about the schedules for the Enhanced FlashCopy group copy operations.

### Notes

This command returns Enhanced FlashCopy group information as shown in this example:
FLASHCOPY GROUP DETAILS

FLASHCOPY GROUPS-----SUMMARY Total Enhanced FlashCopy Groups: 1 Total Enhanced FlashCopy Images: 0 Enabled Enhanced FlashCopy Image Schedules: 0 Status: 1 Optimal, 0 Non Optimal Name Type Status Associated Base Logical drive 2 SG 01 Standard Optimal 2 Total Repository Capacity Available Repository Capacity Enhanced FlashCopy Image Limit 10.318 GB 10.318 GB (100%) 0 Enhanced FlashCopy Images Scheduled 0 No DETAILS Enhanced FlashCopy Group "2 SG 01" Status: Optimal Standard Type: Associated base logical drive: 2 Cluster size: 65,536 bytes Repository Total repository logical drives: 3 Aggregate repository status: Optimal Total repository capacity: 10.318 GB Used capacity: 0.000 MB (0%) Available repository capacity: 10.318 GB (100%) Repository full policy: Auto-purge Enhanced FlashCopy Images Utilization alert threshold: 75%

Enhanced FlashCopy images

Total Enhanced FlashCopy images: 0 Auto-delete Enhanced FlashCopy images: Disabled Enhanced FlashCopy image schedule: Not Applicable

# Show Enhanced FlashCopy Image

This command returns information about the Enhanced FlashCopy images that a user had previously created.

#### Syntax for Showing a Specific Enhanced FlashCopy Image

```
show (enhancedFlashCopyImage ["
enhancedFlashCopyImageName"] | enhancedFlashCopyImages
["enhancedFlashCopyImageName1 ... enhancedFlashCopyImageNamen"] |
allEnhancedFlashCopyImages])
[summary]
```

# Parameters

Parameter	Description
enhancedFlashCopyImage or enhancedFlashCopyImages	The name of the Enhanced FlashCopy image. The name of an Enhanced FlashCopy image is comprised of two parts:
	• The name of the Enhanced FlashCopy group
	• An identifier for the Enhanced FlashCopy image in the Enhanced FlashCopy group.
	The identifier for the Enhanced FlashCopy image can be one of these:
	• An integer value that is the sequence number of the Enhanced FlashCopy in the Enhanced FlashCopy group.
	• <b>NEWEST</b> - Use this option when you want to show the latest Enhanced FlashCopy image created in the Enhanced FlashCopy group.
	• <b>OLDEST</b> - Use this option when you want to show the earliest Enhanced FlashCopy image created in the Enhanced FlashCopy group.
	Enclose the Enhanced FlashCopy image name in double quotation marks (" ") inside square brackets ([ ]).
	You can enter more than one Enhanced FlashCopy name or sequence number. Enclose all of the Enhanced FlashCopy names in one set of double quotation marks (" ") inside square brackets ([ ]). Separate each Enhanced FlashCopy image name with a white space.
allEnhancedFlashCopyImages	The parameter to return information about all of the Enhanced FlashCopy images in the storage subsystem.
summary	This parameter returns a concise list of information about the Enhanced FlashCopy images.

## **Notes**

The name of an Enhanced FlashCopy image has two parts separated by a colon (:):

- The identifier of the Enhanced FlashCopy group
- The identifier of the Enhanced FlashCopy image

For example, if you want to show Enhanced FlashCopy image 12345 in an Enhanced FlashCopy group that has the name **enhancedFlashCopyGroup1**, you would use this command:

show enhancedFlashCopyImage ["enhancedFlashCopyGroup1:12345"];

To show the most recent Enhanced FlashCopy image in an Enhanced FlashCopy group that has the name **enhancedFlashCopyGroup1**, you would use this command:

show enhancedFlashCopyImage ["enhancedFlashCopyGroup1:newest"];

To show the Enhanced FlashCopy images in several Enhanced FlashCopy consistency groups that has the names **enhancedFlashCopyGroup1**, **enhancedFlashCopyGroup2**, and **enhancedFlashCopyGroup3**, you would use this command:

```
show enhancedFlashCopyImages ["enhancedFlashCopyGroup1:12345
enhancedFlashCopyGroup2:newest enhancedFlashCopyGroup3:oldest"];
```

# Show Enhanced FlashCopy Logical Drives

This command returns information about one or more Enhanced FlashCopy logical drives.

### Syntax

```
show (allEnhancedFlashCopyLogicalDrives | enhancedFlashCopyLogicalDrive
[enhancedFlashCopyLogicalDriveName] |
enhancedFlashCopyLogicalDrives
[enhancedFlashCopyLogicalDriveName1 ... enhancedFlashCopyLogicalDriveNamen])
[summary]
```

### **Parameters**

Parameter	Description
allEnhancedFlashCopyLogicalDrives	The parameter to return information about all of the Enhanced FlashCopy logical drives in the storage subsystem.
enhancedFlashCopyLogicalDrive or enhancedFlashCopyLogicalDrives	The name of a specific Enhanced FlashCopy logical drive about which you are retrieving information. Enclose the Enhanced FlashCopy logical drive name in square brackets ([]). When you enter more than one Enhanced FlashCopy logical drive name, separate the names with a space. Enclose all of the names in only one set of square brackets ([]).
summary	The parameter to return a concise list of information about the Enhanced FlashCopy logical drives.

### Notes

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

This command returns Enhanced FlashCopy logical drive information as shown in this example:

```
FLASHCOPY LOGICALDRIVE(FLASHCOPY-IMAGE BASED) SUMMARY
SUMMARY
Total Enhanced FlashCopy Logical drives: 1
Most Recent Enhanced FlashCopy Logical drive: Day month date hh:mm:ss yyyy
Status: 1 Optimal, 0 Non Optimal
Name Type Status Capacity Associated Base Logical drive
2_SV_0001 Standard Optimal 3.000 GB 2
```

Enhanced FlashCopy Logical drive Timestamp Enhanced FlashCopy Image Timestamp Mode

1/23/12 6:44:31 PM IST

1/23/12 6:27:36 PM IST

Total Repository CapacityAvailable Repository Capacity1.199 GB0.125 MB (0%)

The size of your monitor determines how the information wraps and will affect how the information appears.

# Show Enhanced Remote Mirroring Logical Drive Candidates

This command returns information about the candidate logical drives on a remote storage subsystem that you can use as secondary logical drives in a Enhanced Remote Mirroring configuration.

#### Syntax

show remoteMirror candidates primary="logicalDriveName"
remoteStorageSubsystemName="storageSubsystemName"

#### **Parameters**

Parameter	Description
primary	The name of the local logical drive that you want for the primary logical drive in the remote-mirror pair. Enclose the primary logical drive name in double quotation marks ("").
remoteStorageSubsystemName	The remote storage subsystem that contains possible logical drives for a secondary logical drive. If the remote storage subsystem name has special characters, you must also enclose the remote storage subsystem name in double quotation marks ("").

# Show Enhanced Remote Mirroring Logical Drive Synchronization Progress

This command returns the progress of data synchronization between the primary logical drive and the secondary logical drive in a Enhanced Remote Mirroring configuration. This command shows the progress as a percentage of data synchronization that has been completed.

#### **Syntax**

```
show remoteMirror (localLogicalDrive ["logicalDriveName"] |
localLogicalDrives ["logicalDriveName1" ... "logicalDriveNameN"])
synchronizationProgress
```

Parameter	Description
localLogicalDrive or localLogicalDrives	The name of the primary logical drive of the remote mirror pair for which you want to check synchronization progress. Enclose the primary logical drive name in double quotation marks ("") inside of square brackets ([]).

# Show Host Interface Card Diagnostic Status

This command returns the status of running, interrupted, or completed host interface card diagnostics started by the **start host Card diagnostic** command. If the diagnostics have finished, the entire results of the diagnostic tests are shown. If the diagnostics have not finished, only the results of the tests that are finished are shown. The results of the test are shown on the terminal, or you can write the results to a file.

# **Syntax**

show hostCard controller [(a | b)] diagnosticStatus [progressOnly] [file=filename]

Parameter	Description
controller	The controller that has the host interface card on which you are running the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
progress0n1y	The <b>progress0n1y</b> parameter, shows the progress of the diagnostic test without waiting for the diagnostic tests to completely finish.
file	The name of the file that contains the results of the diagnostic tests. This command does not automatically append a file extension to the file name. You must add an extension when you enter the file name.

## **Parameters**

**Note:** The **progress0nly** parameter is useful for seeing the progress of command scripts that need to sequentially complete operations.

# **Show Host Ports**

For all of the HBA host ports that are connected to a storage subsystem, this command returns this information:

- The HBA host port identifier
- The HBA host port name
- The HBA host type

## Syntax

show allHostPorts

# Parameters

None.

## Notes

This command returns HBA host port information similar to this example:

HOST PORT IDENTIFIER	HOST PORT NAME	HOST TYPE
12:34:56:54:33:22:22:22	Jupiter1	Solaris
12:34:56:78:98:98:88:88	Pluto1	Windows 2000/Server 2003 Clustered
54:32:12:34:34:55:65:66	Undefined	Undefined

# **Show Logical Drive**

For the logical drives in a storage subsystem, this command returns the following information:

- The number of logical drives
- The name
- The status
- The capacity
- The RAID level
- The subsystem where the logical drive is located
- Details:
  - The logical drive ID
  - The subsystem ID
  - The disk drive type (Fibre Channel, SATA, SAS)
  - Enclosure loss protection
  - The preferred owner
  - The current owner
  - The segment size
  - The modification priority
  - The read cache status (enabled or disabled)
  - The write cache status (enabled or disabled)
  - The write cache without batteries status (enabled or disabled)
  - The write cache with mirroring status (enabled or disabled)
  - The flush write cache after time
  - The cache read prefetch setting (TRUE or FALSE)
  - The enable background media scan status (enabled or disabled)
  - The media scan with redundancy check status (enabled or disabled)
- The FlashCopy repository logical drives
- The mirror repository logical drives
- The FlashCopy logical drives
- The FlashCopy copies

## Syntax

show (allLogicalDrives | logicalDrive [logicalDriveName] |
logicalDrives [logicalDriveName1 ... logicalDriveNameN]) summary

Parameter	Description
allLogicalDrives	The setting to return information about all of the logical drives in the storage subsystem.
logicalDrive or logicalDrives	The name of the specific logical drive for which you are retrieving information. You can enter more than one logical drive name. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks ("").
summary	The setting to return a concise list of information about the logical drives.

**Note:** For FlashCopy logical drive copies, the show logical drive command returns information about the schedules for the FlashCopy logical drive copies. The schedule information is in this form:

Schedule State: "Active" | "Disabled" | "Completed"
Last Run Time: <mm/dd/yyyy> <hh:mm a.m. | p.m.>
Next Run Time: <mm/dd/yyyy> <hh:mm a.m. | p.m.>
Start Date: <mm/dd/yyyy>End Date: <mm/dd/yyyy> | "No End Date"
Days of Week: <Sunday - Saturday>, <Sunday - Saturday>, ....
Times for flashcopy re-create: <hh:mm a.m. | p.m.>, <hh:mm a.m. | p.m.>

# **Show Logical Drive Action Progress**

**Note:** With firmware version 7.77, the show logicaldrive actionProgress command is deprecated. Replace this command with show storageSubsystem longRunningOperations.

For a long-running operation that is currently running on a logical drive, this command returns information about the logical drive action and amount of the long-running operation that is completed. The amount of the long-running operation that is completed is shown as a percentage (for example, 25 means that 25 percent of the long-running operation is completed).

## Syntax

show logicalDrive ["logicalDriveName"] actionProgress

# **Parameters**

Parameter	Description
logicalDrive	The name of the logical drive that is running the long-running operation. Enclose the logical drive name in double quotation marks ("") inside of square brackets ([]).

# Show Performance Read Cache

This command displays information about the Performance Read Cache.

# **Syntax**

show performanceReadCache [performanceReadCacheName]

Parameter	Description
performanceReadCache	The alphanumeric identifier (including - and _) of the Performance Read Cache for which you want to get information. Enclose the identifier in square brackets ([]). If the Performance Read Cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.

This command returns the Performance Read Cache information similar to this example.

```
Performance Read Cache name: my_cache
Status:
                              Optimal
Type:
                              Read Only
I/O characteristic type:
                             File System
Maximum capacity allowed:
                             1,862.645 GB
Current capacity:
                              557.792 GB
Additional capacity allowed
                             1,304.852 GB
Drive capacities:
                              A11 278.896 GB
Quality of Service (QoS) Attributes
Security capable:
                             No
Secure:
                              No
T10PI capable: No
Associated drives:
enclosure Slot
04
0 11
Volumes using Performance Read Cache:
                                          volume test
```

### **Minimum Firmware Level**

7.84

# **Show Array**

This command returns this information about a Array:

- The status (Online or Offline)
- The drive type (Fibre Channel, SATA, or SAS)
- Enclosure loss protection (yes or no)
- The current owner (the controller in slot A or the controller in slot B)
- The associated volumes and free capacity
- The associated drives

#### **Syntax**

show volumeGroup [volumeGroupName]

#### Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier of the Array (including - and _) for which you want to show information. Enclose the Array identifier in square brackets ([]).

### Notes

This command returns Array information as shown in this example:

```
Array 1 (RAID 5)
Status: Online
Drive type: Fibre Channel
Enclosure loss protection: No
Current owner: Controller in slot A
Associated volumes and free capacities:
1 (1 GB), 1R1 (0.2 GB), Free Capacity (134.533 GB)
```

Associated drives (in piece order): Drive at Enclosure 1, Slot 14 Drive at Enclosure 1, Slot 13 Drive at Enclosure 1, Slot 12

## **Minimum Firmware Level**

6.10

# Show array export dependencies

This command shows a list of dependencies for the drives in an array that you want to move from one storage subsystem to a second storage subsystem.

### Syntax

show array [arrayName] exportDependencies

### Parameter

Parameter	Description
array	The alphanumeric identifier (including - and _) of the array for which you want to show export dependencies. Enclose the array identifier in square brackets ([]).

## Notes

This command spins up the drives in a array, reads the DACstore, and shows a list of import dependencies for the Array. The Array must be in an Exported state or a Forced state.

## **Minimum Firmware Level**

#### 7.10

# Show Array import dependencies

This command shows a list of dependencies for the drives in a Array that you want to move from one storage subsystem to a second storage subsystem.

## **Syntax**

show volumeGroup [volumeGroupName] importDependencies
[cancelImport=(TRUE | FALSE)]

Parameter	Description
volumeGroup	The alphanumeric identifier (including - and _) of the Array for which you want to show import dependencies. Enclose the Array identifier in square brackets ([]).

Parameter	Description
cancelImport	The setting to spin the drives back down after the Array dependencies have been read. To spin down the drives, set this parameter to TRUE. To let the drives stay spinning, set this parameter to FALSE.

This command returns the dependencies of a specific Array, which must be in an Exported state or a Forced state. If a decision is made to retain the listed dependencies, then the cancelImport parameter can be enforced to spin the drives back down.

You must run the show volumeGroup importDependencies command before you run the start volumeGroup import command.

### **Minimum Firmware Level**

7.10

# **Show Logical Drive Performance Statistics**

This command returns information about the performance of the logical drives in a storage subsystem.

#### Syntax

show (alllogicalDrives | logicalDrive [logicalDriveName]
logicalDrives [logicalDriveName1 ... logicalDriveNameN] performanceStats

#### **Parameters**

Parameter	Description
alllogicalDrives	The setting to return performance statistics about all of the logical drives in the storage subsystem.
logicalDrive or logicalDrives	The name of the specific logical drive for which you are retrieving performance statistics. You can enter more than one logical drive name. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks ("").

# **Show Logical Drive Reservations**

This command returns information about the logical drives that have reservations.

#### **Syntax**

show (alllogicalDrives | logicalDrive [logicalDriveName] |
logicalDrives [logicalDriveName1 ... logicalDriveNameN]) reservations

# **Parameters**

Parameter	Description
allLogicalDrives	The setting to return persistent reservation information about all of the logical drives in the storage subsystem.
logicalDrive or logicalDrives	The name of the specific logical drive for which you are retrieving persistent reservation information. You can enter more than one logical drive name. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks ("").

# Show Storage Subsystem

This command returns configuration information about the storage subsystem. The parameters return lists of values for the components and features in the storage subsystem. You can enter the command with a single parameter or more than one parameter. If you enter the command without any parameters, the entire storage subsystem profile is shown (which is the same information as if you entered the **profile** parameter).

# **Syntax**

show storageSubsystem autoSupportConfig profile batteryAge connections
defaultHostType healthStatus hostTypeTable hotSpareCoverage features time
logicalDriveDistribution summary

Parameter	Description
autoSupportConfig	The parameter to show the status of the automatic support data collection when a critical event occurs along with the location where the Support Data bundle was saved.
profile	The parameter to show all of the properties of the logical components and the physical components that comprise the storage subsystem. The information appears in several screens.
batteryAge	The parameter to show the status, the age of the battery in days, and the number of days until the battery needs to be replaced.
connections	The parameter to show a list of where the drive channel ports are located and where the drive channels are connected.
defaultHostType	The parameter to show the default host type and the host type index.
healthStatus	The parameter to show the health, logical properties, and physical component properties of the storage subsystem.
hostTypeTable	The parameter to show a table of all of the host types that are known to the controller. Each row in the table shows a host type index and the platform that the host type index represents.

Parameter	Description
hotSpareCoverage	The parameter to show information about which logical drives of the storage subsystem have hot spare coverage and which logical drives do not.
features	The parameter to show a list of the feature identifiers for all enabled features in the storage subsystem.
time	The parameter to show the current time to which both controllers in the storage subsystem are set.
logicalDriveDistribution	The parameter to show the current controller owner for each logical drive in the storage subsystem.
longRunningOperations	The parameter to show the long running operations for each subsystem and each logical drive in the storage subsystem.
	The <b>longRunningOperation</b> parameter returns this information:
	• Name of the subsystem or logical drive
	Long running operation
	• Status
	• % complete
	Time left
summary	The parameter to show a concise list of information about the storage subsystem configuration.

The **profile** parameter shows detailed information about the storage subsystem. The information appears on several screens on a display monitor. You might need to increase the size of your display buffer to see all of the information. Because this information is so detailed, you might want to save the output of this parameter to a file. To save the output to a file, run the **show storageSubsystem** command that looks like this example:

c:\...\client>smcli 123.45.67.89 -c "show storageSubsystem profile;" -o
"c:\\folder\\storageSubsystemprofile.txt"

The previous command syntax is for a host that is running a Windows operating system. The actual syntax varies depending on your operating system.

The DS3950, DS4800 (models 82, 84, and 88), DS5020, DS5100, and DS5300 controller modules do not support the **show storageSubsystem batteryAge** command.

The **batteryAge** parameter returns information in this form.

```
Battery status: Optimal
Age: 1 day(s)
Days until replacement: 718 day(s
```

The **defaultHostType** parameter returns information in this form. Default host type: Linux (Host type index 6)

The **healthStatus** parameter returns information in this form. Storage Subsystem health status = optimal The **hostTypeTable** parameter returns information in this form.

NVSRAM HOST TYPE INDEX DEFINITIONS INDEX ADT STATUS TYPE 0 Disabled Windows NT Non-Clustered (SP5 or higher) 1 (Default) Disabled Windows 2000/Server 2003 Non-Clustered 2 Disabled Solaris 3 Enabled HP-UX 4 Disabled AIX 5 Disabled Irix 6 Enabled Linux Disabled Windows NT Clustered (SP5 or higher) 7 Disabled Windows 2000/Server 2003 Clustered 8 9 Enabled Netware Non-Failover 10 Enabled PTX 11 Enabled Netware Failover 12 Enabled Solaris (with Veritas DMP)

The **hotSpareCoverage** parameter returns information in this form:

The following subsystems are not protected: 2, 1 Total hot spare drives: 0

Standby: 0 In use: 0

The **features** parameter returns information in this form:

storagePartitionMax flashcopy remoteMirror VolumeCopy

The **time** parameter returns information in this form:

Controller in Slot A Date/Time: Thu Jun 03 14:54:55 MDT 2004 Controller in Slot B Date/Time: Thu Jun 03 14:54:55 MDT 2004

The longRunningOperations parameter returns information in this form:

LOGICAL DEVICES OPERATION STATUS TIME REMAINING Logical Drive-2 Logical drive Disk Copy 10% COMPLETED 5 min

The information fields returned by the **longRunningOperations** parameter have these meanings:

- NAME is the name of a logical drive that is currently in a long running operation. The logical drive name must have the "Logical Drive" as a prefix.
- OPERATION lists the operation being performed on the subsystem or logical drive.
- % COMPLETE shows how much of the long running operation has been performed.
- STATUS can have one of these meanings:
  - Pending The long running operation has not started but will start after the current operation is completed.
  - In Progress The long running operation has started and will run until completed or stopped by user request.
- TIME LEFT indicates the duration remaining to completing the current long running operation. The time is in an "hours minute" format. If less than an hour remains, only the minutes are shown. If less than a minute remains, the message "less than a minute" is shown.

The logicalDriveDistribution parameter returns information in this form:

name•	10
nunc.	10
owner	is controller in slot: A
name:	CTL 0 Mirror Repository
owner	is controller in slot: A
name:	Mirror Repository 1
owner	is controller in slot:A
name:	20
owner	is controller in slot:A
name:	JCG_Remote_MirrorMenuTests
owner	is controller in slot:A
	name: owner name: owner name: owner name: owner

The **feature** parameter returns information about the premium features in this form:

PREMIUM FEATURE	STATUS
driveSlotLimit	Enabled (12 of 192 used)
SSDSupport	Disabled - Feature Key required
highPerformanceTier	Disabled - Feature Key required
remoteMirror	Disabled/Deactivated - Feature Key required
safeStoreExternalKeyMgr	safeStoreExternalKeyMgr
safeStoreSecurity	Disabled - Feature Key require
VolumeCopy	Disabled - Feature Key required

# Show Storage Subsystem Auto Configure

This command shows the default auto-configuration that the storage subsystem creates if you run the **autoConfigure storageSubsystem** command. If you want to determine whether the storage subsystem can support specific properties, enter the parameter for the properties when you run this command. You do not need to enter any parameters for this command to return configuration information.

## **Syntax**

```
show storageSubsystem autoConfiguration
[driveType=(fibre | SATA | SAS)
driveMediaType=(HDD | SSD | unknown | allMedia)
raidLevel=(0 | 1 | 3 | 5 | 6)
subsystemWidth=numberOfDrives
subsystemCount=numberOflogicalDriveGroups
logicalDrivesPersubsystemCount=numberOfLogicalDrivesPersubsystem
hotSpareCount=numberOfHotspares
segmentSize=segmentSizeValue
cacheReadPrefetch=(TRUE | FALSE)]
securityType=(none | capable | enabled)
```

Parameter	Description
driveType	The type of disk drives that you want to use for the storage subsystem. Valid disk drive types are fibre, SATA, or SAS. The <b>driveType</b> parameter is not required if only one type of disk drive is in the storage subsystem.

Parameter	Description
driveMediaType	<ul> <li>The type of drive media that you want to use for the mirror repository subsystem. Valid drive media are these:</li> <li>HDD – Use this option when you have hard drives in the drive expansion enclosure.</li> <li>SSD – Use this option when you have solid state drives in the drive expansion enclosure.</li> <li>unknown – Use if you are not sure what types of drive media are in the drive expansion enclosure.</li> <li>allMedia – Use this option when you want to use all types of drive media that are in the drive expansion enclosure.</li> <li>Use this parameter when you use the repositoryDriveCount parameter.</li> </ul>
	than one type of drive media in your storage subsystem.
raidLevel	The RAID level of the subsystem that contains the disk drives in the storage subsystem. Valid RAID levels are 0, 1, 3, 5, or 6.
subsystemWidth	The number of drives in an subsystem in the storage subsystem, which depends on the capacity of the drives. Use integer values.
logicalDriveGroupCount	The number of subsystems in the storage subsystem. Use integer values.
subsystemCount	The number of subsystems in the storage subsystem. Use integer values.
logicalDrivesPerGroupCount	The number of equal-capacity logical drives per subsystem. Use integer values.
hotSpareCount	The number of hot spares that you want in the storage subsystem. Use integer values.
segmentSize	The amount of data (in KB) that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. Valid values are 8, 16, 32, 64, 128, 256, or 512.
cacheReadPrefetch	The setting to turn on or turn off cache read prefetch. To turn off cache read prefetch, set this parameter to FALSE. To turn on cache read prefetch, set this parameter to TRUE.
securityType	The setting to specify the security level when creating the subsystems and all associated logical drives. These settings are valid:
	• none – The subsystem and logical drives are not secure.
	• capable – The subsystem and logical drives are capable of having security set, but security has not been enabled.
	• enabled – The subsystem and logical drives have security enabled.

**Note:** If you do not specify any properties, this command returns the RAID Level 5 candidates for each drive type. If RAID Level 5 candidates are not available, this command returns candidates for RAID Level 6, RAID Level 3, RAID Level 1, or RAID Level 0. When you specify auto configuration properties, the controllers validate that the firmware can support the properties

## **Drives and subsystem**

A subsystem is a set of drives that are logically grouped together by the controllers in the storage subsystem. The number of drives in an subsystem is a limitation of the RAID level and the controller firmware. When you create an subsystem, follow these guidelines:

- Beginning with firmware version 7.10, you can create an empty subsystem so that you can reserve the capacity for later use.
- You cannot mix drive types, such as SAS and Fibre Channel, within a single subsystem.
- The maximum number of drives in an subsystem depends on these conditions:
  - The type of controller
  - The RAID level
- RAID levels include: 0, 1, 10, 3, 5, and 6.
  - In a DS4700 Model 70 or a DS4700 Model 72 storage subsystem, a subsystem with RAID level 0 and an subsystem with RAID level 10 can have a maximum of 112 drives.
  - In a DS4800 (Models 82, 84, 88) storage subsystem, an subsystem with RAID level 0 and an subsystem with RAID level 10 can have a maximum of 224 drives.
  - An subsystem with RAID level 3, RAID level 5, or RAID level 6 cannot have more than 30 drives.
  - An subsystem with RAID level 6 must have a minimum of five drives.
  - If an subsystem with RAID level 1 has four or more drives, the storage management software automatically converts the subsystem to a RAID level 10, which is RAID level 1 + RAID level 0.
- If an subsystem contains drives that have different capacities, the overall capacity of the subsystem is based on the smallest capacity drive.
- To enable enclosure loss protection, you must create an subsystem that uses drives located in at least three drive expansion enclosures.

## Hot Spares

Hot spare drives can replace any failed drive in the storage subsystem. A hot spare drive must have capacity greater than or equal to any drive that can fail. If a hot spare drive is smaller than a failed drive, you cannot use the hot spare drive to rebuild the data from the failed drive. Hot spare drives are available only for RAID Level 1, RAID Level 3, RAID Level 5, or RAID Level 6.

## Segment Size

The size of a segment determines how many data blocks that the controller writes on a single drive in a logical drive before writing data on the next drive. Each data block stores 512 bytes of data. A data block is the smallest unit of storage. The size of a segment determines how many data blocks that it contains. For example, an 8 KB segment holds 16 data blocks. A 64 KB segment holds 128 data blocks. When you enter a value for the segment size, the value is checked against the supported values that are provided by the controller at run time. If the value that you entered is not valid, the controller returns a list of valid values. Using a single drive for a single request leaves other drives available to simultaneously service other requests. If the logical drive is in an environment where a single user is transferring large units of data (such as multimedia), performance is maximized when a single data transfer request is serviced with a single data stripe. (A data stripe is the segment size that is multiplied by the number of drives in the subsystem that are used for data transfers.) In this case, multiple drives are used for the same request, but each drive is accessed only once.

For optimal performance in a multiuser database or file system storage environment, set your segment size to minimize the number of drives that are required to satisfy a data transfer request.

## **Cache Read Prefetch**

Cache read prefetch lets the controller copy additional data blocks into cache while the controller reads and copies data blocks that are requested by the host from the drive into cache. This action increases the chance that a future request for data can be fulfilled from cache. Cache read prefetch is important for multimedia applications that use sequential data transfers. The configuration settings for the storage subsystem that you use determine the number of additional data blocks that the controller reads into cache. Valid values for the **cacheReadPrefetch** parameter are TRUE or FALSE.

# Show storage subsystem Diagnostic Data Capture

This command show details of the Diagnostic Data Capture on the controller cache, if a Diagnostic Data Capture is available.

#### Syntax

show storageArray coreDump

#### **Parameters**

None

#### Minimum Firmware Level

7.83

# Show Storage Subsystem DBM Database

This command retrieves and shows metadata for the on-board backup locations of a storage subsystem. When there are multiple backup locations, metadata is shown for each location.

#### **Syntax**

show storageSubsystem dbmDatabase

## **Parameters**

None.

# Show Storage Subsystem Diagnostic Data Capture

This command show details of the diagnostic data capture on the controller cache, if a diagnostic data capture is available.

#### Syntax

show storageSubsystem ddc

#### **Parameters**

None.

# Show Storage Subsystem Features

This command lists the feature configuration of a storage subsystem. This command returns information that shows which features are enabled, disabled, evaluation, and available to be installed.

#### Syntax

show storagesubsystem features

#### **Parameters**

None.

## Notes

This command returns the feature information in a format similar to this example.

PREMIUM FEATURE	STATUS
driveSlotLimit	Enabled (12 of 192 used)
SSDSupport	Disabled - Feature Key required
highPerformanceTier	Disabled - Feature Key required
remoteMirror	Disabled/Deactivated - Feature Key require
safeStoreExternalKeyMgr	Disabled - Feature Key required
safeStoreSecurity	Disabled - Feature Key required
VolumeCopy	Disabled - Feature Key required

# Show storage subsystem unconfigured iSCSI Initiators

This command returns a list of initiators that have been detected by the storage subsystem but are not yet configured into the storage subsystem topology.

#### Syntax

show storageArray unconfiguredIscsiInitiators

#### **Parameters**

None

## **Minimum Firmware Level**

7.10

# Show Storage Subsystem Host Topology

This command returns the storage partition topology, the host type labels, and the host type index for the host storage subsystem.

### **Syntax**

show storageSubsystem hostTopology

#### **Parameters**

None.

# Show Storage Subsystem Profile

This command list detailed information about the storage subsystem configuration.

#### Syntax

show storagesubsystem profile

### **Parameters**

Parameter	Description
profile	The parameter to show all of the properties of the logical components and the physical components that comprise the storage subsystem. The information appears in several screens.

#### Notes

The profile parameter shows detailed information about the storage subsystem. The information appears on several screens on a display monitor. You might need to increase the size of your display buffer to see all of the information. Because this information is so detailed, you might want to save the output of this parameter to a file. To save the output to a file, run the show storageSubsystem command that looks like this example.

-c "show storageSubsystem profile;" -o "c:\\folder\\storageSubsystemProfile.txt"

The previous command syntax is for a host that is running a Windows operating system. The actual syntax varies depending on your operating system.

When you save information to a file, you can use the information as a record of your configuration and as an aid during recovery.

# Show Storage Subsystem LUN Mappings

This command returns information from the storage subsystem profile about the LUN mappings in the storage subsystem. Default group LUN mappings are always shown. If you run this command without any parameters, this command returns all of the LUN mappings.

## **Syntax**

show storageSubsystem lunMappings [host ["hostName"] |
hostgroup ["hostGroupName"]]

### **Parameters**

Parameter	Description
host	The name of a specific host for which you want to see the LUN mappings. Enclose the host name in double quotation marks ("") inside of square brackets ([]).
hostGroup	The name of a specific host group for which you want to see the LUN mappings. Enclose the host group name in double quotation marks ("") inside of square brackets ([]).

# Show Storage Subsystem Negotiation Defaults

This statement returns information about connection-level settings that are subject to initiator-target negotiation.

#### Syntax

show storageSubsystem iscsiNegotiationDefaults

## **Parameters**

None.

#### Notes

Information returned includes RAID controller enclosure default settings (that is, those settings that are the starting point for negotiation) and the current active settings.

# Show Storage Subsystem Unreadable Sectors

This command returns a table of the addresses of all of the sectors in the storage subsystem that cannot be read. The table is organized with column headings for the following information:

- 1. Logical drive user label
- 2. LUN
- **3**. Accessible by (host or host group)
- 4. Date/time
- 5. Logical drive-relative logical block address (hexadecimal format 0xnnnnnn)
- 6. Disk drive location (enclosure t, slot s)

- 7. Disk drive-relative logical block address (hexadecimal format 0xnnnnnnn)
- 8. Failure type

The data is sorted first by the logical drive user label and second by the logical block address (LBA). Each entry in the table corresponds to a single sector.

#### **Syntax**

show storageSubsystem unreadableSectors

#### **Parameters**

None.

# Show Enhanced Remote Mirroring volume candidates

This command returns information about the candidate volumes on a remote storage subsystem that you can use as secondary volumes in a Enhanced Remote Mirroring configuration.

**Note:** In previous versions of this command the feature identifier was remoteMirror. This feature identifier is no longer valid and is replaced by syncMirror.

### **Syntax**

show syncMirror candidates primary="volumeName"
remoteStorageArrayName="storageArrayName"

#### Parameter

Parameter	Description
primary	The name of the local volume that you want for the primary volume in the remote-mirrored pair. Enclose the primary volume name in double quotation marks (" ").
remoteStorageArrayName	The remote storage subsystem that contains possible volumes for a secondary volume. If the remote storage subsystem name has special characters, you must also enclose the remote storage subsystem name in double quotation marks (" ").

# **Minimum Firmware Level**

5.40

# Show Enhanced Remote Mirroring volume synchronization progress

This command returns the progress of data synchronization between the primary volume and the secondary volume in a Synchronous Mirroring configuration. This command shows the progress as a percentage of data synchronization that has been completed.

**Note:** In previous versions of this command the feature identifier was remoteMirror. This feature identifier is no longer valid and is replaced by syncMirror.

#### Syntax

```
show syncMirror (localVolume ["volumeName"] |
localVolumes ["volumeName1" ... "volumeNameN"])
synchronizationProgress
```

#### **Parameters**

Parameter	Description
localVolume or localVolumes	The name of the primary volume of the remote mirrored pair for which you want to check synchronization progress. Enclose the primary volume name in double quotation marks (" ") inside of square brackets ([ ]).

## **Minimum Firmware Level**

5.40

# **Show String**

This command shows a string of text from a script file. This command is similar to the **echo** command in MS-DOS and UNIX.

#### Syntax

show "string"

#### **Parameters**

None.

### Notes

Enclose the string in double quotation marks (" ").

# Show Thin Logical Drive

This command returns the expansion history or the consumed capacity for the specified thin logical drive or logical drives.

#### Syntax

```
show (allLogicalDrives | logicaldrive [logicaldriveName] |
logicaldrives [logicaldriveName1 ... logicaldriveNameN])
(consumedCapacity | (expansionHistory [file=fileName]))
```

Parameter	Description
allLogicalDrives	The setting to return information about all of the logical drives in the storage subsystem.

Parameter	Description
logicaldrive or logicaldrives	The name of the specific logical drive for which you are retrieving information. You can enter more than one logical drive name. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks ("").
consumedCapacity	The setting to return a concise list of information about the consumed capacity of the thin logical drives.
expansionHistory	The setting to return a concise list of information about the expansion history of the thin logical drives.
file	The <b>file</b> parameter specifies a file to log the output of the <b>expansionHistory</b> parameter. The file is valid only when used with the <b>expansionHistory</b> parameter. An invalid file name causes the command to fail.

With the **expansionHistory** parameter, the command returns information similar to the example shown below:

Thin logicalDrive I	name: logicalDrive	-nameRepository log	icalDrive Name: REPOS_NNNN
Logged Time	Expansion Type	Start Capacity	End Capacity
MM/DD/YYYY HH:MM:SS	Manual Automatic	NNNNNNN bytes	NNNNNNN bytes

With the **consumedCapacity** parameter, the command returns information similar to the example shown below:

Logical Drive	Provisioned Capacity	Consumed Capacity	Quota	<pre>% Prov. Consumed</pre>
logicaldriveName	500.000 GB	230.000 GB	700.000 GB	46%

# Show Storage Subsystem Unconfigured iSCSI Initiators

This command returns a list of initiators that have been detected by the storage subsystem but are not yet configured into the storage subsystem topology.

## **Syntax**

show storageSubsystem unconfiguredIscsiInitiators

## **Parameters**

None.

# Show VolumeCopy

This command returns this information about VolumeCopy operations:

- The copy status
- The start time stamp
- The completion time stamp
- The copy priority
- The source logical drive World-Wide Identifier (WWID) or the target logical drive WWID
- The target logical drive Read-Only attribute setting

You can retrieve information about a specific VolumeCopy pair or all of the VolumeCopy pairs in the storage subsystem.

#### Syntax

```
show VolumeCopy (allLogicalDrives | source ["sourceName"] |
target ["targetName"])
```

#### **Parameters**

Parameter	Description
allLogicalDrives	The setting to return information about VolumeCopy operations for all of the VolumeCopy pairs.
source	The name of the source logical drive about which you want to retrieve information. Enclose the source logical drive name in double quotation marks ("") inside of square brackets ([]).
target	The name of the target logical drive about which you want to retrieve information. Enclose the target logical drive name in double quotation marks ("") inside of square brackets ([]).

# Show VolumeCopy Source Candidates

This command returns information about the candidate logical drives that you can use as the source for a VolumeCopy operation.

#### Syntax

show VolumeCopy sourcecandidates

#### **Parameters**

None.

#### Notes

This command returns VolumeCopy source information as shown in this example:

```
Logical Drive Name: R1_2drives
Capacity: 40.0 GB
subsystem: 0
Logical Drive Name: R6_5drives_SATA
Capacity: 46.0 GB
```

subsystem: 1 Logical Drive Name: R5 Capacity: 40.0 GB subsystem: 2

# Show VolumeCopy Target Candidates

This command returns information about the candidate logical drives that you can use as the target for a VolumeCopy operation.

## **Syntax**

show VolumeCopy source ["sourceName"] targetCandidates

## Parameter

Parameter	Description
source	The name of the source logical drive for which you are trying to find a candidate target logical drive. Enclose the source logical drive name in double quotation marks ("") inside of square brackets ([]).

# Start subsystem Defragment

This command starts a defragment operation on the specified subsystem.

**Note:** Defragmenting an subsystem starts a long-running operation that you cannot stop.

# Syntax

start subsystem [subsystemNumber] defragment

## Parameter

Parameter	Description
subsystem	The alphanumeric identifier of the subsystem (including - and _) that you want to defragment. Enclose the subsystem identifier in square brackets ([]).

## Notes

Host I/O errors might result in the subsystems with more than 32 logical drives. This operation also might result in internal controller reboots because the timeout period ends before the subsystem definition is set. If you experience this issue, quiesce the host I/O operations, and try the command again.

# Start subsystem Export

This command moves an subsystem into an Exported state. Then you can remove the disk drives that comprise the subsystem and reinstall the disk drives in a different storage subsystem. **Note:** Within the subsystem, you cannot move logical drives that are associated with the premium features from one storage subsystem to another storage subsystem.

#### Syntax

start subsystem [subsystemNumber] export

#### Parameter

Parameter	Description
subsystem	The alphanumeric identifier of the subsystem (including - and _) that you want to export. Enclose the subsystem identifier in square brackets ([]).

#### Notes

When this command is successful, you can run the **start subsystem import** command to finish moving the subsystem to a Complete state, which makes the subsystem available to the new storage subsystem.

If this command is unsuccessful because hardware problems prevented the completion of the export, use the **set subsystem forceState** command. The **set subsystem forceState** command lets you use the **start subsystem import** command to import an subsystem.

After the subsystem is in an Exported state or a Forced state, you can remove the disk drives that comprise the subsystem from the storage subsystem. You can reinstall the disk drives in a different storage subsystem.

# Start array import

This command moves an array into a **Complete** state to make a newly introduced array available to its new storage subsystem. The array must be in an Exported state or a Forced state before you run this command. Upon successfully running the command, the array is operational.

**Note:** Within the subsystem, you cannot move logical drives that are associated with the premium features from one storage subsystem to another storage subsystem.

#### Syntax

start array [arrayName] import

Parameter	Description
array	The alphanumeric identifier of the array (including - and _) that you want to import. Enclose the subsystem identifier in square brackets ([ ]).

Higher-level logical drives that are specifically related to premium features (FlashCopy, Enhanced Remote Mirroring, VolumeCopy, mapping, and persistent reservations) are removed as part of the import operation.

The show subsystem importDependencies command must be run before the start subsystem import command.

# Start subsystem Locate

This command identifies the disk drives that are logically grouped together to form the specified subsystem by flashing the indicator lights on the disk drives. (Use the **stop subsystem locate** command to turn off the indicator lights on the disk drives.)

### **Syntax**

start subsystem [subsystemNumber] locate

### Parameter

Parameter	Description
subsystem	The alphanumeric identifier of the subsystem (including - and _) for which you want to locate the disk drives that belong to that subsystem. Enclose the subsystem identifier in square brackets ([]).

# Start Enhanced Global Mirroring synchronization

This command starts Enhanced Global Mirroring synchronization.

#### Syntax

start enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"]
synchronize

Header	Header
enhancedGlobalMirrorGroup	The name of the Enhanced Global Mirror
	Group for which you want to start
	synchronization. Enclose the Enhanced
	Global Mirror Group name in square
	brackets ([]). If the Enhanced Global Mirror
	Group name has special characters, you also
	must enclose the Enhanced Global Mirror
	Group name in double quotation marks (" ")
	inside square brackets. If the Enhanced
	Global Mirror Group name consists only of
	numbers, such as "1002," you also must
	enclose the Enhanced Global Mirror Group
	name in double quotation marks (" ") inside
	square brackets.

#### 7.84

# Start Cache Backup Device Diagnostic

**Attention:** Before you run this diagnostic test, make sure that the cache backup device has a status of Optimal.

This command runs diagnostic tests to evaluate the functionality of the device that you use to backup the data in the cache if you lose power to the controller. The diagnostic tests are specific to the backup device that is in the controller. Before you run these tests, make these changes to the controller that has the backup device on which you want to run diagnostics:

- Place the controller into service mode (use the set controller [(a | b)]availability=serviceMode command).
- Attach the management client directly to the controller through the management Ethernet port.

**Note:** In a dual-controller configuration, you must run these diagnostic tests through the controller that you want to evaluate. You cannot run these diagnostic tests through the partner controller.

### Syntax

```
start cacheBackupDevice [(1 | n | all)]
controller [(a | b)]
diagnostic diagnosticType=(basic | extended)
[extendedTestID=(writePatterns | random)]
```

Parameter	Description
cacheBackupDevice	The identifier for the cache backup device on which you want to run the diagnostic tests. Valid cache backup device identifiers are 1, 2, 3, 4 or 5.
	• 1 for USB1 on the controller circuit board
	• 2 for USB2 on the controller circuit board
	• 3 for USB3 on the controller circuit board
	• 4 for USB4 on the controller circuit board
	• all for all of the USBs on the controller circuit board
	<b>Note:</b> If you have only one cache backup device, the all identifier does not work. Enclose the identifier for the cache backup device in square brackets ([]).
controller	The controller that has the cache backup device on which you want to run the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

Parameter	Description
diagnosticType	The level of diagnostic testing that you want to run on the cache backup device. You can run one of these levels of testing:
	basic – This option validates the basic operation of the ability of the cache backup device to store cache data. This option determines these capabilities of the cache backup device:
	• Whether the cache backup device is write protected or the cache can write data to the device.
	• If the cache backup device is approaching its write cycle limit.
	extended – This option enables you to run more comprehensive diagnostic tests on the host interface card.
extendTestID	This parameter selects the extended test option that you want to run.
	If you choose the <b>extended</b> parameter, you also must also use the <b>extendedTestID</b> parameter and one of the extended test options.

Extended Test Option	Description
writePatterns	This option writes a predefined pattern of data in blocks to the entire cache backup device. Each block that was written is then read back, and the data is verified for integrity and accuracy.
random	This option writes a random pattern to each flash block in the cache backup device.

- When an unexpected power loss occurs, cache memory can have data that has not been written to the drives. This data must be preserved so that it can be written to the drives when power is restored. The contents of the cache memory are backed up to a persistent storage device, such as a USB flash drive, a SATA drive, or a solid state device (SSD).
- The total storage capacity of the flash drives must be equal to the total cache memory, considering that all storage space in a flash drive is not always usable. For example, in a 1-GB flash drive, approximately 968 MB is usable. Also, in some flash drives, the Cyclic Redundancy Check (CRC) needs to be stored along with the data. Because the metadata region is persisted in these flash drives, the storage capacity for the flash drives must be greater than the size of the cache memory.
- You can run the diagnostic test on only one controller in the storage subsystem at any one time.

# **Start Cache Memory Diagnostic**

This command runs extended diagnostic tests to evaluate the functionality of the cache memory in a controller. Before you run these tests, you must make these changes to the controller on which you want to run diagnostics:

Place the controller into Service mode (use the set controller [(a | b)]availability=serviceMode command).

• Attach the management client directly to the controller through the management Ethernet port.

**Note:** In a dual controller configuration, you must run these diagnostic tests through the controller that you want to evaluate. You cannot run these diagnostic tests through the partner controller.

# **Syntax**

```
start cacheMemory controller [(a | b)] diagnostic
diagnosticType=(basic | extended)
[extendedTestID=(marchC | patterns | psuedoRndm| DMACopy)]
```

Parameter	Description
controller	The controller that has the cache memory on which you want to run the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
diagnosticType	The level of diagnostic testing that you want to run on the host interface card. You can run one of these levels of testing:
	basic – This option validates the ability of the cache memory to address and access data.
	extended – This option enables you to run more comprehensive diagnostic tests on the host interface card.
extendedTestID	This parameter selects the extended test option that you want to run.
	If you choose the <b>extended</b> parameter, you must also used the <b>extended TestID</b> parameter and one of the extended test options.

Extended Test Option	Description
marchC	This option performs a March C test on specific regions of the Reconfigurable Processor Assembly (RPA) memory. This option tests for only one set of inverse patterns.
patterns	This option performs a word pattern test where the test sequence proceeds with a series of read/write operations for all locations in the specified memory region. The test uses a set of special patterns. The test writes and verifies several patterns at 32-bit widths.
pseudoRndm	This option generates a non-repetitive pattern for double word length, writes the pattern to the entire region, and reads back the pattern for verification.
DMAcopy	This option tests the capability of Direct Memory Access (DMA) copy operations across regions in the cache memory. This options uses the RPA hardware capabilities to move the data from one region to another region.

You can run the diagnostic test on only one controller in the storage subsystem at any one time.

# **Start Disk Pool Locate**

This command identifies the drives that are logically grouped together to form the specified disk pool by flashing the indicator lights on the drives. (Use the stop diskPool locate command to turn off the indicator lights on the drives.)

## **Syntax**

start diskPool [diskPoolName] locate

## **Parameters**

Parameter	Description
diskPool	The name of the disk pool that you want to locate. Enclose the disk pool name in square brackets ([]).

# Start Consistency Group Enhanced FlashCopy Rollback

This command starts a rollback operation to the member base logical drives in an Enhanced FlashCopy consistency group. The content of the base logical drives changes immediately to match the point-in-time content of the consistency group Enhanced FlashCopy logical drive. The base logical drives immediately becomes available for read/write requests after the rollback operation has successfully completed. To stop an Enhanced FlashCopy consistency group rollback operation, use the stop rollback command for consistency groups.

The repository logical drive that is associated with the consistency group Enhanced FlashCopy logical drive continues to track any new changes between the base logical drive and the consistency group Enhanced FlashCopy logical drive that occur after the rollback operation is completed.

To stop a rollback operation to the member base logical drives use the stop cgEnhancedFlashCopyImage rollback command.

## Syntax

start cgEnhancedFlashCopyImage ["consistencyGroupName"] rollback
[memberLogicalDriveSet ["memberLogicalDriveName1 ... memberLogicalDriveNamen"]]

Parameter	Description
consistencyGroupName	The name of the consistency group for which you want to start a rollback operation. Enclose the name of the consistency group in double quotation marks (" ") inside square brackets ([ ]). When the <b>memberLogicalDriveSet</b> parameter is not used, the rollback process applies to all member logical drives of the consistency group.

Parameter	Description
memberLogicalDriveSet	The name of one or more member base logical drives in a consistency group that you want to rollback. Enclose each member base logical drive name in double quotation marks (" ") inside parenthesizes (( )). You can enter more than one member logical drive name. Enclose all of the member logical drive names in one set of double quotation marks (" ") inside parenthesizes (( )). Separate each member logical drive name with a white space.

To start a roll back operation for an entire consistency group that has the name **congroup1**, you would use this command:

start cgEnhancedFlashCopyImage ["congroup1"] rollback;

To start a roll back operation for base logical drive members **memVol1**, **memVol2**, and **memVol3** in a consistency group that has the name **congroup2**, you would use this command:

start cgEnhancedFlashCopyImage ["congroup2"]
rollback memberLogicalDriveset=("memVol1 memVol2 memVol3");

# **Start Controller Diagnostic**

This command runs diagnostic tests to evaluate the functionality of the controller card. Before you run these tests, you must make these changes to the controller on which you want to run diagnostics:

- Place the controller into Service Mode (use the set controller [(a | b)]availability=serviceMode command).
- Attach the management client directly to the controller through the management Ethernet port.

**Note:** In a dual controller configuration, you must run these diagnostic tests through the controller that you want to evaluate. You cannot run these diagnostic tests through the partner controller.

#### Syntax

```
start controller [(a | b)] diagnostic diagnosticType=(basic | extended)
[extendedTestID=(SRAM | FIFO | dataCopy| RAID5Parity | RAID6Parity)]
```

Parameter	Description
controller	The controller on which you want to run the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

Parameter	Description
diagnostic	The level of diagnostic testing that you want to run on the host interface card. You can run one of these levels of testing:
	basic – This option validates the ability of the base controller to address and access data.
	extended – This option enables you to run more comprehensive diagnostic tests on the base controller card.
extendedTestID	This parameter selects the extended test option that you want to run.
	If you choose the <b>extended</b> parameter, you must also used the <b>extended TestID</b> parameter and one of the extended test options.

Extended Test Option	Description
SRAM	This option tests for address, data, and data retention. The address test attempts to write to specific address offsets. The data test attempts to write several data patterns to the address offsets. The data retention test attempts to write a data pattern and then read the data pattern back after a delay. The purpose of the SRAM option is to find memory parity or error correcting code (ECC) errors.
FIFO	This option tests the active processor chip (APC) first in, first out (FIFO) data transmission of the Zip chip. The APC FIFO channels are tested concurrently by writing and verifying different patterns to each channel.
dataCopy	This option tests the ability of the Zip chip to support data copy operations that can copy data from one area of the Zip SDRAM to another area of the Zip SDRAM. This test is performed on any available section of the Zip chip that is not busy.
RAID5Parity	This option tests the ability of the Zip APC to generate and verify RAID 5 parity data. Data buffers are setup in processor memory and parity is generated in process or memory. Some data buffers are set up in parallel architecture (RPA) memory and parity is generated for the data within the RPA memory. The parity that is generated within processor memory is then compared with the parity in the Zip APC.
RAID6Parity	This option tests the ability of the Zip APC to generate and verify RAID 6 parity data. Data buffers are setup in processor memory and parity is generated in processor memory. Some data buffers are set up in redundant parallel architecture (RPA) memory and parity is generated for the data within the RPA memory. The parity that is generated within processor memory is then compared with the parity in the Zip APC.

You can run the diagnostic test on only one controller in the storage subsystem at any one time.

# **Start Controller Trace**

# Syntax

This command starts an operation that saves debug trace information to a compressed file. The debug trace information can be used by an IBM Technical Support representative to help analyze how well a storage subsystem is running.

start controller [(a | b | both)] trace dataType=(current | flushed | currentFlushed | all) [forceFlush=(TRUE | FALSE)]

## **Parameters**

Parameter	Description
controller	The controller for which you want to collect the trace debug information. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. You can also simultaneously collect debug for both controllers by entering both. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.
dataType	The type of data that you want to collect:
	<ul> <li>current – Retrieves the current DQ traces</li> </ul>
	<ul> <li>flushed – Retrieves all flushed DQ traces</li> </ul>
	<ul> <li>currentFlushed – Retrieves both the current DQ trace and the flushed DQ trace</li> </ul>
	<ul> <li>all – Retrieves the current DQ trace, flushed DQ trace, and all platform DQ traces</li> </ul>
	<b>Note:</b> If <b>dataType=flushed</b> and <b>forceFlush=True</b> , an error message is returned indicating that only active traces can be flushed to the buffer on retrieval.
forceFlush	The setting to move the DQ information in the current buffer to the flushed buffer when the DQ trace information defined by the <b>dataType</b> parameter is retrieved. To enable force flush, set this parameter to <b>TRUE</b> . To disable <b>force</b> <b>flush</b> , set this parameter to <b>FALSE</b> . <b>Note:</b> If <b>dataType=flushed</b> and <b>forceFlush=True</b> , an error message is returned indicating that only active traces can be flushed to the buffer on retrieval.
file	The file path and the file name to which you want to save the DQ trace information. Refer to the Notes section for information about naming the files.

## Notes

The DQ trace information is written to a compressed file with an extension of .zip. The file name is a combination of a user-defined file name and the storage subsystem identifier (SAID). A constant of "dq" is also added to the file name. The complete file name has this form: user\_defined\_file\_name-SAID-dq.zip

The compressed file contains the information listed in this table.

File Name	Directory	Comments
user_provided_file_name-SAID- A.dq	SAID/timestamp/	The DQ trace data retrieved from controller A.
user_provided_file_name-SAID- B.dq	SAID/timestamp/	The DQ trace data retrieved fromcontroller B.
user_provided_file_name-SAID- trace_description.xm	SAID/timestamp/	The description file in an xml format that describes the DQ file attributes for future data mining.

# **Start Drive Channel Fault Isolation Diagnostics**

This command runs the drive channel fault isolation diagnostics and stores the results.

# **Syntax**

```
start diskDriveChannel [(1 | 2 | 3 | 4 | 5 | 6 | 7 | 8)]
controller [(a | b)] faultDiagnostics
{testDevices=[all | controller=(a | b) |
esms=[enclosureID1 (left | right), ... enclosureIDn (left | right)] |
Drives=[enclosureID1, drawerID1, slotID1, ..., enclosureIDn, drawerIDn, slotIDn]] |
dataPattern=(fixed | pseudoRandom) |
patternNumber=[(@xhexadecimal | number)] |
maxErrorCount=integer |
testIterations=integer |
timeout=timeInterval}
```

Parameter	Description
controller	The identifier letter of the controller that you want to test. Valid controller identifier values are a or b, where a is the controller in slot A, and b is the controller in slot B.
testDevices	The identifiers of the devices (controllers, ESMs, or drives) that you want to test. You can specify <i>all</i> or enter the specific identifiers for the devices that you want to diagnose.
Drives	EnclosureID1,drawerID1, slotID1 enclosureIDn,drawerIDn,slotIDn is for enclosures with multiple disk drive drawers. For enclosures without disk drive drawers, use enclosureID1,slotID1 enclosureIDn,slotIDn instead. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the values in square brackets ([]).
dataPattern	The method of repeatability that you want to test.
patternNumber	The hexadecimal data pattern you want to use to run the test. This number can be any hexadecimal number between 0000 to FFFF.
maxErrorCount	The number of errors that you want to accept before terminating the test.
testIterations	The number of times that you want to repeat the test.

Parameter	Description
timeout	The length of time in minutes that you want to run the test.

Use the save driveChannel faultDiagnostics command and the stop driveChannel faultDiagnostics command in association with the start driveChannel faultDiagnostics command. These commands are needed to save the diagnostic test results to a file and to stop the diagnostic test.

You can also stop this command at any time by pressing Ctrl+C

# **Start Drive Channel Locate**

This command identifies the expansion drawers that are connected to a specific disk drive channel by turning on the indicator lights for the expansion drawer that is connected to the disk drive channel. (Use the **stop driveChannel locate** command to turn off the indicator lights on the expansion drawer.)

#### Syntax

start driveChannel [(1 | 2 | 3 | 4 | 5 | 6 |7 | 8)] locate

## Parameter

Parameter	Description
driveChannel	The identifier number of the disk drive channel that you want to locate. Valid values for the identifier number for the disk drive channel are 1, 2, 3, 4, 5, 6, 7, or 8. Enclose the disk drive channel identifier number in square brackets ([]).

# **Start Disk Drive Initialize**

This command starts disk drive initialization.

**Attention: Possible damage to the storage subsystem configuration** – As soon as you enter this command, all user data is destroyed.

#### Syntax

start drive [enclosureID,drawerID,slotID] initialize

Parameter	Description
drive	EnclosureID1,drawerID1,slotID1 enclosureIDn,drawerIDn,slotIDn is for enclosures with multiple disk drive drawers. For enclosures without disk drive drawers, use enclosureID1,slotID1 enclosureIDn,slotIDn instead. Enclosure ID values are θ to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the values in square brackets ([]).
## Notes

The **drive** parameter supports both high-capacity drive expansion enclosures and low-capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low-capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, you need only specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low-capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to  $\theta$ , and specify the ID of the slot in which a drive resides.

## **Start Drive Locate**

This command locates a disk drive by turning on an indicator light on the disk drive. (Run the **stop drive locate** command to turn off the indicator light on the disk drive.)

### **Syntax**

start drive [enclosureID,drawerID,slotID] locate

### Parameter

Parameter	Description
drive	EnclosureID1,drawerID1,slotID1 enclosureIDn,drawerIDn,slotIDn is for enclosures with multiple disk drive drawers. For enclosures without disk drive drawers, use enclosureID1,slotID1 enclosureIDn,slotIDn instead. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the values in square brackets ([ ]).

## Notes

The **drive** parameter supports both high-capacity drive expansion enclosures and low∆capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low∆capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low∆capacity drive expansion enclosure, you need only specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low∆capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to  $\theta$ , and specify the ID of the slot in which a drive resides.

## **Start Disk Drive Reconstruction**

This command starts reconstructing a disk drive.

## **Syntax**

start Drive [enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn]
reconstruct

#### **Parameter**

Parameter	Description
drive	EnclosureID1,drawerID1,slotID1 enclosureIDn,drawerIDn,slotIDn is for enclosures with multiple disk drive drawers. For enclosures without disk drive drawers, use enclosureID1,slotID1 enclosureIDn,slotIDn instead. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose values in square brackets ([ ]).

### Notes

The **drive** parameter supports both high-capacity drive expansion enclosures and low∆capacity drive expansion enclosures. A high-capacity drive expansion enclosure has drawers that hold the drives. The drawers slide out of the drive expansion enclosure to provide access to the drives. A low∆capacity drive expansion enclosure does not have drawers. For a high-capacity drive expansion enclosure, you must specify the identifier (ID) of the drive expansion enclosure, the ID of the drawer, and the ID of the slot in which a drive resides. For a low∆capacity drive expansion enclosure, you need only specify the ID of the drive expansion enclosure and the ID of the slot in which a drive resides. For a low∆capacity drive expansion enclosure, an alternative method for identifying a location for a drive is to specify the ID of the drive expansion enclosure, set the ID of the drawer to  $\theta$ , and specify the ID of the slot in which a drive resides.

## Start Enclosure Locate

This command locates an enclosure by turning on the indicator light. (Use the **stop enclosure locate** command to turn off the indicator light for the enclosure.)

#### Syntax

start enclosure [enclosureID] locate

#### Parameter

Parameter	Description
enclosure	The enclosure that you want to locate. Enclosure ID values are 0 to 99. Enclose the enclosure ID value in square brackets ([ ]).

## Start Enhanced FlashCopy Image Rollback

This command starts a rollback operation for a set of Enhanced FlashCopy images. The content of the base logical drive changes immediately to match the point-in-time content of the selected Enhanced FlashCopy image logical drive. The base logical drive immediately becomes available for read/write requests after the rollback operation has successfully completed. To stop an Enhanced FlashCopy image rollback operation, use the stop rollback enhancedFlashCopyImage command. The repository logical drive that is associated with the Enhanced FlashCopy image continues to track any new changes between the base logical drive and the Enhanced FlashCopy image logical drive that occur after the rollback operation is completed.

**Note:** You cannot use this command for Enhanced FlashCopy images involved in online VolumeCopy.

## Syntax

start enhancedFlashCopyImage ["enhancedFlashCopyImageName"] rollback

Parameters	
------------	--

Parameter	Description
enhancedFlashCopyImage	The name of the Enhanced FlashCopy image. The name of an Enhanced FlashCopy image is comprised of two parts:
	The name of the Enhanced FlashCopy group
	• An identifier for the Enhanced FlashCopy image in the Enhanced FlashCopy group.
	The identifier for the Enhanced FlashCopy image can be one of these:
	• An integer value that is the sequence number of the Enhanced FlashCopy in the Enhanced FlashCopy group.
	• <b>NEWEST</b> - Use this option when you want to show the latest Enhanced FlashCopy image created in the Enhanced FlashCopy group.
	• <b>OLDEST</b> - Use this option when you want to show the earliest Enhanced FlashCopy image created in the Enhanced FlashCopy group.
	Enclose the Enhanced FlashCopy image name in double quotation marks ("") inside square brackets ([]).

## Notes

The name of an Enhanced FlashCopy image has two parts separated by a colon (:):

- The identifier of the Enhanced FlashCopy group
- The identifier of the Enhanced FlashCopy image

For example, if you want to start a rollback operation for Enhanced FlashCopy image **12345** in an Enhanced FlashCopy group that has the name **enhancedFlashCopyGroup1**, you would use this command:

start enhancedFlashCopyImage ["enhancedFlashCopyGroup1:12345"] rollback;

To start a rollback operation for the most recent Enhanced FlashCopy image in an Enhanced FlashCopy group that has the name **enhancedFlashCopyGroup1**, you would use this command:

start enhancedFlashCopyImage ["enhancedFlashCopyGroup1:newest"]rollback;

# Start Enhanced Remote Mirroring Synchronization

This command starts Enhanced Remote Mirroring synchronization.

#### **Syntax**

start remoteMirror primary ["logicalDriveName"] synchronize

#### Parameter

Parameter	Description
primary	The name of the primary logical drive for which you want to start synchronization. Enclose the primary logical drive name in double quotation marks ("") inside of square brackets ([]).

## Start FlashCopy Rollback

This command starts a rollback operation for one or more FlashCopy logical drives. The content of the base logical drive changes immediately to match the point-in-time content of the selected FlashCopy logical drive. The base logical drive immediately becomes available for read/write requests after the rollback operation has completed. To stop a FlashCopy rollback operation, use the stop rollback logical drive command.

The repository logical drive that is associated with the FlashCopy continues to track any new changes between the base logical drive and the FlashCopy logical drive that occur after the rollback operation is completed.

#### Syntax

start rollback (logicaldrive [flashcopyLogicalDriveName |
logicaldrives [flashcopyLogicalDriveName1 ... flashcopyLogicalDriveNamen])

#### **Parameters**

Parameter	Description
logicaldrive or logicaldrives	The name of the specific FlashCopy logical drive or FlashCopy logical drives for which you want to start a rollback operation. Enclose the FlashCopy logical drive name in square brackets ([]). If the FlashCopy logical drive name has special characters, you must also enclose the FlashCopy logical drive name in double quotation marks (" ").

# Start Host Interface Card Diagnostic

This command runs diagnostic tests to evaluate the functionality of the controller host interface card. The diagnostic tests that this command runs are specific to the host interface card that is in the controller. Before you run these tests, you must make these changes to the controller that has the host interface card on which you want to run diagnostics:

Place the controller into service mode (use the set controller [(a | b)]availability=serviceMode command).

• Attach the management client directly to the controller through the management Ethernet port.

**Note:** In a dual controller configuration, you must run these diagnostic tests through the controller that you want to evaluate. You cannot run these diagnostic tests through the partner controller.

## **Syntax**

```
tart hostCard [(1 | 2 | 3 | 4)] controller [(a | b)] diagnostic
diagnosticType=(basic | extended)
[extendedTestID=(EDC | DMA | RAM | internalLoopback)]
```

## **Parameters**

Parameter	Description		
hostCard	The identifier for host interface card on which you want to run the diagnostic tests. Valid host interface card identifiers are 1, 2, 3, or 4. The value of the identifier is for the position of the host interface card in the controller enclosure or drive enclosure. The position of the host interface card depends on the type of controller enclosure or drive enclosure in your storage subsystem. See the Notes for more information about the host interface card identifier and the position of the host interface cards in a controller enclosure. Enclose the controller identifier in square brackets ([]).		
controller	The controller that has the host interface card on which you want to run the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.		
diagnosticType	The level of diagnostic testing that you want to run on the host interface card. You can run one of these levels of testing:		
	basic – This option validates the ability of the host interface card to transport I/O data. This option takes approximately 30 seconds to complete.		
	extended – This option enables you to run more comprehensive diagnostic tests on the host interface card.		
extendedTestID	This parameter selects the extended test option that you want to run.		
	If you choose the <b>extended</b> parameter, you also must use the <b>extendedTestID</b> parameter and one of the extended test options.		
testIterations	The number of times that you want to repeat the test.		

Extended Test Option for Fibre Channel	Description
EDC	This option tests the Error Detection and Correction (EDC) generation, verification, and deletion functionality of the QE4 chip. This option tests all modes of the EDC operation, such as, insert, verify, and delete EDC data.

Extended Test Option for Fibre Channel	Description
DMA	This option tests the capability of the QE4 chip to take part in a Direct Memory Access (DMA) operation. The DMA can be internal to the chip or can be performed using the services of the raw pool within the Reconfigurable Processor Assembly (RPA) memory.

Extended Test Option for iSCSI	Description	
RAM	This option performs a read/write test for the local RAM, the SRAM, and also performs a checksum test for the NVRAM. This option performs the read/writetest for the RAM and SRAM by writing data to the memory, reading back the data, and comparing the read data to the written data.	
internalLoopBack	This option tests the ability of the physical layer (PHY) to transmit data packets over the physical link. For this test, the PHY is set to an internal loopback mode. Data is then transmitted, received, and compared with the original data. The test is run in two passes:	
	<ul><li>For the first pass, the data is predefined by the firmware.</li><li>For the second pass, the data is generated externally and then transmitted.</li></ul>	

## Notes

You can run the diagnostic test on only one controller in the storage subsystem at any one time.

A controller can have either one or two host interface cards.

- f a controller has one host interface card, the value for the position of each host interface card depends on the position of the controller in the controller enclosure. The host interface card in the controller in controller enclosure slot A has a position value of 1. The host interface card in the controller in controller enclosure slot B has a position value of 2.
- If a controller has two host interface cards, the value for the position of each host interface card depends on the position of the host interface card in the controller and the position of the controller in the controller enclosure. In most cases the position of the host interface card is identified with labels such as Host Card 1 and Host Card 2 on each controller. The position value of the host interface cards are listed in this table.

Controller	Host Card Label	Position
А	Host Card 1	1
А	Host Card 2	2
В	Host Card 1	3
В	Host Card 2	4

# Start iSCSI DHCP Refresh

This command initiates a refresh of the DHCP parameters for the iSCSI interface. If the configuration method for the interface is not set to DHCP, the procedure returns an error.

## Syntax

start storageSubsystem [iscsi-host-port] dhcpRefresh

## Parameter

Parameter	Description
iscsi-host-port	The identifier of the port on the storage subsystem on which you want to refresh the DHCP parameters. Enclose the iSCSI host port name in square brackets ([]).

### Notes

This operation ends the iSCSI connections for the portal and brings down the portal temporarily.

# Start Array defragment

This command starts a defragment operation on the specified Array.

**NOTE** Defragmenting a Array starts a long-running operation that you cannot stop.

## Syntax

start volumeGroup [volumeGroupName] defragment

#### **Parameter**

Parameter	Description
volumeGroup	The alphanumeric identifier of the Array (including - and _) that you want to defragment. Enclose the volume group identifier in square brackets ([]).

## Notes

Host I/O errors might result in the Arrays with more than 32 volumes. This operation also might result in internal controller reboots because the timeout period ends before the Array definition is set. If you experience this issue, quiesce the host I/O operations, and try the command again.

## Minimum Firmware Level

6.10

## Start Array export

This command moves a Array into an Exported state. Then you can remove the drives that comprise the Array and reinstall the drives in a different storage subsystem.

**NOTE** Within the Array, you cannot move volumes that are associated with the premium features from one storage subsystem to another storage subsystem.

### **Syntax**

start volumeGroup [volumeGroupName] export

#### **Parameter**

Parameter	Description
volumeGroup	The alphanumeric identifier of the Array (including - and _) that you want to export. Enclose the Array identifier in square brackets ([]).

### Notes

When this command is successful, you can run the start volumeGroup import command to finish moving the Array to a Complete state, which makes the volume group available to the new storage subsystem.

If this command is unsuccessful because hardware problems prevented the completion of the export, use the set volumeGroup forceState command. The set volumeGroup forceState command lets you use the start volumeGroup import command to import a Array.

After the Array is in an Exported state or a Forced state, you can remove the drives that comprise the volume group from the storage subsystem. You can reinstall the drives in a different storage subsystem.

## **Minimum Firmware Level**

7.10

## Start Array import

This command moves a Array into a Complete state to make a newly introduced Array available to its new storage subsystem. The Array must be in an Exported state or a Forced state before you run this command. Upon successfully running the command, the Array is operational.

**NOTE** Within the Array, you cannot move volumes that are associated with the premium features from one storage subsystem to another storage subsystem.

#### Syntax

start volumeGroup [volumeGroupName] import

Parameter	Description
volumeGroup	The alphanumeric identifier of the Array
	(including - and _) that you want to import. Enclose the Array identifier in square
	brackets ([]).

## Notes

Higher-level volumes that are specifically related to premium features (FlashCopy, Enhanced Remote Mirroring, Volume Copy, mapping, and persistent reservations) are removed as part of the import operation.

You must run the show volumeGroup importDependencies command before you run the start volumeGroup import command.

## **Minimum Firmware Level**

7.10

# **Start Array locate**

This command identifies the drives that are logically grouped together to form the specified Array by flashing the indicator lights on the drives. (Use the stop volumeGroup locate command to turn off the indicator lights on the drives.)

#### Syntax

start volumeGroup [volumeGroupName] locate

#### Parameter

Parameter	Description
volumeGroup	The alphanumeric identifier of the Array (including - and _) for which you want to locate the drives that belong to that Array. Enclose the Array identifier in square brackets ([]).

## **Minimum Firmware Level**

6.16

# **Start Logical Drive Initialization**

This command starts the formatting of a logical drive in a storage subsystem.

**Note:** Formatting a logical drive starts a long-running operation that you cannot stop.

#### Syntax

start logicalDrive [logicalDriveName] initialize

Parameter	Description
logicalDrive	The name of the logical drive for which you are starting the formatting. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you also must enclose the logical drive name in double quotation marks ("").

# Start Secure Disk Drive Erase

This command erases all of the data from one or more Full Disk Encryption (FDE) drives so that they can be reused as FDE drives. Run this command only when the FDE drives are no longer part of a secure subsystem, or when the security key is unknown.

## **Syntax**

start secureErase (Drive [enclosureID,drawer ID,slotID] |
Drives [enclosureID1,drawerID1,slotID1 ... enclosureIDn,drawerIDn,slotIDn])

## **Parameters**

Parameter	Description
Drive or Drives	EnclosureID1,drawerID1,slotID1 enclosureIDn,drawerIDn,slotIDn is for enclosures with multiple disk drive drawers. For enclosures without disk drive drawers, use enclosureID1,slotID1 enclosureIDn,slotIDn instead Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the values in square brackets ([]).

#### Notes

Note:

The controller firmware creates a lock that restricts access to the FDE drives. FDE drives have a state called Security Capable. When you create a security key, the state is set to Security Enabled, which restricts access to all FDE drives that exist within the storage subsystem.

# Start Performance Read Cache performance modeling

This command starts performance modeling for the Performance Read Cache. Performance modeling monitors and measures I/O activity for a period of time and estimates performance for various Performance Read Cache sizes. Performance is estimated using two metrics: cache hit percentage and average response time. The performance modeling data is not available until you stop the operation using the stop performanceReadCache performanceModeling command.

#### **Syntax**

start performanceReadCache [performanceReadCacheName] performanceModeling

Parameter	Description
performanceReadCache	The alphanumeric identifier (including - and _) of the Performance Read Cache that you want to locate. Enclose the identifier in square brackets ([]). If the Performance Read Cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.

## Notes

Performance modeling ends and the performance modeling data is available when one of the following conditions occurs:

- Run the stop performanceReadCache performanceModeling command.
- Retrieve the performance modeling data using IBM System Storage DS Storage Manager.

Performance modeling ends, but no data is available when one of the following conditions occurs:

- You reboot the controller.
- You make any changes to the Performance Read Cache configuration.
- The state of the Performance Read Cache changes.

## Minimum Firmware Level

7.84

# Start Storage Subsystem Locate

This command locates a storage subsystem by turning on the indicator lights for the storage subsystem. (Use the **stop storageSubsystem locate** command to turn off the indicator lights for the storage subsystem.)

#### Syntax

start storageSubsystem locate

#### **Parameters**

None.

# Start Storage Subsystem Configuration Database Diagnostic

This command runs a consistency check against a configuration database.

#### Syntax

```
start storageSubsystem configDbDiagnostic
[sourceLocation=(disk | onboard) |
diagnosticType=fileSystem |
controller[(a | b)]]
```

Parameter	Description
sourceLocation	This parameter specifies the location of the database:
	<ul> <li>disk indicates that data comes directly from the database on the drive</li> </ul>
	• onboard indicates that data comes from the RPA memory location
	The default location is disk.
diagnosticType	The level of diagnostic testing that you want to run on the database. The option for this parameter is <b>fileSystem</b> , which checks the structural integrity of the database .
controller	The controller that has the database on which you want to run the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]).

#### Notes

This command runs a consistency check against a configuration database. All database records are checked. Error data is written to a file in the data folder on disk automatically. You do not need to specify an output file.

Executing a consistency check with the **diagnosticType** parameter set to a value other than filesystem and with the **sourceLocation** parameter set to **onboard** can cause the operation to run for a long time. This can have adverse effects on host I/O processing. This operation should be done only under direction from the support organization.

Upon completion of the diagnostic test, the controller firmware returns one of these results:

- Diagnosis completed without errors. No ZIP file created.
- Diagnosis completed with errors. Refer to the ZIP file created at:
  - ...\Install\_dir\data\FirmwareUpgradeReports\timestamp\_buildNo.zip

If the diagnostic test detects an inconsistency in the configuration database, the controller firmware performs these actions:

- Returns a description of the inconsistency
- · Saves a ZIP file containing raw binary data

The controller firmware saves the ZIP file to this location:

...\Install\_dir\data\FirmwareUpgradeReports\timestamp\_buildNo.zip

You can use the binary data to help determine the cause of the problem, or you can send the file containing the binary data to a IBM Technical Support representative.

To stop the database configuration diagnostic test, use the stop storageSubsystem configDbDiagnostic command.

In addition, you can start the database configuration diagnostic test through the storage management software GUI; however, you cannot stop the database configuration diagnostic test through the storage management software GUI. If you want to stop a running diagnostic test, you must use the stop storageSubsystem configDbDiagnostic command

# Start Storage Subsystem Diagnostic Data Capture

This command produces a storage subsystem diagnostic data capture.

## **Syntax**

start storageSubsystem ddc controller [(a|b)]

### **Parameters**

Parameter	Description
controller	This parameter specifies the controller from which to produce the diagnostic data capture. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

# Start Storage Subsystem iSNS Server Refresh

This command initiates a refresh of the network address information for the iSNS server. If the DHCP server is marginal or unresponsive, the refresh operation can take from two to three minutes to complete.

Note: This command is for IPv4 only.

## **Syntax**

start storageSubsystem isnsServerRefresh

## **Parameters**

None.

## Notes

If you used the set storageSubsystem isnsIPv4ConfigurationMethod command to set the configuration but did not set the configuration to DHCP, running the start storageSubsystem isnsServerRefresh command returns an error.

## Stop array locate

This command turns off the indicator lights on the disk drives that were turned on by the **start array locate** command.

## **Syntax**

stop array locate

## **Parameters**

None.

# Stop Cache Backup Device Diagnostic

This command stops the cache backup device diagnostic tests that were started by the **startcacheBackupDevice diagnostic** command.

#### Syntax

stop cacheBackupDevice controller [(a | b)] diagnostic

### **Parameters**

Parameter	Description
controller	The controller that has the cache backup device on which you are running the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

## **Stop Cache Memory Diagnostic**

This command stops the cache memory diagnostic tests that were started by the **start cacheMemorydiagnostic** command.

#### Syntax

stop cacheMemory controller [(a | b)] diagnostic

#### Parameters

Parameter	Description
controller	The controller that has the cache memory on which you are running the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

# Stop Storage Subsystem Configuration Database Diagnostic

This command stops the diagnostic test to validate the configuration database in the controller firmware that was started by the **start storageSubsystem configDbDiagnostic** command.

#### Syntax

stop storageSubsystem configDbDiagnostic

None.

## Notes

The controller firmware returns a confirmation that the diagnostic test was cancelled.

In addition, you can start the database configuration diagnostic test through the storage management software GUI; however, you cannot stop the database configuration diagnostic test through the storage management software GUI. If you want to stop a running diagnostic test, you must use the **stopstorageSubsystem** configDbDiagnostic command.

If you try to use the **stop storageSubsystem configDbDiagnostic** command after validation of the storage subsystem configuration has finished, you do not receive any message that the validation has finished. This behavior is expected.

# Stop Consistency Group Enhanced FlashCopy Rollback

This command stops a rollback operation to the member base logical drives in an Enhanced FlashCopy consistency group that was initiated by the start rollback cgEnhancedFlashCopy command.

**Note:** Canceling a consistency group Enhanced FlashCopy rollback operation leaves the base logical drive in an indeterminate state with potentially invalid or inconsistent data. The related consistency group Enhanced FlashCopy logical drive becomes disabled and unusable.

## **Syntax**

stop cgEnhancedFlashCopyImage["consistencyGroupName"]| rollback
[memberLogicalDriveSet ["memberLogicalDriveName1 ... memberLogicalDriveNamen"]]

Parameter	Description
consistencyGroupName	The name of the consistency group for which you want to stop a rollback operation. Enclose the name of the consistency group in double quotation marks (" ") inside square brackets ([ ]). When the <b>memberLogicalDriveSet</b> parameter is not used the rollback process applies to all member logical drives of the consistency group.
memberLogicalDriveSet	The name of one or more member base logical drives in a consistency group that you want to stop a rollback operation. Enclose each member base logical drive name in double quotation marks (" ") inside parenthesizes (( )). You can enter more than one member logical drive name. Enclose all of the member logical drive names in one set of double quotation marks (" ") inside parenthesizes (( )). Separate each member logical drive name with a white space.

## Parameters

## Notes

To stop a roll back operation for an entire consistency group that has the name **congroup1**, you would use this command:

stop cgEnhancedFlashCopyImage ["congroup1"] rollback;

To stop a roll back operation for base logical drive members **memVol1**, **memVol2**, and **memVol3** in a consistency group that has the name **congroup2**, you would use this command:

stop cgEnhancedFlashCopyImage ["congroup2"]
rollback memberLogicalDriveset=("memVol1 memVol2 memVol3");

## Stop Consistency Group Enhanced FlashCopy Logical Drive

This command stops the copy-on-write operation for creating a consistency group Enhanced FlashCopy logical drive. To restart the copy-on-write operation use the resume cgEnhancedFlashCopyLogicalDrive command.

**Note:** This command does not delete a consistency group Enhanced FlashCopy logical drive. To delete a consistency group Enhanced FlashCopy logical drive use the delete cgEnhancedFlashCopyLogicalDrive command.

#### Syntax

stop cgEnhancedFlashCopyLogicalDrive [enhancedFlashCopyLogicalDriveName]

#### **Parameters**

Parameter	Description
cgEnhancedFlashCopyLogicalDrive	The name of the consistency group Enhanced FlashCopy logical drive that you want to stop creating. Enclose in square brackets ([]), the name of the consistency group Enhanced FlashCopy logical drive that you want to delete. If the consistency group Enhanced FlashCopy logical drive name has special characters, you also must enclose the Enhanced FlashCopy logical drive name in double quotation marks (" ").

# **Stop Controller Diagnostic**

This command stops the controller diagnostic tests that were started by the **start controllerdiagnostic** command.

#### Syntax

stop controller [(a | b)] diagnostic

Parameter	Description
controller	The setting to return information about a specific controller in the storage subsystem. Valid controller identifiers are aor b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

## **Stop Disk Pool Locate**

This command turns off the indicator lights on the drives that were turned on by the start diskPool locate command.

### **Syntax**

stop diskPool locate

#### **Parameters**

None.

# **Stop Drive Channel Fault Isolation Diagnostics**

This command stops the drive channel fault isolation diagnostics, which stops the **start drive channel fault isolation diagnostics** command before it completes.

See the "Start Drive Channel Fault Isolation Diagnostics" on page 341 command for more information.

### **Syntax**

stop driveChannel faultDiagnostics

### Parameters

None.

#### Notes

Use the start driveChannel faultDiagnostics command and the save driveChannel faultDiagnostics command with the stop driveChannel faultDiagnostics command. These commands are needed to start the diagnostic test and save diagnostic test results to a file.

You can also stop the start driveChannel faultDiagnostics command at any time by pressing Ctrl+C.

# **Stop Drive Channel Locate**

This command turns off the indicator lights on the expansion drawers that were turned on by the **start driveChannel locate** command.

## Syntax

stop driveChannel locate

### Parameters

None.

## **Stop Drive Locate**

This command turns off the indicator light on the disk drive that was turned on by the **start drive locate** command.

#### Syntax

stop drive locate

#### Parameters

None.

## **Stop Enclosure Locate**

This command turns off the indicator light on the enclosure that was turned on by the **start enclosure locate** command.

#### Syntax

stop enclosure locate

#### **Parameters**

None.

# Stop Enhanced FlashCopy Group Pending Enhanced FlashCopy Images

This command cancels all of the pending Enhanced FlashCopy images that are to be created on an Enhanced FlashCopy group. If the Enhanced FlashCopy group does not have any pending Enhanced FlashCopy images, the firmware displays an error message and does not run the command. You can run this command on an Enhanced FlashCopy group or an Enhanced FlashCopy consistency group.

#### Syntax

stop (enhancedFlashCopyGroup [enhancedFlashCopyGroupName] |
consistencyGroup [enhancedFlashCopyConsistencyGroupName])
pendingEnhancedFlashCopyImageCreation

## **Parameters**

Parameter	Description
enhancedFlashCopyGroup	The alphanumeric identifier (including - and
	_) of the Enhanced FlashCopy group for
	which you want to stop pending Enhanced
	FlashCopy images. Enclose the Enhanced
	FlashCopy group identifier in square
	brackets ([ []).

Parameter	Description
consistencyGroup	The alphanumeric identifier (including - and _) of the Enhanced FlashCopy consistency group for which you want to stop pending Enhanced FlashCopy images. Enclose the Enhanced FlashCopy group identifier in square brackets ([ []).

# Stop Enhanced FlashCopy Image

This command cancels a pending Enhanced FlashCopy image copy-on-write operation on an Enhanced FlashCopy group. If the Enhanced FlashCopy group does not have any pending Enhanced FlashCopy images, this command returns an error message.

## **Syntax**

stop(enhancedFlashCopyGroup [enhancedFlashCopyGroupName] |
consistencyGroup [consistencyGroupName])
pendingEnhancedFlashCopyImageCreation

## **Parameters**

Parameter	Description
enhancedF1ashCopyGroup	The alphanumeric identifier (including - and _) of the Enhanced FlashCopy group that has the pending Enhanced FlashCopy image copy-on-write operation that you want to stop. Enclose the Enhanced FlashCopy group name in double quotation marks (" ").
consistencyGroup	The alphanumeric identifier (including - and _) of an Enhanced FlashCopy consistency group that has the pending Enhanced FlashCopy image copy-on-write operation that you want to stop. Enclose the Enhanced FlashCopy consistency group name in double quotation marks (" ").

# Stop Enhanced FlashCopy Image Rollback

This command stops an Enhanced FlashCopy image rollback operation that was initiated by the start enhancedFlashCopyImage rollback command.

#### CAUTION:

Canceling an Enhanced FlashCopy image rollback operation leaves the base logical drive in an indeterminate state with potentially invalid or inconsistent data. The related Enhanced FlashCopy image logical drive becomes disabled and unusable.

## Syntax

stop enhancedFlashCopyImage [enhancedFlashCopyImageName] rollback

Parameter	Description
enhancedFlashCopyImage	The name of the Enhanced FlashCopy image for which you want to stop a rollback operation. The name of an Enhanced FlashCopy image is comprised of two parts:
	• The name of the Enhanced FlashCopy group
	• An identifier for the Enhanced FlashCopy image in the Enhanced FlashCopy group.
	The identifier for the Enhanced FlashCopy image can be one of these:
	• An integer value that is the sequence number of the Enhanced FlashCopy in the Enhanced FlashCopy group.
	• <b>NEWEST</b> - Use this option when you want to show the latest Enhanced FlashCopy image created in the Enhanced FlashCopy group.
	• <b>OLDEST</b> - Use this option when you want to show the earliest Enhanced FlashCopy image created in the Enhanced FlashCopy group.
	Enclose the Enhanced FlashCopy image name in double quotation marks (" ") inside square brackets ([ ]).

## Notes

The name of an Enhanced FlashCopy image has two parts separated by a colon (:):

- The identifier of the Enhanced FlashCopy group
- The identifier of the Enhanced FlashCopy image

For example, if you want to stop a rollback operation for Enhanced FlashCopy image 12345 in an Enhanced FlashCopy group that has the name **enhancedFlashCopyGroup1**, you would use this command:

stop enhancedFlashCopyImage ["enhancedFlashCopyGroup1:12345"] rollback;

To stop a rollback operation for the most recent Enhanced FlashCopy image in an Enhanced FlashCopy group that has the name **enhancedFlashCopyGroup1**, you would use this command:

stop enhancedFlashCopyImage ["enhancedFlashCopyGroup1:newest"] rollback;

# Stop Enhanced FlashCopy Logical Drive

This command stops an Enhanced FlashCopy logical drive operation.

## **Syntax**

stop enhancedFlashCopyLogicalDrive ["enhancedFlashCopyLogicalDriveName"]

Parameter	Description
enhancedFlashCopyLogicalDrive	The name of the Enhanced FlashCopy logical drive that you want to stop. Enclose the Enhanced FlashCopy logical drive name in double quotation marks (" ") inside square brackets ([ ]).

#### Notes

You can use any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#) for the names. Names can have a maximum of 30 characters.

## Stop FlashCopy

This command stops a copy-on-write operation.

## Syntax

stop flashcopy (logicaDrive [logicaDriveName] |
logicaDrives [logicaDriveName1 ... logicaDriveNameN])

## **Parameters**

Parameter	Description
logicaDrive or logicaDrives	The name of the specific logical drive for which you want to stop a copy-on-write operation. You can enter more than one logical drive name. Enclose the logical drive names using one of these forms:
	• On a Windows command line: \" logicaldriveName\"
	• In a Windows script engine window: [" logicaldriveName"]
	• On a Linux command line: \" logicaldriveName\"
	• In a Linux script engine window: [\" logicaldriveName\"]

## Notes

Names can be any combination of alphanumeric characters, underscore (\_), hyphen (-), and pound (#). Names can have a maximum of 30 characters.

One technique for naming the FlashCopy logical drive and the FlashCopy repository logical drive is to add a hyphenated suffix to the original base logical drive name. The suffix distinguishes between the FlashCopy logical drive and the FlashCopy repository logical drive. For example, if you have a base logical drive with a name of Engineering Data, the FlashCopy logical drive can have a name of Engineering Data-S1, and the FlashCopy repository logical drive can have a name of EngineeringData-R1.

If you do not choose a name for either the FlashCopy logical drive or the FlashCopy repository logical drive, the storage management software creates a default name by using the base logical drive name. An example of the FlashCopy logical drive name that the controllers might create is, if the base logical drive name is *aaa* and does not have a FlashCopy logical drive, the default FlashCopy logical drive name is *aaa-1*. If the base logical drive already has *n-1* number of

FlashCopy logical drives, the default name is *aaa-n*. An example of the FlashCopy repository logical drive name that the controller might create is, if the base logical drive name is *aaa* and does not have a FlashCopy repository logical drive, the default FlashCopy repository logical drive name is *aaa-R1*. If the base logical drive already has *n-1* number of FlashCopy repository logical drives, the default name is *aaa-Rn*.

# Stop FlashCopy Rollback

This command stops a FlashCopy rollback operation that was initiated by the start rollback logical drive command.

#### CAUTION:

Canceling a rollback operation leaves the base logical drive in an indeterminate state with potentially invalid or inconsistent data. The related FlashCopy logical drive becomes disabled and unusable.

#### Syntax

stop rollback logicaldrive [flashcopyLogicalDriveName]

#### Parameters

Parameter	Description
logicaldrive	The name of the specific FlashCopy logical drive for which you want to stop a rollback operation. Enclose the FlashCopy logical drive name in square brackets ([]). If the FlashCopy logical drive name has special characters, you must also enclose the FlashCopy logical drive name in double quotation marks (" ").

## **Stop Host Interface Card Diagnostic**

This command stops the host interface card diagnostic tests that were started by the **start host carddiagnostic** command.

#### **Syntax**

stop host card controller [(a | b)] diagnostic

### Parameters

Parameter	Description
controller	The controller that has the host interface card on which you are running the diagnostic tests. Valid controller identifiers are a or b, where a is the controller in slot A, and b is the controller in slot B. Enclose the controller identifier in square brackets ([]). If you do not specify a controller, the storage management software returns a syntax error.

# Stop Storage Subsystem iSCSI Session

This command forces the termination of a storage subsystem iSCSI session.

## **Syntax**

stop storageSubsystem iscsiSession [integer]

## Parameter

Parameter	Description
iscsiSession	The identifier number of the iSCSI session. Enclose the identifier number of the iSCSI session in square brackets ([]).

# Stop Performance Read Cache performance modeling

This command stops the performance modeling operation and displays the performance modeling data for the Performance Read Cache. Before running this command, you must start the performance modeling operation with the start performanceReadCache performanceModeling command. Optionally, you can save the data to a file.

## **Syntax**

stop performanceReadCache [performanceReadCacheName] performanceModeling
[file="filename"]

## **Parameters**

Parameter	Description
performanceReadCache	The alphanumeric identifier (including - and _) of the Performance Read Cache that you want to locate. Enclose the identifier in square brackets ([]). If the Performance Read Cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.
file The fi want i data. 1 quota file= Files defaul perfor readLi name,	The file path and the file name to which you want to save the performance modeling data. Enclose the file name in double quotation marks (" "). For example: file="C:\Program
	Files\CLI\logs\performance.csv" The default name of the file that contains the performance modeling data is readLinkStatus.csv. You can use any file name, but you must use the .csv extension.

## Notes

This command returns the performance modeling information similar to this example. The size of your monitor determines how the information wraps and will affect how the information appears.

Performance Read Cache Name: my\_cache Start time: 4/18/12 2:38:26 PM IST Stop time: 4/18/12 2:38:45 PM IST Duration : 00:00:19 Performance Read Cache Performance Modeling Data (Response Time): SSD Reads HDD Reads HDD Writes

```
Overall Avg. Avg. Avg.
Cache Response Response % of Response % of Response % of
Capacity Time Time I/Os Time I/Os Time I/Os
186 GB 0 ms 0 ms 0.0 % 0 ms 0.0 % 0 ms 0.0 %
372 GB 0 ms 0 ms 0.0 % 0 ms 0.0 % 0 ms 0.0 %
557 GB * 0 ms 0 ms 0.0 % 0 ms 0.0 % 0 ms 0.0 %
558 GB 0 ms 0 ms 0.0 % 0 ms 0.0 % 0 ms 0.0 %
744 GB 0 ms 0 ms 0.0 % 0 ms 0.0 % 0 ms 0.0 %
931 GB 0 ms 0 ms 0.0 % 0 ms 0.0 % 0 ms 0.0 %
1117 GB 0 ms 0 ms 0.0 % 0 ms 0.0 % 0 ms 0.0 %
1303 GB 0 ms 0 ms 0.0 % 0 ms 0.0 % 0 ms 0.0 %
1489 GB 0 ms 0 ms 0.0 % 0 ms 0.0 % 0 ms 0.0 %
1675 GB 0 ms 0 ms 0.0 % 0 ms 0.0 % 0 ms 0.0 %
1862 GB 0 ms 0 ms 0.0 % 0 ms 0.0 % 0 ms 0.0 %
* = Current Performance Read Cache physical capacity.
Performance Read Cache Performance Modeling Data (Cache Hit %):
Cache Capacity Cache Hit %
186 GB 0 %
372 GB 0 %557 GB * 0 %
558 GB 0 %
744 GB 0 %
931 GB 0 %
1117 GB 0 %
1303 GB 0 %
1489 GB 0 %
1675 GB 0 %
1862 GB 0 %
* = Current Performance Read Cache physical capacity.
```

### **Minimum Firmware Level**

7.84

# Stop VolumeCopy

This command stops a VolumeCopy operation. This command is valid for both FlashCopy VolumeCopy pairs and new Enhanced FlashCopy VolumeCopy pairs.

#### Syntax

stop volumeCopy target [targetName] source
[sourceName]

### **Parameters**

Parameter	Description
target	The name of the target volume for which you want to stop a VolumeCopy operation. Enclose the target volume name in square brackets ([]). If the target volume name has special characters, you also must enclose the target volume name in double quotation marks (" ").
source	The name of the source volume for which you want to stop a VolumeCopy operation. Enclose the source volume name in square brackets ([]). If the source volume name has special characters, you also must enclose the source volume name in double quotation marks (" ").

## **Minimum Firmware Level**

5.40

## **Stop Array locate**

This command turns off the indicator lights on the drives that were turned on by the start volumeGroup locate command.

## **Syntax**

stop volumeGroup locate

## **Parameters**

None

## **Minimum Firmware Level**

6.16

# Stop Pending Enhanced FlashCopy Images on Consistency Group

This command stops all of the pending Enhanced FlashCopy images that are to be created on an Enhanced FlashCopy consistency group. If the Enhanced FlashCopy consistency group does not have any pending Enhanced FlashCopy images, the storage management software displays an error message and does not run the command.

## Syntax

stop consistencyGroup [consistencyGroupName] pendingEnhancedFlashCopyImageCreation

## **Parameters**

Parameter	Description
consistencyGroupName	The name of the consistency group for which you want to stop any pending Enhanced FlashCopy operation. Enclose the name of the consistency group in square brackets ([]).

# Stop Storage Subsystem Drive Firmware Download

This command stops a firmware download to the disk drives in a storage subsystem that was started with the **download storageSubsystem driveFirmware** command. This command does not stop a firmware download that is already in progress to a disk drive. This command stops all firmware downloads to disk drives that are waiting for the download.

## Syntax

stop storageSubsystem driveFirmwareDownload

## **Parameters**

None.

# Stop Storage Subsystem Locate

This command turns off the indicator lights on the storage subsystem that were turned on by the **start storageSubsystem locate** command.

#### **Syntax**

stop storageSubsystem locate

#### **Parameters**

None.

# Stop VolumeCopy

This command stops a VolumeCopy operation.

#### Syntax

stop VolumeCopy target [targetName] source [sourceName]

### **Parameters**

Parameter	Description
target	The name of the target logical drive for which you want to stop a VolumeCopy operation. Enclose the target logical drive name in square brackets ([]). If the target logical drive name has special characters, you also must enclose the target logical drive name in double quotation marks ("").
source	The name of the source logical drive for which you want to stop a VolumeCopy operation. Enclose the source logical drive name in square brackets ([]). If the source logical drive name has special characters, you also must enclose the source logical drive name in double quotation marks ("").

## **Suspend Enhanced Global Mirror Group**

This command suspends the synchronization of data on all mirrored pairs at the Enhanced Global Mirror Group level. This suspend command helps to reduce any performance impact to the host application that might occur while any changed data on the local storage subsystem is copied to the remote storage subsystem.

#### Syntax

suspend enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"]

Parameter	Description
enhancedGlobalMirrorGroup	The name of the Enhanced Global Mirror Group that you want to suspend. Enclose the Enhanced Global Mirror Group name in square brackets ([]). If the Enhanced Global Mirror Group name has special characters, you also must enclose the Enhanced Global
	Mirror Group name in double quotation marks (" ") inside square brackets. If the Enhanced Global Mirror Group name consists only of numbers, such as "1002," you also must enclose the Enhanced Global Mirror Group name in double quotation marks (" ") inside square brackets.

### Notes

When an Enhanced Global Mirror Group is in a suspended state, no attempt is made to copy data from the primary volumes to the secondary volumes of the mirrored pairs. Any writes to the primary side of the Enhanced Global Mirror Group are persistently logged in its associated mirror repository volumes. After the Enhanced Global Mirror Group is resumed, only the modified regions of the primary volumes are written to the secondary volumes.

## **Minimum Firmware Level**

7.84

# **Suspend Remote Mirror**

This command suspends a Enhanced Remote Mirroring operation.

### **Syntax**

suspend remoteMirror (primary [primarylogicalDriveName]
primaries [primarylogicalDriveName1 ... primarylogicalDriveNameN])
writeConsistency=(TRUE | FALSE)

### Parameters

Parameter	Description
primary or primaries	The name of the logical drive for which you want to suspend operation. Enclose the logical drive name in square brackets ([]). If the logical drive name has special characters, you must also enclose the logical drive name in double quotation marks ("").
writeConsistency	This parameter defines whether the logical drives identified in this command are in a write-consistency group or are separate. For the logical drives in the same write-consistency group, set this parameter to TRUE. For the logical drives that are separate, set this parameter to FALSE.

### Notes

If you set the **writeConsistency** parameter to TRUE, the logical drives must be in a write-consistency group (or groups). This command suspends all write-consistency groups that contain the logical drives. For example, if logical drives A, B, and C are in a write-consistency group and they have remote counterparts A', B', and C', the command:

suspend
remoteMirror logicalDrive ["A"] writeConsistency=TRUE

suspends A-A', B-B', and C-C'. If you have a write-consistency group 1={A, B, C} and write-consistency group 2={D, E, F}, the command: suspend remoteMirror logicalDrives=["A", "D"] writeConsistency=TRUE

suspends both write-consistency groups.

## Suspend Performance Read Cache

This command temporarily stops caching for all of the volumes that are using the Performance Read Cache. While caching is stopped, host reads are serviced from the base volumes instead of from the Performance Read Cache.

#### Syntax

suspend performanceReadCache [performanceReadCacheName]

#### Parameter

Parameter	Description
performanceReadCache	The alphanumeric identifier (including - and _) of the Performance Read Cache that you want to suspend. Enclose the identifier in square brackets ([]). If the Performance Read Cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks (" ") inside square brackets.

#### Notes

To restart caching, use the resume performanceReadCache command.

### **Minimum Firmware Level**

7.84

## Suspend Enhanced Global Mirror Group

This command suspends the synchronization of data on all mirrored pairs at the Enhanced Global Mirror Group level. This suspend command helps to reduce any performance impact to the host application that might occur while any changed data on the local storage subsystem is copied to the remote storage subsystem.

#### **Syntax**

suspend enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"]

Parameter	Description
enhancedG1oba1MirrorGroup	The name of the Enhanced Global Mirror
	Group that you want to suspend. Enclose
	the Enhanced Global Mirror Group name in
	square brackets ([]). If the Enhanced Global
	Mirror Group name has special characters,
	you also must enclose the Enhanced Global
	Mirror Group name in double quotation
	marks (" ") inside square brackets. If the
	Enhanced Global Mirror Group name
	consists only of numbers, such as "1002,"
	you also must enclose the Enhanced Global
	Mirror Group name in double quotation
	marks (" ") inside square brackets.

### Notes

When an Enhanced Global Mirror Group is in a suspended state, no attempt is made to copy data from the primary volumes to the secondary volumes of the mirrored pairs. Any writes to the primary side of the Enhanced Global Mirror Group are persistently logged in its associated mirror repository volumes. After the Enhanced Global Mirror Group is resumed, only the modified regions of the primary volumes are written to the secondary volumes.

## Minimum Firmware Level

7.84

# **Test Enhanced Global Mirror Group connectivity**

This command tests possible communication problems between the local storage subsystem and the remote the storage subsystem associated with an Enhanced Global Mirror Group.

## **Syntax**

diagnose enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"]
testID=(all | connectivity | latency | bandwidth |
portConnections)

## **Parameters**

Parameter	Description
enhancedGlobalMirrorGroup	The name of an existing Enhanced Global Mirror Group that you want to test. Enclose the Enhanced Global Mirror Group name in square brackets ([]). If the Enhanced Global Mirror Group name has special characters, you also must enclose the Enhanced Global Mirror Group name in double quotation marks (" ").

Parameter	Description
testID	The identifier for the diagnostic test you want to run. The identifier and corresponding tests are as follow:
	• All— Performs all the tests associated with this command .
	<ul> <li>Connectivity — Verifies that the two controllers have a communication path. The connectivity test sends an inter-controller message between the storage subsystems, and then validates that the corresponding Enhanced Global Mirror Group on the remote storage subsystem exists. It also validates that the volume members of the Enhanced Global Mirror Group on the remote system match the volume members of the Enhanced Global Mirror Group on the local system.</li> <li>Latency — Sends a SCSI test unit command to each volume on the remote storage subsystem associated with the Enhanced Global Mirror Group to test the minimum, average, and maximum latency.</li> </ul>
	• Bandwidth — Sends two inter-controller messages to the remote storage subsystem to test the minimum, average, and maximum bandwidth as well as the negotiated link speed of the port on the controller performing the test.
	• Port connections — Shows the port that is being used for mirroring on the local storage subsystem and the port that is receiving the mirrored data on the remote storage subsystem.

## **Minimum Firmware Level**

### 7.84

# Validate Storage Subsystem Security Key

This command validates the security key for a storage subsystem that has full disk encryption (FDE) drives to make sure that the security key is not corrupt.

## **Syntax**

```
validate storageSubsystem securityKey
file="fileName"
passPhrase="passPhraseString"
```

Parameter	Description
file	The file path and the file name that has the security key. Enclose file path and the file name in double quotation marks (" "). For example: file="C:\Program Files\CLI\sup\seckey.slk" Important: The file name must have an extension of <i>.slk</i> .
passPhrase	A character string that encrypts the security key so that you can store the security key in an external file. Enclose the pass phrase in double quotation marks (" ").

## Notes

Your pass phrase must meet these criteria:

- The pass phrase must be between eight and 32 characters long.
- The pass phrase must contain at least one uppercase letter.
- The pass phrase must contain at least one lowercase letter.
- The pass phrase must contain at least one number.
- The pass phrase must contain at least one non-alphanumeric character, for example, <> @

**Note:** If your pass phrase does not meet these criteria, you will receive an error message.

# **Script Commands Listed by Function**

## subsystem Commands

"Create subsystem" on page 84

"Delete subsystem" on page 143

"Enable subsystem Security" on page 168

"Revive array" on page 207

"Set subsystem" on page 224

"Set subsystem Forced State" on page 226

"Show subsystem" on page 293

"Show subsystem Export Dependencies" on page 294

"Show subsystem Import Dependencies" on page 294

"Start subsystem Defragment" on page 331

"Start subsystem Export" on page 331

"Start array import" on page 332

"Start subsystem Locate" on page 333

"Stop array locate" on page 355

### Enhanced Global Mirroring commands

"Activate Enhanced Global Mirroring" on page 61

"Add logical drive to Enhanced Global Mirror Group" on page 66

"Cancel Enhanced Global Mirror Group role reversal" on page 73

"Check Enhanced Global Mirror Group consistency" on page 74

"Clear Enhanced Global Mirroring fault" on page 78

"Check storage subsystem connectivity" on page 77

"Create Enhanced Global Mirror Group" on page 88

"Deactivate Enhanced Global Mirroring" on page 143

"Delete Enhanced Global Mirror Group" on page 144

"Establish Enhanced Global Mirrored Pair" on page 173

"Remove incomplete Enhanced Global Mirrored Pair from an Enhanced Global Mirror Group" on page 189

"Remove volume from Enhanced Global Mirror Group" on page 190

"Reset Enhanced Global Mirror Group statistics" on page 196

"Reset iSCSI IP address" on page 198

"Resume Enhanced Global Mirror Group" on page 201

"Save Enhanced Global Mirror Group statistics" on page 210

"Set Enhanced Global Mirror Group" on page 227

"Show Enhanced Global Mirror Group" on page 295

"Show Enhanced Global Mirror Group synchronization progress" on page 296

"Start Enhanced Global Mirroring synchronization" on page 333

"Suspend Enhanced Global Mirror Group" on page 368

"Test Enhanced Global Mirror Group connectivity" on page 371

## **Controller Commands**

"Clear Drive Channel Statistics" on page 79

"Diagnose Controller" on page 156

"Enable Controller Data Transfer" on page 168 "Reset Controller" on page 197 "Save Controller NVSRAM" on page 212 "Set Controller" on page 232 "Set Controller Service Action Allowed Indicator" on page 235 "Save Disk Drive Channel Fault Isolation Diagnostic Status" on page 212 "Set Drive Channel Status" on page 240 "Set Host Channel" on page 256 "Show Cache Backup Device Diagnostic Status" on page 297 "Show Cache Memory Diagnostic Status" on page 297 "Show Controller" on page 300 "Show Controller Diagnostic Status" on page 301 "Show Controller NVSRAM" on page 301 "Show Drive Channel Statistics" on page 305 "Show Host Interface Card Diagnostic Status" on page 311 "Start Cache Backup Device Diagnostic" on page 334 "Start Cache Memory Diagnostic" on page 335 "Start Storage Subsystem Configuration Database Diagnostic" on page 353 "Start Controller Diagnostic" on page 338 "Start Controller Trace" on page 340 "Start Drive Channel Fault Isolation Diagnostics" on page 341 "Start Drive Channel Locate" on page 342 "Start Host Interface Card Diagnostic" on page 346 "Start iSCSI DHCP Refresh" on page 349 "Stop Cache Backup Device Diagnostic" on page 356 "Stop Cache Memory Diagnostic" on page 356 "Stop Storage Subsystem Configuration Database Diagnostic" on page 356 "Stop Controller Diagnostic" on page 358

"Stop Drive Channel Locate" on page 359

"Stop Host Interface Card Diagnostic" on page 364

### **Database Commands**

"Clear Storage Subsystem Configuration" on page 80

"Clear Storage Subsystem Recovery Mode" on page 82

"Load Storage Subsystem DBM Database" on page 177

"Save Storage Subsystem DBM Database" on page 215

"Save Storage Subsystem DBM Validator Information File" on page 215

"Show Storage Subsystem DBM Database" on page 323

"Start Storage Subsystem Configuration Database Diagnostic" on page 353

"Stop Storage Subsystem Configuration Database Diagnostic" on page 356

## Diagnostic Data Capture Commands

"Clear Storage Subsystem Diagnostic Data Capture" on page 81

"Save Storage Subsystem Diagnostic Data Capture" on page 217

"Show Storage Subsystem Diagnostic Data Capture" on page 324

"Start Storage Subsystem Diagnostic Data Capture" on page 355

## **Disk Pool Commands**

"Create Disk Pool" on page 96

"Delete Disk Pool" on page 146

"Enable Disk Pool Security" on page 169

"Set Disk Pool" on page 236

"Set Disk Pool (Modify Disk Pool)" on page 238

"Show Disk Pool" on page 302

"Start Disk Pool Locate" on page 337

"Stop Disk Pool Locate" on page 359

## **Disk Drive Commands**

"Download Drive Firmware" on page 162

"Replace Disk Drive" on page 195

"Revive Disk Drive" on page 208

"Save Disk Drive Channel Fault Isolation Diagnostic Status" on page 212 "Save Drive Log" on page 213 "Set Disk Drive Hot Spare" on page 241 "Set Disk Drive Service Action Allowed Indicator" on page 242 "Set Drive State" on page 243 "Set Foreign Disk Drive to Native" on page 254 "Show Disk Drive" on page 303 "Show Drive Download Progress" on page 305 "Start Drive Channel Fault Isolation Diagnostics" on page 341 "Start Disk Drive Initialize" on page 342 "Start Drive Locate" on page 343 "Start Disk Drive Reconstruction" on page 343 "Start Secure Disk Drive Erase" on page 352 "Start Drive Channel Locate" on page 342 "Stop Drive Channel Fault Isolation Diagnostics" on page 359 "Stop Drive Locate" on page 360

## **Enclosure Commands**

"Download Enclosure Configuration Settings" on page 164

"Download Environmental Card Firmware" on page 164

"Save Enclosure Log" on page 213

"Set Enclosure Alarm" on page 243

"Set Enclosure Identification" on page 244

"Set Enclosure Service Action Allowed Indicator" on page 244

"Start Enclosure Locate" on page 344

"Stop Enclosure Locate" on page 360

## Enhanced FlashCopy Commands

"Enhanced FlashCopy Group Commands" on page 378

"Enhanced FlashCopy Image Commands" on page 378

"Enhanced FlashCopy Logical Drive Commands"

"Consistency Group Commands" on page 379

### Enhanced FlashCopy Group Commands

"Convert FlashCopy Logical Drive to Enhanced FlashCopy Group" on page 83

"Create Enhanced FlashCopy Group" on page 101

"Delete Enhanced FlashCopy Group" on page 147

"Revive Enhanced FlashCopy Group" on page 208

"Set Enhanced FlashCopy Group Attributes" on page 246

"Set Enhanced FlashCopy Group Media Scan" on page 247

"Set Enhanced FlashCopy Group Repository Logical Drive Capacity" on page 252

"Set Enhanced FlashCopy Group Schedule" on page 247

"Show Enhanced FlashCopy Group" on page 306

"Stop Enhanced FlashCopy Group Pending Enhanced FlashCopy Images" on page 360

#### Enhanced FlashCopy Image Commands

"Create Enhanced FlashCopy Image" on page 107

"Delete Enhanced FlashCopy Image" on page 148

"Resume Enhanced FlashCopy Image Rollback" on page 203

"Show Enhanced FlashCopy Image" on page 307

"Start Enhanced FlashCopy Image Rollback" on page 344

"Stop Enhanced FlashCopy Image" on page 361

"Stop Enhanced FlashCopy Image Rollback" on page 361

### Enhanced FlashCopy Logical Drive Commands

"Create Enhanced FlashCopy Logical Drive" on page 108

"Create Read-Only Enhanced FlashCopy Logical Drive" on page 134

"Delete Enhanced FlashCopy Logical Drive" on page 150

"Rename Enhanced FlashCopy Logical Drive" on page 193

"Resume Enhanced FlashCopy Logical Drive" on page 204

"Revive Enhanced FlashCopy Logical Drive" on page 209

"Set Enhanced FlashCopy Logical Drive Media Scan" on page 251
"Set Enhanced FlashCopy Logical Drive Repository Logical Drive Capacity" on page 252

"Set read-only Enhanced FlashCopy Logical Drive to a Read/Write Logical Drive" on page 278

"Show Enhanced FlashCopy Logical Drives" on page 309

"Stop Enhanced FlashCopy Logical Drive" on page 362

#### **Consistency Group Commands**

"Add Member to Consistency Group" on page 66

"Create Consistency Group" on page 88

"Create Consistency Group Enhanced FlashCopy Image" on page 92

"Create Enhanced FlashCopy Logical Drive" on page 108

"Create Consistency Group Enhanced FlashCopy Logical Drive Mapping" on page 95

"Delete Consistency Group Enhanced FlashCopy Image" on page 145

"Delete Consistency Group Enhanced FlashCopy Logical Drive" on page 145

"Delete Consistency Group" on page 144

"Remove Member Logical Drive from Consistency Group" on page 192

"Resume Consistency Group Enhanced FlashCopy Logical Drive" on page 202

"Resume Enhanced FlashCopy Image Rollback" on page 203

"Set Consistency Group Attributes" on page 230

"Set Consistency Group Enhanced FlashCopy Logical Drive" on page 231

"Show Consistency Group Enhanced FlashCopy Image" on page 298

"Start Consistency Group Enhanced FlashCopy Rollback" on page 337

"Stop Consistency Group Enhanced FlashCopy Logical Drive" on page 358

"Stop Consistency Group Enhanced FlashCopy Rollback" on page 357

"Stop Pending Enhanced FlashCopy Images on Consistency Group" on page 367

#### Enhanced Remote Mirroring Commands

"Activate Enhanced Remote Mirroring Feature" on page 62

"Check Remote Mirror Status" on page 76

"Create Remote Mirror" on page 136

"Deactivate Remote Mirror" on page 143

"Diagnose Remote Mirror" on page 158

"Re-create Enhanced Remote Mirroring Repository Logical Drive" on page 183

"Remove Remote Mirror" on page 192

"Resume Remote Mirror" on page 206

"Set Remote Mirror" on page 278

"Show Enhanced Remote Mirroring Logical Drive Candidates" on page 310

"Show Enhanced Remote Mirroring Logical Drive Synchronization Progress" on page 310

"Start Enhanced Remote Mirroring Synchronization" on page 346

"Suspend Remote Mirror" on page 369

#### Feature Management Commands

"Disable Storage Subsystem Feature" on page 160

"Show Storage Subsystem Features" on page 324

"Show Storage Subsystem Profile" on page 325

## FlashCopy Commands

"Create FlashCopy Logical Drive" on page 110

"Delete FlashCopy Logical Drive" on page 151

"Disable FlashCopy" on page 159

"Resume FlashCopy Rollback" on page 206

"Start FlashCopy Rollback" on page 346

"Re-create FlashCopy" on page 186

"Set FlashCopy Logical Drive" on page 253

"Stop FlashCopy" on page 363

"Stop FlashCopy Rollback" on page 364

#### Host Topology Commands

"Activate Host Port" on page 65

"Activate iSCSI Initiator" on page 65

"Create Host" on page 116

"Create Host Group" on page 117
"Create Host Port" on page 118
"Create iSCSI Initiator" on page 118
"Delete Host" on page 151
"Delete Host Group" on page 152
"Delete Host Port" on page 152
"Delete iSCSI Initiator" on page 152
"Set Host" on page 255
"Set Host Group" on page 257
"Set Host Port" on page 258
"Set iSCSI Initiator" on page 258
"Set iSCSI Target Properties" on page 302
"Show Host Ports" on page 311

#### **iSCSI** Commands

"Create iSCSI Initiator" on page 118 "Delete iSCSI Initiator" on page 152 "Reset Storage Subsystem iSCSI Baseline" on page 200 "Save Storage Subsystem iSCSI Statistics" on page 219 "Set iSCSI Initiator" on page 258 "Set iSCSI Target Properties" on page 259 "Show Current iSCSI Sessions" on page 302 "Show Storage Subsystem Unconfigured iSCSI Initiators" on page 329 "Start iSCSI DHCP Refresh" on page 349 "Stop Storage Subsystem iSCSI Session" on page 364

# Logical Drive Commands (For Standard and Thin Logical Drives and Logical Drives in a Disk pool)

"Check Logical Drive Parity" on page 75 "Clear Logical Drive Reservations" on page 79 "Clear Logical Drive Unreadable Sectors" on page 79 "Create Logical Drive on Disk Pool" on page 131 "Create RAID Logical Drive (Automatic Drive Select)" on page 119 "Create RAID Logical Drive (Free Extent Based Select)" on page 123 "Create RAID Logical Drive (Manual Drive Select)" on page 126 "Delete Logical Drive" on page 153 "Delete Logical Drive on a Disk Pool" on page 154 "Increase Logical Drive Capacity on a Disk Pool" on page 175 "Initialize Thin Logical Drive" on page 176 "Recover RAID Logical Drive" on page 180 "Remove Logical Drive LUN Mapping" on page 187 "Repair Logical Drive Parity" on page 194 "Set Logical Drive Attributes" on page 260 "Set Logical Drive Attributes for a Disk Pool" on page 271 "Set Logical Drive Mapping" on page 276 "Set Thin Logical Drive Attributes" on page 291 "Show Logical Drive" on page 312 "Show Logical Drive Action Progress" on page 313 "Show Logical Drive Performance Statistics" on page 316 "Show Logical Drive Reservations" on page 316 "Show Thin Logical Drive" on page 328 "Start Logical Drive Initialization" on page 351

# **Repository Logical Drive Analysis Commands**

"Check Repository Consistency" on page 76

## **Session Command**

"Set Session" on page 280

## **Performance Read Cache commands**

"Add drives to Performance Read Cache" on page 69

"Change Performance Read Cache application type" on page 74

"Create SSD cache" on page 137

"Delete Performance Read Cache" on page 155

"Enable or disable Performance Read Cache for a volume" on page 170

"Remove drives from Performance Read Cache" on page 188

"Rename Performance Read Cache" on page 194

"Resume Performance Read Cache" on page 207

"Show Performance Read Cache" on page 313

"Start Performance Read Cache performance modeling" on page 352

"Stop Performance Read Cache performance modeling" on page 365

"Suspend Performance Read Cache" on page 370

## Storage Subsystem Commands

"Activate Storage Subsystem Firmware" on page 65 "Autoconfigure Storage Subsystem" on page 69 "Autoconfigure Storage Subsystem Hot Spares" on page 73 "Clear Storage Subsystem Configuration" on page 80 "Clear Storage Subsystem Event Log" on page 82 "Clear Storage Subsystem Firmware Pending Area" on page 82 "Clear Storage Subsystem Recovery Mode" on page 82 "Create Storage Subsystem Security Key" on page 139 "Disable External Security Key Management" on page 159 "Disable Storage Subsystem Feature" on page 160 "Download Storage Subsystem Drive Firmware" on page 166 "Download Storage Subsystem Firmware/NVSRAM" on page 167 "Download Storage Subsystem NVSRAM" on page 167 "Enable External Security Key Management" on page 169 "Export Storage Subsystem Security Key" on page 174 "Import Storage Subsystem Security Key" on page 174 "Load Storage Subsystem DBM Database" on page 177 "Re-create External Security Key" on page 185 "Reset Storage Subsystem Battery Install Date" on page 199 "Reset Storage Subsystem Diagnostic Data" on page 199 "Reset Storage Subsystem iSCSI Baseline" on page 200 "Reset Storage Subsystem Logical Drive Distribution" on page 200 "Reset Storage Subsystem RLS Baseline" on page 201 "Reset Storage Subsystem SAS PHY Baseline" on page 201 "Reset Storage Subsystem SOC Baseline" on page 201 "Save Storage Subsystem Configuration" on page 214 "Save Storage Subsystem DBM Database" on page 215 "Save Storage Subsystem DBM Validator Information File" on page 215 "Save Storage Subsystem Diagnostic Data" on page 216 "Save Storage Subsystem Events" on page 217 "Save Storage Subsystem Firmware Inventory" on page 218 "Save Storage Subsystem iSCSI Statistics" on page 219 "Save Storage Subsystem Performance Statistics" on page 220 "Save Storage Subsystem RLS Counts" on page 220 "Save Storage Subsystem SAS PHY Counts" on page 221 "Save Storage Subsystem SOC Counts" on page 221 "Save Storage Subsystem State Capture" on page 221 "Save Storage Subsystem Support Data" on page 222 "Set Storage Subsystem" on page 281 "Set Storage Subsystem Enclosure Positions" on page 284

"Set Storage Subsystem ICMP Response" on page 285 "Set Storage Subsystem iSNS Server IPv4 Address" on page 285 "Set Storage Subsystem iSNS Server IPv6 Address" on page 286 "Set Storage Subsystem iSNS Server Listening Port" on page 286 "Set Storage Subsystem Learn Cycle" on page 287 "Set Storage Subsystem Redundancy Mode" on page 288 "Set Storage Subsystem Security Key" on page 288 "Set Storage Subsystem Time" on page 289 "Set Storage Subsystem Unnamed Discovery Session" on page 292 "Show Storage Subsystem" on page 317 "Show Storage Subsystem Auto Configure" on page 320 "Show Storage Subsystem DBM Database" on page 323 "Show Storage Subsystem Host Topology" on page 325 "Show Storage Subsystem LUN Mappings" on page 326 "Show Storage Subsystem Negotiation Defaults" on page 326 "Show Storage Subsystem Profile" on page 325 "Show Storage Subsystem Unconfigured iSCSI Initiators" on page 329 "Show Storage Subsystem Unreadable Sectors" on page 326 "Show String" on page 328 "Start Storage Subsystem iSNS Server Refresh" on page 355 "Start Storage Subsystem Locate" on page 353 "Stop Storage Subsystem Drive Firmware Download" on page 367 "Stop Storage Subsystem iSCSI Session" on page 364 "Stop Storage Subsystem Locate" on page 368 "Validate Storage Subsystem Security Key" on page 372

## String Commands

"Show String" on page 328

## **Support Bundle Collection Commands**

"Configure Automatic Support Bundle Collection" on page 83 "Display Automatic Support Bundle Collection Configuration" on page 162

"Display Support Bundle Collection Schedule" on page 162

"Schedule Support Bundle Collection Configuration" on page 223

# VolumeCopy Commands

"Create VolumeCopy" on page 141

"Recopy VolumeCopy" on page 179

"Remove VolumeCopy" on page 193

"Set VolumeCopy" on page 293

"Show VolumeCopy" on page 330

"Show VolumeCopy Source Candidates" on page 330

"Show VolumeCopy Target Candidates" on page 331

"Stop VolumeCopy" on page 368

# Chapter 4. Configuring a Storage Subsystem

Before using script commands to manage a storage subsystem, you should be familiar with these items:

- Controllers
- Disk drives
- Hot spares
- subsystems
- Logical drives
- Disk pools
- RAID technology
- Hosts
- Host groups
- Host bus adapter (HBA) host ports
- Logical unit numbers (LUNs)

Configuring a RAID storage subsystem requires caution and planning to make sure that you define the correct RAID level and configuration for your storage subsystem. You configure a storage subsystem to create logical drives, which are addressable by the hosts, from a collection of disk drives. The commands described in this chapter help you to set up and run a RAID storage subsystem. Additional commands are also available to provide you with more control and flexibility.

**Note:** Many of these commands require a thorough understanding of the firmware as well as an understanding of the network components that need to be mapped. Use the CLI commands and the script commands with caution.

The sections in this chapter show some, but not all, of the CLI commands and the script commands. The commands in this chapter show how you can use the commands to configure a storage subsystem. These presentations do not describe all possible usage and syntax for the commands. (For complete definitions of the commands, including syntax, parameters, and usage notes, see Chapter 3, "Script Commands," on page 27.

This chapter contains examples of CLI command usage and script command usage. The command syntax that is used in the examples is for a host running a Microsoft operating system. As part of the examples, the complete C:\ prompt and the DOS path for the commands are shown. Depending on your operating system, the prompt and path construct will vary.

For most commands, the syntax is the same for all Windows operating systems and UNIX operating systems, as well as for a script file. Windows operating systems, however, have an additional requirement when entering names in a command. On Windows operating systems, you must enclose the name between two back slashes (\ \) in addition to other delimiters. For example, the following name is used in a command running under a Windows operating system: [\"Engineering\"]

For a UNIX operating system, and when used in a script file, the name appears as follows:

["Engineering"]

# **Configuration Concepts**

The following figure shows a host using a computer system [3], and shows the storage subsystem with a controller module [1] and two drive expansion enclosures [2]. The storage subsystem has at least one controller and disk drives. The disk drives are grouped together to form a logical structure. The storage subsystem is physically placed in a cabinet as shown on the left in the figure.





When you configure a storage subsystem, you allocate parts of this logical structure for use by one or more of the hosts that are attached to the storage subsystem. Each disk drive and drive expansion enclosure is numbered to allow the storage to be organized into subsystems and logical drives. This organization helps to provide storage capacity and data protection so that the hosts can safely store and retrieve data from the storage subsystem.

This section defines the physical components and the logical components that are required to organize and configure a storage subsystem. This section also describes how the components relate to each other.

## Controllers

All storage subsystems have one or two controllers. The controllers are circuit-board assemblies that manage data and communication between the hosts and the storage subsystem. The controller manages the data flow between the hosts and the disk drives, keeping track of the logical address of where the data resides. In general, each controller has a processor for performing control operations, NVSRAM for storing the firmware code that operates the storage subsystem, and the buses along which the data flows.

The controllers are located in a controller module, which provides two positions for controllers: slot A and slot B. The script commands identify each controller by the slot in which the controller is installed. If a storage subsystem has only one controller, the controller must be in slot A. Controllers are identified by using model numbers.

Controller models DS4400 and DS4500 use minihubs; two connected to each controller. When viewed from the rear of the controller module, the host-side minihubs are numbered from left-to-right a1, b1, a2, b2. The script commands identify the host channels by using these identifiers. Minihubs also supported the drive-side, where each minihub represents a single channel to the disk drives. When viewed from the rear of the controller module, the disk drive minihubs are numbered from left to right 4, 3, 2, 1. The script commands use these numbers to identify the disk drive channels.

Controllers manage the interface by using the controller firmware to transmit and receive commands between the hosts and the disk drives. Host bus adapters facilitate the communication through whatever interface is selected. Typically, two host bus adapters and two paths are used to optimize redundancy.

These controller modules and controller modules incorporate all host connections and expansion drawer connections into each controller. The host ports must be identified in your command statements to let you complete their network configurations.

The more recent controllers do not use minihubs. These controllers have host ports that are integrated into the controller circuit boards or auxiliary circuit boards that are directly connected to the controller circuit boards.

The following table lists the controller modules that do not use minihubs, the type of host port, and the number of host ports.

Model	Available Host Ports Per Controller	Type of Host Interface
DS3200 controller module (without optional SAS expansion card)	1	SAS
DS3200 controller module (with optional SAS expansion card)	3	SAS
DS3300 controller module	2	iSCSI
DS3400 controller module	2	Fibre Channel
DS3500 series controller module	up to 6	Fibre Channel, SAS, or iSCSI
DCS3700 series controller module	up to 6	Fibre Channel, SAS, iSCSI <sup>1</sup>
DCS3700 storage subsystem with Performance Module Controllers controller module	up to 8	Fibre Channel
DCS3860 series controller module	2	SAS
DS3950 series controller module	4	Fibre Channel and iSCSI
DS4100, DS4200, and DS4300 series controller module	2	Fibre Channel
DS4700 (model 70) series controller module	2	Fibre Channel

Table 9. Host Ports and the Type of Host Interfaces on Controller Modules

Model	Available Host Ports Per Controller	Type of Host Interface
DS4700 (model 72) series controller module	4	Fibre Channel
DS4400, DS4500, and DS4800 series controller module	4	Fibre Channel
DS5020 series controller module	4	Fibre Channel and iSCSI
DS5100 and DS5300 series controller module	up to 8	Fibre Channel and iSCSI

Table 9. Host Ports and the Type of Host Interfaces on Controller Modules (continued)

<sup>1</sup> iSCSI will be supported in a future update.

The DS4800 (models 82, 84, and 88)-series controller module has four host ports and four channel ports on each controller.

The DS3200 controller module has three host ports, which are numbered from left to right: host port 1, host port 2, and host port 3 as shown in the following figure:



Figure 2. DS3200 Host Ports

The DS3300 and DS3400 controller modules have two host ports on each controller, which are numbered from left to right: host port 1 and host port 2 as shown in the following figure:



Figure 3. DS3400 Host Ports

The DS3500 series controller modules can have up to six host ports, depending on the configuration, which are numbered from left to right as shown in the following figure:



Figure 4. DS3500 Series Host Ports

The DCS3700 series controller modules can have up to six host ports, depending on the configuration, as shown in the following figure:



Figure 5. DCS3700 storage subsystem host ports

The DCS3700 storage subsystem with Performance Module Controllers can have up to eight host ports, depending on the configuration, as shown in the following figure:



Figure 6. DCS3700 storage subsystem with Performance Module Controllers series Host Ports with FC HIC installed.

Note: Depending on the HIC, the port number and type differs.

The DCS3860 storage subsystem can have up to two SAS host ports, as shown in the following figure:



Figure 7. DCS3860 storage subsystem with SAS host ports installed.

Label	Controller part
1	Controller A
2	Controller B
3	SAS host port 1
4	SAS host port 2
5	SAS host port 3
6	SAS host port 4

The host ports on the DS4200 and DS4700 (model 72) controller module are numbered from left-to-right on controller B as Ch 1, Ch 2, Ch 3, and Ch 4. Controller A, which is installed upside-down in the enclosure, is numbered from right-to-left in the same sequence, as shown in the following figure:



Figure 8. DS4200 and DS4700 Host Ports

The controller in the DS4800 controller module can have up to four host channels with one port for each channel; up to two drive channels with two ports per channel (for a total of four drive ports); and, up to two Ethernet ports. In the DS4800 (models 82, 84, and 88) controller module, the controllers are stacked one above the other. The top controller is A. The bottom controller is B.

For controller A, the host channel identifiers are a1, a2, a3, and a4 (you will need this reference for the CLI commands and the script commands): the host bus adapter (HBA) host ports are labeled 1, 2, 3, and 4. For controller B, the host channel identifiers are b1, b2, b3, and b4: the host bus adapter (HBA) host ports are labeled 1, 2, 3, and 4.

The HBA host ports are labeled as shown in the following figure:



Figure 9. DS4800 Host Ports

Controller A has drive channels 1 and 2. Controller A drive port pairs labeled 4 and 3 and 2 and 1 connect to controller A drive channels 1, and 2, respectively. Similarly, controller B also has two drive channels, labeled as channels 3 and 4. Controller B drive port pairs labeled 1 and 2 and 3 and 4 connect to controller B drive channel 3 and 4, respectively. Each Ethernet port on a controller can have a unique IP address; however, both Ethernet ports share the same gateway IP address and remote login settings.

The controller in the DS3950 and DS5020 controller module can have up to four host channels with one port for each channel, up to two drive channels with two ports per channel (for a total of four drive ports), and up to two Ethernet ports.

For controller A, the host channel identifiers are a1, a2, a3, and a4 (you will need this reference for the CLI commands and the script commands). For controller B, the host channel identifiers are b1, b2, b3, and b4. The following illustration shows the HBA host ports.



Figure 10. DS3950 and DS5020 Host Ports

Controller A has one drive channel with two drive ports, labeled P1 and P2 from right to left. Controller B has one drive channel with two ports, labeled P1 and P2 from left to right. Each Ethernet port on a controller can have a unique IP address; however, both Ethernet ports share the same gateway IP address and remote login settings.

The DS5300 and DS5100 controller modules can have up to eight host channels with two ports for each channel; up to four drive channels with two ports per channel (for a total of eight drive ports); and, up to two Ethernet ports.

For controller A, the host channel identifiers are a1, a2, a3, a4, a5, a6. a7, and a8 (you will need this reference for the CLI commands and the script commands). The controller A host ports are numbered from the right most port (a1) to the left most port (a8) as shown in Figure 11 on page 395. For controller B, the host channel identifiers are b1, b2, b3, b4, b5, b6, b7, and b8. The controller B host ports are numbered from the left most port (b1) to the right most port (b8) as shown in Figure 11 on page 395.

**Note:** The host ports are physically labeled on the host port covers in groups of 4 (1, 2, 3 and 4).

The HBA host ports are labeled as shown in the following figure:



Figure 11. DS5300 / DS5100 Host Ports

Controller A has drive channels 1, 2, 3, and 4. Controller A drive port pairs, labeled 8 and 7, 6 and 5, 4 and 3, and 2 and 1, connect to controller A drive channels 1, 2, 3 and 4, respectively. Similarly, controller B also has four drive channels, labeled as channels 5, 6, 7, and 8. Controller B drive port pairs, labeled 1 and 2, 3 and 4, 5 and 6, and 7 and 8, connect to controller B drive channels 5, 6, 7 and 8, respectively. Each Ethernet port on a controller can have a unique IP address; however, both Ethernet ports share the same gateway IP address and remote login settings.

## **Disk Drives**

The disk drives provide the actual storage of the data that is sent to the storage subsystem. The disk drives are mounted in an expansion drawer. The expansion drawer is an enclosure that, in addition to the disk drives, contains power supplies, fans, environmental cards, and other supporting components.

Disk drives are located in a storage subsystem by enclosure ID and slot ID. Enclosure ID values are  $\theta$  to 99. In older enclosures, enclosure ID values are set during installation by switches on the rear of the enclosures. In newer enclosures, enclosure ID values are set automatically when the power is applied.

The slot ID is the disk drive position in the expansion drawer. Slot ID values range from 1 to 16. An expansion drawer can contain either 10, 12, 14, or 16 disk drives. The total number of disk drives in a storage subsystem depends on the model of the controller and the type of expansion drawer.

**Note:** A maximum of seven expansion drawers can be on a channel when mixing 14-slot expansion drawers and 16-slot expansion drawers.

The following table lists, by controller number and expansion drawer capacity, the maximum number of disk drives in a storage subsystem.

**Note:** There are maximum limits to the number of Fibre Channel hard drives that you can configure in redundant drive loop pairs. The maximum number of drives dictates a fixed number of supported 14-disk and 16-disk drive enclosure combinations. For more information, please refer to Chapter 1 of the *DS4000 Hard Drive and Storage Expansion Enclosure Installation and Migration Guide*.

Table 10. Maximum Number of Disk Drives

Controller Model	10-Disk Drive Expansion Drawer	12-Disk Drive Expansion Drawer	14-Disk Drive Expansion Drawer	16-Disk Drive Expansion Drawer	12– or 24–Disk Drive Expansion Enclosure	60-Disk Expansion Enclosure
DS3200		48				
DS3300		48				

Controller Model	10-Disk Drive Expansion Drawer	12-Disk Drive Expansion Drawer	14-Disk Drive Expansion Drawer	16-Disk Drive Expansion Drawer	12– or 24–Disk Drive Expansion Enclosure	60-Disk Expansion Enclosure
DS3400		48				
DS3500					192 <sup>1</sup>	
DCS3700						180
DCS3700 storage subsystem with Performance Module Controllers						360
DCS3860						360
DS3950				112		
DS4200			100-110 <sup>1</sup>	112		
DS4300-SCU			14			
DS4300			112	100-110 <sup>1</sup>		
DS4500	220		224	224		
DS4700			112	112		
DS4800			224	224		
DS5020				112		
DS5100				448 <sup>2</sup>		480
DS5300				448		480

Table 1	0. Maximum	Number of	of Disk D	rives i	(continued)
rabie r	0. 1110.11110.111				oon alaada,

<sup>1</sup> Requires a controller for version 7.77.xx.xx and later. The 192 disk drive support premium feature must be purchased to increase the maximum number of drives supported from 96 to 192.

<sup>2</sup> A 448 disk drive support premium feature must be purchased to increase the maximum number of drives supported in a DS5100 subsystem from 256 to 448.

## Hot Spares

A hot spare is a disk drive that acts as a standby in the event that a disk drive containing data fails. The hot spare is a disk drive that has not been assigned to a particular subsystem and, as such, can be used in any subsystem. You can use the hot spare feature with RAID 1, RAID 3, RAID 5, or RAID 6 for subsystems. This concept is not applicable for disk pools. Disk pools have reserve capacity for reconstructed data.

**Note:** Some DS3000 and DS4000 controller firmware versions do not support RAID 6. (Firmware version 7.3x.xx.xx for the DS5000 supports RAID 6.) For more information, see the DS3000 and DS4000 controller firmware readme files or the interoperability matrices.

If a disk drive in an subsystem fails, the controllers automatically replace the failed disk drive with a hot spare. The controllers use redundancy data to reconstruct the data from the failed disk drive onto the hot spare. To be most effective, the disk drive that you assign as a hot spare must have a capacity equal to or greater than the capacity of the largest disk drive in the storage subsystem. The hot spare must

be the same type of disk drive as the disk drive that failed (for example, a Serial Advanced Technology Attachment [SATA] hot spare cannot replace a Fibre Channel hot spare).

You can assign disk drives to act as hot spares manually or have the script commands automatically assign hot spares. If you manually assign a disk drive to be a hot spare, you must identify the disk drive by enclosure ID and slot ID. When you let the script commands automatically assign hot spares, you must enter the number of hot spares that you want in the storage subsystem. The maximum number of hot spare disk drives per storage subsystem is 15.

#### Subsystems

A subsystem is a set of disk drives that are logically grouped together by the controllers in a storage subsystem. After you create an subsystem, you can create one or more logical drives in the subsystem. A subsystem is identified by a sequence number that is defined by the controller firmware when you created the subsystem.

**Note:** You can give a unique name to a subsystem if your Storage Manager version is 10.77 or later with controller firmware version 7.77.xx.xx or later.

To create a subsystem, you must define the capacity and the RAID level.

Capacity is the size of the subsystem. Capacity is determined by the number of disk drives that you assign to the storage subsystem. You can use only unassigned disk drives to create a subsystem. (In this programming guide, the storage space on unassigned disk drives constitutes the *unconfigured capacity* of a storage subsystem.)

*Free capacity* is a contiguous region of unassigned capacity in a designated subsystem. Before you create a new logical drive in an subsystem, you need to know the free capacity space so that you can determine the size of the logical drive.

The number of disk drives you can include in an subsystem is constrained by the capacity of each disk drive. the following table relates the capacity of individual disk drives to the number of disk drives that you can include in a storage subsystem.

	Controller firmware PRIOR to 7.10 (DS4000) or 6.xx (DS3000)	Controller firmware 7.10 or later (DS3000 / DCS3700 / DCS3700 storage subsystem with Performance Module Controllers / DCS3860/ DS4000 / DS5000 only)
Number of drives per subsystem	30 for all RAID levels	<ul> <li>RAID 3 and RAID 5: 30 drives</li> <li>RAID 0 and RAID 1: Maximum of drives supported for a given storage subsystem</li> </ul>
Maximum logical drive size	2 TB	Number of drives in an subsystem multiplied by the size of the smallest capacity drive in the subsystem

Table 11. Maximum Number of Disk Drives in a Subsystem Based on Capacity

## Disk Pools About disk pools

A disk pool is a set of disk drives that is logically grouped together in the storage subsystem. The disk drives in each disk pool must be of the same drive type and drive media type, and they must be similar in size. As with an subsystem, you can create one or more logical drives in the disk pool. However, the disk pool is different from the logical drive group in the way that the data is distributed across the disk drives that comprise the disk pool.

In an subsystem, the data is distributed across the disk drives based on a RAID level. You can specify the RAID level when you create the subsystem. The data for each logical drive is written sequentially across the set of disk drives in the subsystem.

In a disk pool, the storage management software distributes the data for each logical drive randomly across a set the disk drives in the disk pool. Each disk pool must have a minimum of eleven disk drives. Although there is no limit on the maximum number of disk drives that can comprise a disk pool, the disk pool cannot contain more disk drives than the maximum limit for each storage subsystem. The storage management software automatically configures the RAID level (which is RAID 6 - 8D+P+Q) when you create the disk pool. You cannot set or change the RAID level of disk pools or the logical drives in the disk pools.

**Note:** Because disk pools can coexist with subsystems, a storage subsystem can contain both disk pools and subsystems.

#### Benefits of disk pools

• Better utilization of disk drives - When you add disk drives to a storage subsystem, the storage management software automatically detects the disk drives and prompts you to create a single disk pool or multiple disk pools based on the drive type and the current configuration. If disk pools were previously defined, the storage management software provides the option of adding the compatible disk drives to an existing disk pool. When new disk drives are added to an existing disk pool, the storage management software automatically redistributes the data across the new capacity, which now includes the new disk drives that you added. The data in the logical drives remain accessible when you add the disk drives to the disk pool. When you delete disk pool logical drives, the capacity of those logical drives is added to the total usable capacity of the disk pool and therefore it can be reused.

**Note:** You have the option to manually create a disk pool, if you prefer not to proceed with the automatic disk pool creation process. Refer to the "Create Disk Pool" on page 96 topic for more information.

- **Reduced Hot Spots** A host might access some disk drives in the subsystem for data more frequently than other disk drives because of the sequential manner in which the data is written to the disk drives. This frequency of access to disk drives creates hot spots in the subsystem. In a disk pool, the hot spots are significantly reduced because of the random manner in which the data is spread across a large number of disk drives. The reduction of hot spots in the disk pool improves performance of the storage subsystem.
- **Faster Reconstruction of Data** Disk pools do not use hot spare disk drives for data protection like an subsystem does. Instead of hot spare disk drives, disk pools use spare capacity within each drive that comprises the disk pool.

In hot spare drive coverage, the maximum drive IOPS (I/O per second) limits the speed of reconstruction of data from the failed drive to the hot spare drive. In a disk pool, the reconstruction of data is much faster because the spare capacity in all of the disk drives that comprise the disk pool is used. Additionally, the data to reconstruct after a drive failure is reduced because the data is spread randomly across more disk drives in a disk pool.

Faster reconstruction of data in a disk pool also reduces the risk of additional drive failures during a reconstruction operation. For example, consider a drive failure in a RAID Level 5 subsystem that is comprised of three disk drives. The time it takes to reconstruct the data from the failed drive is relatively longer for an subsystem. During the reconstruction of data, if another drive fails in this subsystem, data loss occurs. Unlike subsystems, the time period for which the disk pool is exposed to multiple drive failures during a reconstruction operation is significantly reduced.

• **Reduced Maintenance** – You can configure the storage management software to send alert notifications when the configured capacity of a disk pool is reaching a specified percentage of free capacity. Additionally, you do not need to manage any hot spare disk drives. You can replace a set of disk drives during a scheduled maintenance of the storage subsystem.

#### Limitations of disk pools

The following are limitations of disk pools:

- 10 disk pools per storage subsystem.
- The maximum size of disk pool is 512 TB with controllers having 1 GB cache per controller and FC and SAS host Interface ports installed.
- The maximum size of disk pool is 1024 TB with controllers having 2 GB cache per controller and FC and SAS host Interface ports installed.
- The maximum size of disk pool is 256 TB with controllers having 1 GB or 2 GB cache per controllers and iSCSI host Interface ports installed.
- The maximum size of the logical drive in a disk pool is 64 TB.
- Disk pool cannot be imported from, or exported to another storage subsystem.

## **Logical Drives**

A logical drive is a logical component (object) that is the basic structure that is created on the storage subsystem to store data. A logical drive is a contiguous subsection of an subsystem that is configured to meet application needs for data availability and I/O performance. The IBM Storage Manager client software administers a logical drive as if the logical drive is one "disk drive" for data storage. Logical drives are identified by names or labels that users choose. The logical drive names can be any combination of alphanumeric characters, hyphens (-), and underscores (\_). The maximum length of a logical drive name is 30 characters.

The script commands support the following types of logical drives:

- **Standard logical drive** A logical structure that is the principal type of logical drive for data storage. A standard logical drive is the most common type of logical drive in a storage subsystem.
- Access logical drive A factory-configured logical drive in a storage area network (SAN) environment that is used for communication between the IBM Storage Manager client software and the storage subsystem controller. The access

logical drive uses a logical unit number (LUN) address and consumes 20 MB of storage space. The 20 MB of access logical drive storage space is not available for data storage.

• Thin logical drive – A thin logical drive lets you create large virtual logical drive with small physical storage allocations that can grow over time to meet increased capacity demands. As storage demands increase, you can increase the amount of physical storage capacity as it is needed. You can create a thin logical drive in a disk pool only.

**Note:** You must use the access logical drive only for in-band-managed storage subsystems.

- **FlashCopy logical drive** A logical point-in-time image of another logical drive. A FlashCopy logical drive is the logical equivalent of a complete physical copy; however, it is not an actual, physical copy. Instead, the firmware tracks only the data blocks that are overwritten and copies those blocks to a FlashCopy repository logical drive.
- FlashCopy repository logical drive A special logical drive in the storage subsystem that is created as a resource for a FlashCopy logical drive. A FlashCopy repository logical drive contains FlashCopy data and copy-on-write data for a particular FlashCopy logical drive.
- **Base logical drive** A standard logical drive from which you create a FlashCopy logical drive. The term "base logical drive" is used only to show the relationship between a standard logical drive from which you are taking the point-in-time image and a FlashCopy logical drive.
- **Primary logical drive** A standard logical drive in a Enhanced Remote Mirroring relationship. The primary logical drive accepts host data transfers and stores application data. When you first create the mirror relationship, data from the primary logical drive is copied in its entirety to the associated secondary logical drive.
- Secondary logical drive A standard logical drive in a Enhanced Remote Mirroring relationship that maintains a mirror (or copy) of the data from its associated primary logical drive. The secondary logical drive remains unavailable to host applications while mirroring is underway. In the event of a disaster or a catastrophic failure of the primary site, a system administrator can promote the secondary logical drive to a primary role.
- **Mirror repository logical drive** A special logical drive in a Enhanced Remote Mirroring configuration that is created as a resource for each controller in both the local storage subsystem and the remote storage subsystem. The controller stores mirroring information on this logical drive, including information about remote writes that are not yet complete. A controller can use this information to recover from controller resets and accidental power shutdown of the storage subsystems.

**Note:** FlashCopy Logical Drive and Enhanced Remote Mirroring are premium features that you must activate before you can use them. For more information about FlashCopy logical drives, see Chapter 5, "Using the FlashCopy Premium Feature," on page 443. For more information about Enhanced Remote Mirroring, see Chapter 8, "About Enhanced Remote Mirroring premium feature," on page 495.

The number and capacity of the logical drives in your storage subsystem depends on the type of controller in the storage subsystem. The following tables list the maximum values for the logical drives in a storage subsystem by specific controllers.

Specifications	DS3200, DS3300, DS3400	DS3500 series	DS4200, DS4700	DS4800	DS5100, DS5300	DS3950, DS5020
Maximum logical drives per storage subsystem	256	5121	1024	2048	2048	1024
Maximum logical 256 drivers per partition		256 <sub>2</sub>	256	256	256	256
Maximum logical drive size (less than 12 GB for overhead)	SeeTable 11 on page	GeeTable 11 on page 397.				
Maximum disk drives per subsystem using RAID 5	See Table 11 on page 397.					
Maximum remote N/A mirrors		163	64	128	128	64

Table 12. Logical Drive Specifications by Supported Controller Model (For DS3000, DS4000. DS5000)

Table	13.	Logical	Drive	Specifications	by	Supported	Controller	Model	(For	DCS	Series)
-------	-----	---------	-------	----------------	----	-----------	------------	-------	------	-----	---------

Specifications	DCS3700	DCS3700 subsystem with Performance Module Controllers	DCS3860		
Maximum logical drives per storage subsystem	512	2048	2048		
Maximum logical drivers per partition	256	256	256		
Maximum logical drive size (less than 12 GB for overhead)	SeeTable 11 on page 397.				
Maximum disk drives per subsystem using RAID 5	See Table 11 on page 397.				
Maximum remote mirrors	16 16 16				

**Note:** Addressable capacities greater than 2 TB are possible with some operating systems. Currently these operating systems are Windows, Solaris 9, Solaris 10, AIX 5.2, AIX 5.3, HP-UX 11.23, and Linux 2.6.

A Windows operating system and a Linux 2.6 operating system using 500 GB disk drives can have a maximum logical drive size of 112 TB (224 x 500) when controller firmware version 7.10.xx.xx or later is installed in the DS4000 storage subsystem, or when controller firmware version 7.30.xx.xx is installed in the DS5000 storage subsystem.

<sup>1</sup> Firmware 7.77.xx.xx and later supports 512 logical drives per storage subsystem. Firmware 7.70.xx.xx supports only 25 logical drives per subsystem.

<sup>2</sup> Firmware 7.77.xx.xx and later supports 256 logical drives per partition. 7.70.xx.xx supports only 32 logical drives per partition

<sup>3</sup> Requires firmware 7.77.xx.xx and later and enabling of the 16 remote mirrors premium feature.

## **RAID Levels**

The RAID level defines a storage architecture in which the storage capacity on the disk drives in an subsystem is separated into two parts: part of the capacity stores the user data, and the remainder stores redundant or parity information about the user data. The RAID level that you choose determines how user data is written to and retrieved from the disk drives. Using the script commands, you can define five

RAID levels: RAID 0, RAID 1, RAID 3, RAID 5, and RAID 6. Each level provides different performance and protection features.

RAID 0 provides the fastest storage access but does not provide any redundant information about the stored data. RAID 1, RAID 3, RAID 5, and RAID 6 write redundancy information to the disk drives to provide fault tolerance. The redundancy information might be a copy of the data or an error-correcting code that is derived from the data. In RAID 1, RAID 3, RAID 5, or RAID 6 configurations, if a disk drive fails, the redundancy information can be used to reconstruct the lost data. Regardless of the RAID level that you choose, you can configure only one RAID level across each subsystem. All redundancy information for an subsystem is stored within the subsystem. The following table lists the RAID levels and describes the configuration capabilities of each level.

RAID Level	Configuration
0	<b>Non-redundant striping mode</b> – Use this level for high-performance needs. RAID 0 does not provide any data redundancy. RAID 0 stripes data across all of the disk drives in the subsystem. If a single disk drive fails, all of the associated logical drives fail and all data is lost. RAID 0 is suited for noncritical data. It is not recommended for high-availability needs.
1	<b>Striping mirroring mode</b> – RAID 1 uses disk drive mirroring to create an exact copy from one disk drive to another disk drive. A minimum of two disk drives are required; one for the user data, and one for the mirrored data. RAID 1 offers high performance and the best data availability.
	Data is written to two disk drives simultaneously. If one disk drive in a pair fails, the system can instantly switch to the other disk drive without any loss of data or service. Only half of the disk drives in the subsystem are available for user data. If a single disk drive fails in a RAID 1 subsystem, all of the associated logical drives become degraded, but the mirror disk drive provides access to the data. RAID 1 can survive multiple disk drive failures as long as no more than one failure occurs per mirrored pair. If a disk drive pair fails, all of the associated logical drives fail, and all data is lost.
3	<b>High-bandwidth mode</b> – RAID 3 stripes both user data and redundancy data (in the form of parity) across the disk drives. The equivalent of the capacity of one disk drive is used for the redundancy data. RAID 3 works well for large data transfers in applications, such as multimedia or medical imaging, that write and read large sequential chunks of data.
	If a single disk drive fails in a RAID 3 subsystem, all of the associated logical drives become degraded, but the redundancy data lets the data be reconstructed. If two or more disk drives fail, all of the associated logical drives fail, and all data is lost.

Table 14. RAID Level Configurations

Table 14. RAID Level Configurations (continued)

RAID Level	Configuration
5	<b>High I/O mode</b> – RAID 5 stripes both user data and redundancy data (in the form of parity) across the disk drives. The equivalent of the capacity of one disk drive is used for the redundancy data. RAID 5 works well for multiuser environments, such as databases or file system storage, where typical I/O size is small, and a high proportion of read activity exists.
	If a single disk drive fails in a RAID 5 subsystem, all of the associated logical drives become degraded, and the redundancy data permits the data to be reconstructed. If two or more disk drives fail, all of the associated logical drives fail, and all data is lost.
6	Data protection or continuous access mode – RAID 6 stripes both user data and redundancy data (in the form of parity) across the disk drives. A minimum of five disk drives are required for a RAID 6 subsystem. The equivalent capacity of two disk drives is used for the redundancy data. Two different algorithms calculate redundancy data, which are in the form of both a P parity and a Q parity. RAID 6 works well for larger disk drive sizes. Recovery from a second disk drive failure in the same subsystem is possible. If two disk drives fail in a RAID 6 subsystem, all of the associated logical drives become degraded, but the redundancy data permits the data to be reconstructed. If three or more disk drives fail, all of the associated logical drives fail, and all data is lost.

Note: RAID 6 is supported only with the following requirements:

- Controller firmware version 7.xx.xx.and higher must be installed.
- The controllers must be capable of supporting the P+Q calculation.
- You must purchase and install a premium feature key to use RAID 6 and to use dynamic RAID-level migration. (See "Set subsystem" on page 224 to set your subsystem to RAID 6.)
- The DS4100, DS4300, DS4400, DS4500, and DS4800 storage subsystems do not support RAID 6. The DS3500 series, DS3950, DS4200, DS4700, DS5020, DS5100, and DS5300 storage subsystems do support RAID 6.

## Hosts

A host is a computer that is attached to the storage subsystem for accessing the logical drives in the storage subsystem. The host is attached to the storage subsystem through HBA host ports, which are connectors on host bus adapter circuit boards. You can define specific Volume-to-LUN mappings to an individual host or assign the host to a host group that shares access to one or more logical drives. Hosts are identified by names or labels that users choose. The host name can be any combination of alphanumeric characters, hyphens, and underscores. The maximum length of the host name is 30 characters.

In addition to a host name, some script commands require you to identify a host by its "type." A host type identifies the operating system under which the host is running (such as Windows, Solaris, or Linux). Specifying the host type lets the controllers in the storage subsystem adapt their behavior (such as LUN reporting and error conditions) to the operating system of the host that is sending the information. Host types are identified by a label or an index number that is generated by the controller firmware.

### **Host Groups**

A host group is a topological element that you can define if you want to designate a collection of hosts that will share access to the same logical drives. A host group is a logical entity. Host groups are identified by names or labels that users choose. The host group name can be any combination of alphanumeric characters with a maximum length of 30 characters.

## **Host Bus Adapter Host Ports**

A host bus adapter (HBA) provides the physical connection from the host to the storage subsystem. The host port is a physical connector on an HBA. The HBA is a circuit board that is installed in the host. The HBA can have one or more host ports. Each host port is identified by a unique, 16-byte World-Wide Identifier (WWID). If the HBA has more than one host port, each host port has a unique ID.

When you first turn on the power to a storage subsystem, the Storage Manager software automatically detects the HBA host ports. Initially, all detected host ports belong to a default group. You can use script commands to identify the WWIDs on a storage subsystem and, if you choose, change them. If you move an HBA host port, you must remap any Volume-to-LUN mappings. Access to your data is lost until you remap the logical drives.

The maximum number of HBA host ports that you can logically define for your storage subsystem depends on the type of controller in the storage subsystem. The following table lists the maximum number of HBA host ports that you can define.

Controller Models	Maximum Number of Host Ports
DS3200	64
DS3300	64
DS3400	64
DS3500	256
DCS3700	256
DCS3700 storage subsystem with Performance Module Controllers	1024
DCS3860	1024
DS3950	640
DS4100	256
DS4300	256
DS4200	256 (before controller firmware 7.xx.xx.xx)
	1024 (with controller firmware 7.xx.xx.and higher)
DS4400	512
DS4500	512

Table 15. Maximum Number of HBA Host Ports per Controller

DS4700	256 (before controller firmware 7.xx.xx)
	640 (with controller firmware 7.xx.xx.and higher)
DS4800	512 (before controller firmware 7.xx.xx)
	2048 (with controller firmware 7.xx.xx.and higher)
DS5020	640
DS5100	2048
DS5300	2048

Table 15. Maximum Number of HBA Host Ports per Controller (continued)

## **Logical Unit Numbers**

In the context of the CLI commands and the script commands, a logical unit number (LUN) is a unique value that identifies the logical drives in a storage subsystem. The hosts identify the logical drives that they want to access using the LUN values. When you create a logical drive, the firmware assigns the LUN values, or you can assign LUN values when you enable the Storage Partitioning premium feature. A logical drive can have only one LUN and can be mapped to only one host or host group. Each host has unique addressing capability. That is, when more than one host accesses a storage subsystem, each host might use the same LUN to access different logical drives. The LUNs might be the same, but the logical drives are different. If you are mapping to a host group, the LUN that you specify must be available on every host in the host group.

# **Configuring a Storage Subsystem**

When you configure a storage subsystem, you want to maximize the data availability by making sure that the data is quickly accessible while maintaining the highest level of data protection possible. The speed by which a host can access data is affected by these items:

- The RAID level for the subsystem
- The settings for the segment size and the cache size
- Whether the cache read prefetch capability is turned on or turned off

Data protection is determined by the RAID level, hardware redundancy (such as global hot spares), and software redundancy (such as the Enhanced Remote Mirroring premium feature and the FlashCopy Logical Drive premium feature).

In general, you configure a storage subsystem by defining an subsystem and its associated RAID level, defining the logical drives, and defining which hosts have access to the logical drives. This section explains how to use the script commands to perform the general steps to create a configuration from an subsystem of disk drives.

# **Determining What Is on Your Storage Subsystem**

Even when you create a configuration on a storage subsystem that has never been configured, you still need to determine the hardware features and software features that are to be included with the storage subsystem. When you configure a

storage subsystem that has an existing configuration, you must make sure that your new configuration does not inadvertently alter the existing configuration, unless you are reconfiguring the entire storage subsystem. For example, consider the case where you want to create a new subsystem on unassigned disk drives. Before you create a new subsystem, you must determine which disk drives are available. The commands that are described in this section help you to determine the components and the features in your storage subsystem.

The command that returns general information about the storage subsystem is the **show storageSubsystem** command. This command returns information about the components and properties of your storage subsystem, including these items:

- A detailed profile of the components and features in the storage subsystem
- The age of the battery
- The default host type (which is the current host type)
- Other host types that you can select
- The hot spare locations
- The identifiers for enabled features
- The logical component profiles and the physical component profiles
- The time to which both controllers are set
- · The controller that currently owns each logical drive in the storage subsystem

To return the most information about the storage subsystem, run the **show storageSubsystem** command with the **profile** parameter. This example shows the complete CLI command and script command running on a Windows operating system:

c:\...\client>smcli 123.45.67.89 -c "show storageSubsystem profile;"

This example identifies the storage subsystem by the example IP address 123.45.67.89. You can also identify the storage subsystem by name.

The **show storageSubsystem profile** command returns detailed information about the storage subsystem. The information appears in several display screens. You might need to increase the size of your display buffer to see all of the information. Because this information is so detailed, you might want to save the output to a file. To save the output to a file, enter the command as shown in this example:

c:\...\client>smcli 123.45.67.89 -c "show storageSubsystem profile;"
-o c:\folder\storageSubsystemprofile.txt

In this example, the name folder is the folder in which you choose to place the profile file, and storageSubsystemprofile.txt is the name of the file. You can choose any folder and any file name.

**Note:** Attention: Possible loss of data – When you are writing information to a file, the script engine does not check to determine if the file name already exists. If you choose the name of a file that already exists, the script engine writes over the information in the file without warning.

Chapter 12, "Examples of Information Returned by the Show Commands," on page 551 shows the type of information returned. When you save the information to a file, you can use the information as a record of your configuration and as an aid during recovery.

To return a brief list of the storage subsystem features and components, use the **summary** parameter. The command looks like this example:

c:\...\client>smcli 123.45.67.89 -c "show storageSubsystem summary;"

Following is the type of information that is returned by the **show storageSubsystem** command with the **summary** parameter. PROFILE FOR STORAGE SUBSYSTEM: DS5300 middle (1/27/09 3:23:25 PM) SUMMARY-----Number of controllers: 2 High performance tier controllers: Enabled Number of subsystems: 5 RAID 6: Enabled Total number of logical drives used: 10 Number of standard logical drives: 9 Number of access logical drives: 1 Total number of logical drives allowed: 2048 FlashCopy Logical Drives: Enabled Number of flashcopies used: 0 Number of flashcopies allowed: 2 Number of flashcopies allowed per base logical drive: 2 Disabled/Deactivated Enhanced Remote Mirroring: Number of mirrors used: 0 Number of mirrors allowed: 0 VolumeCopy: Disabled Number of copies used: 0 Number of copies allowed: 0 Number of drives: 32 Mixed drive types: Enabled Current drive type(s): Fibre (5), Fibre/FDE (21), Serial ATA (SATA) (6) Total hot spare drives: 2 Standby: 2 In use: 0 Full Disk Encryption: Enabled Security key identifier: 27000000600A0B800029ECE60000DE9E4940E373 Number of drive enclosures: 2 Number of drive enclosures allowed: 28 Storage Partitioning: Enabled Number of partitions used: 2 Number of partitions allowed: 2 Number of logical drives allowed per partition: 256 Access logical drive: LUN 31,31,31 (see Mappings section for details) Default host OS: Windows 2000/Server 2003/Server 2008 Non-Clustered (Host OS index 2) Current configuration 07.50.00.00 Firmware version:

NVSRAM version: N1818D53R1050V03 EMW version: 10.50.G5.02 AMW version: 10.50.G5.02 Pending configuration Staged firmware download supported: Yes Firmware version: None NVSRAM version: None Transferred on: None Disabled Controller enclosure audible alarm: NVSRAM configured for batteries: Yes Start cache flushing at (in percentage): 80 Stop cache flushing at (in percentage): 80 Cache block size (in KB): 8 Media scan frequency (in days): 1 Failover alert delay (in minutes): 5 Feature enable identifier: 303030323420303035353420493FECB1 Feature pack: DS5300 Feature pack submodel ID: 93 Storage Subsystem world-wide identifier (ID): 600A0B800029ECE60000000493FECAB CONTROLLERS-----Number of controllers: 2 Controller in Enclosure 85, Slot A Status: Online 0 Current configuration Firmware version: 07.50.00.00 Appware version: 07.50.00.00 Bootware version: 07.50.00.00 NVSRAM version: N1818D53R1050V03 Pending configuration Firmware version: None Appware version: None Bootware version: None NVSRAM version: None Transferred on: None Current ID (ALPA) On drive channel 1: 125/0x1 Replacement part number: memSpyTesting 9 Model name: 7000-88 Board ID: 7091 Submodel ID: 93 Product ID: 1818 Revision: 0730 Replacement part number: memSpyTesting 9 Part number: 32847-10 Serial number: SF74700554 Vendor: IBM Date of manufacture: January 24, 2008 Data Cache Total present: 8192 MB 8192 MB Total used: Data Cache Module Status: Optimal Location: Controller A, Slot 1

Capacity:

1,024 MB

Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196510 Manufacturer: Samsung Date of manufacture: October 29, 2007 Data Cache Module Status: Optimal Controller A, Slot 2 Location: 1,024 MB Capacity: Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196542 Manufacturer: Samsung Date of manufacture: October 29, 2007 Data Cache Module Status: Optimal Controller A, Slot 3 Location: Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196568 Manufacturer: Samsung October 29, 2007 Date of manufacture: Data Cache Module Status: Optimal Location: Controller A, Slot 4 Capacity: 1,024 MB Not Available Replacement Part number: Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196524 Manufacturer: Samsung October 29, 2007 Date of manufacture: Data Cache Module Status: Optimal Controller A, Slot 5 Location: Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196503 Manufacturer: Samsung Date of manufacture: October 29, 2007 Data Cache Module Status: Optimal Location: Controller A, Slot 6 Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196519 Manufacturer: Samsung Date of manufacture: October 29, 2007 Data Cache Module Status: Optimal Controller A, Slot 7 Location: 1,024 MB Capacity: Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196523 Manufacturer: Samsung Date of manufacture: October 29, 2007 Data Cache Module

Status: Optimal Location: Controller A, Slot 8 Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196530 Manufacturer: Samsung Date of manufacture: October 29, 2007 Processor cache: Total present: 2044 MB Processor Cache Module Status: Optimal Location: Controller A, Slot 1 2,048 MB Capacity: Replacement Part number: Not Available Manufacturer Part Number: Not Available Part Number: Not Available Serial number: Not Available Manufacturer: Not Available Date of manufacture: Not available Cache Backup Device Status: Optimal USB flash drive Type: Location: Controller A, Connector USB 4 3,919 MB Capacity: Product ID: TS4GUFM-V Part number: Not Available Serial number: 000000000003B Revision level: 0.00 Not Available Manufacturer: Date of manufacture: Not available Optimal Status: Type: USB flash drive Controller A, Connector USB 3 Location: Capacity: 3,919 MB TS4GUFM-V Product ID: Part number: Not Available Serial number: 00000000000007 Revision level: 0.00 Manufacturer: Not Available Date of manufacture: Not available Host Interface Board Optimal Status: Location: Slot 2 Type: Fibre channel Number of ports: 4 0801 Board ID: Replacement part number: Not Available Part number: PN L2-25023-02 Serial number: SN SF84300044 Vendor: VN LSI Date of manufacture: November 1, 2008 Host Interface Board Status: Optimal Location: Slot 1 Type: Fibre channel Number of ports: 4 Board ID: 0801 Replacement part number: Not Available Part number: PN L2-25023-02 Serial number: SN SF84300041 Vendor: VN LSI Date of manufacture: November 1, 2008 Date/Time: Tue Jan 27 15:19:38 EST 2009

Associated Logical Drives (\* = Preferred Owner):

FC\_R3\_A\*, FDE\_R6\_A\*, Sata\_Raid6\_A\*, Secure\_Middle\_R6\_A\* Controller DNS/Network name: XBB2\_MiddleA Remote login: Enabled\* Ethernet port: 1 Up 00:a0:b8:29:ed:8a Auto-negotiate 10 Mbps Link status: MAC address: Negotiation mode: Port speed: Port speed: Duplex mode: Half duplex v4 settings: Enabled IPv4 settings: Network configuration: Static IP address: 192.168.70.70 Subnet mask: 255.255.255.0 Gateway: 0.0.0.0 v6 settings: Enabled Auto-configuration: Stateless Local IP address: FE80:0000:0000:02A0:B8FF:FE29:ED8A IPv6 settings: Link status: MAC address: Port speed: Duplex mode: IPv4 settings: Network configuration Duplex details Port speed: Network configuration Duplex details Port speed: Network configuration Duplex details Network configuration Duplex details Network configuration Duplex details Network configuration Duplex details Network configuration Ethernet port: v4 settings: Enable: Network configuration: Static IP address: 192.168.129.101 255.255.255.0 Gateway: 0.0.0.0 IPv6 settings: Enabled 
 Auto-configuration:
 Static

 Local IP address:
 FE80:0000:0000:0000:0000:0000:0000
 \*NVSRAM setting, may be overridden by DHCP/Bootp server setting Duad . o intonf г: L

Channel:	1
Port:	8, 7
Current ID:	125/0x1
Maximum data rate:	4 Gbps
Current data rate:	2 Gbps
Data rate control:	Auto
Link status:	Up
Drive interface:	Fibre
Channel:	2
Port:	6, 5
Current ID:	125/0x1
Maximum data rate:	4 Gbps
Current data rate:	4 Gbps
Data rate control:	Auto
Link status:	Up
Drive interface:	Fibre
Channel:	3
Port:	4, 3, Out
Current ID:	125/0x1
Maximum data rate:	4 Gbps
Current data rate:	4 Gbps
Data rate control:	Auto
Link status:	Up

Dri	ve interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 4 2, 1, 00 125/0x1 4 Gbps 4 Gbps Auto Up	ut
Dri	ve interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 5 1, 2 125/0x1 4 Gbps 2 Gbps Auto Up	
Dri	ve interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 6 3,4 125/0x1 4 Gbps 4 Gbps Auto Up	
Dri	ve interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 7 5, 6, 00 125/0x1 4 Gbps 4 Gbps Auto Up	ut
Dri	ve interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 8 7, 8, 00 125/0x1 4 Gbps 4 Gbps Auto Up	ut
Hos	t interface: Channel: Current ID: Preferred ID: NL-Port ID: Maximum data rate: Current data rate: Data rate control: Link status: Topology: World-wide port ident World-wide node ident Part type:	tifier: tifier:	Fibre 1 Not applicable/0xFFFFFFF 125/0x1 0x010200 8 Gbps 4 Gbps Auto Up Fabric Attach 20:1e:00:a0:b8:29:ed:8a 20:0e:00:a0:b8:29:ed:8a PM8032 revision 5
Hos	t interface: Channel: Current ID: Preferred ID: NL-Port ID: Maximum data rate: Current data rate: Data rate control:		Fibre 2 Not applicable/0xFFFFFFF 1/0xE8 0xFFFFFF 8 Gbps 8 Gbps Auto

Link status: Down Topology: Not Available 20:2e:00:a0:b8:29:ed:8a World-wide port identifier: World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Host interface: Fibre Channel: 3 Current ID: Not applicable/0xFFFFFFF Preferred ID: 2/0xE4 NL-Port ID: 0xFFFFFF Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available World-wide port identifier: 20:3e:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Host interface: Fibre Channel: 4 Current ID: Not applicable/0xFFFFFFF Preferred ID: 3/0xE2 0xFFFFFF NL-Port ID: Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Not Available Topology: 20:4e:00:a0:b8:29:ed:8a World-wide port identifier: World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a PM8032 revision 5 Part type: Fibre Host interface: Channel: 5 Current ID: Not applicable/0xFFFFFFF Preferred ID: 125/0x1 NL-Port ID: 0x010600 Maximum data rate: 8 Gbps Current data rate: 4 Gbps Data rate control: Auto Link status: Up Topology: Fabric Attach World-wide port identifier: 20:5e:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Fibre Host interface: Channel: 6 Current ID: Not applicable/0xFFFFFFF Preferred ID: 13/0xD2 NL-Port ID: 0xFFFFFF Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available World-wide port identifier: 20:6e:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Host interface: Fibre Channel: 7 Current ID: Not applicable/0xFFFFFFF 14/0xD1 Preferred ID: NL-Port ID: 0xFFFFFF

Maximum data rate: Current data rate: Data rate control: Link status: Topology: World-wide port identifier: World-wide node identifier: Part type:	8 Gbps 8 Gbps Auto Down Not Available 20:7e:00:a0:b8:29:ed:8a 20:0e:00:a0:b8:29:ed:8a PM8032 revision 5
Host interface: Channel: Current ID: Preferred ID: NL-Port ID: Maximum data rate: Current data rate: Data rate control: Link status: Topology: World-wide port identifier: World-wide node identifier: Part type:	Fibre 8 Not applicable/0xFFFFFFF 15/0xCE 0xFFFFF 8 Gbps 8 Gbps Auto Down Not Available 20:8e:00:a0:b8:29:ed:8a 20:0e:00:a0:b8:29:ed:8a PM8032 revision 5
Controller in Enclosure 85, Slot	: В
Status:	Online
Current configuration Firmware version: Appware version: Bootware version: NVSRAM version: Pending configuration Firmware version: Appware version: Bootware version: NVSRAM version: Transferred on: Current ID (ALPA) On drive channel 1: Replacement part number:	07.50.00.00 07.50.00.00 07.50.00.00 N1818D53R1050V03 None None None None None None None
Model name: Board ID: Submodel ID:	7000-88 7091 93
Revision: Replacement part number: Part number: Serial number: Vendor: Date of manufacture:	0730 32847-10 SF74600024 IBM January 23, 2008
Data Cache Total present: Total used:	8192 MB 8192 MB
Data Cache Module Status: Location: Capacity: Replacement Part number: Manufacturer Part Number: Part Number: Serial number: Manufacturer: Date of manufacture: Data Cache Module	Optimal Controller B, Slot 9 1,024 MB Not Available M3 93T2953EZ3-CD5 Not Available 1127652678 Samsung December 31, 2007
Status:	Optimal
Location: Controller B, Slot 10 Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953CZ3-CD5 Part Number: Not Available Serial number: 4076997590 Manufacturer: Samsung Date of manufacture: January 29, 2007 Data Cache Module Status: Optimal Location: Controller B, Slot 11 Capacity: 1,024 MB Not Available Replacement Part number: Manufacturer Part Number: M3 93T2953CZ3-CD5 Part Number: Not Available Serial number: 4076997577 Manufacturer: Samsung Date of manufacture: January 29, 2007 Data Cache Module Status: Optimal Controller B, Slot 12 Location: Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953CZ3-CD5 Part Number: Not Available Serial number: 4076997790 Manufacturer: Samsung Date of manufacture: January 29, 2007 Data Cache Module Status: Optimal Controller B, Slot 13 Location: Capacity: 1.024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953CZ3-CD5 Part Number: Not Available 4076997641 Serial number: Manufacturer: Samsung Date of manufacture: January 29, 2007 Data Cache Module Status: Optimal location: Controller B, Slot 14 1,024 MB Capacity: Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953CZ3-CD5 Part Number: Not Available Serial number: 4110514102 Manufacturer: Samsung Date of manufacture: January 29, 2007 Data Cache Module Status: Optimal Controller B, Slot 15 Location: Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953CZ3-CD5 Part Number: Not Available Serial number: 4076997606 Manufacturer: Samsung Date of manufacture: January 29, 2007 Data Cache Module Status: Optimal Location: Controller B, Slot 16 Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953CZ3-CD5 Part Number: Not Available Serial number: 4076997562 Manufacturer: Samsung

Date of manufacture: January 29, 2007 Processor cache: Total present: 2044 MB Processor Cache Module Status: Optimal Location: Controller B, Slot 2 2,048 MB Capacity: Replacement Part number: Not Available Manufacturer Part Number: Not Available Not Available Part Number: Serial number: Not Available Manufacturer: Not Available Date of manufacture: Not available Cache Backup Device Status: Optimal Type: USB flash drive Location: Controller B, Connector USB 3 Capacity: 3,919 MB Product ID: TS4GUFM-V Part number: Not Available Serial number: 00000000000003 Revision level: 0.00 Manufacturer: Not Available Date of manufacture: Not available Status: Optimal USB flash drive Type: Location: Controller B, Connector USB 4 Capacity: 3.919 MB Product ID: TS4GUFM-V Part number: Not Available Serial number: 000000000000049 Revision level: 0.00 Not Available Manufacturer: Date of manufacture: Not available Host Interface Board Status: Optimal Slot 3 Location: Type: Fibre channel Number of ports: 4 0801 Board ID: Replacement part number: Not Available Part number: PN L2-25023-02 Serial number: SN SF84300021 Vendor: VN LSI Date of manufacture: November 1, 2008 Host Interface Board Status: Optimal Location: Slot 4 Type: Fibre channel Number of ports: 4 0801 Board ID: Replacement part number: Not Available Part number: PN L2-25023-02 Serial number: SN SF84300014 Vendor: VN LSI Date of manufacture: November 1, 2008 Tue Jan 27 15:19:47 EST 2009 Date/Time: Associated Logical Drives (\* = Preferred Owner): FC\_R3\_B\*, FDE\_R1\*, FDE\_R6\_B\*, Sata\_Raid6\_B\*, Secure\_Middle\_R6\_B\* XBB2\_MiddleB Controller DNS/Network name: Remote login: Disabled 1 Ethernet port:

Link status: MAC address: Negotiation mode: Port speed: Duplex mode: IPv4 settings: Network configura IP address: Subnet mask: Gateway: IPv6 settings: Auto-configuratio Local IP address:	tion: n:	Up 00:a0:b8:29:ec:e6 Auto-negotiate 10 Mbps Half duplex Enabled Static 192.168.70.71 255.255.255.0 0.0.0.0 Enabled Stateless FE80:0000:0000:02A0:B8FF:FE29:ECE6
Ethernet port: Link status: MAC address: Negotiation mode: Port speed: Duplex mode: IPv4 settings: Network configura IP address: Subnet mask: Gateway: IPv6 settings: Auto-configuratio Local IP address:	tion: n:	2 Failed 00:a0:b8:29:ec:e7 Auto-negotiate Unknown Half duplex Enabled Static 192.168.129.102 255.255.255.0 0.0.0.0 Enabled Stateless FE80:0000:0000:02A0:B8FF:FE29:ECE7
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 1 8,7 124/0x 4 Gbps 2 Gbps Auto Up	2
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 2 6, 5 124/0x 4 Gbps 4 Gbps Auto Up	2
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 3 4, 3, 124/0x 4 Gbps 4 Gbps Auto Up	Out 2
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 4 2, 1, 124/0x 4 Gbps 4 Gbps Auto Up	Out 2

Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 5 1, 2 124/0x2 4 Gbps 2 Gbps Auto Up		
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 6 3,4 124/0x2 4 Gbps 4 Gbps Auto Up		
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 7 5, 6, 00 124/0x2 4 Gbps 4 Gbps Auto Up	ut	
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 8 7, 8, 00 124/0x2 4 Gbps 4 Gbps Auto Up	ut	
Host interface: Channel: Current ID: Preferred ID: NL-Port ID: Maximum data rate: Current data rate: Data rate control: Link status: Topology: World-wide port ident World-wide node ident Part type:	tifier: tifier:	Fibre 1 Not applicable/0xFFFFFF 8/0xD9 0xFFFFFF 8 Gbps 8 Gbps Auto Down Not Available 20:1f:00:a0:b8:29:ed:8a 20:0e:00:a0:b8:29:ed:8a PM8032 revision 5	
Host interface: Channel: Current ID: Preferred ID: NL-Port ID: Maximum data rate: Current data rate: Data rate control: Link status: Topology: World-wide port ident World-wide node ident Part type:	tifier: tifier:	Fibre 2 Not applicable/0xFFFFFF 9/0xD6 0xFFFFF 8 Gbps 8 Gbps Auto Down Not Available 20:2f:00:a0:b8:29:ed:8a 20:0e:00:a0:b8:29:ed:8a PM8032 revision 5	
Host interface: Channel: Current ID:		Fibre 3 Not applicable/0xFFFFFFFF	

Preferred ID: 10/0xD5 NL-Port ID: 0xFFFFFF Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available World-wide port identifier: 20:3f:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a PM8032 Part type: revision 5 Host interface: Fibre Channel: 4 Current ID: Not applicable/0xFFFFFFF Preferred ID: 11/0xD4 NL-Port ID: 0x010700 Maximum data rate: 8 Gbps Current data rate: 4 Gbps Data rate control: Auto Link status: Up Topology: Fabric Attach World-wide port identifier: 20:4f:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a PM8032 Part type: revision 5 Host interface: Fibre Channel: 5 Current ID: Not applicable/0xFFFFFFF Preferred ID: 12/0xD3 NL-Port ID: 0xFFFFFF Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available World-wide port identifier: 20:5f:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Host interface: Fibre Channel: 6 Current ID: Not applicable/0xFFFFFFF Preferred ID: 13/0xD2 NL-Port ID: 0x010300 Maximum data rate: 8 Gbps Current data rate: 4 Gbps Data rate control: Auto Link status: Up Fabric Attach Topology: World-wide port identifier: 20:6f:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Host interface: Fibre Channel: 7 Current ID: Not applicable/0xFFFFFFF Preferred ID: 14/0xD1 NL-Port ID: 0xFFFFFF 8 Gbps Maximum data rate: Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available World-wide port identifier: 20:7f:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5

Host interface:	Fibre
Channel:	8
Current ID:	Not applicable/0xFFFFFFFF
Preferred ID:	15/0×CE
NL-Port ID:	0xFFFFFF
Maximum data rate:	8 Gbps
Current data rate:	8 Gbps
Data rate control:	Auto
Link status:	Down
Topology:	Not Available
World-wide port identifier:	20:8f:00:a0:b8:29:ed:8a
World-wide node identifier:	20:0e:00:a0:b8:29:ed:8a
Part type:	PM8032 revision 5

The summary information is also returned as the first section of information when you use the **profile** parameter.

The **show** commands return information about the specific components of a storage subsystem. The information returned by each of the **show** commands is the same as the information returned by the **show storageSubsystem profile** command, but the information is constrained to the specific component. For information about a command, refer to the page number shown.

- "Show Controller" on page 300
- "Show Disk Drive" on page 303
- "Show Drive Channel Statistics" on page 305
- "Show Storage Subsystem Host Topology" on page 325
- "Show Storage Subsystem LUN Mappings" on page 326
- "Show Logical Drive" on page 312
- "Show subsystem" on page 293
- "Show Logical Drive Reservations" on page 316

In addition, these commands return information about a storage subsystem:

- "Show Controller" on page 300
- "Show Enhanced Remote Mirroring Logical Drive Candidates" on page 310
- "Show Storage Subsystem Auto Configure" on page 320
- "Show Storage Subsystem Unreadable Sectors" on page 326
- "Show VolumeCopy Source Candidates" on page 330
- "Show VolumeCopy Target Candidates" on page 331
- "Show Logical Drive Performance Statistics" on page 316

There are other commands that can help you learn about your storage subsystem. To see a list of the commands, see Chapter 3, "Script Commands," on page 27. These commands are organized by the storage subsystem activities that the commands support, such as logical drive commands, host commands, enclosure commands, and others.

#### Clearing the Configuration

If you want to create a completely new configuration on a storage subsystem that already has an existing configuration, use the **clear storageSubsystem configuration** command. This command deletes all of the existing configuration information, including all of the subsystems, logical drives, and hot spare definitions from the controller memory.

Attention: Possible damage to the storage subsystem configuration – As soon as you run this command, the existing storage subsystem configuration is deleted.

The command takes this form:

c:\...\client>smcli 123.45.67.89 -c "clear storageSubsystem configuration;"

With controller firmware 7.xx.xx and higher, this command supports two parameters that you can use to limit the amount of configuration information removed:

- **all** Removes the entire configuration of the storage subsystem, including security information and identification information. Removing all of the configuration information returns the storage subsystem to its initial state. (This is what the command will assume if no parameters are specified.)
- **subsystems** Removes the logical drive configuration and the subsystem configuration, but leaves the rest of the configuration intact.

**Note:** With controller firmware earlier than 7.xx.xx.xx,, there are no parameters. The default is the action described for the **all** parameter.

If you want to create new subsystems and logical drives within the storage subsystem, you can use the **clear storageSubsystem configuration** command with the **subsystems** parameter to remove existing subsystems in a pre-existing configuration. This action destroys the pre-existing configuration. Use the **clear storageSubsystem configuration** command only when you create a new configuration.

#### Using the Auto Configure Command

The **autoConfigure storageSubsystem** command creates the subsystems on a storage subsystem, the logical drives in the subsystems, and the hot spares for the storage subsystem. When you use the **autoConfigure storageSubsystem** command, you define these parameters:

- The type of disk drives (Fibre, SATA, SAS)
- The RAID level
- The number of disk drives in an subsystem
- The number of subsystems
- The number of logical drives in each subsystem
- The number of hot spares
- The size of each segment on the disk drives
- A read ahead multiplier

After you define these parameters, the Storage Manager software creates the subsystems, the logical drives, and the hot spares. The controllers assign subsystem numbers and logical drive numbers as they are created. After the Storage Manager software creates the initial configuration, you can use the **set logicalDrive** command to define logical drive labels.

Before you run the **autoConfigure storageSubsystem** command, run the **show storageSubsystem autoConfigure** command. The latter command returns a list of parameter values that the Storage Manager software uses to automatically create a storage subsystem. If you would like to change any of the parameter values, you can do so by entering new values for the parameters when you run the **autoConfigure storageSubsystem** command. If you are satisfied with the parameter values that the **show storageSubsystem autoConfiguration** command returns, run the **autoConfigure storageSubsystem** command without new parameter values.

The **autoConfigure storageSubsystem** command takes this form:

```
autoConfigure storageSubsystem
[driveType=(fibre | SATA | SAS)
raidLevel=(0 | 1 | 3 | 5 | 6)
subsystemWidth=numberOfDrives
subsystemCount=numberOfLogicalDriveGroups |
logicalDrivesPersubsystemCount=numberOfLogicalDrivesPersubsystem |
hotSpareCount=numberOfHotSpares |
segmentSize=segmentSizeValue |
cacheReadPrefetch=(TRUE | FALSE)
securityType=(none | capable | enabled)
T10PI=(none|enabled)]
```

When you use the **autoConfigure storageSubsystem** command, two symbol functions (**getAutoConfigCandidates** and **createAutoConfig**) are used that let the client retrieve default settings for the various automatic configuration parameters, change the settings, query what the results of those changes would be and, finally, apply the desired parameters to create a configuration. The configurability portion of this feature provides enhancements to the automatic subsystem creation algorithms, which produce subsystems with improved performance and more information about disk drive and logical drive attributes so the user can make better choices when configuring logical drives manually.

The **subsystemWidth** parameter defines the number of unassigned disk drives that you want to use for each new subsystem.

The **subsystemCount** parameter defines the number of new subsystems that you want in the storage subsystem.

The **LogicalDrivesPersubsystemCount** parameter defines the number of logical drives that you want in each subsystem.

The **hotSpareCount** parameter defines the number of hot spares that you want in each subsystem.

The **segmentSize** parameter defines the amount of data, in KB, that the controller writes on a single disk drive in a logical drive before writing data on the next disk drive. The smallest units of storage are data blocks. A data block stores 512 bytes of data. The size of a segment determines how many data blocks that it contains. An 8-KB segment holds 16 data blocks. A 64-KB segment holds 128 data blocks.

**Important:** For optimal performance in a multiuser database or file system storage environment, set the segment size to minimize the number of disk drives that are needed to satisfy an I/O request.

Using a single disk drive for a single request leaves other disk drives available to simultaneously service other requests. Valid segment size values are 8, 16, 32, 64, 128, 256, and 512.

**Note:** If you set the cache block size to 16, you cannot create a logical drive with a segment size of 8.

If the logical drive is for a single user with large I/O requests (such as multimedia), performance is maximized when a single I/O request can be serviced

with a single data stripe. A data stripe is the segment size multiplied by the number of disk drives in the subsystem that are used for data storage. In this environment, multiple disk drives are used for the same request, but each disk drive is accessed only once.

The **cacheReadPrefetch** parameter turns on or turns off the ability of the controller to read additional data blocks into the cache. When you turn on cache read prefetch, the controller copies additional data blocks into the cache while it is reading requested data blocks from a disk drive into the cache. This action increases the chance that a future request for data can be fulfilled from the cache, which improves the speed with which data is accessed. The number of additional data blocks that the controller reads into the cache is determined by the configuration settings for the storage subsystem that you use. Cache read prefetch is important for applications that use sequential I/O, such as multimedia applications.

Valid values for the **cacheReadPrefetch** parameter are TRUE or FALSE. If you want to turn on cache read prefetch, set the **cacheReadPrefetch** parameter to TRUE. If you want to turn off cache read prefetch, set the **cacheReadPrefetch** parameter to FALSE.

The following table lists the default values for the segment size and cache read prefetch settings for different storage subsystem uses.

Storage Subsystem Use	Segment Size (KB)	Cache Read Prefetch
File system	128	TRUE
Database	128	TRUE
Multimedia	256	TRUE

Table 16. Default Values for Segment Size and Cache Read Prefetch

After you have finished creating the subsystems and the logical drives by using the **autoConfigure storageSubsystem** command, you can further define the properties of the logical drives in a configuration by using the **set logicalDrive** command. (For a description about how to further define your configuration, see "Modifying Your Configuration" on page 435.)

#### Security Type

The **securityType** parameter is used to specify the security subsystems with Full Disk Encryption (FDE) drives. A Disk Encryption Activation premium feature must be purchased and enabled before you can take advantage of the FDE drive security feature. The minimum version of controller firmware that supports FDE Full Disk Encryption features is 07.5x.xx.xx. If this parameter is set to **none**, the created subsystems and logical drives will not be encrypted. If the parameter is set to **enabled**, the subsystems and logical drives will have security enabled and encrypted.

With FDE, the controller firmware can create a key and activate the Full Disk Encryption feature. The Full Disk Encryption feature encrypts data as the data is written to the drive and decrypts the data as the data is read from the drive. Without the key created by the controller, the data written to the drive is inaccessible.

Before you can set the **securityType** parameter to **capable** or **enabled**, you must create a storage subsystem security key. Use the **create storageSubsystem** 

**securityKey** command to create a storage subsystem security key. These commands are related to the security key:

- create storageSubsystem securityKey
- set storageSubsystem securityKey
- import storageSubsystem securityKey
- export storageSubsystem securityKey
- validate storageSubsystem securityKey

And these commands are about Full Disk Encryption (FDE):

- start secureErase (drive | drives)
- enable subsystem [subsystemName] security
- create hostPort identifier

Note: Firmware version 07.50 adds the securityType parameter.

#### Example of the Auto Configuration Command

c:\...\client>smcli 123.45.67.89 -c "autoConfigure storageSubsystem driveType=fibre raidLevel=5 subsystemWidth=8 subsystemCount=3 logicalDrivesPersubsystemCount=4 hotSpareCount=2 segmentSize=8 cacheReadPrefetch=TRUE securityType=none T10PI=none;"

The command in this example creates a storage subsystem configuration by using Fibre Channel disk drives set to RAID level 5. Three subsystems are created, and each subsystem consists of eight disk drives, which are configured into four logical drives. The storage subsystem has two hot spares. The segment size for each logical drive is 8 KB. The cache read prefetch is turned on, which causes additional data blocks to be written into the cache.

#### T10PI protection

T10PI protection checks and corrects errors that might occur as data is communicated between a host and a storage subsystem. Thus, enabling this feature for a new logical drive ensures that errors are detected and corrected. To enable T10PI, storage subsystem must support T10PI and the disk pool must be T10PI capable.

#### Using the Create LogicalDrive Command

Use the **create logicalDrive** command to create new storage subsystem logical drives in three ways:

- Create a new logical drive while simultaneously creating a new subsystem to which you assign the disk drives
- Create a new logical drive while simultaneously creating a new subsystem to which the Storage Manager software assigns the disk drives
- Create a new logical drive in an existing subsystem

You must have unassigned disk drives in the subsystem. You do not need to assign the entire capacity of the subsystem to a logical drive.

#### Creating Logical Drives with User-Assigned Disk Drives

When you create a new logical drive and assign the disk drives you want to use, the Storage Manager software creates a new subsystem. The controller firmware assigns an subsystem number to the new subsystem. The command takes this form:

```
create logicalDrive drives=(enclosureID1,slotID1...enclosureIDn,slotIDn)
raidLevel=(0 | 1 | 3 | 5 | 6) userLabel=logicalDrivesName
[capacity=logicalDrivesCapacity] [owner=(a | b)
cacheReadPrefetch=(TRUE | FALSE)
segmentSize=segmentSizeValue]
enclosureLossProtect=(TRUE | FALSE)
securityType=(none | capable | enabled) T10PI=(none|enabled)]
```

**Note:** The **capacity** parameter, the **owner** parameter, the **cacheReadPrefetch** parameter, the **segmentSize** parameter, and the **enclosureLossProtect** parameter are optional parameters (indicated by the items inside the square brackets). You can use one or all of the optional parameters as needed to define your configuration. If you choose not to use any of the optional parameters, the default values of the parameters are used for your configuration.

The **userLabel** parameter is the name that you want to give to the logical drive. The logical drive name can be any combination of alphanumeric characters, hyphens, and underscores. The maximum length of the logical drive name is 30 characters. You must enclose the logical drive name with double quotation marks ("").

The **drives** parameter is a list of the disk drives that you want to use for the subsystem. Enter the enclosure ID and the slot ID of each disk drive that you want to use. Enclose the list in parentheses, separate the enclosure ID value and the slot ID value of a disk drive with a comma, and separate each enclosure ID and slot ID pair with a space. This example shows you how to enter enclosure ID values and slot ID values:

(1,1 1,2 1,3 1,4 1,5)

The **capacity** parameter defines the size of the logical drive. You do not need to assign the entire capacity of the disk drives to the logical drive. Later, you can assign any unused space to another logical drive.

The **owner** parameter defines the controller to which you want to assign the logical drive. If you do not specify a controller, the controller firmware determines the logical drive owner.

The **cacheReadPrefetch** parameter and the **segmentSize** parameter are the same as those described for the **autoConfigure storageSubsystem** command.

The **enclosureLossProtect** parameter defines enclosure loss protection for the subsystem. (For a description of how enclosure loss protection works, see "Enclosure Loss Protection" on page 427.)

The **securityType** parameter is used to specify the security level when creating subsystems and logical drives in storage subsystems with Full Disk Encryption (FDE) drives.

#### Example of Creating Logical Drives with User-Assigned Disk Drives

```
c:\...\client>smcli 123.45.67.89 -c "create logicalDrive drives
=(1,1 1,2 1,3 2,1 2,2 2,3) raidLevel=5 userLabel=\"Engineering_1\"
capacity=20GB owner=a cacheReadPrefetch=TRUE segmentSize=128
securityType=enabled T10PI=none;"
```

The command in this example automatically creates a new subsystem and a logical drive with the name Engineering\_1. The subsystem has a RAID level of 5 (RAID 5). The command uses six disk drives to construct the subsystem. The disk drives

have a total logical drive capacity of 20 GB. If each disk drive has a capacity of 18 GB, the total capacity of all the assigned disks is 108 GB. 18 GB x 6 disk drives = 108 GB

Because only 20 GB is assigned to the logical drive, 88 GB remains available (as unconfigured capacity) for other logical drives that a user can add to this subsystem later.

108 GB - 20 GB subsystem size = 88 GB

Cache read prefetch is turned on, which causes additional data blocks to be written into the cache. The segment size for each logical drive is 128 KB. Enclosure loss protection is set to TRUE, which prevents any operations to disk drives in the expansion drawer if the expansion drawer fails. Hot spares are not created for this new subsystem. You must create hot spares after you run this command.

#### Creating Logical Drives with Software-Assigned Disk Drives

If you choose to let the Storage Manager software assign the disk drives when you create the logical drive, you need only to specify the number of disk drives that you want to use. The Storage Manager software then assigns the disk drives. The controller firmware assigns an subsystem number to the new subsystem. To manually create subsystems and logical drives, use the **create logicalDrive** command:

```
create logicalDrive driveCount=numberOfDrives
raidLevel=(0 | 1 | 3 | 5 | 6) userLabel=logicalDriveName
[driveType=(fibre | SATA | SAS)]
[capacity=logicalDriveCapacity | owner=(a | b) |
cacheReadPrefetch=(TRUE | FALSE) |
segmentSize=segmentSizeValue]
[enclosureLossProtect=(TRUE | FALSE)
securityType=(none | capable | enabled)]
```

This command is similar to the previous **create logicalDrive** command in which users assign the disk drives. The difference between this command and the previous one is that this version of the command requires only the number and the type of disk drives you want to use in the subsystem. You do not need to enter a list of disk drives. All of the other parameters are the same. Enclosure loss protection is performed differently when the Storage Manager software assigns the disk drives than when a user assigns the disk drives. (For a description of the difference, see "Enclosure Loss Protection" on page 427.)

# Example of Creating Logical Drives with Software Assigned Disk Drives

c:\...\client>smcli 123.45.67.89 -c "create logicalDrive driveCount=6
raidLevel=5 userLabel=\"Engineering\_1\"
capacity=20GB owner=a cacheReadPrefetch=TRUE segmentSize=128 securityType=capable;"

The command in this example creates the same logical drive as the example for the previous **create logicalDrive** command in which a user assigns the disk drives. The difference is that a user does not know which disk drives are assigned to this subsystem.

#### Creating Logical Drives in an Existing subsystem

If you want to add a new logical drive to an existing subsystem, use this command:

create logicalDrive subsystem=subsystemNumber userLabel=logicalDriveName [freeCapacityArea=freeCapacityIndexNumber |

```
capacity=logicalDriveCapacity | owner=(a | b) |
cacheReadPrefetch=(TRUE | FALSE) |
segmentSize=segmentSizeValue
securityType=(none | capable | enabled)T10PI=(none|enabled)]
```

**Note:** Parameters wrapped in square brackets or curly brackets are optional. You can use one or all of the optional parameters as needed to define your configuration. If you choose not to use any of the optional parameters, the default values of the parameter are provided for your configuration.

The **subsystem** parameter is the number of the subsystem in which you want to create a new logical drive. If you do not know the subsystem numbers on the storage subsystem, you can use the **show allLogicalDrives summary** command to get a list of the logical drives and the subsystems to which the logical drives belong.

The **userLabel** parameter is the name that you want to give to the logical drive. The logical drive name can be any combination of alphanumeric characters, hyphens, and underscores. The maximum length of the logical drive name is 30 characters. You must enclose the logical drive name with double quotation marks ("").

The **freeCapacityArea** parameter defines the free capacity area to use for the logical drive. If a subsystem has several free capacity areas, you can use this parameter to identify which free capacity area to use for logical drive creation. You do not have to assign the entire capacity of the disk drives to the logical drive. Later, you can assign any unused space to another logical drive.

The usage of the **capacity** parameter, the **owner** parameter, the **cacheReadPrefetch** parameter, and the **segmentSize** parameter is the same as described in the previous examples of the **create logicalDrive** command.

#### **Enclosure Loss Protection**

The **enclosureLossProtect** parameter is a boolean switch that you set to turn on or turn off enclosure loss protection. For enclosure loss protection to work, each disk drive in an subsystem must be on a separate enclosure. The way in which enclosure loss protection works depends on the method that you choose to assign the disk drives for an subsystem.

When you assign the disk drives, if you set **enclosureLossProtect**=TRUE and have selected more than one disk drive from any one enclosure, the storage subsystem returns an error. If you set **enclosureLossProtect**=FALSE, the storage subsystem performs operations, but the subsystem that you create does not have enclosure loss protection.

When the controller firmware assigns the disk drives, if **enclosureLossProtect**=TRUE the storage subsystem posts an error if the controller firmware cannot provide disk drives that result in the new subsystem having enclosure loss protection. If **enclosureLossProtect**=FALSE, the storage subsystem performs the operation even if it means that the subsystem might not have enclosure loss protection.

Enclosure loss protection is not valid when creating logical drives on existing subsystems.

#### Configuring a storage subsystem with disk pools

A disk pool is a collection of 11 or more disk drives in a Storage Subsystem that have the same spindle speed, the same security level, and preferably the same capacity to make the most efficient use of the disk drives.

#### Using the create disk pool command

Use the create diskPool command to create a new disk pool in two ways:

- Create a new disk pool automatically by entering the type of drives that you want in the disk pool, and the name that you want to assign to it. Then let the storage management software assign the drives.
- Create a new disk pool manually by selecting and assigning drives to the disk pool.

You do not need to assign the entire capacity of the storage subsystem to the disk pool. You can mix both disk pools and storage subsystems in a storage subsystem.

When you create a new disk pool you want to assign the largest number of drives possible to the disk pool. The more drives that you have in a disk pool, the more robust the disk pool is, with faster rebuild times and simplified management requirements. The minimum number of drives that you can have in a disk pool is 11; the maximum number of drives is as many as required to support the maximum number of logical drives that a controller can manage. The maximum number of logical drives that a DS3500 and DCS3700 storage subsystem supports is 512. The maximum number of logical drives that a DCS3700 storage subsystem supports is 2048. You must have at least 11 unassigned disk drives in the storage subsystem to configure a disk pool.

The only supported RAID level for a disk pool is RAID 6, with a stripe width of 10. The storage subsystem must have a minimum drive count of 11 disk drives to create a pool. The number of drives in the pool also influences how much reserved capacity is needed in order to redistribute data for rebuilds. The following table shows the drive counts and the number of disk drives reserved for reconstruction:

Drive Count	Drives Reserved for Reconstruction	Comments
11	1	The minimum number of drives in a disk pool is 11.
12 -31	2	
32 - 63	3	
64 - 123	4	
128 - 191	6	
192- 255	7	The maximum configuration of DS3500 and DCS3700 is 192 drives
256 - 360	8	The maximum configuration of DCS3860 and the DCS3700 storage system with Performance Module Controllers is 360 drives

Table 17. Drive Counts and the Number of Disk Drives Reserved for Reconstruction

Hot spares are not required or needed for disk pools. Spare capacity for reconstruction is divided among the drives within a disk pool. A small amount of each drive is reserved as reconstruction space to hold reconstructed data in the event of loss of access to a drive or a drive failure. Because of this behavior, the system can sustain drive failures until the capacity is exhausted or the number of drives in the disk pool falls below the minimum drive count. As long as free capacity exists on the system, failed drives are rebuilt and degraded logical drives are brought back to optimal. The only constraint to rebuilding failed drives is the RAID level of the logical drives in the disk pool.

In your storage subsystem with controller firmware version 7.83.xx.xx and later, you can create upto 10 disk pools. Reasons for having several disk pools in a storage subsystem might include separate disk pools to use certain types of drives, or to create different pools for different applications. For the most efficient disk pool, all of the drives need to have the same characteristics:

- Drive type (SAS)
- Drive media (HDD)
- · Spindle speed
- Full-Disk Encryption (FDE)
- T10PI
- Capacity

Before you create a disk pool, run the show storageSubsystem command to determine the drives that are available, and to make sure that you have enough drives in the storage subsystem to create a disk pool.

The create diskPool command has several optional parameters to enable you create the configuration that you want. The use of these parameters is shown in the examples in the following sections. For more information about the use of the optional parameters refer to Chapter 3 of this guide.

In a disk pool, you must create the disk pool first and then the logical drives.

#### Creating disk pools with software-assigned disk drive

If you choose to let the storage management software assign the disk drive when you create a disk pool, you need only to specify the type of disk drive that you want to use and a name (**user label**) for the disk pool. Optionally, you can also specify the number of disk drives that you want to use. The storage management software then assigns the disk drive. The controller firmware assigns a number to the new disk pool. To create a disk pool with a software-assigned disk drive, use this command:

```
create diskPool diskDriveType=(fibre|sas)
userLabel="diskPoolName"
[diskDriveCount=diskDriveCountValue |
warningThreshold=(warningThresholdValue|default) |
criticalThreshold=(criticalThresholdValue|default) |
criticalPriority=(highest|high|medium|low|lowest) |
backgroundPriority=(highest|high|medium|low|lowest) |
degradedPriority=(highest|high|medium|low|lowest) |
securityType=(none|capable|enabled) |
diskDriveMediaType=(hdd | ssd | allMedia | unknown) |
T10PI=(none|enabled)]
```

Example of Creating logical drives with Software-Assigned disk drive

```
c:\...\smX\client>smcli 123.45.67.88 123.45.67.89
-c "create diskpool diskDriveType=sas
userLabel="Engineering_1"
diskDriveCount=64 warningthreshold=65 criticalthreshold=75
criticalpriority=high backgroundpriority=medium
degradedpriority=high securitytype=enabled
diskDrivemediatype=hdd T10PI=enabled;"
```

This command creates a disk pool with these features:

- Type of disk drive is SAS
- The name of the disk pool is **Engineering\_1**. The disk pool name can be any combination of alphanumeric characters, hyphens, and underscores. The maximum length of the disk pool name is 30 characters. You must enclose the disk pool name with double quotation marks ("").
- The storage management software will add 64 disk drives to the disk pool. This assumes that the storage subsystem has a minimum of 64 disk drives that have the same characteristics.
- When the disk pool consumes **65 percent** of its capacity, a warning alert will be posted. The default value is **50 percent**. The **warningthreshold** parameter must always be set to a value lower than the **criticalthreshold** parameter.
- When the disk pool consumes **75 percent** of its capacity, a **critical alert** will be posted. The default value is **85 percent**.
- The priority for reconstruction operations for critical events on the disk pool is set to high. If a condition, such as a two failed disk drive occurs, the storage management software will make the reconstruction of the data a high priority.
- The priority for background operations on this disk pool is set to medium, meaning that background operations (such as reconstruction or formatting) equally share resources with other storage subsystem operations.
- The priority for correcting the disk pool after it has entered a Degraded state is set to high. If a condition, such as a failed disk drive occurs, the storage management software will make the correction of the condition a high priority.
- The securitytype is **enabled**, so the storage management software will use only disk drives that are configured to be Full-Disk Encryption (FDE).
- The type of disk drive to be used is a hard disk drive (hdd).
- The disk pool will use only disk drives with protected data capability.

# Creating disk pools with user-assigned disk drives

In some situations you might be required to create a disk pool by assigning the disk drives to the disk pool, instead of having the software assign the disk Drives. One situation might be when you want to create a small disk pool to test possible configuration parameters. Another situation might be when you do not have enough disk drives with the same characteristics to create a disk pool. If all of the usable disk drives in the storage subsystem do not have the same capacity, you can still configure a disk pool by manually selecting the disk drives. One constraint, however, is that each drive in the disk pool assumes the same capacity as the smallest drive, even if the other disk drives have more capacity.

The command to manually assign disk drives has this form:

```
create diskPool diskDrives=(EnclosureID1,drawerID1,slotID1 ...
EnclosureIDN,drawerIDN,slotIDN
userLabel="diskPoolName"
[driveCount=driveCountValue |
warningThreshold=(warningThresholdValue|default) |
criticalThreshold=(criticalThresholdValue|default) |
criticalPriority=(highest|high|medium|low|lowest) |
```

```
backgroundPriority=(highest|high|medium|low|lowest) |
degradedPriority=(highest|high|medium|low|lowest) |
securityType=(none|capable|enabled) |
driveMediaType=(hdd | ssd | allMedia | unknown) |
T10PI=(none|enabled)]
```

The **diskDrives** parameter is a list of the disk drives that you want to use for the disk pool. Enter the Enclosure ID and the slot ID of each drive that you want to use. For high capacity drive Enclosures that have drawers to hold the disk drives also use the drawer number. For high capacity drive enclosures, the sequence of the location identifiers is drive Enclosure, drawer, slot. Enclose the list in parentheses, separate the Enclosure ID value, drawer ID value, and the slot ID value of a drive with a comma, and separate each Enclosure ID, drawer ID, and slot ID set with a space. This example shows you how to enter enclosure ID values and slot ID values for low capacity drive enclosures:

```
(1,1 1,2 1,3 1,4 1,5)
```

This example shows you how to enter enclosure ID values, drawer ID values, and slot ID values for high capacity drive enclosures:

```
(1,1,1 1,2,3 1,3,5 1,4,6 1,5,8)
```

An example of Creating logical drives with User-Assigned disk drives:

```
c:\...\smX\client>smcli 123.45.67.88 123.45.67.89
-c "create diskpool diskDrives=(1,1,1 1,1,2 1,2,3 ...
2,1,10, 2,2,11)
userLabel="Engineering_1" warningthreshold=65
criticalthreshold=75 criticalpriority=high
backgroundpriority=medium degradedpriority=high
securitytype=enabled drivemediatype=hdd
T10PI=enabled;"
```

This command creates a disk pool with these features:

- The list of disk drives represents the disk drives found in a high capacity drive Enclosure.
- The name of the disk pool is **Engineering\_1**. The disk pool name can be any combination of alphanumeric characters, hyphens, and underscores. The maximum length of the disk pool name is 30 characters. You must enclose the disk pool name with double quotation marks ("").
- When you assign disk drives to a disk pool, you do not need to use the **driveCount** parameter.
- When the disk pool consumes **65 percent** of its capacity, a warning alert will be posted. The default value is **50 percent**. The **warningthreshold** parameter must always be set to a value lower than the **criticalthreshold** parameter.
- When the disk pool consumes **75 percent** of its capacity, a critical alert will be posted. The default value is **85 percent**.
- The priority for reconstruction operations for critical events on the disk pool is set to **high**. If a condition, such as a two failed disk drives occurs, the storage management software will make the reconstruction of the data a high priority.
- The priority for background operations on this disk pool is set to **medium**, meaning that background operations (such as reconstruction or formatting) equally share resources with other storage subsystem operations.
- The priority for correcting the disk pool after it has entered a degraded state is set to **high**. If a condition, such as a failed drive occurs, the storage management software will make the correction of the condition a high priority.

- The **securitytype** parameter is enabled, so the storage management software will use only disk drives that are configured to be Full-Disk Encryption (FDE)s.
- The type of drive to be used is a hard drive (hdd).
- The disk pool will use only disk drives with protected data capability.

# Using the create logicaldrive diskPool command

**Note:** In a disk pool, you must create the disk pool first and then the logical drives.

The create logicalDrive diskPool command enables you to create a logical drive in a disk pool. You can create either one of these types of logical drives:

- Standard logical drive
- Thin logical drive

A standard logical drive has a fixed capacity that you can define when you create the logical drive. The standard logical drive reports only the fixed capacity to the host. In disk pools, the logical drive capacity is distributed across all of the applicable disk drives. You do not need to identify specific disk drives for the logical drive.

A thin logical drive is an expandable logical drive with both physical capacity and virtual capacity. Physical capacity is the size of the logical drive at a given time that is currently allocated for writing data. This size can increase over time. Virtual capacity is capacity reported to the hosts and is the "size" of the logical drive. Thin provisioning enables you to create logical drives with a large virtual capacity and relatively small physical capacity, which is beneficial for storage utilization and efficiency. Thin logical drives can help simplify storage administration because the physical capacity can increase as the application needs change, without disrupting the application, allowing for better storage utilization.

Keep these guidelines in mind when choosing a name for your logical drive:

- A logical drive name can consist of letters, numbers, and the special characters underscore (\_), hyphen (-), and pound (#). If you choose any other characters, an error message appears. You are prompted to choose another name.
- Limit the name to 30 characters. Any leading and trailing spaces in the name are deleted.
- Use a unique, meaningful name that is easy to understand and remember.
- Avoid arbitrary names or names that would quickly lose their meaning in the future.
- If you choose a logical drive name that duplicates that of another logical drive, an error message appears. You are prompted to choose another name.

**Note:** iSCSI host ports do not support T10PI. If there are one or more iSCSI host ports in your storage subsystem, a warning appears.

**T10PI protection** checks for, and corrects errors that might occur as data is communicated between a host and a storage subsystem. Enabling this feature for a new logical drive helps make sure that errors are detected and corrected. To enable T10PI, these conditions must be in place:

- The storage subsystem must be able to support T10PI
- The disk pool must be T10PI capable
- The T10PI premium feature must be enabled

Note: Your storage subsystem must support T10PI feature.

If any of the host connections on the controllers in your storage subsystem do not support T10PI, the associated hosts cannot access data on T10PI-enabled logical drives. Before creating a T10PI-enabled logical drive, make sure that the host connection that you are planning to use supports T10PI.

Dynamic cache read prefetch allows the controller to (optionally) copy additional sequential data blocks into the cache while it is reading data blocks from a drive to cache. This caching increases the chance that future requests for data can be filled from the cache. Cache read-ahead is important for multimedia applications that use sequential I/O. The rate and amount of data that is pre-fetched into cache is self-adjusting based on the rate and request size of the host reads. Random access does not cause data to be pre-fetched into cache. This feature has no effect when read caching is disabled.

Note:

- When you are creating a thin logical drive, the dynamic cache read prefetch option is not available.
- Regardless of the capacity specified, capacity in a disk pool is allocated in 4 GB increments. Any capacity that is not a multiple of 4 GB is allocated, however not usable. To make sure that the entire capacity is usable, specify the capacity in 4 GB increments. If unusable capacity exists, the only way to regain it is to increase the capacity of the logical drive.

# Creating standard logical drives on a disk pool

If you want to add a new standard logical drive to an existing disk pool, use this command:

```
create logicalDrive diskPool="diskPoolName"
userLabel="logicalDriveName"
capacity=logicalDriveCapacity
thinProvisioned=FALSE |
[owner=(a|b) |
mapping=(none|default) |
T10PI=(none|enabled) |
cacheReadPrefetch=(TRUE | FALSE)]
```

The **diskPool** parameter is the name of the disk pool in which you want to create a new logical drive. If you do not know the disk pool names on the storage subsystem, you can use the show storageSubsystem summary command to get a list of the disk pool.

The **userLabel** parameter is the name that you want to give to the logical drive. The logical drive name can be any combination of alphanumeric characters, hyphens, and underscores. The maximum length of the logical drive name is 30 characters. You must enclose the logical drive name with double quotation marks ("").

The **capacity** parameter defines the capacity to use for the logical drive.

The **thinProvisioned** parameter sets the logical drive to either **standard** or **thin**. For a standard logical drive, the **thinProvisioned** parameter must be set to FALSE.

The **owner** parameter defines which controller is designated to be the primary controller of the logical drive. For best performance ensure that the controllers are balanced as closely as possible for total I/Os.

The **mapping** parameter defines whether you want the storage management software to map the logical drive to a host, or if you want to map the logical drive to a host at a later time. To allow the storage management software to map the logical drive to a host use the default parameter. To map the logical drive to a host at a later time, use the **none** parameter. To map a logical drive to a host, use the set logicalDrive logicalUnitNumber command.

To turn on T10PI, set the T10PI parameter to enabled.

To turn on cache read prefetch, set the cacheReadPrefect parameter to TRUE.

The usage of the **owner** parameter, the **cacheReadPrefetch** parameter, and the **segmentSize** parameter is the same as described in the previous examples of the create logicalDrive command.

#### Creating thin logical drives on a disk pool

If you want to add a new thin logical drive to an existing disk pool, use this command:

```
create logicalDrive diskPool="diskPoolName"
userLabel="logicalDriveName" capacity=logicalDriveVirtualCapacity
thinProvisioned=TRUE
[owner=(a|b) |
mapping=(none|default) |
T10PI=(none|enabled) |
(existingRepositoryLabel=existingRepositoryName |
newRepositoryCapacity=newRepositoryCapacityValue [KB | MB | GB | TB | Bytes]) |
repositoryMaxCapacity=repositoryMaxCapacityValue[KB|MB|GB|TB|Bytes] |
repositoryExpansionPolicy=(automatic|manual) |
warningThresholdPercent=warningThresholdPercentValue |
cacheReadPrefetch=(TRUE | FALSE)]
```

The **diskPool** parameter is the name of the disk pool in which you want to create a new thin logical drive. If you do not know the disk pool names on the storage subsystem, you can use the show alllogicalDrives summary command to get a list of the logical drives and the disk pools to which the logical drives belong.

The **userLabel** parameter is the name that you want to give to the disk pool. The disk pool name can be any combination of alphanumeric characters, hyphens, and underscores. The maximum length of the disk pool name is 30 characters. You must enclose the disk pool name with double quotation marks ("").

The **capacity** parameter defines the virtual capacity of the thin logical drive. The capacity is the value that is reported to the host. As users add information to the thin logical drive, the physical size of the logical drive increases. When you define the capacity of the thin logical drive you must define a capacity of at least 32 GB. The maximum capacity that you can define is 63 TB.

The **mapping** parameter defines whether you want the storage management software to map the logical drive to a host, or if you want to map the logical drive to a host at a later time. To allow the storage management software to map the logical drive to a host use the **default** parameter. To map the logical drive to a host at a later time use the **none** parameter. To map a logical drive to a host, use the set logicalDrive logicalUnitNumber command.

The repository capacity is the actual physical capacity of the thin logical drive. The value that you use for the repository capacity is the starting size of the physical component of a thin logical drive. When you define the capacity of the repository

must define a capacity of at least 4 GB. The maximum capacity that you can define is 64 TB. You can use a small starting value for the repository. As data increases in the repository, additional standard logical drives are added to the repository to increase the capacity. You can either use an existing repository or create a new repository with this command. The repository capacity is governed by these parameters:

- **existingRepositoryLabe1** Use this parameter when you want to associate the logical drive with an existing repository logical drive.
- **newRepositoryCapacity** Use this parameter when you want to create a new repository logical drive.
- **repositoryMaxCapacity** Use this parameter to define the maximum size that you want for the repository logical drive.
- **repositoryExpansionPolicy** Use this parameter to define whether the repository logical drive expands automatically or if you need to expand the repository logical drive.

The **warningThresholdPercent** parameter defines when you receive a warning that the repository logical drive is nearing maximum capacity. The value for this parameter is percent of the maximum capacity of the repository logical drive.

To turn on cache read prefetch, set the **cacheReadPrefect** parameter to TRUE.

Note: Thin logical drives do not use T10PI.

# **Modifying Your Configuration**

For most configurations, after you have created your initial configuration by using the **autoConfigure storageSubsystem** command or the **create logicalDrive** command, you must modify the properties of your configuration to make sure that it performs to meet the requirements for data storage. Use the **set** commands to modify a storage subsystem configuration. This section describes how to modify these properties:

- The controller clocks
- The storage subsystem password
- The storage subsystem host type
- The storage subsystem cache
- The global hot spares

# Setting the Controller Clocks

To synchronize the clocks on the controllers with the host, use the **set storageSubsystem time** command. Run this command to make sure that event time stamps that are written by the controllers to the Event Log match the event time stamps that are written to the host log files. The controllers stay available during synchronization. This example shows the command:

c:\...\client>smcli 123.45.67.89 -c "set storageSubsystem time;"

# Setting the Storage Subsystem Password

Use the **set storageSubsystem** command to define a password for a storage subsystem. The command takes this form:

set storageSubsystem password="password"

The **password** parameter defines a password for the storage subsystem. Passwords provide added security to a storage subsystem to help reduce the possibility of implementing destructive commands.

**Attention: Possible data corruption or data loss** – Implementing destructive commands can cause serious damage, including data loss.

Unless you define a password for the storage subsystem, you can run all of the script commands. A password protects the storage subsystem from any command that the controllers consider destructive. A destructive command is any command that can change the state of the storage subsystem, such as logical drive creation; cache modification; or reset, delete, rename, or change commands.

If you have more than one storage subsystem in a storage configuration, each storage subsystem has a separate password. Passwords can have a maximum length of 30 alphanumeric characters. You must enclose the password in double quotation marks (""). This example shows how to use the **set storageSubsystem** command to define a password:

c:\...\client>smcli 123.45.67.89 -c "set storageSubsystem
password="1a2b3c4d5e";"

**Important: Password cannot be reset without current password**– If you forget the storage subsystem password, you cannot reset the password using Storage Manager. To reset the password in that situation, you must contact your IBM support representative.

#### Setting the Storage Subsystem Host Type

Use the **set storageSubsystem** command to define the default host type. The command takes this form:

set storageSubsystem defaultHostType=(hostTypeName | hostTypeIdentifier)

The **defaultHostType** parameter defines how the controllers in the storage subsystem will communicate with the operating system on undefined hosts that are connected to the storage subsystem SAN. This parameter defines the host type only for data I/O activities of the storage subsystem. This parameter does not define the host type for the management station. The operating system can be Windows, Linux, or Solaris.

For example, if you set the **defaultHostType** parameter to Linux, the controller communicates with any undefined host if the undefined host is running a Linux operating system. Typically, you would need to change the host type only when you are setting up the storage subsystem. The only time that you might need to use this parameter is if you need to change how the storage subsystem behaves relative to the hosts that are connected to it.

Before you can define the default host type, you need to determine what host types are connected to the storage subsystem. To return information about host types that are connected to the storage subsystem, use the **show storageSubsystem** command with the **defaultHostType** parameter or the **hostTypeTable** parameter. This command returns a list of the host types with which the controllers can communicate. This command does not return a list of the hosts. These examples show the use of the **show storageSubsystem** command:

c:\...\client>smcli 123.45.67.89 -c "show storageSubsystem defaultHostType;" c:\...\client>smcli 123.45.67.89 -c "show storageSubsystem hostTypeTable;" This example shows how to define a specific default host type after displaying the host type table:

c:\...\client>smcli 123.45.67.89 -c "set storageSubsystem defaultHostType=xx;"

where *xx* is the host type index value. For value *xx*, refer to the previously-displayed host type table. From the table, select the appropriate value for the host server that the LUN from the storage subsystem is mapped to.

#### Setting the Storage Subsystem Cache

The cache is high-speed memory that holds data that is either written to the disk drives or read by the host. A controller has two memory areas used for intermediate storage of read data and write data. The read cache contains data that has been read from the disk drives but not yet transferred to the host. The write cache contains data from the host but not yet written to the disk drives.

The cache acts as a buffer so that data transfers between the host and the disk drive do not need to be synchronized. In read caching, the data for a read operation from the host might already be in the cache from a previous operation, which eliminates the need to access the disk drives. The data stays in the read cache until it is flushed. For write caching, a write operation stores data from the host in cache until it can be written to the disk drives.

**Important:** Changing default cache settings is not recommended – IBM recommends that you do not change the default cache setting values. Incorrect cache settings might severely affect the storage subsystem performance.

The script command set provides two commands to define cache properties:

- set storageSubsystem
- set logicalDrive

Use the **set storageSubsystem** command to change the cache block size, the cache flush start value, and the cache stop value. The command takes this form:

```
set storageSubsystem cacheBlockSize=cacheBlockSizeValue |
cacheFlushStart=cacheFlushStartSize |
cacheFlushStop=cacheFlushStopSize
```

You can enter one, two, or all three of the parameters on the command line.

The cache block size value defines the size of the data block that is used by the controller in transferring data into or out of the cache. You can set the cache block size to either 4KB or 16KB. The value that you use applies to the entire storage subsystem and all of the logical drives in the storage subsystem. For redundant controller configurations, this value includes all logical drives owned by both controllers. Use smaller cache block sizes for systems that require transaction processing requests or I/O streams that are typically small and random. Use larger cache block sizes for large I/O, sequential, high-bandwidth applications. The choice of block sizes than 16-KB block sizes. This example shows how to set the **cacheBlockSize** parameter:

c:\...\client>smcli 123.45.67.89 -c "set storageSubsystem cacheBlockSize=16;"

To prevent data loss or corruption, the controller periodically writes cache data to the disk drives (flushes the cache) when the amount of unwritten data in the cache reaches a predefined level, called a start percentage. The controller also writes cache data to the disk drives when data has been in the cache for a predetermined amount of time. The controller writes data to the disk drives until the amount of data in the cache drops to a stop percentage level. Use the set storageSubsystem command to set the start value and the stop value as percentages of the filled and unfilled, respectively, capacity of the cache. For example, you can specify that the controller start flushing the cache when it reaches 80 percent full (the default, which can be overridden) and stop flushing the cache when it reaches 84 percent empty. This example shows how to set these parameters:

c:\...\client>smcli 123.45.67.89 -c "set storageSubsystem cacheFlushStart=80 cacheFlushStop=84;"

Low start percentages and low stop percentages provide for maximum data protection. For both low start percentages and low stop percentages, the chance that data requested by a read command is not in the cache is increased. When the data is not in the cache, the cache hit percentage for writes and I/O requests decreases. Low start values and low stop values also increase the number of writes that are necessary to maintain the cache level. Increasing the number of writes increases the system overhead and further decreases performance.

Use the set logicalDrive command to change settings for the cache flush modifier, cache without batteries enabled, mirror cache enabled, the read ahead multiplier, read cache enabled, and write cache enabled. Use this command to set properties for all of the logical drives or for a specific logical drive in a subsystem. The command takes this form:

```
set (allLogicalDrives | logicalDrive [logicalDriveName] |
logicalDrives [logicalDriveName1 ... logicalDriveNameN]
logicalDrive <wwID>) |
cacheFlushModifier=cacheFlushModifierValue |
cacheWithoutBatteryEnabled=(TRUE | FALSE) |
mirrorCacheEnabled=(TRUE | FALSE) |
readCacheEnabled=(TRUE | FALSE) |
writeCacheEnabled=(TRUE | FALSE) |
cacheReadPrefetch=(TRUE | FALSE)
```

The **cacheFlushModifier** parameter defines the amount of time that data stays in the cache before it is written to the disk drives. The following table lists the values for the **cacheFlushModifier** parameter.

Value	Meaning
Immediate	Data is flushed as soon as it is placed into the cache.
250	Data is flushed after 250 ms.
500	Data is flushed after 500 ms.
750	Data is flushed after 750 ms.
1	Data is flushed after 1 s.
1500	Data is flushed after 1500 ms.
2	Data is flushed after 2 s.
5	Data is flushed after 5 s.
10	Data is flushed after 10 s.
20	Data is flushed after 20 s.
60	Data is flushed after 60 s (1 min.).
120	Data is flushed after 120 s (2 min.).
300	Data is flushed after 300 s (5 min.).

Table 18. Values for the cacheFlushModifier Parameter

Table 18. Values for the cacheFlushModifier Parameter (continued)

Value	Meaning
1200	Data is flushed after 1200 s (20 min.).
3600	Data is flushed after 3600 s (1 hr.).
Infinite	Data in cache is not subject to any age or time constraints. The data is flushed based on other criteria managed by the controller.

This example shows how to set this parameter value for all of the logical drives in the storage subsystem:

c:\...\client>smcli 123.45.67.89 -c "set alllogicalDrives cacheFlushModifier=10;"

**Important:** Do not set the value of the **cacheFlushModifier** parameter above 10 seconds. An exception is for testing purposes. After running any tests in which you have set the values of the **cacheFlushModifier** parameter above 10 seconds, return the value of the **cacheFlushModifier** parameter to 10 or less seconds.

The **cacheWithoutBatteryEnabled** parameter turns on or turns off the ability of a host to perform write caching without backup batteries in a controller. To enable write caching without batteries, set this parameter to TRUE. To disable write caching without batteries, set this parameter to FALSE. If you set this parameter to TRUE, write caching continues, even when the controller batteries are completely discharged, not fully charged, or not present. If you do not have an uninterruptible power supply (UPS) and you enable this parameter, you can lose data if power to the storage subsystem fails.

**Note:** Attention: Possible loss of data – Because you can lose data when the **cacheWithoutBatteryEnabled** parameter is enabled, IBM recommends that you do not enable this parameter in a production environment.

This example shows how to set this parameter value: c:\...\client>smcli 123.45.67.89 -c "set logicalDrive [\"Engineering\"] cacheWithoutBatteryEnabled=FALSE;"

The **mirrorCacheEnabled** parameter turns on or turns off write caching with mirroring. Write caching with mirroring permits cached data to be mirrored across redundant controllers that have the same cache size. Data written to the cache memory of one controller is also written to the cache memory of the second controller. If one controller fails, the second controller can complete all outstanding write operations.

**Attention: Possible loss of data** – If the **mirrorCacheEnabled** parameter is disabled and one of the controllers fails, you might lose data that was cached in the failed controller memory but not yet written to disk. Therefore, IBM recommends that you do not disable this parameter in a production environment.

To use this option, these conditions must exist:

- The controller pair must be an active/active pair.
- The controllers must have the same size cache.

To enable write caching with mirroring, set this parameter to TRUE. To disable write caching with mirroring, set this parameter to FALSE. This example shows how to set this parameter:

```
c:\...\client>smcli 123.45.67.89 -c "set logicalDrive
[\"Accounting\"] mirrorCacheEnabled=TRUE;"
```

The **readCacheEnabled** parameter turns on or turns off the ability of the host to read data from the cache. Read caching enables read operations from the host to be stored in controller cache memory. If a host requests data that is not in the cache, the controller reads the needed data blocks from the disk drives and places them in the cache. Until the cache is flushed, all of the other requests for this data are fulfilled with cache data rather than from a read, which increases throughput. To enable read caching, set this parameter to TRUE. To disable read caching, set this parameter to FALSE. This example shows how to set this parameter:

c:\...\client>smcli 123.45.67.89 -c "set logicalDrive
[\"Balance\_04\"] readCacheEnabled=TRUE;"

The **writeCacheEnabled** parameter turns on or turns off the ability of the host to write data to the cache. Write caching enables write operations from the host to be stored in cache memory. The logical drive data in the cache is automatically written to the disk drives every 10 seconds. To enable write caching, set this parameter to TRUE. To disable write caching, set this parameter to FALSE. This example shows how to set this parameter:

c:\...\client>smcli 123.45.67.89 -c "set allLogicalDrives
writeCacheEnabled=TRUE;"

The **cacheReadPrefetch** parameter turns on or turns off the ability of the controller to read additional data blocks into cache. When you turn on cache read prefetch, the controller copies additional data blocks into cache while it is reading requested data blocks from a disk drive into cache. This action increases the chance that a future request for data can be fulfilled from the cache, which improves the speed with which data is accessed. The number of additional data blocks that the controller reads into cache is determined by the storage subsystem configuration settings that you use. Cache read prefetch is important for applications that use sequential I/O, such as multimedia applications.

Valid values for the **cacheReadPrefetch** parameter are TRUE or FALSE. If you want to turn on cache read prefetch, set the **cacheReadPrefetch** parameter to TRUE. If you want to turn off cache read prefetch, set the **cacheReadPrefetch** parameter to FALSE. This example shows how to set this parameter:

c:\...\client>smcli 123.45.67.89 -c "set logicalDrive
[\"Engineering\_1\" \"Engineering\_2\"] cacheReadPrefetch=TRUE;"

#### Setting the Modification Priority

Modification priority defines how much processing time is allocated for logical drive modification operations. Time allocated for logical drive modification operations affects system performance. Increases in logical drive modification priority can reduce read/write performance. The modification priority affects these operations:

- Copyback
- Reconstruction
- Initialization
- Changing the segment size
- · Defragmentation of an subsystem

- Adding free capacity to an subsystem
- · Changing the RAID level of an subsystem

The lowest priority rate favors system performance, but the modification operation takes longer. The highest priority rate favors the modification operation, but the system performance might be degraded.

Use the **set logicalDrive** command to define the modification priority for a logical drive. The command takes this form:

set (allLogicalDrives | logicalDrive [logicalDriveName] |
logicalDrives [logicalDriveName1 ... logicalDriveNameN] logicalDrive <wwID> |
accessLogicalDrive)
modificationPriority=(highest | high | medium | low | lowest)

This example shows how to use this command to set the modification priority for logical drives named Engineering\_1 and Engineering\_2:

```
c:\...\client>smcli 123.45.67.89 -c "set logicalDrive
[\"Engineering 1\" \"Engineering 2\" modificationPriority=lowest;"
```

The modification rate is set to lowest so that system performance is not significantly reduced by modification operations.

#### Assigning Global Hot Spares

You can assign or unassign global hot spares by using the **set drive** command. To use this command, you must identify the location of the disk drives by the enclosure ID and the slot ID. Then, you set the **hotSpare** parameter to TRUE to enable the hot spare or FALSE to disable an existing hot spare. The command takes this form:

```
set (drive [enclosureID,slotID] | drives
[enclosureID1,slotID1 ... enclosureIDn,slotIDn]) hotSpare=(TRUE | FALSE)
```

This example shows how to set hot spare disk drives:

c:\...\client>smcli 123.45.67.89 -c "set drives [1,2 1,3] hotSpare=TRUE;"

Enter the enclosure ID and the slot ID of each disk drive that you want to use. Enclose the list in square brackets, separate the enclosure ID value and the slot ID value of a disk drive with a comma, and separate each enclosure ID and slot ID pair with a space.

# Saving a Configuration to a File

After you have created a new configuration or if you want to copy an existing configuration for use on other storage subsystems, you can save the configuration to a file by using the **save storageSubsystem configuration** command. Saving the configuration creates a script file that you can run on the command line. The command takes this form:

```
save storageSubsystem configuration file="filename"
[(allconfig | globalSettings=(TRUE | FALSE))
logicalDriveConfigAndSettings=(TRUE | FALSE)
hostTopology=(TRUE | FALSE) | lunMappings=(TRUE | FALSE)]
```

Attention: Possible loss of data – When information is written to a file, the script engine does not check to determine if the file name already exists. If you choose the name of a file that already exists, the script engine writes over the information in the file without warning.

You can choose to save the entire configuration or specific configuration features. This example shows how to set this parameter value:

c:\...\client>smcli 123.45.67.89 -c "save storageSubsystem configuration file=\"c:\\folder\\storageSubsystemconfig1.scr\";"

In this example, the name folder is the folder in which you choose to place the profile file and storageSubsystemconfig1.scr is the name of the file. You can choose any folder and any file name. The file extension for a configuration file is .scr. The Storage Manager software uses this extension when it creates the configuration file.

# **Chapter 5. Using the FlashCopy Premium Feature**

The FlashCopy premium feature creates a FlashCopy logical drive that you can use as a backup of your data. A FlashCopy logical drive is a logical point-in-time image of a standard logical drive. Because it is not a physical copy, a FlashCopy logical drive is created more quickly than a physical copy and requires less storage space on the disk drive. Typically, you create a FlashCopy logical drive so that an application, such as a backup application, can access the FlashCopy logical drive and read the data while the base logical drive stays online and user accessible. You can also create several FlashCopy logical drives of a base logical drive and write data to the FlashCopy logical drives to perform testing and analysis.

FlashCopy logical drives provide these capabilities:

- Create a complete image of the data on a base logical drive at a particular point in time
- Use only a small amount of storage space
- Provide for quick, frequent, non-disruptive backups, or testing new versions of a database system without affecting real data
- Provide for FlashCopy logical drives to be read, written, and copied
- Use the same availability characteristics of the base logical drive (such as RAID protection and redundant path failover)
- Mapping the FlashCopy logical drive and making it accessible to any host on a storage area network (SAN). You can make FlashCopy data available to secondary hosts for read access and write access by mapping the FlashCopy to the hosts
- Creating up to 16 FlashCopies per logical drive and up to 1024 FlashCopies per storage subsystems (the maximum number of FlashCopy logical drives is one-half of the total number of logical drives that are supported by the controller).
- Increasing the capacity of a FlashCopy logical drive

#### How FlashCopy Works

Three components comprise a FlashCopy logical drive: the base logical drive, the FlashCopy logical drive, and the FlashCopy repository logical drive. The following table lists the components and briefly describes what they do.

Component	Description
Base logical drive	A standard logical drive from which the FlashCopy is created
FlashCopy logical drive	A logical point-in-time image of a standard logical drive
FlashCopy repository logical drive	A logical drive that contains FlashCopy metadata and copy-on-write data for a particular FlashCopy logical drive

Table 19. Components of a FlashCopy Logical Drive

Based on information that you provide through the script commands, the storage management software creates an empty FlashCopy repository logical drive and

defines the mapping from a base logical drive to the FlashCopy repository logical drive. The FlashCopy repository logical drive holds changed data that a host writes to the base logical drive. When the FlashCopy repository logical drive is first created, it holds only the metadata about the FlashCopy logical drive with which it is associated.

**Note:** When you first create a FlashCopy repository logical drive, briefly stop all of the write operations to the base logical drive so that a stable image of the base logical drive is available.

When the host writes to the base logical drive, the new data is also copied to the FlashCopy repository logical drive. This action is called *copy-on-write*. A FlashCopy is constructed by combining the updated data in the FlashCopy repository logical drive with data in the base logical drive that has not been altered. This action creates a complete copy of the base logical drive at a specific point in time. The FlashCopy appears as a logical drive that contains the original data at the time of creation, but the FlashCopy is actually an image that is the combination of the FlashCopy repository logical drive and the original base logical drive. The FlashCopy repository logical drive, which houses original data that has been changed, is the only additional disk drive space that is needed for the FlashCopy logical drive. The additional disk drive space is typically 10 percent to 20 percent of the disk drive space of the base logical drive and varies depending on the amount of changes to the data. The longer a FlashCopy logical drive is active, the larger the FlashCopy repository logical drive must be. The default size of the FlashCopy repository logical drive is 20 percent of the base logical drive; however, you can set the size of the FlashCopy repository logical drive to other values.

You can read, write, and copy a FlashCopy logical drive. Data written by a host to the FlashCopy logical drive is handled in the FlashCopy repository logical drive. When a write occurs to the base logical drive of a FlashCopy logical drive, the new data also overwrites the appropriate FlashCopy repository logical drive data. The following table lists the FlashCopy logical drive commands.

Command	Description
create flashcopyLogicalDrive	This command creates a FlashCopy logical drive.
re-create flashcopy	This command starts a fresh copy-on-write operation by using an existing FlashCopy logical drive.
re-create flashcopy collection	This command restarts multiple FlashCopy logical drives as one batch operation using one or many existing FlashCopy logical drives.
set (flashcopyLogicalDrive)	This command defines the properties for a FlashCopy logical drive and lets you rename a FlashCopy logical drive.
stop flashcopy	This command stops a copy-on-write operation.

Table 20. FlashCopy Logical Drive Commands

# Creating a FlashCopy Logical Drive

The **create FlashCopyLogicalDrive** command provides three methods for defining the disk drives for your FlashCopy repository logical drive:

- Defining the disk drives for the FlashCopy repository logical drive by their enclosure IDs and their slot IDs
- Defining an subsystem in which the FlashCopy repository logical drive resides. In addition, you can define the capacity of the FlashCopy repository logical drive.
- Defining the number of disk drives, but not specific disk drives, for the FlashCopy repository logical drive.

When you use the **create FlashCopyLogicalDrive** command to create a FlashCopy logical drive, the minimum information that you need to provide is the standard logical drive that you want to use for the base logical drive. When you create a FlashCopy logical drive by using minimum information, the storage management software provides default values for the other property parameters that are required for a completely defined FlashCopy logical drive.

# Creating a FlashCopy Logical Drive with User-Assigned Disk Drives

Creating a FlashCopy logical drive by assigning the disk drives provides flexibility in defining your configuration by letting you choose from the available disk drives in your storage subsystem. When you choose the disk drives for your FlashCopy logical drive, you automatically create a new subsystem. You can specify which disk drives to use and the RAID level for the new subsystem. The command takes this form:

```
create flashcopyLogicalDrive baseLogicalDrive="baseLogicalDriveName"
[repositoryRAIDLevel=(1 | 3 | 5 | 6)
(repositoryDrives=(enclosureID1,slotID1 ... enclosureIDn,slotIDn)
userLabel="flashcopyLogicalDriveName"
warningThresholdPercent=percentValue
repositoryPercentOfBase=percentValue
repositoryUserLabel="repositoryName"
repositoryFullPolicy=(failBaseWrites | failFlashCopy)]
[enclosureLossProtect=(TRUE | FALSE)]
```

This example shows a command in which users assign the disk drives:

c:\...\client>smcli 123.45.67.89 -c "create flashcopyLogicalDrive baseLogicalDrive=\"Mars\_Spirit\_4\" repositoryRAIDLevel=5 repositoryDrives=(1,1 1,2 1,3 1,4 1,5);"

The command in this example creates a new FlashCopy of the base logical drive Mars\_Spirit\_4. The FlashCopy repository logical drive consists of five disk drives that form a new subsystem. The new subsystem has a RAID level of 5. This command also takes a FlashCopy of the base logical drive, which starts the copy-on-write operation.

This example shows how to use the command in a script file:

create flashcopyLogicalDrive baseLogicalDrive="Mars\_Spirit\_4" repositoryRAIDLevel=5
repositoryDrives=(1,1 1,2 1,3 1,4 1,5);

This example shows a minimal version of the command:

```
c:\...\client>smcli 123.45.67.89 -c "create flashcopyLogicalDrive
baseLogicalDrive=\"Mars_Spirit_4\";"
```

The command in this example creates a new FlashCopy for the base logical drive Mars\_Spirit\_4. The FlashCopy repository logical drive is created in the same

subsystem as the base logical drive, which means that the FlashCopy repository logical drive has the same RAID level as the base logical drive. This command starts the copy-on-write operation.

This example shows how to use the command in a script file: create flashcopyLogicalDrive baseLogicalDrive="Mars Spirit 4";

# Creating a FlashCopy Logical Drive with Software-Assigned Disk Drives

With this version of the **create FlashCopyLogicalDrive** command, you choose an existing subsystem in which to place the FlashCopy repository logical drive. The storage management software determines which disk drives to use. You can also define how much space to assign to the FlashCopy repository logical drive. Because you are using an existing subsystem, the RAID level for the FlashCopy logical drive defaults to the RAID level of the subsystem in which you place it. You cannot define the RAID level for the FlashCopy logical drive. The command takes this form:

```
create FlashCopyLogicalDrive baseLogicalDrive="baseLogicalDriveName"
[repositorysubsystem=subsystemNumber
freeCapacityArea=freeCapacityIndexNumber
userLabel="FlashCopyLogicalDriveName"
warningThresholdPercent=percentValue
repositoryPercentOfBase=percentValue
repositoryUserLabel=repositoryName
repositoryFullPolicy=(failBaseWrites | failFlashCopy)]
[enclosureLossProtect=(TRUE | FALSE)]
```

This example shows a command in which the storage management software assigns the disk drives:

```
c:\...\client>smcli 123.45.67.89 -c "create
FlashCopyLogicalDrive baseLogicalDrive=\"Mars_Spirit_4\"
repositorysubsystem=2 freeCapacityArea=2;"
```

The command in this example creates a new FlashCopy repository logical drive in subsystem 2. The base logical drive is Mars\_Spirit\_4. The size of the FlashCopy repository logical drive is 4 GB. This command also takes a FlashCopy of the base logical drive, starting the copy-on-write operation.

When you define the capacity of a FlashCopy repository logical drive, specify a size that is 20 percent of the size of the base logical drive. In the previous example, the size of the FlashCopy repository logical drive is set to 4 GB. The underlying assumption is that the base logical drive size is 20 GB ( $0.2 \times 20$  GB= 4 GB).

This example shows how to use the command in a script file: create FlashCopyLogicalDrive baseLogicalDrive="Mars\_Spirit\_4" repositorysubsystem=2 freeCapacityArea=2;

# Creating a FlashCopy Logical Drive by Specifying a Number of Disk Drives

With this version of the **create flashcopyLogicalDrive** command, you must specify the number of disk drives and the RAID level that you want for the FlashCopy repository logical drive. This version of the **create** 

**flashcopyLogicalDrive** command creates a new subsystem. You must have disk drives in the storage subsystem that are not assigned to an subsystem for this command to work.

```
create flashcopyLogicalDrive baseLogicaldrive="baseLogicalDriveName"
[repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDriveCount=numberOfDrives
driveType=(fibre | SATA | SAS)
userLabel="flashcopyLogicalDriveName"
warningThresholdPercent=percentValue
repositoryPercentOfBase=percentValue
repositoryUserLabel="repositoryName"
repositoryFullPolicy=(failBaseWrites | failFlashCopy)]
[enclosureLossProtect=(TRUE | FALSE)]
```

This example shows how to use a command in which users specify the number of disk drives:

```
c:\...\client>smcli 123.45.67.89 -c "create flashcopyLogicalDrive
baseLogicalDrive=\"Mars_Spirit_4\" repositoryRAIDLevel=5
repositoryDriveCount=3;"
```

The command in this example creates a new FlashCopy repository logical drive that consists of three disk drives. Three disk drives comprise a new subsystem that has a RAID level of 5. This command also takes a FlashCopy of the base logical drive, which starts the copy-on-write operation.

This example shows how to use the command in a script file: create flashcopyLogicalDrive baseLogicalDrive= "Mars\_Spirit\_4" repositoryRAIDLevel=5 repositoryDriveCount=3;

# **User-Defined Parameters**

Use the parameters in the **create flashcopyLogicalDrive** command to define the FlashCopy logical drive to suit the requirements of your storage subsystem. The following table lists the parameters and briefly describes what the parameters do.

Parameter	Description
driveType	The type of disk drive that you want to use for the FlashCopy repository logical drive. Valid disk drive types are fibre, SATA, or SAS. This parameter works only with the count-based repository method of defining a FlashCopy logical drive.
repositorysubsystem	The subsystem in which you want to build the FlashCopy repository logical drive. The default value is to build the FlashCopy repository logical drive in the same subsystem as the base logical drive.
freeCapacityArea	The amount of storage space that you want to use for the FlashCopy repository logical drive. Free storage space is defined in units of bytes, KB, MB, GB, or TB.

Table 21. FlashCopy Logical Drive Parameters

Parameter	Description
userLabel	The name that you want to give to the FlashCopy logical drive. If you do not choose a name for the FlashCopy logical drive, the software creates a default name by using the base logical drive name. For example, with a base logical drive name of Mars_Spirit_4:
	FlashCopy logical drive, the default FlashCopy logical drive name is Mars_Spirit_4-1.
	• When the base logical drive already has <i>n</i> -1 number of FlashCopy logical drives, the default name is Mars_Spirit_4- <i>n</i> .
repositoryUserLabel	The name that you want to give to the FlashCopy repository logical drive. If you do not choose a name for the FlashCopy repository logical drive, the software creates a default name by using the base logical drive name. For example, if the base logical drive name is Mars_Spirit_4 and does not have an associated FlashCopy repository logical drive, the default FlashCopy repository logical drive name is Mars_Spirit_4-R1. If the base logical drive already has <i>n</i> -1 number of FlashCopy repository logical drive.
warningThresholdPercent	The percentage of the capacity that you will permit the FlashCopy repository logical drive to get before you receive a warning that the FlashCopy repository logical drive is nearing full. The warning value is a percentage of the total capacity of the FlashCopy repository logical drive. The default value is 50, which represents 50 percent of the total capacity. (You can change this value later by using the <b>set</b> <b>flashcopyLogicalDrive</b> command.)
repositoryPercentOfBase	The size of the FlashCopy repository logical drive as a percentage of the base logical drive size. The default value is 20, which represents 20 percent of the base logical drive size.
repositoryFullPolicy	The type of FlashCopy processing that you want to continue if the FlashCopy repository logical drive is full. You can choose to fail writes to the base logical drive (failBaseWrites) or fail writes to the FlashCopy logical drive (failFlashCopy). The default value is failFlashCopy.

Table 21. FlashCopy Logical Drive Parameters (continued)

This example shows the **create flashcopyLogicalDrive** command that includes user-defined parameters:

```
c:\...\client>smcli 123.45.67.89 -c "create
flashcopyLogicalDrive baseLogicalDrive=\"Mars_Spirit_4\"
repositoryRAIDLevel=5 repositoryDriveCount=5
driveType=fibre userLabel=\"Mars_Spirit_4_EnhFlashCopy1\"
repositoryUserLabel=\"Mars_Spirit_4rep1\"
warningThresholdPercent=75 repositoryPercentOfBase=40
repositoryFullPolicy=failFlashCopy;"
```

This example shows how to use the command in a script file:

```
create flashcopyLogicalDrive baseLogicalDrive="Mars_Spirit_4"
repositoryRAIDLevel=5 repositoryDriveCount=5 driveType=fibre
userLabel="Mars_Spirit_4_EnhFlashCopy1"
repositoryUserLabel="Mars_Spirit_4_rep1"
warningThresholdPercent=75 repositoryPercentOfBase=40
repositoryFullPolicy=failFlashCopy;
```

# FlashCopy Logical Drive Names and FlashCopy Repository Logical Drive Names

The FlashCopy logical drive names and the FlashCopy repository logical drive names can be any combination of alphanumeric characters, hyphens, and underscores. The maximum length of the logical drive names is 30 characters. You must enclose the logical drive name in double quotation marks. The character string cannot contain a new line. Ensure that you use unique names; if you do not use unique names, the controller firmware returns an error.

One technique for naming the FlashCopy logical drive and the FlashCopy repository logical drive is to add a hyphenated suffix to the original base logical drive name. The suffix distinguishes between the FlashCopy logical drive and the FlashCopy repository logical drive. For example, if you have a base logical drive with a name of Engineering Data, the FlashCopy logical drive can have a name of Engineering Data-S1, and the FlashCopy repository logical drive can have a name of Engineering Data-R1.

If you do not choose a unique name for either the FlashCopy logical drive or the FlashCopy repository logical drive, the controllers create a default name by using the base logical drive name. These examples are FlashCopy logical drive names that the controllers might create:

- If the base logical drive name is aaa and does not have a FlashCopy logical drive, the default FlashCopy logical drive name is aaa1.
- If the base logical drive already has *n*-1 number of FlashCopy logical drives, the default name is aaa-*n*.
- If the base logical drive name is aaa and does not have a FlashCopy repository logical drive, the default FlashCopy repository logical drive name is aaa-R1
- If the base logical drive already has *n*-1 number of FlashCopy repository logical drives, the default name is aaa-R*n*.

In the examples from the previous section, the user-defined FlashCopy logical drive name was Mars\_Spirit\_4\_EnhFlashCopy1, and the user-defined FlashCopy repository logical drive name was Mars\_Spirit\_4\_rep1. The default name that was provided by the controller for the FlashCopy logical drive was Mars\_Spirit\_4-1. The default name that was provided by the controller for the FlashCopy repository logical drive was Mars\_Spirit\_4-R1.

# Changing FlashCopy Logical Drive Settings

Use the **set (FlashCopy) logicalDrive** command to change these property settings for a FlashCopy logical drive:

- The FlashCopy logical drive name
- The warning threshold percent
- The FlashCopy repository full policy

This example shows how to change a FlashCopy logical drive name.

c:\...\client>smcli 123.45.67.89 -c "set logicalDrive [\"Mars\_Spirit\_4-1\"] userLabel=\"Mars\_Odyssey\_3-2\";"

This example shows how to use the command in a script file: set logicalDrive ["Mars\_Spirit\_4-1"] userLabel="Mars\_Odyssey\_3-2";

When you change the warning threshold percent and the FlashCopy repository full policy, you can apply the changes to one or several FlashCopy logical drives with this command. This example shows how to use the **set (FlashCopy) logicalDrive** command to change these properties on more than one FlashCopy logical drive:

c:\...\client>smcli 123.45.67.89 -c "set logicalDrive
[\"Mars\_Spirit\_4-1\" \"Mars\_Spirit\_4-2\" \"Mars\_Spirit\_4-3\"
warningThresholdPercent=50
repositoryFullPolicy=failBaseWrites;"

This example shows how to use the command in a script file:

set logicalDrive ["Mars\_Spirit\_4-1" "Mars\_Spirit\_4-2"
"Mars\_Spirit\_4-3"] warningThresholdPercent=50
repositoryFullPolicy=failBaseWrites;

#### Stopping, Restarting, and Deleting a FlashCopy Logical Drive

When you create a FlashCopy logical drive, copy-on-write starts running immediately. As long as a FlashCopy logical drive is enabled, storage subsystem performance is impacted by the copy-on-write operations to the associated FlashCopy repository logical drive.

If you no longer want copy-on-write operations to run, you can use the **stop flashcopy logicalDrive** command to stop the copy-on-write operations. When you stop a FlashCopy logical drive, the FlashCopy logicalDrive and the FlashCopy repository logical drive are still defined for the base logical drive. Only copy-on-write has stopped. This example shows how to stop a FlashCopy logical drive:

c:\...\client>smcli 123.45.67.89 -c "stop flashcopy logicalDrives
[\"Mars\_Spirit\_4-2\" \"Mars\_Spirit\_4-3\"];"

This example shows how to use the command in a script file: stop flashcopy logicalDrives ["Mars\_Spirit\_4-2" "Mars\_Spirit\_4-3"];

When you stop the copy-on-write operations for a specific FlashCopy logical drive, only that FlashCopy logical drive is disabled. All of the other FlashCopy logical drives stay in operation.

When you want to restart a copy-on-write operation, use the **re-create flashcopy logicalDrive** command or the **re-create flashcopy collection** command. The **re-create flashcopy logicalDrive** command starts a fresh copy-on-write operation by using an existing FlashCopy logical drive.

**Note:** The FlashCopy logical drive must be in either an Optimal state or a Disabled state.

When you restart a FlashCopy logical drive, these actions occur:

- All copy-on-write data previously on the FlashCopy repository logical drive is overwritten.
- FlashCopy logical drive parameters and FlashCopy repository logical drive parameters stay the same as the previously disabled FlashCopy logical drive and
the previously disabled FlashCopy repository logical drive. You can also change the **userLabel** parameter, the **warningThresholdPercent** parameter, and the **repositoryFullPolicy** parameter when you restart the FlashCopy logical drive.

• The original names for the FlashCopy repository logical drive are retained.

This example shows how to restart a FlashCopy logical drive:

c:\...\client>smcli 123.45.67.89 -c "re-create flashcopy logicalDrives
[\"Mars\_Spirit\_4-2\" \"Mars\_Spirit\_4-3\"];"

This example shows how to use the command in a script file: re-create flashcopy logicalDrives ["Mars\_Spirit\_4-2" "Mars\_Spirit\_4-3"];

If you do not intend to use a FlashCopy logical drive again, you can delete the FlashCopy logical drive by using the **delete logicalDrive** command. When you delete a FlashCopy logical drive, the associated FlashCopy repository logical drive also is deleted.

# Chapter 6. Using the Enhanced FlashCopy Images premium feature

Controller firmware version 7.83 and later supports a new version of FlashCopy, called Enhanced FlashCopy. Enhanced FlashCopy provides new ways of creating a large number of FlashCopies of a base logical drive, with minimum performance degradation due the **copy-on-write** operations feature. In an existing FlashCopy operation, a **Copy-on-write** operation will be performed for each FlashCopy of the base logical drive. The higher the number of FlashCopies, the higher the number of **copy-on-write** operations. In an Enhanced FlashCopy operation, only one **copy-on-write** operation will be performed for a group of up to 32 FlashCopies of a base logical drive. Up to 32 Enhanced FlashCopy images of a base logical drive can be grouped into an Enhanced FlashCopy group with an Enhanced FlashCopy group repository drive.

An Enhanced FlashCopy image is a logical image of the content of an associated base logical drive created at a specific moment. An Enhanced FlashCopy image can be thought of as a restore point. A host cannot directly read from or write to the Enhanced FlashCopy image because the Enhanced FlashCopy image is used to save only the transient data captured from the base logical drive.

You must create an Enhanced FlashCopy logical drive to enable host access to a complete copy of the data contained in the Enhanced FlashCopy image. The Enhanced FlashCopy logical drive contains its own repository, which is used to save any subsequent modifications made by the host to the base logical drive without affecting the referenced Enhanced FlashCopy image.

Enhanced FlashCopy images are useful any time you need to be able to roll back to a known good data set at a specific point in time. For example, before performing an operation on a logical drive that you might want to reverse, you can create an Enhanced FlashCopy image to enable the reverse operation and restore the entire logical drive to its previous state. An Enhanced FlashCopy image is created almost instantaneously, and initially uses no disk space, because it stores only the incremental changes needed to roll the logical drive back to the point-in-time when the Enhanced FlashCopy image was created.

You can create Enhanced FlashCopy images manually or automate the process with a schedule. With a schedule, an Enhanced FlashCopy image is generated automatically, based on the date and time you choose, and is displayed with its associated Enhanced FlashCopy group. You can create Enhanced FlashCopy images for these storage objects:

• Standard logical drives

A standard logical drive has a fixed physical capacity that is entirely reserved for the data on the logical drive.

• Thin logical drives

A thin logical drive is dynamic. It has a small initial physical capacity and a large virtual capacity. The thin logical drive automatically expands to the virtual capacity as data is added to the logical drive.

Consistency groups

A consistency group is a container that holds several logical drives so that you can manage all of the logical drives as a single entity.

To create an Enhanced FlashCopy image, you must first create an Enhanced FlashCopy group and the associated Enhanced FlashCopy repository logical drive. You can either use an existing repository logical drive or create a repository logical drive by performing these actions:

- Identifying the subsystem or disk pool in which you want to place the repository logical drive
- The capacity for the repository logical drive

You can delete any Enhanced FlashCopy image in an Enhanced FlashCopy group. When an Enhanced FlashCopy image is deleted, it's definition is removed from the system, and the space occupied by the Enhanced FlashCopy image in the repository is released and made available for reuse within the Enhanced FlashCopy group. The Enhanced FlashCopy image can be deleted only in the order of creation, with the oldest one deleted first.

You can roll back data by either:

- Creating an Enhanced FlashCopy logical drive of an Enhanced FlashCopy image, which enables you to retrieve deleted files from that Enhanced FlashCopy logical drive (the base logical drive remains undisturbed).
- Restoring an Enhanced FlashCopy image to the base logical drive, which enables you to roll back the base logical drive to a previous state.

#### Characteristics of Enhanced FlashCopy Images

- Enhanced FlashCopy images are always created inside Enhanced FlashCopy groups.
- Each Enhanced FlashCopy image is associated with exactly one Enhanced FlashCopy group.
- Up to four enhanced FlashCopy groups can be created for a given base logical drive.
- Up to 32 Enhanced FlashCopy images can be created per Enhanced FlashCopy group.
- Up to 256 Enhanced FlashCopy logical drives can be created for Enhanced FlashCopy images per storage subsystem.

## About Enhanced FlashCopy Images Premium Feature Enhanced FlashCopy Groups

An Enhanced FlashCopy group is a collection of Enhanced FlashCopy images of a single associated base logical drive. An Enhanced FlashCopy image is always created within the context of an Enhanced FlashCopy group, so the identity of the Enhanced FlashCopy group is a key component that must be specified when creating an Enhanced FlashCopy image.

The purpose of an Enhanced FlashCopy group is to maintain a sequence of Enhanced FlashCopy images of a given base logical drive without impacting performance. You can set up a schedule for an Enhanced FlashCopy group to automatically create an Enhanced FlashCopy image at a specific time in the future or on a regular basis.

Each Enhanced FlashCopy group has a strict ordering of Enhanced FlashCopy images based on the time an Enhanced FlashCopy image was created. These terms describe the time relationship of Enhanced FlashCopy images:

Successor

An Enhanced FlashCopy image that is created after another Enhanced FlashCopy image has been created, also called newest.

Predecessor

An Enhanced FlashCopy image that is created before another Enhanced FlashCopy image, also called oldest.

The terms newest and oldest are arguments used in the CLI commands.

An Enhanced FlashCopy group uses an Enhanced FlashCopy group repository logical drive to save all data for the Enhanced FlashCopy images contained in the group. An Enhanced FlashCopy image operation uses less disk space than a full physical copy because the data stored in the Enhanced FlashCopy logical drive repository logical drive is only the data that has changed since the Enhanced FlashCopy image has been taken.

An Enhanced FlashCopy group is created initially with one Enhanced FlashCopy group repository logical drive, with a default capacity of 40% the capacity of the base logical drive. The Enhanced FlashCopy group repository logical drive initially contains a small amount of data, which increases with subsequent data updates. If needed, you can increase the size of the Enhanced FlashCopy logical drive repository logical drive by combining unused repository logical drives to the Enhanced FlashCopy group repository logical drives to the Enhanced FlashCopy group repository logical drives to the Enhanced FlashCopy group repository logical drives.

The only types of logical drives that can contain an Enhanced FlashCopy group are these:

- Standard logical drives
- · Thin logical drives

Non standard logical drives, such as Enhanced FlashCopy logical drives, cannot be used for Enhanced FlashCopy groups. The base logical drive can reside on either an subsystem or a disk pool.

**Note:** Enhanced FlashCopy logical drives and Enhanced FlashCopy groups cannot coexist on the same base logical drive.

#### Characteristics of Enhanced FlashCopy Groups:

- Enhanced FlashCopy groups can be initially created with or without Enhanced FlashCopy images.
- Depending on your configuration, a single associated base logical drive has a maximum limit of Enhanced FlashCopy groups.
- Depending on your configuration, an Enhanced FlashCopy group has a maximum limit of Enhanced FlashCopy images.
- Each Enhanced FlashCopy image is associated with exactly one Enhanced FlashCopy group.

## **Repository Logical drives**

Every Enhanced FlashCopy image is created in the context of exactly one Enhanced FlashCopy group. An Enhanced FlashCopy group is a sequence of Enhanced FlashCopy images of a single associated base logical drive. A given Enhanced FlashCopy group has exactly one Enhanced FlashCopy logical drive repository logical drive that is used to save data for all of the Enhanced FlashCopy images that are part of the Enhanced FlashCopy group. When a new Enhanced FlashCopy image is created in an Enhanced FlashCopy group, all previously-created Enhanced FlashCopy images in that Enhanced FlashCopy group become static and stop consuming additional repository space. The newly-created Enhanced FlashCopy image begins consuming repository space during processing of ensuing copy-on-write actions for the associated base logical drive.

The Enhanced FlashCopy group repository logical drive is a standard logical drive that can be expanded to accommodate an increasing number of Enhanced FlashCopy images. The Enhanced FlashCopy group repository logical drive is structured as a concatenated collection of up to 16 standard logical drives. Initially, the repository logical drive has only a single standard logical drive, so its capacity is exactly that of its single standard logical drive. The Enhanced FlashCopy group repository logical drive is expanded by attaching additional standard logical drives. The capacity then becomes the sum of the capacities of all of its concatenated logical drives. Operating in this manner provides a flexible and low-overhead capacity expansion mechanism.

You can create an Enhanced FlashCopy repository logical drive automatically using the default settings or you can manually create an Enhanced FlashCopy repository logical drive by defining the settings for the repository logical drive.

After the Enhanced FlashCopy group and the associated repository logical drive are created, any write operation on the base logical drive necessitates that the affected content of the base logical drive be saved to a different location before it is overwritten by the new data. The Enhanced FlashCopy image copy-on-write mechanism performs this task, and saves the required data to the repository logical drive. An index in the Enhanced FlashCopy logical drive repository logical drive. is then updated to reflect the new location of the Enhanced FlashCopy image content for the affected logical block addressing (LBA) range. An ensuing base logical drive write to an LBA range that has already been preserved by a previous copy-on-write operation does not require an additional copy-on-write action, so the overhead of the copy is not repeated in such cases.

As an Enhanced FlashCopy group attribute, users can choose the way in which Enhanced FlashCopy logical drive repository full conditions are managed through one of these modes of operation when such conditions arises:

#### Auto-Purge Enhanced FlashCopy Images:

Automatically delete the oldest Enhanced FlashCopy images in the Enhanced FlashCopy group to free up space that can be used to satisfy the copy-on-write operation capacity needs in the Enhanced FlashCopy group repository.

#### Fail Base Writes

Fail write requests to the base logical drive that triggered the Enhanced FlashCopy logical drive repository access. In this case, the base logical drive, the Enhanced FlashCopy group and its Enhanced FlashCopy images and Enhanced FlashCopy logical drives remain in their previous states.

Attention: AIX does not support the Fail writes to base logical drive option of the Enhanced FlashCopy logical drive repository full policy. Selecting this option might cause data loss on the base logical drive. If you are using AIX, ensure that the default option Fail FlashCopy logical drive is selected. Ensure you monitor the capacity of the FlashCopy repository logical drive, as you cannot access the FlashCopy if the repository logical drive reaches full capacity.

Each Enhanced FlashCopy group can be configured independently in this regard. Expanding an Enhanced FlashCopy logical drive repository logical drive does not directly change the state of any Enhanced FlashCopy in the Enhanced FlashCopy group. In particular, Enhanced FlashCopy images in the purged state remain in that state even after a repository logical drive expansion creates additional free space for the Enhanced FlashCopy group.

#### **Repository Logical drive Names**

The Storage Manager GUI and CLI automatically provide a name (user label) every time a standard logical drive is created for use as an Enhanced FlashCopy logical drive repository logical drive member. The user label string consists of the prefix "**repos\_**" followed by a four-digit, zero-padded numeric suffix. The numeric value has a minimum value of 1, and is selected using the smallest previously-unused number among all existing logical drives. Users cannot modify or override the name that is automatically provided by IBM DS Storage Manager software.

**Note:** IBM DS Storage Manager software always uses a lower-case "**repos**\_" prefix when it creates Enhanced FlashCopy repository logical drives. IBM DS Storage Manager software tests the logical drive label for conformance to this pattern. The CLI only permits specification of an existing logical drive for use as a Enhanced FlashCopy repository logical drive member if that logical drive user label conforms to the pattern.

## Enhanced FlashCopy Logical drives

An Enhanced FlashCopy logical drive provides the host access to the data contained in an Enhanced FlashCopy image. An Enhanced FlashCopy image cannot be directly read by a hosts because the Enhanced FlashCopy image is only the changed data captured from the base logical drive. Therefore, it is necessary to create a logical drive that a host can access that is a merging of the data from the base logical drive and the Enhanced FlashCopy image.

The Enhanced FlashCopy logical drive can be designated as either read-only or read-write:

- A Read-Only Enhanced FlashCopy logical drive provides a host application with READ access to a copy of the data contained in the Enhanced FlashCopy image, but without the ability to modify the Enhanced FlashCopy image. A Read-Only Enhanced FlashCopy logical drive does not have an associated Enhanced FlashCopy logical drive repository.
- A Read-Write Enhanced FlashCopy logical drive requires an associated Enhanced FlashCopy logical drive repository to provide the host application with WRITE access to a copy of the data contained in the Enhanced FlashCopy image.

An Enhanced FlashCopy logical drive that is designated as read-write must have its own Enhanced FlashCopy logical drive repository logical drive to save any subsequent modifications made by the host application to the base logical drive without affecting the referenced Enhanced FlashCopy image.

The Enhanced FlashCopy logical drive repository logical drive is allocated from the disk pool or subsystem from which the original Enhanced FlashCopy image is allocated. All I/O write operations to the Enhanced FlashCopy image are redirected to the Enhanced FlashCopy logical drive repository that was allocated for saving data modifications. The data of the original Enhanced FlashCopy image remains unchanged.

When you delete an Enhanced FlashCopy logical drive, you can either delete the Enhanced FlashCopy logical drive repository logical drive or retain it as an unmapped logical drive.

## Relationship Between Enhanced FlashCopy Images, Enhanced FlashCopy Groups, and Enhanced FlashCopy Logical drives

The controller firmware and the IBM DS Storage Manager software management software treat Enhanced FlashCopy images, Enhanced FlashCopy groups, and Enhanced FlashCopy logical drives as distinct entities relative to Enhanced FlashCopies:

- An Enhanced FlashCopy group has an association with a standard logical drive that is the base for the Enhanced FlashCopy group.
- All Enhanced FlashCopy images within an Enhanced FlashCopy group have a direct association with that Enhanced FlashCopy group.
- Every Enhanced FlashCopy logical drive of an Enhanced FlashCopy image has a direct association with that Enhanced FlashCopy image. In addition, every Enhanced FlashCopy logical drive has a persistent relationship to the base logical drive of the Enhanced FlashCopy image for which the Enhanced FlashCopy logical drive was initially created. In other words, an Enhanced FlashCopy logical drive is effectively "bound" to its base logical drive. Even if the Enhanced FlashCopy logical drive is in a Stopped state, which detaches from a particular Enhanced FlashCopy image, the Enhanced FlashCopy logical drive retains its association with the base logical drive, and can only be restarted on an Enhanced FlashCopy image for the base logical drive.
- The Enhanced FlashCopy repository logical drives for an Enhanced FlashCopy group have an association with the Enhanced FlashCopy group.

## **Consistency Groups**

A consistency group is a container that holds several logical drives so that you can manage all of the logical drives as one entity. A consistency group allows simultaneous Enhanced FlashCopy images of several logical drives, thus ensuring consistent FlashCopies of a group of logical drives.

If you frequently want to perform the same Enhanced FlashCopy image operations on several logical drives at the same point in time, you can create a consistency group to perform those operations. A consistency group pools several logical drives together so that you can take an Enhanced FlashCopy of all the logical drives at the same point in time. Any operation that you perform on the consistency group is performed simultaneously on all of the logical drives in the consistency group. The FlashCopy image action creates synchronized Enhanced FlashCopy images of all the logical drives and is ideal for applications that span several logical drives, for example, a database application that has the logs on one logical drive and the database on another logical drive. Some of the Enhanced FlashCopy image operations that you can perform on a consistency group include creating, scheduling, and rolling back.

Each logical drive that belongs to a consistency group is referred to as a member logical drive. When you add a logical drive to a consistency group, the system automatically creates a new Enhanced FlashCopy group that corresponds to this member logical drive for the Enhanced FlashCopy images that will be created for this member logical drive. You can set up a schedule for a consistency group to automatically create an Enhanced FlashCopy image of each member logical drive in the group at a specific time in the future or on a regular basis.

You can use either of the following two methods to create a consistency group:

- Create the consistency group, and add the member logical drives in one step.
- Create the consistency group, and then add the member logical drives in a subsequent step.

Standard logical drives and thin logical drives are the only type of logical drives that can be members of a consistency group. Non standard logical drives, such as FlashCopy logical drives, cannot be used for consistency groups. The base logical drive can reside on either an subsystem or a disk pool.

#### Enhanced Remote Mirroring and Enhanced FlashCopy Consistency Groups

- A logical drive can belong to several consistency groups. You can use consistency groups for Enhanced FlashCopies and Synchronous Mirroring. Ensure to define separate and specific consistency groups for Enhanced FlashCopies and Enhanced Remote Mirroring.
- When you add a base logical drive that contains a consistency group to an Enhanced Global Mirror Group, the system automatically changes the Enhanced FlashCopy logical drive repository full policy to automatically purge the oldest Enhanced FlashCopy image and sets the auto-delete limit to the maximum allowable Enhanced FlashCopy limit for a consistency group.
- All member logical drives in a consistency group that also belong to an Enhanced Global Mirror Group, must belong to the same Enhanced Global Mirror Group.

#### **Characteristics of Consistency Groups**

A consistency group has the following characteristics:

- You can initially create consistency groups with or without member logical drives.
- For DS3500 and DCS3700, a consistency group has a maximum 32 allowable number of logical drives.
- For the DCS3700 storage subsystem with Performance Module Controllers and the DCS3860 storage subsystem, a consistency group has a maximum 64 allowable number of logical drives.
- DS3500 and DCS3700 have a maximum 512 allowable number of Enhanced FlashCopy images.
- The DCS3700 storage subsystem with Performance Module Controllers has a maximum of 2048 allowable number of Enhanced FlashCopy images.
- You can schedule an Enhanced FlashCopy image to be created for a consistency group, which causes consistent Enhanced FlashCopy images to be created for all member logical drives.
- You can perform a Rollback operation for a consistency group.

## Working With Enhanced FlashCopy Images Premium Feature

"Creating an Enhanced FlashCopy Image" on page 461

"Canceling a Pending Enhanced FlashCopy Image" on page 461

"Creating an Enhanced FlashCopy Image Schedule" on page 462

"Starting, Stopping, and Resuming an Enhanced FlashCopy Image Rollback" on page 472

"Creating an Enhanced FlashCopy Image Group"

"Deleting an Enhanced FlashCopy Image Group" on page 461

"Deleting an Enhanced FlashCopy Image Group (with options)" on page 464

"Creating an Enhanced FlashCopy Image Logical drive" on page 468

"Deleting an Enhanced FlashCopy Image Logical drive" on page 470

"Changing the Size of a Repository Logical drive" on page 470

"Restarting a Consistency Group Enhanced FlashCopy Image Logical drive" on page 469

"Creating a Consistency Group Enhanced FlashCopy logical drives" on page 465

"Deleting a Consistency Group Enhanced FlashCopy Image" on page 467

## Creating an Enhanced FlashCopy Image Group

Before you can create any Enhanced FlashCopy images you must first create an Enhanced FlashCopy group and the associated repository logical drive.

To create a new Enhanced FlashCopy group use the create enhancedFlashCopyGroup command. This command creates a new Enhanced FlashCopy group that is associated with a specific source logical drive. You also create a relationship to a repository logical drive by one of these methods:

- Use an existing repository logical drive
- · Create a new repository logical drive in an subsystem
- Create a new repository logical drive in a disk pool

To create a new repository logical drive, you identify either an existing subsystem or an existing disk pool and define the size (capacity) for the repository logical drive. You define the size by entering either a percentage of the base logical drive or a specific size that you want to reserve for the repository logical drive.

This example shows a basic command in which a new Enhanced FlashCopy group is being created and associated with an existing repository logical drive:

```
c:\...\smX\client>smcli 123.45.67.88 123.45.67.89
-c "create enhancedFlashCopyGroup userLabel=\"Data_Store_1\"
sourcelogicalDrive=\"Cont_Dev_04\"
repositorylogicalDrive=\"repos_1234\";"
```

Refer to the "Create Enhanced FlashCopy Group" on page 101 section for more information.

Additionally, optional parameters include the following:

- Setting a repository full policy that defines whether you want to fail writes from the base logical drive or to delete (purge) Enhanced FlashCopy images.
- The priority for rollback operations when you want to restore from the Enhanced FlashCopy images. You can choose the amount of system processing to devote to rollback that ranges from a minimal impact to host I/O processing to a high impact that reduces host I/O processing.

- Warning limits for when the repository logical drive is approaching full. The limit is a percentage of the capacity of the repository logical drive.
- A minimum number of Enhanced FlashCopy images that you want to delete if you have chosen to delete the Enhanced FlashCopy images as the method for managing the repository full policy.
- Enabling and defining a schedule for capturing Enhanced FlashCopy images.

## Deleting an Enhanced FlashCopy Image Group

When you delete an Enhanced FlashCopy group the system performs the following actions when an Enhanced FlashCopy group is deleted:

- Deletes all existing Enhanced FlashCopy images from the Enhanced FlashCopy group.
- Deletes the associated repository that exists for the Enhanced FlashCopy group (if selected).
- Disables all the associated Enhanced FlashCopy logical drives that exist for the deleted Enhanced FlashCopy images.

To delete the Enhanced FlashCopy group and its associated repositories, use this command:

delete enhancedFlashCopyGroup
["enhancedFlashCopyGroupName"]
deleteRepositoryMembers=TRUE;

If you want to retain the repository logical drives, set the **deleteRepositoryMembers** parameter of the above command to **FALSE**.

See the "Deleting an Enhanced FlashCopy Image Group" for more information.

## Creating an Enhanced FlashCopy Image

To create a new Enhanced FlashCopy image use the create enhancedFlashCopyImage command. This command creates a new Enhanced FlashCopy image in one or more existing Enhanced FlashCopy groups. Before you can create an Enhanced FlashCopy image, you must first have at least one Enhanced FlashCopy group into which you can place the Enhanced FlashCopy image.

The only parameters that you need to enter with this command are the names of the Enhanced FlashCopy groups.

To create an Enhanced FlashCopy Image of Enhanced FlashCopy Group to **enFCGName**, use the following command:

```
create enhancedFlashCopyImage
enhancedFlashCopyGroup="
enFCGName";
```

See the "Create Enhanced FlashCopy Image" on page 107 for more information.

## Canceling a Pending Enhanced FlashCopy Image

If you attempted to create the Enhanced FlashCopy image in either an Enhanced FlashCopy group or a consistency group, but the Enhanced FlashCopy image that was put in a Pending state, you can cancel the creation of the Enhanced FlashCopy image. The Enhanced FlashCopy image is in a Pending state due to the following concurrent conditions:

- The base logical drive for an Enhanced FlashCopy group or one or more member logical drives of a consistency group that contains this Enhanced FlashCopy image is a member of an Enhanced Global Mirror Group.
- The logical drive or logical drives are currently in a synchronizing operation.

The Enhanced FlashCopy image creation operation completes as soon as the synchronization operation is complete. To cancel the pending Enhanced FlashCopy image creation before the synchronization operation completes, use one of these commands:

- stop EnhancedFlashCopyGroup [enhancedFlashCopyGroupName]pendingEnhancedFlashCopyImageCreation for Enhanced FlashCopy groups
- stop consistencyGroup [enhancedFlashCopyConsistencyGroupName] pendingEnhancedFlashCopyImageCreation for Enhanced FlashCopy consistency groups

When you cancel a pending Enhanced FlashCopy image creation for a specific Enhanced FlashCopy group or Enhanced FlashCopy consistency group, only that group is disabled. All of the other groups stay in operation.

## Creating an Enhanced FlashCopy Image Schedule

You can schedule creating regular Enhanced FlashCopy images to enable file recovery, and scheduled backups. You can create a schedule when you initially create an Enhanced FlashCopy group or consistency group, or you can add one later to an existing Enhanced FlashCopy group or consistency group. You can create a schedule that runs daily or weekly in which you select specific days of the week (Sunday through Saturday). You can temporarily suspend scheduled Enhanced FlashCopy image creation by disabling the schedule:

- You can set up a schedule for an Enhanced FlashCopy group to automatically create an Enhanced FlashCopy image at a specific time in the future or on a regular basis.
- You can set up a schedule for an Enhanced FlashCopy consistency group to automatically create an Enhanced FlashCopy image of each member logical drive in the group at a specific time in the future or on a regular basis.

You can create an Enhanced FlashCopy image schedule in two ways:

- When you create an Enhanced FlashCopy group using the create EnhancedFlashCopyGroup command
- When you create an Enhanced FlashCopy consistency group using the create consistencyGroup command

The following table lists the parameters that you can use to set a schedule for creating an Enhanced FlashCopy image:

Option	Description
enab1eSchedu1e	Use this parameter to turn on or to turn off the ability to schedule a FlashCopy operation. To turn on FlashCopy scheduling, set this parameter to TRUE. To turn off FlashCopy scheduling, set this parameter to FALSE.

Option	Description	
schedule	Use this parameter to schedule a FlashCopy operation.	
	You can use one of these options for setting a schedule for a FlashCopy operation:	
	• immediate	
	• startDate	
	• scheduleDay	
	• startTime	
	• scheduleInterval	
	• endDate	
	• noEndDate	
	• timesPerDay	

### Scheduling Enhanced FlashCopy Images

Use the **enableSchedule** parameter and the **schedule** parameter to schedule automatic creation of Enhanced FlashCopy images. Using these parameters, you can schedule Enhanced FlashCopy images daily, weekly, or monthly (by day or by date). The **enableSchedule** parameter turns on or turns off the ability to schedule Enhanced FlashCopy images. When you enable scheduling, you use the schedule parameter to define when you want the Enhanced FlashCopy images to occur.

This list explains how to use the options for the schedule parameter:

- **immediate** As soon as you enter the command, an Enhanced FlashCopy image is created, and a copy-on-write operation begins.
- **startDate** A specific date on which you want to create an Enhanced FlashCopy image and perform a copy-on-write operation. The format for entering the date is MM:DD:YY. If you do not provide a start date, the current date is used. An example of this option is startDate=06:27:11.
- scheduleDay A day of the week on which you want to create an Enhanced FlashCopy image. You can enter these values: monday, tuesday, wednesday, thursday, friday, saturday, sundayand so on. An example of this option is scheduleDay=wednesday.
- **startTime** The time of a day that you want to create an Enhanced FlashCopy image and start performing a copy-on-write operation. The format for entering the time is HH:MM, where HH is the hour and MM is the minute past the hour. Use a 24-hour clock. For example, 2:00 in the afternoon is 14:00. An example of this option is **startTime=14:27**.
- scheduleInterval An amount of time, in minutes, that you want to have as a minimum between copy-on-write operation. You can possibly create a schedule in which you have overlapping copy-on-write operations because of the duration a copy operation. You can make sure that you have time between copy-on-write operations by using this option. The maximum value for the scheduleInterval option is 1440 minutes. An example of this option is scheduleInterval=180.
- endDate A specific date on which you want to stop creating an Enhanced FlashCopy image and end the copy-on-write operation. The format for entering the date is MM:DD:YY. An example of this option is endDate=11:26:11.
   noEndDate is a value that can be specified for the endDate parameter to indicate that the schedule will be enabled with no end date.

• **timesPerDay** – The number of times that you want the schedule to run in a day. An example of this option is **timesPerDay=4**.

If you also use the **scheduleInterval** option, the firmware chooses between the **timesPerDay** option and the **scheduleInterval** option by selecting the lowest value of the two options. The firmware calculates an integer value for the **scheduleInterval** option by dividing 1440 by the **scheduleInterval** option value that you set. For example, 1440/180 = 8. The firmware then compares the **timesPerDay** integer value with the calculated **scheduleInterval** integer value and uses the smaller value.

To remove a schedule, use the create EnhancedFlashCopyGroup command or the create consistencyGroup command and set the **enableSchedule** parameter to FALSE. The **enableSchedule** parameter set to FALSE turns off the schedule, not the creating of Enhanced FlashCopy images.

### Syntax

The following command will create a schedule for creating Enhanced FlashCopy images for **enhancedFlashCopyGroupName** Enhanced FlashCopy group two times a day, on Mondays and Fridays of the week, starting with the first image at 2:00 AM and the second image 720 minutes after the first image; the start date is May 20th, 2012 with no end date:

```
create enhancedFlashCopyGroup
userLabel="enhancedFlashCopyGroupName"
purceLogicalDrive="BaselogicaldriveName"
repositoryLogicalDrive="repos_0001"
repositoryFullPolicy= purgeEnhancedFlashCopyImages
rollbackPriority=highest
repositoryFullLimit=50
enableSchedule=TRUE
schedule startDate=05:20:12
ScheduleDay=(Monday,Friday) startTime=02:00
timesPerDay=2 scheduleInterval=720
endDate=noEndDate;
```

## Deleting an Enhanced FlashCopy Image Group (with options)

When you delete an Enhanced FlashCopy image from an Enhanced FlashCopy group, the system performs the following actions:

- Deletes the Enhanced FlashCopy image from the storage subsystem.
- Releases the repository's reserve space for reuse within the Enhanced FlashCopy group.
- Disables all the associated Enhanced FlashCopy logical drives that exist for the deleted Enhanced FlashCopy image.

For an Enhanced FlashCopy consistency group you can delete:

- A single Enhanced FlashCopy image.
- Multiple Enhanced FlashCopy images that have the same sequence number and creation timestamp.

When an Enhanced FlashCopy image(s) is deleted from an Enhanced FlashCopy consistency group, the system performs the following actions:

- Deletes the Enhanced FlashCopy images from the storage subsystem.
- Releases the repository's reserve space for reuse within the an Enhanced FlashCopy consistency group.

- Moves any member logical drive, associated with the deleted Enhanced FlashCopy image(s), to a Stopped state.
- Disables the member Enhanced FlashCopy logical drives associated with the deleted Enhanced FlashCopy image(s).

To delete the Enhanced FlashCopy image, use this command: delete EnhancedFlashCopyImage

Optionally you can choose to keep a number of Enhanced FlashCopy images with either of these parameters:

- deleteCount This parameter deletes the oldest Enhanced FlashCopy image first and continues to delete the Enhanced FlashCopy images in the order of creation, until reaching the number that you enter. If the number that you enter is greater than the number of Enhanced FlashCopy images, then all of the Enhanced FlashCopy images are deleted.
- **retainCount** This parameter keeps the most recent Enhanced FlashCopy images in the consistency group, and delete the older FlashCopy images.
- **enhancedFlashCopyImageId** The only value that be used with this parameter is **OLDEST** to delete only the oldest Enhanced FlashCopy image.

If you have an Enhanced FlashCopy logical drive associated with the to-be-deleted Enhanced FlashCopy images, you can choose to either stop the deletion operation or continue with the deletion anyway. To continue with the deletion whether the Enhanced FlashCopy images have associated Enhanced FlashCopy logical drives or not, set the **ignoreEnhancedFlashCopylogicalDrive** parameter to TRUE. To stop the deletion when the Enhanced FlashCopy image has the associated Enhanced FlashCopy logical drive, set the **ignoreEnhancedFlashCopylogicalDrive** parameter to FALSE. The default is FALSE. This parameter applies only if the Enhanced FlashCopy image is associated with an Enhanced FlashCopy logical drive.

## Creating a Consistency Group Enhanced FlashCopy logical drives

A consistency group Enhanced FlashCopy logical drive contains several Enhanced FlashCopy logical drives to provide host access to an Enhanced FlashCopy image that has been taken for each selected member logical drive at the same moment in time.

The consistency group Enhanced FlashCopy logical drive can be designated as either read-only or read-write. Read-write consistency group Enhanced FlashCopy logical drives require a repository for each member logical drive that you select to save any subsequent modifications made by the host application to the base logical drive without affecting the referenced Enhanced FlashCopy image. Each member repository is created at the same time the consistency group Enhanced FlashCopy logical drive is created.

**Note:** A read-write Enhanced FlashCopy logical drive requires an overall repository. The overall repository is created initially with one individual repository logical drive. However, the overall repository can contain multiple repository logical drives in the future for expansion purposes.

#### Prerequisites

• The Enhanced FlashCopy premium feature must be enabled on the storage subsystem.

• The consistency group must contain at least one member logical drive before you can create a consistency group Enhanced FlashCopy logical drive.

#### Guidelines

Keep these guidelines in mind when creating a consistency group Enhanced FlashCopy logical drive:

- There is a maximum of 512 Enhanced FlashCopy images for the DS3500 and DCS3700 controller subsystems.
- There is a maximum of 2048 Enhanced FlashCopy images for the DCS3700 storage subsystem with Performance Module Controllers.
- You cannot create an Enhanced FlashCopy logical drive of a failed logical drive.
- Enhanced FlashCopy logical drive repositories are fully resize-able. If you have the storage capacity you can increase the size of the Enhanced FlashCopy repository to avoid a repository full message. Conversely, if you find that the Enhanced FlashCopy logical drive repository is larger than you need, you can reduce its size to free up space that is needed by other logical drives.
- If you attempt to create an Enhanced FlashCopy logical drive for an Enhanced FlashCopy image and that Enhanced FlashCopy image is in a pending Enhanced FlashCopy image creation operation it is due to the following conditions:
  - The base logical drive that contains this Enhanced FlashCopy image is a member of an Enhanced Global Mirror Group AND
  - The base logical drive is currently in a synchronizing operation. The Enhanced FlashCopy image creation will complete as soon as the synchronization operation is complete.
- The Enhanced FlashCopy logical drive repository must have the same T10PI and Quality of Service (QoS) settings as the associated base logical drive. For example, if a base logical drive that is used for an Enhanced FlashCopy logical drive is T10PI enabled, then each member logical drive in the associated Enhanced FlashCopy logical drive repository must be T10PI enabled.

Example commands to create a consistency group Enhanced FlashCopy logical drive are as follows:

• To create consistency group Enhanced FlashCopy logical drive with read/write capability for Enhanced FlashCopy consistency group **CG\_001** using the newest Enhanced FlashCopy image for all members of the Enhanced FlashCopy consistency group:

create cgEnhancedFlashCopyLogicalDrive userLabel="cgEnhFCLDName3" cgEnhancedFlashCopyImageID="CG\_001:newest"

• To create consistency group Enhanced FlashCopy logical drive with read-only capability for Enhanced FlashCopy consistency group **CG\_001** using the Enhanced FlashCopy image with the creation order 6 for all members of the Enhanced FlashCopy consistency group:

```
create cgEnhancedFlashCopyLogicalDrive
userLabel="cgEnhFCLDName3"
cgEnhancedFlashCopyImageID="CG_001:6" readOnly
```

**Note:** To determine the creation order of the consistency group Enhanced FlashCopy image in an Enhanced FlashCopy consistency group, right-click the consistency group Enhanced FlashCopy image and select Properties in the pull-down menu. When the properties window opens, look for the value of the **creation order** field. • To create consistency group Enhanced FlashCopy logical drive with read/write capability for Enhanced FlashCopy consistency group CG\_001 using the oldest Enhanced FlashCopy image for only member logical drive 1 of subsystem 0 and member logical drive 5 of disk pool Disk\_Pool\_1. Note that the "(" and ")" for the members field indicated that new repository logical drives will be created:

```
create cgEnhancedFlashCopyLogicalDrive
userLabel="cgEnhFCLDName4"
cgEnhancedFlashCopyImageID="CG_001:oldest"
members=("1:0" "5:Disk Pool 1")
```

To create the same thing as above, using the existing repository logical drives of Repos\_0113 in subsystem 0 for member logical drive 1 and Repos\_0112 for member logical drive 5 in disk pool Disk\_pool\_1, use the following command:

**Note:** The name of the existing repositories must be in double quotation marks "" as shown

```
create cgEnhancedFlashCopyLogicalDrive
userLabel="cgEnhFCLDName4"
cgEnhancedFlashCopyImageID="CG_001:6"
members=("1":"repos_0113" "5":"repos_0112");
```

When you create an Enhanced FlashCopy image of a consistency group, the result is an Enhanced FlashCopy image of every member logical drive of the consistency group. To create a consistency group Enhanced FlashCopy image, use this command:

create cgEnhancedFlashCopyImage consistencyGroup="
consistencyGroupName"

When you enter this command you must give the consistency group Enhanced FlashCopy logical drive a unique name and identify the consistency group Enhanced FlashCopy image to be associated with the consistency group Enhanced FlashCopy logical drive.

Optionally you can choose the name of the repository logical drive associated with the Enhanced FlashCopy logical drive and set warning limit for that repository logical drive.

## Deleting a Consistency Group Enhanced FlashCopy Image

When you delete an Enhanced FlashCopy consistency group the system performs the following actions:

- Deletes all existing Enhanced FlashCopy images from the consistency group.
- Deletes all existing Enhanced FlashCopy logical drives from the consistency group.
- Deletes all the associated Enhanced FlashCopy images that exist for each member logical drive in the consistency group
- Deletes all the associated Enhanced FlashCopy logical drives that exist for each member logical drive in the consistency group
- Deletes all associated repositories that exist for each member logical drive in the consistency group (if selected).

To delete the Enhanced FlashCopy consistency group and its associated repositories, use this command:

delete consistencyGroup [consistencyGroupName] deleteRepositoryMembers=True;

To keep the repositories for later use, set the **deleteRepositoryMembers** parameter to False instead.

Optionally you can choose to keep a number of Enhanced FlashCopy images with these parameters:

- deleteCount This parameter deletes the oldest Enhanced FlashCopy image first and continues to delete the oldest Enhanced FlashCopy images until reaching the number that you enter. If the number that you enter is greater than the number of Enhanced FlashCopy images, then all of the Enhanced FlashCopy images are deleted.
- retainCount This parameter keeps the most recent Enhanced FlashCopy images in the consistency group.

If you have an Enhanced FlashCopy logical drive associated with the to-be-deleted Enhanced FlashCopy images, you can choose to either stop the deletion operation or continue with the deletion anyway. To continue with the deletion whether the Enhanced FlashCopy images have associated Enhanced FlashCopy logical drives or not, set the **ignoreEnhancedFlashCopylogical drive** parameter to TRUE. To stop the deletion when the Enhanced FlashCopy image has the associated Enhanced FlashCopy logical drive, set the **ignoreEnhancedFlashCopylogical drive** parameter to FALSE. The default is FALSE. This parameter applies only if the Enhanced FlashCopy image is associated with an Enhanced FlashCopy logical drive.

## Creating an Enhanced FlashCopy Image Logical drive

You create an Enhanced FlashCopy logical drive to provide host access to an Enhanced FlashCopy image within an Enhanced FlashCopy group. A read-write Enhanced FlashCopy logical drive has its own repository that is used to save any subsequent modifications made by the host application to the base logical drive without affecting the referenced Enhanced FlashCopy image.

You create an Enhanced FlashCopy logical drive to provide host access to an Enhanced FlashCopy image within an Enhanced FlashCopy group. A read-write Enhanced FlashCopy logical drive has its own repository that is used to save any subsequent modifications made by the host application to the base logical drive without affecting the referenced Enhanced FlashCopy image.

The Enhanced FlashCopy logical drive can be designated as either read-only or read-write.

- A read-only Enhanced FlashCopy logical drive provides a host application with READ access to a copy of the data contained in the Enhanced FlashCopy image, but without the ability to modify the Enhanced FlashCopy image. A read-only Enhanced FlashCopy logical drive does not have an associated repository.
- A read-write Enhanced FlashCopy logical drive requires an associated repository to provide the host application with WRITE access to a copy of the data contained in the Enhanced FlashCopy image.

#### Prerequisites

The Enhanced FlashCopy premium feature must be enabled on the local storage subsystems.

#### Guidelines

Keep these guidelines in mind when creating an Enhanced FlashCopy logical drive:

- You cannot create an Enhanced FlashCopy logical drive of a Failed base logical drive.
- Enhanced FlashCopy repositories are fully resize-able. If you have the storage capacity you can increase the size of the Enhanced FlashCopy repository to avoid a repository full message. Conversely, if you find that the Enhanced FlashCopy repository is larger than you need, you can reduce its size to free up space that is needed by other logical drives.
- If you create an Enhanced FlashCopy logical drive for an Enhanced FlashCopy image and that Enhanced FlashCopy image creation operation remains in a Pending state it is due to the following conditions:
  - The base logical drive that contains this Enhanced FlashCopy image is a member of an Enhanced Global Mirror Group AND
  - The base logical drive is currently in a synchronizing operation. The Enhanced FlashCopy image creation will complete as soon as the synchronization operation is complete.
- The Enhanced FlashCopy logical drive repository must have the same T10PI and Quality of Service (QoS) settings as the associated base logical drive. For example, if a base logical drive that is used for an Enhanced FlashCopy logical drive is T10PI enabled, then each repository logical drive in the overall repository must be T10PI enabled.

#### Creating an Enhanced FlashCopy Logical drive

To create a read-only Enhanced FlashCopy logical drive for the oldest enhanced FlashCopy image of the Enhanced FlashCopy group **SG\_05** using the default Enhanced FlashCopy logical drive repository drive settings, use this command: create enhancedFlashCopyLogicalDrive userLabel="enhFlashCopyLDName" enhancedFlashCopyImageID="SG\_05:oldest" readOnly;

To create a read-write Enhanced FlashCopy logical drive for the newest enhanced FlashCopy image of the Enhanced FlashCopy group **SG\_05** using the default Enhanced FlashCopy logical drive repository drive settings, use this command: create enhancedFlashCopyLogicalDrive userLabel="enhFlashCopyLDName" enhancedFlashCopyImageID="SG\_05:newest";

## Restarting a Consistency Group Enhanced FlashCopy Image Logical drive

If an Enhanced FlashCopy logical drive has entered a Stopped state, this cause the Enhanced FlashCopy logical drive to become inaccessible for read or write operations. The Enhanced FlashCopy logical drive is also detached from the Enhanced FlashCopy with which it was previously associated. The Enhanced FlashCopy logical drive still remains logically bound to the associated base logical drive.

To restart an Enhanced FlashCopy logical drive that has entered a Stopped state, use this command:

```
resume cgEnhancedFlashCopyLogicalDrive ["
enhancedFlashCopyLogicalDriveName"]
cgEnhancedFlashCopyImage="enhancedFlashCopyImageName"
```

When you run this command, you must identify the Consistency Group Enhanced FlashCopy logical drive and the consistency group Enhanced FlashCopy image used.

## Deleting an Enhanced FlashCopy Image Logical drive

If you no longer need an Enhanced FlashCopy logical drive, you can delete the logical drive using this command:

delete EnhancedFlashCopylogicalDrive

When you delete an Enhanced FlashCopy logical drive the system performs the following actions:

- Deletes all existing links with the Enhanced FlashCopy image.
- Deletes the associated repository that exists for the Enhanced FlashCopy group (if selected).

If you want to retain the repository members, set the **deleteRepositoryMembers** parameter to FALSE.

## Changing the Size of a Repository Logical drive

You can increase or decrease the size of a repository logical drive.

## Increasing the Size of a Repository Logical drive

Typically, you use increase capacity when you receive a warning that the repository is becoming full. You can increase the repository capacity by performing one of these tasks:

- · Adding one or more existing repository logical drives
- Creating a new repository logical drive using free capacity that is available on an subsystem or disk pool.

If any subsystem or disk pool does not have free capacity, you can add unconfigured capacity in the form of unused drives to an subsystem or disk pool.

#### Prerequisites

You cannot increase the storage capacity of a repository logical drive if one of these conditions exists:

- The logical drive that you want to add does not have an Optimal status.
- Any logical drive in the subsystem or disk pool that you want to add is in any state of modification.
- The subsystem or disk pool does not have any free capacity.
- The subsystem or disk pool does not have any unconfigured capacity.
- Eligible logical drives are not available.

#### Guidelines

- Each of the individual logical drives in the overall repository must have the same T10PI settings as the associated base logical drive for the storage object. For example, if a base logical drive for a storage object is T10PI enabled, then the associated overall repository must be T10PI enabled.
- Ensure that a base logical drive and each of the individual logical drives in the overall repository have the same Quality of Service (QoS) attributes, specifically for the following characteristics:

- RAID Level. A repository in a disk pool is considered to have a matching RAID Level for any base logical drive on an subsystem, regardless of the base logical drive's actual RAID Level. However, a repository on an subsystem is considered to have a matching RAID Level only if that RAID Level is identical to the RAID Level of the base logical drive.
- Drive Type. A match requires that the base logical drive and the repository logical drive reside on either an subsystem or disk pool with identical drive type attributes.
- You cannot increase or decrease the repository capacity for an Enhanced FlashCopy logical drive that is read-only because it does not have an associated repository. Only Enhanced FlashCopy logical drives that are read-write require a repository

To increase the size of an Enhanced FlashCopy logical drive repository logical drive using existing repository of **repos\_xxxx**, use this command:

```
    set EnhancedFlashCopyLogicalDrive ["EnhancedFlashCopyLogicalDriveName"]
increaseRepositoryCapacity
repositorylogicaldrives=(repos_xxxx)
```

• To increase the size of an Enhanced FlashCopy group repository logical drive using existing repository of **repos\_xxxx** use this command:

```
set EnhancedFlashCopyGroup ["EnhancedFlashCopyGroupName"]
increaseRepositoryCapacity
repositorylogicaldrives=(repos xxxx)
```

Note:

- repositoryLogicalDrive Can be either the name of existing repository like repos\_xxxx or the name of the subsystem or disk pool that the Enhanced FlashCopy group resides, and the capacity of the repository, for example (subsystemName capacity=00)
- **repositoryFullLimit=percentValue** Repository capacity at which you will receive a warning that the Enhanced FlashCopy repository logical drive is nearing full capacity.

## Decreasing the Size of a Repository Logical drive

If a repository logical drive has more capacity than is needed, you can reduce the capacity of the repository logical drive by removing member logical drives from the end of its concatenated set. A repository logical drive must always retain at least one member after such a reduction. The member logical drives that are removed in such an operation are effectively "detached" from the repository logical drive, thus reducing the repository logical drive capacity. The detached logical drives can then be reused to provide additional capacity for the same or a different repository logical drive.

You can reduce the storage capacity of an existing repository logical drive for the following storage objects:

- Enhanced FlashCopy group
- Enhanced FlashCopy logical drive
- Consistency group member logical drive
- · Consistency group member Enhanced FlashCopy logical drive

#### Prerequisites

You cannot decrease the storage capacity of the overall repository if one of these conditions exists:

- The overall repository contains only one repository member logical drive.
- If there are one or more Enhanced FlashCopy images associated with the overall repository.
- If an Enhanced FlashCopy logical drive or a consistency group member Enhanced FlashCopy logical drive is disabled.

#### Guidelines

- You can only remove repository member logical drives only in the reverse order that they were added.
- An overall repository must have at least one repository member logical drive.
- You cannot increase or decrease the repository capacity for an Enhanced FlashCopy logical drive that is read-only because it does not have an associated repository. Only Enhanced FlashCopy logical drives that are read-write require a repository logical drive.
- When you decrease capacity for an Enhanced FlashCopy logical drive or a consistency group member Enhanced FlashCopy logical drive, the system automatically migrations the logical drive to a Disabled state.

To decrease the size of an Enhanced FlashCopy logical drive repository logical drive by the number of logical drives as defined by the count parameter repository logical drive use this command:

 set EnhancedFlashCopyLogicalDrive ["EnhancedFlashCopyLogicalDriveName"] decreaseRepositoryCapacity count=numberOfLogicalDrives

To decrease the size of an Enhanced FlashCopy group repository logical drive by the number of logical drives as defined by the **count** parameter, use this command:

set EnhancedFlashCopyGroup ["EnhancedFlashCopyGroupName"]
decreaseRepositoryCapacity count=numberOfLogicalDrives

## Starting, Stopping, and Resuming an Enhanced FlashCopy Image Rollback

Enhanced FlashCopy images are useful any time you need to be able to roll back to a known good data set at a specific point in time. For example, before performing a risky operation on a logical drive, you can create an Enhanced FlashCopy image to enable "undo" capability for the entire logical drive. You can start a rollback from the following types of Enhanced FlashCopy images:

- Enhanced FlashCopy image of a base logical drive, which allows you to roll back the base logical drive associated with an Enhanced FlashCopy group to a previous state.
- Consistency group Enhanced FlashCopy image, which allows you to roll back all or select member logical drives of the consistency group to a previous state.

The Enhanced FlashCopy image rollback commands provide a way for you to manage the content that is maintained by Enhanced FlashCopy groups and Enhanced FlashCopy consistency groups. An Enhanced FlashCopy image rollback operation enables you to restore the contents of a base logical drive to a point-in-time image that was captured when an Enhanced FlashCopy image was created. The base logical drive is then immediately accessible for read/write operation using the rolled-back content after the controller firmware accepts the rollback request. The read/write operation can take place while the rolled-back content is being transferred to the base logical drive.

You can roll back data by either:

- Creating an Enhanced FlashCopy logical drive of an Enhanced FlashCopy image, which allows you to retrieve deleted files from that Enhanced FlashCopy logical drive (the base logical drive remains undisturbed).
- Restoring an Enhanced FlashCopy image to the base logical drive, which allows you to roll back the base logical drive to a previous point-in-time.

**Note:** The host will have immediate access to the new-rolled-back base logical drive, but the existing base logical drive will not allow the host read-write access once the rollback is initiated. You can create an Enhanced FlashCopy of the base logical drive just prior to initiating the rollback to preserve the pre-rollback base logical drive for recovery purposes.

Keep these guidelines in mind before you start a rollback operation:

- The rollback operation does not change the content of the Enhanced FlashCopy images that are associated with the base logical drive.
- You cannot perform the following actions when a rollback operation is in progress:
  - Delete the Enhanced FlashCopy image that is being used for the rollback.
  - Create a new Enhanced FlashCopy image for a base logical drive that is participating in a rollback operation.
  - Change the associated Enhanced FlashCopy group's Repository-Full Policy
- You cannot start a rollback operation when any of these operations are in progress in the storage subsystem:
  - Dynamic Capacity Expansion (DCE) to increase the capacity of an subsystem.
  - Dynamic Volume Expansion (DVE) to increase the capacity of a logical drive.
  - Dynamic RAID Migration (DRM) to change the RAID level of an subsystem.
  - Dynamic Segment Size (DSS) to change the segment size of a logical drive.
- You cannot start a rollback operation if the base logical drive is participating in a VolumeCopy.
- You cannot start a rollback operation if the base logical drive is a secondary logical drive in a remote mirror. However, if the base logical drive is the primary logical drive in a remote mirror, you can start a rollback operation. Additionally, you cannot perform a role reversal in a remote mirror if the primary logical drive is participating in a rollback operation.
- A rollback operation fails if any of the used capacity in the associated Enhanced FlashCopy repository logical drive has unreadable sectors.

## Setting Enhanced FlashCopy Image Rollback Priority

Rollback operations require some level of system overhead, which can reduce overall system performance. You can define the level of overhead the system devotes to a rollback operation using the create EnhancedFlashCopyGroup command or the create consistencyGroup command. These commands have a **rollbackPriority** parameter. Values for the **rollbackPriority** parameter range from highest through lowest. The highest value means that the rollback operation has priority over all other host I/O. The lowest value means that the rollback operation should be performed with minimal impact to host I/O.

## Starting an Enhanced FlashCopy Image Rollback

When you start an Enhanced FlashCopy image rollback, the contents of the base logical drive immediately start to change to the contents of the Enhanced FlashCopy image. To start an Enhanced FlashCopy image rollback use the start EnhancedFlashCopyImage rollback command or the start cgEnhancedFlashCopyImagerollback command. These commands accepts the name of one or more member logical drives. The start EnhancedFlashCopyImage rollback command works with specific EnhancedFlashCopy images. The start cgEnhancedFlashCopyImagerollback command works with specific member logical drives in the consistency group.

## Stopping an Enhanced FlashCopy Image Rollback

Attention: Possible loss of data access – Stopping an Enhanced FlashCopy image rollback can leave the base logical drive and the Enhanced FlashCopy image unusable.

Stopping an Enhanced FlashCopy image rollback leaves the base logical drive in an indeterminate state with potentially invalid or inconsistent data that is typically unusable by a host system. The base logical drive appears as failed in the storage management software. Stop an Enhanced FlashCopy image rollback only in cases where recovery options exist for restoring the data in a base logical drive. If you must stop an Enhanced FlashCopy image rollback use the stop EnhancedFlashCopyImage rollback command or the stop cgEnhancedFlashCopyImage rollback command.

You can cancel an active rollback that is in progress (actively copying data), a pending rollback (in a pending queue awaiting resources to start), or a rollback that has been paused due to an error.

After you cancel a rollback operation, you must take one of the following actions:

- Reinitialize the content of the base logical drive.
- Perform a new rollback operation to restore the base logical drive (using either the same Enhanced FlashCopy image that was used in the Cancel Rollback operation or a different Enhanced FlashCopy image to perform the new rollback operation).

**Note:** If the Enhanced FlashCopy group on which the Enhanced FlashCopy image resides has one or more Enhanced FlashCopy images that have been automatically purged, the Enhanced FlashCopy image used for the rollback operation might not be available for future rollbacks.

## **Resuming an Enhanced FlashCopy Image Rollback**

In some cases a rollback operation might pause because of a condition or action of the controller. If this occurs, you will see a status of Paused. After the controller is operating normally, you can resume an Enhanced FlashCopy image rollback by using either the resume EnhancedFlashCopyImage rollback command or the resume cgEnhancedFlashCopylogicalDrive command.

## Enhanced FlashCopy Image Rollback Status

You can see the status of an Enhanced FlashCopy image rollback operation by running either the show EnhancedFlashCopyImage command or the show

cgEnhancedFlashCopyImage command. These commands return one of these statuses during an Enhanced FlashCopy image rollback operation:

Table 22. FlashCopy Rollback Operation Status

Status	Description	
None	No Enhanced FlashCopy image rollback operations are running.	
In Progress	An Enhanced FlashCopy image rollback operation is running.	
	When an Enhanced FlashCopy image rollback operation is running the amount of the rollback operation finished is shown as a percentage and an estimate of the time remaining is also shown.	
Paused	An Enhanced FlashCopy image rollback operation was started, but has been paused due to an error condition.	
	If an Enhanced FlashCopy image rollback operation has a status of Paused, the completion percentage shows the amount of work completed, and the estimated time until completion will be <b>-1</b> .	
Pending	An Enhanced FlashCopy image rollback operation request was accepted, but the rollback operation is currently waiting for previously scheduled Enhanced FlashCopy image rollback operations to finish.	
	The percentage complete is <b>-1</b> , and the estimated time until completion is <b>-1</b> .	

## Chapter 7. About Enhanced Global Mirroring premium feature

When you set up Enhanced Global Mirroring, you create a one-to-one relationship between the logical drives in a remote-mirrored pair. The remote-mirrored pair consists of a *primary logical drive* on a local storage subsystem and a *secondary logical drive* on a storage subsystem at another site. The primary-secondary role is defined in an Enhanced Global Mirror Group. Mirrored relationships are created by adding a logical drive to the primary Enhanced Global Mirror Group and adding the corresponding logical drive to the secondary Enhanced Global Mirror Group on the remote storage subsystem. The maximum numbers of supported Enhanced Global Mirrored Pairs are listed in the following table.

**Note:** Remember, most mirror operations (such as creating relationships, synchronization settings, and mirroring role) are managed through an Enhanced Global Mirror Group.

Controller Model	Per Storage subsystem	Per Enhanced Global Mirroring Group
DCS3700 storage subsystem	32	32
DCS3700 storage subsystem with Performance Module Controllers	128	64

The Enhanced Global Mirroring premium feature is supported only on DCS3700 and the DCS3700 storage subsystem with Performance Module Controllers.

The primary logical drive is the logical drive that accepts host I/O activity and stores application data. When the mirror relationship is first created, data from the primary logical drive is copied in its entirety to the secondary logical drive. This process is known as a *full synchronization* and is directed by the controller owner of the primary logical drive. During a full synchronization, the primary logical drive remains fully accessible for all normal I/O operations.

The controller owner of the primary logical drive initiates remote writes to the secondary logical drive to keep the data on the two logical drives synchronized.

The secondary logical drive maintains a mirror (or copy) of the data on its associated primary logical drive. The controller owner of the secondary logical drive receives remote writes from the controller owner of the primary logical drive but does not accept host write requests. Hosts are able to read from the secondary logical drive, which appears as read-only.

In the event of a disaster or a catastrophic failure at the primary site, you can perform a role reversal to promote the secondary logical drive to a primary role. Hosts then are able to read from and write to the newly promoted logical drive, and business operations can continue.

When write operations are performed to the primary logical drive of an Enhanced Global mirrored pair, the modified data region of the primary logical drive is tracked. Periodically, the firmware creates a new point-in-time image of the primary logical drive and sends the changed data regions to the secondary logical

drive. When data synchronization completes, the system uses the point-in-time images on the secondary logical drive to ensure that the data is maintained in a consistent state during subsequent synchronization operations to the secondary logical drive.

The automatic cycle in an active Enhanced Global Mirroring session works as follows to maintain data at a remote site to be a point-in-time consistent copy of data at the local site.

- 1. Enhanced Global Mirror Groups consisting of logical drives are created at the local site.
- 2. Increments of consistent data are sent to the remote site.
- 3. 3. Point-in-time copy operations are performed at the remote site.

These steps are repeated according to the defined mirroring intervals.

The fundamental assumptions of the Enhanced Global Mirroring premium feature are these:

- The data link between the primary storage subsystem and the secondary storage subsystem has significantly lower bandwidth and higher latency than the host interface. For this reason, the mirrored data movement must be decoupled from the primary host data requests to minimize the performance impact to the host application. Additionally, with a restricted data link unsynchronized regions between the primary and secondary logical drives must be tracked in sufficiently small granularity to minimize data movement. Long distance data links to remote storage subsystems require additional networking equipment, which can result in varying throughput performance.
- Data on the secondary logical drive must support a site-level failover for disaster recovery. For this reason, the data on the secondary logical drive is protected during the synchronization process so that writes to the secondary logical drive do not render the logical drive data unusable. Additionally, many applications require the use of more than one logical drive, each of which must be mirrored in order to support a site-level failover. In these cases, the set of logical drives must be mirrored as a set, and the synchronization process must coordinate the data movement and synchronization intervals to create a consistent, usable data set on the secondary storage subsystem.

## **Configuring controllers for Enhanced Global Mirroring**

Enhanced Global Mirroring operations are performed between controllers that have same IDs on the primary and secondary storage subsystems; that is, controller A on the primary storage subsystem interacts only with controller A on the secondary storage subsystem and controller B on the primary storage subsystem interacts only with controller B on the secondary storage subsystem. Enhanced Global Mirroring operations are not attempted between controllers that have different IDs. The network environment does not need to provide connectivity between the controllers that have different IDs; however, it is likely to be more cost effective to use a fabric/switch configuration that does provide such connectivity, even though the firmware does not require it.

For Fibre Channel topology the Enhanced Global Mirroring premium feature requires that one host-side Fibre Channel port of each controller be dedicated to mirroring operations. An additional requirement for connectivity is that the dedicated Enhanced Global Mirroring ports must be attached to a Fibre Channel Fabric. The dedicated port is activated by the storage system administrator. Activating the dedicated port is a different operation than enabling the premium feature. After you activate the dedicated port you must enable the Enhanced Global Mirroring premium feature. Enabling the Enhanced Global Mirroring premium feature performs these actions:

- · Activates the storage subsystem for Fibre Channel mirroring
- Enables the controller firmware to support your creation of mirror groups and mirror pairs

When mirror groups or mirror pairs exist, you cannot deactivate the Enhanced Global Mirroring premium feature.

An iSCSI topology does not require a dedicated port as with Fibre Channel. The activation step is not required when setting up Enhanced Global Mirroring over iSCSI. The controller firmware maintains a list of remote storage subsystems that the iSCSI initiator attempts to establish sessions using all portals of the portal group. The first portal that successfully establishes an iSCSI connection is used for all subsequent communication with that remote storage subsystem. If communication fails, a new session is attempted using all portals in the portal group. iSCSI ports are configured at the system level on a port-by-port basis. For configuration messaging and data transfer, the communication between the controller uses these global settings:

- VLAN Both local and remote systems must have the same VLAN setting in order to communicate
- iSCSI listening port
- Jumbo Frames
- Ethernet Priority

### Enhanced Global Mirror Groups

Some applications, such as file systems and databases, distribute data storage across many logical drives. To create a failover site for such applications, all of the application data logical drives must be replicated on a storage subsystem at a remote site. The synchronization between the original data logical drives and the replicated data logical drives must be coordinated so that writes on the local storage subsystem are accurately reflected on the remote storage subsystem. The purpose of an Enhanced Global Mirror Group is to associate a set of logical drives so that mirror synchronization can be coordinated for the logical drive set, creating a consistent data set on the remote storage system.

- The local storage subsystem is the primary side of the Enhanced Global mirror group. The remote storage subsystem is the secondary side of the Enhanced Global Mirror Group.
- All logical drives added to the Enhanced Global Mirror Group on the local storage subsystem hold the primary role in the mirrored relationship.
- All logical drives added to the Enhanced Global Mirror Group on the remote storage subsystem hold the secondary role in the mirrored relationship.

During the synchronization, point-in-time images for all primary logical drives are created at the same time. Data is replicated for all logical drives in the group to the remote system within the same synchronization interval. The synchronization progress for a given resynchronization operation varies among member logical drives. However, point-in-time images of the Secondary logical drives at the completion of data synchronization are created at the same time after all members of the group have completed the synchronization. If resynchronization fails (even due to just one member), prior consistent images remain on the Secondary AMG to preserve the last known consistent data.

The properties associated with the Enhanced Global Mirror Group include the following:

- World Wide Name
- User label
- Role: primary or secondary

**Note:** The role is an Enhanced Global Mirror Group attribute, not a mirrored logical drive attribute. Role reversal affects all logical drives in the Enhanced Global Mirror Group.

- Mirrored I/O Channel Type: Fibre Channel or iSCSI
- Timestamp of the last known recovery point
- Resynchronization interval

Resynchronization interval is the amount of time between automatically sending updates of modified data from the primary storage subsystem to the secondary storage subsystem. The interval, expressed in minutes, represents the time between the starting points of sending updates from the primary to the secondary. A resynchronization interval of zero means that synchronization is manual.

- Synchronization progress: The synchronization progress is the percent complete and an estimated time to completion of the current synchronization activity if in progress.
- Alert threshold for resynchronization completion time (for example, alert if the resynchronization takes too long)
- Alert threshold for recovery point degradation (such as, age of the Secondary PiTs)
- Repository utilization warning threshold
- Member logical drive list

The following properties are shared between the local storage subsystem and remote storage subsystem:

- World Wide Name
- User label
- Mirrored I/O Channel Type: Fibre Channel or iSCSI
- Resynchronization interval
- Alert threshold for resynchronization completion time (for example, alert if the resynchronization takes too long)
- Alert threshold for recovery point degradation (such as, age of the Secondary PiTs)
- Repository utilization warning threshold

The shared properties are communicated to the remote storage subsystem when you create the Enhanced Global Mirror Group. Changes to a shared property are made first on the local storage subsystem then communicated to the remote storage subsystem. If a shared property fails change on both the local storage subsystem and the remote storage subsystem, the local storage subsystem is considered to be the master setting. When the two storage subsystems reconnect after a communication outage or reboot, the Enhanced Global Mirror Group properties are reconciled from the local storage subsystem.

If the two storage subsystems detect a dual role conflict after a communication outage is resolved, the Enhanced Global Mirror Group properties are not

reconciled until the dual-role conflict is resolved. To resolve the role conflict you must determine which storage subsystem is the primary storage subsystem. At that time, the Enhanced Global Mirror Group properties are then assumed from the side that you have defined to be the primary side.

## Mirror repository logical drives

A mirror repository logical drive is a special logical drive in the storage subsystem that is created as a resource for the controller owner of the primary logical drive in a remote mirrored pair. The controller stores mirror information on this logical drive, including information about remote writes that are not yet complete. The controller can use this information to recover from controller resets and the accidental powering down of the storage subsystems.

Mirror repository logical drives are used to manage mirror data synchronization. Mirror repository logical drives are required for both the primary and secondary logical drives in a mirror pair. This mirror repository logical drive management is different from the Enhanced Remote Mirroring premium feature where a single mirror repository is set up once and then used for all Enhanced Remote Mirroring relationships.

The mirror repository pair stores three types of data:

- Copy-on-write repository data that is used for preserving resynchronization images on the mirror-primary and recovery point images on the mirror-secondary. The copy-on-write repository is structurally similar to the repository use for Enhanced FlashCopy groups.
- A pair of delta logs that are used to track regions of the mirror-primary logical drive that are written between synchronization intervals. Even though the delta logs are only used on the primary side of the mirror, they are also allocated on the secondary side to support a role-reversal.
- A log that tracks synchronization statistics on each mirror pair.

Enhanced Global Mirroring repository logical drives are expandable repository logical drives (ERV). The minimum mirror repository size is the greater of either:

- 0.02 percent of the base logical drive capacity (regardless of role)
- 32 MB

You can expand the capacity of the mirror repository logical drives using the standard rules for ERV expansion. You cannot reduce the size of the mirror repository logical drives.

The maximum repository capacity is 101 percent of the base logical drive capacity. Primary and secondary mirror repositories are not required to be the same size. Mirror repository logical drives are independent of the associated primary mirror or secondary mirror logical drive. They can be created on separate logical drive groups with different RAID levels. Mirror repository logical drives must have compatible security and T10PI quality of service as the associated mirror logical drive. For example, if a mirror logical drive has data security enabled, the associated mirror repository must have data security enabled. Similarly, if the mirror logical drive has T10PI enabled, the associated mirror repository logical drive must have T10PI enabled.

Because of the critical nature of the data being stored, do not use RAID Level 0 as the RAID level of a mirror repository logical drive.

## **Creating an Enhanced Global Mirrored Pair**

## Prerequisites

Before you create any mirror relationships, you must create an Enhanced Global Mirror Group. The Enhanced Global Mirror Group is a logical entity that spans a local storage subsystem and a remote storage subsystem that is used for mirroring and that contains one or more mirrored pairs. The mirrored pairs consist of two logical drives: a primary logical drive on the local storage subsystem and a secondary logical drive on the remote storage subsystem. If neither the primary or secondary logical drives exist, you must create these logical drives.

Follow these guidelines when you create an Enhanced Global Mirrored Pair:

- The Enhanced Global Mirroring premium feature must be enabled and activated on the local and remote storage subsystems that you want to use for mirroring.
- The local and remote storage subsystems must be connected through a proper Fibre Channel fabric or iSCSI interface.
- The remote storage subsystem must contain a logical drive that has a capacity that is greater than or equal to the capacity of the logical drive that is to be used as the primary logical drive on the local storage subsystem.
- The RAID level of the secondary logical drive does not have to be the same as the primary logical drive.
- Ensure you know the passwords for both the local and remote storage subsystems.

Use these steps to create the an Enhanced Global Mirror Group pair.

- 1. Enable the Enhanced Global Mirroring premium feature.
- 2. Activate the Enhanced Global Mirroring premium feature.
- 3. Create the Enhanced Global Mirror Group.

## Enabling the Enhanced Global Mirroring premium feature

The first step in creating an Enhanced Global Mirrored Pair is to make sure that the Enhanced Global Mirroring premium feature is enabled on both storage subsystems. Because Enhanced Global Mirroring is a premium feature, you need a feature key file to enable the premium feature. The command for enabling the feature key file is as follows:

enable storagesubsystem feature file="filename"

In this command, the filename is the complete file path and file name of a valid feature key file. Enclose the file path and the file name in double quotation marks (""). Valid file names for feature key files end with a.keyextension.

You can use the show storageArray features command to list the premium features installed on the storage subsystem.

## Activating the Enhanced Global Mirroring premium feature

After you activate the premium feature, the secondary ports for each controller are reserved and dedicated to remote mirror use. Any host-initiated I/O operation is not accepted by the dedicated port, and any request received on the dedicated port is accepted only from another controller participating in the mirror relationship.

To activate the Enhanced Global Mirror premium feature, use this command:

activate storagesubsystem feature=enhancedGlobalMirror

The storage subsystem performs the following actions when you activate the Enhanced Global Mirroring premium feature:

- Logs out all hosts currently using the highest numbered Fibre Channel host port on the controllers.
- Reserves the highest numbered Fibre Channel host port on the controllers for mirror data transmissions.
- Rejects all host communication to this controller host port as long as the Enhanced Global Mirroring premium feature is active.

After you activate the Enhanced Global Mirroring premium feature, you must set up an Enhanced Global Mirror Group and an Enhanced Global Mirror pair.

## Creating the Enhanced Global Mirroring group

An Enhanced Global Mirror Group contains several mirrored pairs so that they can be managed as a single entity. You create an Enhanced Global Mirror Group to define the synchronization settings for all mirrored pairs within the group. Each mirrored pair in an Enhanced Global Mirror Group shares the same synchronization settings, primary role, secondary role, and write mode.

The Enhanced Global Mirror Group is associated with the local storage subsystem and remote storage subsystem that is used for mirroring. The local storage subsystem is the primary side of the Enhanced Global Mirror Group, while the remote storage subsystem is the secondary side of the Enhanced Global Mirror Group. All logical drives added to the Enhanced Global Mirror Group on the local storage subsystem hold the primary role in the mirror relationship. Subsequently, all logical drives added to the Enhanced Global Mirror Group on the remote storage subsystem hold the secondary role in the mirror relationship.

Guidelines to follow when creating an Enhanced Global Mirror Group:

- A storage subsystem has a maximum number of Enhanced Global Mirror Group. The maximum number of Enhanced Global Mirror Groups depends on your configuration
- Enhanced Global Mirror Groups are created empty and Enhanced Global Mirror pairs are added to the groups later. You can add only Enhanced Global Mirror pairs to an Enhanced Global Mirror Group. Each mirrored pair is associated with exactly one Enhanced Global Mirror Group.
- All storage subsystems with the Enhanced Global Mirroring premium feature activated are listed and can be used for mirror activities.
- Storage subsystems are displayed by their storage subsystem name. If a storage subsystem does not have a name, the storage subsystem is displayed as "Unnamed".

Ensure that you run the Create Enhanced Global Mirror Group command on the local storage subsystem. Enhanced Global Mirror Group creation is initiated from the storage subsystem that contains the logical drives that hold the primary role in the mirror relationship. You use the Create Enhanced Global Mirror Group command to specify the remote storage subsystem that contains the logical drives that will provide the secondary role in the mirror relationship.

The command has this form:

```
create enhancedGlobalMirrorGroup userLabel="enhancedGlobalMirrorGroupName"
(remoteStoragesubsystemName="storagesubsystemName" | remoteStoragesubsystemWwn="wwID")
interfaceType=(FC | iSCSI)
[remotePassword="password"
syncInterval=integer (minutes | hours | days)
warningSyncThreshold=integer (minutes | hours | days)
warningRecoveryThreshold=integer (minutes | hours | days)
warningThresholdPercent=percentValue
autoResync=(TRUE | FALSE)]
```

When you run this command you create a new, empty Enhanced Global Mirror Group on both the local storage subsystem and the remote storage subsystem.

This example shows how to create an Enhanced Global Mirror Group on a Windows command prompt:

```
c:\...\smX\client>smcli 123.45.67.88 123.45.67.89
-c "create enhancedGlobalMirrorGroup userLabel="\EngDevData\"
remoteStoragesubsystemName="\Eng_Backup\"
interfaceType=iSCSI
remotePassword="\xxxxx\"
syncInterval=8 hours
warningSyncThreshold=1 hours
warningRecoveryThreshold=2 hours
warningThresholdPercent=80
autoResync=TRUE]
```

The command in this example creates a new Enhanced Global Mirror Group with a repository logical drive on a remote storage subsystem that already has the name "Eng\_Backup." The interface between the local storage subsystem and the remote storage subsystem is iSCSI. The remote storage subsystem is password protected and the password must be used to create the Enhanced Global Mirror Group. The synchronization between the local storage subsystem and the remote storage subsystem automatically takes place once every eight hours. If the synchronization cannot be completed successfully the administrator receives a message one hour after the synchronization did not work. When the repository logical drive has reached 80 percent of capacity, a warning alert is sent to the administrator.

This example shows how to use the command in a script file:

```
create enhancedGlobalMirrorGroup userLabel="EngDevData"
remoteStoragesubsystemName="Eng_Backup"
interfaceType=iSCSI
remotePassword="xxxxx"
syncInterval=8 hours
warningSyncThreshold=1 hours
warningRecoveryThreshold=2 hours
warningThresholdPercent=80
autoResync=TRUE]
```

After you create the Enhanced Global Mirror Group, you can create the Enhanced Global Mirrored Pair to start performing Enhanced Remote Mirroring operations.

#### Creating the Enhanced Global Mirroring pair

After you have created the Enhanced Global Mirror Group you can create the mirrored pairs that you want in the Enhanced Global Mirror Group. Creating a mirrored pair involves these steps:

- 1. Determining if you have a candidate secondary mirror logical drive on the remote storage subsystem.
- **2**. Adding a primary logical drive that allows you to mirror to the Enhanced Global Mirror Group on the local storage subsystem.

**3**. Enabling the relationship between the primary logical drive on the local storage subsystem and the repository logical drive on the remote storage subsystem.

When you add a member logical drive to the to the primary Enhanced Global Mirror Group on the local storage subsystem, the firmware creates a place holder object in the secondary Enhanced Global Mirror Group on the remote storage subsystem. The placeholder object represents the associated mirror logical drive. When you add a member logical drive to the secondary Enhanced Global Mirror Group on the remote storage subsystem, the member logical drive replaces the placeholder object and resolves the mirror configuration.

After you have created the Enhanced Global Mirrored Pair you will actually have three logical drives in the relationship:

- Primary logical drive on the local storage subsystem that holds the original data
- Secondary logical drive on the remote storage subsystem that holds the duplicate data
- Repository logical drive on the local storage subsystem that holds the data that was changed on the primary logical drive

If any member logical drives of the Enhanced Global Mirror Group are actively synchronizing data from a periodic synchronization interval when the new member logical drive is added, the synchronization process is halted. Existing recovery points from a past completed synchronization process, if any, are preserved. After the new Enhanced Global Mirror Group mirror-pair has been initialized, a new consistent set of recovery points are taken of all member logical drives of the Enhanced Global Mirror Group.

Guidelines to follow when creating the Enhanced Global Mirror pairs:

- Primary logical drives and secondary mirror repository logical drives do not need to be the same size.
- Mirror repository logical drives are independent of the associated primary logical drive and secondary logical drive so that they can be created in separate arrays with different RAID levels.
- Mirror repository logical drives must have the same T10PI and Quality of Service (QoS) settings as the associated logical drives. For example, if a mirrored pair is T10PI enabled, the associated mirror repository logical drive must be T10PI enabled.
- Member logical drives on both the local storage subsystem and the remote storage subsystem must be standard RAID logical drives. They cannot be Enhanced FlashCopy logical drives, views, or repository logical drives.

#### Determining if You Have a Secondary Mirror logical drive.

To determine if you have secondary mirror logical drive candidate that is appropriate to your needs on the remote storage subsystem, run this command on the remote storage subsystem:

show alllogicalDrives summary

This command returns information about the number of logical drives on the storage subsystem, the names of the logical drives, the capacity, RAID level, and other information. If you run the command without the summary parameter, you receive several pages of detailed information. However, running with the summary

parameter provides you with enough information to determine which of the logical drives you might be able to use for a repository logical drive.

#### Adding a Primary logical drive to the Enhanced Global Mirror Group

**Note:** You cannot add a primary logical drive to the Enhanced Global Mirror Group if the Enhanced Global Mirror Group contains orphaned mirror-pair members.

A mirrored relationship between the primary logical drive on the local storage subsystem and the secondary logical drive on the remote storage subsystem is established by adding a member logical drive to the Enhanced Global Mirror Group on the local storage subsystem. To add a member logical drive to the Enhanced Global Mirror Group on the local storage subsystem run this command on the local storage subsystem:

```
add logicalDrive="logicalDriveName"
enhancedGlobalMirrorGroup="enhancedGlobalMirrorGroupName"
remotePassword="password"
(repositorylogicalDrive="repos_xxxx" |
repositorylogicalDrive=(logicalDriveGroupName [capacity=capacityValue])
repositorylogicalDrive=(diskPoolName [capacity=capacityValue]))
```

When you run this command, you must perform these actions:

- Identify the logical drive on the local storage subsystem that you want to mirror to a repository logical drive on the remote storage subsystem.
- Identify the Enhanced Global Mirror Group in which you want to place the logical drive that you want to mirror.
- Identify an existing repository logical drive or create a new repository logical drive.

If an unused repository logical drive already exists on the remote storage subsystem, you can reuse that repository logical drive. Otherwise you must create the repository logical drive. With this command you have the option of creating a repository logical drive in either a logical drive group or a disk pool. You identify in which array or a disk pool you want to place the repository logical drive and the size of the repository logical drive. The storage management software and firmware then creates a repository logical drive with a name "repos\_xxxx" where "xxxx" is a numerical identifier. After the repository logical drive is created, you cannot rename it.

Establishing the Link Between the primary logical drive and the secondary logical drive

After you have identified a qualified secondary logical drive on the remote storage subsystem and added the primary logical drive on the local storage subsystem, you must link the two logical drives. Use this command to link the two logical drives:

establish enhancedGlobalMirror logicalDrive="secondarylogicalDriveName"
enhancedGlobalMirrorGroup="enhancedGlobalMirrorGroupName"
primarylogicalDrive="primarylogicalDriveName"

- secondarylogicalDriveName is the member logical drive on the remote storage subsystem.
- enhancedGlobalMirrorGroupNameNameis the member Enhanced Global Mirror Group that contains the mirror-pair.
• primarylogicalDriveName is the member logical drive on the local storage subsystem.

After you run this command, the Enhanced Global Mirrored Pair is linked and the initial mirror starts. During the initial synchronization, all of the data is copied from the primary logical drive to the secondary logical drive. During the copy operation, the primary logical drive is accessible by all hosts for write operations, but the secondary logical drive is not ready to use for recovery. During the initial synchronization, the repository logical drive delta log tracks write requests to the primary logical drive. At the conclusion of the initial synchronization between the primary logical drive and the secondary logical drive. Because the initial synchronization is a complete copy between the primary logical drive and the secondary logical drive and the secondary logical drive. Enhanced Global Mirroring operations can now be performed as you defined when you created the Enhanced Global Mirror Group.

If you want to make any changes to the synchronization settings between the local, primary logical drive and the remote secondary logical drive, use the set enhancedGlobalMirrorGroup command.

# **Changing Enhanced Global Mirroring settings**

The set enhancedGlobalMirrorGroupcommand enables you to change the property settings for an Enhanced Global Mirrored Pair. Use this command to change these property settings:

- Synchronization interval The time between automatically sending updates of modified data from the local storage subsystem to the remote storage subsystem.
- Synchronization warning threshold Time to wait until a warning is triggered when the synchronization of logical drives takes longer than the defined time.
- Recovery warning threshold Time to wait until a warning is triggered when the automatic data update for the point-in-time image on the remote storage subsystem is older than a defined time.
- Warning threshold percent Percentage of capacity of the mirror repository logical drive at which a warning is sent if the mirror repository logical drive is nearing full.
- User label A new name for the Enhanced Global Mirror Group.
- Automatic resynchronization A setting to enable or disable automatic resynchronization between the primary logical drive and the secondary logical drive.
- Logical drive A repository logical drive for which you want to increase the capacity.
- Repository logical drive An unused repository logical drive that you want to add to another repository logical drive to increase capacity.
- Role A setting to change the role, primary or secondary, of the logical drives in an Enhanced Global Mirror Group.
- Force The setting to force a role change on the local storage subsystem if the link between the local storage subsystem and the remote storage subsystem is not available.
- No synchronization A setting to force synchronization before a role change.

This example shows how to use the set enhancedGlobalMirrorGroup command to increase the capacity of a repository logical drive:

c:\...\smX\client>smcli 123.45.67.88 123.45.67.89
-c "set enhancedGlobalMirrorGroup \"amg\_001\"
logicalDrive="\repos\_006\" increaseRepositoryCapacity
repositorylogicalDrives=("\repos\_0020\" "\repos\_0021\");"

This example shows how to use the command in a script file:

set enhancedGlobalMirrorGroup "amg\_001"
logicalDrive="repos\_006" increaseRepositoryCapacity
repositorylogicalDrives=("repos 0020" "repos 0021");

#### Suspending and resuming the Enhanced Global Mirror Group

Use the suspend enhancedGlobalMirrorGroup command to stop data transfer between all of the primary logical drives and all of the secondary logical drives in an Enhanced Global Mirror Group without disabling the Enhanced Global Mirroring relationships. Suspending the Enhanced Global Mirror relationship lets you control when the data on the primary logical drive and data on the secondary logical drive are synchronized. Suspending the Enhanced Global Mirroring helps to reduce any performance impact to the host application that might occur while any changed data on the primary logical drive is copied to the secondary logical drive. Suspending the Enhanced Global Mirroring is particularly useful when you want to run a backup of the data on the secondary logical drive.

When the Enhanced Global Mirroring relationship is in a suspended state, the primary logical drive does not make any attempt to contact the secondary logical drive. Any writes to the primary logical drive are persistently logged in the Enhanced Global Mirroring repository logical drives. After the Enhanced Global Mirror relationship resumes, any data that is written to the primary logical drive is automatically written to the secondary logical drive. Only the modified data blocks on the primary logical drive are written to the secondary logical drive. Full synchronization is not required. This example shows the suspend enhancedGlobalMirrorGroup command:

c:\...\smX\client>smcli 123.45.67.88 123.45.67.89
-c "suspend enhancedGlobalMirrorGroup ["amg\_001"];"

The Enhanced Global Mirror Group name is amg\_001. You must enclose the name in square brackets. In this example the double quotation marks are optional. The double quotation marks must be used if the Enhanced Global Mirror Group name has special characters, such as a colon as used in this example "amg:001". The double quotation marks also must be used if the Enhanced Global Mirror Group name is only a number, as in this example "001". Using double quotation marks when not needed does not prevent the command from running successfully. If you have any doubts about the Enhanced Global Mirror Group name, use the double quotation marks inside the square brackets.

This example shows how to use the command in a script file: suspend enhancedGlobalMirrorGroup ["amg 001"];

The mirror relationship remains suspended until you use the resume enhancedGlobalMirrorGroup command to restart synchronization activities. This command restarts data transfers between a primary logical drive and a secondary logical drive in a mirror relationship after the mirror has been suspended or unsynchronized. This example shows the resume enhancedGlobalMirrorGroup command: c:\...\smX\client>smcli 123.45.67.88 123.45.67.89 -c "resume enhancedGlobalMirrorGroup ["amg\_001"];"

This example shows how to use the command in a script file: resume enhancedGlobalMirrorGroup ["amg\_001"];

# Manually resynchronizing logical drives in an Enhanced Global Mirror Group

Manually resynchronizing the logical drives in an Enhanced Global Mirror Group immediately resynchronizes all of the mirror relationships within the Enhanced Global Mirror Group. You cannot perform this operation if one of these conditions exists:

- The Enhanced Global Mirror Group has failed because any dependent component of the mirror configuration is in a Failed state.
- The Enhanced Global Mirror Group is in a Suspended state.

The manual resynchronization request must be issued to the primary Enhanced Global Mirror Group.

You can run a manual resynchronization on an Enhanced Global Mirror Group that is configured with a periodic synchronization interval. A manual resynchronization does not affect the periodic resynchronization schedule. If you manually resynchronize the Enhanced Global Mirror Group when it is performing an initial synchronization, the manual resynchronization command is rejected.

If you run a manual resynchronization when an Enhanced Global Mirror Group is in the midst of a periodic synchronization operation, the current process is stopped and the point in time synchronization points on the primary are discarded. The delta log tracking new host writes and the delta log being used for the synchronization process are merged. New point in time synchronization points are created and a new resynchronization is started.

To manually run a resynchronization, run this command: start enhancedGlobalMirrorGroup synchronize

# **Changing Enhanced Global Mirroring roles**

You can change the mirroring roles of an Enhanced Global Mirror Group, promoting the secondary storage subsystem to the primary storage subsystem. Each storage subsystem changes roles. First the original primary becomes the secondary, and then the original secondary becomes the primary. As soon as the role change completes, all Enhanced Global Mirror Group member logical drives on the new-primary storage subsystem become fully accessible for host I/O operations. Follow these guidelines when changing roles:

- You can run the command to change roles from either primary or secondary storage subsystem.
- All of the mirror pairs of the Enhanced Global Mirror Group must have a valid recovery point before performing a role reversal.
- The Enhanced Global Mirror Group must be in a non-failed state and must not already be in role changing operation.
- You cannot change roles if the Enhanced Global Mirror Group contains incomplete or orphaned member logical drives.

When the communication between the Enhanced Global Mirroring storage subsystems is operating normally, role reversal is coordinated between the local and remote storage subsystems. During role reversal these two events occur:

- The original primary logical drives are protected from new write requests just as if they were secondary logical drives.
- A resynchronization process from the original primary logical drives to the original secondary logical drives starts.

The resynchronization operation completes after all mirror-pairs of the Enhanced Global Mirror Group are fully synchronized. If a scheduled synchronization operation is in progress when a controller receives the role change request, the resynchronization operation is stopped and restarted. This results in all regions flagged in both the host-write tracking delta log and the synchronization delta log being copied to the secondary logical drives.

After the synchronization finishes and the roles are changed these actions take place:

- Recovery points are created on the original primary storage subsystem that is now the secondary storage subsystem.
- The logical drives on the new primary storage subsystem change state to Optimal and operate normally: servicing read and write requests, tracking write requests, and periodically synchronizing to the secondary logical drives.

To change roles, use this form of the set enhancedGlobalMirrorGroup command set enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"] role=(primary | secondary)

The role parameter enables you to define the role you want from either the primary storage subsystem or the secondary storage subsystem. For example, if you are on the primary storage subsystem and you want to change roles, use the secondary parameter. The command would look like this:

set enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"] role=secondary

As soon as you run this command, the storage subsystems start changing roles.

Two optional parameters support changing the roles of the primary storage subsystem and the secondary storage subsystem:

- force
- nosync

The force parameter forces a role change if the communications link between the storage subsystems is down and promotion or demotion on the local side results in a dual-primary condition or a dual-secondary condition. To force a role reversal, set this parameter to TRUE. The default value is FALSE.

The nosync parameter defines whether to perform an initial synchronization between the primary storage subsystem and the secondary storage subsystem before a role reversal operation is performed. To perform an initial synchronization, set this parameter to TRUE. The default value is FALSE.

If the role change is interrupted because of a communication failure between the storage subsystems, the mirror roles can possibly end as two secondary roles. This role conflict does not compromise the data synchronization state.

# Canceling a pending Enhanced Global Mirror Group role change

You can cancel a pending role change by running this command: stop enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"] rolechange

This command restores the Enhanced Global Mirror Group to the normal operating state including write requests to the logical drives on the primary logical drive. The synchronization process that was started as part of the role change operation is allowed to finish. The next periodic resynchronization is scheduled based on the most recently completed periodic synchronization and the current Enhanced Global Mirror Group synchronization settings.

# **Resolving role conflicts**

Because you can force a change in the role of the storage subsystems in an Enhanced Global Mirror Group, sometimes this might result in a condition where you have two primary or two secondary roles at the same time. This usually occurs when communications between the storage subsystems cannot take place. For example, the original primary side might be operational, but cannot be reached because of a link failure. In this case, a forced promotion of the secondary to become a primary results in both sides being able to receive new data write requests after the most recent resynchronization. Later, the original primary site might be reactivated or connectivity reestablished resulting in both local and remote storage subsystems viewing themselves as the primary.

If you have a role conflict, your Enhanced Global Mirror Group might have two primary logical drives or two secondary logical drives in an Enhanced Global Mirror Group. If you have two primary logical drives, both logical drives can accept host writes, but you do not have any mirror logical drives. If you have two secondary logical drives, neither logical drive can accept host writes. In either case, the user no longer has a valid Enhanced Global Mirror Group.

If such a role conflict occurs, you need to migrate the mirrored pair back to a state in which one side is clearly recognized as primary and the other as secondary. To migrate back to a valid Enhanced Global Mirror Group, use this command. set enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"] role=(primary | secondary)

Run this command on one of the storage subsystems in the Enhanced Global Mirror Group and define the role of that storage subsystem as you need to meet your mirroring requirements. Define the storage subsystem as either the storage subsystem with the primary logical drives or the storage subsystem with the secondary logical drives.

## Removing logical drives from the Enhanced Global Mirror group

When you have an Enhanced Global Mirror Group you have three logical drives that you need to manage:

- Primary on the local storage subsystem
- Secondary on the remote storage subsystem
- Repository on both storage subsystems

Use the remove logicalDrive enhancedGlobalMirrorGroup command to remove the link between a primary logical drive and a secondary logical drive.

Removing a logical drive from an Enhanced Global Mirror Group disassociates the logical drive from the Enhanced Global Mirror Group and also disassociates the corresponding logical drive on the remote storage subsystem from the Enhanced Global Mirror Group. Removing a logical drive disassociates the mirror repositories from the affected logical drives on both of the local and remote storage subsystems. You can then either delete the repository logical drive or you can keep the repository logical drive for later use as a repository logical drive for a different configuration. If you choose to keep the repository logical drive, you can use it as the principal repository logical drive for another Enhanced Global Mirrored Pair or to increase the capacity of another repository logical drive.

The logical drive on the primary Enhanced Global Mirror Group is removed first to halt any in-progress synchronization I/O operations. If the remote storage subsystem is not accessible because of an inter-storage subsystem communication problem, you can force the removal operation so that only the logical drive on the local Enhanced Global Mirror Group is removed. The corresponding logical drive on the remote storage subsystem remains in the remote Enhanced Global Mirror Group.

Removing the link between a primary logical drive and a secondary logical drive does not affect any of the existing data on either logical drive. The link between the logical drives is removed, but the primary logical drive still continues normal I/O operations. Later, you can establish the mirror relationship between the two logical drives and resume normal mirror operations. You can remove the mirror relationship for one or several remote mirrored pairs with this command.

This example shows the remove logicalDrive enhancedGlobalMirrorGroup command:

```
c:\...\smX\client>smcli 123.45.67.88 123.45.67.89
-c "remove logicalDrive [\"Jan_04_Account\"] enhancedGlobalMirrorGroup=\"amg_001\"
deleteRepositoryMembers=TRUE;"
```

The command in this example removes a logical drive with the name Jan\_04\_Accountfrom the Enhanced Global Mirror Group with the name amg\_001. Because the deleteRepositoryMembers is set to TRUE, the repository logical drive is deleted. You must enclose the logical drive name in double quotation marks inside square brackets. Also, you must enclose the Enhanced Global Mirror Group name in double quotation marks.

This example shows how to use the command in a script file: remove logicalDrive ["Jan\_04\_Account"] enhancedGlobalMirrorGroup="amg\_001" deleteRepositoryMembers=TRUE;

You must run this command on the local storage subsystem to remove the primary logical drive.

If the logical drive is not successfully removed from both sides of the Enhanced Global Mirror Group, the mirror logical drive that was not removed becomes an orphan. Orphans are detected when communications between the controller on the local storage subsystem and the corresponding controller on the remote storage subsystem is restored. At that time the two sides of the mirror configuration reconcile mirror parameters. Some synchronization operations function normally when a Enhanced Global Mirror Group contains orphaned members, but some configuration operations, such as role change, cannot be performed when the Enhanced Global Mirror Group contains orphaned members. To avoid this situation, remove the incomplete mirror logical drives by running this command:

# **Deleting an Enhanced Global Mirror Group**

If you need to remove an Enhanced Global Mirror Group for any reason, you can remove it using this command:

delete enhancedGlobalMirrorGroup ["enhancedGlobalMirrorGroupName"]

Before you can delete an Enhanced Global Mirror Group you must remove the primary logical drive, the secondary logical drive, and the repository logical drive from the Enhanced Global Mirror Group. The Enhanced Global Mirror Group must be completely empty before you can delete it. You can remove these logical drives using this command:

remove logicalDrive=["logicalDrivename"]
enhancedGlobalMirrorGroup="enhancedGlobalMirrorGroupName"

Removing these logical drives from the Enhanced Global Mirror Group does not delete the logical drives. The data is not lost.

Deleting the Enhanced Global Mirror Group first removes the remote Enhanced Global Mirror Group then removes the local Enhanced Global Mirror Group. You can run the delete enhancedGlobalMirrorGroup command from either the primary or secondary Enhanced Global Mirror Group.

# Chapter 8. About Enhanced Remote Mirroring premium feature

The Enhanced Remote Mirroring premium feature provides for online, real-time replication of data between storage subsystems over a remote distance. In the event of a disaster or a catastrophic failure on one storage subsystem, you can promote the second storage subsystem to take over responsibility for computing services. Enhanced Remote Mirroring is designed for extended storage environments in which the storage subsystems that are used for Enhanced Remote Mirroring are maintained at separate sites. Logical drives on one storage subsystem are mirrored to logical drives on another storage subsystem across a fabric SAN. Data transfers can be synchronous or asynchronous. You choose the method when you set up the remote-mirror pair. The data transfers occur at Fibre Channel speeds to maintain data on the different storage subsystems. Because Enhanced Remote Mirroring is storage based, it does not require any server overhead or application overhead.

You can use Enhanced Remote Mirroring for these functions:

- **Disaster recovery** Enhanced Remote Mirroring lets you replicate data from one site to another site, which provides an exact mirror duplicate at the remote (secondary) site. If the primary site fails, you can use mirrored data at the remote site for failover and recovery. You can then shift storage operations to the remote site for continued operation of all of the services that are usually provided by the primary site.
- **Data vaulting and data availability** Enhanced Remote Mirroring lets you send data off site where it can be protected. You can then use the off site copy for testing or to act as a source for a full backup to avoid interrupting operations at the primary site.
- **Two-way data protection** Enhanced Remote Mirroring provides the ability to have two storage subsystems back up each other by mirroring critical logical drives on each storage subsystem to logical drives on the other storage subsystem. This action lets each storage subsystem recover data from the other storage subsystem in the event of any service interruptions.

# How Enhanced Remote Mirroring Works

When you create a remote-mirror pair, the remote-mirror pair consists of a *primary logical drive* on a local storage subsystem and a *secondary logical drive* on a storage subsystem at another site. A standard logical drive might only be included in one mirrored logical drive pair.

**Note:** Enhanced Remote Mirroring is not supported in DS3000 storage subsystem configurations.

Controller Model	Maximum Number of Defined Mirrors
DS3500 series	8 (16 with controller firmware 7.77 and later)
DCS3700	16

Table 23. Maximum Number of Defined Mirrors

Controller Model	Maximum Number of Defined Mirrors
DCS3700 storage subsystem with Performance Module Controllers	16
DS3950	64
DS4300	32
DS4400	64
DS4500	64
DS4700	64
DS4800	128
DS5020	64
DS5100	128
DS5300	128

Table 23. Maximum Number of Defined Mirrors (continued)

The primary logical drive is the logical drive that accepts host I/O activity and stores application data. When the mirror relationship is first created, data from the primary logical drive is copied in its entirety to the secondary logical drive. This process is known as a *full synchronization* and is directed by the controller owner of the primary logical drive. During a full synchronization, the primary logical drive remains fully accessible for all normal I/O operations.

The controller owner of the primary logical drive initiates remote writes to the secondary logical drive to keep the data on the two logical drives synchronized.

The secondary logical drive maintains a mirror (or copy) of the data on its associated primary logical drive. The controller owner of the secondary logical drive receives remote writes from the controller owner of the primary logical drive but will not accept host write requests. Hosts are able to read from the secondary logical drive, which appears as read-only.

In the event of a disaster or a catastrophic failure at the primary site, you can perform a role reversal to promote the secondary logical drive to a primary role. Hosts then are able to read from and write to the newly promoted logical drive, and business operations can continue.

## Mirror Repository Logical Drives

A mirror repository logical drive is a special logical drive in the storage subsystem that is created as a resource for the controller owner of the primary logical drive in a remote-mirror pair. The controller stores mirroring information on this logical drive, including information about remote writes that are not yet complete. The controller can use this information to recover from controller resets and the accidental powering down of the storage subsystems.

When you activate the Enhanced Remote Mirroring premium feature on the storage subsystem, you create two mirror repository logical drives, one for each controller in the storage subsystem. An individual mirror repository logical drive is not needed for each remote mirror.

When you create the mirror repository logical drives, you specify the location of the logical drives. You can either use existing free capacity, or you can create an subsystem for the logical drives from unconfigured capacity and then specify the RAID level.

Because of the critical nature of the data being stored, do not use RAID 0 as the RAID level of mirror repository logical drives. The required size of each logical drive is 128 MB, or 256 MB total for both mirror repository logical drives of a dual-controller storage subsystem. In previous versions of the Enhanced Remote Mirroring feature, the mirror repository logical drives required less disk storage space and needed to be upgraded to use the maximum amount of mirror relationships.

#### **Mirror Relationships**

Before you create a mirror relationship, you must enable the Enhanced Remote Mirroring premium feature on both the primary storage subsystem and the secondary storage subsystem. You must also create a secondary logical drive on the secondary site if one does not already exist. The secondary logical drive must be a standard logical drive of equal or greater capacity than the associated primary logical drive.

When secondary logical drives are available, you can establish a mirror relationship in the storage management software by identifying the primary logical drive and the storage subsystem that contains the secondary logical drive.

When you first create the mirror relationship, a full synchronization automatically occurs, with data from the primary logical drive copied in its entirety to the secondary logical drive.

## **Data Replication**

The controllers manage data replication between the primary logical drive and the secondary logical drive. This process is transparent to host machines and applications. This section describes how data is replicated between the storage subsystems that are participating in Enhanced Remote Mirroring. This section also describes the actions taken by the controller owner of the primary logical drive if a link interruption occurs between storage subsystems.

#### Write Modes

When the controller owner of the primary logical drive receives a write request from a host, the controller first logs information about the write to a mirror repository logical drive, and then writes the data to the primary logical drive. The controller then initiates a remote write operation to copy the affected data blocks to the secondary logical drive at the secondary storage subsystem.

The Enhanced Remote Mirroring feature provides two write mode options that affect when the I/O completion indication is sent back to the host: Synchronous and Asynchronous.

#### **Metro Mirror**

Metro mirror provides the highest level security for full data recovery from the secondary storage subsystem in the event of a disaster. Metro mirror does reduce host I/O performance. When this write mode is selected, host write requests are written to the primary logical drive and then copied to the secondary logical drive.

After the host write request has been written to the primary logical drive and the data has been successfully copied to the secondary logical drive, the controller removes the log record on the mirror repository logical drive. The controller then sends an I/O completion indication back to the host system. Metro mirror is selected as the default value and is the recommended write mode.

#### **Global Copy**

Global copy offers faster host I/O performance but does not guarantee that a copy operation has successfully completed before processing the next write request. When you use Global copy, host write requests are written to the primary logical drive. The controller then sends an "I/O complete" indication back to the host system, without acknowledging that the data has been successfully copied to the secondary (remote) storage subsystem.

When using Global copy, write requests are not guaranteed to be completed in the same order on the secondary logical drive as they are on the primary logical drive. If the order of write requests is not retained, data on the secondary logical drive might become inconsistent with the data on the primary logical drive. This event could jeopardize any attempt to recover data if a disaster occurs on the primary storage subsystem.

#### **Global Mirror**

When multiple mirror relationships exist on a single storage subsystem and have been configured to use Global copy and to preserve consistent write order, they are considered to be an interdependent group that is in the *Global mirror*. The data on the secondary, remote storage subsystem cannot be considered fully synchronized until all of the remote mirrors that are in the Global mirror are synchronized.

If one mirror relationship in the group becomes unsynchronized, all of the mirror relationships in the group become unsynchronized. Any write activity to the remote, secondary storage subsystems is prevented to protect the consistency of the remote data set.

# Link Interruptions or Secondary Logical Drive Errors

When processing write requests, the primary controller might be able to write to the primary logical drive, but a link interruption might prevent communication with the remote (secondary) controller.

In this case, the remote write operation cannot be completed to the secondary logical drive, and the primary logical drive and the secondary logical drive are no longer correctly mirrored. The primary controller migrations the mirrored pair into an Unsynchronized state and sends an I/O completion to the primary host. The primary host can continue to write to the primary logical drive, but remote writes do not take place.

When communication is restored between the controller owner of the primary logical drive and the controller owner of the secondary logical drive, a resynchronization takes place. This resynchronization happens automatically, or it must be started manually, depending on which write mode you chose when setting up the mirror relationship. During the resynchronization, only the blocks of data that have changed on the primary logical drive during the link interruption are copied to the secondary logical drive. After the resynchronization starts, the mirrored pair migrations from an Unsynchronized status to a Synchronization in Progress status.

The primary controller also marks the mirrored pair as unsynchronized when a logical drive error on the secondary side prevents the remote write from completing. For example, an offline secondary logical drive or a failed secondary logical drive can cause the remote mirror to become unsynchronized. When the logical drive error is corrected (the secondary logical drive is placed online or recovered to an Optimal status), then synchronization is required. The mirrored pair then migrations to a Synchronization in Progress status.

## Resynchronization

Data replication between the primary logical drive and the secondary logical drive in a mirror relationship is managed by the controllers and is transparent to host machines and applications. When the controller owner of the primary logical drive receives a write request from a host, the controller first logs information about the write to a mirror repository logical drive. The controller then writes the data to the primary logical drive. The controller then initiates a write operation to copy the affected data to the secondary logical drive on the remote storage subsystem.

If a link interruption or a logical drive error prevents communication with the secondary storage subsystem, the controller owner of the primary logical drive migrations the mirrored pair into an Unsynchronized status. The controller owner then sends an I/O completion to the host sending the write request. The host can continue to issue write requests to the primary logical drive, but remote writes to the secondary logical drive do not take place.

When connectivity is restored between the controller owner of the primary logical drive and the controller owner of the secondary logical drive, the logical drives must be resynchronized by copying the blocks of data that changed during the interruption to the secondary logical drive. Only the blocks of data that have changed on the primary logical drive during the link interruption are copied to the secondary logical drive.

Attention: Possible loss of data access – Any communication disruptions between the primary storage subsystem and the secondary storage subsystem while resynchronization is underway could result in a mix of new data and old data on the secondary logical drive. This condition would render the data unusable in a disaster recovery situation.

# **Creating a Remote Mirror Pair**

Before you create any mirror relationships, logical drives must exist at both the primary site and the secondary site. The logical drive that resides on the local storage subsystem is the primary logical drive. Similarly, the logical drive that resides on the remote storage subsystem is the secondary logical drive. If neither the primary logical drive or the secondary logical drive exist, you must create these logical drives. When you create the secondary logical drive, you must consider these items:

- The secondary logical drive must be of equal or greater size than the primary logical drive.
- The RAID level of the secondary logical drive does not have to be the same as the primary logical drive.

Use these steps to create the logical drive.

- 1. Enable the Enhanced Remote Mirroring Feature.
- 2. Activate the Enhanced Remote Mirroring Feature.
- 3. Determine Candidates for a Remote Mirror Pair.
- 4. Create the Remote Mirror Relationship.

#### **Performance Considerations**

Consider these performance issues when you create mirror relationships:

- The controller owner of a primary logical drive performs a full synchronization in the background while processing local I/O writes to the primary logical drive and associated remote writes to the secondary logical drive. Because the full synchronization diverts controller processing resources from I/O writes, full synchronization can have a performance impact to the host application.
- To reduce the performance impact, you can set the synchronization priority level to determine how the controller owner will prioritize the full synchronization relative to other I/O activity. To set the synchronization priority level, consider these guidelines:
  - A full synchronization at the lowest synchronization priority level takes approximately eight times as long as a full synchronization at the highest synchronization priority level.
  - A full synchronization at the low synchronization priority level takes approximately six times as long as a full synchronization at the highest synchronization priority level.
  - A full synchronization at the medium synchronization priority level takes approximately three-and-a-half times as long as a full synchronization at the highest synchronization priority level.
  - A full synchronization at the high synchronization priority level takes approximately twice as long as a full synchronization at the highest synchronization priority level.
- When the mirrored logical drive pair is in a Synchronization in Progress state, all host write data is copied to the remote system. Both controller I/O bandwidth and I/O latency can affect host write performance. Host read performance is not affected by the mirroring relationship.
- The time that it takes for data to be copied from the primary logical drive to the secondary logical drive might impact overall performance. This impact is primarily caused by the delay and system resource required for copying data to the remote mirror. Some delay might also occur because of the limit to the number of simultaneous writes.

## **Enabling the Enhanced Remote Mirroring Feature**

The first step in creating a remote mirror is to make sure that the Enhanced Remote Mirroring feature is enabled on both storage subsystems. Because Enhanced Remote Mirroring is a premium feature, you need a Feature Key file to enable the feature. The command for enabling the Feature Key file is as follows: enable storageSubsystem feature file="filename"

where the **file** parameter is the complete file path and file name of a valid Feature Key file. Enclose the file path and the file name in double quotation marks (""). Valid file names for Feature Key files end with a .key extension.

# Activating the Enhanced Remote Mirroring Feature

Activating the Enhanced Remote Mirroring feature prepares the storage subsystems to create and configure mirror relationships. After you activate the feature, the secondary ports for each controller are reserved and dedicated to remote mirror use. In addition, a mirror repository logical drive is automatically created for each controller in the storage subsystem. As part of the activation process, you can decide where the mirror repository logical drives will reside, free capacity on an existing subsystem or in a newly created subsystem, and the RAID level for the mirror repository logical drives.

The free capacity that you select for the mirror repository logical drive must have a total of 256 MB of capacity available. Two mirror repository logical drives are created on this capacity, one for each controller. If you enter a value for the repository storage space that is too small for the mirror repository logical drives, the firmware returns an error message that gives the amount of space needed for the mirror repository logical drives. The command does not try to activate the Enhanced Remote Mirroring feature. You can re-enter the command using the value from the error message for the repository storage space value.

The RAID level that you choose for the mirror repository logical drive has these constraints:

- **RAID 0** You cannot use RAID 0.
- **RAID 1** The number of drives must be an even number. If you select an odd number of drives, the controller firmware returns an error.
- **RAID 3 or RAID 5** You must have a minimum of three disk drives in the subsystem.
- RAID 6 You must have a minimum of five disk drives in the subsystem.

To activate the Enhanced Remote Mirroring feature, use this command: activate storageSubsystem feature=remoteMirror

The activate storageSubsystem feature=remoteMirror command provides three methods for defining the disk drives for your mirror repository logical drive:

- You define each disk drive for the mirror repository logical drive by its enclosure ID and its slot ID.
- You define an subsystem in which the mirror repository logical drive resides. You can optionally define the capacity of the mirror repository logical drive.
- You define the number of disk drives, but not specific disk drives, for the mirror repository logical drive.

# Activating the Enhanced Remote Mirroring Feature with User-Assigned Disk Drives

Activating the Enhanced Remote Mirroring feature by assigning the disk drives provides flexibility in defining your configuration by letting you choose from the available disk drives in your storage subsystem. Choosing the disk drives for your remote mirror automatically creates a new subsystem. You can specify which disk drives to use and the RAID level for the new subsystem.

The command takes this form:

activate storageSubsystem feature=remoteMirror repositoryRAIDLevel=(1 | 3 | 5 | 6) repositoryDrives= (enclosureID1,slotID1 ... enclosureIDn,slotIDn) enclosureLossProtect=(TRUE | FALSE)

This example shows a command in which you assign the disk drives:

```
c:\...\client>smcli 123.45.67.89 -c "activate
storageSubsystem feature=remoteMirror repositoryRAIDLevel=5
repositoryDrives=(1,1 1,2 1,3 1,4 1,5);"
```

The command in this example creates a new mirror repository logical drive consisting of five disk drives that forms a new subsystem. The new subsystem has a RAID level of 5.

This example shows how to use the command in a script file:

activate storageSubsystem feature=remoteMirror repositoryRAIDLevel=5 repositoryDrives=(1,1 1,2 1,3 1,4 1,5);

# Activating the Enhanced Remote Mirroring Feature with Software-Assigned Disk Drives

With this version of the activate storageSubsystem feature=remoteMirror command, you choose an existing subsystem in which to place the mirror repository logical drive. The storage management software then determines which disk drives to use. You can also define how much space to assign to the mirror repository logical drive. Because you are using an existing subsystem, the RAID level for the mirror repository logical drive defaults to the RAID level of the subsystem in which you place it. You cannot define the RAID level for the mirror repository logical drive.

The command takes this form:

activate storageSubsystem feature=remoteMirror repositorysubsystem=subsystemNumber
[freeCapacityArea=freeCapacityIndexNumber]

This example shows a command in which the software assigns the disk drives: c:\...\client>smcli 123.45.67.89 -c "activate storageSubsystem feature=remoteMirror repositorysubsystem=2 freeCapacityArea=2;"

The command in this example creates a new mirror repository logical drive in subsystem 2 using the second free capacity area.

This example shows how to use the command in a script file: activate storageSubsystem feature=remoteMirror repositorysubsystem=2 freeCapacityArea=2;

# Activating the Enhanced Remote Mirroring Feature by Specifying a Number of Disk Drives

With this version of the activate storageSubsystem feature=remoteMirror command, you must specify the number of disk drives and the RAID level that you want for the mirror repository logical drive. This version of the command creates a new subsystem. For this command to work, you must have disk drives in the storage subsystem that are not assigned to an subsystem.

```
activate storageSubsystem feature=remoteMirror
repositoryRAIDLevel=(1 | 3 | 5 | 6)
repositoryDriveCount=numberOfDrives
[driveType=(fibre | SATA | SAS |)]
[enclosureLossProtect=(TRUE | FALSE)]
```

This example shows a command in which you specify the number of disk drives:

```
c:\...\client>smcli 123.45.67.89 -c "activate storageSubsystem f
eature=remoteMirror repositoryRAIDLevel=5 repositoryDriveCount=5
driveType=fibre;"
```

The command in this example creates a new mirror repository logical drive by using five software is selected disk drives for the mirror repository logical drive. The mirror repository logical drive has a RAID level of 5. The type of drive for the mirror repository logical drive is Fibre Channel.

This example shows how to use the command in a script file:

```
activate storageSubsystem feature=remoteMirror
repositoryRAIDLevel=5 repositoryDriveCount=5
driveType=fibre;
```

#### Determining Candidates for a Remote-Mirror Pair

All of the logical drives and disk drives on the remote storage subsystem might not be available for use as secondary logical drives. To determine which logical drives on a remote storage subsystem that you can use as candidates for secondary logical drives, use the **show remoteMirror candidates** command. This command returns a list of the logical drives that you can use when creating a remote mirror.

The command takes this form:

```
c:\...\client>smcli 123.45.67.89 -c "show
remoteMirror candidates primary=\"logicalDriveName\"
remoteStorageSubsystemName=\"storageSubsystemName\";"
```

where *logicalDriveName* is the name of the logical drive that you want to use for the primary logical drive, and *StorageSubsystemName* is the remote storage subsystem that contains possible candidates for the secondary logical drive. Enclose both the logical drive name and the storage subsystem name in double quotation marks ("").

#### Creating a Remote Mirror Pair

When you create a new remote mirror, you must define which logical drives that you want to use for the primary (local) logical drive and the secondary (remote) logical drive. You define the primary logical drive by the name of the logical drive. You define the secondary logical drive by name with either the name or the World-Wide Identifier (WWID) of the storage subsystem on which the secondary logical drive resides. The primary logical drive name, the secondary logical drive name, and the remote storage subsystem name (or WWID) are the minimum information that you need to provide. Using this command, you can also define synchronization priority, write order, and write mode.

The command takes this form:

```
create remoteMirror primary="primaryLogicalDriveName"
secondary="secondaryLogicalDriveName"
(remoteStorageSubsystemName="storageSubsystemName" |
```

```
remoteStorageSubsystemWwn="wwID") remotePassword=password
syncPriority=(highest | high | medium | low | lowest)
writeOrder=(preserved | notPreserved)
writeMode=(synchronous | asynchronous)
```

**Note:** You can use the optional parameters as needed to help define your configuration.

This example shows the create remoteMirror command:

```
c:\...\client>smcli 123.45.67.89 -c "create
remoteMirror primary=\"Jan_04_Account\"
secondary=\"Jan_04_Account_B\" remoteStorageSubsystemName=\"Tabor\"
remotePassword=\"jdw2ga05\" syncPriority=highest
writeMode=synchronous;"
```

The command in this example creates a remote mirror in which the primary logical drive is named Jan\_04\_Account on the local storage subsystem. The secondary logical drive is named Jan\_04\_Account\_B on the remote storage subsystem that is named Tabor. The names used in this example are similar, but that is not a requirement for the logical drive names in a remote-mirror pair. In this example, the remote storage subsystem has a password that you must enter when making any change to the storage subsystem configuration. Creating a remote-mirror pair is a significant change to a storage subsystem configuration. Setting the write mode to synchronous and the synchronization priority to highest means that host write requests are written to the primary logical drive and then immediately copied to the secondary logical drive. These actions help to make sure that the data on the secondary logical drive is as accurate a copy of the data on the primary logical drive as possible. The highest synchronization priority does, however, use more system resources, which can reduce system performance.

This example shows how to use the command in a script file:

```
create remoteMirror primary="Jan_04_Account"
secondary="Jan_04_Account_B" remoteStorageSubsystemName="Tabor"
remotePassword="jdw2ga05" syncPriority=highest
writeMode=synchronous;
```

After you have created a remote mirror, you can see the progress of data synchronization between the primary logical drive and the secondary logical drive by running the show remoteMirror synchronizationProgress command. This command shows the progress as a percentage of data synchronization that has completed.

### Changing Enhanced Remote Mirroring Settings

The set remoteMirror command lets you change the property settings for a remote mirror pair. Use this command to change these property settings:

- The logical drive role (either primary or secondary)
- The synchronization priority
- The write order
- The write mode

You can apply the changes to one or several remote mirror pairs by using this command. Use the primary logical drive name to identify the remote mirror pairs for which you are changing the properties.

This example shows how to use the set remoteMirror command:

c:\...\client>smcli 123.45.67.89 -c "set remoteMirror localLogicalDrive [Jan\_04\_Account] syncPriority=medium writeOrder=notpreserved writeMode=asynchronous;"

This example shows how to use the command in a script file:

set remoteMirror localLogicalDrive [Jan\_04\_Account]
syncPriority=medium writeOrder=notpreserved
writeMode=asynchronous;

#### Suspending and Resuming a Mirror Relationship

Use the suspend remoteMirror command to stop data transfer between a primary logical drive and a secondary logical drive in a mirror relationship without disabling the mirror relationship. Suspending a mirror relationship lets you control when the data on the primary logical drive and data on the secondary logical drive are synchronized. Suspending a mirror relationship helps to reduce any performance impact to the host application that might occur while any changed data on the primary logical drive is copied to the secondary logical drive. Suspending a mirror relationship is particularly useful when you want to run a backup of the data on the secondary logical drive.

When a mirror relationship is in a suspended state, the primary logical drive does not make any attempt to contact the secondary logical drive. Any writes to the primary logical drive are persistently logged in the mirror repository logical drives. After the mirror relationship resumes, any data that is written to the primary logical drive is automatically written to the secondary logical drive. Only the modified data blocks on the primary logical drive are written to the secondary logical drive. Full synchronization is not required.

**Important:** If you suspend a remote mirror that is set up in the Global mirror, you suspend all remote mirror pairs within the group. You can then resume mirror operations for any of the individual remote-mirror pairs that are in the group.

This example shows the suspend remoteMirror command:

c:\...\client>smcli 123.45.67.89 -c "suspend remoteMirror primary Jan\_04\_Account
writeConsistency=false;"

The **writeConsistency** parameter defines whether the logical drives identified in this command are in a write-consistency group or are separate. For the logical drives in a write $\triangle$  consistency group, set this parameter to TRUE. For the logical drives that are not in a write-consistency group, set this parameter to FALSE.

This example shows how to use the command in a script file: suspend remoteMirror logicalDrive Jan\_04\_Account writeConsistency=false;

The mirror relationship remains suspended until you use the resume remoteMirror command to restart synchronization activities. This command restarts data transfers between a primary logical drive and a secondary logical drive in a mirror relationship after the mirror has been suspended or unsynchronized.

This example shows the resume remoteMirror command: c:\...\client>smcli 123.45.67.89 -c "resume remoteMirror logicalDrive Jan 04 Account writeConsistency=false;"

The **writeConsistency** parameter in this command operates the same as in the previous command.

This example shows how to use the command in a script file: resume remoteMirror logicalDrive Jan\_04\_Account writeConsistency=false;

# **Removing a Mirror Relationship**

Use the remove remoteMirror command to remove the link between a primary logical drive and a secondary logical drive. (Removing a mirror relationship is similar to deleting a mirror relationship.) Removing the link between a primary logical drive and a secondary logical drive does not affect any of the existing data on either logical drive. The link between the logical drives is removed, but the primary logical drive still continues normal I/O operations. Later, you can establish the mirror relationship between the two logical drives and resume normal mirror operations. You can remove the mirror relationship for one or several remote mirror pairs with this command.

This example shows the remove remoteMirror command:

c:\...\client>smcli 123.45.67.89 -c "remove remoteMirror localLogicalDrive
[Jan\_04\_Account];"

When you run this command, use the name of the primary logical drive of the remote mirror pair.

This example shows how to use the command in a script file: remove remoteMirror localLogicalDrive [Jan 04 Account];

To re-establish the link between a primary logical drive and a secondary logical drive, use the create remoteMirror command.

## Deleting a Primary Logical Drive or a Secondary Logical Drive

Use the **delete logicalDrive** command to remove a primary logical drive or a secondary logical drive from a storage subsystem. Deleting a logical drive in a mirror relationship removes the mirror relationship and completely deletes the logical drive from the storage subsystem. You cannot redefine the mirror relationship until you create a new logical drive or choose an alternate logical drive to replace the deleted logical drive.

**Attention: Possible loss of data access** – Deleting a primary logical drive or a secondary logical drive permanently removes the data from the storage subsystem.

# **Disabling the Enhanced Remote Mirroring Feature**

You disable the Enhanced Remote Mirroring feature to prevent the new mirror relationship from being created. When you disable the Enhanced Remote Mirroring feature, the feature is in a Disabled/Active state. In this state, you can maintain and manage previously existing mirror relationships; however, you cannot create new relationships. To disable the Enhanced Remote Mirroring feature, use this command:

disable storageSubsystem feature=remoteMirror

# **Deactivating the Enhanced Remote Mirroring Feature**

If you no longer require the Enhanced Remote Mirroring feature and you have removed all of the mirror relationships, you can deactivate the feature. Deactivating the feature redestablishes the normal use of dedicated ports on both storage subsystems and deletes both mirror repository logical drives. To deactivate the Enhanced Remote Mirroring feature, use this command:

deactivate storageSubsystem feature=remoteMirror

# Interaction with Other Features

You can run the Enhanced Remote Mirroring feature while running these premium features:

- Storage Partitioning
- FlashCopy Logical Drive
- VolumeCopy

When you run the Enhanced Remote Mirroring feature with other premium features, you must consider the requirements of the other premium features to help make sure that you set up a stable storage subsystem configuration.

In addition to running with the premium features, you can also run the Enhanced Remote Mirroring feature while running Dynamic Volume Expansion (DVE).

# **Storage Partitioning**

Storage Partitioning is a premium feature that lets hosts share access to logical drives in a storage subsystem. You create a storage partition when you define any of these logical components in a storage subsystem:

- A host
- A host group
- A Volume-to-LUN mapping

The Volume-to-LUN mapping lets you define which host group or host has access to a particular logical drive in the storage subsystem.

When you create storage partitions, define the partitions after you have created the primary logical drive and the secondary logical drive in a Enhanced Remote Mirroring configuration. The storage partition definitions for the primary storage subsystem and the secondary storage subsystem are independent of each other. If these definitions are put in place while the logical drive is in a secondary role, the administrative effort associated with the site recovery is reduced if it becomes necessary to promote the logical drive to a primary role.

# FlashCopy Logical Drives

A FlashCopy logical drive is a point-in-time image of a logical drive. Typically, it is created so that an application, such as a backup application, can access the FlashCopy logical drive and read the data while the base logical drive stays online and is accessible to hosts.

The logical drive for which the point-in-time image is created is known as the base logical drive and must be a standard logical drive in the storage subsystem. The FlashCopy repository logical drive stores information about all data that changed since the FlashCopy was created.

In this version of the storage management software, you can create FlashCopy logical drives based on the primary logical drive or secondary logical drive of a remote mirror.

## VolumeCopy

The VolumeCopy premium feature copies data from one logical drive (the source logical drive) to another logical drive (the target logical drive) within a single storage subsystem. You can use this feature to perform these functions:

- Copy data from subsystems that use smaller-capacity disk drives to subsystems that use larger-capacity disk drives
- Back up data
- Restore FlashCopy logical drive data to the base logical drive.

You can use a primary logical drive in a remote mirror as a source logical drive or a target logical drive in a VolumeCopy. You cannot use a secondary logical drive as a source logical drive or a target logical drive.

**Note:** If you start a role reversal during a copy-in-progress, the copy fails and cannot be restarted.

# Dynamic Logical Drive Expansion (DVE)

A Dynamic Volume Expansion (DVE) is a modification operation that increases the capacity of a standard logical drive or a FlashCopy repository logical drive. The increase in capacity is achieved by using the free capacity that is available in the subsystem of the standard logical drive or the FlashCopy repository logical drive.

This modification operation is considered to be "dynamic" because you can continually access data on subsystems, logical drives, and disk drives throughout the entire operation.

A DVE operation can be performed on a primary logical drive or a secondary logical drive of a mirror relationship.

**Note:** Although the storage management software indicates that the logical drive has increased capacity, its usable capacity is the size of the smaller of the primary logical drive or the secondary logical drive.

You cannot perform a DVE operation on a mirror repository logical drive.

#### Enhanced Remote Mirroring utility

This section describes the host utility to achieve periodic consistency with Enhanced Remote Mirroring configurations. This section also describes how to run the Enhanced Remote Mirroring utility.

## Description of the Enhanced Remote Mirroring Utility

The Enhanced Remote Mirroring utility lets you periodically synchronize the Enhanced Remote Mirroring pairs in your storage subsystem. When defining a Enhanced Remote Mirroring configuration, you have the option to set the write modes to either Synchronous or Asynchronous. Metro mirror provides the highest level security for full data recovery from the secondary storage subsystem in the event of a disaster. Metro mirror does, however, reduce host I/O performance.

Global copy offers faster host I/O performance, but it does not guarantee that a copy operation has successfully completed before processing the next write request. With Global copy, you cannot make sure that a logical drive, or collection of logical drives, at a secondary site ever reach a consistent, recoverable state.

The Enhanced Remote Mirroring utility enables you to bring a collection of remote logical drives into a mutually consistent and recoverable state. You can choose to run the utility based on application demands, link state and speed, and other factors that are relevant to your environment.

The Enhanced Remote Mirroring utility has these characteristics:

- The utility is implemented as a command line-invoked Java-based application.
- The utility is bundled as part of the DS5000 Storage Manager installation package.
- The utility accepts a command line argument that lets you specify the name of a configuration file that contains a complete specification of the work to be carried out by the utility.
- More than one instance of the utility can run concurrently, as long as the utilities do not try to process any of the same logical drives and mirrors.

**Note:** The Enhanced Remote Mirroring utility does not check to make sure that concurrently running instances of the utility are not trying to process the same logical drives and mirrors. If you choose to simultaneously run more than one instance of the Enhanced Remote Mirroring utility, you must make sure that the configuration files that you choose to run do not list the same logical drives and mirrors.

# **Operation of the Enhanced Remote Mirroring Utility**

The Enhanced Remote Mirroring utility performs steps that generate a recoverable state for multiple mirror logical drives at a secondary site. The utility runs these steps to create consistent, recoverable images of a set of logical drives:

- 1. **On the primary storage subsystem** The utility reconfigures all of the participating logical drives from Enhanced Global Mirror Group to Enhanced Remote Mirroring. This action makes sure that the stream of write operations becomes recoverable on the secondary side.
- 2. On the primary storage subsystem The utility polls all of the participating logical drives until the associated mirror states all have the Optimal state. In cases where the remote link is slow or the primary host I/O activity is high, one or more mirrors are likely to be in the Unsynchronized state before they migrate to the Synchronized state. By waiting until all of the mirrors have Optimal status, the utility makes sure that all of the delta logs for the affected logical drives are cleared, and the secondary logical drives are recoverable.
- 3. On the primary storage subsystem The utility suspends the mirrored pairs for all of the participating logical drives. This action causes updates to stop on the secondary side, leaving the secondary logical drives in a recoverable state because they were being updated in Synchronous mode immediately before the suspension. By separating the mirrors in this manner, the primary-side applications run faster, while leaving the secondary logical drives in a recoverable state. The delta log tracks changes made because of application writes on the primary side while in this state.
- On the secondary storage subsystem The utility generates a FlashCopy of each participating logical drive on the secondary side, which creates point-in-time images that are recoverable.

- 5. On the primary storage subsystem The utility resumes the mirroring operations for all of the participating logical drives. This action causes the mirrors to migrate to the Synchronized state and start the process of restoring coherency between the primary site and the secondary site.
- 6. **On the primary storage subsystem** The utility reconfigures all of the affected logical drives for Asynchronous mode.

# **Running the Enhanced Remote Mirroring Utility**

The Enhanced Remote Mirroring utility uses a command line argument that lets you specify the name of a configuration file. The configuration file contains a complete specification of the input parameters that are needed by the utility. To run the utility, enter this syntax:

asyncRVMUtil configuration\_file -d debug\_file

where *configuration\_file* is the file that you provide as input. The configuration file specifies the Enhanced Remote Mirroring logical drives that you want to synchronize by using the utility. When you create the configuration file, use these conditions to define the logical drives in the file:

- All the primary logical drives in a logical drive set must belong to the same storage subsystem.
- The maximum number of logical drive sets that you can specify in the file is four.
- The maximum number of mirrored pairs that you can specify as part of a consistency group is eight.

The optional parameter, **-d**, lets you specify a file to which you can send information regarding how the utility runs. In this example, the file name is *debug\_file*. The debug file contains trace information that can be reviewed by an IBM Technical Support representative to determine how well the Enhanced Remote Mirroring utility has run.

**Note:** Depending on the location of the configuration file and the debug file, you must specify the complete path with the file name.

To run the Enhanced Remote Mirroring utility, you must enter the **asyncRVMUtil** command from the command line. Because UNIX operating systems are case sensitive, you must type the command exactly as shown. On Windows operating systems, you can type the command in all uppercase, in all lowercase, or in mixed case.

**Note:** To use the Enhanced Remote Mirroring utility, you must be managing the storage subsystem by using the command line interface, not the graphical user interface of DS Storage Manager.

# **Configuration Utility**

The configuration file is an ASCII flat text file that provides the information for the Enhanced Remote Mirroring synchronization used by the Enhanced Remote Mirroring utility. The file defines the mirror logical drive sets to be synchronized. All of the mirror logical drives in the logical drive sets defined in the configuration file are run collectively to create a recoverable image. If any one of the mirrors in the logical drive set fails, the operation is stopped for this logical drive set and carried on to the next logical drive set that is listed in the configuration file.

The configuration file supports this syntax:

```
content ::= {spec}
spec ::= logSpec | logicalDriveSetSpec
logSpec ::= "Log" "{" {logAttribute} "}"
logAttribute ::= fileSpec
fileSpec ::= "file" "=" fileName
LogicalDriveSetSpec ::= "LogicalDriveSet" LogicalDriveSetName
"{" {LogicalDriveSetAttribute} "}"
LogicalDriveSetAttribute ::= timeoutSpec | mirrorSpec
timeoutSpec ::= "OptimalWaitTimeLimit" "=" integer
mirrorSpec ::= "Mirror" "{" {mirrorAttribute} "}"
mirrorAttribute ::= primarySpec | secondarySpec |
FlashCopySpec
primarySpec ::= "Primary" "=" LogicalDriveSpec
secondarySpec ::= "Secondary" "=" LogicalDriveSpec
FlashCopySpec ::= "Copy" "=" logicalDriveSpec
logicalDriveSpec ::= storageSubsystemName"."logicalDriveUserLabel
```

In this syntax, items enclosed in double quotation marks ("") are terminal symbols. Items separated by a vertical bar (1) are alternative values (enter one or the other, but not both). Items enclosed in curly braces ({ }) are optional (you can use the item zero or more times).

These definitions are provided for non-terminals in the syntax:

- *integer* The timeout value must be an integer (decimal digits from 0–9).
- *logicalDriveSetName* The name of the set of logical drives on which you want to run the Enhanced Remote Mirroring utility.
- *fileName* The name of a file, using characters and conventions that are appropriate for the system on which the application is running.
- storageSubsystemName The label that you have assigned for a storage subsystem, as would be used in the CLI to specify the name of the storage subsystem.
- *logicalDriveUserLabel* The label that you have assigned for a logical drive that uniquely identifies the logical drive within the storage subsystem.

**Note:** Names and labels can be any characters that are defined as appropriate for your operating system. The maximum length for a name or label is 30 characters. If the name or label contains special characters (as defined by the operating system) or period characters, you must enclose the name or label in double quotation marks (""). You can, optionally, enclose the name or label in double quotation marks at any time.

The following items are considered syntax errors:

- More than one logSpec command in the input file
- Zero or more than one **fileSpec** attribute in a **logSpec** command (you must include exactly one **fileSpec** attribute in the **logSpec** command)
- More than one timeoutSpec attribute in a logicalDriveSetSpec command
- Zero or more than one primarySpec attribute in a mirrorSpec command (you must include exactly one primarySpec attribute in the mirrorSpec command)
- Zero or more than one secondarySpec attribute in a mirrorSpec command (you
  must include exactly one secondarySpec attribute in the mirrorSpec command)
- Zero or more than one FlashCopySpec attribute in a mirrorSpec command (you
  must include exactly one FlashCopySpec attribute in the mirrorSpec command)

**Important:** In the Enhanced Remote Mirroring utility configuration file, you must specify the primary logical drive, the secondary logical drive, and the copy (FlashCopy ) logical drive. The utility does not make sure that the secondary logical drive is correct for the Enhanced Remote Mirroring relationship. The utility also does not make sure that the FlashCopy logical drive is actually a FlashCopy for the secondary logical drive. *You must make sure that these logical drives are correct*. If the logical drives are not correct, the utility will run, but the logical drives will not be consistent. For each mirror, the secondary logical drive and the copy logical drive must reside on the same storage subsystem.

This example shows a configuration file for the Enhanced Remote Mirroring utility.

```
Log{ file="d:\rvm-consistency.log" }
LogicalDriveSet "set1" {
optimalWaitTimeLimit = 15
Mirror {
 Primary = LosAngelesSubsystem.PayrollLogicalDrive
 Secondary = NewYorkSubsystem.PayrollLogicalDrive
 Copy = NewYorkSubsystem.PayrollLogicalDriveImage
}
Mirror {
 Primary = LosAngelesSubsystem.PayrollLogicalDrive
 Secondary = BostonSubsystem.PayrollLogicalDrive
 Copy = BostonSubsystem.PayrollLogicalDriveImage
LogicalDriveSet "set2" {
Mirror {
 Primary = BostonSubsystem.HRLogicalDrive
 Secondary = LosAngelesSubsystem.HRLogicalDrive
 Copy = LosAngelesSubsystem.HRLogicalDriveImage
}
```

# Chapter 9. Using the VolumeCopy Premium Feature

The VolumeCopy premium feature lets you copy data from one logical drive (the source) to another logical drive (the target) in a single storage subsystem. You can use this feature to perform these tasks:

- Back up data
- Copy data from subsystems that use smaller-capacity disk drives to subsystems using greater-capacity disk drives
- Restore FlashCopy logical drive data to the associated base logical drive

# How VolumeCopy Works

When you create a VolumeCopy, you create a copy pair that consists of a source logical drive and a target logical drive. Both the source logical drive and the target logical drive are located on the same storage subsystem. During a VolumeCopy, the controllers manage copying the data from the source logical drive to the target logical drive. The VolumeCopy is transparent to the host machines and applications, except that you cannot write to the source logical drive during a VolumeCopy operation.

**Tip:** Because of this restriction, IBM strongly recommends using a FlashCopy version of the logical drive as the source for the VolumeCopy, rather than using the original logical drive as the source. If you use the FlashCopy as the source for the logical VolumeCopy operation, then I/O operations to the original logical drive can continue.

While a VolumeCopy is In Progress, the same controller must own both the source logical drive and the target logical drive. If one controller does not own both the source logical drive and the target logical drive before creating the VolumeCopy, ownership of the target logical drive is automatically transferred to the controller that owns the source logical drive. When the VolumeCopy is finished or stopped, ownership of the target logical drive is restored to its preferred controller. If ownership of the source logical drive changes while a VolumeCopy is running, ownership of the target logical drive also changes.

## **Source Logical Drive**

The source logical drive is the logical drive that accepts host I/O and stores data. When you start a VolumeCopy, data from the source logical drive is copied in its entirety to the target logical drive. While a VolumeCopy has a status of In Progress, Pending, or Failed, the source logical drive is available only for read activity.

After the VolumeCopy completes, the source logical drive becomes available to host applications for write requests. The target logical drive automatically becomes read-only to hosts, and write requests to the target logical drive are rejected.

The following are valid source logical drives:

- A standard logical drive
- A FlashCopy logical drive
- The base logical drive of a FlashCopy logical drive

· A primary logical drive that is participating in a remote-mirror pair

The following are not valid source logical drives:

- · A secondary logical drive that is participating in a remote-mirror pair
- A FlashCopy repository logical drive
- A mirror repository logical drive
- A failed logical drive
- A missing logical drive
- A logical drive currently in a modification operation
- A logical drive that is holding a Small Computer System Interface-2 (SCSI-2) reservation or a persistent reservation
- A logical drive that is a source logical drive or a target logical drive in another VolumeCopy that has a status of In Progress, Pending, or Failed

### **Target Logical Drive**

A target logical drive contains a copy of the data from the source logical drive. When a VolumeCopy is started, data from the source logical drive is copied in its entirety to the target logical drive.

**Attention: Possible loss of data access** – A VolumeCopy overwrites data on the target logical drive. Before you start a new operation, make sure that you no longer need the old data, or you have backed up the old data on the target logical drive.

While the VolumeCopy has a status of In Progress, Pending, or Failed, the controllers reject read and write requests to the target logical drive. After the VolumeCopy operation is finished, the target logical drive automatically becomes read-only to the hosts, and write requests to the target logical drive are rejected. You can change the Read-Only attribute after the VolumeCopy has completed or has been stopped. (For more information about the Read-Only attribute, see "Viewing logical VolumeCopy Properties" on page 518.)

The following logical drives are valid target logical drives:

- A standard logical drive
- The base logical drive of a disabled FlashCopy logical drive or failed FlashCopy logical drive
- · A primary logical drive that is participating in a remote-mirror pair

The following logical drives are not valid target logical drives:

- The base logical drive of an active FlashCopy logical drive
- A FlashCopy logical drive
- A mirror repository logical drive
- A FlashCopy repository logical drive
- · A secondary logical drive in a remote-mirror pair
- A failed logical drive
- A missing logical drive
- A logical drive with a status of Degraded
- A logical drive that is currently in a modification operation
- A logical drive that is holding a SCSI-2 reservation or a persistent reservation

• A logical drive that is a source logical drive or a target logical drive in another VolumeCopy that has a status of In Progress, Pending, or Failed

# VolumeCopy and Persistent Reservations

You cannot use logical drives that hold persistent reservations for either a source logical drive or a target logical drive. Persistent reservations are configured and managed through the server cluster software and prevent other hosts from accessing the reserved logical drive. Unlike other types of reservations, a persistent reservation reserves host access to the logical drive across multiple HBA host ports, which provides various levels of access control.

To determine which logical drives have reservations, run the **show (logicalDrive) reservations** command. To remove a reservation, run the **clear (logicalDrive) reservations** command.

# **Storage Subsystem Performance**

During a VolumeCopy operation, the resources of the storage subsystem might be diverted from processing I/O activity to completing a VolumeCopy, which might affect the overall performance of the storage subsystem.

These factors contribute to the performance of the storage subsystem:

- The I/O activity
- The logical drive RAID level
- The logical drive configuration (number of disk drives in the subsystems and cache parameters)
- The logical drive type (FlashCopy logical drives might take more time to copy than standard logical drives)

When you create a new VolumeCopy, you define the copy priority to determine how much controller processing time is allocated for a VolumeCopy compared with I/O activity.

Copy priority has five relative settings ranging from highest to lowest. The highest priority rate supports the logical VolumeCopy, but I/O activity might be affected. The lowest priority rate supports I/O activity, but the VolumeCopy takes longer. You define the copy priority when you create the VolumeCopy pair. You can redefine the copy priority later by using the **set VolumeCopy** command. You can also redefine the VolumeCopy priority when you recopy a logical drive.

## **Restrictions**

These restrictions apply to the source logical drive, the target logical drive, and the storage subsystem:

- While a VolumeCopy operation has a status of In Progress, Pending, or Failed, the source logical drive is available for read activity only. After the VolumeCopy finishes, read activity from and write activity to the source logical drive is permitted.
- A logical drive can be selected as a target logical drive for only one VolumeCopy at a time.
- The maximum allowable number of VolumeCopies per storage subsystem depends upon the storage subsystem configuration.
- A logical drive that is reserved by the host cannot be selected as a source logical drive or as a target logical drive.

- A logical drive with a status of Failed cannot be used as a source logical drive or as a target logical drive.
- A logical drive with a status of Degraded cannot be used as a target logical drive.
- You cannot select a logical drive that is participating in a modification operation as a source logical drive or as a target logical drive. Modification operations include Dynamic Capacity Expansion (DCE), Dynamic RAID Level Migration (DRM), Dynamic Segment Sizing (DSS), Dynamic Volume Expansion (DVE), and defragmenting an subsystem.

#### VolumeCopy Commands

The following table lists the VolumeCopy commands and briefly describes what the commands do.

Command	Description
create VolumeCopy	Creates a VolumeCopy and starts the VolumeCopy operation.
disable storageSubsystem feature VolumeCopy	Turns off the current VolumeCopy operation.
enable storageSubsystem feature	Activates the VolumeCopy premium feature.
recopy VolumeCopy	Re-initiates a VolumeCopy operation using an existing VolumeCopy pair.
remove VolumeCopy	Removes a VolumeCopy pair.
set VolumeCopy	Defines the properties for a VolumeCopy pair.
show VolumeCopy	Returns information about VolumeCopy operations. You can retrieve information about a specific VolumeCopy pair, or all of the VolumeCopy pairs in the storage subsystem.
show VolumeCopy sourceCandidates	Returns information about the candidate logical drives that you can use as the source for a VolumeCopy operation.
show VolumeCopy targetCandidates	Returns information about the candidate logical drives that you can use as the target for a VolumeCopy operation.
stop VolumeCopy	Stops a logical VolumeCopy operation.

Table 24. VolumeCopy Commands

# Creating a logical VolumeCopy

Before you create a VolumeCopy, make sure that a suitable target logical drive exists on the storage subsystem, or create a new target logical drive specifically for the VolumeCopy. The target logical drive that you use must have a capacity equal to or greater than the source logical drive.

You can have a maximum of eight VolumeCopies with a status of In Progress at one time. Any VolumeCopy greater than eight has a status of Pending until one of the VolumeCopies with a status of In Progress has completed the VolumeCopy process.

To create a VolumeCopy, perform these general steps:

1. Enable the VolumeCopy premium feature.

- 2. Determine the candidates for a VolumeCopy.
- 3. Create the target logical drive and the source logical drive for the VolumeCopy.

# Enabling the logical VolumeCopy Feature

The first step in creating a VolumeCopy is to make sure that the feature is enabled on the storage subsystem. Because VolumeCopy is a premium feature, you need a Feature Key file to enable the feature. This command enables the Feature Key file: enable storageSubsystem feature file="filename"

- where the **file** parameter is the complete file path and file name of a valid Feature Key file. Enclose the file path and file name in double quotation marks (" "). Valid file names for Feature Key files usually end with a .key extension.

# Determining logical VolumeCopy Candidates

All logical drives and disk drives might not be available for use in VolumeCopy operations. To determine which candidate logical drives on the storage subsystem that you can use as a source logical drive, use the **show VolumeCopy sourceCandidates** command. To determine which candidate logical drives on the storage subsystem that you can use as a target logical drive, use the **show VolumeCopy targetCandidates** command. These commands return a list of the expansion drawer, slot, and capacity information for the source logical drive candidates and the target logical drive candidates. You can use the **show VolumeCopy sourceCandidates** command and the **show VolumeCopy targetCandidates** command only after you have enabled the VolumeCopy premium feature.

# Creating a logical VolumeCopy

**Attention: Possible loss of data access** – A VolumeCopy overwrites data on the target logical drive. Ensure that you no longer need the data or have backed up the data on the target logical drive before you start a VolumeCopy operation.

When you create a VolumeCopy, you must define which logical drives that you want to use for the source logical drive and the target logical drive. You define the source logical drive and the target logical drive by the name of each logical drive. You can also define the copy priority and choose whether you want the target logical drive to be read-only after the data is copied from the source logical drive.

The command takes this form:

create VolumeCopy source="sourceName" target="targetName" [copyPriority=(highest | high | medium | low | lowest) targetReadOnlyEnabled=(TRUE | FALSE)]

**Important:** Setting the targetReadOnlyEnabled attribute to false breaks the one-to-one relationship between the source and target VolumeCopy logical drives. Therefore, any data that are written to the target VolumeCopy logical drive will not be copied to the source VolumeCopy logical drive. In addition, these data will be lost if you recopy of the source VolumeCopy logical drive to the target VolumeCopy logical drive.

Before you run the create VolumeCopy command, perform these actions:

- Stop all I/O activity to the source logical drive and the target logical drive.
- Dismount any file systems on the source logical drive and the target logical drive, if applicable.

This example shows the **create VolumeCopy** command:

c:\...\client>smcli 123.45.67.89 -c "create VolumeCopy source=\"Jaba\_Hut\" target=\"Obi\_1\" copyPriority=medium targetrReadOnlyEnabled=TRUE;"

The command in this example copies the data from the source logical drive named Jaba\_Hut to the target logical drive named Obi\_1. Setting the copy priority to medium provides a compromise between how quickly the data is copied from the source logical drive to the target logical drive and the amount of processing resources that are required for data transfers to other logical drives in the storage subsystem. Setting the **targetReadOnlyEnabled** parameter to TRUE means that write requests cannot be made to the target logical drive, making sure that the data on the target logical drive stays unaltered.

This example shows how to use the command in a script file: create VolumeCopy source="Jaba\_Hut" target="Obi\_1" copyPriority=medium targetReadOnlyEnabled=TRUE;

After the VolumeCopy operation is completed, the target logical drive automatically becomes read-only to hosts. Any write requests to the target logical drive are rejected, unless you disable the Read-Only attribute by using the **set VolumeCopy** command.

To view the progress of a VolumeCopy, use the **show logicalDrive actionProgress** command. This command returns information about the logical drive action, the percentage completed, and the time remaining until the VolumeCopy is complete.

# Viewing logical VolumeCopy Properties

Use the **show VolumeCopy** command to view information about one or more selected source logical drives or target logical drives. This command returns these values:

- The role
- The copy status
- The start time stamp
- The completion time stamp
- The copy priority
- The Read-Only attribute setting for the target logical drive
- The source logical drive World-Wide Identifier (WWID) or the target logical drive WWID

If a logical drive is participating in more than one VolumeCopy (it can be a source logical drive for one VolumeCopy operation and a target logical drive for another VolumeCopy operation), the details are repeated for each associated copy pair.

The command takes this form: show VolumeCopy (allLogicalDrives | source [sourceName] | target [targetName])

This example shows the **show VolumeCopy** command:

c:\...\client>smcli 123.45.67.89 -c "show VolumeCopy source [\"JabaHut\"];"

The command in this example is requesting information about the source logical drive JabaHut. If you wanted information about all of the logical drives, you would use the **allLogicalDrives** parameter. You can also request information about a specific target logical drive.

This example shows how to use the command in a script file: show VolumeCopy source ["JabaHut"];

# Changing logical VolumeCopy Settings

The **set VolumeCopy** command lets you change these property settings for a VolumeCopy pair:

- · The copy priority
- The target logical drive read/write permission

Copy priority has five relative settings ranging from highest to lowest. The highest priority supports the VolumeCopy, but I/O activity might be affected. The lowest priority supports I/O activity, but the VolumeCopy takes longer. You can change the copy priority at these times:

- Before the VolumeCopy operation starts
- While the VolumeCopy operation has a status of In Progress
- After the logical VolumeCopy operation has completed when re-creating a VolumeCopy operation by using the **recopy VolumeCopy** command

When you create a VolumeCopy pair and after the original VolumeCopy has completed, the target logical drive is automatically defined as read-only to the hosts. The read-only status of the target logical drive helps to make sure that the copied data on the target logical drive is not corrupted by additional writes to the target logical drive after the VolumeCopy is created. You want to maintain the read-only status when you are performing these tasks:

- Using the target logical drive for backup purposes
- Copying data from one subsystem to a larger subsystem for greater accessibility
- Planning to use the data on the target logical drive to copy back to the base logical drive in case of a disabled FlashCopy logical drive or failed FlashCopy logical drive

At other times, you might want to write additional data to the target logical drive. You can use the **set VolumeCopy** command to reset the read/write permission for the target logical drive.

**Note:** If you have set the VolumeCopy parameters to enable host writes to the target logical drive, the read request and the write request to the target logical drive are rejected while the VolumeCopy operation has a status of In Progress, Pending, or Failed.

The command takes this form:

```
set VolumeCopy target [targetName] [source [sourceName]]
copyPriority=(highest | high | medium | low | lowest)
targetReadOnlyEnabled=(TRUE | FALSE)
```

**Important:** Setting the targetReadOnlyEnabled attribute to false breaks the one-to-one relationship between the source and target VolumeCopy logical drives. Therefore, any data that are written to the target VolumeCopy logical drive will

not be copied to the source VolumeCopy logical drive. In addition, these data will be lost if you recopy of the source VolumeCopy logical drive to the target VolumeCopy logical drive.

Note: You can use the parameters as needed to help define your configuration.

This example shows the **set VolumeCopy** command:

c:\...\client>smcli 123.45.67.89 -c "set VolumeCopy target [\"Obi\_1\"] copyPriority=highest targetReadOnlyEnabled=FALSE;"

This example shows how to use the command in a script file: set VolumeCopy target ["Obi 1"] copyPriority=highest targetReadOnlyEnabled=FALSE;

#### **Recopying a Logical Drive**

Use the **recopy VolumeCopy** command to create a new VolumeCopy for a previously defined copy pair that has a status of Stopped, Failed, or Completed. You can use the **recopy VolumeCopy** command to create backups of the target logical drive. Then, you can copy the backup to tape for off-site storage.

When you use the **recopy VolumeCopy** command to make a backup, you cannot write to the source logical drive while the recopy operation is running. The recopy operation might take a long time.

**Restriction:** Because of this restriction, IBM strongly recommends using a FlashCopy version of the logical drive as the source for the VolumeCopy, rather than using the original logical drive as the source. If you use the FlashCopy as the source for the VolumeCopy operation, then I/O operations to the original logical drive can continue.

When you run the **recopy VolumeCopy** command, the data on the source logical drive is copied in its entirety to the target logical drive.

Attention: Possible loss of data access – The recopy VolumeCopy command overwrites existing data on the target logical drive and makes the target logical drive read-only to hosts. The recopy VolumeCopy command fails all of the FlashCopy logical drives that are associated with the target logical drive, if any exist.

You can also reset the copy priority by using the **recopy VolumeCopy** command if you want to change the copy priority for the recopy operation. The higher priorities allocate storage subsystem resources to the VolumeCopy at the expense of storage subsystem performance.

The command takes this form:

recopy VolumeCopy target [targetName] [source [sourceName] copyPriority=(highest | high | medium | low | lowest) targetReadOnlyEnabled=(TRUE | FALSE)]

**Note:** You can use the optional parameters as needed to help define your configuration.

This example shows the **show VolumeCopy** command:

c:\...\client>smcli 123.45.67.89 -c "recopy VolumeCopy target [\"Obi\_1\"]
copyPriority=highest;"

The command in this example copies data from the source logical drive that is associated with the target logical drive Obi\_1 to the target logical drive again. The copy priority is set to the highest value to complete the VolumeCopy as quickly as possible. The underlying consideration for using this command is that you have already created the VolumeCopy pair, which has already created one VolumeCopy. By using this command, you are copying the data from the source logical drive to the target logical drive with the assumption that the data on the source logical drive has changed since the previous copy was made.

This example shows you how to use the command in a script file: recopy VolumeCopy target ["Obi\_1"] copyPriority=highest;

# Stopping a logical VolumeCopy

The **stop VolumeCopy** command lets you stop a VolumeCopy that has a status of In Progress, Pending, or Failed. After you have stopped a logical VolumeCopy, you can use the **recopy VolumeCopy** command to create a new VolumeCopy by using the original VolumeCopy pair. After you stop a VolumeCopy operation, all of the mapped hosts will have write access to the source logical drive.

The command takes this form: stop VolumeCopy target [targetName] [source [sourceName]]

This example shows the **show VolumeCopy** command: c:\...\client>smcli 123.45.67.89 -c "stop VolumeCopy target [\"Obi\_1\"];"

This example shows how to use the command in a script file: stop VolumeCopy target ["Obi\_1"];

## **Removing Copy Pairs**

The **remove VolumeCopy** command lets you remove a VolumeCopy pair from the storage subsystem configuration. All of the VolumeCopy-related information for the source logical drive and the target logical drive is removed from the storage subsystem configuration. The data on the source logical drive or the target logical drive is not deleted. Removing a VolumeCopy from the storage subsystem configuration also removes the Read-Only attribute for the target logical drive.

**Important:** If the VolumeCopy has a status of In Progress, you must stop the VolumeCopy before you can remove the VolumeCopy pair from the storage subsystem configuration.

The command takes this form: remove VolumeCopy target [targetName] [source [sourceName]]

This example shows the **remove VolumeCopy** command: c:\...\client>smcli 123.45.67.89 -c "remove VolumeCopy target [\"Obi\_1\"];"

This example shows how to use the command in a script file: remove VolumeCopy target ["Obi\_1"];

#### Interaction with Other Features

You can run the VolumeCopy premium feature while running the following premium features:

- Storage Partitioning
- FlashCopy Logical Drive
- Enhanced Remote Mirroring

When you are running the VolumeCopy feature with other premium features, you must consider the requirements of other premium features to help make sure that you set up a stable storage subsystem configuration.

In addition to the premium features, you also can run the VolumeCopy feature while running Dynamic Logical Drive Expansion (DVE).

#### Storage Partitioning

Storage Partitioning is a premium feature that lets hosts share access to logical drives in a storage subsystem. You create a storage partition when you define any of these logical components in a storage subsystem:

- A host
- A host group
- A Volume-to-LUN mapping

The Volume-to-LUN mapping lets you define which host group or host has access to a particular logical drive in the storage subsystem.

After you create a logical VolumeCopy, the target logical drive automatically becomes read-only to hosts to make sure that the data is preserved. Hosts that have been mapped to a target logical drive do not have write access to the logical drive, and any attempt to write to the read-only target logical drive results in a host I/O error.

If you want hosts to have write access to the data on the target logical drive, use the **set VolumeCopy** command to disable the Read-Only attribute for the target logical drive.

#### FlashCopy Logical Drives

A FlashCopy logical drive is a point-in-time image of a logical drive. It is usually created so that an application, such as a backup application, can access the FlashCopy logical drive and read the data while the base logical drive stays online and is accessible to hosts.

The logical drive for which the point-in-time image is created is known as the base logical drive and must be a standard logical drive in the storage subsystem. The FlashCopy repository logical drive stores information about all of the data that changed since the FlashCopy was created.

**Tip:** You can select a FlashCopy logical drive as the source logical drive for a VolumeCopy. This selection is a good use of this feature, because it performs complete backups without significant impact to the storage subsystem availability. Some I/O processing resources are lost to the copy operation.
**Important:** If you choose the base logical drive of a FlashCopy logical drive as your target logical drive, you must disable all of the FlashCopy logical drives that are associated with the base logical drive before you can select it as a target logical drive.

When you create a FlashCopy logical drive, a FlashCopy repository logical drive is automatically created. The FlashCopy repository logical drive stores information about the data that has changed since the FlashCopy logical drive was created. You cannot select a FlashCopy repository logical drive as a source logical drive or a target logical drive in a VolumeCopy.

You can use the FlashCopy Logical Drive premium feature with the VolumeCopy premium feature to back up data on the same storage subsystem and to restore the data on the FlashCopy logical drive back to its original base logical drive.

### Enhanced Remote Mirroring

The Enhanced Remote Mirroring premium feature provides for online, real-time replication of data between storage subsystems over a remote distance. In the event of a disaster or a catastrophic failure of one storage subsystem, you can promote a secondary storage subsystem to take over responsibility for data storage.

When you create a remote mirror, a remote mirror pair is created, which consists of a primary logical drive at the primary storage subsystem and a secondary logical drive at a remote storage subsystem.

The primary logical drive is the logical drive that accepts host I/O and stores data. When the mirror relationship is initially created, data from the primary logical drive is copied in its entirety to the secondary logical drive. This process is known as a full synchronization and is directed by the controller owner of the primary logical drive. During a full synchronization, the primary logical drive remains fully accessible for all normal I/O activity.

The controller owner of the primary logical drive starts remote writes to the secondary logical drive to keep the data on the two logical drives synchronized. Whenever the data on the primary logical drive and the secondary logical drive becomes unsynchronized, the controller owner of the primary logical drive starts a resynchronization, where only the data that changed during the interruption is copied.

The secondary logical drive maintains a mirror of the data on its associated primary logical drive. The controller owner of the secondary logical drive receives remote writes from the controller owner of the primary logical drive but does not accept host write requests.

The secondary logical drive stays available to host applications as read-only while mirroring is underway. In the event of a disaster or a catastrophic failure at the primary site, you can perform a role reversal to promote the secondary logical drive to a primary role. Hosts are then able to access the newly promoted logical drive, and business operations can continue.

You can select a primary logical drive that is participating in a remote mirror pair to be used as the source logical drive or a target logical drive for a VolumeCopy. A secondary logical drive that is participating in a remote-mirror pair cannot be selected as a source logical drive or a target logical drive.

#### **Role Reversals**

A role reversal is the act of promoting the secondary logical drive to be the primary logical drive of the remote mirror pair, and demoting the primary logical drive to be the secondary logical drive.

In the event of a disaster at the storage subsystem that contains the primary logical drive, you can fail over to the secondary site by performing a role reversal to promote the secondary logical drive to the primary logical drive role. This action lets hosts continue to access data, and business operations can continue.

Trying a role reversal in which the original primary logical drive is the source logical drive for an active VolumeCopy (the status is In Progress or Pending) causes the VolumeCopy to fail. The failure occurs when the original primary logical drive becomes the new secondary logical drive. You cannot restart the VolumeCopy until you return the roles of the logical drives back to their original state; that is, the logical drive that was originally the primary logical drive is set once again to be the primary logical drive.

If the primary storage subsystem is recovered but is unreachable due to a link failure, a forced promotion of the secondary logical drive will result in both the primary logical drive and the secondary logical drive viewing themselves in the primary logical drive role (dual-primary condition). If this condition occurs, the VolumeCopy in which the primary logical drive is participating is unaffected by the role change.

You can perform a role reversal by using the **set remoteMirror** command. (For information about the **set remoteMirror** command, see "Set Remote Mirror" on page 278.)

• To change a secondary logical drive to a primary logical drive, use this command, which promotes the selected secondary logical drive to become the primary logical drive of the remote mirror pair. Use this command after a catastrophic failure has occurred.

set remoteMirror role=primary

• To change a primary logical drive to a secondary logical drive, use this command, which demotes the selected primary logical drive to become the secondary logical drive. Use this command after a catastrophic failure has occurred.

set remoteMirror role=secondary

# Chapter 10. Using the Performance Read Cache premium feature

**Note:** The Performance Read Cache premium feature is available only on the DS3500 storage subsystem controller, the DCS3700 storage subsystem with Performance Module Controllers, and the DCS3860 storage subsystem.

The Performance Read Cache premium feature provides a way to improve read-only performance. Performance Read Cache is a set of Solid-State Disk (SSD) drives that you logically group together in your storage subsystem to implement a read cache for end-user volumes. Performance Read Cache is a secondary cache for use with the primary cache in the controller DRAM. In controller cache, the data is stored in DRAM after a host read. In Performance Read Cache, the data is copied from user-specified base volumes and stored on two internal RAID volumes (one per controller) that are automatically created when you create an Performance Read Cache. These volumes are used for internal cache processing purposes. These volumes are not accessible or displayed in the user interface. However, these two volumes do count against the total number of volumes allowed in the storage subsystem.

**Note:** The Performance Read Cache feature is available only when the Performance Read Cache premium feature is enabled, an Performance Read Cache entity is created, and each base volume for which you want to use the Performance Read Cache has the Performance Read Cache attributes enabled.

Using high performance SSDs to cache the data from the base volumes improves the application I/O performance and response times, and delivers sustained performance improvement across different workloads, especially for high-IOP workloads. Simple volume I/O mechanisms are used to move data to and from the Performance Read Cache. After data is cached and stored on the SSDs, subsequent reads of that data are performed on the Performance Read Cache, thereby eliminating the need to access the base volume.

These are the characteristics of a workload that benefits from using Performance Read Cache:

- Performance is limited by hard disk drive IOPs.
- There are a high percentage of reads relative to writes.
- A large number of reads that are repeat reads to the same or adjacent areas of the drive.
- The capacity of the data that is repeatedly accessed is smaller than the cache capacity. The more volumes being cached, the higher the probability that the capacity of the data accessed exceeds the capacity of the cache.

The Performance Read Cache premium feature moves data from the hard drives to SSDs following a host read or write so that a subsequent host read of the same logical block addressing (LBA) can be read directly from SSD with a much lower response time than rereading from hard drives.

When you create an Performance Read Cache, you choose which I/O type (file system, database, or web server) that most closely matches the applications and volumes that will use the Performance Read Cache. In addition, you specify the capacity of the Performance Read Cache from a list of possible candidates

consisting of different counts of SSD drives. You also have the option to enable Performance Read Cache on all eligible volumes that are currently mapped to hosts. Lastly, after you create the Performance Read Cache, you can enable or disable it on existing volumes or as part of a new volume creation.

#### Performance Read Cache operations

The following list includes common tasks that you can perform with the Performance Read Cache feature.

- Create
- Performance Modeling
- Locate
- · View Associated Physical Components
- Add Drives (Capacity)
- Remove Drives (Capacity)
- Suspend
- Resume
- Delete
- Enabling/Disabling Performance Read Cache
- Change I/O Type
- Rename

#### Performance Read Cache restrictions

- Performance Read Cache is not supported on Enhanced FlashCopys (legacy) volumes and Enhanced FlashCopy images.
- If you import or export base volumes that are Performance Read Cache enabled or disabled, the cached data is not imported or exported.
- The maximum usable Performance Read Cache capacity on a storage subsystem is dependent on the controller's primary cache capacity.
- When you create an Performance Read Cache, T10PI is automatically enabled if all the drives are T10PI-capable and the T10PI premium feature is enabled. T10PI cannot be disabled on the Performance Read Cache, nor can you add non-T10PI capable drives to a T10PI-enabled Performance Read Cache.
- You cannot remove the last drive in an Performance Read Cache without first deleting the Performance Read Cache.
- Only one Performance Read Cache is supported per storage subsystem.

# Creating the Performance Read Cache, adding volumes, and removing volumes

Before you create the Performance Read Cache, make sure that suitable SSD drives are available on the storage subsystem. You can achieve the best performance when the working set of the data fits in the Performance Read Cache so that most host reads can be serviced from the lower latency solid state disks instead of the higher latency hard drives (HDDs).

To create the Performance Read Cache run this command: create performanceReadCache userLabel="performanceReadCacheName"drives=(enclosureID1,drawerID1,slot ID1 ... enclosureIDn,drawerIDn,slotIDn)[updateExistingVolumes=(TRUE

```
FALSE)
```

526

To use this command you need only to create a name for the SSD cache and identify the specific SSD drives that you want to include in the cache. After you run the command, all of the volumes in the storage subsystem can use the Performance Read Cache.

If you do not want all of the volumes in the storage subsystem to use the Performance Read Cache, set the updateExistingVolumes parameter to FALSE. When you create the Performance Read Cache without volumes, you can later assign specific volumes to the Performance Read Cache by running this command for each volume that you want to add:

set volume ["volumeName"] performanceReadCacheEnabled=TRUE

You can assign standard volumes, Enhanced FlashCopy volumes, or consistency group Enhanced FlashCopy volumes. You can assign only volume at a time to the Performance Read Cache.

If you do not want a specific volume to use the Performance Read Cache, you can remove the volume from the Performance Read Cache by running this command: set volume ["volumeName"] performanceReadCacheEnabled=FALSE

After you have created the Performance Read Cache, you can model the performance of the Performance Read Cache to make sure that the Performance Read Cache is running as required for you storage subsystem

# Performance Read Cache performance modeling

After you create the Performance Read Cache, you can run a performance modeling tool to help determine the best cache size for your storage subsystem.

The performance modeling tool helps you determine the performance improvement for Performance Read Cache capacity by running a workload on the SSD cache that has the same characteristics as the workload that you run in production. Performance modeling monitors and measures I/O activity for a period of time and estimates performance for various Performance Read Cache sizes. The performance modeling tool then shows the actual performance for the Performance Read Cache that you created.

Depending on the cache capacity and workload, it might take 10 to 20 hours to fully populate the cache. The cache has valid information even after a run of a few minutes, but the longer that you can run the performance modeling tool the more accurate the results of the metric data.

The performance modeling tool provides an estimate of performance using these metrics:

- Cache hit percentage
- Average response time

**Note:** Performance modeling does not survive a controller reboot.

#### Starting and Stopping Performance Read Cache performance modeling

To start a performance modeling operation, use this command:

start performanceReadCache [performanceReadCacheName] performanceModeling

Enclose the identifier in square brackets ([]). If the Performance Read Cache name contains special characters or consists only of numbers, you also must enclose the identifier in double quotation marks ("") inside square brackets.

After you start the performance modeling operation it will continue to run until you stop the operation. To stop the performance modeling operation, use this command:

stop performanceReadCache [performanceReadCacheName] performanceModeling

This command immediately stops collecting the performance modeling data and displays the data that has been collected. The data is displayed in the same window from which you ran the command. If you prefer, you can save the performance modeling data to a file by running this command:

stop performanceReadCache [performanceReadCacheName] performanceModeling
file="fileName"

fileName is the file path and the file name to which you want to save the performance modeling data. Enclose the file name in double quotation marks (" "), as shown in this example:

file="C:\Program Files\CLI\logs\performance.csv"

The default name of the file that contains the performance modeling data is readLinkStatus.csv. You can use any file name, but you must use the .csv extension.

The performance modeling metric data are not available until you stop the performance modeling operation.

You can see the performance modeling data only after you have stopped the performance modeling operation. You cannot get intermediate results while the modeling is being run. When you stop the performance modeling tool, a graphical representation of the results appears, but you can view the results in tabular format by saving the data to a .csv file. To compare different results, you can run several performance modeling operations and save each result to a .csv file for each result.

Then you can use a spreadsheet program outside of the storage management software to compare the data from the .csv file. The performance modeling tool does not support the loading of saved files.

The .csv File Information

The .csv file shows the following information:

- Performance Read Cache Capacity (GB) The amount of Performance Read Cache capacity.
- Performance Read Cache Hits (%) Derived from cache hits and total reads.
- Average Overall Response Time This value is calculated by the software and is not the result of averaging the response times for SSD Reads (Cache Hits), HDD Reads, and HDD Writes.
- Average Response Time (milliseconds) This value is the same as the Average Overall Response Time in the next paragraph.
- SSD Reads (Cache Hits) The total number of host reads of SSD cache-enabled volumes that were satisfied from the Performance Read Cache.

- Average Response Time The average response time of the SSD reads (Cache Hits).
- Percentage of I/Os An indication of the percentage of SSD reads (Cache Hits).
- HDD Reads The total number of host reads of Performance Read Cache-enabled volumes.

Compare the reads relative to writes. The reads need to be greater than the writes for effective Performance Read Cache operation. The greater the ratio of reads to writes the better the operation of the cache.

- Average Response Time The average response time of the HDD reads.
- Percentage of I/Os An indication of the percentage of HDD reads.
- HDD Writes The total number of host writes to Performance Read Cache-enabled volumes. Compare the reads relative to writes. The reads need to be greater than the writes for effective Performance Read Cache operation. The greater the ratio of reads to writes the better the operation of the cache.
- Average Response Time The average response time of the HDD Writes.
- Percentage of I/Os An indication of the percentage of HDD Writes.

#### Understanding Performance Read Cache performance modeling results

When you run the performance modeling tool with a workload that has the same characteristics as you run in production, the performance modeling tool collects this type of information:

- △Cache hit percentage The cache-hit percentage indicates the percentage of all read commands that find data in the Performance Read Cache for each of the cache capacities. For almost all workloads, a cache-hit percentage around 75 percent indicates that you have sufficient capacity. A cache-hit percentage much higher than this means that the workload performance is likely to be improved with more Performance Read Cache capacity. However, for some workloads, a cache-hit percentage around 50 percent or even lower indicates that you have sufficient capacity.
- Average response time The performance modeling tool uses calculated response times from the operation of the Performance Read Cache to estimate the time it takes to run during each external and internal I/O operation. The tool uses these response time measurements along with measurements of I/O operations performed during the run of a workload that is applied to the base volumes with Performance Read Cache enabled. Keep in mind that workload on other volumes can affect these results due to increased response time. The resulting estimates are of the average response time for external I/O operations. When making a decision to increase or decrease the capacity of your current Performance Read Cache, look for an acceptable response time, and then compare that with the response time of your current Performance Read Cache capacity.

You can make changes by using options outside of the performance modeling tool. For example, you can change the I/O characteristic types and enable or disable Performance Read Cache on base volumes. These two parameters control the working set for the performance modeling. You change the I/O characteristic types by selecting these menu options: **Storage >> Performance Read Cache >> Change I/O Type**. You enable or disable Performance Read Cache on base volumes by selecting these menu options: **Storage >> Volume >> Performance Read Cache**.

#### Thrashing

Thrashing is a condition where data is constantly being read from the base volume into the Performance Read Cache, but not accessed again before it is cleared from the Performance Read Cache. The data is cleared from the SSD cache because all Performance Read Cache blocks are allocated and another volume or logical block addressing (LBA) range needs to allocate a cache block that currently had data. The Performance Read Cache processes do not move data from the base volume into the Performance Read Cache if only one read occurs into a cache block; therefore, a purely random read workload does not cause thrashing and unnecessary overhead with any working set size.

To determine whether cache thrashing might be occurring, look at the Recycle Actions statistic in the detailed statistics that you can view when you save the Performance Read Cache statistics to a ..csv file

### Performance Read Cache management tasks

As soon as you create the Performance Read Cache all of the volumes assigned to the Performance Read Cache can start using it. The Performance Read Cache is a non-addressable volume that requires little maintenance or management. You can, however, perform these tasks to make sure that the Performance Read Cache runs best to support your storage subsystem:

- View information about the drives, status, and capacity of the Performance Read Cache.
- Locate the drives that physically comprise the Performance Read Cache.
- Adding drives to and removing drives from the Performance Read Cache.
- Suspend and resume Performance Read Cache operation.
- Rename the Performance Read Cache.

#### Viewing information about the Performance Read Cache

When you create your Performance Read Cache, you provide a name and select the solid state disk (SSD) drives for the cache. However, the firmware maintains more information about the Performance Read Cache, such as maximum capacity, consumed capacity, size of the drives, and other information. You can see all of the information about the Performance Read Cache by running this command: show performanceReadCache [performanceReadCacheName]

This command returns information about the Performance Read Cache that is similar to this example.

```
Performance Read Cache name: my_cache
Status:
                                   Optimal
                                   Read Only
Type:
I/O characteristic type:
                                    File System
Maximum capacity allowed:
                                 1,862.645 GB
Current capacity:
                          557.792 GB
                               1,304.852 GB
Additional capacity allowed
Drive capacities: All 278.896 GB
Quality of Service (QoS) Attributes
 Security capable:
                           No
 Secure:
                      No
 T10PI capable:
                  No
Associated drives:
Enclosure
            Slot
0
      4
      11
0
Volumes using Performance Read Cache:
                                           volume_test
```

By reviewing this information, you can make sure that your Performance Read Cache is running as you want it to run. For example, if the current capacity is close to the maximum capacity, you might want to add another drive to your Performance Read Cache.

#### Adding drives to and removing drives from the Performance Read Cache

After you have created your Performance Read Cache, you might decide that you do not have enough capacity and want to add a drive. Conversely, you might decide that you have too much capacity and that you want to remove a drive to make the drive available for a Array or disk pool.

You can increase the capacity of an existing Performance Read Cache by using this command to add solid SSDs:

set performanceReadCache [performanceReadCacheName]
addDrives=(enclosureID1,drawerID1,slotID1 ...
enclosureIDn,drawerIDn,slotIDn)

You can add one or more SSDs by specifying the location of the drives that you want to add. For high-capacity drive enclosures, specify the enclosure ID value, the drawer ID value, and the slot ID value for each SSD that you want to add. For low-capacity drive enclosures, specify the enclosure ID value and the slot ID value for each SSD that you want to add. Enclosure ID values are 0 to 99. Drawer ID values are 1 to 5. Slot ID values are 1 to 32. Enclose the enclosure ID values, the drawer ID values, and the slot ID values in parentheses.

If you want to remove drives, you can use this command:

```
set performanceReadCache [performanceReadCacheName]
removeDrives=(enclosureID1,drawerID1,slotID1 ...
enclosureIDn,drawerIDn,slotIDn)
```

You cannot remove all of the SSDs from the Performance Read Cache using this command; at least one SSD must remain in the Performance Read Cache. If you want to completely remove the Performance Read Cache, use this command: delete performanceReadCache [performanceReadCacheName]

When you run this command, all data in the Performance Read Cache is purged.

#### Locating drives in the Performance Read Cache

The drives that comprise the Performance Read Cache might be dispersed throughout the storage subsystem. For maintenance or other reasons, you might be required to find the drives in the Performance Read Cache. You can locate the drives in the Performance Read Cache using this command:

start performanceReadCache [performanceReadCacheName] locate

This command identifies the drives that are logically grouped together to form the Performance Read Cache by flashing the indicator lights on the drives in the Performance Read Cache. To turn off the indicator lights on the drives, use this command:

stop performanceReadCache locate

#### Suspending and Resuming Performance Read Cache Operation

At times you will need to perform maintenance on the drives in the Performance Read Cache. Such maintenance might include replacing non-optimal drives or upgrading drives. To perform maintenance on the Performance Read Cache, you will need to first suspend the operations of the cache. To suspend operations use this command:

suspend performanceReadCache [performanceReadCacheName]

This command temporarily stops caching for all of the volumes that are using the Performance Read Cache. While caching is stopped, host reads are serviced from the base volumes instead of from the Performance Read Cache.

After performing maintenance, you can restart the Performance Read Cache by using this command:

resume performanceReadCache [performanceReadCacheName]

#### **Renaming the Performance Read Cache**

If you want to change the name of the Performance Read Cache, you can use this command:

set performanceReadCache [old\_performanceReadCacheName]
userLabel="new\_performanceReadCacheName"

The old Performance Read Cache name must be enclosed in square brackets. The new Performance Read Cache name, however, must be enclosed in double quotation marks.

# Chapter 11. Maintaining a storage subsystem

Maintenance covers a broad spectrum of activity with the goal of keeping a storage subsystem operational and available to all hosts. This chapter provides descriptions of commands you can use to perform storage subsystem maintenance. The commands are organized into four sections:

- Routine maintenance
- Performance tuning
- Troubleshooting and diagnostics
- Recovery operations

The organization is not a rigid approach, and you can use the commands as appropriate for your storage subsystem. The commands listed in this chapter do not cover the entire subsystem of commands you can use for maintenance. Other commands, particularly the **set** commands, can provide diagnostic or maintenance capabilities.

# **Routine Maintenance**

Routine maintenance involves those tasks that you might perform periodically to make sure that the storage subsystem is running as well as possible or to detect conditions before they become problems.

### **Running a Media Scan**

Media scan provides a way of detecting disk drive media errors before they are found during a normal read from or write to the disk drives. Any media scan errors that are detected are reported to the Event Log. The Event Log provides an early indication of an impending disk drive failure and reduces the possibility of encountering a media error during host operations. A media scan is performed as a background operation and scans all data and redundancy information in defined user logical drives.

A media scan runs on all of the logical drives in the storage subsystem that have these conditions:

- Has Optimal status
- Has no modification operations in progress
- Has media scan enabled

Errors that are detected during a scan of a user logical drive are reported to the Major Event Log (MEL) and handled as follows:

- Unrecovered media error The disk drive could not read the requested data on its first try or on any subsequent retries. The result of this action is that for logical drives with redundancy protection, the data is reconstructed, rewritten to the disk drive, and verified, and the error is reported to the Event Log. For logical drives without redundancy protection, the error is not corrected, but it is reported to the Event Log.
- **Recovered media error** The disk drive could not read the requested data on its first attempt. The result of this action is that the data is rewritten to the disk drive and verified. The error is reported to the Event Log.

- **Redundancy mismatches** Redundancy errors are found, and a media error is forced on the block stripe so that it is found when the disk drive is scanned again. If redundancy is repaired, this forced media error is removed. The result of this action is that the first 10 redundancy mismatches found on a logical drive are reported to the Event Log.
- Unfixable error The data could not be read, and parity information or redundancy information could not be used to regenerate it. For example, redundancy information cannot be used to reconstruct data on a degraded logical drive. The result of this action is that the error is reported to the Event Log.

The script command set provides two commands to define media scan properties:

- set logicalDrive
- set storageSubsystem

The **set logicalDrive** command enables a media scan for the logical drive. The command takes this form:

```
set (allLogicalDrives | logicalDrive [logicalDriveName] | logicalDrives
[logicalDriveName1 ... LogicalDriveNameN] | logicalDrive <wwID>)
mediaScanEnabled=(TRUE | FALSE)
```

The **set storageSubsystem** command defines how frequently a media scan is run on a storage subsystem. The command takes this form:

```
set storageSubsystem mediaScanRate=(disabled | 1-30)
```

The **mediaScanRate** values define the number of days over which the media scan runs. Valid values are disabled, which turns off the media scan; or 1 day to 30 days, where 1 day is the fastest scan rate, and 30 days is the slowest. A value other than what is shown will not allow the media scan to function.

### **Running a Redundancy Check**

Redundancy checks are performed when media scans are run. (For a description about how to set up and run media scans, see "Running a Media Scan" on page 533.) During a redundancy check, all of the data blocks in a logical drive are scanned, and, depending on the RAID level, deteriorated data is corrected. Correction is performed as follows:

- For RAID 3, RAID 5, or RAID 6 logical drives, redundancy is checked and repaired.
- For RAID 1 logical drives, the data is compared between the mirrored disk drives and data inconsistencies are repaired.
- RAID 0 logical drives have no redundancy.

Before you can run a redundancy check, you must enable redundancy checking by using the **set logicalDrive** command. The command takes this form:

```
set (allLogicalDrives | logicalDrive [logicalDriveName] | logicalDrives
[logicalDriveName1 ... logicalDriveNameN] | logicalDrive <wwID>)
redundancyCheckEnabled=(TRUE | FALSE)
```

# **Resetting a Controller**

**Important:** When you reset a controller, the controller is no longer available for I/O operations until the reset is complete. If a host is using logical drives that are owned by the controller being reset, the I/O that is directed to the controller is

rejected. Before resetting the controller, either make sure that the logical drives that are owned by the controller are not in use, or make sure that a multi-path driver is installed on all of the hosts that are using these logical drives.

Resetting a controller is the same as rebooting the controller processors. To reset a controller, use this command: reset controller [(a | b)]

# **Enabling a Controller Data Transfer**

At times, a controller might become quiescent while running diagnostics. If this condition occurs, the controller might become unresponsive. To revive a controller that has become quiescent while running diagnostics, use this command: enable controller [(a | b)] dataTransfer

# **Resetting the Battery Age**

After you have replaced the batteries in the storage subsystem, you must reset the age of the battery. You can reset either the batteries for an entire storage subsystem or a battery in a specific controller. To reset the age of the batteries to zero days, use this command:

reset storageSubystem batteryInstallDate [controller=(a | b)]

**Note:** The DS3950, DS4800 (models 82, 84, and 88), DS5020, DS5100, and DS5300 controller modules do not support this command. The batteries in the DS3950, DS4800 (models 82, 84, and 88), DS5020, DS5100, and DS5300 controller modules do not require that you reset the battery age after you have replaced the batteries.

# **Removing Persistent Reservations**

Persistent reservations preserve logical drive registrations, and they prevent hosts, other than the host defined for the logical drive, from accessing the logical drive. You must remove persistent reservations before you make these changes to your configuration:

- Change or delete LUN mappings on a logical drive holding a reservation
- · Delete subsystems or logical drives that have any reservations

To determine which logical drives have reservations, use this command:

show (allLogicalDrives | logicalDrive [logicalDriveName] | logicalDrives
[logicalDriveName1 ... logicalDriveNameN]) reservations

To clear persistent logical drive reservations, use this command:

clear (allLogicalDrives | logicalDrive [logicalDriveName] | logicalDrives
[logicalDriveName1 ... logicalDriveNameN]) reservations

# Synchronizing the Controller Clocks

To synchronize the clocks on both controllers in a storage subsystem with the host clock, use this command:

set storageSubsystem time

# **Locating Disk Drives**

At times, you might need to locate a specific disk drive. In very large storage subsystem configurations, this task can sometimes be awkward. If you need to

locate a specific disk drive, you can do so by turning on the indicator light on the front of the disk drive. To locate a disk drive, use this command: start drive [*enclosureID*,*slotID*] locate

To turn off the indicator light after locating the disk drive, use this command: stop drive locate

### Relocating an subsystem

subsystem relocation describes the action of moving disk drives within the same storage subsystem. This is a supported capability; however, any relocation of storage subsystem components must be completed under the guidance of an IBM Technical Support representative.

This section describes the commands that you use to remove a set of disk drives and then reinstall them into a different storage subsystem.

#### Hot and Cold subsystem Relocation

There are two methods you can use to move subsystems: hot subsystem relocation and cold subsystem relocation.

**Note:** For more detailed information about relocating subsystems, see one of the following documents:

- DS3000: IBM System Storage DS3000 Installation, User's and Maintenance Guide
- **DS4000 / DS5000**: IBM System Storage DS4000/DS5000 Hard Drive and Storage Expansion Enclosure Installation and Migration Guide
- Hot subsystem relocation lets you add or move storage without reconfiguring the storage subsystem and, in some cases, without rebooting. During hot subsystem relocation, the storage subsystem power is not turned off.
- Cold subsystem relocation requires that the power to the source storage subsystem and the destination storage subsystem be turned off before moving the subsystems from one storage subsystem to another. Then the power to the storage subsystems can be turned on.

To ensure that any subsystem being moved to a different destination storage subsystem is correctly recognized and managed by the new storage subsystem, use hot subsystem relocation whenever possible.

**Attention: Possible loss of data access** – You must move a single subsystem at a time, and it must go into a storage subsystem with the same level of controller firmware.

#### **Basic Process Steps**

Relocating an subsystem from one storage subsystem to another includes these procedures:

- 1. Verifying the status of the storage subsystems
- **2.** Ensuring that both storage subsystems have the latest controller firmware versions installed
- 3. Locating the drives in the subsystem
- 4. Placing the subsystem offline
- 5. Removing drives from the storage subsystem
- 6. Replacing an subsystem into the new storage subsystem

**Attention:** Check firmware versions – Ensure that both existing and new storage subsystems have the latest controller firmware versions installed.

• See the IBM Disk Support website to find the latest firmware versions for your storage subsystems:

www.ibm.com/systems/support/storage/disk

- Also check the following documentation for more information:
  - DS3000: IBM System Storage DS3000 Installation, User's and Maintenance Guide
  - **DS4000 / DS5000**: IBM System Storage DS4000/DS5000 Hard Drive and Storage Expansion Enclosure Installation and Migration Guide

To perform these steps, you must be familiar with the following CLI commands. The command syntax is provided to assist in your use of these new commands.

#### subsystem Relocation Commands

**Note:** subsystem relocation commands are supported with the following controller firmware only:

- DS5000: 07.30.xx.xx
- DS4000: 07.1x.xx.xx
- DS3000: Not supported.

Use the following command to place a specific storage subsystem into an exported state so that its disk drives may be removed.

start subsystem [user-label] export

At this point you are allowed to remove the disk drives that comprise the subsystem, and physically reinstall them into a different storage subsystem.

Use the following command to logically move a specific storage subsystem from an exported state to the complete state. start subsystem [user-label] import

Your relocated subsystem is now available for use.

For additional information, refer to these commands:

- "Show subsystem Export Dependencies" on page 294
- "Show subsystem Import Dependencies" on page 294
- "Start subsystem Export" on page 331
- "Start array import" on page 332

# **Performance Tuning**

Over time, as a storage subsystem exchanges data between the hosts and the disk drives, its performance can degrade. You can monitor the performance of a storage subsystem and make adjustments to the operational settings on the storage subsystem to help improve performance.

### Monitoring the Performance

You can monitor the performance of a storage subsystem by using the **save storageSubsystem performanceStats** command. This command saves performance information to a file that you can review to help determine how well the storage subsystem is running. The following table lists the performance information that is saved to the file.

Table 25. Information About Storage Subsystem Performance

Type of Information	Description		
Devices	These devices are included in the file:		
	• <b>Controllers</b> – The controller in slot A or slot B and a list of the logical drives that are owned by the controller		
	• Logical drives – A list of the logical drive names		
	• <b>Storage subsystem totals</b> – A list of the totals for both controllers in an active/active controller pair, regardless if one, both, or neither are selected for monitoring		
Total I/Os	The number of total I/Os performed since the storage subsystem was started		
Read Percentage	The percentage of total I/Os that are read operations (calculate the write percentage by subtracting the read percentage from 100 percent)		
Cache Hit Percentage	The percentage of reads that are fulfilled by data from the cache rather than requiring an actual read from a disk drive		
Current <sup>®</sup> KB per second	The current transfer rate in kilobytes per second (current means that the number of kilobytes per second since the last time that the polling interval elapsed, causing an update to occur)		
Maximum KB per second	The highest data transfer value that is achieved in the current kilobyte-per-second statistic block		
Current I/O per second (IOPS)	The current number of I/Os per second (current means the number of I/Os per second since the last time that the polling interval elapsed, causing an update to occur)		
Maximum I/O per second	The highest number of I/Os achieved in the current I/O-per-second statistic block		

The command takes this form:

save storageSubsystem performanceStats file="filename"

where *filename* is the name of the file in which you want to save the performance statistics. You can use any file name that your operating system can support. The default file type is .csv. The performance information is saved as a comma-delimited file.

Before you use the save storageSubsystem performanceStats command, run the set session performanceMonitorInterval command and the set session performanceMonitorIterations command to specify how often statistics are collected.

### Changing the RAID Levels

When you create an subsystem, you can define the RAID level for the logical drives in that subsystem. You can change the RAID level later to improve performance or provide more secure protection for your data.

RAID 6 is supported in DS4000 and DS5000 storage products only, with controller firmware version 7.1x.xx.xx or higher.

To change the RAID level, use this command: set subsystem [subsystemNumber] raidLevel=(0 | 1 | 3 | 5 | 6)

where *subsystemNumber* is the number of the subsystem for which you want to change the RAID level.

# **Changing the Segment Size**

When you create a new logical drive, you can define the segment size for that logical drive. In addition, you can change the segment size later to optimize performance. In a multiuser database or file system storage environment, set your segment size to minimize the number of disk drives that are needed to satisfy an I/O request. Use larger values for the segment size. Using a single disk drive for a single request leaves other disk drives available to simultaneously service other requests. If the logical drive is in a single-user large I/O environment, performance is maximized when a single I/O request is serviced with a single data stripe; use smaller values for the segment size. To change the segment size, use this command:

```
set logicalDrive ([logicalDriveName] | <wwID>) segmentSize=segmentSizeValue
```

where *segmentSizeValue* is the new segment size that you want to set. Valid segment size values are 8, 16, 32, 64, 128, 256, and 512. You can identify the logical drive by name or by WWID. (For usage information, see "Set Logical Drive Attributes" on page 260.)

# **Changing the Cache Parameters**

The script command set provides two commands that you can use to change cache parameter settings:

- set storageSubsystem
- set logicalDrive

The set storageSubsystem command lets you change settings for these items:

- The cache block size
- The cache flush start percentage
- The cache flush stop percentage

The **set logicalDrive** command lets you change settings for these items:

- The cache flush modifier
- The cache without batteries enabled or disabled
- The mirror cache enabled or disabled
- The read cache enabled or disabled
- The write cache enabled or disabled
- The read ahead multiplier
- The redundancy check enabled or disabled

# Defragmenting an subsystem

When you defragment an subsystem, you consolidate the free capacity in the subsystem into one contiguous area. Defragmentation does not change the way in which the data is stored on the logical drives. As an example, consider an subsystem with five logical drives. If you delete logical drive 1 and logical drive 3, your subsystem is configured as follows:

space, logical drive 2, space, logical drive 4, logical drive 5, original unused space

When you defragment this subsystem, the space (free capacity) is consolidated into one contiguous location after the logical drives. After being defragmented, the subsystem appears as follows:

logical drive 2, logical drive 4, logical drive 5, consolidated unused space

To defragment an subsystem, use this command: start subsystem [*subsystemNumber*] defragment

where *subsystemNumber* is the identifier for the subsystem.

**Note:** If you need to defragment the data stored in the logical drive, you must use the appropriate defragmenting tool for the installed operating system.

### **Troubleshooting and Diagnostics**

If a storage subsystem exhibits abnormal operation or failures, you can use the commands that are described in this section to help determine the cause of the problem.

### **Collecting All Support Data**

To gather the most comprehensive information about a storage subsystem, run the **save storageSu bsystem supportData** command. This command collects data for remote troubleshooting and analysis of problems with the storage management software. All of the files gathered are compressed into a single archive in a zipped file format. The following table lists the type of support data that is collected.

Type of Data	Description and File Name
Storage subsystem profile	A list of all components and properties of a storage subsystem. storageSu bsystemProfile.txt
Major Event Log	A detailed list of errors that occur on the storage subsystem. The list is stored in reserved areas on the disk drives in the storage subsystem. The list records configuration events and failures with storage subsystem components. majorEventLog.txt
Read link status	A detailed list of errors that have been detected in the traffic flow between the devices on the Fibre Channel loop. A file of historical read link status data might also be included in the archive. readLinkStatus.csv
Switch-on-a-chip (SOC) error statistics	Information from the loop-switch ports that are connected to Fibre Channel devices. socStatistics.csv

Table 26. Support Data for the Storage Subsystem

	Description and The Name
NVSRAM	A controller file that specifies the default settings for the controllers.
	NVSRAMdata.txt
Performance statistics	A detailed description of how a storage subsystem is performing. Collected data includes the I/O activity of specific controllers or logical drives, the transfer rate of the controller, the current I/Os per second, and the maximum I/Os per second.
	performanceStatistics.csv
Persistent reservations and persistent registrations	A detailed list of logical drives on the storage subsystem and persistent reservations and persistent registrations.
	persistentRegistrations.txt
Object bundle	A detailed description of the status of the storage subsystem and its components, which was valid at the time that the file was generated. The object bundle file is a binary file and does not contain human-readable information.
	objectBundle
Drive diagnostic data	A detailed list of log sense data from all of the disk drives in the storage subsystem.
	driveDiagnosticData.txt
Recovery profile	A detailed description of the latest recovery profile record and historical data.
	recoveryProfile.csv
Unreadable sectors	A detailed list of all of the unreadable sectors that have been logged to the storage subsystem.
	badBlocksData.txt
State capture data	A detailed description of the current state of the storage subsystem.
	stateCaptureData.dmp
Storage subsystem	A detailed listing of the hardware components and the software components that comprise the storage subsystem configuration. storageSubsystemConfiguration.cfg

Table 26. Support Data for the Storage Subsystem (continued)

# **Collecting Drive Data**

To gather information about all of the disk drives in a storage subsystem, use the **save allDrives** command. This command collects sense data and saves the data to a file. The sense data consists of statistical information that is maintained by each of the disk drives in the storage subsystem.

# **Diagnosing a Controller**

The **diagnose controller** command provides these tests that help you make sure that a controller is functioning correctly:

- The read test
- The write test
- The data-loopback test

The read test initiates a read command as it would be sent over an I/O data path. The read test compares data with a known, specific data pattern, and the read test checks for data integrity and errors. If the read command is unsuccessful or the data compared is not correct, the controller is considered to be in error and is placed offline.

The write test initiates a write command as it would be sent over an I/O data path to the diagnostics region on a specified disk drive. This diagnostics region is then read and compared to a specific data pattern. If the write fails or the data compared is not correct, the controller is considered to be in error, and it is failed and placed offline.

Run the data-loopback test only on controllers that have connections between the controller and the disk drives. The test passes data through each controller disk drive-side channel, the minihub, out onto the loop, and back again. Enough data is transferred to determine error conditions on the channel. If the test fails on any channel, this status is saved so that it can be returned if all of the other tests pass.

For best results, run all three tests after you first install the storage subsystem and any time that you that have made changes to the storage subsystem or the components that are connected to the storage subsystem (such as hubs, switches, and host adapters).

A custom data pattern file called diagnosticsDataPattern.dpf is included in the root directory of the installation CD. You can modify this file, but the file must have these properties to work correctly for the tests:

- The file values must be entered in hexadecimal format (00 to FF) with only one space between the values.
- The file must be no larger than 64 bytes in size. Smaller files will work, but larger files can cause an error.

The test results contain a generic, overall status message and a set of specific test results. Each test result contains these items:

- Test (read, write, or data loopback)
- Port (read or write)
- Level (internal or external)
- Status (pass or fail)

Events are written to the Event Log when the diagnostics are started and when testing is completed. These events help you to evaluate whether diagnostics testing was successful or failed and the reason for the failure.

# **Running Read Link Status Diagnostics**

Read link status (RLS) error counts refer to link errors that have been detected in the traffic flow of a Fibre Channel loop. The errors detected are represented as a count (32-bit field) of error occurrences that are accumulated over time. The counts provide coarse measure of the integrity of the components and devices on the loop. By analyzing the error counts that are retrieved, you can determine the components or devices within the Fibre Channel loop that might be experiencing problems communicating with the other devices on the loop. A high error count for a particular component or device indicates that it might be experiencing problems and should be given immediate attention.

Error counts are calculated from the current baseline. The baseline describes the error count values for each type of device in the Fibre Channel loop, either when the controller goes through its start-of-day sequence or when you reset the baseline. The baseline indicates the difference in error counts from the time the baseline was established to the time you request the read link status data.

The script command set provides two commands for running RLS diagnostics:

- **reset storageSubsystem RLSBaseline** Resets the RLS baseline for all devices by setting all of the counts to 0.
- **save storageSubsystem RLSCounts** Saves the RLS counters to a file that you can review later. The default file name is readLinkStatus.csv.

Run the **reset storageSubsystem RLSBaseline** command before you run the **save storageSubsystem RLSBaseline** command.

The following table lists the type of data contained by the file that is generated by the **save storageSubsystem RLSBaseline** command.

Type of Data	Description
Devices	A list of all devices on the Fibre Channel loop. The devices appear in channel order. Within each channel, the devices are sorted according to the device position in the loop.
Baseline time	The date and time when the baseline was set.
Elapsed time	The time that has elapsed from when the baseline time was set to when the read link status was gathered.
Invalid transmission word (ITW)	The total number of ITW errors that were detected on the Fibre Channel loop from the baseline time to the current date and time. ITW might also be referred to as the Received Bad Character Count. ITW counts indicate that in decoding a read/write transmission, the mapping did not exist and the running disparity of the transmission word is invalid. This data is the key error count to be used when analyzing the error count data.

Table 27. RLS Baseline Data for the Storage Subsystem

Type of Data	Description
Link failure (LF)	The total number of LF errors that were detected on the Fibre Channel loop from the baseline time to the current date and time.
	An LF condition is either a link fault signal, a loss of signal, or a loss of synchronization condition. The LF signal indicates a failure with the media module laser operation.
Loss of synchronization (LOS)	The total number of LOS errors that were detected on the Fibre Channel loop from the baseline time to the current date and time.
	LOS errors indicate that the receiver cannot acquire symbol lock with the incoming data stream due to a degraded input signal. If this condition persists, the number of LOS errors increases.
Loss of signal (LOSG)	The total number of LOSG errors that were detected on the Fibre Channel loop from the baseline date to the current date and time.
	LOSG errors typically indicate a loss of signal from the transmitting node or the physical component within the Fibre Channel loop. Physical components where a loss of signal typically occurs include the gigabit interface converters (GBICs), the Small Form-factor Pluggable (SFP) transceivers, and the Fibre Channel fiber-optic cable.
Primitive sequence protocol (PSP)	The total number of PSP errors that were detected on the Fibre Channel loop from the baseline date to the current date and time. PSP refers to the number of NPort protocol errors that were detected and Link Reset Response (LRR) primitive sequences that were received while the link is up. An LRR is issued by another NPort in response to a link reset.
	An NPort is a Fibre Channel-defined port at the end of a link, such as a server or a workstation. Each port can act as an originator or a responder (or both) and contains a transmitter and receiver. Each port is given a unique name, called an NPort or an NLPort identifier. If an NPort is connected to a loop, it becomes an NLPort. An NLPort is a Fibre Channel controller ID in a hexadecimal number. The hexadecimal number varies depending on the topology:
	• For a private arbitrated loop, the ID is a 1-byte arbitrated loop physical address (ALPA).
	• For all other arbitrated loops, it appears as a single 24-bit hexadecimal number (a triplet of domain, area, and ALPA where each field is 1 byte).
	• For fabric and point-to-point, the ID is a 3-byte hexadecimal number used in the DID and SID (destination identifier and source identifier) fields of Fibre Channel frames.

Table 27. RLS Baseline Data for the Storage Subsystem (continued)

Table 27.	RLS	Baseline	Data	for	the	Storage	Subsystem	(continued)	)
-----------	-----	----------	------	-----	-----	---------	-----------	-------------	---

Type of Data	Description
Invalid cyclic redundancy check (ICRC)	The total number of ICRC errors that were detected on the Fibre Channel loop from the baseline date to the current date and time.
	An ICRC count indicates that a frame has been received with an invalid cyclic redundancy check value. A cyclic redundancy check reads the data, calculates the cyclic redundancy check character, and compares the calculated cyclic redundancy check character with a cyclic check character already present in the data. If they are equal, the new data is presumed to be the same as the old data. If the calculated characters and the old characters do not match, an error is posted, and the data is re-sent.

#### Interpreting the RLS Results

The way that you interpret the RLS results is based on the concept that the device immediately following the problematic component will have the largest number of invalid transition word (ITW) error counts. The process is to obtain the ITW count for every component and device on the loop, analyze the data in loop order, and identify any large increases in the ITW counts.

#### **Important:**

The current error counting standard for when to calculate the ITW error count is not well defined. Different vendor devices calculate at different rates. Analysis of the data must take this discrepancy into consideration.

# **Collecting Switch-on-a-Chip Error Statistics**

Switch-on-a-chip (SOC) error statistics provide information about the loop-switch ports that are connected to the Fibre Channel devices in a storage subsystem. (RLS counts provide information about the Fibre Channel devices.) Reporting SOC error statistics is available only on storage subsystems that have SOC loop-switch devices that are incorporated into the controller disk drive channel or the ESM circuitry. SOC devices are integrated circuits that join together Fibre Channel devices in arbitrated loop topologies. SOC devices automatically collect statistical information for each SOC port that is connected to a controller port, an ESM port, a disk drive port, or an expansion connector. An IBM Technical Support representative can use the statistical information with RLS counts to identify problems with Fibre Channel devices that are attached to the loop.

SOC error statics include this information:

- The port state
- The port insertion count
- The loop state
- The loop up count
- The CRC error count
- The relative frequency drift error average
- The loop cycle count
- The operating system (OS) error count

- The port connections attempted count
- · The port connections held off count
- The port utilization

The method for collecting error statistics starts by establishing a baseline for the SOC error statistics. The baseline consists of SOC error statistics that are established at a set time for each SOC device on the loop. The baseline is set by clearing the error counters in each SOC device. You can set a device baseline by performing one of these actions:

- Turning on the power to the device or resetting the device
- Running the reset storageSubsystem SOCBaseline command

In addition, each controller also initializes the SOC error counters in all of the expansion drawers that are attached to the controller following a cold boot (power-on or hot insertion). If you add an expansion drawer while the power is turned on to the storage subsystem, a new baseline is established for any device on the expansion drawer.

After you have established the baseline for the SOC devices, you run the storage subsystem for a predetermined amount of time (for example, two hours). At the end of the run time, you collect the SOC error statistics by saving the information to a file. To save the information, run the **save storageSubsystem SOCCounts file** *filename* command. The default name of the file that contains the SOC error statistics is socStatistics.csv. You can use any file name that has the .csv extension.

Analyzing the SOC error statistics is beyond the scope of normal storage subsystem management. After you have collected the SOC error statistics in a file, send the file to an IBM Technical Support representative.

### **Recovery Operations**

Recovery operations include repairing the storage subsystem and returning it to an operational state. This might involve replacing a failed CRU, a failed controller, or a disk drive or restoring data or the storage subsystem to operation. For information about when it is appropriate to replace a CRU, see "Routine Maintenance" on page 533.

### Setting the Controller Operational Mode

A controller has three operational modes:

- Online
- Offline
- Service

Placing a controller online sets it to the Optimal state and makes it active and available for I/O operations. Placing a controller offline makes it unavailable for I/O operations and moves its subsystems to the other controller if failover protection is enabled.

Taking a controller offline can seriously impact data integrity and storage subsystem operation.

• If you do not use write cache mirroring, data in the cache of the controller you place offline is lost.

• If you take a controller offline and you have controller failover protection through a host multi-path driver, the other controller in the pair takes over. subsystems and their associated logical drives that were assigned to the offline controller are automatically reassigned to the remaining controller. If you do not have a multi-path driver installed on the application host and you take a controller offline while the application is using associated logical drives, application errors will occur.

**Attention: Possible loss of data access** – Placing a controller offline can cause loss of data.

Use Service mode to replace CRUs, such as a controller. Placing a controller in Service mode makes it unavailable for I/O operations and moves its subsystems to the second controller without affecting the preferred path of the subsystem. This action might significantly reduce performance. The subsystems are automatically transferred back to the preferred controller when it is placed back online.

If you change a controller to Service mode while an application is using the associated logical drives on the controller, the change causes I/O errors unless a multi-path driver is installed on the host. Before you place a controller in Service mode, make sure that the logical drives are not in use, or a multi-path driver is installed on all of the hosts that are using these logical drives.

In addition, if you do not have a multi-path driver, you must make appropriate operating system-specific modifications to make sure that the subsystems moved are accessed on the new path when you change to Service mode.

**Important:** Place a controller in Service mode only under the direction of an IBM Technical Support representative.

To change the operational mode of a controller, use this command: set controller [(a | b)] availability=(online | offline | serviceMode)

# **Changing the Controller Ownership**

You can change which controller is the owner of a logical drive by using the **set logical drive** command. The command takes this form:

set(allLogicalDrives|logicalDrive [logicalDriveName]|logicalDrives
[logicalDriveName1...logicalDriveNameN]
llogicalDrive <wwID>)owner=(a|b)

# **Initializing a Drive**

Attention: Possible loss of data access – When you initialize a disk drive, all data on the disk drive is lost.

You must initialize a disk drive when you have moved a disk drive that was previously part of a multidisk subsystem from one storage subsystem to another. If you do not move the entire set of disk drives, the subsystem information and the logical drive information on the disk drives that you move are incomplete. Each disk drive that you move contains only part of the information that is defined for the logical drive and the subsystem. To be able to reuse the disk drives to create a new subsystem and logical drive, you must delete all of the old information from the disk drives by initializing the disk drive.

When you initialize a disk drive, all of the old subsystem information and logical drive information are deleted, and the disk drive is returned to an unassigned

state. Returning a disk drive to an unassigned state adds unconfigured capacity to a storage subsystem. You can use this capacity to create additional subsystems and logical drives.

To initialize a disk drive, use this command: start drive [*enclosureID*, *slotID*] initialize

where *enclosureID* and *slotID* are the identifiers for the disk drive.

### **Reconstructing a Drive**

If two or more of the disk drives in a subsystem have failed, the logical drive shows a status of Failed. All of the logical drives in the subsystem are no longer operating. To return the subsystem to an Optimal status, you must replace the failed disk drives. Then, you must reconstruct the data on the new disk drives. The data that you reconstruct is the data as it would appear on the failed disk drives.

**Important:** You can use this command only when the disk drive is assigned to a RAID 1, RAID 3, RAID 5, or RAID 6 subsystem.

To reconstruct a disk drive, use this command: start drive [*enclosureID*, *slotID*] reconstruct

where *enclosureID* and *slotID* are the identifiers for the disk drive.

### **Initializing a Logical Drive**

**Attention: Possible loss of data** – When you initialize a logical drive, all existing data on the logical drive and all of the information about the logical drive are destroyed. Do not use this command unless you want to destroy all the data in the given disk drive.

A logical drive is automatically initialized when you first create it. If the logical drive starts showing failures, you might be required to re-initialize the logical drive to correct the failure condition.

Consider these restrictions when you initialize a logical drive:

- You cannot cancel the operation after it begins.
- You cannot use this option if any modification operations are in progress on the logical drive or the subsystem.
- You cannot change the cache parameters of the logical drive while the initialization operation is in progress.

To initialize a logical drive, use this command: start logicalDrive [logicalDriveName] initialize

where *logicalDriveName* is the identifier for the logical drive.

### **Redistributing Logical Drives**

When you redistribute logical drives, you return the logical drives to their preferred controller owners. The preferred controller ownership of a logical drive or an subsystem is the controller of an active-active pair that is designated to own the logical drives. The preferred owner for a logical drive is initially designated when the logical drive is created. If the preferred controller is being replaced or undergoing a firmware download, ownership of the logical drives is automatically shifted to the other controller. That controller becomes the current owner of the logical drives. This change is considered to be a routine ownership change and is reported in the Event Log.

To redistribute logical drives to their preferred controllers, use this command: reset storageSubsystem logicalDriveDistribution

**Important:** If you run this command without a multi-path driver on the hosts, stop I/O activity to the logical drives to prevent application errors.

**Important:** You cannot run this command if all of the logical drives are currently owned by their preferred controller, or the storage subsystem does not have defined logical drives.

Under some host operating systems, you must reconfigure the multi-path host driver. You might also need to make operating system modifications to recognize the new I/O path to the logical drive.

### **Replacing CRUs**

The DS3950, DS3000, DS4000, DS5000, and DS5020 controller module components, such as the controller CRUs, the power-fan CRUs, and the interconnect-battery CRUs, have a Service Action Allowed (SAA) indicator light. This indicator light is a blue LED. The Service Action Allowed indicator light helps to make sure that you do not remove a CRU before it is safe to do so.

Attention: Possible loss of data access – Never remove a component that has a Service Action Required indicator light on unless the Service Action Allowed (SAA) indicator light is on.

If a component fails and must be replaced, the Service Action Required indicator light on that CRU comes on to indicate that service action is required, provided no data availability dependencies or other conditions exist that dictate the CRU should not be removed. The Service Action Allowed indicator light automatically comes on or goes off when conditions change. In most cases, the Service Action Allowed indicator light comes on steadily when the Service Action Required indicator light comes on for the CRU.

If the interconnect-battery CRU must be replaced, the Service Action Allowed indicator light does not come on automatically. Before the Service Action Allowed indicator light on the interconnect-battery CRU can come on, you must place controller CRU in slot B into Service mode. This action routes all control and I/O activity through one controller to help make sure that data access is maintained while the interconnect-battery CRU is removed. The Service Action Allowed indicator light comes on after the new CRU has been installed.

The ability to remove a CRU depends on the data availability dependencies of the controller module or the controller module. The Service Action Allowed indicator light does not come on if removing a CRU jeopardizes data on the expansion drawers or current I/O activity. An example of limiting when you can remove a CRU is when one controller CRU has a Service Action Required indicator light on. You cannot remove the other controller CRU (the Service Action Allowed indicator light does not come on), because doing so would jeopardize the data either on the expansion drawers or migrating through the controllers.

A less obvious example, as in the DS4800 storage subsystem, is when the power supply for the controller CRU in slot A has failed, and the controller CRU in slot B has failed. Removing the controller CRU in slot B before replacing the failed power-fan CRU causes the controller CRU in slot A to lose power, which results in a loss of data access. This action occurs because power distribution from each power-fan CRU is through the controller CRU that is physically connected to that power-fan CRU.

So, in the preceding example, these actions occur:

- The power-fan CRU has both its Service Action Required indicator light and its Service Action Allowed indicator light on.
- The controller CRU in slot B has only its Service Action Required indicator light on, but its Service Action Allowed indicator light is off.
- After the failed power-fan CRU has been replaced, the Service Action Allowed indicator light comes on for the controller CRU in slot B.

For more information about the Service Action Not Allowed indicator, please see the *IBM System Storage DSxxxx Storage Subsystem Installation, User's and Maintenance Guide* for your storage subsystem product.

# Chapter 12. Examples of Information Returned by the Show Commands

This chapter provides examples of information that is returned by the **show** commands. These examples show the type of information and the information detail. This information is useful in determining the components, features, and identifiers that you might need when you configure or maintain a storage subsystem.

### Show Storage Subsystem Command

The **show storageSubsystem** command returns information about the components and the features in a storage subsystem. If you run the command with the **profile** parameter, the command returns information in the form shown by this example. This information is the most detailed report that you can receive about the storage subsystem. After you have configured a storage subsystem, save the configuration description to a file as a reference.

PROFILE FOR STORAGE SUBSYSTEM: DS5300\_middle (1/27/09 3:23:25 PM)

```
SUMMARY-----
  Number of controllers: 2
  High performance tier controllers: Enabled
  Number of subsystems: 5
  RAID 6: Enabled
  Total number of logical drives used:
                                          10
     Number of standard logical drives:
                                          9
     Number of access logical drives:
                                          1
  Total number of logical drives allowed: 2048
                                                           Enabled
  FlashCopy Logical Drives:
     Number of flashcopies used:
                                                           0
     Number of flashcopies allowed:
                                                           2
     Number of flashcopies allowed per base logical drive:
                                                           2
                                         Disabled/Deactivated
  Enhanced Remote Mirroring:
     Number of mirrors used:
                                 0
     Number of mirrors allowed: 0
  VolumeCopy:
                                Disabled
     Number of copies used:
                                0
     Number of copies allowed: 0
  Number of drives:
                           32
  Mixed drive types:
                           Enabled
  Current drive type(s): Fibre (5), Fibre/FDE (21), Serial ATA (SATA) (6)
  Total hot spare drives: 2
     Standby:
                           2
```

In use: 0	)
Full Disk Encryption: Enab Security key identifier: 2	01ed 27000000600A0B800029ECE60000DE9E4940E373
Number of drive enclosures Number of drive enclosures	s: 2 s allowed: 28
Storage Partitioning: Number of partitions us Number of partitions al Number of logical drives a	Enabled sed: 2 lowed: 2 allowed per partition: 256
Access logical drive: LUN Default host OS: Wir (Ho	V 31,31,31 (see Mappings section for details) ndows 2000/Server 2003/Server 2008 Non-Clustered ost OS index 2)
Current configuration Firmware version: NVSRAM version: EMW version: AMW version: Pending configuration Staged firmware downloa Firmware version: NVSRAM version: Transferred on: Controller enclosure audib	07.50.00.00 N1818D53R1050V03 10.50.G5.02 10.50.G5.02 None None None None Disabled
NVSRAM configured for batt	ceries: Yes
Start cache flushing at (i Stop cache flushing at (ir Cache block size (in KB):	n percentage): 80 n percentage): 80 8
Media scan frequency (in c	lays): 1
Failover alert delay (in m	ninutes): 5
Feature enable identifier:	303030323420303035353420493FECB1
Feature pack: DS5300 Feature pack submodel ID:	93
Storage Subsystem world-wi	de identifier (ID): 600A0B800029ECE60000000493FECAB
CONTROLLERS Number of controllers: 2	
Controller in Enclosure	2 85, Slot A
Status:	Online
Current configuration Firmware version Appware version Bootware version NVSRAM version: Pending configuration Firmware version: Appware version	on 07.50.00.00 on: 07.50.00.00 on: 07.50.00.00 N1818D53R1050V03 on c None on: None

Bootware version: None NVSRAM version: None Transferred on: None Current ID (ALPA) On drive channel 1: 125/0x1 memSpyTesting 9 Replacement part number: Model name: 7000-88 Board ID: 7091 Submodel ID: 93 1818 Product ID: Revision: 0730 Replacement part number: memSpyTesting 9 32847-10 Part number: Serial number: SF74700554 Vendor: ΙBΜ Date of manufacture: January 24, 2008 Data Cache 8192 MB Total present: Total used: 8192 MB Data Cache Module Status: Optimal Location: Controller A, Slot 1 1,024 MB Capacity: Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196510 Samsung Manufacturer: Date of manufacture: October 29, 2007 Data Cache Module Status: Optimal Location: Controller A, Slot 2 1,024 MB Capacity: Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196542 Manufacturer: Samsung Date of manufacture: October 29, 2007 Data Cache Module Status: Optimal Controller A, Slot 3 Location: Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196568 Manufacturer: Samsung Date of manufacture: October 29, 2007 Data Cache Module Status: Optimal Controller A, Slot 4 Location: Capacity: 1,024 MB Not Available Replacement Part number: Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196524 Manufacturer: Samsung October 29, 2007 Date of manufacture: Data Cache Module Status: Optimal Location: Controller A, Slot 5 Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196503

Manufacturer: Samsung Date of manufacture: October 29, 2007 Data Cache Module Status: Optimal Location: Controller A, Slot 6 Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196519 Manufacturer: Samsung October 29, 2007 Date of manufacture: Data Cache Module Status: Optimal Controller A, Slot 7 Location: Capacity: 1,024 MB Not Available Replacement Part number: Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196523 Manufacturer: Samsung October 29, 2007 Date of manufacture: Data Cache Module Status: Optimal Location: Controller A, Slot 8 1,024 MB Capacity: Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Part Number: Not Available Serial number: 1426196530 Manufacturer: Samsung Date of manufacture: October 29, 2007 Processor cache: Total present: 2044 MB Processor Cache Module Status: Optimal Controller A, Slot 1 Location: Capacity: 2,048 MB Replacement Part number: Not Available Manufacturer Part Number: Not Available Part Number: Not Available Serial number: Not Available Manufacturer: Not Available Date of manufacture: Not available Cache Backup Device Status: Optimal Type: USB flash drive Location: Controller A, Connector USB 4 Capacity: 3,919 MB Product ID: TS4GUFM-V Part number: Not Available Serial number: 000000000003B Revision level: 0.00 Manufacturer: Not Available Date of manufacture: Not available Status: Optimal Type: USB flash drive Controller A, Connector USB 3 Location: Capacity: 3,919 MB Product ID: TS4GUFM-V Part number: Not Available Serial number: 00000000000007 Revision level: 0.00 Manufacturer: Not Available Date of manufacture: Not available Host Interface Board Optimal Status:

Slot 2 Location: Type: Fibre channel Number of ports: 4 Board ID: 0801 Replacement part number: Not Available Part number: PN L2-25023-02 Serial number: SN SF84300044 Vendor: VN LSI Date of manufacture: November 1, 2008 Host Interface Board Status: Optimal Slot 1 Location: Fibre channel Type: Number of ports: 4 Board ID: 0801 Replacement part number: Not Available PN L2-25023-02 Part number: Serial number: SN SF84300041 VN LSI Vendor: Date of manufacture: November 1, 2008 Tue Jan 27 15:19:38 EST 2009 Date/Time: Associated Logical Drives (\* = Preferred Owner): FC R3 A\*, FDE R6 A\*, Sata Raid6 A\*, Secure Middle R6 A\* Controller DNS/Network name: XBB2 MiddleA Enabled\* Remote login: Ethernet port: 1 Link status: Up MAC address: 00:a0:b8:29:ed:8a Negotiation mode: Auto-negotiate 10 Mbps Port speed: Duplex mode: Half duplex IPv4 settings: Enabled Network configuration: Static 192.100., 255.255.255.0 IP address: Subnet mask: Gateway: 0.0.0.0 IPv6 settings: Enabled Auto-configuration:StatelessLocal IP address:FE80:0000:0000:02A0:B8FF:FE29:ED8A Ethernet port: 2 Link status: Failed MAC address: 00:a0:b8:29:ed:8b Negotiation mode: Auto-negotiate Port speed: Unknown Half duplex Duplex mode: IPv4 settings: Enabled Network configuration: Static IP address: 192.168.129.101 Subnet mask: 255.255.255.0 Gateway: 0.0.0.0 IPv6 settings: Enabled Auto-configuration: Static Local IP address: FE80:0000:0000:0000:0000:0000:0000 \*NVSRAM setting, may be overridden by DHCP/Bootp server setting

Fibre
1
8,7
125/0x1

Maximum data rate:	4 Gbps
Current data rate:	2 Gbps
Data rate control:	Auto
Link status:	Up
Drive interface:	Fibre
Channel:	2
Port:	6,5
Current ID:	125/0x1
Maximum data rate:	4 Gbps
Current data rate:	4 Gbps
Data rate control:	Auto
Link status:	Up
Drive interface:	Fibre
Channel:	3
Port:	4, 3, Out
Current ID:	125/0x1
Maximum data rate:	4 Gbps
Current data rate:	4 Gbps
Data rate control:	Auto
Link status:	Up
Drive interface:	Fibre
Channel:	4
Port:	2, 1, Out
Current ID:	125/0x1
Maximum data rate:	4 Gbps
Current data rate:	4 Gbps
Data rate control:	Auto
Link status:	Up
Drive interface:	Fibre
Channel:	5
Port:	1, 2
Current ID:	125/0x1
Maximum data rate:	4 Gbps
Current data rate:	2 Gbps
Data rate control:	Auto
Link status:	Up
Drive interface:	Fibre
Channel:	6
Port:	3, 4
Current ID:	125/0x1
Maximum data rate:	4 Gbps
Current data rate:	4 Gbps
Data rate control:	Auto
Link status:	Up
Drive interface:	Fibre
Channel:	7
Port:	5, 6, Out
Current ID:	125/0x1
Maximum data rate:	4 Gbps
Current data rate:	4 Gbps
Data rate control:	Auto
Link status:	Up
Drive interface:	Fibre
Channel:	8
Port:	7, 8, Out
Current ID:	125/0x1
Maximum data rate:	4 Gbps
Current data rate:	4 Gbps
Data rate control:	Auto
Link status:	Up

Host interface: Fibre Channel: 1 Not applicable/0xFFFFFFF Current ID: Preferred ID: 125/0x1 NL-Port ID: 0x010200 Maximum data rate: 8 Gbps Current data rate: 4 Gbps Data rate control: Auto Link status: Up Topology: Fabric Attach World-wide port identifier: 20:1e:00:a0:b8:29:ed:8a 20:0e:00:a0:b8:29:ed:8a World-wide node identifier: Part type: PM8032 revision 5 Host interface: Fibre Channel: 2 Current ID: Not applicable/0xFFFFFFF Preferred ID: 1/0xE8 NL-Port ID: 0xFFFFFF Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available World-wide port identifier: 20:2e:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a PM8032 Part type: revision 5 Host interface: Fibre Channel: 3 Current ID: Not applicable/0xFFFFFFF Preferred ID: 2/0xE4 NL-Port ID: **0xFFFFF** Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available World-wide port identifier: 20:3e:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Host interface: Fibre Channel: 4 Current ID: Not applicable/0xFFFFFFF Preferred ID: 3/0xE2 NL-Port ID: 0xFFFFFF Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available World-wide port identifier: 20:4e:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Host interface: Fibre Channel: 5 Current ID: Not applicable/0xFFFFFFF Preferred ID: 125/0x1 NL-Port ID: 0x010600 Maximum data rate: 8 Gbps Current data rate: 4 Gbps Data rate control: Auto Link status: Up Topology: Fabric Attach

World-wide port identifier: 20:5e:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Host interface: Fibre Channel: 6 Not applicable/0xFFFFFFF Current ID: Preferred ID: 13/0xD2 0xFFFFFF NL-Port ID: Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available World-wide port identifier: 20:6e:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a PM8032 revision 5 Part type: Host interface: Fibre Channel: 7 Current ID: Not applicable/0xFFFFFFF Preferred ID: 14/0xD1 NL-Port ID: 0xFFFFFF Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Not Available Topology: World-wide port identifier: 20:7e:00:a0:b8:29:ed:8a 20:0e:00:a0:b8:29:ed:8a World-wide node identifier: PM8032 Part type: revision 5 Host interface: Fibre Channel: 8 Current ID: Not applicable/0xFFFFFFF Preferred ID: 15/0xCE NL-Port ID: 0xFFFFFF Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available World-wide port identifier: 20:8e:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a revision 5 PM8032 Part type: Controller in Enclosure 85, Slot B Status: Online 0 Current configuration 07.50.00.00 Firmware version: Appware version: 07.50.00.00 Bootware version: 07.50.00.00 NVSRAM version: N1818D53R1050V03 Pending configuration Firmware version: None Appware version: None Bootware version: None NVSRAM version: None Transferred on: None Current ID (ALPA) On drive channel 1: 124/0x2 Replacement part number: Not Available Model name: 7000-88 Board ID: 7091
Submodel ID: 93 Product ID: 1818 Revision: 0730 Replacement part number: Part number: 32847-10 Serial number: SF74600024 Vendor: ΙBΜ Date of manufacture: January 23, 2008 Data Cache 8192 MB Total present: Total used: 8192 MB Data Cache Module Status: Optimal Location: Controller B, Slot 9 1,024 MB Capacity: Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953EZ3-CD5 Not Available Part Number: Serial number: 1127652678 Manufacturer: Samsung Date of manufacture: December 31, 2007 Data Cache Module Status: Optimal Controller B, Slot 10 Location: Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953CZ3-CD5 Part Number: Not Available Serial number: 4076997590 Manufacturer: Samsung Date of manufacture: January 29, 2007 Data Cache Module Status: Optimal Location: Controller B, Slot 11 Capacity: 1,024 MB Not Available Replacement Part number: Manufacturer Part Number: M3 93T2953CZ3-CD5 Part Number: Not Available Serial number: 4076997577 Manufacturer: Samsung Date of manufacture: January 29, 2007 Data Cache Module Status: Optimal Controller B, Slot 12 Location: 1,024 MB Capacity: Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953CZ3-CD5 Part Number: Not Available Serial number: 4076997790 Manufacturer: Samsung January 29, 2007 Date of manufacture: Data Cache Module Status: Optimal Location: Controller B, Slot 13 Capacity: 1.024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953CZ3-CD5 Part Number: Not Available Serial number: 4076997641 Manufacturer: Samsung Date of manufacture: January 29, 2007 Data Cache Module Status: Optimal Location: Controller B, Slot 14 Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953CZ3-CD5

Part Number: Not Available Serial number: 4110514102 Manufacturer: Samsung Date of manufacture: Data Cache Module Status: Optimal Location: Capacity: 1,024 MB Replacement Part number: Not Available Manufacturer Part Number: M3 93T2953CZ3-CD5 Part Number: Not Available Serial number: 4076997606 Manufacturer: Samsung Date of manufacture: Data Cache Module Status: Optimal Location: 1.024 MB Capacity: Replacement Part number: Not Available Manufacturer Part Number: Part Number: Not Available Serial number: 4076997562 Manufacturer: Samsung Date of manufacture: Processor cache: 2044 MB Total present: Processor Cache Module Status: Optimal Location: 2,048 MB Capacity: Not Available Replacement Part number: Manufacturer Part Number: Not Available Part Number: Not Available Serial number: Not Available Not Available Manufacturer: Date of manufacture: Not available Cache Backup Device Status: Optimal Type: Location: Capacity: 3,919 MB Product ID: TS4GUFM-V Part number: Not Available Serial number: 00000000000003 Revision level: 0.00 Manufacturer: Not Available Date of manufacture: Not available Status: Optimal Type: Location: Capacity: 3,919 MB Product ID: TS4GUFM-V Part number: Not Available Serial number: 00000000000049 Revision level: 0.00 Manufacturer: Not Available Date of manufacture: Not available Host Interface Board Status Optimal Slot 3 Location: Type: Fibre channel Number of ports: Δ 0801 Board ID: Replacement part number: Not Available Part number: PN L2-25023-02 Serial number: SN SF84300021 VN LSI Vendor:

January 29, 2007 Controller B, Slot 15 January 29, 2007 Controller B, Slot 16 M3 93T2953CZ3-CD5 January 29, 2007 Controller B, Slot 2 USB flash drive Controller B, Connector USB 3 USB flash drive Controller B, Connector USB 4

Date of manufacture: November 1, 2008 Host Interface Board Status: Optimal Slot 4 Location: Type: Fibre channel Number of ports: 4 0801 Board ID: Replacement part number: Not Available Part number: PN L2-25023-02 Serial number: SN SF84300014 VN LSI Vendor: Date of manufacture: November 1, 2008 Tue Jan 27 15:19:47 EST 2009 Date/Time: Associated Logical Drives (\* = Preferred Owner): FC\_R3\_B\*, FDE\_R1\*, FDE\_R6\_B\*, Sata\_Raid6\_B\*, Secure\_Middle\_R6\_B\* XBB2 MiddleB Controller DNS/Network name: Remote login: Disabled Ethernet port: 1 Link status: Up MAC address: 00:a0:b8:29:ec:e6 Negotiation mode: Auto-negotiate Port speed: 10 Mbps Duplex mode: Half duplex IPv4 settings: Enabled Network configuration: Static IP address: 192.168.70.71 Subnet mask: 255.255.255.0 Gateway: 0.0.0.0 IPv6 settings: Enabled Auto-configuration: Stateless Local IP address: FE80:0000:0000:02A0:B8FF:FE29:ECE6 Ethernet port: 2 Link status: Failed MAC address: 00:a0:b8:29:ec:e7 Negotiation mode: Auto-negotiate Port speed: Unknown Duplex mode: Half duplex IPv4 settings: Enabled Network configuration: Static IP address: 192.168.129.102 Subnet mask: 255.255.255.0 Gateway: 0.0.0.0 IPv6 settings: Enabled Auto-configuration: Stateless Local IP address: FE80:0000:0000:02A0:B8FF:FE29:ECE7 Drive interface: Fibre Channel: 1 8,7 Port: Current ID: 124/0x2 Maximum data rate: 4 Gbps Current data rate: 2 Gbps Data rate control: Auto Link status: Up Drive interface: Fibre Channel: 2 6, 5 Port: 124/0x2 Current ID:

Maximum data rate: Current data rate: Data rate control: Link status:	4 Gbps 4 Gbps Auto Up	
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 3 4, 3, 0 124/0x2 4 Gbps 4 Gbps Auto Up	ut
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 4 2, 1, 0 124/0x2 4 Gbps 4 Gbps Auto Up	ut
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 5 1, 2 124/0x2 4 Gbps 2 Gbps Auto Up	
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 6 3, 4 124/0x2 4 Gbps 4 Gbps Auto Up	
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 7 5, 6, 0 124/0x2 4 Gbps 4 Gbps Auto Up	ut
Drive interface: Channel: Port: Current ID: Maximum data rate: Current data rate: Data rate control: Link status:	Fibre 8 7, 8, 0 124/0x2 4 Gbps 4 Gbps Auto Up	ut
Host interface: Channel: Current ID: Preferred ID: NL-Port ID: Maximum data rate: Current data rate: Data rate control:		Fibre 1 Not applicable/0xFFFFFFF 8/0xD9 0xFFFFFF 8 Gbps 8 Gbps Auto

Link status: Down Topology: Not Available World-wide port identifier: 20:1f:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Host interface: Fibre Channel: 2 Current ID: Not applicable/0xFFFFFFF Preferred ID: 9/0xD6 NL-Port ID: 0xFFFFFF Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available World-wide port identifier: 20:2f:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Host interface: Fibre Channel: 3 Current ID: Not applicable/0xFFFFFFF Preferred ID: 10/0xD5 NL-Port ID: 0xFFFFFF Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available 20:3f:00:a0:b8:29:ed:8a World-wide port identifier: World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a PM8032 Part type: revision 5 Host interface: Fibre Channel: 4 Current ID: Not applicable/0xFFFFFFF Preferred ID: 11/0xD4 NL-Port ID: 0x010700 Maximum data rate: 8 Gbps Current data rate: 4 Gbps Data rate control: Auto Link status: Up Topology: Fabric Attach World-wide port identifier: 20:4f:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Fibre Host interface: Channel: 5 Current ID: Not applicable/0xFFFFFFF Preferred ID: 12/0xD3 NL-Port ID: 0xFFFFFF Maximum data rate: 8 Gbps Current data rate: 8 Gbps Data rate control: Auto Link status: Down Topology: Not Available World-wide port identifier: 20:5f:00:a0:b8:29:ed:8a World-wide node identifier: 20:0e:00:a0:b8:29:ed:8a Part type: PM8032 revision 5 Host interface: Fibre Channel: 6 Current ID: Not applicable/0xFFFFFFF Preferred ID: 13/0xD2 NL-Port ID: 0x010300

Maximum data rate: Current data rate: Data rate control: Link status: Topology: World-wide port identifi World-wide node identifi Part type:	8 Gbps 4 Gbps Auto Up Fabric Attach er: 20:6f:00:a0:b8:29:ed:8a er: 20:0e:00:a0:b8:29:ed:8a PM8032 revision 5
Host interface: Channel: Current ID:	Fibre 7 Not applicable/0xFFFFFFFF
Preferred ID: NL-Port ID: Maximum data rate: Current data rate: Data rate control: Link status: Topology: World-wide port identifi World-wide node identifi Part type:	14/0xD1 0xFFFFF 8 Gbps 8 Gbps Auto Down Not Available er: 20:7f:00:a0:b8:29:ed:8a er: 20:0e:00:a0:b8:29:ed:8a PM8032 revision 5
Host interface: Channel: Current ID: Preferred ID: NL-Port ID: Maximum data rate: Current data rate: Data rate control: Link status: Topology: World-wide port identifi World-wide node identifi Part type:	Fibre 8 Not applicable/0xFFFFFFF 15/0xCE 0xFFFFF 8 Gbps 8 Gbps Auto Down Not Available er: 20:8f:00:a0:b8:29:ed:8a er: 20:0e:00:a0:b8:29:ed:8a PM8032 revision 5
subsystemS Number of subsystems: 5	
Name:	1
Status:	Optimal
Capacity RAID level:	1.064 TB 6
Drive type: Enclosure loss protection:	Fibre Channel/Full Disk Encryption(FDE) No
Security Capable: Secure:	Yes Yes
Current owner:	Controller in slot A,B
Associated logical drives a	nd free capacity
Logical Drive C Secure_Middle_R6_A Secure_Middle_R6_B Free Capacity	apacity 500.000 GB 510.000 GB 79.857 GB

### Associated drives - present (in piece order)

Enclosure	Slot
2	3
2	4
2	5
2	6
2	7
2	12
4	2
4	4
4	7
4	8

Name:	FC_R3
Status:	Optimal
Capacity	836.688 GB
RAID level:	3
Drive type:	Fibre channel
Enclosure loss protection:	No
Security Capable:	No
Secure:	No
Current owner:	Controller in slot A,B

Associated logical drives and free capacity

Logical Drive	Capacity
FC_R3_A	400.000 GB
FC R3 B	420.000 GB
Free Capacity	16.688 GB

Associated drives - present (in piece order)

Enclosure	Slot
4	9
4	12
4	14
4	16

Na	me:	FDE_R6
	Status:	Optimal
	Capacity RAID level:	817.381 GB 6
	Drive type: Enclosure loss protection:	Fibre Channel/Full Disk Encryption(FDE) No

Security	Capable:	Yes
Secure:		Yes

Current owner:

Controller in slot A,B

Associated logical drives and free capacity

Logical Drive	Capacity
FDE R6 A	400.000 GB
FDE_R6_B	410.000 GB
Free Capacity	7.381 GB

## Associated drives - present (in piece order)

Enclosure	Slot
2	1
2	2
2	13
2	15
2	16
4	5
4	6
4	3

Na	me:	Sata_Raid6
	Status:	Optimal
	Capacity RAID level:	2.727 TB 6
	Drive type: Enclosure loss protection:	Serial ATA (SATA) No
	Security Capable: Secure:	No No
	Current owner:	Controller in slot A,B

Associated logical drives and free capacity

Logical Drive	Capacity
Sata_Raid6_A	1,000.000 GB
Sata_Raid6_B	1,000.000 GB
Free Capacity	792.550 GB

Associated drives - present (in piece order)

Slot
8
9
10
11
1

Name: Secure R1 Status: Optimal 278.897 GB Capacity RAID level: 1 Fibre Channel/Full Disk Encryption(FDE) Drive type: Enclosure loss protection: No Security Capable: Yes Secure: Yes Controller in slot B Current owner: Associated logical drives and free capacity Logical Drive Capacity FDE R1 278.000 GB Free Capacity 918.295 MB Associated drives - present (in piece order) Enclosure Slot 4 15 [mirrored pair with drive at enclosure 4, slot 10] 4 10 [mirrored pair with drive at enclosure 4, slot 15] STANDARD LOGICAL DRIVES-----SUMMARY Number of standard logical drives: 9 See other Logical Drives sub-tabs for premium feature information. NAME STATUS CAPACITY RAID LEVEL subsystem DRIVE TYPE FC R3 A Optimal 400.0 GB 3 FC R3 Fibre channel FC R3 B Optimal 420.0 GB FC R3 3 Fibre channel FDE R1 Optimal 278.0 GB Secure\_R1 Fibre, FDE 1 FDE R6 A Optimal 400.0 GB FDE R6 Fibre, FDE 6 FDE R6 B Optimal 410.0 GB 6 FDE R6 Fibre, FDE Optimal 1,000.0 GB 6 Sata\_Raid6\_A Sata\_Raid6 Serial ATA (SATA) Optimal 1,000.0 GB 6 Sata Raid6 B Sata\_Raid6 Serial ATA (SATA) Secure Middle R6 A Optimal 500.0 GB 6 1 Fibre, FDE Secure Middle R6 B Optimal 510.0 GB 6 1 Fibre, FDE DETAILS Logical Drive name: FC R3 A Logical Drive status: Optimal Capacity: 400.0 GB Logical Drive ID: 60:0a:0b:80:00:29:ed:8a:00:00:cd:50:49:41:11:ff Subsystem ID (SSID): 6

13

4

	Associated subsystem: RAID level:	FC_R3 3		
	Secure:	No		
	Drive type: Enclosure loss protection:	Fibre chan No	nel	
	Preferred owner: Current owner:	Controller Controller	in slot A in slot A	
	Segment size: Capacity reserved for future Maximum future segment size Modification priority:	e segment s :	ize changes:	256 KB Yes 2,048 KB High
	Read cache: Write cache: Write cache without batte Write cache with mirrorin Flush write cache after (in Dynamic cache read prefetch	eries: ng: seconds): :	Enabled Enabled Disabled Enabled 10.00 Enabled	
	Enable background media sca Media scan with redundancy o	n: check:	Enabled Enabled	
	Pre-Read redundancy check:		Disabled	
Log	gical Drive name:	FDE_R1		
	Logical Drive status:	Optimal		
	Capacity: Logical Drive ID: Subsystem ID (SSID): Associated subsystem: RAID level:	278.0 GB 60:0a:0b:80 8 Secure 1	0:00:33:98:48 _R1	:00:00:c9:73:49:41:0a:a1
	Secure:	Yes		
	Drive type: Enclosure loss protection:	Fibre Chan No	nel/Full Disk	Encryption(FDE)
	Preferred owner: Current owner:	Controller Controller	in slot B in slot B	
	Segment size: Capacity reserved for future Maximum future segment size Modification priority:	e segment s :	ize changes:	128 KB Yes 2,048 KB High
	Read cache: Write cache: Write cache without batte Write cache with mirrorin Flush write cache after (in Dynamic cache read prefetch	eries: ng: seconds): :	Enabled Enabled Disabled Enabled 10.00 Enabled	
	Enable background media scan Media scan with redundancy o	n: check:	Enabled Disabled	
	Pre-Read redundancy check:		Disabled	

Logical Drive name:	FDE_R6_B					
Logical Drive status:	Optimal					
Capacity: Logical Drive ID: Subsystem ID (SSID): Associated subsystem: RAID level:	410.0 GB 60:0a:0b:80 5 FDE_R6 6	0:00:29:ec:e6	:00:00:df:13:49:41:13:4f			
Secure:	Yes					
Drive type: Enclosure loss protection:	Fibre Chan No	nel/Full Disk	Encryption(FDE)			
Preferred owner: Current owner:	Controller Controller	in slot B in slot B				
Segment size: Capacity reserved for future Maximum future segment size: Modification priority:	e segment s :	ize changes:	128 KB Yes 2,048 KB High			
Read cache: Write cache: Write cache without batte Write cache with mirrorin Flush write cache after (in Dynamic cache read prefetch:	eries: ng: seconds): :	Enabled Enabled Disabled Enabled 10.00 Enabled				
Enable background media scan Media scan with redundancy o	n: check:	Enabled Enabled				
Pre-Read redundancy check:		Disabled				
Logical Drive name:	Sata_Raid6	_A				
Logical Drive status:	Optimal					
Capacity: Logical Drive ID: Subsystem ID (SSID): Associated subsystem: RAID level:	1,000.0 GB 60:0a:0b:80 0 Sata_Ra 6	0:00:29:ed:8a aid6	:00:00:cc:90:49:3f:f7:b5			
Secure:	No					
Drive type: Enclosure loss protection:	Serial ATA No	(SATA)				
Preferred owner: Current owner:	Controller Controller	in slot A in slot A				
Segment size: Capacity reserved for future Maximum future segment size: Modification priority:	e segment s :	ize changes:	128 KB Yes 2,048 KB High			
Read cache: Write cache: Write cache without batte Write cache with mirrorin	eries: ng:	Enabled Enabled Disabled Enabled				

Flush write cache after (in Dynamic cache read prefetch	<pre>seconds): :</pre>	10.00 Enabled			
Enable background media sca Media scan with redundancy	n: check:	Enabled Enabled			
Pre-Read redundancy check:		Disabled			
Logical Drive name:	Sata_Raid6	_B			
Logical Drive status:	Optimal				
Capacity: Logical Drive ID: Subsystem ID (SSID): Associated subsystem: RAID level:	1,000.0 GB 60:0a:0b:8 1 Sata_R 6	0:00:29:ec:e6 aid6	:00:00:dd:b1:49:3f:f9:87		
Secure:	No				
Drive type: Enclosure loss protection:	Serial ATA No	(SATA)			
Preferred owner: Current owner:	Controller Controller	in slot B in slot B			
Segment size: Capacity reserved for futur Maximum future segment size Modification priority:	e segment s :	ize changes:	128 KB Yes 2,048 KB High		
Read cache: Write cache: Write cache without batt Write cache with mirrori Flush write cache after (in Dynamic cache read prefetch	eries: ng: seconds): :	Enabled Enabled Disabled Enabled 10.00 Enabled			
Enable background media sca Media scan with redundancy	n: check:	Enabled Enabled			
Pre-Read redundancy check: Logical Drive name:	Secure_Mid	Disabled dle_R6_B			
Logical Drive status:	Optimal				
Capacity: Logical Drive ID: Subsystem ID (SSID): Associated subsystem: RAID level:	510.0 GB 60:0a:0b:8 3 1 6	0:00:29:ec:e6	:00:00:de:a0:49:40:e6:2d		
Secure:	Yes				
Drive type: Enclosure loss protection:	Fibre Chan No	nel/Full Disk	Encryption(FDE)		
Preferred owner: Current owner:	Controller Controller	in slot B in slot B			
Segment size: Capacity reserved for futur Maximum future segment size	e segment s :	ize changes:	128 KB Yes 2,048 KB		

Modification priority:

Read cache: Write cache: Write cache without batteries: Write cache with mirroring: Flush write cache after (in seconds): Dynamic cache read prefetch:							Enabled Enabled Disabled Enabled 10.00 Enabled					
	Enable b Media sc	ackground m an with rec	nedia sca dundancy	n: chec	:k:	Enal Enal	ble ble	ed ed				
	Pre-Read	redundancy	/ check:			Disa	abl	led				
FLASH	COPY REPO	SITORY LOG	ICAL DRIV	'ES								
Nu	mber of f	lashcopy re	epositori	es:	0							
FLASH	COPY LOGI	CAL DRIVES-										
Nu	mber of f	lashcopy lo	ogical dr	ives	s: 0							
MISSI	NG LOGICA	L DRIVES										
Nu	mber of m	issing logi	ical driv	es:	0							
DRIVE	S											
SUMMA Nu	RY mber of d Current	rives: 32 drive types	s: Fibre	(5),	, Fibre,	/FDE	(21	l), Serial	ATA (S	ATA)	(6)	
BA	SIC:											
e	nclosure,	SLOT STAT	TUS CAP	ACIT	יד אי	YPE		CURRENT	DATA	RATE	PRODUCT	ID
2	<b>,</b> 1	Optimal	136.732	GB	Fibre,	FDE	4	Gbps		ST31	46756FC	F
2	, 2	Optimal	136.732	GB	Fibre,	FDE	4	Gbps		ST31	46756FC	F
2	, 3	Optimal	136.732	GB	Fibre,	FDE	4	Gbps		ST31	46756FC	F
2	, 4	Optimal	136.732	GB	Fibre,	FDE	4	Gbps		ST31	46756FC	F
2	<b>,</b> 5	Optimal	136.732	GB	Fibre,	FDE	4	Gbps		ST31	46756FC	F
2	<b>,</b> 6	E097 Optimal	136.732	GB	Fibre,	FDE	4	Gbps		ST31	46756FC	F
2	<b>,</b> 7	E097 Optimal	279.397	GB	Fibre,	FDE	4	Gbps		ST33	00056FC	F
2	, 8	E097 Optimal	698.638	GB	SATA		4	Gbps		ST37	50330NS	
43W97 2	15 42C041 , 9	7IBM BB10 Optimal	9 698.638	GB	SATA		4	Gbps		ST37	50330NS	
43W97 2	15 42C041 , 10	7IBM BB10 Optimal	9 698.638	GB	SATA		4	Gbps		ST37	50330NS	
43W97 2	15 42C041 , 11	7IBM BB10 Optimal	9 698.638	GB	SATA		4	Gbps		ST37	50330NS	
43W97 2	15 42C041 , 12	7IBM BB10 Optimal	) 136.732	GB	Fibre,	FDE	4	Gbps		ST31	46756FC	F
2	. 13	E097 Optimal	279.397	GB	Fibre.	FDE	4	Gbps		ST33	00056FC	F
2	14	E097 Optimal	410 127	GR	Fibre	FDF	Д	Ghns		STRA	50056FC	F
2	, 1T	E097	270 207		Libra		-+	Chas		CT22		' -
2	, 10	E097	2/9.39/	uВ	ribre,	FUE	4	ups		2122	00000FL	F

High

2,	16	Optimal E097	279.397 GE	B Fibre,	FDE	4 Gbps	ST3300056FC	F
4,	1	Optimal	698.638 GE	3 SATA		4 Gbps	ST3750330NS	
43W971	5 42C0417	7IBM BB10	)					
4,	2	Optimal E097	136.732 GE	3 Fibre,	FDE	4 Gbps	ST3146756FC	F
4,	3	Optimal E097	136.732 GE	B Fibre,	FDE	4 Gbps	ST3146756FC	F
4,	4	Optimal F097	136.732 GE	B Fibre,	FDE	4 Gbps	ST3146756FC	F
4,	5	Optimal F097	136.732 GE	B Fibre,	FDE	4 Gbps	ST3146756FC	F
4,	6	Optimal F097	136.732 GE	B Fibre,	FDE	4 Gbps	ST3146756FC	F
4,	7	Optimal F097	136.732 GE	B Fibre,	FDE	4 Gbps	ST3146756FC	F
4,	8	Optimal F097	136.732 GE	B Fibre,	FDE	4 Gbps	ST3146756FC	F
4,	9	Optimal	279.397 GE	3 Fibre		4 Gbps	HUS1545300FC	
4,	10	Optimal F097	279.397 GE	B Fibre,	FDE	4 Gbps	ST3300056FC	F
4,	11	Optimal	419.187 GE	3 Fibre		4 Gbps	HUS1545450FC	
4,	12	Optimal	279.397 GE	3 Fibre		4 Gbps	HUS1545300FC	
4.	13	Ontimal	698,638 GF	SATA		4 Ghns	ST3750330NS	
43W971	5 420041	7TBM BB16	)					
4,	14	Optimal	279.397 GE	3 Fibre		4 Gbps	HUS1545300FC	
4,	15	Optimal F097	279.397 GE	B Fibre,	FDE	4 Gbps	ST3300056FC	F
4,	16	Optimal JG01	279.397 GE	8 Fibre		4 Gbps	HUS1545300FC	

DRIVE CHANNELS:

enclosure,	SLOT	PREFERRED	CHANNEL	REDUNDANT	CHANNEL
2, 1	3		8		
2,2	8		3		
2,3	3		8		
2,4	8		3		
2,5	3		8		
2,6	8		3		
2, 7	3		8		
2,8	8		3		
2, 9	3		8		
2, 10	8		3		
2, 11	3		8		
2, 12	8		3		
2, 13	3		8		
2, 14	8		3		
2, 15	3		8		
2, 16	8		3		
4, 1	7		4		
4, 2	4		7		
4, 3	7		4		
4, 4	4		7		
4, 5	7		4		
4, 6	4		7		
4, 7	7		4		
4, 8	4		7		
4, 9	7		4		
4, 10	4		7		
4, 11	7		4		
4, 12	4		7		
4, 13	7		4		
4, 13	/		4		

4, 1 4, 1 4, 1	14 4 15 7 16 4		7 4 7							
HOT SPARE The fo Total H	COVERAGE: llowing subsystems not spare drives:	are n 2	ot protected: Sata_Raid6							
Star In u	ndby: 2 use: 0									
Standby Prot	/ drive at enclosu cects the followin	ure 2, ng subs	slot 14 (Fibre, 419.187 GB) ystems: 1, FDE_R6, FC_R3, Secure_R1							
Standby Prot	Standby drive at enclosure 4, slot 11 (Fibre, 419.187 GB) Protects the following subsystems: 1, FDE_R6, FC_R3, Secure_R1									
DETAILS Drive a	S at Enclosure 2, Sl	ot 1								
Stat	cus:	0p	timal							
Mode Raw Usal Wor Asso	e: capacity: ble capacity: ld-wide identifier ociated subsystem:	As 13 13 20	Assigned 136.732 GB 136.232 GB 20:00:00:1d:38:1d:1d:73:00:00:00:00:00:00:00:00 FDE_R6							
Port 0 1	: Channel 3 8	ID 0/0x 0/0x	EF EF							
Secu Secu Read Secu	urity Capable: ure: d/write accessible urity key identifi	e: er:	Yes Yes Yes 27000000600A0B800029ECE60000DE9E4940E373							
Spee Curr Proo Firr Ser Veno Date	ed: rent data rate: duct ID: nware version: ial number: dor: e of manufacture:		15,015 RPM 4 Gbps ST3146756FC F E097 3QN089B500009912TLGR IBM-SSG October 16, 2008							
Drive a	at Enclosure 2, Sl	ot 2								
Stat	cus:	0p	timal							
Mode Raw Usal Wor Asso	e: capacity: ble capacity: d-wide identifier ociated subsystem:	As 13 13 20	signed 6.732 GB 6.232 GB :00:00:1d:38:1d:1d:e1:00:00:00:00:00:00:00:00 FDE_R6							
Port 0 1	: Channel 8 3	ID 1/0x 1/0x	E8 E8							

	Security Secure: Read/writ Security	Capable: e accessible: key identifier	` <b>:</b>	Yes Yes 27000000600A0B800029ECE60000DE9E4940E373
	Speed: Current d Product I Firmware Serial num Vendor: Date of m	ata rate: D: version: mber: anufacture:		15,015 RPM 4 Gbps ST3146756FC F E097 3QN03KCT00009912TM1L IBM-SSG October 16, 2008
Dr	ive at Enc	losure 2, Slot	: 3	
	Status:		0pt	timal
	Mode: Raw capac Usable ca World-wid Associate	ity: pacity: e identifier: d subsystem:	Ass 130 130 201	signed 6.732 GB 6.232 GB :00:00:1d:38:1d:1d:7c:00:00:00:00:00:00:00:00 1
	Port 0 1	Channel 3 8	ID 2/0x8 2/0x8	E4 E4
	Security Secure: Read/writ Security	Capable: e accessible: key identifier	`:	Yes Yes Yes 2700000600A0B800029ECE60000DE9E4940E373
Dr	Speed: Current d Product I Firmware Serial nuu Vendor: Date of m ive at Enc	ata rate: D: version: mber: anufacture: losure 2, Slot	: 6	15,015 RPM 4 Gbps ST3146756FC F E097 3QN0899100009903RE2S IBM-SSG October 16, 2008
	Status:		0pt	timal
	Mode: Raw capac Usable ca World-wid Associate	ity: pacity: e identifier: d subsystem:	Ass 130 130 203	signed 6.732 GB 6.232 GB :00:00:1d:38:1d:1d:5c:00:00:00:00:00:00:00:00 1
	Port O 1	Channel 8 3	ID 5/0x1 5/0x1	E0 E0
	Security Secure: Read/writ Security	Capable: e accessible: key identifier	`:	Yes Yes Yes 27000000600A0B800029ECE60000DE9E4940E373
	Speed: Current d Product I	ata rate: D:		15,015 RPM 4 Gbps ST3146756FC F

	Firmware Serial nur Vendor: Date of ma	version: nber: anufacture:		E097 3QN07L7H00009912TL9U IBM-SSG October 16, 2008				
Dri	ive at Enc	losure 2, Slot	: 7					
	Status:		0p	timal				
	Mode: Raw capac Usable ca World-widd Associated	ity: pacity: e identifier: d subsystem:	As: 279 278 20	signed 0.397 GB 3.897 GB :00:00:1d:38:1b:16:1d 1	:00:00:00:00:00:00:00:00			
	Port O 1	Channel 3 8	ID 6/0xI 6/0xI	0C 0C				
	Security ( Secure: Read/write Security	Capable: e accessible: key identifier	`:	Yes Yes Yes 2700000600A0B800029	ECE60000DE9E4940E373			
	Speed: Current data rate: Product ID: Firmware version: Serial number: Vendor: Date of manufacture:			15,015 RPM 4 Gbps ST3300056FC F E097 3QP0A4C000009912TLAE IBM-SSG October 16, 2008				
Dri	ive at Enc	losure 2, Slot	: 8					
	Status:		0p	timal				
	Mode: Raw capac Usable ca World-widd Associated	ity: pacity: e identifier: d subsystem:	As: 698 698 20	ssigned 98.638 GB 98.138 GB 0:00:00:a0:b8:41:47:0e:00:00:00:00:00:00:00:00 Sata_Raid6				
	Port O 1	Channel 8 3	ID 7/0x 7/0x	[10PI [10PI				
	Security ( Secure: Read/write Security	Capable: e accessible: key identifier	`:	No No Yes Not Applicable				
	Speed: Current da Product II Package vo Firmware v Serial nur Vendor: Date of ma	ata rate: D: ersion: version: mber: anufacture:		7,200 RPM 4 Gbps ST3750330NS BB10 BB10 5QK04MJZ Not Available Not Available	43W9715 42C0417IBM			
	ATA Trans	lator						

Product ID:	BR-2401-3.0
Vendor:	SLI
Firmware Version:	LP1158
Location:	Individual Drive Controller
Drive at Enclosure 2, Slot	9
Status:	Optimal
Mode:	Assigned
Raw capacity:	698.638 GB
Usable capacity:	698.138 GB
World-wide identifier:	20:00:00:a0:b8:41:27:26:00:00:00:00:00:00:00:00
Associated subsystem:	Sata_Raid6
Port Channel	ID
0 3	8/0xD9
1 8	8/0xD9
Security Capable:	No
Secure:	No
Read/write accessible:	Yes
Security key identifien	*: Not Applicable
Speed:	7,200 RPM
Current data rate:	4 Gbps
Product ID:	ST3750330NS 43W9715 42C0417IBM
Package version:	BB10
Firmware version:	BB10
Serial number:	5QK04MLD
Vendor:	Not Available
Date of manufacture:	Not Available
ATA Translator Product ID: Vendor: Firmware Version: Location:	BR-2401-3.0 SLI LP1158 Individual Drive Controller
Drive at Enclosure 4, Slo	5 8
Status:	Optimal
Mode:	Assigned
Raw capacity:	136.732 GB
Usable capacity:	136.232 GB
World-wide identifier:	20:00:00:1d:38:1d:1e:18:00:00:00:00:00:00:00:00
Associated subsystem:	1
Port Channel	ID
0 4	7/0xT10PI
1 7	7/0xT10PI
Security Capable:	Yes
Secure:	Yes
Read/write accessible:	Yes
Security key identifien	27000000600A0B800029ECE60000DE9E4940E373
Speed:	15,015 RPM
Current data rate:	4 Gbps

	Product I Firmware Serial nu Vendor: Date of m	ID: version: umber: nanufactur	re:	ST3146 E097 3QN080 IBM-SS Octobe	756FC DY0000 G r 16,	F 19903R 2008	DM	IQ	
Dri	ive at End	closure 4	, Slot 9						
	Status:			Optimal					
	Mode: Raw capacity: Usable capacity: World-wide identifier: Associated subsystem:			Assigned 279.397 GB 278.897 GB 50:00:cc:a0:08:02:58:0c:00:00:00:00:00:00:00:00 FC_R3					
	Port O 1	Channel 7 4	ID 8/ 8/	0xD9 0xD9					
	Security Secure: Read/wrii Security	Capable: te access key iden <sup>:</sup>	ible: tifier:	No No Yes Not Ap	plicab	ole			
	Speed: Current of Product D Firmware Serial nu Vendor: Date of m	data rate ID: version: umber: nanufactur	: re:	15,000 4 Gbps HUS154 JG01 JLV18Y IBM-SS July 4	RPM 5300FC UC G , 2008	3			
DRIVE	CHANNELS-								
SUN	IMARY								
	CHANNEL 1 2	PORT 8,7 6 5	STATUS Optimal	CTRL A Up	LINK	CTRL Up	В	LINK	

1	8,7	Optimal	Up	Up
2	6,5	Optimal	Up	Up
3	4,3,0ut	Optimal	Up	Up
4	2,1,Out	Optimal	Up	Up
5	1,2	Optimal	Up	Up
6	3,4	Optimal	Up	Up
7	5,6,Out	Optimal	Up	Up
8	7,8,0ut	Optimal	Up	Up

DETAILS

DRIVE CHANNEL 1

Port: 8, 7 Status: Optimal Max. Rate: 4 Gbps Current Rate: 2 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up

#### DRIVE COUNTS

Total # of attached drives: 0

## CUMULATIVE ERROR COUNTS

Controller A

Baseline time set: Sample period (days, hh:mm:ss): Controller detected errors: Drive detected errors: Timeout errors: Link down errors: Total I/O count:	1/22/09 5 days, 0 0 0 N/A 330846	6:51:27 AM 08:28:31
Controller B		
Baseline time set: Sample period (days, hh:mm:ss): Controller detected errors: Drive detected errors: Timeout errors: Link down errors: Total I/O count:	1/22/09 5 days, 0 0 0 N/A 0	6:51:27 AM 08:28:31
DRIVE CHANNEL 2		
Port: 6, 5 Status: Optimal Max. Rate: 4 Gbps Current Rate: 4 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up		
DRIVE COUNTS		
Total # of attached drives: 0		
CUMULATIVE ERROR COUNTS		
Controller A		
Baseline time set: Sample period (days, hh:mm:ss): Controller detected errors: Drive detected errors: Timeout errors: Link down errors: Total I/O count:	1/22/09 5 days, 0 0 0 N/A 330856	6:51:27 AM 08:28:31
Controller B		
Baseline time set: Sample period (days, hh:mm:ss): Controller detected errors: Drive detected errors: Timeout errors: Link down errors: Total I/O count:	1/22/09 5 days, 0 0 0 N/A 0	6:51:27 AM 08:28:31
DRIVE CHANNEL 3		
Port: 4, 3, Out Status: Optimal Max. Rate: 4 Gbps Current Rate: 4 Gbps		

Rate Control: Auto Controller A link status: Up Controller B link status: Up DRIVE COUNTS Total # of attached drives: 16 Connected to: A, Port 4 Attached drives: 16 Drive enclosure: 2 (16 drives) CUMULATIVE ERROR COUNTS Controller A 1/22/09 6:51:27 AM Baseline time set: Sample period (days, hh:mm:ss): 5 days, 08:28:31 Controller detected errors: 0 Drive detected errors: 17 Timeout errors: 0 Link down errors: N/A Total I/O count: 15610992 Controller B Baseline time set: 1/22/09 6:51:27 AM Sample period (days, hh:mm:ss): 5 days, 08:28:31 Controller detected errors: 0 Drive detected errors: 0 Timeout errors: 0 Link down errors: N/A Total I/O count: 0 DRIVE CHANNEL 4 Port: 2, 1, Out Status: Optimal Max. Rate: 4 Gbps Current Rate: 4 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up DRIVE COUNTS Total # of attached drives: 16 Connected to: A, Port 2 Attached drives: 16 Drive enclosure: 4 (16 drives) CUMULATIVE ERROR COUNTS Controller A Baseline time set: 1/22/09 6:51:27 AM Sample period (days, hh:mm:ss): 5 days, 08:28:31 Controller detected errors: 0 Drive detected errors: 17 Timeout errors: 0 Link down errors: N/A Total I/O count: 10253384 Controller B 1/22/09 6:51:27 AM Baseline time set:

```
Sample period (days, hh:mm:ss): 5 days, 08:28:31
         Controller detected errors:
                                          0
         Drive detected errors:
                                          0
         Timeout errors:
                                          0
         Link down errors:
                                          N/A
         Total I/O count:
                                          0
DRIVE CHANNEL 5
   Port: 1, 2
   Status: Optimal
   Max. Rate: 4 Gbps
   Current Rate: 2 Gbps
   Rate Control: Auto
   Controller A link status: Up
   Controller B link status: Up
   DRIVE COUNTS
      Total # of attached drives: 0
   CUMULATIVE ERROR COUNTS
      Controller A
         Baseline time set:
                                         1/22/09 6:51:27 AM
         Sample period (days, hh:mm:ss): 5 days, 08:28:31
         Controller detected errors:
                                          0
         Drive detected errors:
                                          0
         Timeout errors:
                                          0
         Link down errors:
                                         N/A
         Total I/O count:
                                          330848
      Controller B
         Baseline time set:
                                          1/22/09 6:51:27 AM
         Sample period (days, hh:mm:ss): 5 days, 08:28:31
         Controller detected errors:
                                          0
         Drive detected errors:
                                          0
         Timeout errors:
                                          0
         Link down errors:
                                         N/A
         Total I/O count:
                                          0
DRIVE CHANNEL 6
   Port: 3, 4
Status: Optimal
   Max. Rate: 4 Gbps
   Current Rate: 4 Gbps
   Rate Control: Auto
   Controller A link status: Up
   Controller B link status: Up
   DRIVE COUNTS
      Total # of attached drives: 0
   CUMULATIVE ERROR COUNTS
      Controller A
         Baseline time set:1/22/09 6:51:27 ASample period (days, hh:mm:ss):5 days, 08:28:31
                                         1/22/09 6:51:27 AM
         Controller detected errors:
                                          0
         Drive detected errors:
                                          0
```

```
0
         Timeout errors:
        Link down errors:
                                         N/A
        Total I/O count:
                                         330848
     Controller B
                                         1/22/09 6:51:27 AM
        Baseline time set:
         Sample period (days, hh:mm:ss): 5 days, 08:28:31
         Controller detected errors:
                                         0
         Drive detected errors:
                                         0
         Timeout errors:
                                         0
        Link down errors:
                                        N/A
        Total I/O count:
                                         0
DRIVE CHANNEL 7
   Port: 5, 6, Out
   Status: Optimal
  Max. Rate: 4 Gbps
   Current Rate: 4 Gbps
  Rate Control: Auto
  Controller A link status: Up
   Controller B link status: Up
  DRIVE COUNTS
     Total # of attached drives: 16
     Connected to: B, Port 5
        Attached drives: 16
           Drive enclosure: 4 (16 drives)
   CUMULATIVE ERROR COUNTS
     Controller A
         Baseline time set:
                                         1/22/09 6:51:27 AM
         Sample period (days, hh:mm:ss): 5 days, 08:28:31
        Controller detected errors:
                                         0
        Drive detected errors:
                                        17
        Timeout errors:
                                         0
        Link down errors:
                                        N/A
        Total I/O count:
                                        11605011
     Controller B
         Baseline time set:
                                         1/22/09 6:51:27 AM
        Sample period (days, hh:mm:ss): 5 days, 08:28:31
        Controller detected errors:
                                         0
         Drive detected errors:
                                         0
                                         0
        Timeout errors:
        Link down errors:
                                         N/A
        Total I/O count:
                                         0
DRIVE CHANNEL 8
   Port: 7, 8, Out
   Status: Optimal
  Max. Rate: 4 Gbps
   Current Rate: 4 Gbps
  Rate Control: Auto
   Controller A link status: Up
   Controller B link status: Up
  DRIVE COUNTS
     Total # of attached drives: 16
```

Connected to: B, Port 7 Attached drives: 16 Drive enclosure: 2 (16 drives) CUMULATIVE ERROR COUNTS Controller A Baseline time set: 1/22/09 6:51:27 AM Sample period (days, hh:mm:ss): 5 days, 08:28:31 Controller detected errors: 0 17 Drive detected errors: Timeout errors: 0 Link down errors: N/A Total I/O count: 13646672 Controller B Baseline time set: 1/22/09 6:51:27 AM Sample period (days, hh:mm:ss): 5 days, 08:28:31 Controller detected errors: 0 Drive detected errors: 0 Timeout errors: 0 Link down errors: N/A Total I/O count: 0 ENCLOSURES------Drive Enclosure 2 Overall Component Information Enclosure audible alarm: **Disabled** Enclosure path redundancy: OK Current drive types: Fibre channel/Full Disk Encryption (FDE), Serial ATA (SATA) Fibre channel/Full Disk Encryption (FDE), Serial ATA (SATA) Part number: PN 39M5707 SN 1T54402443 Serial number: Vendor: VN IBM January 1, 2006 Date of manufacture: ESM Controllers Detected: 2 ESM card status: Optimal Firmware version: 98C3 Configuration settings version: FD 00.50 01/22/2008 Maximum data rate: 4 Gbps Current data rate: Not Available Location: A (left) Card communication: 0K Product ID: EXP5000 Part number: PN 39M5570 Serial number: SN 1T54136450 Vendor TBM Date of manufacture: December 1, 2005 ESM card status: Optimal Firmware version: 98C3 Configuration settings version: FD 00.50 01/22/2008 Maximum data rate: 4 Gbps

Not Available

Current data rate:

Location:	B (right)
Card communication:	OK
Product ID:	EXP5000
Part number:	PN 39M5570
Serial number:	SN 1T54721521
Vendor:	IBM
Date of manufacture:	December 1, 2005

SFPs Detected: 2

SFP status:	Optimal
Attached to:	ESM Controller A (left)
Location:	Unknown
Supported data rate(s):	1 Gbps, 2 Gbps, 4 Gbps
Link length:	Intermediate
Connector:	LC
Transmitter type:	Shortwave Laser w/o OFC
Transmission media:	TM Multi-mode 62.5m(M6)
IEEE company ID:	00 04 85
Revision:	1
Part number:	PLRXPLVCSG324N
Serial number:	C549RQ0DK
Vendor:	PICOLIGHT
Date of manufacture:	November 29, 2005

SFP status:	Optimal
Attached to:	ESM Controller B (right)
Location:	Unknown
Supported data rate(s):	1 Gbps, 2 Gbps, 4 Gbps
Link length:	Intermediate
Connector:	LC
Transmitter type:	Shortwave Laser w/o OFC
Transmission media:	TM Multi-mode 62.5m(M6)
IEEE company ID:	00 04 85
Revision:	1
Part number:	PLRXPLVCSG324N
Serial number:	C549RQ0ER
Vendor:	PICOLIGHT
Date of manufacture:	November 29, 2005

Power-Fan Controllers Detected: 2

Power-fan Controller (left) status:OptimalPart number:PN 15240-01Serial number:SN YPT054300403Vendor:VN ENGENIODate of manufacture:October 1, 2005

Power-fan Controller (right) status:OptimalPart number:PN 15240-01Serial number:SN YPT054300407Vendor:VN ENGENIODate of manufacture:October 1, 2005

Power Supplies Detected: 2

Power supply status:	Optimal
Location:	Power supply Controller (left)
Part number:	PN 15240-01
Serial number:	SN YPT054300403
Vendor:	VN ENGENIO
Date of manufacture:	October 1, 2005
Power supply status:	Optimal
Location:	Power supply Controller (right)
Part number:	PN 15240-01
Serial number:	SN YPT054300407
Vendor:	VN ENGENIO
Date of manufacture:	October 1, 2005

Fans Detected: 2

Fan Status: Optimal Location: Power-fan Controller (left)

Fan Status: Optimal Location: Power-fan Controller (right)

Temperature Sensors Detected: 4

Temperature sensor status: Optimal Location: ESM Controller A (left)

Temperature sensor status: Optimal Location: ESM Controller B (right)

Temperature sensor status: Optimal Location: Power-fan Controller (left)

Temperature sensor status: Optimal Location: Power-fan Controller (right)

Drive Enclosure 4 Overall Component Information

Enclosure audible alarm: Disabled Enclosure path redundancy: OK Current drive types: Fibre channel, Fibre channel/Full Disk Encryption (FDE), Serial ATA (SATA) Fibre channel, Fibre channel/Full Disk Encryption (FDE), Serial ATA (SATA)

Part number:	PN 16490-00
Serial number:	SN 1T53023041
Vendor:	VN ENGENIO
Date of manufacture:	November 1, 2005

ESM Controllers Detected: 2

ESM card status: Optimal Firmware version: 98C3 Configuration settings version: FD 00.50 01/22/2008 Maximum data rate: 4 Gbps Current data rate: Not Available Location: A (left) Card communication: 0K Product ID: EXP5000 Part number: PN 39M5570 Serial number: SN 1T54136308 Vendor: ΙBΜ Date of manufacture: November 1, 2005

ESM card status: Optimal 98C3 Firmware version: Configuration settings version: FD 00.50 01/22/2008 Maximum data rate: 4 Gbps Current data rate: Not Available Location: B (right) Card communication: 0K Product ID: EXP5000 Part number: PN 39M5570 Serial number: SN 1T54136298 ΙBΜ Vendor: Date of manufacture: November 1, 2005

SFPs Detected: 2

SFP status:	Optimal
Attached to:	ESM Controller A (left)
Location:	Unknown
Supported data rate(s):	1 Gbps, 2 Gbps, 4 Gbps
Link length:	Intermediate
Connector:	LC
Transmitter type:	Shortwave Laser w/o OFC
Transmission media:	TM Multi-mode 62.5m(M6)
IEEE company ID:	00 90 65
Revision:	A
Part number:	FTLF8524P2BNV
Serial number:	PAC62UB
Vendor:	FINISAR CORP.
Date of manufacture:	September 23, 2006

SFP status:OptimalAttached to:ESM Controller B (right)Location:UnknownSupported data rate(s):1 Gbps, 2 Gbps, 4 Gbps

Intermediate
LC
Shortwave Laser w/o OFC
TM Multi-mode 62.5m(M6)
00 17 6a
Not Available
AFBR-57R5AEZ
A8063602FJ
AVAGO
September 6, 2006

Power-Fan Controllers Detected: 2

Power-fan Controller	(left)	status: Optimal
Part number:		PN 15240-05
Serial number:		SN YPT060110704
Vendor:		VN ENGENIO
Date of manufacture:		January 1, 2006

Power-fan Controller (right) status: Optimal Part number: PN 15240-05 Serial number: SN YPT060110720 Vendor: VN ENGENIO Date of manufacture: January 1, 2006

Power Supplies Detected: 2

Power supply status:OptimalLocation:Power supply Controller (left)Part number:PN 15240-05Serial number:SN YPT060110704Vendor:VN ENGENIODate of manufacture:January 1, 2006

Power supply status:OptimalLocation:Power supply Controller (right)Part number:PN 15240-05Serial number:SN YPT060110720Vendor:VN ENGENIODate of manufacture:January 1, 2006

Fans Detected: 2

Fan Status: Optimal Location: Power-fan Controller (left)

Fan Status: Optimal Location: Power-fan Controller (right)

Temperature Sensors Detected: 4

Temperature sensor status:	Optimal
Location:	ESM Controller A (left)
Temperature sensor status:	Optimal
Location:	ESM Controller B (right)
Temperature sensor status:	Optimal
Location:	Power-fan Controller (left)
Temperature sensor status:	Optimal
Location:	Power-fan Controller (right)

Controller Enclosure 85 Overall Component Information

Enclosure audible alarm:	Disabled
Part number:	PN 24643-02
Serial number:	SN SP80961772
Vendor:	VN LSI
Date of manufacture:	February 1, 2008

Interconnect-Battery Controllers Detected: 1

Interconnect-battery Controller status: Optimal Replacement part number: PN 24643-02 Serial number: SN SP80961772 Vendor: VN LSI Date of manufacture: February 1, 2008

Interconnect-Battery Controllers Detected: 1

Interconnect status: Optimal Location: Interconnect-battery Controller

Batteries Detected: 2

Battery status: Optimal Location: Interconnect-battery Controller (left pack) Age: 361 days Days until replacement: 3,329 days Last learn cycle: Not Availal Last learn cycle: Next learn cycle: Not Available February 4, 2009 Weeks between learn cycles: 8 Part number: PN 24543-03 Serial number: SN SE80100511SA VN LSI Vendor: Date of manufacture: January 1, 2008

timal
<pre>terconnect-battery Controller (right pack)</pre>
1 days
329 days
t Available
bruary 4, 2009
24543-05
SE81600979PS
LSI
ril 1, 2008

SFPs Detected: 8

SFP status:	Optimal
Attached to:	Drive-side of controller A
Location:	Channel 3, Port 1
Supported data rate(s):	1 Gbps, 2 Gbps, 4 Gbps
Link length:	Intermediate
Connector:	LC
Transmitter type:	Shortwave Laser w/o OFC
Transmission media:	TM Multi-mode 62.5m(M6)
IEEE company ID:	00 01 9c
Revision:	Not Available
Part number:	JSH-42S3AR3
Serial number:	F6336794105B
Vendor:	JDS UNIPHASE
Date of manufacture:	August 19, 2006
SFP status:	Optimal
Attached to:	Drive-side of controller A
Location:	Channel 4, Port 1
Supported data rate(s):	1 Gbps, 2 Gbps, 4 Gbps
Link length:	Intermediate
Connector:	LC
Transmitter type:	Shortwave Laser w/o OFC
Transmission media:	TM Multi-mode 62.5m(M6)
IEEE company ID:	00 01 9c
Revision:	Not Available
Part number:	JSH-42S3AR3
Serial number:	F6336794060C
Vendor:	JDS UNIPHASE
Date of manufacture:	August 19, 2006
SFP status:	Optimal
Attached to:	Host-side of controller A
Location:	Channel 5
Supported data rate(s):	1 Gbps, 2 Gbps, 4 Gbps
Link length:	Intermediate
Connector:	LC
Transmitter type:	Shortwave Laser w/o OFC
Transmission media:	TM Multi-mode 62.5m(M6)
IEEE company ID:	00 04 85
Revision:	1
Part number:	PLRXPLVESG462N
Serial number:	C807SP1NN
Vendor:	PICOLIGHT
Date of manufacture:	February 15, 2008

SFP st Attach Locati Suppor Link Connec Transm Transm IEEE o Revisi Part m Serial Vendor Date o	tatus: ned to: ion: rted data rate length: ctor: nitter type: nission media company ID: ion: number: I number: r: of manufacture	e(s): : e:	Optimal Host-side Channel 1 1 Gbps, 2 Intermedia LC Shortwave TM Multi- 00 17 6a Not Availa AFBR-57R5 A20650500 AVAGO December	of controlle Gbps, 4 Gbps ate Laser w/o OF mode 62.5m(M6 able AQZ 5 14, 2006	r A C )		
Power-Fan Cont	trollers Deteo	cted:	2				
Power- Part r Serial Vendor Date c Power-	: Optimal PN 24744-01 SN I9T074300 DT 10/2007 Not Availabl s: Optimal	543 e					
Part r Serial Vendor Date c	number: l number: r: of manufacture	e:	,	PN 24744-01 SN I9T07430 DT 10/2007 Not Availab	0541 le		
Power Supplies	s Detected: 2	2					
	Power supply Location:	status	: Optima Power-	l fan Controlle	r (left)		
	Power supply Location:	status	: Optima Power-	l fan Controlle	r (right)		
Fans Detected: 4							
	Fan Status: Location:	Optima Power-	l fan Contro	oller (left)			
	Fan Status: Location:	Optima Power-	l fan Contro	oller (left)			
	Fan Status: Location:	Optima Power-	l fan Contro	oller (right)			

Fan Status: Optimal Location: Power-fan Controller (right)

Temperature Senso	ors Detected:	2					
Temperati Location:	ure sensor sta :	atus: Optima Power-	l fan Controller (left)				
Temperatu Location:	ure sensor sta :	atus: Optima Power-	ll fan Controller (right)				
MAPPINGS (Storage	Partitioning	- Enabled (2	2 of 2 used))				
Logical [ Drive status	Drive Name	LUN Control	ler Accessible by	Logical			
Access Lo FC_R3_A FC_R3_B FDE_R1 Sata_Raid Sata_Paid	bgical Drive	<ul> <li>31 A,B</li> <li>0 A</li> <li>1 B</li> <li>4 B</li> <li>2 A</li> <li>3 B</li> </ul>	Host Left Host Left Host Left Host Left Host Left	Optimal Optimal Optimal Optimal Optimal			
Access Lo FDE_R6_A FDE_R6_B Secure_M	iddle_R6_A	31 A,B 0 A 1 B 2 A	Host Right Host Right Host Right Host Right	Optimal Optimal Optimal Optimal			
Secure_M Access Lo	iddle_R6_B ogical Drive	3 B 31 A,B	Host Right Storage Subsystem	Optimal Optimal			
TOPOLO	)GY DEFINITION	٧S					
ST( Non-Clustered	DRAGE SUBSYSTE Default type:	EM :	Windows 2000/Server 20	03/Server 2008			
	Default Group	0					
Non-Clustered	Host: Host type:	:	Right Windows 2000/Server 20	03/Server 2008			
	Host port Alias: Host type	identifier:	21:00:00:e0:8b:95:c7:1e QLogic_1				
Non-Clustered	Host port	identifier:	21:01:00:e0:8b:b5:c7:1	е			
Non-Clustered	Host type:	:	Windows 2000/Server 20	03/Server 2008			
	Host port Alias: Host:	identifier:	10:00:00:05:1e:56:67:9 Brocade1 Left	8			
Non-Clustered	Host type:	: identifier:	Windows 2000/Server 20	03/Server 2008			
N 01	Alias: Host type:	:	Qlogic1 Windows 2000/Server 20	03/Server 2008			
Non-Clustered	Host port Alias:	identifier:	21:01:00:e0:8b:a3:fd:c Qlogic2	с			

NVSRAM HOST TYPE DEFINITIONS

NOTE: The following indexes are not used: 16, 18 - 31

	ADT STATUS	
AIX	Disabled	6
AIX-ADT/AVT	Enabled	4
DEFAULT	Disabled	0
HP-UX	Enabled	7
HPXTPGS	Disabled	17
IBM TS SAN VCE	Enabled	12
Irix	Disabled	10
LNXCLVMWARE	Disabled	13
Linux	Enabled	5
Mac OS X	Enabled	1
Netware Failover	Enabled	11
Solaris (with Veritas DMP)	Enabled	14
Solaris (with or without MPXIO)	Disabled	8
Windows 2000/Server 2003/Server 2008 Clustered	Disabled	3
Windows 2000/Server 2003/Server 2008 Clustered (supports DMP)	Enabled	15
Windows 2000/Server 2003/Server 2008 Non-Clustered	Disabled	2
Windows 2000/Server 2003/Server 2008 Non-Clustered (supports DMP)	Enabled	9

# Show Controller NVSRAM

The **show controller NVSRAM** command returns a table of values in the controller NVSRAM that is similar to that shown in this example. With the information from the table, you can modify the contents of the NVSRAM by using the **set controller** command. This example shows information for a controller in slot A in a controller module. You can produce a similar table for a controller in slot B, or you can produce a table for both controllers.

**Attention:** Incorrectly setting the NVSRAM can severely affect the performance and the availability of the subsystems. You should only modify the NVSRAM settings according to the latest information documented for your configuration in DS4000 publications, readmes, or by your IBM support representative.

(	0010:	5500	0000	0000	0000	0000	0000	0000	0000	U
Regior 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6	233 0000: 0010: 0020: 0030: 0040: 0050:	: 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000	
Regior 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6	n 234 0000: 0010: 0020: 0030: 0040: 0050: 0060: 0070:	: 1020 1121 1222 1323 1424 0000 0000 0000	3040 3141 3242 3343 3444 0000 0000 0000	5060 5161 5262 5363 5464 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000 0000	0.P. 1AQa 2BRb 3CSc 4DTd
Regior 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6	n 236 )000: )010: )020: )030: )040: )050: )060: )060:	: 0000 0000 0000 0000 0000 0000 0000 0	0000 0000 0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000 0000 0000	0000 0000 0000 0000 0000 0000 0000 0000	
Regior (	n 237 0000: 0010:	: 0000 0000	0000 0000	0000 007d	0000 0d0e	007d 0f00	0000 0000	0000 0000	0000 0000	
Regior ( ( ( (	n 238 )000: )010: )020: )030:	: 0000 0000 0000 80be	ca32 0000 0000 9f45	0000 0000 0000 7300	0000 0000 0000 2000	0050 f001 cc8c 0f00	0600 0000 008a 1400	0000 8480 0029 0000	0000 0000 fe00 0000	2P Es
Region 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6	n 240 )000: )010: )020: )030: )040: )040: )050: )040: )0	: 0c12 3030 fc01 0183 04c0 6f75 00000 00000 0000 0000 0000 0000 00000 0000 0000 0000 0000	4453 5f41 83b5 a846 73a2 0000 0000 0000 0000 0000 0000 0000 0	3530 ab04 04c0 726f 01fb 0870 0000 0000 0000 0000 0000 0000 000	3030 c0a8 a846 6269 0180 6173 0000 0000 0000 0000 0000 0000 0000 0	5f41 8032 32bf 6eac a109 7377 0000 0000 0000 0000 0000 0000 00	7f7f 0104 04ff 04c0 616e 6f72 0000 0000 0000 0000 0000 0000 0000 0	7f7f ffff a801 6f6e 64ff 0000 0000 0000 0000 0000 0000 0000	7f33 ff00 00fd 0103 796d 0000 0000 0000 0000 0000 0000 0000 0	DS5000.A3 002 F2 robin Fanonym ouspassword
Regior 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6 6	241 0000: 0010: 0020: 0030: 0040: 0050: 0060: 0070:	: 0000 0000 80be 4465 0000 0000 0000	ca32 0000 9f45 6661 0000 0000 0000	0000 0000 7300 756c 0000 0000 0000	0000 0000 2000 7400 0000 0000 0000	0050 f001 cc8c 0f00 0000 0000 0000 0000	0600 0000 008a 1400 0000 0000 0000 0000	0000 8480 0029 0000 0000 0000 0000 0000	0000 0000 fe00 0000 0000 0000 0000 0000	2P Es Default

0080: 0090:	0000 0000	0000 0000	0000 0000	0000 0000	0000 0000	0000 0000	0000 0000	0000 0000	
Region 242 0000: 0010: 0020: 0030:	: 4445 0020 0001 0000	4641 0000 0100 0000	554c 0100 0000 0000	5400 0001 0008 0000	0000 0000 0100 0000	0000 0000 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	DEFAULT
Region 242 0000: 0010: 0020: 0030:	: 4d61 017f 0000 0000	6320 0000 0100 0000	4f53 0100 0100 0000	2058 0001 0008 0000	0000 0001 0100 0000	0000 0000 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	Mac.OS.X
Region 242 0000: 0010: 0020: 0030:	: 5732 0020 0001 0000	4b4e 0000 0100 0000	4554 0100 0000 0000	4e43 0001 0008 0000	4c00 0000 0100 0000	0000 0000 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	W2KNETNCL
Region 242 0000: 0010: 0020: 0030:	: 5732 0020 0001 0000	4b4e 0000 0100 0000	4554 0100 0000 0000	434c 0001 0008 0000	0000 0001 0100 0000	0000 0001 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	W2KNETCL
Region 242 0000: 0010: 0020: 0030:	: 4149 0100 0000 0000	582d 0000 0100 0000	4144 0100 0100 0000	542f 0001 0008 0000	4156 0100 0100 0000	5400 0000 0000 0000	0000 0001 0000 0000	0000 0000 0000 0000	AIX.ADT.AVT
Region 242 0000: 0010: 0020: 0030:	+c4e 017f 0000 0000	5800 0000 0100 0000	0000 0100 0100 0000	0000 0001 0008 0000	0000 0001 0100 0000	0000 0000 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	LNX
Region 242 0000: 0010: 0020: 0030:	: 4149 0100 0000 0000	5800 0000 0100 0000	0000 0100 0000 0000	0000 0001 0008 0000	0000 0100 0100 0000	0000 0000 0000 0000	0000 0001 0000 0000	0000 0000 0000 0000	AIX
Region 242 0000: 0010: 0020: 0030:	: 4850 017f 0100 0000	5800 0000 0000 0000	0000 0100 0100 0000	0000 0001 001c 0000	0000 0101 0100 0000	0000 0001 0000 0000	0000 0101 0000 0000	0000 0001 0000 0000	HPX
Region 242 0000: 0010: 0020: 0030:	: 534f 0120 0000 0000	4c00 0000 0000 0000	0000 0100 0000 0000	0000 0001 0008 0000	0000 0001 0100 0000	0000 0000 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	SOL
Region 242 0000: 0010: 0020: 0030:	: 5732 0020 0000 0000	4b4e 0000 0100 0000	4554 0100 0180 0000	4e43 0001 0008 0000	4c44 0000 0100 0000	4d50 0100 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	W2KNETNCLDMP
Region 242 0000: 0010: 0020:	: 4952 007f 0000	5800 0000 0100	0000 0100 0000	0000 2001 0008	0000 0000 0100	0000 0000 0000	0000 0000 0000	0000 0000 0000	IRX

	0030:	0000	0000	0000	0000	0000	0000	0000	0000	•••••
Regi	on 242 0000: 0010: 0020: 0030:	4e57 007f 0100 0000	5246 0000 0100 0000	4f00 0100 0100 0000	0000 0001 0008 0000	0000 0000 0100 0000	0000 0000 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	NWRF0
Regi	on 242 0000: 0010: 0020: 0030:	: 4942 0000 0000 0000	4d20 0000 0100 0000	5453 0100 0100 0000	2053 0000 0008 0000	414e 0000 0100 0000	2056 0001 0000 0000	4345 0000 0000 0000	0000 0000 0000 0000	IBM.TS.SAN.VCE.
Regi	on 242 0000: 0010: 0020: 0030:	4c4e 017f 0000 0000	5843 0000 0100 0000	4c56 0100 0000 0000	4d57 0001 0008 0000	4152 0001 0100 0000	4500 0001 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	LNXCLVMWARE
Regi	on 242 0000: 0010: 0020: 0030:	534f 0120 0000 0000	4c41 0000 0001 0000	5654 0100 0180 0000	0000 0001 0000 0000	0000 0001 0100 0000	0000 0100 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	SOLAVT
Regi	on 242 0000: 0010: 0020: 0030:	: 5732 0020 0000 0000	4b4e 0000 0101 0000	4554 0100 0180 0000	434c 0001 0008 0000	444d 0001 0100 0000	5000 0101 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	W2KNETCLDMP
Regi	on 242 0000: 0010: 0020: 0030:	: 0075 0000 0000 0000	6e6e 0000 0000 0000	616d 0000 0000 0000	6564 0000 0000 0000	2031 0000 0000 0000	3600 0000 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	.unnamed.16
Regi	on 242 0000: 0010: 0020: 0030:	: 4850 017f 0001 0000	5854 0000 0000 0000	5047 0100 0000 0000	5300 0001 0008 0000	0000 0101 0200 0000	0000 0001 0000 0000	0000 0101 0000 0000	0000 0000 0000 0000	HPXTPGS
Regi	on 242 0000: 0010: 0020: 0030:	: 0075 0000 0000 0000	6e6e 0000 0000 0000	616d 0000 0000 0000	6564 0000 0000 0000	2031 0000 0000 0000	3800 0000 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	.unnamed.18
Regi	on 242 0000: 0010: 0020: 0030:	: 0075 0000 0000 0000	6e6e 0000 0000 0000	616d 0000 0000 0000	6564 0000 0000 0000	2031 0000 0000 0000	3900 0000 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	.unnamed.19
Regi	on 242 0000: 0010: 0020: 0030:	: 0075 0000 0000 0000	6e6e 0000 0000 0000	616d 0000 0000 0000	6564 0000 0000 0000	2032 0000 0000 0000	3000 0000 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	.unnamed.20
Regi	on 242 0000: 0010: 0020: 0030:	: 0075 0000 0000 0000	6e6e 0000 0000 0000	616d 0000 0000 0000	6564 0000 0000 0000	2032 0000 0000 0000	3100 0000 0000 0000	0000 0000 0000 0000	0000 0000 0000 0000	.unnamed.21
Region 242: 0000: 0075 6e6e 616d 6564 2032 3200 0000 0000 .unnamed.22..... Region 242: 0000: 0075 6e6e 616d 6564 2032 3300 0000 0000 .unnamed.23.... Region 242: 0000: 0075 6e6e 616d 6564 2032 3400 0000 0000 .unnamed.24.... Region 242: 0000: 0075 6e6e 616d 6564 2032 3500 0000 0000 .unnamed.25..... Region 242: 0000: 0075 6e6e 616d 6564 2032 3600 0000 0000 .unnamed.26.... Region 242: 0000: 0075 6e6e 616d 6564 2032 3700 0000 0000 .unnamed.27.... Region 242: 0000: 0075 6e6e 616d 6564 2032 3800 0000 0000 .unnamed.28.... Region 242: 0000: 0075 6e6e 616d 6564 2032 3900 0000 0000 .unnamed.29.... Region 242: 0000: 0075 6e6e 616d 6564 2033 3000 0000 0000 .unnamed.30..... Region 242: 0000: 0075 6e6e 616d 6564 2033 3100 0000 0000 .unnamed.31.... .

## **Show Drive**

The **show drive** command returns information about the disk drives in a storage subsystem.

DRIVES				-								
SUMMARY Number of Current	drives: 32 drive typ	es: Fibre	e (16	5), Fib	re/FD	E	(14),	Serial	ATA	(SATA	) (2)	
BASIC:												
enclosure,	SLOT STA	TUS CAP	ACI	רא די	YPE		CI	JRRENT I	DATA	RATE	PRODUCT	ID
5, 1	Optimal	419.187	GB	Fibre		4	Gbps			HUS1	545450FC	
5,2	Optimal F097	136.732	GB	Fibre,	FDE	4	Gbps			ST31	46756FC	F
5,3	Optimal JG01	419.187	GB	Fibre		4	Gbps			HUS1	545450FC	
5,4	Optimal E097	279.397	GB	Fibre,	FDE	4	Gbps			ST33	00056FC	F
5,5	Optimal JG01	419.187	GB	Fibre		4	Gbps			HUS1	545450FC	
5,6	Optimal E097	136.732	GB	Fibre,	FDE	4	Gbps			ST31	46756FC	F
5,7	Optimal JG01	419.187	GB	Fibre		4	Gbps			HUS1	545450FC	
5,8	Optimal E097	419.187	GB	Fibre,	FDE	4	Gbps			ST34	50056FC	F
5,9 43W9715 42C04	Optimal 17IBM BB	698.638 10	GB	SATA		4	Gbps			ST37	50330NS	
5, 10	Optimal JG01	279.397	GB	Fibre		4	Gbps			HUS1	545300FC	
5, 11	Optimal JG01	419.187	GB	Fibre		4	Gbps			HUS1	545450FC	
5, 12	Optimal JG01	279.397	GB	Fibre		4	Gbps			HUS1	545300FC	
5,13	Optimal E3BE	419.187	GB	Fibre,	FDE	4	Gbps			S5EG	450	
5, 14	Optimal E097	279.397	GB	Fibre,	FDE	4	Gbps			ST33	00056FC	F
5, 15	Optimal E097	279.397	GB	Fibre,	FDE	4	Gbps			ST33	00056FC	F
5, 16	Optimal E097	279.397	GB	Fibre,	FDE	4	Gbps			ST33	00056FC	F
7,1	Optimal E097	279.397	GB	Fibre,	FDE	4	Gbps			ST33	00056FC	F
7,2	Optimal E097	279.397	GB	Fibre,	FDE	4	Gbps			ST33	00056FC	F
7,3	Optimal E097	279.397	GB	Fibre,	FDE	4	Gbps			ST33	00056FC	F
7,4	Optimal E097	279.397	GB	Fibre,	FDE	4	Gbps			ST33	00056FC	F
7,5	Optimal JG01	419.187	GB	Fibre		4	Gbps			HUS1	545450FC	
7,6	Optimal JG01	279.397	GB	Fibre		4	Gbps			HUS1	545300FC	
7,7	Optimal JG01	419.187	GB	Fibre		4	Gbps			HUS1	545450FC	
7,8 43W9715 42C04	Optimal 17IBM BB	698.638 10	GB	SATA		4	Gbps			ST37	50330NS	
7,9	Optimal E097	279.397	GB	Fibre,	FDE	4	Gbps			ST33	00056FC	F
7, 10	Optimal JG01	419.187	GB	Fibre		4	Gbps			HUS1	545450FC	
7,11	Optimal JG01	419.187	GB	Fibre		4	Gbps			HUS1	545450FC	
7,12	Optimal E097	279.397	GB	Fibre,	FDE	4	Gbps			ST33	00056FC	F
7,13	Optimal	419.187	GB	Fibre		4	Gbps			HUS1	545450FC	

		.1601				
7,	14	Optimal JG01	279.397 GB	Fibre	4 Gbps	HUS1545300FC
7,	15	Optimal JG01	419.187 GB	Fibre	4 Gbps	HUS1545450FC
7,	16	Optimal JG01	279.397 GB	Fibre	4 Gbps	HUS1545300FC

DRIVE CHANNELS:

enclosure 5, 1 5, 2 5, 3 5, 4 5, 5 5, 6 5, 7 5, 8 5, 9 5, 10 5, 11 5, 12 5, 13 5, 14 5, 15 5, 16 7, 1 7, 2 7, 3 7, 4 7, 5 7, 6 7, 7 7, 8 7, 9 7, 10 7, 11 7, 12 7, 13 7, 14 7, 15 7, 16	, SLOT 5 3 5 3 5 3 5 3 5 3 5 3 5 3 5 3 5 3 5	PREFERRED	CHANNEL 3 5 3 5 3 5 3 5 3 5 3 5 3 5 3 5 3 5 3	REDUNDA	ANT CHANI	NEL.
HOT SPARE COVER The followin Total hot sp	AGE: g subsys are driv	stems are n ves: 5	not prote	cted: Sa	ata_R1, I	FDE_R0
In use:	0					
Standby driv	e at end	closure 5,	slot 13	(Fibre,	419.187	GB)
Protects	the foll	lowing subs	systems:	FDE_R6,	Mix_R5,	FC_R3
Standby driv	e at end	closure 7,	slot 12	(Fibre,	279.397	GB)
Protects	the fol <sup>1</sup>	lowing subs	systems:	FDE_R6,	FC_R3	
Standby driv	e at end	closure 7,	slot 11	(Fibre,	419.187	GB)
Protects	the fol <sup>1</sup>	lowing subs	systems:	FDE_R6,	Mix_R5,	FC_R3
Standby driv	e at end	closure 5,	slot 11	(Fibre,	419.187	GB)
Protects	the fol <sup>1</sup>	lowing subs	systems:	FDE_R6,	Mix_R5,	FC_R3
Standby driv	e at end	closure 7,	slot 10	(Fibre,	419.187	GB)
Protects	the fol'	lowing subs	systems:	FDE_R6,	Mix_R5,	FC_R3

DETAILS Drive at Enclosure 5, Slot 1

Status: Optimal Mode: Assigned 419.187 GB Raw capacity: Usable capacity: 418.687 GB World-wide identifier: 50:00:cc:a0:08:02:93:00:00:00:00:00:00:00:00:00 Associated subsystem: Mix R5 Port Channel ID 0 5 0/0xEF 1 3 0/0xEF Security Capable: No Secure: No Read/write accessible: Yes Security key identifier: Not Applicable 15,000 RPM Speed: Current data rate: 4 Gbps Product ID: HUS1545450FC JG01 Firmware version: Serial number: JMV1DWJC IBM-SSG Vendor: Date of manufacture: July 4, 2008 Drive at Enclosure 5, Slot 2 Status: Optimal Mode: Assigned 136.732 GB Raw capacity: 136.232 GB Usable capacity: World-wide identifier: 20:00:00:1d:38:1d:1d:74:00:00:00:00:00:00:00:00 Associated subsystem: FDE R0 Port Channe] ID 1/0xE8 0 3 1 5 1/0xE8 Security Capable: Yes Secure: Yes Read/write accessible: Yes 2700000600A0B800029EE020000C8594941050A Security key identifier: Speed: 15.015 RPM Current data rate: 4 Gbps Product ID: ST3146756FC F Firmware version: E097 Serial number: 3QN0897K00009912TLN9 Vendor: IBM-SSG Date of manufacture: October 16, 2008 Drive at Enclosure 5, Slot 9

Status:

Optimal

Mode: Raw capaci Usable cap World-wide Associated	ity: bacity: e identifier: d subsystem:	Ass 698 698 209	signed 3.638 GB 3.138 GB :00:00:a0:b8:41:20:c8 Sata_R1	:00:00:00	0:00:00:00:00:00
Port O 1	Channel 5 3	ID 8/0xI 8/0xI	09 09		
Security ( Secure: Read/write Security H	Capable: e accessible: key identifien	r:	No No Yes Not Applicable		
Speed: Current da Product II Package ve Firmware v Serial num Vendor: Date of ma	ata rate: ): ersion: version: nber: anufacture:		7,200 RPM 4 Gbps ST3750330NS BB10 BB10 5QK04MDK Not Available Not Available	43W9715	42C0417IBM
ATA Trans Product II Vendor: Firmware V Location:	lator ): /ersion:		BR-2401-3.0 SLI LP1158 Individual Drive Cont	troller	

## **Show Drive Channel Status**

The **show drive channel stat** command returns information about the drive channels in a storage subsystem. Use this information to determine how well the channels are running and errors that might be occurring on the channels. DRIVE CHANNELS------

SUMMARY

CHANNEL PORT STATUS CTRL A LINK CTRL B LINK 1 8,7 Optimal Up Up 2 6,5 Optimal Up Up 3 4,3,Out Optimal Up Up 4 2,1,Out Optimal Up Up 5 1,2,Out Optimal Up Up 6 3,4,Out Optimal Up Up Optimal Up 7 5,6 Up 8 Optimal Up 7,8 Up

DETAILS

DRIVE CHANNEL 1

Port: 8, 7 Status: Optimal Max. Rate: 4 Gbps Current Rate: 4 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up DRIVE COUNTS

Total # of attached drives: 0

CUMULATIVE ERROR COUNTS

Controller A

Baseline time set:	1/22/09	6:51:12 AM
Sample period (days, hh:mm:ss):	5 days,	08:31:15
Controller detected errors:	0	
Drive detected errors:	0	
Timeout errors:	0	
Link down errors:	N/A	
Total I/O count:	314641	

Controller B

Baseline time set:	1/22/09	6:51:11 AM
Sample period (days, hh:mm:ss):	5 days,	08:31:20
Controller detected errors:	0	
Drive detected errors:	0	
Timeout errors:	0	
Link down errors:	N/A	
Total I/O count:	315842	

DRIVE CHANNEL 2

Port: 6, 5 Status: Optimal Max. Rate: 4 Gbps Current Rate: 4 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up

DRIVE COUNTS

Total # of attached drives: 0

CUMULATIVE ERROR COUNTS

#### Controller A

Baseline time set:	1/22/09	6:51:12 AM
Sample period (days, hh:mm:ss):	5 days,	08:31:15
Controller detected errors:	0	
Drive detected errors:	0	
Timeout errors:	0	
Link down errors:	N/A	
Total I/O count:	314640	

Controller B

Baseline time set:	1/22/09 6:51:11 AM
Sample period (days, hh:mm:ss):	5 days, 08:31:20
Controller detected errors:	Θ
Drive detected errors:	0
Timeout errors:	Θ
Link down errors:	N/A
Total I/O count:	315844

DRIVE CHANNEL 3

```
Port: 4, 3, Out
Status: Optimal
   Max. Rate: 4 Gbps
   Current Rate: 4 Gbps
   Rate Control: Auto
   Controller A link status: Up
   Controller B link status: Up
   DRIVE COUNTS
      Total # of attached drives: 16
      Connected to: A, Port 4
         Attached drives: 16
            Drive enclosure: 5 (16 drives)
   CUMULATIVE ERROR COUNTS
      Controller A
                                          1/22/09 6:51:12 AM
         Baseline time set:
         Sample period (days, hh:mm:ss): 5 days, 08:31:15
         Controller detected errors:
                                         1
         Drive detected errors:
                                         17
         Timeout errors:
                                          0
         Link down errors:
                                         N/A
         Total I/O count:
                                          11543881
      Controller B
                                          1/22/09 6:51:11 AM
         Baseline time set:
         Sample period (days, hh:mm:ss): 5 days, 08:31:20
         Controller detected errors:
                                          0
         Drive detected errors:
                                          17
         Timeout errors:
                                          0
         Link down errors:
                                          N/A
         Total I/O count:
                                         11819839
DRIVE CHANNEL 4
   Port: 2, 1, Out
   Status: Optimal
  Max. Rate: 4 Gbps
   Current Rate: 4 Gbps
   Rate Control: Auto
   Controller A link status: Up
   Controller B link status: Up
   DRIVE COUNTS
      Total # of attached drives: 16
      Connected to: A, Port 2
         Attached drives: 16
           Drive enclosure: 7 (16 drives)
   CUMULATIVE ERROR COUNTS
      Controller A
         Baseline time set:
                                          1/22/09 6:51:12 AM
         Sample period (days, hh:mm:ss): 5 days, 08:31:15
         Controller detected errors:
                                          0
         Drive detected errors:
                                          17
         Timeout errors:
                                          0
         Link down errors:
                                          N/A
         Total I/O count:
                                          18518855
```

Controller B

Baseline time set:	1/22/09 6:51:11 AM
Sample period (days, hh:mm:ss):	5 days, 08:31:20
Controller detected errors:	0
Drive detected errors:	17
Timeout errors:	0
Link down errors:	N/A
Total I/O count:	17271049

DRIVE CHANNEL 5

Port: 1, 2, Out Status: Optimal Max. Rate: 4 Gbps Current Rate: 4 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up

DRIVE COUNTS

Total # of attached drives: 16 Connected to: B, Port 2 Attached drives: 16 Drive enclosure: 5 (16 drives)

CUMULATIVE ERROR COUNTS

Controller A

Baseline time set:	1/22/09 6:51:12 AM
Sample period (days, hh:mm:ss):	5 days, 08:31:15
Controller detected errors:	0
Drive detected errors:	18
Timeout errors:	0
Link down errors:	N/A
Total I/O count:	12263553

Controller B

Baseline time set:	1/22/09 6:51:11 AM
Sample period (days, hh:mm:ss):	5 days, 08:31:20
Controller detected errors:	Θ
Drive detected errors:	18
Timeout errors:	0
Link down errors:	N/A
Total I/O count:	10936652

DRIVE CHANNEL 6

Port: 3, 4, Out Status: Optimal Max. Rate: 4 Gbps Current Rate: 4 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up

DRIVE COUNTS

Total # of attached drives: 16 Connected to: B, Port 4 Attached drives: 16 Drive enclosure: 7 (16 drives)

#### CUMULATIVE ERROR COUNTS

Controller A

Baseline time set:	1/22/09	6:51:12 AM
Sample period (days, hh:mm:ss):	5 days,	08:31:15
Controller detected errors:	0	
Drive detected errors:	17	
Timeout errors:	0	
Link down errors:	N/A	
Total I/O count:	8459801	

### Controller B

Baseline time set:	1/22/09 6:51:11 AM
Sample period (days, hh:mm:ss):	5 days, 08:31:20
Controller detected errors:	Θ
Drive detected errors:	17
Timeout errors:	0
Link down errors:	N/A
Total I/O count:	8735592

#### DRIVE CHANNEL 7

```
Port: 5, 6
Status: Optimal
Max. Rate: 4 Gbps
Current Rate: 4 Gbps
Rate Control: Auto
Controller A link status: Up
Controller B link status: Up
```

## DRIVE COUNTS

Total # of attached drives: 0

#### CUMULATIVE ERROR COUNTS

#### Controller A

Baseline time set:	1/22/09 6:51:12 AM
Sample period (days, hh:mm:ss):	5 days, 08:31:15
Controller detected errors:	0
Drive detected errors:	0
Timeout errors:	0
Link down errors:	N/A
Total I/O count:	314636

#### Controller B

Baseline time set:	1/22/09 6:51:	11 AM
Sample period (days, hh:mm:ss):	5 days, 08:31	:20
Controller detected errors:	Θ	
Drive detected errors:	Θ	
Timeout errors:	Θ	
Link down errors:	N/A	
Total I/O count:	315839	

DRIVE CHANNEL 8

Port: 7, 8 Status: Optimal Max. Rate: 4 Gbps Current Rate: 4 Gbps

Rate Control: Auto Controller A link status: Up Controller B link status: Up		
DRIVE COUNTS		
Total # of attached drives: 0		
CUMULATIVE ERROR COUNTS		
Controller A		
Baseline time set: Sample period (days, hh:mm:ss): Controller detected errors: Drive detected errors: Timeout errors: Link down errors: Total I/O count:	1/22/09 5 days, 0 0 0 N/A 314636	6:51:12 AM 08:31:15
Controller B		
Baseline time set: Sample period (days, hh:mm:ss): Controller detected errors: Drive detected errors: Timeout errors: Link down errors: Total I/O count:	1/22/09 5 days, 0 0 0 N/A 315839	6:51:11 AM 08:31:20

# **Show Logical Drives**

The **show logicalDrives** command returns information about the logical drives in a storage subsystem.

STANDARD LOGICAL DRIVES------

SUMMARY

Number of standard logical drives: 8

See other Logical Drives sub-tabs for premium feature information.

NAME	STATUS	CAPACITY	RAID LEVEL	subsyste	m DRIVE TYPE
FC_R3_A	Optimal	500.0 GB	3	FC_R3	Fibre channel
FC_R3_B	Optimal	550.0 GB	3	FC_R3	Fibre channel
DE_R6_A	Optimal	400.0 GB	6	FDE_R6	Fibre, FDE
FDE_R6_E	3 Optimal	440.0 GB	6	FDE_R6	Fibre, FDE
Mix_R5_/ /Full Di	A Optimal isk Encrypti	500.0 GB on(FDE)	5	Mix_R5	Fibre channel, Fibre channel
Mix_R5_E	B Optimal	520.0 GB	5	Mix_R5	Fibre channel, Fibre channel
/Full Di	isk Encrypti	on(FDE)			
Sata_R1_	_A Optimal	500.0 GB	1	Sata_R1	Serial ATA (SATA)
Sata_R1_	_B Optimal	190.0 GB	1	Sata_R1	Serial ATA (SATA)

Logical Drive name:	FC_R3_A		
Logical Drive status:	Optimal		
Capacity: Logical Drive ID: Subsystem ID (SSID): Associated subsystem: RAID level:	500.0 GB 60:0a:0b:8 6 FC_R3 3	0:00:29:ee:02	:00:00:c8:64:49:41:0a:7a
Secure:	No		
Drive type: Enclosure loss protection:	Fibre chan No	nel	
Preferred owner: Current owner:	Controller Controller	in slot A in slot A	
Segment size: Capacity reserved for futur Maximum future segment size Modification priority:	e segment s :	ize changes:	128 KB Yes 2,048 KB High
Read cache: Write cache: Write cache without batt Write cache with mirrori Flush write cache after (in Dynamic cache read prefetch	eries: ng: seconds): :	Enabled Enabled Disabled Enabled 10.00 Enabled	
Enable background media sca Media scan with redundancy	n: check:	Enabled Disabled	
Pre-Read redundancy check:		Disabled	
Logical Drive name:	FC_R3_B		
Logical Drive status:	Optimal		
Capacity: Logical Drive ID: Subsystem ID (SSID): Associated subsystem: RAID level:	550.0 GB 60:0a:0b:8 7 FC_R3 3	0:00:33:98:48	:00:00:c9:76:49:41:0c:4c
Secure:	No		
Drive type: Enclosure loss protection:	Fibre chan No	nel	
Preferred owner: Current owner:	Controller Controller	in slot B in slot B	
Segment size: Capacity reserved for futur Maximum future segment size Modification priority:	e segment s :	ize changes:	128 KB Yes 2,048 KB High
Read cache: Write cache: Write cache without batt	eries:	Enabled Enabled Disabled	

	Write cache with mirrori Flush write cache after (in Dynamic cache read prefetch	ng: seconds): :	Enabled 10.00 Enabled	
	Enable background media sca Media scan with redundancy	n: check:	Enabled Disabled	
	Pre-Read redundancy check:		Disabled	
	Logical Drive name:	FDE_R6_	B	
	Logical Drive status:	Optimal		
	Capacity: Logical Drive ID: Subsystem ID (SSID): Associated subsystem: RAID level:	440.0 GB 60:0a:0b:8 2 FDE_R6 6	0:00:29:ee:02	:00:00:c8:5b:49:41:06:db
	Secure:	Yes		
	Drive type: Enclosure loss protection:	Fibre Chan No	nel/Full Disk	Encryption(FDE)
	Preferred owner: Current owner:	Controller Controller	in slot B in slot B	
	Segment size: Capacity reserved for futury Maximum future segment size Modification priority:	e segment s :	ize changes:	128 KB Yes 2,048 KB High
	Read cache: Write cache: Write cache without batt Write cache with mirrori Flush write cache after (in Dynamic cache read prefetch	eries: ng: seconds): :	Enabled Enabled Disabled Enabled 10.00 Enabled	
	Enable background media sca Media scan with redundancy	n: check:	Enabled Disabled	
	Pre-Read redundancy check:		Disabled	
Log	gical Drive name:	Mix_R5_A		
	Logical Drive status:	Optimal		
	Capacity: Logical Drive ID: Subsystem ID (SSID): Associated subsystem: RAID level:	500.0 GB 60:0a:0b:8 3 Mix_R5 5	0:00:33:98:48	:00:00:c9:70:49:41:09:a6
	Secure:	No		
	Drive type:	Fibre chan	nel, Fibre ch	annel/Full Disk Encryption
	(FUE) Enclosure loss protection:	No		
	Preferred owner: Current owner:	Controller Controller	in slot A in slot A	

	Segment size: Capacity reserved for futur Maximum future segment size Modification priority:	e segment s :	ize changes:	128 KB Yes 2,048 KB High
	Read cache: Write cache: Write cache without batt Write cache with mirrori Flush write cache after (in Dynamic cache read prefetch	eries: ng: seconds): :	Enabled Enabled Disabled Enabled 10.00 Enabled	
	Enable background media sca Media scan with redundancy	n: check:	Enabled Disabled	
	Pre-Read redundancy check:		Disabled	
Lo	gical Drive name:	Mix_R5_B		
	Logical Drive status:	Optimal		
	Capacity: Logical Drive ID: Subsystem ID (SSID): Associated subsystem: RAID level:	520.0 GB 60:0a:0b:8 4 Mix_R5 5	0:00:29:ee:02	:00:00:c8:5e:49:41:09:04
	Secure:	No		
(	Drive type: FDE)	Fibre chan	nel, Fibre ch	annel/Full Disk Encryption
	Enclosure loss protection:	No		
	Preferred owner: Current owner:	Controller Controller	in slot B in slot B	
	Segment size: Capacity reserved for futur Maximum future segment size Modification priority:	e segment s :	ize changes:	128 KB Yes 2,048 KB High
	Read cache: Write cache: Write cache without batt Write cache with mirrori Flush write cache after (in Dynamic cache read prefetch	eries: ng: seconds): :	Enabled Enabled Disabled Enabled 10.00 Enabled	
	Enable background media sca Media scan with redundancy	n: check:	Enabled Disabled	
	Pre-Read redundancy check:		Disabled	
Lo	gical Drive name:	Sata_R1_A		
	Logical Drive status:	Optimal		
	Capacity: Logical Drive ID: Subsystem ID (SSID): Associated subsystem:	500.0 GB 60:0a:0b:80 0 Sata_R	0:00:29:ee:02 1	:00:00:c8:53:49:40:03:e2

	RAID level:	1		
	Secure:	No		
	Drive type: Enclosure loss protection:	Serial ATA Yes	(SATA)	
	Preferred owner: Current owner:	Controller Controller	in slot A in slot A	
	Segment size: Capacity reserved for futur Maximum future segment size Modification priority:	e segment s :	ize changes:	128 KB Yes 2,048 KB High
	Read cache: Write cache: Write cache without batt Write cache with mirrori Flush write cache after (in Dynamic cache read prefetch	eries: ng: seconds): :	Enabled Enabled Disabled Enabled 10.00 Enabled	
	Enable background media sca Media scan with redundancy	n: check:	Enabled Disabled	
	Pre-Read redundancy check:		Disabled	
Log	gical Drive name:	Sata_R1_B		
	Logical Drive status:	Optimal		
	Capacity: Logical Drive ID: Subsystem ID (SSID): Associated subsystem: RAID level:	190.0 GB 60:0a:0b:8 8 Sata_R 1	0:00:33:98:48 1	:00:00:c9:93:49:42:37:1f
	Secure:	No		
	Drive type: Enclosure loss protection:	Serial ATA Yes	(SATA)	
	Preferred owner: Current owner:	Controller Controller	in slot B in slot B	
	Segment size: Capacity reserved for future Maximum future segment size Modification priority:	e segment s :	ize changes:	128 KB Yes 2,048 KB High
	Read cache: Write cache: Write cache without batt Write cache with mirrori Flush write cache after (in Dynamic cache read prefetch	eries: ng: seconds): :	Enabled Enabled Disabled Enabled 10.00 Enabled	
	Enable background media sca Media scan with redundancy	n: check:	Enabled Disabled	
	Pre-Read redundancy check:		Disabled	

FLASHCOPY REPOSITORY LOGICAL DRIVES------Number of flashcopy repositories: 0 FLASHCOPY LOGICAL DRIVES------Number of flashcopy logical drives: 0 MISSING LOGICAL DRIVES------Number of missing logical drives: 0

# Chapter 13. Example Script Files

This appendix provides example scripts for configuring a storage subsystem. These examples show how the script commands appear in a complete script file. Also, you can copy these scripts and modify them to create a configuration unique to your storage subsystem.

You can create a script file in two ways:

- Using the save storageSubsystem configuration command
- Writing a script

By using the save storageSubsystem configuration command, you can create a file that you can use to copy an existing configuration from one storage subsystem to other storage subsystems. You can also use this file to restore an existing configuration that has become corrupted. You also can copy an existing file to serve as a pattern from which you create a new script file by modifying portions of the original file. The default file extension is .scr.

You can create a new script file by using a text editor, such as Microsoft Notepad. The maximum line length is 256 characters. The command syntax must conform to the usage guidelines that are described in Chapter 2, "About the Script Commands," on page 15 and the command formatting rules that are listed in "Formatting rules for script commands" on page 29. When you create a new script file, you can use any file name and extension that will run on the host operating system.

This example shows how to run a script file from the command line. c:\...\client>smcli 123.45.67.89 -f scriptfile.scr;

## **Configuration Script Example 1**

This example creates a new logical drive by using the **create logicalDrive** command in the free space of an subsystem.

Show "Create RAID 5 Logical Drive 7 on existing subsystem 1";

//Create logical drive on subsystem created by the create logical drives command

//Note: For subsystems that use all available capacity, the last logical drive on the group is created using all remaining capacity by omitting the capacity=logical drive creation parameter

create logicalDrive subsystem=1 RAIDLevel=5 userLabel="7" owner=A segmentSize=16 cacheReadPrefetch=TRUE capacity=2GB securityType=true;

show "Setting additional attributes for logical drive 7"; //Configuration settings that cannot be set during logical drivecreation

set logicalDrive["7"] cacheFlushModifier=10; set logicalDrive ["7"] cacheWithoutBatteryEnabled=false;

set logicalDrive ["7"] mirrorEnabled=true;

set logicalDrive ["7"] readCacheEnabled=true;

set logicalDrive ["7"] writeCacheEnabled=true; set logicalDrive ["7"] mediaScanEnabled=false;

set logicalDrive ["7"] redundancyCheckEnabled=false;

set logicalDrive ["7"] modificationPriority=high;

This example shows blank lines between the lines beginning with **Show**, **Create**, **//Note**, and **create**. The blank lines are included in this example only for clarity. Each command is actually written on one line in the script file; however, the size of this page has caused the command text to wrap. You might want to include blank lines in your script files to separate blocks of commands or make a comment that stands out. To include a comment, enter two forward slashes (//), which causes the script engine to treat the line as a comment.

The first line of text is the **show string** command. This command shows text that is bounded by double quotation marks ("") on a display monitor when the script file runs. In this example, the text **Create RAID 5 Logical Drive 7 on existing subsystem 1** serves as a title that describes the expected results of running this script file.

The line beginning with **//Create** is a comment that explains that the purpose of this script file is to create a new logical drive by using the **create logicalDrive** command on an existing subsystem.

The line beginning **//Note:** is a comment in the script file that explains that the size of the last logical drive created that uses all of the available capacity because the **capacity** parameter is not used.

The command in this example creates a new logical drive in subsystem 1. The logical drive has a RAID level of 5. The logical drive name (user label) is 7. (Note the double quotation marks around the 7. The double quotation marks define that the information in the double quotation marks is a label.) The new logical drive is assigned to the controller in slot A in the controller module. The segment size is set to 16. The logical drive has a read ahead multiplier value of 256. The capacity of the logical drive is 2 GB.

The command takes this form:

```
create logicalDrive subsystem=subsystemNumber userLabel=logicalDriveName
[freeCapacityArea=freeCapacityIndexNumber] [capacity=logicalDriveCapacity
| owner=(a | b) | cacheReadPrefetch=(TRUE | FALSE) | segmentSize=segmentSizeValue]
[enclosureLossProtect=(TRUE | FALSE)] | [dssPreAllocate = (true | false) ]
| [ securityType= (true | capable | none) ]
```

The general form of the command shows the optional parameters in a different sequence than the optional parameters in the example command. You can enter optional parameters in any sequence. You must enter the required parameters in the sequence shown in the command descriptions.

The line showing "Setting additional attributes for logical drive 7" is another example of using the **show** "*string*" command. The reason for placing this command here is to tell the user that the **create logicalDrive** command ran successfully and that properties that could not be set by the **create logicalDrive** command are now set.

The **set logicalDrive** parameters are shown on separate lines. You do not need to use separate lines for each parameter. You can enter more than one parameter with the **set logicalDrive** command by leaving a space between the parameters, as in this example:

set logicalDrive["7"] cacheFlushModifier=10 cacheWithoutBatteryEnabled=false
modificationPriority=high;

By using separate lines, you can see more clearly the parameters that you are setting and the values to which you are setting the parameters. Blocking the

parameters in this manner makes it easier to either edit the file or copy specific parameter settings for use in another script file.

## **Configuration Script Example 2**

This example creates a new logical drive by using the **create logicalDrive** command with user-defined disk drives in the storage subsystem.

Show "Create RAID3 Logical Drive 2 on existing subsystem 2";

//This command creates the subsystem and the initial logical drive on that subsystem.

//Note: For subsystems that use all available capacity, the last logical drive on the subsystem is created using all remaining capacity by omitting the capacity=logical drive creation parameter create logicalDrive RAIDLevel=3 userLabel="2" drives=[0,1 0,6 1,7 1,3 2,3 2,6] owner=B segmentSize=16 capacity=2GB; show "Setting additional attributes for logical drive 7"'//Configuration settings that cannot be set during logical drive creation set logicalDrive ["7"] cacheFlushModifier=10; set logicalDrive ["7"] cacheWithoutBatteryEnabled=false; set logicalDrive ["7"] mirrorEnabled=true; set logicalDrive ["7"] readCacheEnabled=true; set logicalDrive ["7"] writeCacheEnabled=true; set logicalDrive ["7"] mediaScanEnabled=false; set logicalDrive ["7"] redundantCheckEnabled=false;

set logicalDrive ["7"] modificationPriority=high;

The command in this example, like the **create logicalDrive** command in the previous example, creates a new logical drive. The significant difference between these two examples is that this example shows how you can define specific disk drives to include in the logical drive. Use the **show storageSubsystem profile** command to find out what disk drives are available in a storage subsystem.

## The **create logicalDrive** command takes this form:

```
create logicalDrive raidLevel=(0 | 1 | 3 | 5 | 6) userLabel=logicalDriveName
drives=(enclosureID1,slotID1...enclosureIDn,slotIDn)
[capacity=logicalDriveCapacity | owner=(a | b) |
cacheReadPrefetch=(TRUE | FALSE) | segmentSize=segmentSizeValue]
[enclosureLossProtect=(TRUE | FALSE)] | [dssPreAllocate = (true | false) ]
[ [ securityType= (true | capable | none) ]
```

# **Chapter 14. Deprecated Commands and Parameters**

This appendix lists the commands, command formats, and parameters that are no longer supported by this level of software. The information is presented in two tables. "Deprecated Commands" lists commands no longer supported in this level of software and the new commands that replaced them. "Deprecated Parameters" lists the parameters no longer supported in this level of software and the new parameters that replaced them.

## **Deprecated Commands**

Deprecated Command	New Command
start rollback (volume [flashcopyVolumeName   volumes [flashcopyVolumeName1 flashcopyVolumeNamen])	<pre>start (volume [flashcopyVolumeName   volumes [flashcopyVolumeName1 flashcopyVolumeNamen]) rollback The order of the terms in the syntax is changed to be consistent with other commands.</pre>

## **Deprecated Parameters**

Table 29. Deprecated Parameters

Old Syntax	New Syntax
<pre>copyType=(online   offline)</pre>	This parameter is removed from these commands: • recopy volumeCopy • remove volumeCopy
	<ul> <li>set volumeCopy</li> </ul>

# Appendix A. Examples of information returned by the show commands

This appendix provides examples of information that is returned by the show commands. These examples show the type of information and the information detail. This information is useful in determining the components, features, and identifiers that you might need when you configure or maintain a storage subsystem.

## Show storage subsystem

The show storageArray command returns information about the components and the features in a storage subsystem. If you run the command with the profile parameter, the command returns information in the form shown by this example. This information is the most detailed report that you can receive about the storage subsystem. After you have configured a storage subsystem, save the configuration description to a file as a reference. The show storageArray command returns information about the components and the features in a storage subsystem. If you run the command with the profile parameter, the command returns information in the form shown by this example. This information is the most detailed report that you can receive about the storage subsystem. After you have configured a storage subsystem, save the configuration description to a file as a reference. Storage array profile PROFILE FOR STORAGE ARRAY: example9 (x/xx/xx 1:06:05 PM) SUMMARY-----Number of controllers: 1 Number of volume groups: 1 Total number of volumes (includes an access volume): 5 of 2048 used Number of standard volumes: 2 Number of access volumes: 1 Number of mirror repositories: 2 Number of snapshot repositories: 0 Number of snapshot volumes: 0 Number of mirrored pairs: 0 of 64 used Number of copies: 0 Number of drives: 9 Supported drive types: Fibre (9) Total hot spare drives: 0 Standby: 0 In use: 0 Access volume: LUN 7 (see Mappings section for details) Default host type: Linux (Host type index 6) Current configuration Firmware version: PkgInfo 96.10.21.00 NVSRAM version: N4884-610800-001 Pending configuration Staged firmware download supported?: No Firmware version: Not applicable NVSRAM version: Not applicable Transferred on: Not applicable NVSRAM configured for batteries?: Yes Start cache flushing at (in percentage): 80 Stop cache flushing at (in percentage): 80 Cache block size (in KB): 4 Media scan frequency (in days): Disabled Failover alert delay (in minutes): 5 Feature enable identifier: 1234567891011121314151617181ABCD

```
CONTROLLERS-----
                            Number of controllers: 1
   Controller in Slot A
      Status: Online
      Current configuration
         Firmware version: 96.10.21.00
            Appware version: 96.10.21.00
            Bootware version: 96.10.21.00
         NVSRAM version: N4884-610800-001
      Pending configuration
         Firmware version: Not applicable
            Appware version: Not applicable
            Bootware version: Not applicable
         NVSRAM version: Not applicable
Transferred on: Not applicable
      Board ID: 4884
      Product ID: INF-01-00
      Product revision: 9610
      Serial number: 1T14148766
      Date of manufacture: October 14, 2001
      Cache/processor size (MB): 1024/128
      Date/Time: Fri Feb 13 13:06:05 MST 2004
      Associated Volumes (* = Preferred Owner):
         1*, CTL 0 Mirror Repository*, Mirror Repository 1*, Primary_Volume*
      Ethernet port: 1
         MAC address: 00:a0:b8:0c:c3:f5
         Host name: ausctlr9
         Network configuration: Static
         IP address: 172.22.4.249
         Subnet mask: 255.255.255.0
         Gateway: 172.22.4.1
         Remote login: Enabled
      Drive interface: Fibre
         Channel: 1
         Current ID: 125/0×1
         Maximum data rate: 2 Gbps
         Current data rate: 1 Gbps
         Data rate control: Switch
         Link status: Up
      Drive interface: Fibre
         Channel: 2
         Current ID: 125/0×1
         Maximum data rate: 2 Gbps
         Current data rate: 1 Gbps
         Data rate control: Switch
         Link status: Up
      Drive interface: Fibre
         Channel: 3
         Current ID: 125/0×1
         Maximum data rate: 2 Gbps
         Current data rate: 1 Gbps
         Data rate control: Switch
         Link status: Up
      Drive interface: Fibre
         Channel: 4
         Current ID: 125/0×1
         Maximum data rate: 2 Gbps
         Current data rate: 1 Gbps
         Data rate control: Switch
         Link status: Up
```

```
Host interface: Fibre
          Port: 1
          Current ID: Not applicable/0xFFFFFFFF
          Preferred ID: 126/0×0
          NL-Port ID: 0x011100
          Maximum data rate: 2 Gbps
          Current data rate: 1 Gbps
          Data rate control: Switch
          Link status: Up
          Topology: Fabric Attach
          World-wide port name: 20:2c:00:a0:b8:0c:c3:f6
          World-wide node name: 20:2c:00:a0:b8:0c:c3:f5
          Part type: HPFC-5200
                                     revision 10
      Host interface: Fibre
          Port: 2
          Current ID: Not applicable/0xFFFFFFF
          Preferred ID: 126/0x0
          NL-Port ID: 0x011000
          Maximum data rate: 2 Gbps
          Current data rate: 1 Gbps
          Data rate control: Switch
          Link status: Up
          Topology: Fabric Attach
          World-wide port name: 20:2c:00:a0:b8:0c:c3:f7
          World-wide node name: 20:2c:00:a0:b8:0c:c3:f5
          Part type: HPFC-5200
                                     revision 10
VOLUME GROUPS-----
   Number of volume groups: 1
   Volume Group 1 (RAID 5)
       Status: Online
       Drive type: Fibre Channel
      Tray loss protection: No
      Current owner: Controller in slot A
      Associated volumes and free capacities:
      CTL 0 Mirror Repository (0.125 GB), Mirror Repository 1 (0.125 GB),
Primary_Volume (5 GB), 1 (1 GB), Free Capacity (129.113 GB)
Associated drives (in piece order):
          Drive at Tray 0, Slot 1
Drive at Tray 0, Slot 2
          Drive at Tray 0, Slot 3
          Drive at Tray 0, Slot 5
          Drive at Tray 0, Slot 6
          Drive at Tray 0, Slot 7
Drive at Tray 0, Slot 10
Drive at Tray 0, Slot 12
          Drive at Tray 0, Slot 14
```

STANDARD VOLUMES-----SUMMARY Number of standard volumes: 2 See other Volumes sub-tabs for premium feature information. NAME STATUS CAPACITY RAID LEVEL VOLUME GROUP Optimal 1 GB 1 1 5 Primary\_Volume Optimal 5 GB 5 1 DETAILS Volume name: 1 Volume ID: 60:0a:0b:80:00:0c:c3:f5:00:00:00:23:40:22:86:ea Subsystem ID (SSID): 3 Status: Optimal Drive type: Fibre Channel Tray loss protection: No Preferred owner: Controller in slot A Current owner: Controller in slot A Capacity: 1 GB RAID level: 5 Segment size: 64 KB Modification priority: High Associated volume group: 1 Read cache: Enabled Write cache: Enabled Write cache without batteries: Disabled Write cache with mirroring: Enabled Flush write cache after (in seconds): 10.00 Cache read ahead multiplier: 1 Enable background media scan: Disabled Media scan with redundancy check: Disabled Volume name: Primary\_Volume Volume ID: 60:0a:0b:80:00:0c:c3:f5:00:00:00:1c:40:22:45:ca Subsystem ID (SSID): 2 Status: Optimal Drive type: Fibre Channel Tray loss protection: No Preferred owner: Controller in slot A Current owner: Controller in slot A Capacity: 5 GB RAID level: 5 Segment size: 64 KB Modification priority: High Associated volume group: 1 Read cache: Enabled Write cache: Enabled Write cache without batteries: Disabled Write cache with mirroring: Enabled Flush write cache after (in seconds): 10.00 Cache read ahead multiplier: 1 Enable background media scan: Disabled Media scan with redundancy check: Disabled

```
SNAPSHOT REPOSITORY VOLUMES -----
                                    ------
  Number of snapshot repositories: 0
MIRROR REPOSITORY VOLUMES-----
  Number of mirror repositories: 0
SNAPSHOT VOLUMES-----
  Number of snapshot volumes: 0
MIRRORED PAIRS-----
SUMMARY
  Number of mirrored pairs: 0 of 64 used
   See also Standard sub-tab for standard volume properties
COPIES
SUMMARY
   Number of copies: 0
  See also Standard sub-tab for standard volume properties
MISSING VOLUMES-----
   Number of missing volumes: 0
  See other Volumes sub-tabs for premium feature information
DRIVES
SUMMARY
  Number of drives: 9
     Supported drive types: Fibre (9)
  BASIC:
                      CAPACITY
                                CURRENT DATA RATE PRODUCT ID
                                                                  FIRMWARE VERSION
  TRAY, SLOT STATUS
                     16.959 GB 1 Gbps
   0, 1
             Optimal
                                                  ST318451FC
                                                                  F28D
                      16.959 GB 1 Gbps
16.959 GB 1 Gbps
   0, 2
             Optimal
                                                 ST318451FC
                                                                  F28D
   0, 3
             Optimal
                                                 ST318451FC
                                                                  F28D
   0, 5
             Optimal
                      16.959 GB 1 Gbps
                                                 ST318451FC
                                                                  F28D
   0, 6
                      16.959 GB
             Optimal
                                1 Gbps
                                                 ST318451FC
                                                                  F28D
   0, 7
                      16.959 GB 1 Gbps
             Optimal
                                                 ST318451FC
                                                                  F28D
                     16.959 GB 1 Gbps
16.959 GB 1 Gbps
             Optimal
                                                 ST318451FC
   0, 10
                                                                  F28D
             Optimal
                                                 ST318451FC
                                                                  F28B
   ο,
     12
             Optimal 16.959 GB 1 Gbps
                                                 ST318451FC
  0, 14
                                                                  F28B
  DRIVE CHANNELS:
  TRAY, SLOT PREFERRED CHANNEL REDUNDANT CHANNEL
  0, 1
0, 2
0, 3
              4
                               2
             2
                               4
             4
                               2
  0, 5
0, 6
             4
                               2
             2
                               4
   ο,
     7
              4
                               2
   0, 10
             2
                               4
```

4

4

ο,

12 Ο,

14

2

z

HOT SPARE COVERAGE: The following volume groups are not protected: 1 Total hot spare drives: 0 Standby: 0 In use: 0 DETAILS Drive at Tray 0, Slot 1 Drive port: 1, Channel: 4, ID: 0/0xEF Drive port: 2, Channel: 2, ID: 0/0×EF Drive path redundancy: OK Status: Optimal Raw capacity: 16.959 GB Usable capacity: 16.92 GB Current data rate: 1 Gbps Product ID: ST318451FC Firmware version: F28D Serial number: 3CC052RK00007116BMW7 Vendor: SEAGATE Date of manufacture: July 19, 2001 World-wide name: 20:00:00:20:37:f5:b3:b6 Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 1 Drive at Tray 0, Slot 2 Drive port: 1, Channel: 2, ID: 1/0×E8 Drive port: 2, Channel: 4, ID: 1/0×E8 Drive path redundancy: OK Status: Optimal Raw capacity: 16.959 GB Usable capacity: 16.92 GB Current data rate: 1 Gbps Product ID: ST318451FC Firmware version: F28D Serial number: 3CC0Q06K0000720291J9 Vendor: SEAGATE Date of manufacture: July 18, 2001 World-wide name: 20:00:00:20:37:f5:ac:cd Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 1 Drive at Tray 0, Slot 3 Drive port: 1, Channel: 4, ID: 2/0xE4 Drive port: 2, Channel: 2, ID: 2/0xE4 Drive path redundancy: OK Status: Optimal Raw capacity: 16.959 GB Usable capacity: 16.92 GB Current data rate: 1 Gbps Product ID: ST318451FC Firmware version: F28D Serial number: 3CC052MA00007116DB89 Vendor: SEAGATE Date of manufacture: July 19, 2001 World-wide name: 20:00:00:20:37:f5:b3:fc Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 1

Drive at Tray 0, Slot 5 Drive port: 1, Channel: 4, ID: 4/0×E1 Drive port: 2, Channel: 2, ID: 4/0×E1 Drive path redundancy: OK Status: Optimal Raw capacity: 16.959 GB Usable capacity: 16.92 GB Current data rate: 1 Gbps Product ID: ST318451FC Firmware version: F28D Serial number: 3CC0526N00007143A10K Vendor: SEAGATE Date of manufacture: July 19, 2001 World-wide name: 20:00:00:20:37:f5:b4:5b Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 1 Drive at Tray 0, Slot 6 Drive port: 1, Channel: 2, ID: 5/0×E0 Drive port: 2, Channel: 4, ID: 5/0×E0 Drive path redundancy: OK Status: Optimal Raw capacity: 16.959 GB Usable capacity: 16.92 GB Current data rate: 1 Gbps Product ID: ST318451FC Firmware version: F28D Serial number: 3CC03R58000071123U2D Vendor: SEAGATE Date of manufacture: July 19, 2001 World-wide name: 20:00:00:20:37:f5:b3:06 Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 1 Drive at Tray 0, Slot 7 Drive port: 1, Channel: 4, ID: 6/0xDC Drive port: 2, Channel: 2, ID: 6/0xDC Drive path redundancy: OK Status: Optimal Raw capacity: 16.959 GB Usable capacity: 16.92 GB Current data rate: 1 Gbps Product ID: ST318451FC Firmware version: F28D Serial number: 3CC052J400007116BSJV Vendor: SEAGATE Date of manufacture: July 19, 2001 World-wide name: 20:00:00:20:37:f5:b3:0b Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 1

Drive at Tray 0, Slot 10 Drive port: 1, Channel: 2, ID: 96/0x3A Drive port: 2, Channel: 4, ID: 96/0x3A Drive path redundancy: OK Status: Optimal Raw capacity: 16.959 GB Usable capacity: 16.92 GB Current data rate: 1 Gbps Product ID: ST318451FC Firmware version: F28D Serial number: 3CCOP5Q80000711552CZ Vendor: SEAGATE Date of manufacture: June 30, 2001 World-wide name: 20:00:00:20:37:f5:69:50 Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 1 Drive at Tray 0, Slot 12 Drive port: 1, Channel: 2, ID: 72/0x67 Drive port: 2, Channel: 4, ID: 72/0x67 Drive path redundancy: OK Status: Optimal Raw capacity: 16.959 GB Usable capacity: 16.92 GB Current data rate: 1 Gbps Product ID: ST318451FC Firmware version: F28B Serial number: 3CCOP7TS00007116DL9D vendor: LSILOGIC Date of manufacture: Not available World-wide name: 20:00:00:20:37:f5:68:76 Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 1 Drive at Tray 0, Slot 14 Drive port: 1, Channel: 2, ID: 104/0x2E Drive port: 2, Channel: 4, ID: 104/0x2E Drive path redundancy: OK Status: Optimal Raw capacity: 16.959 GB Usable capacity: 16.92 GB Current data rate: 1 Gbps Product ID: ST318451FC Firmware version: F28B Serial number: 3CC03STH00007112JNM4 Vendor: LSILOGIC Date of manufacture: Not available World-wide name: 20:00:00:20:37:f5:b3:69 Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 1

```
DRIVE CHANNEL INFORMATION - Channel 1
   SUMMARY-----
      Channel 1 status: Optimal
      Controller A link status: Up
   CUMULATIVE ERROR COUNTS-----
      Controller A Cumulative Error Counts
         Baseline time set: 2/9/04 10:46:37 AM
         Sample period (days, hh:mm:ss): 4 days, 02:06:52
Controller detected errors: 0
         Drive detected errors: 0
         Timeout errors: 0
         Link down errors: 3
         Total I/O count: 11
DRIVE CHANNEL INFORMATION - Channel 2
   SUMMARY-----
      Channel 2 status: Optimal
      Controller A link status: Up
   CUMULATIVE ERROR COUNTS-----
      Controller A Cumulative Error Counts
Baseline time set: 2/9/04 10:46:42 AM
         Sample period (days, hh:mm:ss): 4 days, 02:06:47
Controller detected errors: 0
         Drive detected errors: 0
         Timeout errors: 0
         Link down errors: 2
         Total I/O count: 239075
DRIVE CHANNEL INFORMATION - Channel 3
   SUMMARY-----
      Channel 3 status: Optimal
      Controller A link status: Up
   CUMULATIVE ERROR COUNTS-----
      Controller A Cumulative Error Counts
         Baseline time set: 2/9/04 10:46:42 AM
Sample period (days, hh:mm:ss): 4 days, 02:06:47
Controller detected errors: 0
         Drive detected errors: 0
         Timeout errors: 0
         Link down errors: 3
         Total I/O count: 11
DRIVE CHANNEL INFORMATION - Channel 4
   SUMMARY-----
      Channel 4 status: Optimal
      Controller A link status: Up
   CUMULATIVE ERROR COUNTS-----
      Controller A Cumulative Error Counts
         Baseline time set: 2/9/04 10:46:42 AM
Sample period (days, hh:mm:ss): 4 days, 02:06:47
Controller detected errors: 0
         Drive detected errors: 0
         Timeout errors: 0
         Link down errors: 2
         Total I/O count: 45807
```

```
TRAYS-----
   Controller Tray Overall Component Information
      Drive technology: Unknown
      Minihub data rate mismatch: No
      Fan canister: Optimal
      Fan canister: Optimal
      Battery status: Optimal
Age: 0 day(s)
Days until replacement: 719 day(s)
      Power supply canister
         Status: Optimal
      Power supply canister
         Status: Optimal
      Temperature: Optimal
      Host mini-hub canister
          Status: Optimal
         Location: Controller A - Port 1
         Serial number: SN 1T13742151
         Part number: PN 348-0043290
         Vendor: VN LSILOGIC
Date of manufacture: Not available
         SFP
             Status: Optimal
             Location: In connection
Supported data rate(s): 1 Gbps, 2 Gbps
             Link length: Intermediate
             Connector: LC
             Transmitter type: Shortwave Laser w/o OFC
             Transmission media: TM Multi-mode 50m(M5) TM Multi-mode 62.5m(M6)
             IEEE company ID: 08 00 5a
             Revision: AA10
             Part number: IBM42P21SNY
            Serial number: 21P7053000M6R
Vendor: IBM
             Date of manufacture: July 1, 2001
         SFP
             Status: Optimal
             Location: Out connection
             Supported data rate(s): 1 Gbps, 2 Gbps
             Link length: Intermediate
             Connector: LC
             Transmitter type: Shortwave Laser w/o OFC
             Transmission media: TM Multi-mode 50m(M5) TM Multi-mode 62.5m(M6)
             IEEE company ID: 08 00 5a
             Revision: AA10
             Part number: IBM42P21SNY
             Serial number: 21P7053000M54
             Vendor: IBM
             Bate of manufacture: July 1, 2001
```

```
Host mini-hub canister
   Status: Optimal
   Location: Controller A - Port 2
   Serial number: SN 1T13742219
   Part number: PN 348-0043290
   Vendor: VN LSILOGIC
Date of manufacture: Not available
   SFP
      Status: Optimal
      Location: In connection
      Supported data rate(s): 1 Gbps, 2 Gbps
      Link length: Intermediate
      Connector: LC
      Transmitter type: Shortwave Laser w/o OFC
      Transmission media: TM Multi-mode 50m(M5) TM Multi-mode 62.5m(M6)
      IEEE company ID: 08 00 5a
      Revision: AA10
      Part number: IBM42P215NY
      Serial number: 21P7053000M3Z
Vendor: IBM
      Date of manufacture: July 1, 2001
   SFP
      Status: Optimal
      Location: Out connection
      Supported data rate(s): 1 Gbps, 2 Gbps
      Link length: Intermediate
      Connector: LC
      Transmitter type: Shortwave Laser w/o OFC
      Transmission media: TM Multi-mode 50m(M5) TM Multi-mode 62.5m(M6)
      IEEE company ID: 08 00 5a
      Revision: AA10
      Part number: IBM42P21SNY
      Serial number: 21P7053000M3J
      Vendor: IBM
      Date of manufacture: July 1, 2001
Drive mini-hub canister
   Status: Optimal
   Location: Channel 1
   Serial number: SN 1T13438343
   Part number: PN 348-0043290
   Vendor: VN LSILOGIC
Date of manufacture: Not available
   SFP
      Status: Optimal
      Location: Out connection
      Supported data rate(s): 1 Gbps, 2 Gbps
      Link length: Intermediate
Connector: LC
      Transmitter type: Shortwave Laser w/o OFC
      Transmission media: TM Multi-mode 50m(M5) TM Multi-mode 62.5m(M6)
      IEEE company ID: 08 00 5a
      Revision: AA10
      Part number: IBM42P21SNY
      Serial number: 21P7053000M6Z
      Vendor: IBM
      Date of manufacture: July 1, 2001
```

```
Drive mini-hub canister
   Status: Optimal
   Location: Channel 2
   Serial number: SN 1T13844441
Part number: PN 19K1270
   Vendor: VN IBM
Date of manufacture: Not available
   SEP
      Status: Optimal
      Location: In connection
      Supported data rate(s): 1 Gbps, 2 Gbps
      Link length: Intermediate
      Connector: LC
      Transmitter type: Shortwave Laser w/o OFC
      Transmission media: TM Multi-mode 50m(M5) TM Multi-mode 62.5m(M6)
      IEEE company ID: 08 00 5a
      Revision: AA10
      Part number: IBM42P21SNY
      Serial number: 21P7053000M4E
      Vendor: IBM
Date of manufacture: July 1, 2001
   SEP
      Status: Optimal
      Location: Out connection
      Supported data rate(s): 1 Gbps, 2 Gbps
      Link length: Intermediate
      Connector: LC
      Transmitter type: Shortwave Laser w/o OFC
      Transmission media: TM Multi-mode 50m(M5) TM Multi-mode 62.5m(M6)
      IEEE company ID: 08 00 5a
      Revision: AA10
      Part number: IBM42P21SNY
      Serial number: 21P7053000M40
      Vendor: IBM
      Date of manufacture: July 1, 2001
Drive mini-hub canister
   Status: Optimal
   Location: Channel 3
   Serial number: SN 1T13844373
   Part number: PN 348-0043290
   Vendor: VN LSILOGIC
Date of manufacture: Not available
   SFP
      Status: Optimal
      Location: In connection
      Supported data rate(s): 1 Gbps, 2 Gbps
      Link length: Intermediate
      Connector: LC
      Transmitter type: Shortwave Laser w/o OFC
      Transmission media: TM Multi-mode 50m(M5) TM Multi-mode 62.5m(M6)
      IEEE company ID: 08 00 5a
      Revision: AA10
      Part number: IBM42P215NY
      Serial number: 21P7053000M4R
      Vendor: IBM
      Date of manufacture: July 1, 2001
   SEP
      Status: Optimal
      Location: Out connection
Supported data rate(s): 1 Gbps, 2 Gbps
      Link length: Intermediate
      Connector: LC
      Transmitter type: Shortwave Laser w/o OFC
      Transmission media: TM Multi-mode 50m(M5) TM Multi-mode 62.5m(M6)
      IEEE company ID: 08 00 5a
      Revision: AA10
      Part number: IBM42P21SNY
      Serial number: 21P7053000LXS
      Vendor: IBM
      Date of manufacture: July 1, 2001
```

```
Drive mini-hub canister
   Status: Optimal
   Location: Channel 4
   Serial number: SN 1T13844386
   Part number: PN 348-0043290
   Vendor: VN LSILOGIC
   Date of manufacture: Not available
   SFP
      Status: Optimal
      Location: In connection
Supported data rate(5): 1 Gbps, 2 Gbps
      Link length: Intermediate
      Connector: LC
      Transmitter type: Shortwave Laser w/o OFC
      Transmission media: TM Multi-mode 50m(M5) TM Multi-mode 62.5m(M6)
      IEEE company ID: 08 00 5a
      Revision: AA10
      Part number: IBM42P21SNY
      Serial number: 21P7053000M6E
      Vendor: IBM
      Date of manufacture: July 1, 2001
   SEP
      Status: Optimal
      Location: Out connection
      Supported data rate(s): 1 Gbps, 2 Gbps
      Link length: Intermediate
      Connector: LC
      Transmitter type: Shortwave Laser w/o OFC
      Transmission media: TM Multi-mode 50m(M5) TM Multi-mode 62.5m(M6)
      IEEE company ID: 08 00 5a
      Revision: AA10
      Part number: IBM42P21SNY
      Serial number: 21P7053000M57
      Vendor: IBM
      Date of manufacture: July 1, 2001
```
Drive Tray 0 Overall Component Information Drive technology: Fibre channel Part number: PN 348-0042807 Serial number: SN 0146001584 Vendor: VN LSILOGIC Date of manufacture: November 1, 2001 Tray path redundancy: OK Fan canister: Optimal Fan canister: Optimal Power supply canister Status: Optimal Power supply canister Status: Optimal Temperature: Optimal Temperature: Optimal ESM card Status: Optimal Firmware version: 9163 Maximum data rate: 1 Gbps Current data rate: 1 Gbps Location: A (left canister) Card communication: OK Product ID: SYM2500-FC Part number: PN 348-0043832 Serial number: SN 1T22941398 Vendor: LSILOGIC Date of manufacture: July 1, 2002 GBIC Status: Optimal Location: In connection Type: Optical Shortwave ESM card Status: Optimal Firmware version: 9163 Maximum data rate: 1 Gbps Current data rate: 1 Gbps Location: B (right canister) Card communication: OK Product ID: SYM2500-FC Part number: PN 348-0043832 Serial number: SN 1T13945863 Vendor: LSILOGIC Date of manufacture: October 1, 2001 GBIC Status: Optimal Location: In connection Type: Optical Shortwave

```
MAPPINGS (SANshare STORAGE PARTITIONING - ENABLED (0 of 64 used))------
   Volume-To-LUN Mappings
     Access Volume, LUN 7, Default Group
   Topology
      [Undefined Host Ports]
         10:00:00:a0:b8:04:29:3f
         21:00:00:e0:8b:01:34:c1
         21:01:00:e0:8b:22:80:38
      [Default Group]
   NVSRAM Host Type Internal Definitions
      Index 0
         Name: Windows NT Non-Clustered (SP5 or higher)
         AVT status: Disabled
      Index 1
         Name: Windows 2000/Server 2003 Non-Clustered
         AVT status: Disabled
      Index 2
         Name: Solaris
         AVT status: Disabled
      Index 3
         Name: HP-UX
         AVT status: Enabled
      Index 4
         Name: AIX
         AVT status: Disabled
      Index 5
         Name: Irix
         AVT status: Disabled
      Index 6 (DEFAULT)
         Name: Linux
         AVT status: Disabled
      Index 7
         Name: Windows NT Clustered (SP5 or higher)
         AVT status: Disabled
      Index 8
         Name: Windows 2000/Server 2003 Clustered
         AVT status: Disabled
      Index 9
         Name: Netware Non-Failover
         AVT status: Enabled
      Index 10
         Name: PTX
         AVT status: Enabled
      Index 11
         Name: Netware Failover
         AVT status: Enabled
      Index 12
         Name: Solaris (with Veritas DMP)
         AVT status: Enabled
```

#### Show controller NVSRAM

The show controller NVSRAM command returns a table of values in the controller NVSRAM that is similar to that shown in this example. With the information from the table, you can modify the contents of the NVSRAM by using the set controller command. This example shows information for a controller in slot A in a controller enclosure. You can produce a similar table for a controller in slot B, or you can produce a table for both controllers.

#### Show volume

The show volume command returns information about the volumes in a storage subsystem.

```
STANDARD VOLUMES-----
SUMMARY
Number of standard volumes: 5
```

See other Volumes sub-tabs for premium feature information. NAME STATUS CAPACITY RAID LEVEL VOLUME GROUP LUN Optimal 5,120.000 1 GB 10 6 13 2 Optimal 1.000 GB 1 Volume-Group-2 14 3 Optimal 10.000 GB 1 Volume-Group-2 15 4 Optimal 3.000 GB 1 Volume-Group-2 16 Unnamed Optimal 100.004 MB 0 Volume-Group-1 0 DETAILS Volume name: 1 Volume status: Optimal Capacity: 5,120.000 GB Volume world-wide identifier: 60:0a:0b:80:00:29:ed:12:00:00 Subsystem ID (SSID): 14 Associated Array: 6 RAID level: 10 LUN: 13 Accessible By: Default Group Media type: Hard Disk Drive Interface type: Serial ATA (SATA) Enclosure loss protection: Yes Secure: No Preferred owner: Controller in slot A Current owner: Controller in slot A Segment size: 128 KB Capacity reserved for future Yes segment size changes: Maximum future segment size: 2,048 KB Modification priority: High Read cache: Enabled Write cache: Enabled Write cache without batteries: Disabled Write cache with mirroring: Enabled Flush write cache after (in seconds): 10.00 Dynamic cache read prefetch: Enabled Enable background media scan: Disabled Media scan with redundancy check: Disabled Pre-Read redundancy check: Disabled Volume name: 2 Volume status: Optimal Capacity: 1.000 GB Volume world-wide identifier: 60:0a:0b:80:00:29:ed:12:00:00 Subsystem ID (SSID): 15 Associated Array: Volume-Group-2 RAID level: 1 LUN: 14 Accessible By: Default Group Media type: Hard Disk Drive Interface type: Fibre Channel Enclosure loss protection: Yes Secure: No Controller in slot B Preferred owner: Current owner: Controller in slot B Segment size: 128 KB Capacity reserved for future segment size changes: Yes Maximum future segment size: 2,048 KB Modification priority: High Read cache: Enabled Write cache: Enabled Disabled Write cache without batteries:

Write cache with mirroring: Enabled Flush write cache after (in seconds): 10.00 Dynamic cache read prefetch: Enabled Enable background media scan: Disabled Media scan with redundancy check: Disabled Pre-Read redundancy check: Disabled Volume name: 3 Optimal Volume status: 10.00 GB Capacity: Volume world-wide identifier: 60:0a:0b:80:00:29:ed:12:00:00 Subsystem ID (SSID): 16 Associated Array: Volume-Group-2 RAID level: 1 LUN: 15 Accessible By: Default Group Hard Disk Drive Media type: Interface type: Fibre Channel Enclosure loss protection: Yes Secure: No Preferred owner: Controller in slot A Current owner: Controller in slot A Segment size: 128 KB Capacity reserved for future segment size changes: Yes Maximum future segment size: 2,048 KB Modification priority: High Read cache: Enabled Write cache: Enabled Write cache without batteries: Disabled Write cache with mirroring: Enabled Flush write cache after 10.00 (in seconds): Dynamic cache read prefetch: Enabled Enable background media scan: Disabled Media scan with redundancy check:Disabled Pre-Read redundancy check: Disabled Volume name: 4 Volume status: Optimal 3.000 GB Capacity: Volume world-wide identifier: 60:0a:0b:80:00:29:ed:12:00:00 Subsystem ID (SSID): 17 Associated Array: Volume-Group-2 RAID level: 1 IUN: 16 Accessible By: Default Group Media type: Hard Disk Drive Interface type: Fibre Channel Appendix A: Examples of Information Returned by the Show Commands 225 Enclosure loss protection: Yes Secure: No Preferred owner: Controller in slot B Current owner: Controller in slot B Segment size: 128 KB Capacity reserved for future segment size changes: Yes Maximum future segment size: 2,048 KB Modification priority: High Read cache: Enabled Write cache: Enabled Write cache without batteries: Disabled Write cache with mirroring: Enabled Flush write cache after (in seconds): 10.00 Dynamic cache read prefetch: Enabled Enable background media scan: Disabled Media scan with redundancy check: Disabled

Pre-Read redundancy check: Disabled Volume name: Unamed Volume status: Optimal 100.004 GB Capacity: Volume world-wide identifier: 60:0a:0b:80:00:29:ed:12:00:00 Subsystem ID (SSID): 0 Associated Array: Volume-Group-1 RAID level: 0 0 IUN: Accessible By: Default Group Media type: Hard Disk Drive Interface type: Serial ATA (SATA) Enclosure loss protection: No Secure: No Preferred owner: Controller in slot B Current owner: Controller in slot B 16 KB Segment size: Capacity reserved for future segment size changes: Yes Maximum future segmentsize: Not Appl Modification priority: Low Read cache: Enabled Write cache: Disabled Write cache without batteries: Disabled Write cache with mirroring: Disabled Flush write cache after (in seconds): 10.00 Dynamic cache read prefetch: Disabled Enable background media scan: Enabled Media scan with redundancy check: Enabled Enhanced FlashCopy REPOSITORY VOLUMES-----SUMMARY Number of Enhanced FlashCopy repositories: NAME CAPACITY USAGE(%) THRESHOLD WARNING FULL POLICY DAE1-1 0 50% full Fail Enhanced FlashCopy DETAILS Enhanced FlashCopy REPOSITORY VOLUME NAME: DAE1-1 Enhanced FlashCopy repository volume status: Optimal Capacity usage (%): 0 Notify when capacity reaches: 50% full Enhanced FlashCopy repository full policy: Fail Enhanced FlashCopy volume Associated base volume (standard): Unnamed Associated Enhanced FlashCopy volume: DAE1 Volume name: DAE1-1 Volume status: Optimal Capacity: 20.000 MB Volume world-wide identifier: 60:0a:0b:80:00:29:ed:12 Subsystem ID (SSID): 11 RAID level: 0 Hard Disk Drive Media type: Interface type: Serial ATA (SATA) Enclosure loss protection: No Secure: No Controller in slot B Preferred owner: Current owner: Controller in slot B Segment size: 64 KB Capacity reserved for future segment size changes: No Maximum future segment size: Not ap Modification priority: High Read cache: Enabled Write cache: Enabled Disabled Write cache without batteries:

Write cache with mirroring: Enabled Flush write cache after (in seconds): 10.00 Dynamic cache read prefetch: Disabled Disabled Enable background media scan: Media scan with redundancy check: Disabled MIRROR REPOSITORY VOLUMES-----SUMMARY Number of mirror repositories: 2 NAME STATUS CAPACITY RAID LEVEL VOLUME Mirror Repository 2 Optimal 129.093 MB 10 6 Mirror Repository 1 Optimal 129.093 MB 10 6 DETAILS MIRROR REPOSITORY VOLUME NAME: Mirror Repository 2 Mirror repository volume status: Optimal Mirror Repository 2 Volume name: Volume status: Optimal 129.093 MB Capacity: Volume world-wide identifier: 60:0a:0b:80:00:29:ed Subsystem ID (SSID): 12 Associated Array: 6 RAID level: 10 Media type: Hard Disk Drive Interface type: Serial ATA (SATA) Enclosure loss protection: Yes Secure: No Preferred owner: Controller inslot B Current owner: Controller in slot B Segment size: 32 KB Capacity reserved for future segment size changes: No Maximum future segment size: Not applicable Modification priority: High MIRROR REPOSITORY VOLUME NAME: Mirror Repository 1 Mirror repository volume status: Optimal Volume name: Mirror Repository 1 Volume status: Optimal 129.093 MB Capacity: Volume world-wide identifier: 60:0a:0b:80:00:29:ed Subsystem ID (SSID): 13 Associated Array: 6 RAID level: 10 Hard Disk Drive Media type: Interface type: Serial ATA (SATA) Enclosure loss protection: Yes Secure: No Preferred owner: Controller in slot A Current owner: Controller in slot A Segment size: 32 KB Capacity reserved for future segment size changes: No Maximum future segment size: Not applicable Modification priority: High Enhanced FlashCopy VOLUMES-----SUMMARY Number of Enhanced FlashCopy volumes: 1 NAME STATUS CREATION TIMESTAMP DAE1 Optimal 9/24/10 8:54 AM DETAILS Enhanced FlashCopy VOLUME NAME: DAF1 Enhanced FlashCopy status: Optimal 9/24/10 8:54 AM Creation timestamp: Associated base volume (standard): Unnamed Associated Enhanced FlashCopy repository volume: DAE1-1 Volume world-wide identifier: 60:0a:0b:80:00:29:ed:12:00 Capacity: 100.004 MB Preferred owner: Controller in slot B

```
Current owner:
                           Controller in slot B
COPIES-----
SUMMARY
Number of copies: 10
S = Source volume
T = Target volume
COPY PAIR
               STATUS
                       COMPLETION TIMESTAMP
5 (S), 10 (T)
               Completed 10/14/10 3:16:27 PM
5 (S), 8 (T)
                 Completed 10/18/10 9:46:45 AM
                Stopped None
Completed 10/14/10 3:13:37 PM
10 (S), 9 (T)
(S), 7 (T)
5 (S), 4 (T)
                 Completed 10/14/10 3:18:23 PM
                 Completed 10/14/10 3:22:56 PM
1 (S), 3 (T)
Unnamed (S), 5 (T) Completed 9/16/10 2:30:06 PM
Unnamed (S), 11 (T) Stopped None
Unnamed (S), 6 (T) Completed 9/2/10 10:03:56 AM
Unnamed (S), 1 (T) Completed 9/16/10 12:41:14 PM
DETAILS
Copy pair: Unnamed and 4
Copy status:
                    Copy pair: Unnamed and 4
Copy status:
                    Completed
Start timestamp:
                      9/16/10 2:29:23 PM
Completion timestamp:
                         9/16/10 2:30:06 PM
Copy priority:
                     Lowest
Source volume:
                     Unnamed
Volume world-wide
identifier: 60:0a:0b:80:00:29:ed:12
Target volume: 5
Volume world-wide
identifier: 60:0a:0b:80:00:47:5b:8a
Read-only:
                   Disabled
Copy pair:
                   Unnamed and 3
Copy status:
                    Stopped
Start timestamp:
                      None
Completion timestamp:
                         None
Copy priority:
                     Lowest
Source volume:
                     Unnamed
Volume world-wide
identifier: 60:0a:0b:80:00:29:ed:12
Target volume:
                     11
Volume world-wide
identifier: 60:0a:0b:80:00:29:ed:12
Read-only:
                   Enabled
Copy pair:
                   Unnamed and 2
Copy status:
                   Completed
Start timestamp:
                      9/2/10 10:03:41 AM
Completion timestamp:
                        9/2/10 10:03:56 AM
Co py priority:
                      Medium
Source volume:
                     Unnamed
Volume world-wide
identifier: 60:0a:0b:80:00:29:ed:12
Target volume:
                     6
Volume world-wide
identifier: 60:0a:0b:80:00:29:ed:12
Read-only:
                   Enabled
Copy pair:
                   Unnamed and 1
Copy status:
                    Completed
                    9/16/10 12:40:58 PM
Start timestamp:
Completion timestamp:
                         9/16/10 12:41:14 PM
Copy priority:
                     Medium
Source volume:
                     Unnamed
Volume world-wide
identifier: 60:0a:0b:80:00:29:ed:1
Target volume:
                     1
Volume world-wide
identifier: 60:0a:0b:80:00:47:5b:8
Read-only:
                   Enabled
```

MIRRORED PAIRS-----Number of mirrored pairs: 0 of 64 used MISSING VOLUMES-----Number of missing volumes: 0

#### Show drive channel stat

The show drive channel stat command returns information about the drive channels in a storage subsystem. Use this information to determine how well the channels are running and errors that might be occurring on the channels.

```
DRIVE CHANNELS-----
SUMMARY
CHANNEL PORT STATUS
1 8,7,ESM A 1A,ESM A 1B,ESM A 1A,ESM A 1B,ESM A 1B Optimal
2 6,5
                          Optimal
                          Optimal
3 4,3
4 2,1
                          Optimal
5 1,2,ESM B 1B,ESM B 1A,ESM B 1B,ESM B 1A,ESM B 1B Optimal
6 3,4
                         Optimal
7 5,6
                          Optimal
8 7,8
                          Optimal
DETAILS
DRIVE CHANNEL 1
Port: 8, 7, ESM A 1A, ESM A 1B, ESM A 1A, ESM A 1B, ESM A 1B
Status: Optimal
Max. Rate: 4 Gbps
Current Rate: 4 Gbps
Rate Control: Auto
Controller A link status: Up
Controller B link status: Up
Trunking active: No
DRIVE COUNTS
Total # of attached drives: 44
Connected to: Controller A, Port 8
 Attached drives: 44
 Drive enclosure: 3 (14 drives)
 Drive enclosure: 1 (15 drives)
 Drive enclosure: 2 (15 drives)
CUMULATIVE ERROR COUNTS
 Controller A
                                10/30/10 1:15:59 PM
  Baseline time set:
   Sample period (days, hh:mm:ss): 32 days, 00:55:04
  Controller detected errors:
                                     0
  Drive detected errors:
                                  48
  Timeout errors:
                              1
  Link down errors:
                               N/A
   Total I/O count:
                               199070838
 Controller B
  Baseline time set: 10/30/10 1:15:59 PM
   Sample period (days, hh:mm:ss): 32 days, 00:53:22
   Controller detected errors:
                                    0
   Drive detected errors:
                                  52
  Timeout errors:
                              0
                              N/A
  Link down errors:
  Total I/O count:
                              198778804
DRIVE CHANNEL 2
Port: 6, 5
Status: Optimal
Max. Rate: 4 Gbps
Current Rate: 4 Gbps
 Rate Control: Auto
Controller A link status: Up
```

Controller B link status: Up Trunking active: No DRIVE COUNTS Total # of attached drives: 0 CUMULATIVE ERROR COUNTS Controller A 10/30/10 1:15:59 PM Baseline time set: Sample period (days, hh:mm:ss): 32 days,00:55:04 Controller detected errors: 0 Drive detected errors: 0 Timeout errors: 2 Link down errors: N/A Total I/O count: 14238433 Controller B Baseline time set: 10/30/10 1:15:59 PM Sample period (days, hh:mm:ss): 32 days, 00:53:22 Controller detected errors: 0 0 Drive detected errors: Timeout errors: 0 Link down errors: N/A Total I/O count: 13470436 DRIVE CHANNEL 3 Port: 6, 5 Status: Optimal Max. Rate: 4 Gbps Current Rate: 4 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up Trunking active: No DRIVE COUNTS Total # of attached drives: 0 CUMULATIVE ERROR COUNTS Controller A Baseline time set: 10/30/10 1:15:59 PM Sample period (days, hh:mm:ss): 32 days, 00:55:04 Controller detected errors: 0 Drive detected errors: 0 Timeout errors: 0 Link down errors: N/A 13414513 Total I/O count: Controller B Baseline time set: 10/30/10 1:15:59 PM Sample period (days, hh:mm:ss): 32 days, 00:53:22 Controller detected errors: 0 Drive detected errors: 0 Timeout errors: 0 Link down errors: N/A Total I/O count: 13201515 DRIVE CHANNEL 4 Port: 2, 1 Status: Optimal Max. Rate: 4 Gbps Current Rate: 2 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up Trunking active: No DRIVE COUNTS Total # of attached drives: 0 CUMULATIVE ERROR COUNTS

Controller A Baseline time set: 10/30/10 1:15:59 PM Sample period (days, hh:mm:ss): 32 days, 00:55:04 Controller detected errors: 111 Drive detected errors: 0 Timeout errors: 0 Link down errors: N/A Total I/O count: 13093814 Controller B Baseline time set: 10/30/10 1:15:59 PM Sample period (days, hh:mm:ss): 32 days, 00:53:22 Controller detected errors: 54 Drive detected errors: 0 Timeout errors: 0 Link down errors: N/A Total I/O count: 13039285 DRIVE CHANNEL 5 Port: 1, 2, ESM B 1B, ESM B 1A, ESM B 1B, ESM B 1A, ESM B 1B Status: Optimal Max. Rate: 4 Gbps Current Rate: 4 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up Trunking active: No DRIVE COUNTS Total # of attached drives: 44 Connected to: Controller B, Port 1 Attached drives: 44 Drive enclosure: 3 (14 drives) Drive enclosure: 1 (15 drives) Drive enclosure: 2 (15 drives) CUMULATIVE ERROR COUNTS Controller A Baseline time set: 10/30/10 1:15:59 PM Sample period (days, hh:mm:ss): 32 days, 00:55:04 Controller detected errors: 0 Drive detected errors: 49 Timeout errors: 1 Link down errors: N/A Total I/O count: 183366503 Controller B Baseline time set: 10/30/10 1:15:59 PM 32 days, 00:53:22 Sample period (days, hh:mm:ss): Controller detected errors: 1 Drive detected errors: 52 Timeout errors: 0 Link down errors: N/A 182512319 Total I/O count: DRIVE CHANNEL 6 Port: 3, 4 Status: Optimal Max. Rate: 4 Gbps Current Rate: 2 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up Trunking active: No DRIVE COUNTS Total # of attached drives: 0 CUMULATIVE ERROR COUNTS

Controller A Baseline time set: 10/30/10 1:15:59 PM Sample period (days, hh:mm:ss): 32 days, 00:55:04 Controller detected errors: 0 Drive detected errors: 0 Timeout errors: 0 Link down errors: 0 Total I/O count: 13296480 Controller B Baseline time set: 10/30/10 1:15:59 PM Sample period (days, hh:mm:ss): 32 days, 00:53:22 Controller detected errors: 0 0 Drive detected errors: Timeout errors: 0 Link down errors: N/A Total I/O count: 13275865 DRIVE CHANNEL 7 Port: 5, 6 Status: Optimal Max. Rate: 4 Gbps Current Rate: 2 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up Trunking active: No DRIVE COUNTS Total # of attached drives: 0 CUMULATIVE ERROR COUNTS Controller A Baseline time set: 10/30/10 1:15:59 PM Sample period (days, hh:mm:ss): 32 days, 00:55:04 Controller detected errors: 0 0 Drive detected errors: Timeout errors: 0 Link down errors: 0 Total I/O count: 131818784 Controller B Baseline time set: 10/30/10 1:15:59 PM Sample period (days, hh:mm:ss): 32 days, 00:53:22 0 Controller detected errors: Drive detected errors: 0 Timeout errors: 0 Link down errors: N/A Total I/O count: 13171844 DRIVE CHANNEL 8 Port: 7, 8 Status: Optimal Max. Rate: 4 Gbps Current Rate: 4 Gbps Rate Control: Auto Controller A link status: Up Controller B link status: Up Trunking active: No DRIVE COUNTS Total # of attached drives: 0 CUMULATIVE ERROR COUNTS Controller A Baseline time set: 10/30/10 1:15:59 PM Sample period (days, hh:mm:ss): 32 days, 00:55:04 Controller detected errors: 44 0 Drive detected errors:

Timeout errors:	0
Link down errors:	0
Total I/O count:	13067464
Controller B	
Baseline time set: 10/30/10 1:15:59 PM	
Sample period (days, hh:mm:ss):	32 days, 00:53:22
Controller detected errors:	25
Drive detected errors:	0
Timeout errors:	Θ
Link down errors:	N/A
Total I/O count:	12987004

## Show drive

The show drive command returns information about the drives in a storage subsystem. The show storageArray command returns information about the components and the features in a storage subsystem. If you run the command with the profile parameter, the command returns information in the form shown by this example. This information is the most detailed report that you can receive about the storage subsystem. After you have configured a storage subsystem, save the configuration description to a file as a reference. DRIVES-----SUMMARY Number of drives: 14 Supported drive types: Fibre (14) BASIC:

TRAY,	SLOT	STATUS	CAPACITY	$\subset l$	URRENT DATA RATE	Ξ	PRODUCT ID	FIRMWARE VERSI	ION
1, 1		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 2		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 3		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 4		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 5		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 6		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 7		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 8		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 9		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 10	b.	Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 11		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 12		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 13		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	
1, 14		Optimal	68.366 GB	2	Gbps		ST373453FC	MS03	

#### DRIVE CHANNELS:

TR/	ΑY,	SLOT	PREFERRED	CHANNEL	REDUNI	DANT (	THANNEL	
1,	1		3		Drive	path	redundancy:	Lost
1,	2		3		Drive	path	redundancy:	Lost
1,	3		3		Drive	path	redundancy:	Lost
1,	4		3		Drive	path	redundancy:	Lost
1,	5		3		Drive	path	redundancy:	Lost
1,	6		3		Drive	path	redundancy:	Lost
1,	7		3		Drive	path	redundancy:	Lost
1,	8		3		Drive	path	redundancy:	Lost
1,	9		3		Drive	path	redundancy:	Lost
1,	10		3		Drive	path	redundancy:	Lost
1,	11		3		Drive	path	redundancy:	Lost
1,	12		3		Drive	path	redundancy:	Lost
1,	13		3		Drive	path	redundancy:	Lost
1,	14		3		Drive	path	redundancy:	Lost

HOT SPARE COVERAGE: The following volume groups are not protected: 3

Total hot spare drives: 2 Standby: 2 In use: 0 Standby drive at tray 1, slot 8 (Fibre, 68.366 GB) Protects the following volume groups: 2, 1

DETAILS Drive at Tray 1, Slot 1 Drive port: 1, Channel: 3, ID: 8/0xD9 Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW0535L00007329LFD8 Vendor: SEAGATE Date of manufacture: January 20, 2003 World-wide name: 20:00:00:04:cf:0b:c6:0d Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 2 Drive at Tray 1, Slot 2 Drive port: 1, Channel: 3, ID: 9/0xD6 Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW056D500007329VVAP Vendor: SEAGATE Date of manufacture: January 23, 2003 World-wide name: 20:00:00:04:cf:0b:dc:29 Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 2 Drive at Tray 1, Slot 3 Drive port: 1, Channel: 3, ID: 10/0xD5 Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW04LMK00007329DYZK Vendor: SEAGATE Date of manufacture: January 17, 2003 World-wide name: 20:00:00:04:cf:0b:b1:f6 Drive type: Fibre Channel Speed: 15015 RPM Mode: Unassigned Associated volume group: Volume group None

Drive at Tray 1, Slot 4 Drive port: 1, Channel: 3, ID: 11/0xD4 Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW04N2T00007329DZGZ Vendor: SEAGATE Date of manufacture: January 17, 2003 World-wide name: 20:00:00:04:cf:0b:b1:81 Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 3 Drive at Tray 1, Slot 5 Drive port: 1, Channel: 3, ID: 12/0xD3 Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW04YN500007329JQYG Vendor: SEAGATE Date of manufacture: January 23, 2003 World-wide name: 20:00:00:04:cf:0b:db:2c Drive type: Fibre Channel Speed: 15015 RPM Mode: Unassigned Associated volume group: Volume group None Drive at Tray 1, Slot 6 Drive port: 1, Channel: 3, ID: 13/0xD2 Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MSO3 Serial number: 3HW056N300007328ZC57 Vendor: SEAGATE Date of manufacture: January 23, 2003 World-wide name: 20:00:00:04:cf:0b:db:66 Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 2

Drive at Tray 1, Slot 7 Drive port: 1, Channel: 3, ID: 14/0xD1 Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW04×KQ00007329VUR8 Vendor: SEAGATE Date of manufacture: January 23, 2003 World-wide name: 20:00:00:04:cf:0b:db:a1 Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 2 Drive at Tray 1, Slot 8 Drive port: 1, Channel: 3, ID: 15/0xCE Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW04QL300007328H6BU Vendor: SEAGATE Date of manufacture: January 23, 2003 World-wide name: 20:00:00:04:cf:0b:db:6f Drive type: Fibre Channel Speed: 15015 RPM Mode: Hot spare standby Associated volume group: Volume group None Drive at Tray 1, Slot 9 Drive port: 1, Channel: 3, ID: 81/0x54 Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW04W5E00007329EQX2 Vendor: SEAGATE Date of manufacture: January 18, 2003 World-wide name: 20:00:00:04:cf:0b:b2:5b Drive type: Fibre Channel Speed: 15015 RPM Mode: Unassigned Associated volume group: Volume group None

Drive at Tray 1, Slot 10 Drive port: 1, Channel: 3, ID: 97/0x39 Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW04V5P00007329DZPH Vendor: SEAGATE Date of manufacture: January 17, 2003 World-wide name: 20:00:00:04:cf:0b:b1:97 Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 2 Drive at Tray 1, Slot 11 Drive port: 1, Channel: 3, ID: 65/0x71 Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW056N700007329T7EG Vendor: SEAGATE Date of manufacture: January 23, 2003 World-wide name: 20:00:00:04:cf:0b:db:69 Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 2 Drive at Tray 1, Slot 12 Drive port: 1, Channel: 3, ID: 73/0x66 Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW04DL400007329ERF1 Vendor: SEAGATE Date of manufacture: January 17, 2003 World-wide name: 20:00:00:04:cf:0b:b1:84 Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 1

Drive at Tray 1, Slot 13 Drive port: 1, Channel: 3, ID: 89/0x4A Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW04DJT00007329DZJF Vendor: SEAGATE Date of manufacture: January 18, 2003 World-wide name: 20:00:00:04:cf:0b:b0:ff Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 1 Drive at Tray 1, Slot 14 Drive port: 1, Channel: 3, ID: 105/0x2D Drive path redundancy: Lost Status: Optimal Raw capacity: 68.366 GB Usable capacity: 67.866 GB Current data rate: 2 Gbps Product ID: ST373453FC Firmware version: MS03 Serial number: 3HW04VAK00007329ER3P Vendor: SEAGATE Date of manufacture: January 17, 2003 World-wide name: 20:00:00:04:cf:0b:b1:6e Drive type: Fibre Channel Speed: 15015 RPM Mode: Assigned Associated volume group: Volume group 1

# Appendix B. Example script files

This appendix provides example scripts for configuring a storage subsystem. These examples show how the script commands appear in a complete script file. Also, you can copy these scripts and modify them to create a configuration unique to your storage subsystem.

You can create a script file in two ways:

- Using the save storageArray configuration command
- Writing a script

By using the save storageArray configuration command, you can create a file that you can use to copy an existing configuration from one storage subsystem to other storage subsystems. You can also use this file to restore an existing configuration that has become corrupted. You also can copy an existing file to serve as a pattern from which you create a new script file by modifying portions of the original file. The default file extension is .scr.

You can create a new script file by using a text editor, such as Microsoft Notepad. The maximum line length is 256 characters. The command syntax must conform to the guidelines in the Chapter 1, "About the Command Line Interface," on page 1 topic, and the Chapter 2, "About the Script Commands," on page 15 topic. When you create a new script file, you can use any file name and extension that will run on the host operating system

This example shows how to run a script file from the command line.

c:\...\smX\client>smcli 123.45.67.88 123.45.67.89
-f scriptfile.scr;

## **Configuration script example 1**

This example creates a new volume by using the create volume command in the free space of a Array.

```
Show "Create RAID 5 Volume 7 on existing Array 1";
//Create volume on Array created by the create
volume drives command
//Note: For Arrays that use all available
capacity, the last volume on the group is created using
all remaining capacity by omitting the capacity=volume
creation parameter
create volume volumeGroup=1 RAIDLevel=5 userLabel="7"
owner=A segmentSize=16 cacheReadPrefetch=TRUE capacity=2GB;
show "Setting additional attributes for volume 7";
//Configuration settings that cannot be set during volume
creation
set volume["7"] cacheFlushModifier=10;
set volume["7"] cacheWithoutBatteryEnabled=false;
set volume["7"] mirrorEnabled=true;
set volume["7"] readCacheEnabled=true;
set volume["7"] writeCacheEnabled=true;
set volume["7"] mediaScanEnabled=false;
set volume["7"] redundancyCheckEnabled=false;
set volume["7"] modificationPriority=high;
```

This example shows blank lines between the lines beginning with Show, Create,//Note, and create. The blank lines are included in this example only for clarity. Each command is actually written on one line in the script file; however, the size of this page has caused the command text to wrap. You might want to include blank lines in your script files to separate blocks of commands or make a comment that stands out. To include a comment, enter two forward slashes (//), which causes the Script Engine to treat the line as a comment.

The first line of text is the show string command. This command shows text that is bounded by double quotation marks (" ") on a display monitor when the script file runs. In this example, the text Create RAID 5 Volume 7 on existing Volume Group 1 serves as a title that describes the expected results of running this script file.

The line beginning with //Create is a comment that explains that the purpose of this script file is to create a new volume by using the create volume command on an existing Array.

The line beginning //Note: is a comment in the script file that explains that the size of the last volume created that uses all of the available capacity because the capacity parameter is not used.

The command in this example creates a new volume in Array 1. The volume has RAID Level 5. The volume name (user label) is 7. (Note the double quotation marks around the 7. The double quotation marks define that the information in the double quotation marks is a label.) The new volume is assigned to the controller in slot A in the controller enclosure. The segment size is set to 16. The volume has a read ahead multiplier value of 256. The capacity of the volume is 2 GB.

The command takes this form: create volume volumeGroup=volumeGroupNumber userLabel=volumeName [freeCapacityArea=freeCapacityIndexNumber] [capacity=volumeCapacity | owner=(a | b) | cacheReadPrefetch=(TRUE | FALSE) | segmentSize=segmentSizeValue]

[enclosureLossProtect=(TRUE | FALSE)]

The general form of the command shows the optional parameters in a different sequence than the optional parameters in the example command. You can enter optional parameters in any sequence. You must enter the required parameters in the sequence shown in the command descriptions.

The line showing "Setting additional attributes for volume 7" is another example of using the show "string" command. The reason for placing this command here is to tell the user that the create volume command ran successfully and that properties that could not be set by the create volume command are now set.

The set volume parameters are shown on separate lines. You do not need to use separate lines for each parameter. You can enter more than one parameter with the set volume command by leaving a space between the parameters, as in this example:

```
set volume["7"] cacheFlushModifier=10
cacheWithoutBatteryEnabled=false
modificationPriority=high;
```

By using separate lines, you can see more clearly the parameters that you are setting and the values to which you are setting the parameters. Blocking the parameters in this manner makes it easier to either edit the file or copy specific parameter settings for use in another script file.

## Configuration script example 2

This example creates a new volume by using the create volume command with user-defined drives in the storage subsystem.

```
Show "Create RAID3 Volume 2 on existing Array 2";
//This command creates the Array and the initial volume on that group.
//Note: For Arrays that use all available capacity, the last volume
on the Array is created using all remaining capacity by omitting the
capacity=volume creation parameter
250 Configuration Script Example 2
create volume RAIDLevel=3 userLabel="2"
drives=[0,1 0,6 1,7 1,3 2,3 2,6] owner=B
segmentSize=16 capacity=2GB;
show "Setting additional attributes for voluem 7"'
//Configuration settings that cannot be set during volume creation
set volume ["7"] cacheFlushModifier=10;
set volume ["7"] cacheWithoutBatteryEnabled=false;
set volume ["7"] readCacheEnabled=true;
```

set volume ["7"] writeCacheEnabled=true;

set volume ["7"] mediaScanEnabled=false; set volume ["7"] redundantCheckEnabled=false;

set volume ["7"] modificationPriority=high;

The command in this example, like the create volume command in the previous example, creates a new volume. The significant difference between these two examples is that this example shows how you can define specific drives to include in the volume. Use the show storageArray profile command to find out what drives are available in a storage subsystem.

The create volume command takes this form:

```
create volume raidLevel=(0 | 1 | 3 | 5 | 6) userLabel=volumeName
drives=(enclosureID1,slotID1...enclosureIDn,slotIDn)
[capacity=volumeCapacity | owner=(a | b) |
cacheReadPrefetch=(TRUE | FALSE) |
segmentSize=segmentSizeValue]
[enclosureLossProtect=(TRUE | FALSE)]
```

# Appendix C. Global Copy mirror utility

This appendix describes the host utility to achieve periodic consistency with Global Copy Mirror configurations. This appendix also describes how to run the Global Copy utility.

**Note:** The Global Copy Mirror utility works only with the synchronous remote mirror commands. This utility does not work with the Enhanced Global Remote Mirror commands.

# **Description of the Global Copy Mirror Utility**

The Global Copy Mirror utility enables you to periodically synchronize the Enhanced Remote Mirroring pairs in your storage subsystem. When defining a Enhanced Remote Mirroring configuration, you have the option to set the write modes to either Synchronous or Asynchronous. Metro Mirror provides the highest level security for full data recovery from the secondary storage subsystem in the event of a disaster. Metro Mirror does, however, reduce host I/O performance. Global Copy offers faster host I/O performance, but it does not guarantee that a copy operation has successfully completed before processing the next write request. With Global Copy, you cannot make sure that a volume, or collection of volumes, at a secondary site ever reach a consistent, recoverable state.

The Global Copy Mirror utility enables you to bring a collection of Enhanced Global remote logical drives into a mutually consistent and recoverable state. You can choose to run the utility based on application demands, link state and speed, and other factors that are relevant to your environment.

The Global Copy Mirror utility has these characteristics:

- The utility is implemented as a command line-invoked Java-based application.
- The utility is bundled as part of the IBM DS Storage Manager 10 installation package.
- The utility accepts a command line argument that lets you specify the name of a configuration file that contains a complete specification of the work to be carried out by the utility.
- More than one instance of the utility can run concurrently, as long as the utilities do not try to process any of the same volumes and mirrors.

**Note:** The Global Copy Mirror utility does not check to make sure that concurrently running instances of the utility are not trying to process the same volumes and mirrors. If you choose to simultaneously run more than one instance of the Global Copy Mirror utility, you must make sure that the configuration files that you choose to run do not list the same volumes and mirrors.

## **Operation of the Enhanced Remote Mirroring utility**

The Global Copy Mirror utility performs steps that generate a recoverable state for multiple mirror volumes at a secondary site. The utility runs these steps to create consistent, recoverable images of a set of volumes:

- 1. On the primary storage subsystem The utility reconfigures all of the participating volumes from Global Copy mirroring to Metro Mirror mirroring. This action makes sure that the stream of write operations becomes recoverable on the secondary side.
- 2. On the primary storage subsystem The utility polls all of the participating volumes until the associated mirror states all have the Optimal state. In cases where the remote link is slow or the primary host I/O activity is high, one or more mirrors are likely to be in the Unsynchronized state before they migrate to the Synchronized state. By waiting until all of the mirrors have Optimal status, the utility makes sure that all of the delta logs for the affected volumes are cleared, and the secondary volumes are recoverable.
- 3. On the primary storage subsystem The utility suspends the mirrored pairs for all of the participating volumes. This action causes updates to stop on the secondary side, leaving the secondary volumes in a recoverable state because they were being updated in Synchronous mode immediately before the suspension. By separating the mirrors in this manner, the primary-side applications run faster, while leaving the secondary volumes in a recoverable state. The delta log tracks changes made because of application writes on the primary side while in this state.
- 4. **On the secondary storage subsystem** The utility generates a FlashCopy of each participating volume on the secondary side, which creates point-in-time images that are recoverable.
- 5. On the primary storage subsystem The utility resumes the mirroring operations for all of the participating volumes. This action causes the mirrors to migrate to the Synchronized state and start the process of restoring coherency between the primary site and the secondary site.
- 6. On the primary storage subsystem The utility reconfigures all of the affected volumes for Global Copy mode.

# **Running the Enhanced Remote Mirroring Utility**

The Global Copy Mirror utility uses a command line argument that enables you to specify the name of a configuration file. The configuration file contains a complete specification of the input parameters that are needed by the utility. To run the utility, enter this syntax:

asyncRVMUtil configuration\_file -d debug\_file

In this command, configuration\_file is the file that you provide as input. The configuration file specifies the Enhanced Remote Mirroring volumes that you want to synchronize by using the utility. When you create the configuration file, use these conditions to define the volumes in the file:

- All the primary volumes in a volume set must belong to the same storage subsystem.
- The maximum number of volume sets that you can specify in the file is four.
- The maximum number of mirrored pairs that you can specify as part of a consistency group is eight.

The optional parameter, -d, lets you specify a file to which you can send information regarding how the utility runs. In this example, the file name is debug\_file. The debug file contains trace information that can be reviewed by your Technical Support representative to determine how well the Global Copy Mirror utility has run.

**Note:** Depending on the location of the configuration file and the debug file, you must specify the complete path with the file name.

To run the Global Copy Mirror utility, you must enter the asyncRVMUtil command from the command line. Because UNIX operating systems are case sensitive, you must type the command exactly as shown. On Windows operating systems, you can type the command in all uppercase, in all lowercase, or in mixed case.

**Note:** To use the Global Copy Mirror utility, you must be managing the storage subsystem by using the command line interface, not the graphical user interface of IBM DS Storage Manager 10.

#### Configuration utility

The configuration file is an ASCII flat text file that provides the information for the Enhanced Remote Mirroring synchronization used by the Global Copy Mirror utility. The file defines the mirror volume sets to be synchronized. All of the mirror volumes in the volume sets defined in the configuration file are run collectively to create a recoverable image. If any one of the mirrors in the volume set fails, the operation is stopped for this volume set and carried on to the next volume set that is listed in the configuration file.

The configuration file supports this syntax:

```
content ::= {spec}
spec ::= logSpec | volumeSetSpec
logSpec ::= "Log" "{" {logAttribute} "}"
logAttribute ::= fileSpec
fileSpec ::= "file" "=" fileName
volumeSetSpec ::= "VolumeSet" volumeSetName
"{" {volumeSetAttribute} "}"
volumeSetAttribute ::= timeoutSpec | mirrorSpec
timeoutSpec ::= "OptimalWaitTimeLimit" "=" integer
mirrorSpec ::= "Mirror" "{" {mirrorAttribute} "}"
mirrorAttribute ::= primarySpec | secondarySpec |
Enhanced FlashCopySpec
primarySpec ::= "Primary" "=" volumeSpec
secondarySpec ::= "Secondary" "=" volumeSpec
Enhanced FlashCopySpec ::= "Copy" "=" volumeSpec
volumeSpec ::= storageArrayName"."volumeUserLabel
```

In this syntax, items enclosed in double quotation marks ("") are terminal symbols. Items separated by a vertical bar (1) are alternative values (enter one or the other, but not both). Items enclosed in curly braces ({ }) are optional (you can use the item zero or more times).

These definitions are provided for non-terminals in the syntax:

- integer The timeout value must be an integer (decimal digits from 0–9).
- volumeSetName The name of the set of volumes on which you want to run the Global Copy Mirror utility.
- fileName The name of a file, using characters and conventions that are appropriate for the system on which the application is running.
- storageArrayName The label that you have assigned for a storage subsystem, as would be used in the CLI to specify the name of the storage subsystem.
- volumeUserLabel The label that you have assigned for a volume that uniquely identifies the volume within the storage array.

**Note:** Names and labels can be any characters that are defined as appropriate for your operating system. The maximum length for a name or label is 30 characters. If the name or label contains special characters (as defined by the operating system) or period characters, you must enclose the name or label in double quotation marks (""). You can, optionally, enclose the name or label in double quotation marks at any time.

These items are considered syntax errors:

- More than one logSpec command in the input file
- Zero or more than one fileSpec attribute in a logSpec command (you must include exactly one fileSpec attribute in the logSpec command)
- More than one timeoutSpec attribute in a volumeSetSpec command
- Zero or more than one primarySpec attribute in a mirrorSpec command (you must include exactly one primarySpec attribute in the mirrorSpec command)
- Zero or more than one secondarySpec attribute in a mirrorSpec command (you
  must include exactly one secondarySpec attribute in the mirrorSpec command)
- Zero or more than one Enhanced FlashCopySpec attribute in a mirrorSpec command (you must include exactly one Enhanced FlashCopySpec attribute in the mirrorSpec command)

**Note:** In the Global Copy Mirror utility configuration file, you must specify the primary volume, the secondary volume, and the copy (FlashCopy) volume. The utility does not make sure that the secondary volume is correct for the Enhanced Remote Mirroring relationship. The utility also does not make sure that the Enhanced FlashCopy (legacy) volume is actually a FlashCopy for the secondary volume. *You must make sure that these volumes are correct.* If the volumes are not correct, the utility will run, but the volumes will not be consistent. For each mirror, the secondary volume and the copy volume must reside on the same storage subsystem.

This example shows a configuration file for the Enhanced Remote Mirroring utility

```
Log{ file="d:\rvm-consistency.log" }
VolumeSet "set1" {
optimalWaitTimeLimit = 15
Mirror {
Primary = LosAngelesArray.PayrollVolume
Secondary = NewYorkArray.PayrollVolume
Copy = NewYorkArray.PayrollVolumeImage
Mirror {
Primary = LosAngelesArray.PayrollVolume
Secondary = BostonArray.PayrollVolume
Copy = BostonArray.PayrollVolumeImage
VolumeSet "set2" {
Mirror {
Primary = BostonArray.HRVolume
Secondary = LosAngelesArray.HRVolume
Copy = LosAngelesArray.HRVolumeImage
```

# Appendix D. Simplex-to-Duplex Conversion

Some models of controller enclosures and drive enclosures are available in either a simplex configuration (one controller) or a duplex configuration (two controllers). You can convert a simplex configuration to a duplex configuration by installing new nonvolatile static random access memory (NVSRAM) and a second controller. This appendix explains how to convert a simplex configuration to a duplex configuration by using CLI commands or by using the storage management software.

# **General steps**

You can upgrade a controller enclosure or a drive enclosure that has a simplex configuration to a duplex configuration by performing these tasks:

- 1. Install new NVSRAM on the existing controller in your controller enclosure or drive enclosure.
- 2. Revise the controller enclosure configuration or the drive enclosure configuration to run with two controllers.
- 3. Install a second controller.
- 4. Connect the host cables.
- 5. Connect the drive enclosure cables.
- 6. Run diagnostics to make sure that your new configuration is running correctly.

# **Tools and Equipment**

The procedures in this appendix require these items:

- Anti-static protection
- A No. 2 Phillips screwdriver
- A second controller
- Small Form-factor Pluggable (SFP) transceivers (for Fibre Channel configurations)
- Host-to-controller cables
- Controller-to-environmental services module (ESM) cables

# Step 1 – Installing the duplex NVSRAM

**Note:** Before trying to download NVSRAM, you must contact your Technical Support representative to make sure that you are downloading the NVSRAM that is appropriate for the controller in your storage subsystem.

NVSRAM files specify the default settings for the controller enclosure controllers or drive enclosure controllers. Follow the instructions in this step to upgrade the NVSRAM on the controller in your controller enclosure or your drive enclosure.

To get a copy of the latest NVSRAM, perform one of these tasks:

- Download the duplex NVSRAM by using the command line interface.
- Download the duplex NVSRAM by using the graphical user interface (GUI) of the storage management software.
- Copy the duplex NVSRAM from the installation CD in the conversion kit.

Ensure that the controller enclosure or the drive enclosure has an Optimal status. If one or more managed devices has a Needs Attention status, determine and correct the condition that created the Needs Attention status before proceeding with this conversion instruction.

#### Downloading the NVSRAM by using the command line interface

- 1. Make a copy of your storage subsystem profile, and save it in the event that you might need to restore the storage subsystem.
- 2. Start the command line interface.
- 3. On the command line, type this command, and press Enter. In this command, ctlr-A\_IP\_address is the IP address of the of the original simplex controller and filename is the complete file path and name of the file that contains the new NVSRAM. Valid file names must end with a .dlp extension. Enclose the file name in double quotation marks ("").

smcli ctlr-A\_IP\_address -c "download storageArray
NVSRAM file="filename";"

#### Downloading the NVSRAM by using the GUI

- 1. Make a copy of your storage subsystem profile, and save it in the event that you might need to restore the storage subsystem.
- 2. At the storage management station, start the SMclient software.
- In the Subsystem Management Window, select Advanced >> Maintenance >> Download >> Controller NVSRAM.
- 4. In the Download NVSRAM dialog, enter the NVSRAM file name in the Selected NVSRAM text box. If you do not know the file name, click Browse, and navigate to a folder with the NVSRAM files.
- 5. Select the file that corresponds to your storage subsystem type.
- 6. Click OK.

The **Confirm Download** dialog appears.

- 7. To start the download, click **Yes**. 8. Based on the dialog that appears after the download has completed, perform one of these actions:
  - Download Successful dialog Click Done.

— **Error dialog** – Read the information in the dialog, and take the appropriate action.

#### Copying NVSRAM from the installation CD

- 1. Make a copy of your storage subsystem profile, and save it in the event that you might need to restore the storage subsystem.
- 2. Insert the Installation CD into the CD-ROM drive.
- **3**. At the storage management station, start the SMclient software.
- In the Subsystem Management Window, select Advanced >> Maintenance >> Download >> Controller NVSRAM.
- 5. In the **Download NVSRAM** dialog, select the CD-ROM drive and the */nvsram* folder. Either double-click the folder or type the folder name in the **Selected NVSRAM** file text box.
- 6. Select the file that corresponds to your storage subsystem type.
- 7. Click OK.

The Confirm Download dialog appears.

8. To start the download, click Yes.

- **9**. Based on the dialog that appears after the download is completed, perform one of these actions:
  - Download Successful dialog Click Done.

— **Error dialog** – Read the information in the dialog, and take the appropriate action.

# Step 2 – Setting the configuration to duplex

After rebooting the controller enclosure or the drive enclosure, an "alternate controller missing" error message appears. This message indicates that the controller in slot A has successfully converted to Duplex mode. This message persists until you have completed the tasks to install the second controller, installed the host cables, and installed the drive enclosure cables.

- 1. Start the command line interface.
- On the command line, type this command, and press Enter. In this command, ctlr-A\_IP\_address is the IP address of the of the original simplex controller.
   smcli ctlr-A\_IP\_address -c "set storageArray redundancyMode=duplex;"
- 3. Reboot the controller enclosure or the drive enclosure.

# Step 3 – Installing the second controller

**Attention: Possible hardware damage** – To prevent electrostatic discharge damage to the enclosure, use proper anti-static protection when handling enclosure components.

**Note:** For best operation, the new controller must have a part number identical to the existing controller, or the new controller must be a certified substitute. The part number is on a label on the controller. To provide full functionality in dual-controller configurations, make sure that both controllers in the controller enclosure or the drive enclosure have the same memory capacity. Although you can install two controllers of different memories in a controller enclosure or a drive enclosure, the mismatch disables some functions, such as cache mirroring.

1. Put on anti-static protection.

Attention: Possible damage to the controller – Do not remove the electrostatic protection until you have finished installing the controller and you have connected the host cables and the drive enclosure cables.

2. Unpack the new controller.

**Attention:** Possible damage to the controller – Bumping the controller against another surface might damage the data connectors on the rear of the controller. Use caution when handling the controller.

- **3**. Remove the blank controller CRU/FRU from the enclosure by releasing the handle, and pulling the blank controller CRU/FRU out of the enclosure.
- 4. Slide the new controller CRU/FRU into the empty slot by pushing the controller CRU/FRU until it snaps into place, and locking the handle into the closed position.

# Step 4 – Connecting the host cables

The steps in this procedure describe how to attach Fibre Channel host cables. The steps for connecting other types of host cables are similar, but they do not require the installation of Small Form-factor Pluggable (SFP) transceivers.

1. If there is a black plastic plug in the host port, remove it.

**2.** Install an SFP transceiver into the controller by pushing the SFP transceiver into the host port until it snaps into place.

Attention: Possible degraded performance – To prevent degraded performance, do not twist, fold, pinch, or step on fiber-optic cables. Do not bend fiber-optic cables tighter than a 5-cm (2-in.) radius.

- 3. Plug one end of the fiber-optic cable into the SFP transceiver in the host port.
- 4. Plug the other end of the fiber-optic cable into one of the HBAs in the host (direct topology) or into a switch (switch topology).
- 5. Attach a label to each end of the cable by using this scheme. A label is very important if you need to disconnect the cables later to service a controller.
  - The host name and the host bus adapter (HBA) port (if direct topology)
  - The switch name and port (if switch topology)
  - The controller ID (for example, controller A)

— The host channel ID (for example, host channel 1)**Example label abbreviation** – Assume that a cable is connected between port 1 in HBA 1 of a host named Engineering and host channel 1 of controller A. A label abbreviation could be as follows:

Heng-ABA1/P1, CtA-Hch1

6. Repeat step 1 through step 5 for each host channel that you intend to use.

#### Step 5 – Connecting the controller to a drive enclosure

The steps in this procedure describe how to attach Fibre Channel cables to a drive enclosure. The steps for connecting other types of drive enclosure cables are similar, but they do not require the installation of SFP transceivers.

- 1. If there is a black plastic plug in the drive port of the new controller CRU/FRU, remove it.
- 2. Insert an SFP transceiver into the drive port on a controller CRU/FRU.
- **3**. Plug one end of the cable into the SFP transceiver.
- 4. Plug the other end of the cable into the appropriate in port or out port on the environmental services module (ESM) in the drive enclosure as applicable for your cabling configuration.
- 5. Attach a label to each end of the cable by using this scheme. A label is very important if you need to disconnect the cables later to service a controller.
  - The controller ID (for example, controller A)
  - The drive channel number and port ID (for example, drive channel 1, port 2)
  - The ESM ID (for example, ESM A)
  - The ESM port ID (for example, In, Out, 1A, or 1B)

— The drive enclosure ID**Example label abbreviation** – Assume that a cable is connected between drive channel 1, port 2 of controller A to the out port of the left ESM (A) in drive enclosure 1. A label abbreviation could be as follows: CtA-Dch1/P2, Dm1-ESM A (left), Out

6. Repeat step 1 through step 5 for each drive enclosure.

7. Remove the anti-static protection.

#### Step 6 – Running diagnostics

- 1. Using the LEDs on the storage subsystem and information provided by the storage management software, check the status of all enclosures in the storage subsystem.
- **2**. Does any component have a Needs Attention status?

— Yes – Click the Recovery Guru toolbar button in the Subsystem Management Window, and complete the recovery procedure. If a problem is still indicated, contact your Technical Support representative.

— No – Go to step 3.

3. Create, save, and print a new storage subsystem profile.

# **Appendix E. Additional Documentation**

The following tables present an overview of the IBM System Storage DS Storage Manager, storage subsystem, and storage expansion enclosure product libraries, and other related documents. Each table lists documents that are included in the libraries and what common tasks they address.

To access the documents, go to http://www.ibm.com/servers/storage/support/disk/ or http://www.ibm.com/shop/publications/order/.

# DS3000 storage subsystem documents

The following table lists the DS3000 storage subsystem documents and the common user tasks addressed by each document.

Title	User Tasks								
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance			
IBM System Storage DS3000 Storage Subsystem Installation, User's and Maintenance Guides	-	-		~	~	~			
IBM System Storage DS3000 Quick Start Guides		100							
IBM System Storage DS3500 and EXP3500 Installation, User's and Maintenance Guide	-	-		~	~	~			
IBM System Storage DS3500 and EXP3500 Rack Installation and Quick Start Guide		-							

## DS4000 storage subsystem documents

The following table lists the DS4000 storage subsystem documents and the common user tasks addressed by each document.

Title	User Tasks							
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance		
IBM System Storage DS4000 Storage Subsystem Installation, User's and Maintenance Guides	-	100		~	~			
IBM System Storage DS4000 Quick Start Guides								

# DS3000, DS4000, and DS5000 storage expansion enclosure documents

The following table lists the DS3000, DS4000, and DS5000 storage expansion enclosure documents and the common user tasks addressed by each document.

Title	User Tasks								
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance			
IBM System Storage EXPxxx Storage Expansion Enclosure Installation, User's, and Maintenance Guide	~	~		~	~				
IBM System Storage Quick Start Guide, Quick Reference for the DS4700 and DS4200, Sections 2, 3, and 4 also for installing the EXP810 and EXP420				~					
IBM System Storage Hard Drive and Storage Expansion Enclosures Installation and Migration Guide	4	~							
IBM System Storage Quick Start Guide, Quick Reference for the DS5100 and DS5300, and the EXP5000		~	L	~					

# Other DS3000 and DS4000-related documents

The following table lists other DS3000 and DS4000-related documents and the common user tasks addressed by each document.

Title	User Tasks								
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance			
IBM Safety Information					~				
IBM TotalStorage Hardware Maintenance Manual						~			
IBM System Storage Problem Determination Guide						~			
IBM Fibre Channel Planning and Integration: User's Guide and Service Information	~	-			~	~			
IBM TotalStorage FC2-133 Host Bus Adapter Installation and User's Guide		~			~				
IBM TotalStorage FC2-133 Dual Port Host Bus Adapter Installation and User's Guide		-			~				
IBM Netfinity <sup>®</sup> Fibre Channel Cabling Instructions		~							
IBM Fibre Channel SAN Configuration Setup Guide	~		-	~	~				

#### Note:

1. The *IBM TotalStorage DS4000 Hardware Maintenance Manual* does not contain maintenance information for the IBM System Storage DS4100, DS4200, DS4300, DS4500, DS4700, DS4800, DS5100, or DS5300 storage subsystems. You can find maintenance information for these products in the *IBM System Storage DSx000 Storage Subsystem Installation, User's, and Maintenance Guide* for the particular subsystem.

# DS5100 and DS5300 storage subsystem documents

The following table lists the DS5100 and DS5300 storage subsystem documents and the common user tasks addressed by each document.

Title	User Tasks							
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance		
IBM System Storage DS5100 and DS5300 Storage Subsystem Installation, User's and Maintenance Guide		~						

Title	User Tasks								
	Planning	Hardware Installation	Software Installation	Configuration	Operation and Administration	Diagnosis and Maintenance			
IBM System Storage Quick Start Guide, Quick Reference for DS5100 and DS5300 Storage Subsystems, and for the EXP5060 Storage Expansion Enclosure		~	4	4					
IBM System Storage DS5000 EXP5060 Storage Expansion Enclosure Installation, User's, and Maintenance Guide									
Installing or replacing a DS5000 Cache and Flash Memory Card	1	~							
Installing or replacing a DS5000 Host Interface Card	1								
### Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan, Ltd. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATIONS "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publications. IBM may make improvements or changes (or both) in the product(s) or program(s) (or both), described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Some software may differ from its retail version (if available), and may not include user manuals or all program functionality.

#### Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (<sup>®</sup> or <sup>™</sup>), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at http://www.ibm.com/ legal/copytrade.shtml.

Adobe and PostScript are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States and/or other countries.

Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc., in the United States, other countries, or both and is used under license therefrom.

Intel, Intel Xeon, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc., in the United States, other countries, or both.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

#### Important notes

Processor speed indicates the internal clock speed of the microprocessor; other factors also affect application performance.

CD or DVD drive speed is the variable read rate. Actual speeds vary and are often less than the possible maximum.

When referring to processor storage, real and virtual storage, or channel logical drive, KB stands for 1024 bytes, MB stands for 1 048 576 bytes, and GB stands for 1 073 741 824 bytes.

When referring to hard disk drive capacity or communications logical drive, MB stands for 1 000 000 bytes, and GB stands for 1 000 000 bytes. Total user-accessible capacity can vary depending on operating environments.

Maximum internal hard disk drive capacities assume the replacement of any standard hard disk drives and population of all hard disk drive bays with the largest currently supported drives that are available from IBM.

Maximum memory might require replacement of the standard memory with an optional memory module.

IBM makes no representation or warranties regarding non-IBM products and services that are ServerProven<sup>®</sup>, including but not limited to the implied warranties of merchantability and fitness for a particular purpose. These products are offered and warranted solely by third parties.

IBM makes no representations or warranties with respect to non-IBM products. Support (if any) for the non-IBM products is provided by the third party, not IBM.

Some software might differ from its retail version (if available) and might not include user manuals or all program functionality.

#### **Documentation format**

The publications for this product are in Adobe Portable Document Format (PDF) and should be compliant with accessibility standards. If you experience difficulties when you use the PDF files and want to request a Web-based format or accessible PDF document for a publication, direct your mail to the following address:

Information Development IBM Corporation 205/A015 3039 E. Cornwallis Road P.O. Box 12195 Research Triangle Park, North Carolina 27709-2195 U.S.A.

In the request, be sure to include the publication part number and title.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

### Index

### Α

Activate Host Port 65 Activate iSCSI Initiator 65 Add Member to Consistency Group 66 address of the IBM director of licensing 667 AIX host

support xvi

### С

Clear Storage Subsystem Diagnostic Data Capture 81 Clear Storage Subsystem Recovery F Mode 82 Configure Automatic Support Bundle Collection 83 Convert FlashCopy Logical Drive to Enhanced FlashCopy Group 83 Create Consistency Group 88 Create Consistency Group Enhanced FlashCopy Image 92 Create Consistency Group Enhanced I FlashCopy Logical Drive 93 Create Consistency Group Enhanced FlashCopy Logical Drive Mapping 91, 95 Create Disk Pool 96 Create Enhanced FlashCopy Group 101 Create Enhanced FlashCopy Image 107 Create Enhanced FlashCopy Logical Drive 108 Create Logical Drive on Disk Pool 131 Create Read-Only Enhanced FlashCopy

Logical Drive 134

## D

- dcs xiii dcs3700 xiii dcs3860 xiii Delete Consistency Group 144 Delete Consistency Group Enhanced FlashCopy Image 145 Delete Consistency Group Enhanced FlashCopy Logical Drive 145 Delete Disk Pool 147 Delete Enhanced FlashCopy Group 147 Delete Enhanced FlashCopy Image 148 Delete Enhanced FlashCopy Logical Drive 150 Delete FlashCopy Logical Drive 151 Delete Logical Drive on a Disk Pool 154 device drivers downloading latest versions xv Disable FlashCopy 159
- Display Automatic Support Bundle Collection Configuration 162

**Display Support Bundle Collection** Schedule 162

documentation websites xv documentation, Storage Subsystem, Storage Subsystem library 663 drivers xv DS3000 software/firmware support xiii DS4000 software/firmware support xiii DS5000 software/firmware support xiii

### Ε

Enable Disk Pool Security 169

fire suppression xvii

### н

hardware service and support xvii

#### IBM

director of licensing address 667 Increase Logical Drive Capacity on a Disk Pool 175 Initialize Thin Logical Drive 176 Intel and AMD-based host support xvi

#### L

Linux host support xvi

### Ν

notes, important 668 notices general 667

products, developed 667

### R

Remove Member Logical Drive from Consistency Group 192 Rename Enhanced FlashCopy Logical Drive 193 resources websites xv Resume Consistency Group Enhanced FlashCopy Logical Drive 202

Resume Enhanced FlashCopy Image Rollback 203 Resume Enhanced FlashCopy Logical Drive 204 Resume FlashCopy Rollback 206 Revive Enhanced FlashCopy Group 208 Revive Enhanced FlashCopy Logical Drive 209

### S

Save Storage Subsystem Diagnostic Data Capture 217 Schedule Support Bundle Collection Configuration 223 services offered in the U.S.A. 667 Set Consistency Group Attributes 230 Set Consistency Group Enhanced FlashCopy Logical Drive 231 Set Disk Pool 236 Set Disk Pool (Modify Disk Pool) 238 Set Enhanced FlashCopy Group Attributes 246 Set Enhanced FlashCopy Group Media Scan 247 Set Enhanced FlashCopy Group Repository Logical Drive Capacity 253 Set Enhanced FlashCopy Group Schedule 248 Set Enhanced FlashCopy Logical Drive Media Scan 251 Set Enhanced FlashCopy Logical Drive Repository Logical Drive Capacity 252 Set Logical Drive Attributes 266 Set Logical Drive Attributes for a Disk Pool 271 Set Logical Drive Mapping 276 Set read-only Enhanced FlashCopy Logical Drive to a Read/Write Logical Drive 278 Set Thin Logical Drive Attributes 291 Show Consistency Group Enhanced FlashCopy Image 298 Show Disk Pool 302 Show Enhanced FlashCopy Group 306 Show Enhanced FlashCopy Image 307 Show Enhanced FlashCopy Logical Drives 309 Show Storage Subsystem DBM Database 323 Show Storage Subsystem Diagnostic Data Capture 324 Show Storage Subsystem Features 324 Show Storage Subsystem Profile 325 Show Thin Logical Drive 328 Start Consistency Group Enhanced FlashCopy Rollback 337 Start Disk Pool Locate 337 Start Enhanced FlashCopy Image Rollback 345 Start FlashCopy Rollback 346

Start Storage Subsystem Configuration Database Diagnostic 353 Start Storage Subsystem Diagnostic Data Capture 355 Start Storage Subsystem iSNS Server Refresh 355 Stop Consistency Group Enhanced FlashCopy Logical Drive 358 Stop Consistency Group Enhanced FlashCopy Rollback 357 Stop Disk Pool Locate 359 Stop Enhanced FlashCopy Group Pending Enhanced FlashCopy Images 360 Stop Enhanced FlashCopy Image 361 Stop Enhanced FlashCopy Image Rollback 361 Stop Enhanced FlashCopy Logical Drive 362 Stop FlashCopy Rollback 364 Stop Pending Enhanced FlashCopy Images on Consistency Group 367 storage area network (SAN) technical support website xvi Storage Manager software where to obtain xv switch technical support website xvi System p host support xvi System Storage Interoperation Center (SSIC) xv System x host support xvi

### T

trademarks 668

### W

Web sites AIX xvi SAN support xvi websites Fix central xvi IBM publications center xvi IBM System Storage product information xvi interoperability matrix xvi list xv premium feature activation xvi SSIC xv switch support xvi System p xvi System p xvi

# 

Part Number: 00Y8828

Printed in USA



(1P) P/N: 00Y8828

GA32-0961-06

