



User's Guide



User's Guide

Note

Before using this information and the product it supports, read the information in "Notices" on page 337.

July 2009

This edition applies to version 4, release 2, modification 0, Interim Feature 1 of IBM Tivoli OMEGAMON XE for Storage on z/OS (product number 5698-A37) and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright International Business Machines Corporation 2000, 2009.

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

Contents

| | | | |
|--|-------------|--|-----------|
| Tables | ix | Chapter 7. Monitoring logical control units. | 37 |
| Figures | xi | Navigating the workspaces of the Logical Control Unit node | 37 |
| About this guide | xiii | Logical Control Unit workspace | 39 |
| Chapter 1. Overview of Tivoli OMEGAMON XE for Storage on z/OS | 1 | Summary information for the workspaces in this node. | 40 |
| About Tivoli OMEGAMON XE for Storage on z/OS | 1 | Chapter 8. Monitoring tape groups. | 43 |
| Overview of attributes | 1 | Navigating the workspaces of the Tape Group node | 43 |
| Overview of product workspaces | 2 | Tape Group workspace | 43 |
| Creating Take Action commands. | 4 | Summary information for the workspaces in this node. | 44 |
| New in this release 4.2.0 | 4 | Chapter 9. Monitoring Virtual Tape Subsystems (VTS). | 45 |
| Chapter 2. Workspaces | 7 | Navigating the workspaces of the Virtual Tape Subsystems node | 45 |
| Nodes in the Navigator, which are linked to workspaces | 7 | Virtual Tape Subsystems workspace | 45 |
| Accessing workspaces | 8 | Summary information for the workspaces in this node. | 46 |
| Standard path and method for navigating workspaces | 9 | Chapter 10. Monitoring SMS storage groups performance | 51 |
| Integration with other products | 9 | Navigating the workspaces of the SMS Storage Groups Performance node | 51 |
| About dynamic linking to cross-product workspaces | 11 | SMS Storage Groups Performance workspace | 53 |
| Chapter 3. Monitoring applications. | 13 | Summary information for the workspaces in this node. | 54 |
| Navigating the workspaces of the Application Summary node | 13 | Chapter 11. Monitoring SMS storage groups space | 59 |
| Application Summary workspace | 13 | Navigating the workspaces of the SMS Storage Groups Space node | 59 |
| Summary information for the workspaces in this node. | 15 | SMS Storage Groups Space workspace | 60 |
| Chapter 4. Monitoring channel paths | 17 | Summary information for the workspaces in this node. | 60 |
| Navigating the workspaces of the Channel Path node. | 17 | Chapter 12. Monitoring user DASD groups performance | 65 |
| Channel Path workspace | 17 | Navigating the workspaces of the User DASD Groups Performance node | 65 |
| Summary information for the workspaces in this node. | 18 | User DASD Groups Performance workspace | 67 |
| Chapter 5. Monitoring cache CU performance | 21 | Summary information for the workspaces in this node. | 67 |
| Navigating the workspaces of the Cache CU Performance node | 21 | Chapter 13. Monitoring user DASD groups space | 69 |
| Cache CU Performance workspace | 23 | Navigating the workspaces of the User DASD Groups Space node | 69 |
| Summary information for the workspaces in this node. | 24 | User DASD Groups Space workspace. | 69 |
| Chapter 6. Monitoring cache CU status | 33 | Summary information for the workspaces in this node. | 70 |
| Navigating the workspaces of the Cache CU Status node. | 33 | | |
| Cache CU Status workspace. | 34 | | |
| Summary information for the workspaces in this node. | 35 | | |

Chapter 14. Monitoring DFSMSHsm status 73

Navigating the workspaces of the DFSMSHsm Status node. 73
DFSMSHsm Status workspace 73
Summary information for the workspaces in this node. 74

Chapter 15. Monitoring DFSMSrmm status 77

Navigating the workspaces of the DFSMSrmm Status node 77
DFSMSrmm System Summary workspace 79
Summary information for the workspaces in this node. 81
DFSMSrmm Dataset Details workspaces. 85

Chapter 16. Monitoring Dataset Attributes System Summary data . . . 87

Navigating the workspaces of the Dataset Attributes System Summary node 87
Dataset Attributes System Summary workspace . . . 88
Summary information for the workspaces in this node. 89

Chapter 17. Monitoring dataset groups 93

Navigating the workspaces of the Dataset Group Summary node 93
Dataset Group Summary workspace 93
Summary information for the workspaces in this node. 94

Chapter 18. Monitoring SMS configuration 97

Navigating the workspaces of the SMS Configuration node. 97
SMS Configuration workspace 97
Summary information for the workspaces in this node. 98

Chapter 19. Monitoring Storage Toolkit action requests 101

Navigating the workspaces of the Storage Toolkit node 101
Storage Toolkit Action Requests workspace . . . 101
Summary information for the workspaces in this node 102

Chapter 20. Directly monitoring situations 103

Direct situation analysis for the Application Summary workspace 103
Direct situation analysis for the Channel Path workspace 103
Direct situation analysis for the Cache CU Performance workspace 104
Direct situation analysis for the Logical Control Unit workspace 107

Direct situation analysis for the SMS Storage Groups Performance workspace 108
Direct situation analysis for the SMS Storage Groups Space workspace 109
Direct situation analysis for the HSM Status workspace 110

Chapter 21. Situations. 111

Overview of situations 111
Predefined situations descriptions 112
 KS3_Applic_Resp_Time_Critical 112
 KS3_Applic_Resp_Time_Warning 112
 KS3_Cachecu_Cache_Stat_Critical 112
 KS3_Cachecu_DFW_Retry_Critical 112
 KS3_Cachecu_DFW_Retry_Warning 112
 KS3_Cachecu_Inact_Vols_Critical 112
 KS3_Cachecu_Inact_Vols_Warning 113
 KS3_Cachecu_NVS_Stat_Critical 113
 KS3_Cachecu_Read_HitP_Critical. 113
 KS3_Cachecu_Read_HitP_Warning 113
 KS3_Cachecu_Trk_Dstg_Critical 113
 KS3_Cachecu_Trk_Dstg_Warning 113
 KS3_Cachecu_Trk_Stag_Critical 113
 KS3_Cachecu_Trk_Stag_Warning 114
 KS3_Cachecu_Write_HitP_Critical 114
 KS3_Cachecu_Write_HitP_Warning 114
 KS3_Channel_Busy_Pct_Critical 114
 KS3_Channel_Busy_Pct_Warning 114
 KS3_HSM_Backup_Held_Critical 114
 KS3_HSM_Backup_Queue_Critical 114
 KS3_HSM_Backup_Queue_Warning 115
 KS3_HSM_Dump_Held_Critical 115
 KS3_HSM_Dump_Queue_Critical. 115
 KS3_HSM_Dump_Queue_Warning 115
 KS3_HSM_Inactive_Host_Warning 115
 KS3_HSM_Migrate_Held_Critical 115
 KS3_HSM_Migrate_Queue_Critical 116
 KS3_HSM_Migrate_Queue_Warning. 116
 KS3_HSM_Recall_Held_Critical 116
 KS3_HSM_Recall_Queue_Critical 116
 KS3_HSM_Recall_Queue_Warning 116
 KS3_HSM_Recovery_Held_Critical 116
 KS3_HSM_Recovery_Queue_Critical. 117
 KS3_HSM_Recovery_Queue_Warning 117
 KS3_HSM_Status_Inactive_Crit 117
 KS3_LCU_Av_Delay_Q_Critical 117
 KS3_LCU_Av_Delay_Q_Warning 117
 KS3_LCU_Cont_Rate_Critical 117
 KS3_LCU_Cont_Rate_Warning. 118
 KS3_LCU_IO_Rate_Sec_Critical 118
 KS3_LCU_IO_Rate_Sec_Warning 118
 KS3_RMM_CDS_Backup_Critical 118
 KS3_RMM_CDS_Backup_Warning 118
 KS3_RMM_CDS_Space_Critical 118
 KS3_RMM_CDS_Space_Warning 118
 KS3_RMM_Exit_Status_Critical 118
 KS3_RMM_Journal_Status_Critical 119
 KS3_RMM_Operating_Mode_Warning 119
 KS3_RMM_Scratch_Tape_Critical 119
 KS3_RMM_Scratch_Tape_Warning 119
 KS3_RMM_Inactive_Critical 119

| | | | |
|--|------------|--|-----|
| KS3_Stg_Toolkit_Result_Critical | 119 | Dataset Attributes SMS Class Detail | 164 |
| KS3_Stg_Toolkit_Result_Warning | 119 | Dataset Attributes SMS Class Summary | 165 |
| KS3_Storage_Gr_Pct_Free_Crit. | 119 | Dataset Attributes Space Allocated Summary | 166 |
| KS3_Storage_Gr_Pct_Free_Warning | 120 | Dataset Attributes Space Unused Summary | 166 |
| KS3_Vol_Cache_DFW_Retry_Critical. | 120 | Dataset Attributes System Summary. | 167 |
| KS3_Vol_Cache_DFW_Retry_Warning | 120 | Dataset Attributes Uncataloged Detail | 169 |
| KS3_Vol_Cache_Read_HitP_Critical | 120 | Dataset Attributes Uncataloged Summary | 170 |
| KS3_Vol_Cache_Read_HitP_Warning | 120 | Dataset Group Defs | 170 |
| KS3_Vol_Cache_Writ_HitP_Critical | 121 | Dataset Group Details | 170 |
| KS3_Vol_Cache_Writ_HitP_Warning | 121 | Dataset Group Detail Attrs | 174 |
| KS3_Vol_Disabled_VTOC_Critical | 121 | Dataset Group Details Volume. | 175 |
| KS3_Vol_EAV_Fragment_Index_Crit. | 121 | Dataset Group Summary | 177 |
| KS3_Vol_EAV_Fragment_Index_Warn | 121 | Dataset Performance | 178 |
| KS3_Vol_EAV_Free_Space_Pct_Crit | 121 | Dataset SMS. | 180 |
| KS3_Vol_EAV_Free_Space_Pct_Warn. | 121 | Dataset Space | 180 |
| KS3_Vol_Fragment_Index_Critical | 122 | Dataset Users | 182 |
| KS3_Vol_Fragment_Index_Warning | 122 | HSM CDS | 182 |
| KS3_Vol_Free_Space_Pct_Critical | 122 | HSM Common Storage | 183 |
| KS3_Vol_Free_Space_Pct_Warning | 122 | HSM Dataset Recover Version | 184 |
| KS3_Vol_Perf_Resp_Time_Critical | 122 | HSM Function Details | 184 |
| KS3_Vol_Perf_Resp_Time_Warning | 122 | HSM Function Statistics | 185 |
| KS3_VTS_Disconnect_Time_Crit | 122 | HSM Function Summary | 187 |
| KS3_VTS_Host_GB_Warning | 123 | HSM Host Function Status | 187 |
| KS3_VTS_Pct_Copy_Throt_Warn | 123 | HSM Host Request Summary | 189 |
| KS3_VTS_Pct_Wr_Over_Throt_Warn. | 123 | HSM Host Status | 189 |
| KS3_VTS_Recall_Pct_Warning | 123 | HSM Private Storage | 191 |
| KS3_VTS_Virt_MtPend_Av_Warning. | 123 | HSM Requests | 191 |
| KS3_VTS_Virt_MtPend_Mx_Warning | 123 | HSM Status | 192 |
| | | Logical Control Unit | 194 |
| | | Logical Control Unit Channel Paths | 195 |
| | | LCU Control Unit | 196 |
| | | RMM Configuration | 197 |
| | | RMM Control | 200 |
| | | RMM Control Data Set | 200 |
| | | RMM Dataset Details. | 201 |
| | | RMM Location Details | 203 |
| | | RMM Owner Details | 204 |
| | | RMM Pending Actions | 204 |
| | | RMM System Summary | 205 |
| | | RMM Volume Details. | 208 |
| | | RMM VRS Details | 214 |
| | | SMS Cache Sets | 215 |
| | | SMS Configuration | 215 |
| | | SMS Data Class | 216 |
| | | SMS Management Class | 217 |
| | | SMS Storage Class. | 219 |
| | | SMS Storage Group | 220 |
| | | SMS Storage Group Status | 222 |
| | | SMS Systems | 222 |
| | | Storage Toolkit Requests. | 222 |
| | | Storage Toolkit Result Summary | 224 |
| | | Storage Toolkit Result Detail | 224 |
| | | Storage Toolkit Result Resources | 225 |
| | | Symmetrix Configuration | 226 |
| | | Symmetrix Devices | 226 |
| | | Symmetrix Director Summary | 227 |
| | | Symmetrix Directors | 227 |
| | | Symmetrix SSIDs | 229 |
| | | Tape Device | 229 |
| | | Tape Group | 230 |
| | | TotalStorageDS Array Configuration. | 231 |
| Chapter 22. Attributes. | 125 | | |
| List of attribute groups and attributes | 125 | | |
| Application Dataset Details. | 127 | | |
| Application Dataset Monitor | 128 | | |
| Application Monitoring | 129 | | |
| Application Volume Performance. | 130 | | |
| Application Volume Space | 132 | | |
| Cache Control Unit | 133 | | |
| Cache Devices | 134 | | |
| Cache Raid Rank | 138 | | |
| Channel Path | 138 | | |
| Cross System Volume Summary | 140 | | |
| Cross System Volumes | 141 | | |
| DASD Physical Group | 142 | | |
| DASD Volume Performance | 145 | | |
| DASD Volume Space | 147 | | |
| Dataset Attributes Blocksize Summary | 149 | | |
| Dataset Attributes CA Split Summary | 149 | | |
| Dataset Attributes Catalog Summary | 150 | | |
| Dataset Attributes CI Split Summary | 151 | | |
| Dataset Attributes Creation Date Summary | 152 | | |
| Dataset Attributes Defs | 152 | | |
| Dataset Attributes Dataset Detail | 153 | | |
| Dataset Attributes DSORG Detail. | 157 | | |
| Dataset Attributes DSORG Summary | 158 | | |
| Dataset Attributes Extent Summary | 159 | | |
| Dataset Attributes Inefficient Blocksize Summary | 160 | | |
| Dataset Attributes Mature Dataset Summary | 161 | | |
| Dataset Attributes Never Referenced Dataset Summary. | 162 | | |
| Dataset Attributes Newborn Dataset Summary | 163 | | |
| Dataset Attributes Reference Date Summary | 163 | | |

| | |
|---|-----|
| TotalStorageDS Configuration | 232 |
| TotalStorageDS Extent Pool | 232 |
| TotalStorageDS Rank | 234 |
| TotalStorageDS SSIDs. | 234 |
| Volume Group Summary | 235 |
| Volume Status | 238 |
| Volume Users | 238 |
| VTS Cache | 238 |
| VTS Capacity | 239 |
| VTS Composite Library | 241 |
| VTS Overview | 242 |
| VTS Physical Devices. | 243 |
| VTS TS7700-series Cache Container | 244 |
| VTS TS7700-series Cache Partition Container | 245 |
| VTS TS7700-series Cache Preference Group | 246 |
| VTS TS7700-series Cluster | 247 |
| VTS Virtual Devices | 248 |
| VTS Virtual Devices Summary. | 250 |

Chapter 23. Storage Toolkit 251

| | |
|---|-----|
| Storage Toolkit menus and dialog boxes | 251 |
| How the Storage Toolkit processes action requests | 253 |
| Using the Storage Toolkit | 254 |
| Standard tabs in the dialog boxes of the Storage Toolkit. | 258 |
| Options tab | 258 |
| General tab | 258 |
| Scheduling tab. | 259 |
| JCL tab | 259 |
| Command tab | 260 |
| Authentication and authorization for running Storage Toolkit commands | 261 |
| Data View dialog box | 261 |
| Using substitution variables in the Storage Toolkit. | 262 |
| Custom actions in the Storage Toolkit | 265 |
| Create Batch Job dialog box | 265 |
| Issue Command dialog box. | 268 |
| Submit Command or Job dialog box. | 269 |
| Volume actions in the Storage Toolkit | 270 |
| Volume Backup dialog box | 271 |
| Volume Conversion dialog box | 271 |
| Volume Migration dialog box | 271 |
| Volume Space Release dialog box. | 271 |
| Compress dialog box | 272 |
| Defrag dialog box | 272 |
| Data set actions in the Storage Toolkit | 273 |
| Dataset Backup dialog box | 273 |
| Dataset Migration dialog box | 273 |
| Dataset Copy and Move dialog box | 274 |
| Dataset Recall dialog box | 274 |
| Dataset Recover dialog box. | 274 |
| Dataset Space Release dialog box. | 274 |
| Compress dialog box | 275 |
| Dataset Allocate dialog box. | 275 |
| Dataset Alter dialog box. | 276 |
| Dataset Define Cluster dialog box | 277 |
| Dataset Catalog dialog box | 280 |
| Dataset Delete dialog box | 281 |
| Dataset Listcat dialog box | 281 |
| Print dialog box | 283 |

| | |
|--|-----|
| Dataset Rename dialog box. | 284 |
| Dataset Repro dialog box | 284 |
| Dataset Uncatalog dialog box | 286 |
| Dataset Verify dialog box | 286 |
| DFSMSHsm actions in the Storage Toolkit | 286 |
| Cancel HSM Request dialog box | 286 |
| DFSMSrmm actions in the Storage Toolkit. | 287 |
| Add Dataset dialog box | 288 |
| Change Dataset dialog box | 288 |
| Delete Dataset dialog box | 288 |
| Add Owner dialog box | 289 |
| Change Owner dialog box | 289 |
| Delete Owner dialog box | 289 |
| Add Volume dialog box | 289 |
| Change Volume dialog box. | 290 |
| Confirm Actions for All Volumes dialog box | 290 |
| Delete Volume dialog box | 290 |
| Add VRS dialog box | 290 |
| Change VRS dialog box | 291 |
| Delete VRS dialog box | 291 |
| Overview of common action commands | 291 |
| Managing actions | 292 |
| Storage Toolkit limitations and hints. | 292 |

Chapter 24. Dataset Group feature 297

| | |
|---|-----|
| The Dataset Group feature versus the Dataset Attributes feature | 297 |
| Working with data set groups | 297 |
| Dataset Group Properties dialog box | 297 |
| Dataset Group Data Collection dialog box. | 298 |
| Adding a data set group | 298 |
| Editing and deleting data set groups | 298 |
| Setting the collection interval | 298 |
| Data set group masking rules | 299 |
| Using data set groups | 300 |
| Obtaining an overview of space usage | 300 |
| Identifying data sets with excess space | 301 |
| Create an alert when response time exceeds specific thresholds. | 302 |
| Identify VSAM data sets with excessive CA/CI splits | 304 |
| Cataloged and not-allocated data sets | 306 |
| Notify the Production Scheduling group when a critical data set exceeds a specified number of extents | 308 |
| Automatically migrate infrequently used data sets. | 308 |

Chapter 25. Dataset Attributes Database 311

| | |
|---|-----|
| Benefits | 311 |
| Optimization | 311 |
| Attribute derivation and definition | 311 |
| Use cases. | 313 |
| Configuring the Dataset Attribute Database feature | 313 |
| Optimizing data collection for the Dataset Attributes Database feature. | 314 |
| Command-line option | 315 |
| Custom queries. | 315 |

| | |
|---|------------|
| Chapter 26. Launching of IBM TotalStorage Productivity Center . . . | 317 |
| About the IBM TotalStorage Productivity Center | 317 |
| About the launch feature in Tivoli OMEGAMON XE for Storage on z/OS | 317 |
| Launching the IBM TotalStorage Productivity Center. | 317 |
| Workspaces that provide the launch feature . . . | 318 |
| Configuring the launch feature | 318 |
| | |
| Chapter 27. Reports available for IBM Tivoli OMEGAMON XE for Storage on z/OS | 321 |
| Using these reports | 321 |
| Typical types of reports available with Tivoli Common Reporting | 322 |
| OMEGAMON XE for Storage on z/OS reports . . . | 322 |
| | |
| Chapter 28. Working with the DFSMSrmm Status feature. | 323 |
| Benefits | 323 |
| Optimization | 323 |
| Attribute derivation and definition | 324 |
| Scenarios for using the DFSMSrmm feature . . . | 324 |
| Optimizing data collection for the DFSMSrmm Status feature | 325 |
| Command-line option | 325 |

| | |
|--|-----|
| Custom queries. | 326 |
| RMM Collection Control dialog box | 326 |
| Details regarding the Collection Control dialog box. | 326 |

| | |
|---|------------|
| Appendix A. Documentation library | 329 |
| Tivoli OMEGAMON XE for Storage on z/OS library. | 329 |
| Shared OMEGAMON XE publications | 329 |
| IBM Tivoli Monitoring library | 330 |
| Related publications | 331 |
| Other sources of documentation | 331 |

| | |
|---|------------|
| Appendix B. Support information . . . | 333 |
| Using IBM Support Assistant | 333 |
| Obtaining fixes | 333 |
| Receiving weekly support updates | 334 |
| Contacting IBM Software Support | 335 |
| Determining the business impact | 335 |
| Describing problems and gathering information | 336 |
| Submitting problems | 336 |

| | |
|--------------------------|------------|
| Notices | 337 |
| Trademarks | 338 |

| | |
|------------------------|------------|
| Index | 339 |
|------------------------|------------|

Tables

| | | | |
|--|-----|---|-----|
| 1. Links to workspaces of related products | 10 | 19. Command reference guides | 258 |
| 2. Application Summary node: bulleted list showing the hierarchy of linked workspaces | 13 | 20. Sample settings for "Files whose contents should be copied for later viewing" | 266 |
| 3. Channel Path node: bulleted list showing the hierarchy of workspace links | 17 | 21. IDCAMS ALLOCATE options that are available in the Allocate dialog box of the Storage Toolkit | 275 |
| 4. Cache CU Performance node: bulleted list showing the hierarchy of linked workspaces | 21 | 22. IDCAMS DEFINE CLUSTER options that are available in the Dataset Define Cluster dialog box of the Storage Toolkit. Part 1 of 3, Cluster options | 277 |
| 5. Cache CU Status node: bulleted list showing the hierarchy of linked workspaces | 34 | 23. IDCAMS DEFINE CLUSTER options that are available in the Dataset Define Cluster dialog box of the Storage Toolkit. Part 2 of 3, (optional) Data component options | 278 |
| 6. Logical Control Unit node: bulleted list showing the hierarchy of linked workspaces | 38 | 24. IDCAMS DEFINE CLUSTER options that are available in the Dataset Define Cluster dialog box of the Storage Toolkit. Part 3 of 3, (optional) Index component options | 279 |
| 7. Tape Group node: bulleted list showing the hierarchy of linked workspaces | 43 | 25. IDCAMS LISTCAT options that are available in the Dataset Listcat dialog box of the Storage Toolkit | 282 |
| 8. Virtual Tape Subsystems node: bulleted list showing the hierarchy of linked workspaces | 45 | 26. IDCAMS REPRO options that are available in the Dataset Repro dialog box of the Storage Toolkit, Part 1 of 2, Repro tab | 284 |
| 9. SMS Storage Groups Performance node: bulleted list showing the hierarchy of linked workspaces | 52 | 27. IDCAMS REPRO options that are available in the Dataset Repro dialog box of the Storage Toolkit, Part 2 of 2, Merge/DBCS tab | 285 |
| 10. SMS Storage Groups Space node: bulleted list showing the hierarchy of linked workspaces | 59 | 28. Common commands that are enabled by the Storage Toolkit | 291 |
| 11. User DASD Groups Performance node: bulleted list showing the hierarchy of linked workspaces | 66 | 29. Mask characters | 299 |
| 12. User DASD Groups Space node: bulleted list showing the hierarchy of linked workspaces | 69 | 30. Sample masks | 299 |
| 13. DFSMSHsm Status node: bulleted list showing the hierarchy of linked workspaces | 73 | 31. Using the Tivoli Enterprise Portal to manage data sets | 313 |
| 14. DFSMSrmm Status node: bulleted list showing the hierarchy of linked workspaces | 78 | 32. Dataset Attribute Data Collection dialog box | 313 |
| 15. Dataset Attributes System Summary node: bulleted list showing the hierarchy of workspace links | 88 | 33. Using the Tivoli Enterprise Portal to manage DFSMSrmm tape volumes and data sets | 325 |
| 16. Dataset Group Summary node: bulleted list showing the hierarchy of linked workspaces | 93 | 34. RMM Collection Control dialog box | 326 |
| 17. SMS Configuration node: bulleted list showing the hierarchy of linked workspaces | 97 | | |
| 18. Storage Toolkit Action Requests node: bulleted list showing the hierarchy of linked workspaces | 101 | | |

Figures

1. Variable elements for JCL files used with the Storage Toolkit 267
2. Dataset Group Summary workspace showing four new data set groups 301
3. Properties panel of Dataset Group Summary workspace showing a new threshold. . . . 302
4. Creation of an alert in the Dataset Group Summary workspace 304
5. Creating a situation for a data set group for VSAM attributes workspace. 306
6. Creating a situation in the Dataset Group Details workspace to monitor cataloged and non-cataloged data sets 307
7. Partial view of a formula that monitors the migration of infrequently used data sets . . 310

About this guide

IBM Tivoli OMEGAMON XE for Storage on z/OS: User's Guide introduces the features, workspaces, predefined situations, and attributes, for the IBM® Tivoli® OMEGAMON® XE for Storage on z/OS® product and supplements the user assistance provided with this product. For a list of what is new in this release of the product, see “New in this release 4.2.0” on page 4.

Users of this book need to be familiar with performance monitoring concepts. If you use IBM Tivoli Data Warehouse, you need to be familiar with the operating system that hosts the warehouse.

The document assumes no previous experience with Tivoli OMEGAMON products or with IBM Tivoli Monitoring. To learn about this family of products: <http://www-306.ibm.com/software/tivoli/solutions/availability/products.html>.

Chapter 1. Overview of Tivoli OMEGAMON XE for Storage on z/OS

This section introduces the features and functionality of Tivoli OMEGAMON XE for Storage on z/OS.

About Tivoli OMEGAMON XE for Storage on z/OS

Tivoli OMEGAMON XE for Storage on z/OS is the comprehensive monitor for z/OS I/O subsystem performance and storage availability. The product combines comprehensive storage performance monitoring with a flexible, easy-to-use browser interface that helps you clearly understand storage conditions and ensure optimal performance. Specifically, you use Tivoli OMEGAMON XE for Storage on z/OS to manage the performance and availability of mainframe attached storage (including disk, tape devices, and virtual tape devices) and the data sets that are located on them. With version 4.2.0 of the product, support for virtual tape subsystems includes integration and compatibility with Generation 2 of the VTS (Virtualization Engine™ device model TS7700) and its new architecture for collecting and presenting performance and availability metrics.

The product also features in-depth analysis of the following three key components of IBM storage software:

- Data Facility Systems Managed Storage (DFSMS™), which manages the service levels and priorities of data sets, based on user created storage goals.
- Data Facility Hierarchical Storage Manager (DFSMSHsm™), which provides functions for backing up and recovering data, and managing space on volumes in the storage hierarchy.
- Data Facility Removable Media Manager (DFSMSrmm™), which is a functional component of, and shipped as part of, DFSMS. DFSMSrmm manages removable media resources for the z/OS environment, including automatic libraries like the IBM Virtual Tape Server.

The features of Tivoli OMEGAMON XE for Storage on z/OS include the ability to perform trending and other historical analysis, and the ability to monitor user-created groups of volumes.

Overview of attributes

Tivoli OMEGAMON XE for Storage on z/OS gathers data from remote agents located on the managed systems of your network and stores this data in system elements called attributes. You can use these attributes to build situations to monitor the performance of the managed systems you are concerned with.

Attributes correspond to the column names contained on the table views in associated Tivoli OMEGAMON XE for Storage on z/OS workspaces. Related attributes are grouped into attribute groups. A given table view contains information provided by a single attribute group.

To use attributes effectively, you must first understand the structure of an attribute.

- An attribute is made up of an attribute group and an attribute item.
- An attribute group associates attributes that refer to some unique characteristic of the data that is being stored.
- An attribute item stores data for a particular property of an attribute group.

For example, in the attribute

```
TOTALSTORAGE_RANKS.Total_Bytes_Read
```

Total Bytes Read is an attribute item that stores a value representing the total number of bytes read from the rank over the interval in the following attribute group: “TotalStorageDS Rank” on page 234.

You use attributes to create situations that monitor the state of your storage environment. A situation describes a condition you want to test. When you start a situation, the monitoring server compares the values that you have assigned for the situation's attributes with the values collected by Tivoli OMEGAMON XE for Storage on z/OS and registers a situation event if the condition is met. You are alerted to situation events by indicator icons that appear in the Navigator.

Attributes are listed and described in the following documentation:

- Online help for Tivoli OMEGAMON XE for Storage on z/OS
- *IBM Tivoli OMEGAMON XE for Storage on z/OS: User's Guide*

Overview of product workspaces

When using Tivoli Enterprise Portal, information is displayed in workspaces. Within a given workspace, information is displayed in tabular form. Tivoli Enterprise Portal refers to this tabular format for information as a table view. Information can also be displayed in the workspace as a chart, graph, or other format you can specify.

About workspaces

A workspace is the working area of the Tivoli Enterprise Portal application window. At the left of the workspace is a Navigator that permits you to select the workspace you want to display. As part of the application window, the right side of the status bar shows the Tivoli Enterprise Portal server name to which the displayed information applies, as well as the ID of the current user. Some key points about workspaces are as follows:

- As you select items in the Navigator, the workspace presents views pertinent to your selection. Each workspace has at least one view.
- Every workspace has a set of properties associated with it. You can customize the workspace by working in the Properties editor to change the style and content of each view.
- Another way to customize the workspace is to change the type of view or to add views to the workspace.
- Always save your changes before moving to a new workspace. Otherwise, the changes you make to the workspace are lost.
- Some workspaces, such as those dedicated to tracking the metrics of a DS8000® storage device, are not available in selection menus and lists when DS8000 storage devices do not exist in the monitoring environment.

Formats for information

Information can be presented to you in any of the following views:

- Table view
- Pie chart view
- Bar chart view
- Plot chart view
- Needle gauge view
- Thermometer gauge view
- Notepad view
- Situation Event console view, which shows the status of the situations associated with the system.
- Take Action view, which is used to send a command to the system.
- Terminal view, which enables you to start a 3270 or 5250 work session.
- Browser view, which permits you to open a browser to see HTML pages and Web sites.

Use the Export function of the Tivoli Enterprise Portal to store product data as text or as comma separated values (CSV format). For example, you can right-click on a table view and select **Export** to export the data in the view.

Associating workspaces with attributes

There is a direct relationship between attributes and workspaces. An attribute group typically corresponds to a table view within a named workspace and attribute items correspond to columns in the table view.

Each Tivoli Enterprise Portal workspace displays real-time information or historical data for many of the attributes. The information is available to you, independent of whether you are using Tivoli OMEGAMON XE for Storage on z/OS to monitor situations.

For descriptions of the individual workspaces, see Chapter 2, “Workspaces,” on page 7.

Using Tivoli OMEGAMON XE for Storage on z/OS information

You can view information about each managed system that you are monitoring. Use this information to perform the following tasks:

- Monitor the performance of each managed system, helping you to identify bottlenecks and evaluate tuning decisions
- Select the most effective threshold values for situations you create
- Review status information when a change in the state of a given resource occurs, such as from OK to Warning or Critical

Defining workspace properties

Every workspace has a set of properties associated with it. You can customize the workspace by working in the Properties editor to change the style and content of each view. Changes you make to workspace properties, such as adding or editing a view are only temporary. They are lost when you exit Tivoli Enterprise Portal unless you save the workspace.

The properties of a workspace can include some or all of the following:

Query Specify what data must be specified in the chart or table

Filters Refine the view by filtering out unwanted data from the chart or table

Thresholds

Establish threshold values and color indicators for a table view

Configuration

Specify the script to run or the connection to make whenever you open the terminal view

Style Change the behavior and appearance of the view

Investigating a situation event

When the conditions of a situation have been met, the situation evaluates to True, causing a situation event indicator to be displayed in the Navigator. You can investigate the situation event by opening its workspace.

The situation event workspace shows two table views, one with the values of the attributes when the situation evaluated to True, and the other with the attributes' current values.

The situation event workspace can also display a text view with any expert advice written by the author of the situation, and the Take Action view so you can send a command to the application started on that system.

Filtering information

To manually set up filtering for a given table view, place the cursor on the table view, press the right mouse button, and select **Properties**. (The equivalent keystroke is **Ctrl+R**.) From the

displayed dialog box, select the **Filters** tab. Here you can select the columns to display as well as set up the criteria for which rows to display. To save your filtering specifications, you must save the workspace before exiting.

Sorting Information

Sorting is handled by simply clicking on a column heading. Click once and the report is sorted in ascending order. Click a second time to resort the report into descending order. A third click returns you to the report's default sort order.

Adding a workspace to your favorites

When using Tivoli Enterprise Portal in browser mode, you can start it from any workstation by entering the URL of the Tivoli Enterprise Portal server where the browser mode client is installed. Each Tivoli Enterprise Portal workspace also has a URL so that you can save the workspace to your Favorites list or specify it as your home page.

Creating Take Action commands

The Take Action feature lets you issue a command to any system in your network where one or more monitoring agents are installed. You can create a Take Action command and run it immediately from a workspace, or you can run a saved Take Action command.

The Tivoli Enterprise Portal is always hosted on a distributed computer system. When you submit a Take Action command from the Tivoli Enterprise Portal to a z/OS system, a return code of zero displayed in the portal interface. This code indicates successful submission of the command but gives no indication of the result of the command. You can find the command output in the z/OS SYSLOG.

Note: The Chapter 23, "Storage Toolkit," on page 251 provides functionality similar to the Take Action commands. You create action requests in the Storage Toolkit that issue DFSMShsm and DFSMSdss™ commands in target z/OS systems.

Here is an example of using a saved Take Action command:

1. Right-click an item in the Navigator.
2. From the pop-up menu, select **Take Action** to display the Take Action dialog box.
3. In the Take Action dialog box's Action area, select a saved Take Action command. To edit the command's argument values, select the Arguments button, click **OK**.
4. In the Take Action dialog box Destination Systems area, select the target system or systems to which you want to send the Take Action command and click **OK**. This causes the command to be sent to the system or systems where it is run.

New in this release 4.2.0

Tivoli OMEGAMON XE for Storage on z/OS now offers the following functionality:

New conventions and features in all Tivoli OMEGAMON XE products

- **Version numbering:** All components are identified by the version number in the Managed System Status table view of the Tivoli Enterprise Portal. For example, the Version column for rows in this table view might display the following entries:
 - 04.20.00 with an entry for a monitoring agent indicates that the version of this monitoring agent is 4.2.0.
 - 06.21.00 with an entry for a monitoring server indicates that the version of this monitoring server is IBM Tivoli Monitoring 6.2.1.

The new version numbering feature helps you to quickly identify the version of each monitoring agent in an environment that comprises multiple versions.

In addition, the version number appears in the description of queries in the Situation Editor. Multiple queries with the same name and description are required to support upgrade scenarios. The addition of a version number enables you to identify which query is appropriate for a particular agent at a particular version.

- **FMIDs merged:** Starting with Version 4.2.0, delivery and installation of the OMEGAMON II interface (sometimes called the CUA interface) for this OMEGAMON XE product has changed. The OMEGAMON II FMID has been merged with the OMEGAMON XE FMID. The OMEGAMON II FMID has been eliminated. This change should shorten the required installation time for these products and simplify the configuration and maintenance work that is required.

The OMEGAMON II manuals continue to be provided with the current version of Tivoli OMEGAMON XE for Storage on z/OS. No enhancements have been made to these documents since version 5.2.0 of the OMEGAMON II interface.

New conventions and features in Tivoli OMEGAMON XE for Storage on z/OS

- **Situation startup:** All product-provided situations are configured so that they do *not* run at startup. The autostart status of existing situations (situations from Version 410 or previous versions of the product that you are migrating from) is not affected.
- **DFSMSrmm monitoring:** New DFSMSrmm monitoring and analysis function helps you manage the critical DFSMSrmm subsystem. The DFSMSrmm Status workspace provides a high-level overview of DFSMSrmm statistics and links to related workspaces that track data for common aspects of DFSMSrmm management and administration.
- **Storage Toolkit enhancements:** New capabilities of the Storage Toolkit enable you to issue a broad range of common commands from the following facilities in a z/OS system:
 - DFSMSdss (Device Support Services)
 - DFSMSHsm (Hierarchical Storage Manager)
 - ICKDSF (Device Support Facilities)
 - IDCAMS (Access Method Services)
 - DFSMSrmm (DFSMS Removable Media Manager)
 - Mainframe (z/OS console command)
 - TSO (Time Sharing Option)

The actions target resources such as volumes on DASD and mainframe data sets. For example, you can use the actions to migrate data sets or release free space.

The Storage Toolkit now also enables you to issue custom commands and create batch jobs, as described in Chapter 23, “Storage Toolkit,” on page 251.

- **ICAT:** The Installation and Configuration Assistance Tool (ICAT) now enables you to configure the product as follows:
 - Configure parameters for collection of metrics related to TS7700 Virtual Tape servers
 - Set the following type of request parameters for the Storage Toolkit:
 - Size of Result data set
 - Size of JCL data set
 - Specifications for the jobcard for a generated JCL
- **TS7700-based VTS:** Product functionality for Virtual Tape Subsystems is now extended to include support for integration and compatibility with Generation 2 of the VTS (Virtualization Engine device model TS7700) and its new architecture for collecting and presenting performance and availability metrics.
- **Support for Extended Address Volume (EAV):** The product now enables you to monitor the EAV device model. You can measure whether the EAVs are being properly utilized and have been defined in the correct Storage Groups and measure track-managed space availability on EAVs.

- **Log and Trace Analyzer:** The product provides integration with the Log and Trace Analyzer, a tool that links messages that are displayed in the RKLVLLOG to information regarding workarounds and solutions that are available as APARs (Authorized Problem Analysis Reports). See the IBM Tivoli OMEGAMON XE for Storage on z/OS: Troubleshooting Guide for more information regarding the Log and Trace Analyzer.
- **Troubleshooting:** The name of the *Problem Determination Guide* has changed to *Troubleshooting Guide*.
- **New name for an existing node:** The name of the *HSM Status node* has changed to *DFSMSHsm Status node* to be consistent with general naming conventions. The names of the associated workspaces also begin with “DFSMSHsm” (not “HSM”).
- **New reporting capability through Tivoli Common Reporting:** Your Product now provides reports that run under Tivoli Common Reporting. See Chapter 27, “Reports available for IBM Tivoli OMEGAMON XE for Storage on z/OS,” on page 321.
- **Support for event forwarding:** If you are using IBM Tivoli Enterprise Console® (TEC) or IBM Tivoli Netcool®/OMNIBUS, in addition to IBM Tivoli Monitoring, to manage events in your enterprise, you can now forward events reported by Tivoli OMEGAMON XE for Storage on z/OS monitoring agents to these event management products. Refer to the following documents for further information:
 - See the *IBM Tivoli OMEGAMON XE for Storage on z/OS: Planning and Configuration Guide* for a discussion of issues related to event forwarding in Tivoli OMEGAMON XE for Storage on z/OS.
 - See the “Event integration scenarios” section of the *IBM Tivoli Monitoring: Installation and Setup Guide* for a discussion of the benefits of these products and the details of how they can be integrated with IBM Tivoli Monitoring.
 - See *IBM Tivoli Monitoring: Configuring the Tivoli Enterprise Monitoring Server on z/OS* to learn how to configure and optimize event forwarding.

New in release 4.2.0 Interim Feature 1.0

- The historical trending functionality for several attribute groups has been enhanced to support Tivoli Data Warehouse (TDW), summarization, and pruning; thereby extending historical capability for capacity planning and troubleshooting purposes. For a list of the enhanced attribute groups, a list of workspaces with the extended reporting capabilities, and information on how to define historical configuration collection intervals, see Chapter 7 in the *Planning and Configuration Guide*. Also, please review Table 7 in Appendix C in the *Planning and Configuration Guide* for details on changes to this enhancement.
- This interim feature provides storage group and user DASD group space metrics in cylinders, tracks, and gigabytes (in addition to megabytes), including the percent utilized space and the percent free space. Look for this data in the User DASD Groups Space, User DASD Groups Space Trend, SMS Storage Groups Space, and SMS Storage Groups Space Trend workspaces.
- This interim feature provides you with the performance and response time data you need to monitor the Solid State Device (SSD). This enhancement allows you to determine whether the device is functioning optimally and to identify SSD arrays, ranks, extent pools, and the logical volumes that reside on the SSD arrays.
- The Edit Data Set dialog box now displays the name of the data set being edited in the title of the dialog box. It also provides several visual cues that indicate your editing position, your editing mode, and the logical record length of the data set. See “Edit Data Set dialog box” on page 268 for more information.
- **Troubleshooting Guide:** New messages were added to the *Troubleshooting Guide*.

Chapter 2. Workspaces

The predefined workspaces in IBM Tivoli OMEGAMON XE for Storage on z/OS are described in this section.

Each workspace contains a table view that provides information you can use to monitor a particular resource in your system. Column headings in the table view correspond to attributes that you can use to create situations. Each of the attributes in a specific table view belongs to a single attribute group.

This document has a section for each of the nodes that you see in the Navigator. These nodes are listed in “Nodes in the Navigator, which are linked to workspaces.” Each node links to a workspace that has the same name. Each workspace has links to related workspaces. Each link in “Nodes in the Navigator, which are linked to workspaces” goes to a section that describes the corresponding workspace. The sections also describe all linked workspaces.

Nodes in the Navigator, which are linked to workspaces

Note: In the Navigator and in this document, nodes are listed in functional order, not alphabetically.

Enterprise

z/OS Systems

System_name

Storage Subsystem

SYSTEM:SUBSYSTEM:AGENT_NAME

Application Summary

Channel Path

Cache CU Performance

Cache CU Status

Logical Control Unit

Tape Group

Virtual Tape Subsystems

SMS Storage Groups Performance

SMS Storage Groups Space

User DASD Groups Performance

User DASD Groups Space

DFSMSHsm Status

DFSMSrmm Status

Dataset Attributes System Summary

Dataset Group Summary

SMS Configuration

Storage Toolkit Action Requests

This document does not attempt to describe all details for a workspace. For details regarding the column headings (which correspond to attribute names) in a given table view, see Chapter 22, “Attributes,” on page 125. For additional details regarding any workspace, go to the workspace itself and inspect the tables, views, and other features.

Accessing workspaces

This topic describes the security features and tells how to access workspaces in the Tivoli Enterprise Portal.

Security features for user accounts

Security features such as the following regulate access to workspaces or access to the functionality in workspaces:

- The typical user of the Tivoli Enterprise Portal has a standard level of privileges to view and modify the portal. For example, the **sysadmin** default user can see the views and modify the portal. Be aware that if you use a user account that has the **Workspace administration** level of privileges, the views that you see differ, as in these examples:
 - The title and status bars display ***ADMIN MODE*** after the user name to indicate that the user ID is in workspace administration mode.
 - Comprehensive lists of workspace links are displayed in some pop-up menus, instead of the standard menus that are nested within specific workspaces.

The **Workspace administration** mode has additional privileges and purposes that are described in the *IBM Tivoli Monitoring: User's Guide*.

- “Authentication and authorization for running Storage Toolkit commands” on page 261 describes the authorization that is required for the action requests of the Storage Toolkit.

Accessing Tivoli OMEGAMON XE for Storage on z/OS in the portal

Access Tivoli OMEGAMON XE for Storage on z/OS in the portal through the Navigator, as in the following example.

1. Click the plus sign (+) beside the **z/OS Systems** icon to see available systems.
2. Click the plus sign (+) beside the name of the z/OS system for which you want information, for example, **SYSG**. The available subsystems are displayed.
3. Click the plus sign (+) beside the **Storage Subsystem** icon.
4. Click the plus sign (+) beside the **STORAGE** managed system for which you want information, for example, **VTCTH@@L:SYSG:STORAGE**. A list of product nodes is displayed.
5. Click the node of a workspace, for example, **Application Summary**. The portal displays the Application Summary workspace.

“Nodes in the Navigator, which are linked to workspaces” on page 7 shows the product nodes within the hierarchy of a typical Navigator.

Accessing the workspace that corresponds to a node in the Navigator

To access a workspace that corresponds to a node in the Navigator, click the node. The workspace is displayed in the portal.

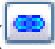
Accessing workspace links from a Navigator node

1. Click the node. The node is activated and the corresponding workspace is displayed in the portal.
2. Right-click the node to display the pop-up menu.
3. Select the workspace you want in the **Workspace** submenu.

Note: Some workspace links that are mentioned in this document are not always visible. The Tivoli Enterprise Portal displays some workspace links only when the target workspaces are relevant in the monitoring environment. For example, if you do not monitor DS8000 storage facilities, a link is not available for a workspace that tracks attributes of DS8000 storage facilities.

Accessing workspace links from views in a workspace

1. You can access related workspaces from table views and bar charts as follows:

- **Table views:** Right-click on the link icon () located at the left of a row in a table view.


- **Bar charts:** Right-click a bar in a bar chart view and select **Link To** in the menu to access the list of links.
2. Select the name of a workspace in the pop-up menu.

Accessing cross-product workspace links

You access cross-product workspace links like any other workspace link. You use this type of link to go to workspaces in other Tivoli products, as described in “Integration with other products.”


Standard path and method for navigating workspaces

This document presents the hierarchy of linked workspaces for each workspace node in Tivoli OMEGAMON XE for Storage on z/OS. For example, the list in Table 2 on page 13 shows the hierarchy for the **Application Summary** workspace node in the Navigator of the portal. The bulleted list illustrates the standard *path* for navigating the workspaces. The standard *method* of navigating workspaces is as follows:

1. Right-click the link icon () that is located at the left side of each row in a table view. (Most workspaces have at least one table view.)
2. Select a new workspace to go to in the pop-up menu.

Note: Some link names in the pop-up menu do not exactly match the name that you see in the title bar of the target workspace.

Keep the following tips in mind as you a navigate workspaces:

- When no link icon is displayed in the table view of a workspace, you have reached the end of the standard flow of navigation for a set of workspaces.
- Most workspace nodes in Tivoli OMEGAMON XE for Storage on z/OS have multiple subpaths.
- Some subpaths terminate in the navigation tree structure of a different, but related, workspace node of Tivoli OMEGAMON XE for Storage on z/OS. For example, many nodes ultimately link to the workspaces of the SMS Storage Groups Performance node, where relevant performance data is available.
- The **Back** button in the browser client or the **Back** arrow icon () in the Desktop client are essential navigation tools as you traverse multiple subpaths. You move backward to leave one subpath and enter another. Alternatively, you can click the workspace node in the Navigator to return to the first workspace of the node in which you are navigating.

Integration with other products

IBM Tivoli Monitoring provides the dynamic workspace linking feature to enable you to quickly access workspaces in multiple products. Tivoli OMEGAMON XE for Storage on z/OS uses that feature to provide predefined links that go to storage-related workspaces in Tivoli OMEGAMON XE on z/OS and in HSM Advanced Reporter. For example, the Application Summary workspace provides a predefined link to the Address Space Bottleneck Detail workspace in Tivoli OMEGAMON XE on z/OS.

Note: You can also modify the links that are provided by the product by default. See *Tivoli Enterprise Portal Help* for information on the Link Wizard that helps you modify the default linkages provided in the popup menus of rows in a workspace.

Note: Dynamic linking is available only for instances of OMEGAMON XE products that report to the same Tivoli Enterprise Monitoring Server as IBM Tivoli OMEGAMON XE for Storage on z/OS and that run in the same logical partition.

Links to workspaces of related products

Table 1 lists the links that are available from the workspaces of Tivoli OMEGAMON XE for Storage on z/OS to the workspaces of related products. "Accessing workspaces" on page 8 provides general information about navigating workspaces.

Table 1. Links to workspaces of related products

| Product to which Tivoli OMEGAMON XE for Storage on z/OS is linked | Workspace where the link is located in Tivoli OMEGAMON XE for Storage on z/OS | Name of the link in the pop-up menu |
|---|---|---|
| IBM Tivoli Advanced Audit for DFSMSHsm | DFSMSHsm Status | Backup Control Dataset Error Summary |
| | | Migration Control Dataset Error Summary |
| | | Offline Control Dataset Error Summary |
| IBM Tivoli Advanced Reporter for DFSMSHsm | DFSMSHsm Status | BCDS Volume Summary |
| | | BCDS Details |
| | | Daily Space Management Detail |
| | | Daily Space Management Summary |
| | | DFSMSHsm Return Code Summary |
| | | Extent Reduction Summary |
| | | Log Migration Detail |
| | | Log Recall Detail |
| | | MCDS Details |
| | MCDS Volume Summary | |
| | SMS Configuration | BCDS Details |
| MCDS Details | | |
| IBM Tivoli Advanced Catalog Management for z/OS | Catalog Summary | Catalog Dataset Attributes |
| | | Catalog Details Link |
| | | Last Catalog Backup |
| | Dataset Group Details Primary | Catalog Dataset Attributes |
| | | Catalog Details |
| | | Last Catalog Backup |
| | Dataset Group Details VSAM | Catalog Dataset Attributes |
| | | Catalog Details |
| | | Last Catalog Backup |
| | Dataset Group VSAM Attributes | Catalog Dataset Attributes |
| | | Catalog Details |
| | | Last Catalog Backup |

Table 1. Links to workspaces of related products (continued)

| Product to which Tivoli OMEGAMON XE for Storage on z/OS is linked | Workspace where the link is located in Tivoli OMEGAMON XE for Storage on z/OS | Name of the link in the pop-up menu |
|---|---|-------------------------------------|
| IBM Tivoli OMEGAMON XE on z/OS | Application Summary | Address Space CPU Usage Details |
| | | Address Space CPU Usage Enclaves |
| | | Address Space Impact Analysis |
| | | Address Space Storage for Job |
| | | Owned Enclaves |
| | DFSMSHsm Function Details | Address Space Bottleneck Details |
| | | Address Space CPU Usage Details |
| | | Address Space CPU Usage Enclaves |
| | | Address Space Impact Analysis |
| | | Address Space Storage for Job |
| | | Owned Enclaves |
| | DFSMSHsm Host Details | Address Space Bottleneck Details |
| | | Address Space CPU Usage Details |
| | | Address Space CPU Usage Enclaves |
| | | Address Space Impact Analysis |
| | | Address Space Storage for Job |
| | | Owned Enclaves |
| | DFSMSHsm Status | Address Space Bottleneck Details |
| | | Address Space CPU Usage Details |
| | | Address Space CPU Usage Enclaves |
| Address Space Impact Analysis | | |
| Address Space Storage for Job | | |
| Owned Enclaves | | |
| DFSMSHsm Storage Usage | Address Space Storage | |

About dynamic linking to cross-product workspaces

Dynamic workspace linking (sometimes referred to as DWL) permits you to easily navigate between workspaces that are provided by multiple products. This feature aids problem determination and improves integration across the monitoring products, permitting you to quickly determine the root cause of a problem. Predefined cross-product links provided by the OMEGAMON XE products permits you to obtain additional information about systems, subsystems, resources, or network components that are being monitored by other monitoring agents.

When you right-click on a link, the list of links is displayed. This list might contain links to workspaces provided by one or more monitoring products. The product you are linking to must be installed and configured.

Note: Your Tivoli Enterprise Portal user ID must be authorized to access the target product. Otherwise links to workspaces in the targeted product are not included in the list.

Choose a workspace from the list to navigate to that workspace. By linking to the target workspace in context, you receive additional information that is related to the system, subsystem, or resource you are currently viewing.

If you choose a workspace from the list and the target workspace is not available, you receive message KFWITM081E. Refer to the *Problem Determination Guide* for more information. Also, the "Target Not Found" topic in the online help in the Tivoli Enterprise Portal describes options for responding to failed links. The following scenario describes how this problem can arise:

- You are upgrading from OMEGAMON XE version 4.1.0 products to OMEGAMON XE version 4.2.0 products, and you have a combination of version 4.1.0 and version 4.2.0 monitoring agents installed in your environment. For example, you might have a Tivoli OMEGAMON XE for Storage on z/OS version 4.2.0 monitoring agent and an OMEGAMON XE on z/OS version 4.1.0 monitoring agent running on the same z/OS system during the migration period.
- After this upgrade, the use of dynamic workspace linking to link from an OMEGAMON XE version 4.2.0 workspace to workspaces in an OMEGAMON XE version 4.1.0 product works correctly, as long as the target workspace exists in the version 4.1.0 product. If the target workspace does not exist, you receive the KFWITM081E message.
- However, if a target workspace for version 4.2.0 has been modified (for example to accept link parameters to limit the data displayed) you might notice different behavior during the upgrade.

Chapter 3. Monitoring applications

This topic describes the workspaces that monitor the status of applications. You access these workspaces through the Application Summary node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the Application Summary node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Application Summary node to navigate to the workspaces that you see in Table 2. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 2. Application Summary node: bulleted list showing the hierarchy of linked workspaces

- | |
|---|
| Application Summary <ul style="list-style-type: none">• Application Dataset Performance<ul style="list-style-type: none">– Dataset Details• Application Volume Performance<ul style="list-style-type: none">– Application Volume Dataset Performance– Physical Volume Group Performance– Device Performance Details– Cross System Volume Performance– Volume Performance Trend<ul style="list-style-type: none">- Dataset Performance Summary History<ul style="list-style-type: none">• Dataset Performance Historical Detail• Application Volume Cache<ul style="list-style-type: none">– Application Volume Dataset Performance– Physical Volume Group Cache– Device Performance Details– Cross System Volume Performance– Volume Cache Trend• Application Dataset Space Usage<ul style="list-style-type: none">– Dataset Details• Application Volume Space Usage<ul style="list-style-type: none">– Application Dataset Volume Space Utilization– Device Space Details |
|---|

Application Summary workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help, including in the *hover help* (a form of online help that can be viewed by moving a cursor over a GUI item such as an icon or a field). The direct situation analysis feature is

available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See Chapter 20, “Directly monitoring situations,” on page 103.

Views The workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading. For information about related workspaces, see “Summary information for the workspaces in this node” on page 15.

Application table view

The table view shows one row of data for each active address space for which monitoring has been requested. The view contains the following information the following information:

- The name of the job or started task
- The decimal ASID number
- The I/O rate for the job
- The number of data sets allocated for the job
- The worst response time suffered by any of those data sets
- The name of the data set suffering the worst response time and the volume on which it is located
- The number of volumes in use by the job and the worst response time on any of those volumes along with the name of the volume incurring that response time

Applications with Dataset MSR > 40 bar graph

This graph contains a bar for each monitored application that has a data set with an average millisecond response time since the beginning of the RMF™ interval exceeded 40 milliseconds. The height of the bar indicates the worst data set response time for the application.

Applications with volume MSR > 25 bar graph

This graph contains a bar for each monitored application that has a volume with an average millisecond response time in the previous RMF interval exceeding 25 milliseconds. The height of the bar indicates the worst volume response time for the application.

Defining applications to monitor

You define the list of monitored applications by setting up a situation. The situation can be named anything, however, it must be written against the Application Monitoring group. The attributes that must be set in the situation to control monitoring of the application include the attributes and values shown in the following table.

| Attribute | Value |
|--------------------|--|
| Application | Name of the application for which monitoring is to be started or stopped. To specify a group of applications to monitor whose names all begin with a common string (for example, CICSxxx), use the Value of Expression function in the Situation Editor. Type the application name with an asterisk as the final character, for example, CICS*. This specification enables monitoring for all applications that start with the string that precedes the asterisk. |
| Monitor Status | Y = Begin monitoring this application. N = Stop monitoring this application. |
| I/O Monitor Status | Y = Turn ON data set-level I/O monitoring for this application. N = Turn OFF data set-level I/O monitoring for this application. |

To collect data set performance statistics for all data sets accessed by an application, set the 'I/O Monitor Status' attribute to 'Y' when defining the situation that starts monitoring for the application. The following additional guidelines can help you customize data collection:

- Data set performance statistics are gathered for all data sets accessed by the application regardless of the volume on which the data sets are located and the data set performance options in effect for that volume.
- To discontinue the collection of data set performance statistics for the application, you change the value of the 'I/O Monitor Status' attribute to 'N'. Depending on the data set collection status of the volumes being accessed by the application, some data set performance statistics continue to be collected.
- This workspace is one of two places that display statistics from the level of data sets. These statistics are available only when the Tivoli Enterprise subsystem is running.

Workspaces that are linked to this workspace

See "Navigating the workspaces of the Application Summary node" on page 13.

Learn how to navigate among workspaces here: "Accessing workspaces" on page 8. Also see "Integration with other products" on page 9.

See Chapter 2, "Workspaces," on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Application Summary node.

See Chapter 2, "Workspaces," on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: "Accessing workspaces" on page 8. Also see "Integration with other products" on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See "Application Summary workspace" on page 13 for more information on the primary workspace for this node.

| Application Summary workspace | | |
|--------------------------------------|------------------|------------------------|
| View | View type | Attribute group |
| Application Summary Report | TABLE | Application Monitoring |
| Applications with Dataset MSR > 40 | BAR | Application Monitoring |
| Applications with Volume MSR > 25 | BAR | Application Monitoring |

| Application Dataset Performance workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Dataset Performance Report | TABLE | Dataset Performance |
| Datasets with MSR > Storage Class Objective | TABLE | Dataset Performance |

| Application Volume Performance workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |

| | | |
|---------------------------|-------|--------------------------------|
| Volume Performance Report | TABLE | Application Volume Performance |
| Volume Response Time > 25 | BAR | Application Volume Performance |

| Application Volume Dataset Performance workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Dataset Performance Report | TABLE | Dataset Performance |
| Datasets with MSR > Storage Class Objective | TABLE | Dataset Performance |

| Application Volume Cache workspace | | |
|---|------------------|--------------------------------|
| View | View type | Attribute group |
| Volume Cache Report | TABLE | Application Volume Performance |
| Read Hit Percent < 55 | BAR | Application Volume Performance |
| Write Hit Percent < 30 | BAR | Application Volume Performance |

| Application Dataset Space Usage workspace | | |
|--|------------------|-----------------------------|
| View | View type | Attribute group |
| Dataset Space Report | TABLE | Application Dataset Details |
| Volumes with Free Space % < 10 | BAR | Application Volume Space |

| Application Volume Space Usage workspace | | |
|---|------------------|--------------------------|
| View | View type | Attribute group |
| Volumes with Free Space % < 10 | BAR | Application Volume Space |
| Volumes with Fragmentation Index > 650 | BAR | Application Volume Space |
| Volume Space Report | TABLE | Application Volume Space |

| Application Dataset Volume Space Usage workspace | | |
|---|------------------|-----------------------------|
| View | View type | Attribute group |
| Dataset Space Report | TABLE | Application Dataset Details |
| Volume Space Allocation | PIE | DASD Volume Space |

| Application Dataset Response Time workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Dataset Performance Report | TABLE | Dataset Performance |
| Datasets with MSR > Storage Class Objective | TABLE | Dataset Performance |

Chapter 4. Monitoring channel paths

This topic describes the workspaces that monitor the status of channel paths. You access these workspaces through the Channel Path node in the Tivoli Enterprise Portal.

Navigating the workspaces of the Channel Path node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Channel Path node to navigate to the workspaces that you see in Table 3. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 3. Channel Path node: bulleted list showing the hierarchy of workspace links

| |
|---|
| <p>Channel Path</p> <ul style="list-style-type: none">• Channel Volume Performance (restricted to channels with hard disk drives attached)<ul style="list-style-type: none">- Dataset Performance Summary<ul style="list-style-type: none">- Dataset Performance Detail<ul style="list-style-type: none">• Dataset Details• Dataset Performance Detail Trend- Dataset Details- Dataset Performance Summary Trend<ul style="list-style-type: none">• Dataset Performance Historical Detail- Device Performance Details- Cross System Volume Performance- Volume Performance Trend<ul style="list-style-type: none">- Dataset Historic Performance Summary<ul style="list-style-type: none">• Dataset Performance Historical Detail• Channel Volume Cache (restricted to channels with hard disk drives attached)<ul style="list-style-type: none">- Dataset Performance Summary<ul style="list-style-type: none">- (as above)- Device Performance Details- Cross System Volume Performance- Volume Cache Trend• Channel Path Trend<ul style="list-style-type: none">- Channel Path Historical Volume Performance<ul style="list-style-type: none">- Dataset Historic Performance Summary<ul style="list-style-type: none">• Dataset Performance Historical Detail- Channel Path Historical Volume Cache Performance<ul style="list-style-type: none">- Dataset Historic Performance Summary<ul style="list-style-type: none">• Dataset Performance Historical Detail <p>The following workspaces display the volumes that have the least favorable readings for the attribute within the SSID selected for navigation:</p> <ul style="list-style-type: none">• Highest Channel Volumes Response Time (Top 10)• Highest Channel Volumes MPL (Top 10)• Highest Channel Volumes I/O Rate (Top 10)• Highest Channel Volumes Busy Percent (Top 10)• Lowest Channel Volumes Connect Percent (Top 10) |
|---|

Channel Path workspace

This topic describes the main components of the principal workspace for this node.

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See Chapter 20, “Directly monitoring situations,” on page 103.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Channel path table view

The table view displays one row of data for each channel path. The view contains the following information:

- The channel path ID
- Its status
- Whether or not it is a shared path
- Utilization percentage across the complex
- Utilization percentage within this LPAR
- The type of connection it is
- Whether or not it is DCM managed
- The BUS utilization percentage
- Various read and write metrics for both the complex and the reporting LPAR
- The number of SMS volumes, NON-SMS volumes and the total number of volumes configured in the LPAR
- The channel path management facility group number for the path

Complex % utilization bar graph

This graph contains a bar for each channel path that had any utilization since the start of the RMF interval. The height of the bar indicates the percent to which the path was utilized in the complex.

LPAR % utilization DASD bar graph

This graph contains a bar for each channel path for which there are any volumes configured in the reporting LPAR and for which there was any utilization within the reporting LPAR since the beginning of the RMF interval.

LPAR MB/Sec WRITE - DASD bar graph

This graph contains a bar for each channel path for which there are any volumes configured in the reporting LPAR and for which there were any write operations that were performed to DASD in the reporting LPAR since the beginning of the RMF interval.

LPAR MB/Sec READ - DASD bar graph

This graph contains a bar for each channel path for which there are any volumes configured in the reporting LPAR and for which there were any read operations to DASD that were performed in the reporting LPAR since the beginning of the RMF interval.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Channel Path node” on page 17.

Learn how to navigate to workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Channel Path node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “Channel Path workspace” on page 17 for more information on the primary workspace for this node.

| Channel Path workspace | | |
|-------------------------------|------------------|------------------------|
| View | View type | Attribute group |
| Complex % Utilization | BAR | Channel Path |
| LPAR % Utilization - DASD | BAR | Channel Path |
| LPAR MB/Sec Written - DASD | BAR | Channel Path |
| LPAR MB/Sec Read - DASD | BAR | Channel Path |
| Channel Path Report | TABLE | Channel Path |

| Channel Volume Performance workspace | | |
|---|------------------|-------------------------|
| View | View type | Attribute group |
| Volume Response Time > 25 | BAR | DASD Volume Performance |
| Volume Performance Report | TABLE | DASD Volume Performance |

| Channel Volume Cache workspace | | |
|---------------------------------------|------------------|------------------------|
| View | View type | Attribute group |
| Volume Cache Report | TABLE | Cache Devices |
| Read Hit Percent < 55 | BAR | Cache Devices |
| Write Hit Percent < 30 | BAR | Cache Devices |

| Channel Path Trend workspace | | |
|-------------------------------------|------------------|------------------------|
| View | View type | Attribute group |
| Channel Utilization | PLOT | Channel Path |
| Channel Path Trend | TABLE | Channel Path |

| Channel Volume Response Time workspace | | |
|---|------------------|-------------------------|
| View | View type | Attribute group |
| Volume Performance Report | TABLE | DASD Volume Performance |
| Volume Response Time > 25 | BAR | DASD Volume Performance |

Chapter 5. Monitoring cache CU performance

This topic describes the workspaces that monitor the status of the performance of cache CU. You access these workspaces through the Cache CU Performance node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the Cache CU Performance node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Cache CU Performance node to navigate to the workspaces that you see in Table 4. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 4. Cache CU Performance node: bulleted list showing the hierarchy of linked workspaces

- | |
|--|
| Cache CU Performance |
| • Cache CU Volume Cache |
| – Dataset Performance Summary ¹ |
| – Physical Volume Group Cache |
| – Dataset Performance Summary ¹ |
| – Device Performance Details |
| – Cross System Volume Performance |
| – Volume Cache Trend |

Table 4. Cache CU Performance node: bulleted list showing the hierarchy of linked workspaces (continued)

| |
|---|
| <ul style="list-style-type: none"> • Cache CU Volume Performance <ul style="list-style-type: none"> – Dataset Performance Summary ¹ – Physical Volume Group Performance <ul style="list-style-type: none"> - Dataset Performance Summary ¹ - Volume Performance Trend <ul style="list-style-type: none"> • Dataset Performance Summary History – Device Performance Details – Cross System Volume Performance – Volume Performance Trend <ul style="list-style-type: none"> - Dataset Historic Performance Summary <ul style="list-style-type: none"> • Dataset Performance Historical Detail • Symmetrix Configuration. <p>Note: Obtaining EMC Symmetric data can only be done in conjunction with the applicable EMC hardware and software. The hardware required is an EMC Symmetrix at the appropriate microcode level. Also required are the EMC Symmetrix Control Facility runtime libraries available from EMC.</p> • Cache CU Raid Rank <ul style="list-style-type: none"> – Physical Volume Group Cache <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Cache CU Raid Rank Trend • Cache CU Performance Trend <ul style="list-style-type: none"> – Cache CU Volume Historical Performance <ul style="list-style-type: none"> - Dataset Performance Summary History <ul style="list-style-type: none"> • Dataset History Performance Detail – CCU Volume Cache History <ul style="list-style-type: none"> - Dataset Historic Performance Summary <ul style="list-style-type: none"> • Dataset Performance Historical Detail – Cache CU Raid Rank Historical Data • TotalStorage Configuration <ul style="list-style-type: none"> – TotalStorage Ranks <ul style="list-style-type: none"> - TotalStorage Array Configuration - TotalStorage Ranks Trend – TotalStorage Extent Pool Volumes <ul style="list-style-type: none"> - TotalStorage Volume Trend – TotalStorage Extent Pool Trend <ul style="list-style-type: none"> - TotalStorage Rank History - TotalStorage Volume History • TotalStorage CU Volumes <ul style="list-style-type: none"> – TotalStorage Volume Trend <p>The following workspaces display the volumes that have the least favorable readings for the attribute within the SSID selected for navigation:</p> <ul style="list-style-type: none"> • Lowest Volumes CU Read Hit Percent (Top 10) ² • Lowest Volumes CU Write Hit Percent (Top 10) ² • Lowest Volumes DFW Hit Percent (Top 10) ² • Highest Volumes Destaging Rate (Top 10) ² • Highest Volumes CU Staging Rate (Top 10) ² |
|---|

Footnotes for Table 4 on page 21:

1. This group of workspaces is organized in the following nested pattern:
 - Dataset Performance Summary
 - Dataset Performance Detail
 - Dataset Details
 - Dataset Performance Detail Trend
 - Dataset Details
 - Dataset Performance Summary Trend

- Dataset Performance Historical Detail
- 2. This workspace links to the following set of related workspaces:
 - Dataset Performance Summary (See the preceding note ⁽¹⁾ for details.)
 - Physical Volume Group Performance
 - Dataset Performance Summary (See the preceding note ⁽¹⁾ for details.)
 - Volume Performance Trend
 - Dataset Performance Summary History
 - Device Performance Details
 - Cross System Volume Performance
 - Volume Cache Trend
 - Dataset Historic Performance Summary
 - Dataset Performance Historical Detail

Access the following linked workspaces by selecting the Cache CU Performance node in the Navigator. Right-click the node. Select the workspace that you want to access in the pop-up menu.

These workspaces list the volumes with the least favorable readings for the attribute across the z/OS image.

- Lowest Volume Read Hit Percent
- Lowest Volume Write Hit Percent
- Lowest Volume DFW Hit Percent
- Highest Volume Destaging Rate
- Highest Volume Staging Rate

Cache CU Performance workspace

This topic describes the main components of the principal workspace for this node.

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See Chapter 20, “Directly monitoring situations,” on page 103.

Views The workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading. For information about related workspaces, see “Summary information for the workspaces in this node” on page 24.

Cache performance table view

The table view displays one row of data for each cache control unit configured to the reporting LPAR. The view contains the following information:

- The subsystem ID of the control unit
- The control unit type
- The read and write hit percentages for the control unit
- The read I/O percent
- The bypass and inhibit percents
- The DASD Fast Write hit and retry percents
- The Cache Fast Write read and write percents
- The track destaging rate
- The track staging rate
- Storage facility series

Read Hit Percent bar graph

This graph contains a bar for each cache control unit that had a read hit percent greater than zero. The height of the bar indicates the read hit percentage.

Write Hit Percent bar graph

This graph contains a bar for each cache control unit that had a write hit percent greater than zero. The height of the bar indicates the write hit percentage.

DFW Hit Percent bar graph

This graph contains a bar for each cache control unit with a DASD Fast Write hit percent greater than zero. The height of the bar indicates the DASD Fast Write hit percentage.

DFW Retry Percent bar graph

This graph contains a bar for each cache control unit with a DASD Fast Write retry percent greater than zero. The height of the bar indicates the DASD Fast Write retry percentage.

Track Destaging Rate bar graph

This graph contains a bar for each cache control unit with a track destaging rate greater than zero. The height of the bar indicates the track destaging rate.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Cache CU Performance node” on page 21.

Learn how to navigate to workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Cache CU Performance node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “Cache CU Performance workspace” on page 23 for more information on the primary workspace for this node.

| Cache CU Performance workspace | | |
|---------------------------------------|-----------|--------------------|
| View | View type | Attribute group |
| Cache Control Unit Performance Report | TABLE | Cache Control Unit |
| DFW Hit Percent | BAR | Cache Control Unit |
| DFW Retry Percent | BAR | Cache Control Unit |
| Track Destaging Rate | BAR | Cache Control Unit |
| Read Hit Percent | BAR | Cache Control Unit |
| Write Hit Percent | BAR | Cache Control Unit |

| Cache CU Volume Cache workspace | | |
|---------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| Read Hit Percent < 55 | BAR | Cache Devices |
| Write Hit Percent < 30 | BAR | Cache Devices |

| | | |
|---------------------|-------|---------------|
| Volume Cache Report | TABLE | Cache Devices |
|---------------------|-------|---------------|

| Cache CU Volume Performance workspace | | |
|---------------------------------------|-----------|-------------------------|
| View | View type | Attribute group |
| Volume Performance Report | TABLE | DASD Volume Performance |
| Volume Response Time > 25 | BAR | DASD Volume Performance |

| Cache CU Raid Rank workspace | | |
|------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| Raid Rank Report | TABLE | Cache Raid Rank |
| Read Response Time | BAR | Cache Raid Rank |
| Write Response Time | BAR | Cache Raid Rank |

| Symmetrix Configuration workspace | | |
|-----------------------------------|-----------|-------------------------|
| View | View type | Attribute group |
| Subsystem IDs | TABLE | Symmetrix SSIDs |
| Symmetrix Configuration Report | TABLE | Symmetrix Configuration |
| Symmetrix Directors Report | TABLE | Symmetrix Directors |

| Symmetrix Disk Director Summary workspace | | |
|---|-----------|----------------------------|
| View | View type | Attribute group |
| Logical Volumes Per Disk | BAR | Symmetrix Director Summary |
| Cache Efficiency | BAR | Symmetrix Directors |

| Symmetrix Devices workspace | | |
|-------------------------------|-----------|-------------------|
| View | View type | Attribute group |
| Devices with Mirrors in Error | TABLE | Symmetrix Devices |
| Symmetrix Devices | TABLE | Symmetrix Devices |

| Cache CU Performance Trend workspace | | |
|--------------------------------------|-----------|--------------------|
| View | View type | Attribute group |
| Cache CU Performance | PLOT | Cache Control Unit |
| Cache Control Unit Performance Trend | TABLE | Cache Control Unit |

| Volume Cache Trend workspace | | |
|------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| Volume Cache Trend | TABLE | Cache Devices |
| Cache Hit Percent | PLOT | Cache Devices |

| Volume Cache History workspace | | |
|--------------------------------|-----------|-----------------|
| View | View type | Attribute group |

| | | |
|------------------------|-------|---------------|
| Volume Cache Report | TABLE | Cache Devices |
| Read Hit Percent < 55 | BAR | Cache Devices |
| Write Hit Percent < 30 | BAR | Cache Devices |

| Cache CU Historic Raid Rank workspace | | |
|---------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| Raid Rank Report | TABLE | Cache Raid Rank |
| Read Response Time | BAR | Cache Raid Rank |
| Write Response Time | BAR | Cache Raid Rank |

| Cache CU Raid Rank Trend workspace | | |
|------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| Raid Rank Report | TABLE | Cache Raid Rank |
| Response Time | PLOT | Cache Raid Rank |

| TotalStorage® Ranks workspace | | |
|-------------------------------|-----------|---------------------|
| View | View type | Attribute group |
| TotalStorage Ranks | TABLE | TotalStorageDS Rank |
| Operations Per Second | BAR | TotalStorageDS Rank |
| Average Response Time | BAR | TotalStorageDS Rank |

| TotalStorage Array Configuration workspace | | |
|--|-----------|----------------------|
| View | View type | Attribute group |
| DDM Capacity | BAR | TotalStorageDS Array |
| TotalStorage Arrays | TABLE | TotalStorageDS Array |
| DDM Speed | BAR | TotalStorageDS Array |

| TotalStorage CU Volumes workspace | | |
|-----------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| TotalStorage CU Volumes | TABLE | Cache Devices |
| Physical I/O Rate | BAR | Cache Devices |
| Physical Response Time | BAR | Cache Devices |

| TotalStorage Extent Pool Volumes workspace | | |
|--|-----------|-----------------|
| View | View type | Attribute group |
| TotalStorage Extent Pool Volumes | TABLE | Cache Devices |
| Physical I/O Rate | BAR | Cache Devices |
| Average Physical Response Time | BAR | Cache Devices |

| TotalStorage Extent Pool Trend workspace | | |
|--|-----------|----------------------------|
| View | View type | Attribute group |
| TotalStorage Extent Pool Trend | TABLE | TotalStorageDS Extent Pool |

| | | |
|---------------|------|----------------------------|
| Response Time | PLOT | TotalStorageDS Extent Pool |
|---------------|------|----------------------------|

| TotalStorage Ranks Trend workspace | | |
|------------------------------------|-----------|---------------------|
| View | View type | Attribute group |
| TotalStorage Ranks Trend | TABLE | TotalStorageDS Rank |
| Average Response Time | PLOT | TotalStorageDS Rank |

| TotalStorage Volume Trend workspace | | |
|-------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| TotalStorage Volume Trend | TABLE | Cache Devices |
| Physical Response Time | PLOT | Cache Devices |

| TotalStorage Rank History workspace | | |
|-------------------------------------|-----------|---------------------|
| View | View type | Attribute group |
| TotalStorage Ranks | TABLE | TotalStorageDS Rank |
| Operations per Second | BAR | TotalStorageDS Rank |
| Average Response Time | BAR | TotalStorageDS Rank |

| TotalStorage Volume History workspace | | |
|---------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| TotalStorage Volumes | TABLE | Cache Devices |
| Physical I/O Rate | BAR | Cache Devices |
| Physical Response Time | BAR | Cache Devices |

| TotalStorage Configuration workspace | | |
|--------------------------------------|-----------|------------------------------|
| View | View type | Attribute group |
| Average Read Response Time | BAR | TotalStorageDS Extent Pool |
| Average Write Response Time | BAR | TotalStorageDS Extent Pool |
| TotalStorage Extent Pools | TABLE | TotalStorageDS Extent Pool |
| TotalStorage SSIDs | TABLE | TotalStorageDS SSIDs |
| TotalStorage Configuration | TABLE | TotalStorageDS Configuration |

| Cache CU Destaging Trend workspace | | |
|--------------------------------------|-----------|--------------------|
| View | View type | Attribute group |
| Cache Control Unit Performance Trend | TABLE | Cache Control Unit |
| Cache CU Physical Activity | PLOT | Cache Control Unit |

| Cache CU Volume Destaging workspace | | |
|-------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| Physical Volume Activity | BAR | Cache Devices |
| Volume Cache Destaging Report | TABLE | Cache Devices |

| Volume Destaging Trend workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Volume Destaging Trend | TABLE | Cache Devices |
| Physical Volume Activity | PLOT | Cache Devices |

| Volume Destaging History workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Physical Volume Activity | BAR | Cache Devices |
| Volume Historical Destaging Report | TABLE | Cache Devices |

| Cache CU Staging Trend workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Cache CU Physical Activity | PLOT | Cache Control Unit |
| Cache Control Unit Performance Trend | TABLE | Cache Control Unit |

| Cache CU Volume Staging workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Volume Cache Staging Report | TABLE | Cache Devices |
| Physical Volume Activity | BAR | Cache Devices |

| Volume Staging Trend workspace | | |
|---------------------------------------|------------------|------------------------|
| View | View type | Attribute group |
| Volume Staging Trend | TABLE | Cache Devices |
| Physical Volume Activity | PLOT | Cache Devices |

| Volume Staging History workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Volume Historical Staging Report | TABLE | Cache Devices |
| Physical Volume Activity | BAR | Cache Devices |

| Cache CU Volume DFW Retry workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Volume Cache DFW Retry Report | TABLE | Cache Devices |
| DFW Retry % | BAR | Cache Devices |
| DFW Hit % | BAR | Cache Devices |

| Cache CU DFW Retry Trend workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| DFW Retry % | PLOT | Cache Control Unit |
| Cache Control Unit DFW Retry Trend | TABLE | Cache Control Unit |

| Cache CU Volume DFW Retry History workspace | | |
|---|-----------|-----------------|
| View | View type | Attribute group |
| Volume Cache Report | TABLE | Cache Devices |
| DFW Retry % | BAR | Cache Devices |
| DFW Hit % | BAR | Cache Devices |

| Cache CU Volume Read Hit Percent workspace | | |
|--|-----------|-----------------|
| View | View type | Attribute group |
| Volume Cache Report | TABLE | Cache Devices |
| Read Hit Percent | BAR | Cache Devices |
| Write Hit Percent | BAR | Cache Devices |

| Cache CU Read Hit Percent Trend workspace | | |
|---|-----------|--------------------|
| View | View type | Attribute group |
| Cache CU Read Hit Percent | PLOT | Cache Control Unit |
| Cache Control Unit Performance Trend | TABLE | Cache Control Unit |

| Cache CU Volume Read Hit Percent History workspace | | |
|--|-----------|-----------------|
| View | View type | Attribute group |
| Volume Cache Report | TABLE | Cache Devices |
| Read Hit Percent | BAR | Cache Devices |
| Write Hit Percent | BAR | Cache Devices |

| Cache CU Volume Write Hit Percent workspace | | |
|---|-----------|-----------------|
| View | View type | Attribute group |
| Volume Cache Report | TABLE | Cache Devices |
| Write Hit Percent | BAR | Cache Devices |
| Read Hit Percent | BAR | Cache Devices |

| Cache CU Write Hit Percent Trend workspace | | |
|--|-----------|--------------------|
| View | View type | Attribute group |
| Cache CU Write Hit Percent | PLOT | Cache Control Unit |
| Cache Control Unit Performance Trend | TABLE | Cache Control Unit |

| Cache CU Volume Write Hit Percent History workspace | | |
|---|-----------|-----------------|
| View | View type | Attribute group |
| Volume Cache Report | TABLE | Cache Devices |
| Write Hit Percent | BAR | Cache Devices |
| Read Hit Percent | BAR | Cache Devices |

| Cache CU Volume Inactive Status workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Volume Cache Report | TABLE | Cache Devices |
| I/O Count | BAR | Cache Devices |

| Lowest Volume Read Hit Percent workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Volumes With Lowest Read Hit Percent | TABLE | Cache Devices |
| Lowest Read Hit Percent | BAR | Cache Devices |
| Staging Rate | BAR | Cache Devices |

| Lowest Volume Write Hit Percent workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Volumes With Lowest Write Hit Percent | TABLE | Cache Devices |
| Lowest Write Hit Percent | BAR | Cache Devices |
| Staging Rate | BAR | Cache Devices |

| Lowest Volume DFW Hit Percent workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Volumes With Lowest DFW Hit Percent | TABLE | Cache Devices |
| Lowest DFW Hit Percent | BAR | Cache Devices |
| DFW Retry Percent | BAR | Cache Devices |

| Highest Volume Destaging Rate workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Volumes With Highest Destaging Rate | TABLE | Cache Devices |
| Highest Destaging Rate | BAR | Cache Devices |
| Volume Staging Rate | BAR | Cache Devices |

| Highest Volume Staging Rate workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Volumes With Highest Staging Rate | TABLE | Cache Devices |
| Highest Staging Rate | BAR | Cache Devices |
| Volume Destaging Rate | BAR | Cache Devices |

| Lowest Volume CU Read Hit Percent workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Lowest Volume Read Hit Percent Report | TABLE | Cache Devices |

| | | |
|-------------------------|-----|---------------|
| Lowest Read Hit Percent | BAR | Cache Devices |
| Staging Rate | BAR | Cache Devices |

| Lowest Volume CU Write Hit Percent workspace | | |
|--|-----------|-----------------|
| View | View type | Attribute group |
| Lowest Volume Write Hit Percent Report | TABLE | Cache Devices |
| Lowest Write Hit Percent | BAR | Cache Devices |
| Staging Rate | BAR | Cache Devices |

| Lowest Volume CU DFW Hit Percent workspace | | |
|--|-----------|-----------------|
| View | View type | Attribute group |
| Lowest Volume DFW Hit Percent Report | TABLE | Cache Devices |
| Lowest DFW Hit Percent | BAR | Cache Devices |
| DFW Retry Percent > 0 | BAR | Cache Devices |

| Highest Volume CU Destaging Rate workspace | | |
|--|-----------|-----------------|
| View | View type | Attribute group |
| Highest Volume Destaging Rate Report | TABLE | Cache Devices |
| Highest Destaging Rate | BAR | Cache Devices |
| Volume Staging Rate > 0 | BAR | Cache Devices |

| Highest Volume CU Staging Rate workspace | | |
|--|-----------|-----------------|
| View | View type | Attribute group |
| Highest Volume Staging Rate Report | TABLE | Cache Devices |
| Highest Staging Rate | BAR | Cache Devices |
| Volume Destaging Rate > 0 | BAR | Cache Devices |

Chapter 6. Monitoring cache CU status

This topic describes the workspaces that monitor the status of cache CU. You access these workspaces through the Cache CU Status node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the Cache CU Status node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Cache CU Status node to navigate to the workspaces that you see in Table 5 on page 34. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 5. Cache CU Status node: bulleted list showing the hierarchy of linked workspaces

| |
|---|
| <p>Cache CU Status</p> <ul style="list-style-type: none"> • Cache CU Volume Cache Performance <ul style="list-style-type: none"> – Dataset Performance Summary ¹ – Physical Volume Group Cache <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Device Performance Details – Cross System Volume Performance <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Volume Cache Trend • Cache CU Volume Performance <ul style="list-style-type: none"> – Dataset Performance Summary ¹ – Physical Volume Group Performance <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Volume Performance Trend <ul style="list-style-type: none"> - Dataset Performance Summary History – Device Performance Details – Cross System Volume Performance <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Volume Performance Trend <ul style="list-style-type: none"> - Dataset Performance Summary History <ul style="list-style-type: none"> • Dataset Performance Historical Detail • Symmetrix Configuration Note: Obtaining EMC Symmetric data can only be done in conjunction with the applicable EMC hardware and software. The hardware required is an EMC Symmetrix at the appropriate microcode level. Also required are the EMC Symmetrix Control Facility runtime libraries available from EMC. • Cache CU Raid Rank <ul style="list-style-type: none"> – Physical Volume Group Cache <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Cache CU Raid Rank Trend • Cache CU Status Trend <ul style="list-style-type: none"> – CCU Volume Cache History <ul style="list-style-type: none"> - Dataset Historic Performance <ul style="list-style-type: none"> • Dataset Performance Historical Detail • TotalStorage Configuration <ul style="list-style-type: none"> – TotalStorage Ranks <ul style="list-style-type: none"> - TotalStorage Array Configuration - TotalStorage Ranks Trend – TotalStorage Extent Pool Volumes <ul style="list-style-type: none"> - TotalStorage Volume Trend – TotalStorage Extent Pool Trend <ul style="list-style-type: none"> - TotalStorage Rank History - TotalStorage Volume History • TotalStorage CU Volumes <ul style="list-style-type: none"> – TotalStorage Volume Trend |
|---|

Footnotes for Table 5:

1. This group of workspaces is organized in the following nested pattern:
 - Dataset Performance Summary
 - Dataset Performance Detail
 - Dataset Details
 - Dataset Performance Detail Trend
 - Dataset Details
 - Dataset Performance Summary Trend
 - Dataset Performance Historical Detail

Cache CU Status workspace

This topic describes the main components of the principal workspace for this node.

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See Chapter 20, “Directly monitoring situations,” on page 103.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Cache status table view

The table view displays one row of data for each cache control unit configured to the reporting LPAR. The view contains the following information:

- The subsystem ID of the control unit
- The control unit type
- The number of active volumes
- The number of deactivated volumes
- The cache status
- The amount of cache configured in the control unit and the amount currently available
- The nonvolatile storage status, the amount of NVS configured and the amount pinned
- Storage facility series

Available cache bar graph

This graph contains a bar for each cache control unit. The height of the bar indicates the amount of cache (in megabytes) available in the unit.

NVS pinned bar graph

This graph contains a bar for each cache control unit. The height of the bar indicates the amount of NVS (nonvolatile storage, in kilobytes) pinned in the unit.

Deactivated volumes bar graph

This graph contains a bar for each cache control unit. The height of the bar indicates the number of deactivated volumes on that control unit.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Cache CU Status node” on page 33.

Learn how to navigate to workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Cache CU Status node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “Cache CU Status workspace” on page 34 for more information on the primary workspace for this node.

| |
|---------------------------|
| Cache CU Status workspace |
|---------------------------|

| View | View type | Attribute group |
|----------------------------------|------------------|------------------------|
| Cache Control Unit Status Report | TABLE | Cache Control Unit |
| Available Cache | BAR | Cache Control Unit |
| NVS Pinned | BAR | Cache Control Unit |
| Deactivated Volumes | BAR | Cache Control Unit |

| Cache CU Status Trend workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Cache Control Unit Status Trend | TABLE | Cache Control Unit |
| NVS Pinned | PLOT | Cache Control Unit |
| Deactivated Volumes | PLOT | Cache Control Unit |
| Cache Available | PLOT | Cache Control Unit |

| TotalStorage DS Configuration workspace | | |
|--|------------------|------------------------------|
| View | View type | Attribute group |
| Report | TABLE | TotalStorageDS Configuration |
| Report | TABLE | TotalStorageDS SSIDs |
| Report | TABLE | TotalStorageDS Extent Pool |

Chapter 7. Monitoring logical control units

This topic describes the workspaces that monitor the status of logical control units. You access these workspaces through the Logical Control Unit node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the Logical Control Unit node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Logical Control Unit node to navigate to the workspaces that you see in Table 6 on page 38. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 6. Logical Control Unit node: bulleted list showing the hierarchy of linked workspaces

| |
|--|
| <p>Logical Control Unit</p> <ul style="list-style-type: none"> • LCU Volume Performance <ul style="list-style-type: none"> – Dataset Performance Summary ¹ – Physical Volume Group Performance <ul style="list-style-type: none"> - Dataset Performance Summary ¹ - Volume Performance Trend <ul style="list-style-type: none"> • Dataset Performance Summary History – Device Performance Details – Cross System Volume Performance <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Volume Performance Trend <ul style="list-style-type: none"> - Dataset Historic Performance Summary <ul style="list-style-type: none"> • Dataset Performance Historical Detail • LCU Channel Detail • LCU Volume Cache <ul style="list-style-type: none"> – Dataset Performance Summary ¹ – Physical Volume Group Cache <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Device Performance Details – Cross System Volume Performance <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Volume Cache Trend • LCU DCM Channel Summary • Logical Control Unit Trend <ul style="list-style-type: none"> – LCU Volume Historical Performance <ul style="list-style-type: none"> - Dataset Historic Performance Summary <ul style="list-style-type: none"> • Dataset Performance Historical Detail – LCU Historical Volume Cache Performance <ul style="list-style-type: none"> - Dataset Historic Performance Summary <ul style="list-style-type: none"> • Dataset Performance Historical Detail <p>These workspaces display the volumes with the least favorable readings for the attribute within the logical control unit selected for navigation.</p> <ul style="list-style-type: none"> • Highest LCU Volumes Response Time (Top 10) ² • Highest LCU Volumes MPL (Top 10) ² • Highest LCU Volumes I/O Rate (Top 10) ² • Highest LCU Volumes Busy Percent (Top 10) ² • Lowest LCU Volumes Connect Percent (Top 10)² |
|--|

Footnotes for Table 6:

1. This group of workspaces is organized in the following nested pattern:
 - Dataset Performance Summary
 - Dataset Performance Detail
 - Dataset Details
 - Dataset Performance Detail Trend
 - Dataset Details
 - Dataset Performance Summary Trend
 - Dataset Performance Historical Detail
2. This workspace links to the following set of related workspaces:

- Dataset Performance Summary (See the preceding note (¹) for details.)
- Physical Volume Group Performance
 - Dataset Performance Summary (See the preceding note (¹) for details.)
 - Volume Performance Trend
 - Dataset Performance Summary History
- Device Performance Details
- Cross System Volume Performance
 - Dataset Performance Summary ¹
- Volume Performance Trend
 - Dataset Historic Performance Summary
 - Dataset Performance Historical Detail

Logical Control Unit workspace

This topic describes the main components of the principal workspace for this node.

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See Chapter 20, “Directly monitoring situations,” on page 103.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Logical control unit table view

The table view displays one row of data for each logical control unit configured in the reporting LPAR. The view contains the following information:

- The logical control unit number
- The average delay queue
- The contention rate
- The director port busy percentage
- The control unit busy percentage
- The channel path I/O rate
- The number of attached channels and whether there are DCM channels attached to the control unit
- The high volume utilization rate
- The high volume millisecond response time
- The low read hit percentage
- The low write hit percentage
- The high DFW retry percentage
- The number of SMS and non-SMS volumes
- The total number of volumes

Average delay queue bar graph

This graph contains one bar for each logical control unit. The height of the bar indicates the average number of delayed requests on the control unit header.

Average contention rate bar graph

This graph contains one bar for each logical control unit. The height of the bar indicates the contention rate for the unit.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Logical Control Unit node” on page 37.

Learn how to navigate to workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Logical Control Unit node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “Logical Control Unit workspace” on page 39 for more information on the primary workspace for this node.

| Logical Control Unit workspace | | |
|--------------------------------|-----------|----------------------|
| View | View type | Attribute group |
| Logical Control Unit Report | TABLE | Logical Control Unit |
| Average Delay Queue | BAR | Logical Control Unit |
| Average Contention Rate | BAR | Logical Control Unit |

| LCU Volume Performance workspace | | |
|----------------------------------|-----------|-------------------------|
| View | View type | Attribute group |
| Volume Response Time > 25 | BAR | DASD Volume Performance |
| Volume Performance Report | TABLE | DASD Volume Performance |

| LCU Channel Detail workspace | | |
|------------------------------|-----------|------------------------------------|
| View | View type | Attribute group |
| Control Unit Percent Busy | BAR | Logical Control Unit Channel Paths |
| Complex Channel Percent Busy | BAR | Logical Control Unit Channel Paths |
| Director Port Percent Busy | BAR | Logical Control Unit Channel Paths |
| Logical Control Unit Details | TABLE | Logical Control Unit Channel Paths |

| LCU Volume Cache workspace | | |
|----------------------------|-----------|-----------------|
| View | View type | Attribute group |
| Volume Cache Report | TABLE | Cache Devices |
| Read Hit Percent < 55 | BAR | Cache Devices |
| Write Hit Percent < 30 | BAR | Cache Devices |

| LCU DCM Channels Summary workspace | | |
|------------------------------------|--|--|
|------------------------------------|--|--|

| View | View type | Attribute group |
|-----------------------------|-----------|----------------------|
| DCM Configuration | BAR | Logical Control Unit |
| DCM Channels Percent Busy | BAR | Logical Control Unit |
| DCM Channels I/O per Second | BAR | Logical Control Unit |

| Logical Control Unit Trend workspace | | |
|--------------------------------------|-----------|----------------------|
| View | View type | Attribute group |
| Logical Control Unit Trend Report | TABLE | Logical Control Unit |
| LCU Trend | PLOT | Logical Control Unit |

| LCU Volume Event Performance workspace | | |
|--|-----------|-------------------------|
| View | View type | Attribute group |
| Volume Performance Report | TABLE | DASD Volume Performance |
| Volume Response Time | BAR | DASD Volume Performance |

| LCU HyperPAV Details workspace | | |
|---|-----------|------------------|
| View | View type | Attribute group |
| HyperPAV Report | TABLE | LCU Control Unit |
| I/Os Queued - Alias Unavailable | BAR | LCU Control Unit |
| Percent I/Os Queued - Alias Unavailable | BAR | LCU Control Unit |

| LCU HyperPAV Details Trend workspace | | |
|---|-----------|------------------|
| View | View type | Attribute group |
| HyperPAV Report | TABLE | LCU Control Unit |
| Percent I/Os Queued - Alias Unavailable | PLOT | LCU Control Unit |

Chapter 8. Monitoring tape groups

This topic describes the workspaces that monitor the status of tape groups. You access these workspaces through the Tape Group node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the Tape Group node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Tape Group node to navigate to the workspaces that you see in Table 7. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 7. Tape Group node: bulleted list showing the hierarchy of linked workspaces

| |
|----------------|
| Tape Group |
| • Tape Devices |

Tape Group workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Tape group table view

The table view displays one row of data for each tape group (device or tape library) configured in the reporting LPAR. The view contains the following information:

- The library name or device type
- The library ID number
- The library type (or Non_Library for non-library devices)
- The library device type
- The console name
- The library description
- The number of configured drives, online drives and available drives
- The number of pending mounts and the length of time the current mount has been pending
- The total number of slots in the device and the number currently available

- The number of scratch volumes for 3480 type magnetic tape cartridges (Media1), for 3490E type enhanced capacity magnetic tape cartridges (Media2), for 3590 type enhanced capacity magnetic tape cartridges (Media3), and for 3590 type extended tape cartridges (Media 4)
- The I/O rate
- The number temporary and permanent errors encountered on the device

Mount pend time > 5 minutes bar graph

This graph contains one bar for each tape group which has been waiting more than 5 minutes for a tape mount. The bar height indicates how long the mount has been pending.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Tape Group node” on page 43.

Learn how to navigate to workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Tape Group node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “Tape Group workspace” on page 43 for more information on the primary workspace for this node.

| Tape Group workspace | | |
|-----------------------------|-----------|-----------------|
| View | View type | Attribute group |
| Tape Group Report | TABLE | Tape Group |
| Mount Pend Time > 5 Minutes | BAR | Tape Group |

| Tape Devices workspace | | |
|-----------------------------|-----------|-----------------|
| View | View type | Attribute group |
| Tape Device Report | TABLE | Tape Device |
| Mount Pend Time > 5 Minutes | BAR | Tape Device |

Chapter 9. Monitoring Virtual Tape Subsystems (VTS)

This topic describes the workspaces that monitor the status of virtual tape subsystems. You access these workspaces through the Virtual Tape Subsystems node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the Virtual Tape Subsystems node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Virtual Tape Subsystems node to navigate to the workspaces that you see in Table 8. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 8. Virtual Tape Subsystems node: bulleted list showing the hierarchy of linked workspaces

- | |
|--|
| Virtual Tape Subsystems <ul style="list-style-type: none">• VTS Physical Devices• VTS Virtual Devices• VTS Capacity• VTS Tape Volume Cache• TS7700 Series Cluster<ul style="list-style-type: none">– TS7700 Series Cache<ul style="list-style-type: none">- TS7700 Series Cache Partition<ul style="list-style-type: none">• TS7700 Series Cache Preference Group• VTS Virtual Drive Summary<ul style="list-style-type: none">– VTS Composite Details |
|--|

Virtual Tape Subsystems workspace

This workspace is the principal workspace for this node, and it enables you to monitor both types of virtual tape subsystems, 3494 and 7700.

In a table view, when you right-click the row for a 3494 subsystem, you see standard (non-7700) options. When you select a 7700 subsystem, you see only the link to the TS7700 Series Cluster Workspace.

When you click the Virtual Tape Subsystems node in the Navigator of the Tivoli Enterprise Portal, you see the Virtual Tape Subsystems workspace. The main components of the workspace are as follows:

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Virtual tape subsystems table view

The table view contains one row of data for virtual tape subsystem. The view contains the following information:

- The VTS name and library type
- The number of physical drives available
- The number of virtual devices online
- The average and maximum virtual mount pending times
- The virtual disconnect time
- The volume data (in gigabytes) read and written through the host channels to virtual volumes associated with this VTS
- The recall percent
- Number of scratch volumes, library type, model and serial number for this VTS

Volume recall % > 20 bar graph

This graph contains one bar for each VTS for which the percentage of virtual mounts to actual physical mounts exceeds 20%. The bar height indicates the recall percentage.

Average virtual mount pending time > 300 bar graph

This graph contains one bar for each VTS for which the average time spent waiting for a virtual mount exceeded 300 seconds. The height of the bar indicates the actual time spent waiting.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Virtual Tape Subsystems node” on page 45.

Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Virtual Tape Subsystems node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “Virtual Tape Subsystems workspace” on page 45 for more additional information on the primary workspace for this node.

| Virtual Tape Subsystems workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Virtual Tape Subsystems Report | TABLE | VTS Overview |
| Volume Recall % > 20 | BAR | VTS Overview |
| Average Virtual Mount Pend Time > 300 | BAR | VTS Overview |

| VTS Physical Devices workspace | | |
|---------------------------------------|------------------|------------------------|
| View | View type | Attribute group |
| Mount Pend Time | BAR | VTS Physical Devices |

| | | |
|------------------------|-----|----------------------|
| Concurrent Mounts | BAR | VTS Physical Devices |
| Physical Drives Status | BAR | VTS Physical Devices |
| Physical Mount Reason | BAR | VTS Physical Devices |

| VTS Virtual Devices workspace | | |
|--------------------------------------|------------------|------------------------|
| View | View type | Attribute group |
| Concurrent Mounts | BAR | VTS Virtual Devices |
| Mount Pend Time | BAR | VTS Virtual Devices |
| Mount Use Time | BAR | VTS Virtual Devices |
| Virtual Device Activity Report | TABLE | VTS Virtual Devices |
| Recall Mount Time | BAR | VTS Virtual Devices |

| VTS Capacity workspace | | |
|-------------------------------|------------------|------------------------|
| View | View type | Attribute group |
| VTS Capacity Report | TABLE | VTS Capacity |
| Host Channel Activity | BAR | VTS Capacity |
| VTS Channel Activity | BAR | VTS Capacity |
| Physical Tape Usage | BAR | VTS Capacity |

| VTS Tape Volume Cache workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| VTS Tape Volume Cache Report | TABLE | VTS Cache |
| Throttling Delay | BAR | VTS Cache |
| TVC Volume Age | BAR | VTS Cache |
| Virtual Mounts | BAR | VTS Cache |
| Throttling Percent | PIE | VTS Cache |

| VTS Composite Details workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| VTS Controller Report | TABLE | VTS Composite Library |
| Mounts by Unit | BAR | VTS Virtual Devices |
| Logical Volumes to be Copied | BAR | VTS Virtual Devices |
| Data to Copy | BAR | VTS Virtual Devices |

| VTS Virtual Drive Summary workspace | | |
|--|------------------|-----------------------------|
| View | View type | Attribute group |
| Virtual Tape Subsystems Virtual Tape Drives | TABLE | VTS Virtual Devices Summary |
| Tape Drive Response Time | BAR | VTS Virtual Devices Summary |

Workspaces related to TS7700-series VTS

The following attribute groups supply data to the workspaces that are related to TS7700-series VTS:

- “VTS TS7700-series Cache Container” on page 244
- “VTS TS7700-series Cache Partition Container” on page 245
- “VTS TS7700-series Cache Preference Group” on page 246
- “VTS TS7700-series Cluster” on page 247

| TS7700 Series Cluster workspace | | | |
|---------------------------------|-----------|-----------------|--|
| View | View type | Attribute group | Description |
| Cluster Summary | TABLE | VTS CLUSTER | Contains basic information about a cluster in a TS7700 tape virtualization system, such as ID and machine type and model. The view also contains metrics that are derived from standard sources, such as the virtualization (vNode) or hierarchical data storage management (hNode). |
| Volume Recall Percent | BAR | VTS CLUSTER | Charts the following metric for each cluster in the grid: Volume Recall Percent |
| Average Virtual Mount Pend Time | BAR | VTS CLUSTER | Charts the following metric for each cluster in the grid: Average Virtual Mount Pend Time. |

| TS7700 Series Cache workspace | | | |
|-------------------------------|-----------|-----------------|--|
| View | View type | Attribute group | Description |
| Cache Throttling Percentages | BAR | VTS TPVOLC | Charts the following two metrics: Percent Write Overrun Throttle and Percent Copy Throttle. |
| Cache Throttling Averages | BAR | VTS TPVOLC | Charts the following three metrics: Average Write Overrun Throttle, Average Copy Throttle, and Average Overall Throttle. |
| Tape Volume Cache Summary | TABLE | VTS TPVOLC | Contains basic information about a cache in a TS7700 tape virtualization system, based on metrics that are derived from the “HSM - Cache Container” record of the “Hnode HSM Historical Record.” |

| TS7700 Series Cache Partition workspace | | | |
|---|-----------|-----------------|-------------|
| View | View type | Attribute group | Description |

| | | | |
|--------------------------------|-------|-------------|---|
| Average Cache Miss Mount Times | BAR | VTS CACHEPC | Charts the following metric for each of the 8 partitions: Average Cache Miss Mount Time. |
| Cache Miss Mounts | BAR | VTS CACHEPC | Charts the following metric for each of the 8 partitions: Cache Miss Mounts. |
| Cache Partition Summary | TABLE | VTS CACHEPC | Contains basic information about a cache partition in a TS7700 tape virtualization system. All attributes in this table are derived from the "HSM - Cache - Partition Container" record of the "Hnode HSM Historical Record." |

| TS7700 Series Cache Preference Group workspace | | | |
|---|------------------|------------------------|--|
| View | View type | Attribute group | Description |
| Average Cache Age | BAR | VTS CHPGRP | Charts the following three metrics: 4 Hour Average Cache Age, 48 Hour Average Cache Age, and 35 Day Average Cache Age. |
| Number of Volumes Migrated | BAR | VTS CHPGRP | Charts the following three metrics: Volumes Migrated Last 4 Hours, Volumes Migrated Last 48 Hours, and Volumes Migrated Last 35 Days. |
| Cache Preference Group Summary | TABLE | VTS CHPGRP | Contains basic information about a cache preference group in a TS7700 tape virtualization system, based on metrics that are derived from the HSM HSM - Cache - Partition - Preference Group Container record of the "Hnode HSM Historical Record." |

Chapter 10. Monitoring SMS storage groups performance

This topic describes the workspaces that monitor the status of SMS storage groups performance. You access these workspaces through the SMS Storage Groups Performance node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the SMS Storage Groups Performance node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the SMS Storage Groups Performance node to navigate to the workspaces that you see in Table 9 on page 52. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 9. SMS Storage Groups Performance node: bulleted list showing the hierarchy of linked workspaces

| |
|--|
| <p>SMS Storage Groups Performance</p> <ul style="list-style-type: none"> • SMS Storage Group Volume Performance <ul style="list-style-type: none"> – Dataset Performance Summary ¹ – Volume Performance Trend <ul style="list-style-type: none"> - Dataset Historic Performance Summary <ul style="list-style-type: none"> • Dataset Performance Historical Detail – Physical Volume Group Performance <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Volume Performance Trend <ul style="list-style-type: none"> - Dataset Historic Performance Summary <ul style="list-style-type: none"> • Dataset Performance Historical Detail – Device Performance Details – Cross System Volume Performance <ul style="list-style-type: none"> - Dataset Performance Summary ¹ • SMS Storage Group Volume Cache <ul style="list-style-type: none"> – Dataset Performance Summary ¹ – Physical Volume Group Cache <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Device Performance Details – Cross System Volume Performance <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Volume Cache Trend • Cross System Group Summary <ul style="list-style-type: none"> – Cross System Volume Performance <ul style="list-style-type: none"> - Dataset Performance Summary ¹ • Storage Group Performance Trend <ul style="list-style-type: none"> – SMS Storage Group Historical Volume Performance <ul style="list-style-type: none"> - Dataset Historic Performance Summary <ul style="list-style-type: none"> • Dataset Performance Historical Detail – SMS Storage Group Historic Volume Cache Performance <ul style="list-style-type: none"> - Dataset Historic Performance Summary <ul style="list-style-type: none"> • Dataset Performance Historical Detail <p>These workspaces list the volumes with the least favorable readings for the attribute within the SMS storage group selected for navigation:</p> <ul style="list-style-type: none"> • Highest Volume Response Time (Top 10) ² • Highest Volume MPL (Top 10) ² • Highest Volume I/O Rate (Top 10) ² • Highest Volume Busy Percent (Top 10) ² • Lowest Volume Connect Percent (Top 10) ² • Highest Volume Group Response Time (Top 10) ² • Highest Volume Group MPL (Top 10) ² • Highest Volume Group I/O Rate (Top 10) ² • Highest Volume Group Busy Percent (Top 10) ² • Lowest Volume Group Connect Percent (Top 10) ² |
|--|

Footnotes for Table 9:

1. This group of workspaces is organized in the following nested pattern:
 - Dataset Performance Summary
 - Dataset Performance Detail
 - Dataset Details
 - Dataset Performance Detail Trend
 - Dataset Details
 - Dataset Performance Summary Trend
 - Dataset Performance Historical Detail
2. This workspace links to the following set of related workspaces:
 - Dataset Performance Summary (See the preceding note (¹) for details.)
 - Physical Volume Group Performance

- Dataset Performance Summary ¹
- Device Performance Details
- Cross System Volume Performance
 - Dataset Performance Summary ¹
- Volume Performance Trend
 - Dataset Historic Performance Summary
 - Dataset Performance Historical Detail

Access the following linked workspaces by selecting the SMS Storage Groups Performance node in the Navigator. Right-click the node. Select the workspace that you want to access in the pop-up menu.

These workspaces list the volumes with the least favorable readings for the attribute across the z/OS image.

- Highest Volume Response Time *
- Highest Volume MPL *
- Highest Volume I/O Rate *
- Highest Volume Busy Percent *
- Lowest Volume Connect Percent *

Each of the workspaces in the preceding bulleted list link to the following set of related workspaces:

- Dataset Performance Summary
 - Dataset Performance Detail
 - Dataset Details
 - Dataset Performance Detail Trend
 - Dataset Details
 - Dataset Performance Summary Trend
 - Dataset Performance Historical Detail
- Physical Volume Group Performance
 - Dataset Performance Summary ¹
- Device Performance Details
- Cross System Volume Performance
 - Dataset Performance Summary ¹
- Volume Performance Trend
 - Dataset Historic Performance Summary
 - Dataset Performance Historical Detail

SMS Storage Groups Performance workspace

This topic describes the main components of the principal workspace for this node.

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See Chapter 20, “Directly monitoring situations,” on page 103.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

SMS storage groups performance table view

The table view contains one row of data for each SMS storage group defined in the reporting LPAR. The view contains the following information:

- The name of the group
- The type

- The SMS status
- The number of volumes in the group
- The current highest response time for any volume in the group
- The current high busy percent for any volume in the group
- The device multi-processing level
- The low read and write hit percents among all volumes in the group
- The highest DFW retry percentage in the group

High volume response time > 25 bar graph

This graph contains one bar for each SMS storage group with a highest volume response time that exceeds 25 milliseconds. The height of the bar indicates the actual response time for that volume.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the SMS Storage Groups Performance node” on page 51.

Learn how to navigate to workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the SMS Storage Groups Performance node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “SMS Storage Groups Performance workspace” on page 53 for more information on the primary workspace for this node.

| SMS Storage Groups Performance workspace | | |
|--|-----------|----------------------|
| View | View type | Attribute group |
| High Volume Response Time > 25 | BAR | Volume Group Summary |
| SMS Storage Group Performance Report | TABLE | Volume Group Summary |

| SMS Storage Group Volume Performance workspace | | |
|--|-----------|-------------------------|
| View | View type | Attribute group |
| Volume Response Time > 25 | BAR | DASD Volume Performance |
| Volume Performance Report | TABLE | DASD Volume Performance |

| Dataset Performance Summary workspace | | |
|---------------------------------------|-----------|---------------------|
| View | View type | Attribute group |
| Dataset Performance Summary Report | TABLE | Dataset Performance |

| | | |
|---|-------|---------------------|
| Datasets with MSR > Storage Class Objective | TABLE | Dataset Performance |
|---|-------|---------------------|

| Dataset Performance Detail workspace | | |
|---|-----------|---------------------|
| View | View type | Attribute group |
| Dataset Performance Detail Report | TABLE | Dataset Performance |
| Component Percent of Dataset MSR for All Jobs | PIE | Dataset Performance |

| SMS Storage Group Volume Cache workspace | | |
|--|-----------|-----------------|
| View | View type | Attribute group |
| Read Hit Percent < 55 | BAR | Cache Devices |
| Write Hit Percent < 30 | BAR | Cache Devices |
| Volume Cache Report | TABLE | Cache Devices |

| Volume Performance Trend workspace | | |
|------------------------------------|-----------|-------------------------|
| View | View type | Attribute group |
| Volume Performance Trend Report | TABLE | DASD Volume Performance |
| Response Time Analysis | PLOT | DASD Volume Performance |

| Physical Volume Group Performance workspace | | |
|---|-----------|---------------------|
| View | View type | Attribute group |
| Volume Response Time > 25 | BAR | DASD Physical Group |
| Physical Volume Group Performance Report | TABLE | DASD Physical Group |

| Physical Volume Group Cache workspace | | |
|---------------------------------------|-----------|---------------------|
| View | View type | Attribute group |
| Read Hit Percent < 55 | BAR | DASD Physical Group |
| Write Hit Percent < 30 | BAR | DASD Physical Group |
| Physical Volume Group Cache Report | TABLE | DASD Physical Group |

| Dataset Details workspace | | |
|-----------------------------|-----------|---------------------|
| View | View type | Attribute group |
| Dataset Space Attributes | TABLE | Dataset Space |
| Dataset Users | TABLE | Dataset Users |
| Dataset SMS Constructs | TABLE | Dataset SMS |
| Dataset Performance Summary | TABLE | Dataset Performance |

| Device Performance Details workspace | | |
|---|-----------|-----------------|
| View | View type | Attribute group |
| Volume Cache Performance Details Report | TABLE | Cache Devices |

| | | |
|-----------------------------------|-------|-------------------------|
| Volume Status | TABLE | Volume Status |
| Volume Details | TABLE | DASD Volume Performance |
| Volume Performance Details Report | TABLE | DASD Volume Performance |
| Volume Users | TABLE | Volume Users |

| Device Performance Details workspace | | |
|---|-----------|-------------------------|
| View | View type | Attribute group |
| Volume Cache Performance Details Report | TABLE | Cache Devices |
| Volume Status | TABLE | Volume Status |
| Volume Details | TABLE | DASD Volume Performance |
| Volume Performance Details Report | TABLE | DASD Volume Performance |
| Volume Users | TABLE | Volume Users |

| Cross System Group Summary workspace | | |
|--------------------------------------|-----------|--------------------------|
| View | View type | Attribute group |
| Device Contention Index | BAR | Cross System Vol Summary |
| Cross System Performance Report | TABLE | Cross System Vol Summary |

| Cross System Volume Performance workspace | | |
|---|-----------|----------------------|
| View | View type | Attribute group |
| Response Time | BAR | Cross System Volumes |
| Percent Busy | BAR | Cross System Volumes |
| Cross System Volume Performance Report | TABLE | Cross System Volumes |

| SMS Storage Group Performance Trend workspace | | |
|---|-----------|----------------------|
| View | View type | Attribute group |
| High Volume Response | BAR | Volume Group Summary |
| SMS Storage Group Performance Trend Report | TABLE | Volume Group Summary |

| Group Volume Historical Performance workspace | | |
|---|-----------|-------------------------|
| View | View type | Attribute group |
| Volume Performance Report | TABLE | DASD Volume Performance |
| Volumes with Response Time > 25 | BAR | DASD Volume Performance |

| Dataset Performance Summary Trend workspace | | |
|---|-----------|---------------------|
| View | View type | Attribute group |
| Dataset Performance Summary Trend | TABLE | Dataset Performance |
| Response Time Analysis | PLOT | Dataset Performance |

| Dataset Historic Performance Summary workspace | | |
|--|--|--|
|--|--|--|

| View | View type | Attribute group |
|---|-----------|---------------------|
| Dataset Performance Summary History Report | TABLE | Dataset Performance |
| Datasets with MSR > Storage Class Objective History | TABLE | Dataset Performance |

| Dataset Performance Historical Detail workspace | | |
|---|-----------|---------------------|
| View | View type | Attribute group |
| Dataset Performance Detail Report | TABLE | Dataset Performance |
| Component Percent of Dataset MSR for All Jobs | PIE | Dataset Performance |

| Dataset Performance Detail Trend workspace | | |
|--|-----------|---------------------|
| View | View type | Attribute group |
| Dataset Performance Application Trend | TABLE | Dataset Performance |
| Response Time Analysis | PLOT | Dataset Performance |

| Dataset Performance Event workspace | | |
|-------------------------------------|-----------|---------------------|
| View | View type | Attribute group |
| Dataset Performance Summary Report | TABLE | Dataset Performance |
| Dataset Response Time | BAR | Dataset Performance |

| Highest Volume Response Time workspace | | |
|--|-----------|-------------------------|
| View | View type | Attribute group |
| Highest Volume Response Time Report | TABLE | DASD Volume Performance |
| Volume Response Time | BAR | DASD Volume Performance |

| Highest Volume MPL workspace | | |
|------------------------------|-----------|-------------------------|
| View | View type | Attribute group |
| Highest Volume MPL Report | TABLE | DASD Volume Performance |
| Volume MPL | BAR | DASD Volume Performance |
| Volume Response Time | BAR | DASD Volume Performance |

| Highest Volume I/O Rate workspace | | |
|---|-----------|-------------------------|
| View | View type | Attribute group |
| Highest Volume I/O Report | TABLE | DASD Volume Performance |
| Provide summary performance information for the datasets on a specified volume. | BAR | DASD Volume Performance |
| Volume Response Time | BAR | DASD Volume Performance |

| Highest Volume Busy Percent workspace | | |
|--|------------------|-------------------------|
| View | View type | Attribute group |
| Highest Volume Busy Percent Report | TABLE | DASD Volume Performance |
| Volume Busy Percent | BAR | DASD Volume Performance |
| Volume Response Time | BAR | DASD Volume Performance |

| Lowest Volume Connect Percent workspace | | |
|--|------------------|-------------------------|
| View | View type | Attribute group |
| Lowest Volume Connect Time Percent Report | TABLE | DASD Volume Performance |
| Volume Connect Time Percent | BAR | DASD Volume Performance |
| Volume Response Time | BAR | DASD Volume Performance |

| Highest Volume Group Response Time workspace | | |
|---|------------------|-------------------------|
| View | View type | Attribute group |
| Highest Volume Performance Report | TABLE | DASD Volume Performance |
| Volume Response Time | BAR | DASD Volume Performance |

| Highest Volume Group I/O Rate workspace | | |
|--|------------------|-------------------------|
| View | View type | Attribute group |
| Highest Volume I/O Report | TABLE | DASD Volume Performance |
| Volume I/O Rate | BAR | DASD Volume Performance |
| Volume Response Time | BAR | DASD Volume Performance |

| Highest Volume Group MPL workspace | | |
|---|------------------|-------------------------|
| View | View type | Attribute group |
| Highest Volume MPL Report | TABLE | DASD Volume Performance |
| Volume MPL | BAR | DASD Volume Performance |
| Volume Response Time | BAR | DASD Volume Performance |

| Highest Volume Group Busy Percent workspace | | |
|--|------------------|-------------------------|
| View | View type | Attribute group |
| Highest Volume Busy Percent Report | TABLE | DASD Volume Performance |
| Volume Busy Percent | BAR | DASD Volume Performance |
| Volume Response Time | BAR | DASD Volume Performance |

| Lowest Volume Group Connect Percent workspace | | |
|--|------------------|-------------------------|
| View | View type | Attribute group |
| Volume Connect Time Percent | BAR | DASD Volume Performance |
| Lowest Volume Connect Time Percent Report | TABLE | DASD Volume Performance |
| Volume Response Time | BAR | DASD Volume Performance |

Chapter 11. Monitoring SMS storage groups space

This topic describes the workspaces that monitor the status of SMS Storage Groups Space. You access these workspaces through the SMS Storage Groups Space node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the SMS Storage Groups Space node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the SMS Storage Groups Space node to navigate to the workspaces that you see in Table 10. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 10. SMS Storage Groups Space node: bulleted list showing the hierarchy of linked workspaces

| |
|---|
| SMS Storage Groups Space |
| • SMS Storage Group Volume Space |
| – Dataset Space Summary |
| – Dataset Details |
| – Device Space Details |
| – Volume Space Trend |
| • SMS Storage Group Space Trend |
| – SMS Storage Group Volume Space History |
| These workspaces list the volumes with the least favorable readings for the attribute within the SMS Storage Group: |
| • Lowest Volume Free Space |
| • Lowest Volume Free Space Percent |
| • Highest Volume Fragmentation Index |
| • Lowest Volume Group Free Space |
| • Lowest Volume Group Free Space Percent |
| • Highest Volume Group Fragmentation Index |

Access the following linked workspaces by selecting the SMS Storage Groups Space node in the Navigator. Right-click the node. Select the workspace that you want to access in the pop-up menu.

These workspaces list the volumes with the least favorable readings for the attribute within the z/OS image.

- Lowest Volume Free Space
- Lowest Volume Free Space Percent
- Highest Volume Fragmentation Index

Each of the workspaces in the preceding bulleted list link to the following set of related workspaces:

- Dataset Space Summary
 - Dataset Details

- Device Space Details
- Volume Space Trend

SMS Storage Groups Space workspace

This topic describes the main components of the principal workspace for this node.

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See Chapter 20, “Directly monitoring situations,” on page 103.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

SMS storage groups space table view

The table view contains one row of data for each SMS storage group. The view contains the following information:

- The group name
- The storage group status
- The type and whether or not it is enabled
- The number of volumes in the group
- The status of the VTOC index
- The lowest freespace percent of all volumes in the group
- The highest fragmentation index of all volumes
- The largest free extent
- The total free space and the percent of free space for the group
- The total amount of space for the entire storage group

Storage groups with free space % < 10 bar graph

This graph contains one bar for each storage group with a combined free space percentage that falls below 10%.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the SMS Storage Groups Space node” on page 59.

Learn how to navigate to workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the SMS Storage Groups Space node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “SMS Storage Groups Space workspace” on page 60 for more information on the primary workspace for this node.

| SMS Storage Groups Space workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| SMS Storage Groups Space Summary | TABLE | Volume Group Summary |
| Storage Groups with Free Space % < 10 | BAR | Volume Group Summary |

| SMS Storage Group Volume Space workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Volume Space Report | TABLE | DASD Volume Space |
| Volumes with Free Space % < 10 | BAR | DASD Volume Space |
| Volumes with Fragmentation Index > 650 | BAR | DASD Volume Space |

| Dataset Space Summary workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Dataset Space Summary Report | TABLE | Dataset Space |
| Volume Space Allocation | PIE | DASD Volume Space |

| Device Space Details workspace | | |
|---------------------------------------|------------------|-------------------------|
| View | View type | Attribute group |
| Volume Free Space Details | TABLE | DASD Volume Space |
| Track-Managed Free Space Details | TABLE | DASD Volume Space |
| Volume Allocated Space Details | TABLE | DASD Volume Space |
| Volume Status | TABLE | Volume Status |
| Volume Details | TABLE | DASD Volume Performance |

| SMS Storage Group Space Trend workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| SMS Storage Group Space Trend | TABLE | Volume Group Summary |
| Megabytes of Free Space | BAR | Volume Group Summary |

| Volume Group Space History workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Volume Space Report | TABLE | DASD Volume Space |
| Volumes with Free Space % < 10 | BAR | DASD Volume Space |
| Volumes with Fragmentation Index > 650 | BAR | DASD Volume Space |

| Storage Group Percent Free workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Volume Free Space Percent | BAR | DASD Volume Space |

| | | |
|----------------------------|-------|-------------------|
| Volume Space Report | TABLE | DASD Volume Space |
| Volume Fragmentation Index | BAR | DASD Volume Space |

| Storage Group Percent Free Trend workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Storage Group Free Space Percent | PLOT | Volume Group Summary |
| SMS Storage Group Space Trend | TABLE | Volume Group Summary |

| Dataset Space Event workspace | | |
|--------------------------------------|------------------|------------------------|
| View | View type | Attribute group |
| Volume Space Allocation | PIE | DASD Volume Space |
| Dataset Space Summary Report | TABLE | Dataset Space |

| Volume Space Trend workspace | | |
|-------------------------------------|------------------|------------------------|
| View | View type | Attribute group |
| Volume Space Trend | TABLE | DASD Volume Space |
| Percent Free Space | PLOT | DASD Volume Space |

| Lowest Volume Freespace workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Volume Free Space Report | TABLE | DASD Volume Space |
| Free Space | BAR | DASD Volume Space |
| Free Space Percent | BAR | DASD Volume Space |

| Lowest Volume Free Space Percent workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Volume Free Space Percent Report | TABLE | DASD Volume Space |
| Free Space Percent | BAR | DASD Volume Space |
| Free Space | BAR | DASD Volume Space |

| Highest Volume Fragmentation Index workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Volume Fragmentation Index Report | TABLE | DASD Volume Space |
| Fragmentation Index | BAR | DASD Volume Space |
| Free Space | BAR | DASD Volume Space |

| Lowest Volume Group Freespace workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Volume Free Space Report | TABLE | DASD Volume Space |
| Free Space | BAR | DASD Volume Space |
| Free Space Percent | BAR | DASD Volume Space |

| Lowest Volume Group Free Space Percent workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Volume Free Space Percent Report | TABLE | DASD Volume Space |
| Free Space Percent | BAR | DASD Volume Space |
| Free Space | BAR | DASD Volume Space |

| Highest Volume Group Fragmentation Index workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Volume Fragmentation Index Report | TABLE | DASD Volume Space |
| Fragmentation Index | BAR | DASD Volume Space |
| Free Space | BAR | DASD Volume Space |

Chapter 12. Monitoring user DASD groups performance

This topic describes the workspaces that monitor the status of user DASD groups performance. You access these workspaces through the User DASD Groups Performance node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the User DASD Groups Performance node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the User DASD Groups Performance node to navigate to the workspaces that you see in Table 11 on page 66. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 11. User DASD Groups Performance node: bulleted list showing the hierarchy of linked workspaces

| |
|--|
| <p>User DASD Groups Performance</p> <ul style="list-style-type: none"> • User DASD Group Volume Performance <ul style="list-style-type: none"> – Dataset Performance Summary ¹ – Physical Volume Group Performance <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Device Performance Details – Cross System Volume Performance – Volume Performance Trend <ul style="list-style-type: none"> - Dataset Historic Performance Summary <ul style="list-style-type: none"> • Dataset Performance Historical Detail • User DASD Group Volume Cache <ul style="list-style-type: none"> – Dataset Performance Summary ¹ – Physical Volume Group Cache <ul style="list-style-type: none"> - Dataset Performance Summary ¹ – Device Performance Details – Cross System Volume Performance – Volume Cache Trend • Cross System Group Summary <ul style="list-style-type: none"> – Cross System Volume Performance • User DASD Performance Trend <ul style="list-style-type: none"> – User DASD Group Historic Volume Performance <ul style="list-style-type: none"> - Dataset Historic Performance Summary <ul style="list-style-type: none"> • Dataset Performance Historical Detail – User DASD Group Historic Volume Cache Performance <ul style="list-style-type: none"> - Dataset Historic Performance Summary <ul style="list-style-type: none"> • Dataset Performance Historical Detail <p>The following workspaces display the volumes that have the least favorable readings for the attribute within the user DASD group selected for navigation:</p> <ul style="list-style-type: none"> • Highest User DASD Group Volumes Response Time (Top 10) ² • Highest User DASD Group Volumes MPL (Top 10) ² • Highest User DASD Group Volumes I/O Rate (Top 10) ² • Highest User DASD Group Volumes Busy Percent (Top 10) ² • Highest User DASD Group Volumes Connect Percent (Top 10) ² |
|--|

Footnotes for Table 11:

1. This group of workspaces is organized in the following nested pattern:
 - Dataset Performance Summary
 - Dataset Performance Detail
 - Dataset Details
 - Dataset Performance Detail Trend
 - Dataset Details
 - Dataset Performance Summary Trend
 - Dataset Performance Historical Detail
2. This workspace links to the following set of related workspaces:
 - Dataset Performance Summary (See the preceding note ⁽¹⁾ for details.)
 - Physical Volume Group Performance
 - Dataset Performance Summary (See the preceding note ⁽¹⁾ for details.)
 - Volume Performance Trend
 - Dataset Performance Summary History
 - Device Performance Details
 - Cross System Volume Performance
 - Volume Performance Trend
 - Dataset Historic Performance Summary
 - Dataset Performance Historical Detail

User DASD Groups Performance workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

User DASD groups performance table view

The table view contains one row of data for each defined user DASD group. Each row lists:

- The name of the group
- The number of SMS and non-SMS volumes
- The total number of volumes in the group
- The current highest response time for any volume in the group
- The current high busy percent for any volume in the group
- The device multi-processing level
- The low read and write hit percentages among all volumes in the group
- The highest DFW retry percentage in the group

High volume response time > 25 bar graph

This graph contains one bar for each user DASD group with a highest volume response time that exceeds 25 milliseconds. The height of the bar indicates the actual response time for that volume.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the User DASD Groups Performance node” on page 65.

Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the User DASD Groups Performance node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “User DASD Groups Performance workspace” for more information on the primary workspace for this node.

| User DASD Groups Performance workspace | | |
|--|-----------|----------------------|
| View | View type | Attribute group |
| User DASD Groups Performance Report | TABLE | Volume Group Summary |

| | | |
|--------------------------------|-----|----------------------|
| High Volume Response Time > 25 | BAR | Volume Group Summary |
|--------------------------------|-----|----------------------|

| User DASD Group Volume Performance workspace | | |
|--|-----------|-------------------------|
| View | View type | Attribute group |
| Volume Performance Report | TABLE | DASD Volume Performance |
| Volume Response Time > 25 | BAR | DASD Volume Performance |

| User DASD Group Volume Cache workspace | | |
|--|-----------|-----------------|
| View | View type | Attribute group |
| Volume Cache Report | TABLE | Cache Devices |
| Read Hit Percent < 55 | BAR | Cache Devices |
| Write Hit Percent < 30 | BAR | Cache Devices |

| User DASD Group Performance Trend workspace | | |
|---|-----------|----------------------|
| View | View type | Attribute group |
| User DASD Groups Performance Trend Report | TABLE | Volume Group Summary |
| High Volume Response Time | BAR | Volume Group Summary |

Chapter 13. Monitoring user DASD groups space

This topic describes the workspaces that monitor the status of user DASD groups space. You access these workspaces through the User DASD Groups Space node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the User DASD Groups Space node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the User DASD Groups Space node to navigate to the workspaces that you see in Table 12. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 12. User DASD Groups Space node: bulleted list showing the hierarchy of linked workspaces

| |
|---|
| User DASD Groups Space |
| • User DASD Group Volume Space |
| – Dataset Space Summary |
| – Dataset Details |
| – Device Space Details |
| • User DASD Group Space Trend |
| – User DASD Group Volume Historical Space |
| The following workspaces display the volumes that have the least favorable readings for the attribute within the User DASD Group: |
| • Lowest User DASD Group Volumes Free Space (Top 10) |
| • Lowest User DASD Group Volumes Free Space Percent (Top 10) |
| • Highest User DASD Group Volumes Fragmentation Index (Top 10) |

User DASD Groups Space workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

User DASD groups space table view

The table view contains one row of data for each user DASD group. The view contains the following information:

- The group name

- The number of SMS and non-SMS volumes and the total number of volumes in the group
- The status of the VTOC index
- The lowest freespace percent of all volumes in the group
- The highest fragmentation index of all volumes
- The largest free extent
- The total free space
- The percentage of free space for the group and the total amount of space for the entire storage group

User DASD groups with free space % < 10 bar graph

This graph contains one bar for each user DASD group with a combined free space percentage that falls below 10%.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the User DASD Groups Space node” on page 69.

Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the User DASD Groups Space node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “User DASD Groups Space workspace” on page 69 for more information on the primary workspace for this node.

| User DASD Groups Space workspace | | |
|---|-----------|----------------------|
| View | View type | Attribute group |
| User DASD Groups Space Report | TABLE | Volume Group Summary |
| User DASD Groups with Free Space % < 10 | BAR | Volume Group Summary |

| User DASD Group Volume Space workspace | | |
|--|-----------|-------------------|
| View | View type | Attribute group |
| Volume Space Report | TABLE | DASD Volume Space |
| Volumes with Free Space % < 10 | BAR | DASD Volume Space |
| Volumes with Fragmentation Index > 650 | BAR | DASD Volume Space |

| User DASD Group Space Trend workspace |
|---------------------------------------|
|---------------------------------------|

| View | View type | Attribute group |
|-------------------------------------|------------------|------------------------|
| User DASD Groups Space Trend Report | TABLE | Volume Group Summary |
| Megabytes of Free Space | BAR | Volume Group Summary |

| DFSMSHsm Status workspace | | |
|----------------------------------|------------------|------------------------|
| View | View type | Attribute group |
| HSM Functions Status Report | TABLE | HSM Function Summary |
| Waiting Requests | BAR | HSM Function Summary |
| HSM Control Dataset Report | TABLE | HSM CDS |
| HSM Status Report | TABLE | HSM Status |

Chapter 14. Monitoring DFSMShsm status

This topic describes the workspaces that monitor the status of DFSMShsm activity. You access these workspaces through the DFSMShsm Status node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the DFSMShsm Status node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the DFSMShsm Status node to navigate to the workspaces that you see in Table 13. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 13. DFSMShsm Status node: bulleted list showing the hierarchy of linked workspaces

| |
|-----------------------------|
| DFSMShsm Status |
| • DFSMShsm Request Details |
| – DFSMShsm Function Details |
| • DFSMShsm Storage Usage |
| • DFSMShsm Host Details |

DFSMShsm Status workspace

This topic describes the main components of the principal workspace for this node.

When you click the DFSMShsm Status node in the Navigator of the Tivoli Enterprise Portal, you see the DFSMShsm Status workspace. This workspace monitors DFSMShsm activity for the entire HSM subsystem, including utilization of volumes and CDS space. The main components of the workspace are as follows:

The direct situation analysis feature is available for this workspace. The feature can help you analyze and understand a problem that is displayed in this workspace. See Chapter 20, “Directly monitoring situations,” on page 103.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

| DFSMShsm Status workspace | |
|----------------------------------|-----------|
| View | View type |
| DFSMShsm Functions Summary | TABLE |
| DFSMShsm Control Data Set Report | TABLE |
| DFSMShsm Waiting Requests | BAR |

| | |
|--------------------------|-------|
| DFSMSHsm Function Status | TABLE |
| DFSMSHsm Status Report | TABLE |

DFSMSHsm Functions Status Report

The tabular report in the top center portion of this workspace contains one row for each DFSMSHsm function. Each row displays:

- The status of that function for the reporting LPAR
- The number of data set requests and volume requests for that function and how many are active and how many waiting

DFSMSHsm Control Dataset Report

The tabular report covering the entire width of the workspace in the vertical center contains one row for each of the DFSMSHsm control data sets. Each row displays:

- The DD name of the data set
- The number of extents for the data and index components
- The percent of Available Space for the data and index components of the data set
- The percent of Free Space for the data and index components of the data set
- The total size of the data and the index components of the data set

DFSMSHsm Status Report

The tabular report covering the entire width of the workspace at the bottom of the screen contains a single row of data describing DFSMSHsm status in the reporting LPAR. This row indicates:

- Whether DFSMSHsm is active
- The address space ID
- The version of DFHSM
- DFSMSHsm start time
- The interval start time
- The number of DFSMSHsm requests during this collection interval
- The CPU usage during this interval
- The number of requests that abended during this interval
- The number of volume migrations, backups and the number of migration and backup volumes recycled during this interval

Waiting requests bar graph

This graph contains one bar for each DFSMSHsm function which has at least one waiting request. The height of the bar indicates the total number of waiting requests for this function.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the DFSMSHsm Status node” on page 73.

Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the DFSMSHsm Status node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “DFSMSHsm Status workspace” on page 73 for more information on the primary workspace for this node.

| DFSMSHsm Status workspace | | |
|----------------------------------|-----------|--------------------------|
| View | View type | Attribute group |
| DFSMSHsm Functions Summary | TABLE | HSM Function Summary |
| DFSMSHsm Control Data Set Report | TABLE | HSM CDS |
| DFSMSHsm Waiting Requests | BAR | HSM Host Request Summary |
| DFSMSHsm Function Status | TABLE | HSM Host Function Status |
| DFSMSHsm Status Report | TABLE | HSM Status |

| DFSMSHsm Request Details workspace | | |
|------------------------------------|-----------|-------------------------|
| View | View type | Attribute group |
| DFSMSHsm Function Statistics | TABLE | HSM Requests |
| DFSMSHsm Requests | TABLE | HSM Function Statistics |

| DFSMSHsm Storage Usage workspace | | |
|-------------------------------------|-----------|---------------------|
| View | View type | Attribute group |
| DFSMSHsm Host Storage Utilization | TABLE | HSM Private Storage |
| DFSMSHsm Host Frame Usage | BAR | HSM Private Storage |
| DFSMSHsm Common Storage | TABLE | HSM Common Storage |
| DFSMSHsm Common Storage Utilization | BAR | HSM Common Storage |

| DFSMSHsm Host Details workspace | | |
|---------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSHsm Host Details Report | TABLE | HSM Host Status |
| DFSMSHsm Requests | BAR | HSM Host Status |
| DFSMSHsm Failed Requests | BAR | HSM Host Status |

| DFSMSHsm Function Details workspace | | |
|-------------------------------------|-----------|----------------------|
| View | View type | Attribute group |
| DFSMSHsm Host Function Details | TABLE | HSM Function Details |
| DFSMSHsm Function Response Time | BAR | HSM Function Details |

Chapter 15. Monitoring DFSMSrmm status

This topic describes the workspaces that monitor the status of DFSMSrmm resources. You access these workspaces through the DFSMSrmm Status node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Related information

Chapter 28, “Working with the DFSMSrmm Status feature,” on page 323

This topic describes key concepts and capabilities of the DFSMSrmm Status feature.

“RMM Collection Control dialog box” on page 326

You use this dialog box to configure the data collection settings for the DFSMSrmm feature.

Navigating the workspaces of the DFSMSrmm Status node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources. See “Accessing workspaces” on page 8 to learn more about navigating from workspace to workspace.

You access the primary workspaces of this node as follows:

1. Click the node in the Tivoli Enterprise Portal. The node is activated and the first workspace is displayed in the portal.
2. Right-click the node to display the pop-up menu.
3. Select the **Workspace** submenu to see a list of the available linked workspaces. These linked workspaces, in turn, link to many other, related workspaces as shown in Table 14 on page 78

Workspaces that are available in the “Workspaces” pop-up menu of the Navigator node:

DFSMSrmm Location Summary

DFSMSrmm Error Summary

DFSMSrmm Owner Summary

DFSMSrmm Pending Actions

DFSMSrmm Account Summary

DFSMSrmm Volume Search

DFSMSrmm Job Summary

DFSMSrmm Dataset Search

DFSMSrmm Program Summary

DFSMSrmm Configuration

DFSMSrmm VRS Summary

You start in the “DFSMSrmm System Summary workspace” on page 79 to access all the other workspaces that are listed in Table 14 on page 78. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Note: The following italic phrase is repeated in Table 14 on page 78: *“DFSMSrmm Dataset Details workspaces are provided here.”* A set of standard workspaces is available where you see this phrase in the workspace hierarchy. The phrase includes a link to a description of the set of workspaces, which detail the volume, owner, VRS, and SMS construct of the relevant data set.

Table 14. DFSMSrmm Status node: bulleted list showing the hierarchy of linked workspaces

| |
|--|
| <p>DFSMSrmm Status</p> <ul style="list-style-type: none"> • DFSMSrmm Location Summary <ul style="list-style-type: none"> - DFSMSrmm Location Details - DFSMSrmm Location Summary Volumes <ul style="list-style-type: none"> - DFSMSrmm Volume Details <ul style="list-style-type: none"> • DFSMSrmm Realtime Datasets <ul style="list-style-type: none"> - <i>DFSMSrmm Dataset Details workspaces are provided here.</i> - DFSMSrmm Realtime Datasets <ul style="list-style-type: none"> • <i>DFSMSrmm Dataset Details workspaces are provided here.</i> - DFSMSrmm Location Realtime Volumes <ul style="list-style-type: none"> - DFSMSrmm Volume Details <ul style="list-style-type: none"> • DFSMSrmm Realtime Datasets <ul style="list-style-type: none"> - <i>DFSMSrmm Dataset Details workspaces are provided here.</i> - DFSMSrmm Owner Details - DFSMSrmm Realtime Datasets <ul style="list-style-type: none"> • <i>DFSMSrmm Dataset Details workspaces are provided here.</i> • DFSMSrmm Owner Summary <ul style="list-style-type: none"> - DFSMSrmm Owner Details - DFSMSrmm Owner Summary Volumes <ul style="list-style-type: none"> - DFSMSrmm Volume Details <ul style="list-style-type: none"> • DFSMSrmm Realtime Datasets <ul style="list-style-type: none"> - <i>DFSMSrmm Dataset Details workspaces are provided here.</i> - DFSMSrmm Realtime Datasets <ul style="list-style-type: none"> • <i>DFSMSrmm Dataset Details workspaces are provided here.</i> - DFSMSrmm Owner Realtime Volumes <ul style="list-style-type: none"> - DFSMSrmm Volume Details <ul style="list-style-type: none"> • DFSMSrmm Realtime Datasets <ul style="list-style-type: none"> - <i>DFSMSrmm Dataset Details workspaces are provided here.</i> • DFSMSrmm Owner Details - DFSMSrmm Realtime Datasets <ul style="list-style-type: none"> • <i>DFSMSrmm Dataset Details workspaces are provided here.</i> • DFSMSrmm Account Summary <ul style="list-style-type: none"> - DFSMSrmm Account Summary Volumes <ul style="list-style-type: none"> - DFSMSrmm Volume Details <ul style="list-style-type: none"> • DFSMSrmm Realtime Datasets <ul style="list-style-type: none"> - <i>DFSMSrmm Dataset Details workspaces are provided here.</i> - DFSMSrmm Realtime Datasets <ul style="list-style-type: none"> • <i>DFSMSrmm Dataset Details workspaces are provided here.</i> • DFSMSrmm Job Summary <ul style="list-style-type: none"> - DFSMSrmm Job Summary Datasets <ul style="list-style-type: none"> - <i>DFSMSrmm Dataset Details workspaces are provided here.</i> - DFSMSrmm Volume Details <ul style="list-style-type: none"> • DFSMSrmm Realtime Datasets <ul style="list-style-type: none"> - <i>DFSMSrmm Dataset Details workspaces are provided here.</i> • DFSMSrmm Owner Details |
|--|

Table 14. DFSMSrmm Status node: bulleted list showing the hierarchy of linked workspaces (continued)

- DFSMSrmm Program Summary
 - DFSMSrmm Program Summary Datasets
 - *DFSMSrmm Dataset Details workspaces are provided here.*
 - DFSMSrmm Volume Details
 - DFSMSrmm Realtime Datasets
 - *DFSMSrmm Dataset Details workspaces are provided here.*
 - DFSMSrmm Owner Details
- DFSMSrmm VRS Summary
 - DFSMSrmm VRS Summary Datasets
 - *DFSMSrmm Dataset Details workspaces are provided here.*
 - DFSMSrmm Volume Details
 - DFSMSrmm Realtime Datasets
 - *DFSMSrmm Dataset Details workspaces are provided here.*
 - DFSMSrmm Owner Details
 - DFSMSrmm VRS Details
- DFSMSrmm Error Summary
 - DFSMSrmm Volume Details
 - DFSMSrmm Realtime Datasets
 - *DFSMSrmm Dataset Details workspaces are provided here.*
 - DFSMSrmm Owner Details
- DFSMSrmm Pending Actions
 - DFSMSrmm Pending Action Volumes
 - DFSMSrmm Realtime Datasets
 - *DFSMSrmm Dataset Details workspaces are provided here.*
 - DFSMSrmm Volume Details
 - DFSMSrmm Realtime Datasets
 - *DFSMSrmm Dataset Details workspaces are provided here.*
 - DFSMSrmm Owner Details
 - DFSMSrmm Pending Move Volumes
 - DFSMSrmm Realtime Datasets
 - *DFSMSrmm Dataset Details workspaces are provided here.*
 - DFSMSrmm Volume Details
 - DFSMSrmm Realtime Datasets
 - *DFSMSrmm Dataset Details workspaces are provided here.*
 - DFSMSrmm Owner Details
- DFSMSrmm Volume Search
 - DFSMSrmm Volume Details
 - DFSMSrmm Realtime Datasets
 - *DFSMSrmm Dataset Details workspaces are provided here.*
 - DFSMSrmm Realtime Datasets
 - *DFSMSrmm Dataset Details workspaces are provided here.*
- DFSMSrmm Dataset Search
 - *DFSMSrmm Dataset Details workspaces are provided here.*
 - DFSMSrmm Volume Details
 - DFSMSrmm Realtime Datasets
 - *DFSMSrmm Dataset Details workspaces are provided here.*
- DFSMSrmm Configuration
 - Dataset Details

DFSMSrmm System Summary workspace

This topic describes the main components of the principal workspace for this node.

When you click the DFSMSrmm Status node in the Navigator of the Tivoli Enterprise Portal, you see the DFSMSrmm System Summary workspace. The main components of the workspace are listed here.

Note: DFSMSRmm Data collection is turned off by default, so this workspace is initially empty. In addition, if your installation does not utilize DFSMSRmm, this workspace is always empty.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

DFSMSRmm Control Dataset Space bar graph

This graph contains bars whose height represents the percentage utilization of the Control Data Set data component and the Journal, as reported by DFSMSRmm.

DFSMSRmm Pending Action Summary Volume Counts bar graph

This graph contains bars representing the total number of volumes in each of the possible DFSMSRmm Pending Action states: Scratch, Replace, Return, Init, Erase and Notify.

DFSMSRmm Volume Type Distribution Summary pie graph

This graph shows the number of non-scratch volumes of each type: Master, User, Init and Entry.

DFSMSRmm Volume Utilization Summary bar graph

This graph contains shows the distribution of volumes based on their percentage utilization. The height of each bar represents the number of volumes in that utilization range (for example, utilization between 21% and 30%).

DFSMSRmm System Summary Total table view

This table contains a single row that summarizes the entire DFSMSRmm environment and contains the following attributes. Notice that many of these same attributes are shown on the linked-to Summary workspaces described below. In those cases the attribute values are relative to the specific object that you are summarizing (for example, a location, an owner, a creating job, and so on). The attribute values are not relative to the entire DFSMSRmm environment:

- Collection Status - whether a DFSMSRmm collection cycle has been run, if it ran to completion or was stopped, and so on.
- Summary Collection Start Time - the time that the last collection cycle started. All summary data displayed on this and other summary workspaces is based on the contents of the DFSMSRmm CDS between the indicated Collection Start and End times
- Summary Collection End Time
- Total Volumes - the total number of volumes found in the CDS
- Total Datasets - the total number of data sets contained on all volumes
- Physical Volumes - the number of volumes that are neither Logical nor Stacked
- Scratch Volumes
- Master Volumes
- User Volumes
- Init Volumes
- Entry Volumes
- New Scratch Volumes - the number of volumes returned to scratch status during or since the last DFSMSRmm expiration processing
- Total Errors - the total number of temporary and permanent read and write errors on all volumes
- Total Permanent Errors
- Total Temporary Errors

DFSMSRmm System Usage and Activity Summary table view

This table contains a single row that summarizes the entire DFSMSRmm environment, and contains the following attributes. Some of these attributes can be calculated only when

you have selected the Summarize RMM Dataset Attributes option in the “RMM Collection Control dialog box” on page 326. Notice that many of these same attributes are shown on the linked-to Summary workspaces described below. In those cases the attribute values are relative to the specific object that you are summarizing (for example, a location, an owner, a creating job, and so on). The attribute values are not relative to the entire DFSMSrmm environment:

- Total Capacity - capacity of all volumes, in gigabytes
- Total Usage - the total usage on all volumes, in gigabytes
- Average Volume Utilization - average utilization across all volumes, expressed as percentage
- Average Compression Ratio - calculated percentage compression achieved between the data written by the application and the data recorded on the tape media
- Average Volume Use Count - average number of times each volume was read or written
- Volumes R/W Previous Day - the number of volumes that were read or written on the previous Julian day
- % Volumes R/W Previous Day - the percentage of the total volumes that were read or written on the previous Julian day
- Average Datasets per Volumes - the average number of data sets contained on each tape volume
- Average Dataset Size - the average size of a data set, in megabytes to 1 decimal place
- Datasets R/W Previous Day - the number of data sets that were read or written on the previous Julian day
- Datasets Created Previous Day - the number of data sets that were created on the previous Julian day
- Data Created Previous Day - the amount of data that was created on the previous Julian day in all newly created data sets, expressed as megabytes to 1 decimal place

Workspaces that are linked to this workspace

See “Navigating the workspaces of the DFSMSrmm Status node” on page 77.

Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

See Chapter 28, “Working with the DFSMSrmm Status feature,” on page 323 for additional information.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the DFSMSrmm Status node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

This workspace monitors DFSMSrmm activity for the entire RMM subsystem, including utilization of volumes and CDS space.

| DFSMSrmm Status workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| DFSMSrmm Control Dataset Space | BAR | RMM Control Dataset |
| DFSMSrmm Pending Action Summary Volume Counts | BAR | RMM Summary |
| DFSMSrmm Volume Utilization Summary | BAR | RMM Summary |
| DFSMSrmm Volume Type Distribution Summary | PIE | RMM Summary |
| DFSMSrmm System Summary Totals | TABLE | RMM Summary |
| DFSMSrmm System Usage & Activity Summary | TABLE | RMM Summary |

This workspace contains views displaying summary attributes for all Locations that were found during the last Summary Collection cycle:

| DFSMSrmm Location Summary workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| DFSMSrmm Location Summary Totals | TABLE | RMM Summary |
| DFSMSrmm Location Usage & Activity Summary | TABLE | RMM Summary |
| DFSMSrmm Top 10 Locations by Total Volumes | BAR | RMM Summary |

This workspace contains both static and dynamic values from the control data set (CDS), which is a data set that contains configuration, operational, and communication information for the DFSMSrmm functions.

| DFSMSrmm Configuration workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| DFSMSrmm Control Dataset | TABLE | RMM Control Dataset |
| DFSMSrmm Control Dataset Space | BAR | RMM Control Dataset |
| DFSMSrmm Days Since Last Utility Processing | BAR | RMM Config |
| DFSMSrmm Configuration Options | TABLE | RMM Config |
| DFSMSrmm Configuration State | TABLE | RMM Config |

This workspace displays the top volumes with non-zero temporary and permanent error counts:

| DFSMSrmm Error Summary workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| DFSMSrmm Top 100 Temporary Error Volumes | TABLE | RMM Summary |
| DFSMSrmm Top 100 Permanent Error Volumes | TABLE | RMM Summary |
| DFSMSrmm Top 10 Permanent Error Volumes | BAR | RMM Summary |
| DFSMSrmm Top 10 Temporary Error Volumes | BAR | RMM Summary |

| DFSMSrmm Realtime Datasets workspace | | |
|--------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Datasets | TABLE | RMM Dataset |
| DFSMSrmm Dataset Block Count | BAR | RMM Dataset |

This workspace displays attributes for all actions. The metrics that are displayed for each action status correspond to the last collection cycle:

| DFSMSrmm Pending Actions workspace | | |
|---------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Action Status | TABLE | RMM Action |
| DFSMSrmm Move Status | TABLE | RMM Action |
| DFSMSrmm Action Summary Volume Counts | BAR | RMM Summary |

This workspace contains views displaying summary attributes for all Owners that were found during the last Summary Collection cycle:

| DFSMSrmm Owner Summary workspace | | |
|---|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Top 10 Owners by Total Volumes | BAR | RMM Summary |
| DFSMSrmm Owner Usage & Activity Summary | TABLE | RMM Summary |
| DFSMSrmm Owner Summary Totals | TABLE | RMM Summary |

This workspace contains views displaying summary attributes for all accounts that were found during the last Summary Collection cycle:

| DFSMSrmm Account Summary workspace | | |
|---|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Account Summary Activity | TABLE | RMM Summary |
| DFSMSrmm Account Summary Totals | TABLE | RMM Summary |
| DFSMSrmm Top 10 Summary Accounts by Total Volumes | BAR | RMM Summary |

This workspace displays all attributes for all DFSMSrmm Locations.

| DFSMSrmm Location Details workspace | | |
|-------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Location Details | TABLE | RMM Location |
| DFSMSrmm Location Summary Volumes | BAR | RMM Summary |

This workspace contains views displaying summary attributes for all Job objects that were found during the last Summary Collection cycle:

| DFSMSrmm Job Summary workspace | | |
|--------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Job Summary | TABLE | RMM Summary |

| | | |
|-------------------------------------|-----|-------------|
| DFSMSrmm Top Jobs by Total Datasets | BAR | RMM Summary |
|-------------------------------------|-----|-------------|

This workspace contains views displaying summary attributes for all Program objects that were found during the last Summary Collection cycle:

| DFSMSrmm Program Summary workspace | | |
|---|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Program Summary | TABLE | RMM Summary |
| DFSMSrmm Top Programs by Total Datasets | BAR | RMM Summary |

This workspace contains views displaying summary attributes for all VRS objects that were found during the last Summary Collection cycle:

| DFSMSrmm VRS Summary workspace | | |
|---------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Top 10 VRS by Total Datasets | BAR | RMM Summary |
| DFSMSrmm VRS Summary | TABLE | RMM Summary |

This workspace displays the volumes that are in the selected action state.

| DFSMSrmm Pending Action Volumes workspace | | |
|---|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Pending Action Volumes | TABLE | RMM Volume |
| DFSMSrmm Pending Action Summary Volume Counts | BAR | RMM Summary |

This workspace displays a subset of attributes for a list of volumes that were obtained during the last summary collection cycle. These volumes are associated with a specific object that was summarized in another workspace.

| DFSMSrmm Summary Volumes workspace | | |
|------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Summary Volumes | TABLE | RMM Volume |
| DFSMSrmm Summary Volume Total | BAR | RMM Summary |

This workspace displays selected attributes for a list of data sets.

| DFSMSrmm Dataset Search workspace | | |
|-----------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Dataset Block Count | BAR | RMM Dataset |
| DFSMSrmm Datasets | TABLE | RMM Dataset |

| DFSMSrmm Volume Search workspace | | |
|--------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Permanent Volume Errors > 0 | BAR | RMM Volume |
| DFSMSrmm Volumes | TABLE | RMM Volume |

| DFSMSrmm Summary Datasets workspace | | |
|-------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Summary Datasets | TABLE | RMM Volume |
| DFSMSrmm Summary Data Set Total | BAR | RMM Summary |

This workspace displays the volumes that are in the move state.

| DFSMSrmm Pending Move Volumes workspace | | |
|---|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Pending Move Volumes | TABLE | RMM Volume |
| DFSMSrmm Pending Action Summary Volume Counts | BAR | RMM Summary |

This workspace displays a filler set of attributes for one or more volumes obtained in real-time from the DFSMSrmm facility. The list of volumes that this workspace displays is associated with a selected Owner or Location, with a selected Pending Action, or is specified through the pop-up menu the DFSMSrmm Volume Search link.

| DFSMSrmm Realtime Volumes workspace | | |
|--------------------------------------|-----------|-----------------|
| View | View type | Attribute group |
| DFSMSrmm Permanent Volume Errors > 0 | BAR | RMM Volume |
| DFSMSrmm Realtime Volumes | TABLE | RMM Volume |

DFSMSrmm Dataset Details workspaces

The workspaces in this set of workspaces provide information about a data set and its volume, owner, VRS, and SMS construct. The set workspaces is provided in many places in the hierarchy of workspaces, and represents the deepest level of detail that is available.

| DFSMSrmm Dataset Details workspace | | |
|--|-----------|-----------------|
| This workspace displays all attributes for a single DFSMSrmm data set. | | |
| View | View type | Attribute group |
| DFSMSrmm Dataset Information | TABLE | RMM Dataset |
| DFSMSrmm Dataset Usage | BAR | RMM Dataset |
| DFSMSrmm Datasets Statistics | TABLE | RMM Dataset |
| DFSMSrmm Dataset Dates | TABLE | RMM Dataset |
| DFSMSrmm Dataset VRS | TABLE | RMM Dataset |

| DFSMSrmm Volume Details workspace | | |
|--|-----------|-----------------|
| This workspace displays all attributes for a single DFSMSrmm volume. | | |
| View | View type | Attribute group |
| DFSMSrmm Volume Information | TABLE | RMM Volume |
| DFSMSrmm Volume Date Information | TABLE | RMM Volume |
| DFSMSrmm Volume Statistics | TABLE | RMM Volume |

| | | |
|--------------------------------------|-------|------------|
| DFSMSrmm Volume Storage and Movement | TABLE | RMM Volume |
| Volume Errors | BAR | RMM Volume |

| | | |
|---|------------------|------------------------|
| DFSMSrmm Owner Details workspace | | |
| This workspace displays all attributes for a single DFSMSrmm Owner. | | |
| View | View type | Attribute group |
| DFSMSrmm Owner Details | TABLE | RMM Owner |
| DFSMSrmm Owner Volumes | BAR | RMM Owner |

| | | |
|--|------------------|------------------------|
| DFSMSrmm VRS Details workspace | | |
| This workspace displays all attributes for a single DFSMSrmm Vital Record Specification. | | |
| View | View type | Attribute group |
| DFSMSrmm VRS Details | TABLE | RMM VRS |
| DFSMSrmm VRS Priority | BAR | RMM VRS |
| DFSMSrmm VRS Count | BAR | RMM VRS |

Chapter 16. Monitoring Dataset Attributes System Summary data

This topic describes the workspaces that monitor the status of Dataset Attributes System Summary data. You access these workspaces through the Dataset Attributes System Summary node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the Dataset Attributes System Summary node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Dataset Attributes System Summary node to navigate to the workspaces that you see in Table 15 on page 88. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 15. Dataset Attributes System Summary node: bulleted list showing the hierarchy of workspace links

| |
|--|
| <p>Dataset Attributes System Summary</p> <ul style="list-style-type: none">• Dataset Space Summary<ul style="list-style-type: none">- Dataset Attribute Details<ul style="list-style-type: none">- Dataset Details• DCB Summary<ul style="list-style-type: none">- Largest Datasets in DSORG- Largest Inefficient Blocksize Datasets<ul style="list-style-type: none">- Dataset Attribute Details<ul style="list-style-type: none">• Dataset Details- Newborn Dataset Summary<ul style="list-style-type: none">- Dataset Attribute Details<ul style="list-style-type: none">• Dataset Details- Largest Never Referenced Datasets<ul style="list-style-type: none">- Dataset Attribute Details<ul style="list-style-type: none">• Dataset Details- Mature Dataset Summary – Largest Datasets Unreferenced > 366 Days<ul style="list-style-type: none">- Dataset Attribute Details<ul style="list-style-type: none">• Dataset Details• Catalog Summary<ul style="list-style-type: none">- Datasets in Catalog<ul style="list-style-type: none">- Dataset Attribute Details<ul style="list-style-type: none">• Dataset Details• DFSMS Class Summary<ul style="list-style-type: none">- SMS Data Class Details- SMS Management Class Details- SMS Storage Class Details- SMS Storage Group Details- Largest Datasets in SMS Class (appears for each class)• Uncataloged Dataset Summary<ul style="list-style-type: none">- Largest Uncataloged Datasets on Volume<ul style="list-style-type: none">- Dataset Attribute Details<ul style="list-style-type: none">• Dataset Details |
|--|

Dataset Attributes System Summary workspace

This topic describes the main components of the principal workspace for this node.

This workspace provides the principal view of the Dataset Attribute Database, which maintains data on data set attributes related to space, DCB, and cataloging. Also see “Configuring the Dataset Attribute Database feature” on page 313.

Note: Data collection is turned off by default, so this workspace is initially empty.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Table view

The table view displays one row of data for each resource. The columns show the following attributes:

- Volumes online
- Volumes processed
- Datasets processed
- Available tracks
- Allocated tracks

- Used tracks
- Unused tracks
- Percent of available allocated
- Percent of allocated used
- Percent of allocated free
- Name of largest data set
- Size of largest data set
- Name of data set with largest unused space
- Amount of space unused by largest data set
- Name of data set with oldest reference date
- Number of unreferenced days for oldest data set
- Number of uncataloged data sets
- Collection start & end date/time

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Dataset Attributes System Summary node” on page 87.

Learn how to navigate to workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Dataset Attributes System Summary node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “Dataset Attributes System Summary workspace” on page 88 for more information on the primary workspace for this node.

| Dataset Attributes System Summary workspace | | |
|---|-----------|-----------------------------------|
| View | View type | Attribute group |
| Dataset Attribute Summary | TABLE | Dataset Attributes System Summary |
| Overall Space Utilization | PIE | Dataset Attributes System Summary |
| Dataset Attribute Dataset Extremes | TABLE | Dataset Attributes System Summary |
| Overall Space Allocation | BAR | Dataset Attributes System Summary |

| Dataset Space Summary workspace | | |
|---------------------------------|-----------|--|
| View | View type | Attribute group |
| Top Datasets by Allocated Space | TABLE | Dataset Attributes Space Allocated Summary |
| Top Datasets by Extents | TABLE | Dataset Attributes Extent Summary |

| | | |
|------------------------------|-------|---|
| Top Datasets by Unused Space | TABLE | Dataset Attributes Space Unused Summary |
| Top Datasets by CA Splits | TABLE | Dataset Attributes CA Split Summary |
| Top Datasets by CI Splits | TABLE | Dataset Attributes CI Split Summary |

| DCB Summary workspace | | |
|--|------------------|---|
| View | View type | Attribute group |
| DSORG Summary Report | TABLE | Dataset Attributes DSORG Summary |
| Blocksize Distribution Report | TABLE | Dataset Attributes Blocksize Summary |
| Blocksize Track Allocation Distribution | BAR | Dataset Attributes Blocksize Summary |
| Creation Date Distribution Report | TABLE | Dataset Attributes Creation Date Summary |
| Creation Date Track Allocation Distribution | BAR | Dataset Attributes Creation Date Summary |
| Reference Date Distribution Report | TABLE | Dataset Attributes Reference Date Summary |
| Reference Date Track Allocation Distribution | BAR | Dataset Attributes Reference Date Summary |

| Catalog Summary workspace | | |
|----------------------------------|------------------|------------------------------------|
| View | View type | Attribute group |
| Catalog Summary Report | TABLE | Dataset Attributes Catalog Summary |
| Datasets per Catalog | BAR | Dataset Attributes Catalog Summary |

| DFSMS Class Summary workspace | | |
|--------------------------------------|------------------|--------------------------------------|
| View | View type | Attribute group |
| Storage Group Allocation | BAR | Dataset Attributes SMS Class Summary |
| Storage Groups | TABLE | Dataset Attributes SMS Class Summary |
| Data Classes | TABLE | Dataset Attributes SMS Class Summary |
| Data Class Allocation | BAR | Dataset Attributes SMS Class Summary |
| Management Classes | TABLE | Dataset Attributes SMS Class Summary |
| Management Class Allocation | BAR | Dataset Attributes SMS Class Summary |
| Storage Classes | TABLE | Dataset Attributes SMS Class Summary |
| Storage Class Allocation | BAR | Dataset Attributes SMS Class Summary |

| Dataset Attribute Details workspace |
|--|
|--|

| View | View type | Attribute group |
|-----------------------------|-----------|-----------------------------------|
| Space Attributes | TABLE | Dataset Attributes Dataset Detail |
| DCB and Extended Attributes | TABLE | Dataset Attributes Dataset Detail |
| GDG Information | TABLE | Dataset Attributes Dataset Detail |
| Dataset SMS Constructs | TABLE | Dataset Attributes Dataset Detail |
| Dataset Space | BAR | Dataset Attributes Dataset Detail |
| VSAM Attributes | TABLE | Dataset Attributes Dataset Detail |

| Largest Datasets in SMS Class workspace | | |
|--|-----------|-------------------------------------|
| View | View type | Attribute group |
| Largest Datasets in SMS Class by Allocated Space | TABLE | Dataset Attributes SMS Class Detail |
| Space Allocation for Largest Datasets in SMS Class | BAR | Dataset Attributes SMS Class Detail |

| Largest Datasets in DSORG workspace | | |
|---|-----------|---------------------------------|
| View | View type | Attribute group |
| Space Allocated for Largest Datasets in DSORG | BAR | Dataset Attributes DSORG Detail |
| Largest Datasets in DSORG by Allocated Space | TABLE | Dataset Attributes DSORG Detail |

| Largest Inefficient Blocksize Datasets workspace | | |
|---|-----------|--|
| View | View type | Attribute group |
| Largest Datasets where Blocksize < 128 by Allocated Space | TABLE | Dataset Attributes Inefficient Blocksize Summary |
| Space Allocation for Largest Datasets where Blocksize < 128 | BAR | Dataset Attributes Inefficient Blocksize Summary |

| Newborn Dataset Summary - Datasets 0 or 1 Day Old workspace | | |
|---|-----------|--|
| View | View type | Attribute group |
| Largest Datasets 0 or 1 Day Old by Allocated Space | TABLE | Dataset Attributes Newborn Dataset Summary |
| Space Allocation for Largest Datasets 0 or 1 Day Old | BAR | Dataset Attributes Newborn Dataset Summary |

| Mature Dataset Summary - Largest Datasets Unreferenced > 366 Days workspace | | |
|---|-----------|---|
| View | View type | Attribute group |
| Largest Datasets Unreferenced for > 366 Days by Allocated Space | TABLE | Dataset Attributes Mature Dataset Summary |
| Space Allocation for Largest Datasets Unreferenced for > 366 Days | BAR | Dataset Attributes Mature Dataset Summary |

| Largest Never Referenced Datasets workspace | | |
|---|-----------|-----------------|
| View | View type | Attribute group |

| | | |
|--|-------|---|
| Largest Never Referenced Datasets by Allocated Space | TABLE | Dataset Attributes Never Referenced Dataset Summary |
| Space Allocation for Largest Never Referenced Datasets | BAR | Dataset Attributes Never Referenced Dataset Summary |

| Datasets in Catalog workspace | | |
|---|------------------|------------------------------------|
| View | View type | Attribute group |
| All Datasets in Catalog | TABLE | Dataset Attributes Dataset Detail |
| Space Utilization for All Datasets in Catalog | PIE | Dataset Attributes Catalog Summary |

| Uncataloged Dataset Summary workspace | | |
|--|------------------|--|
| View | View type | Attribute group |
| Uncataloged Dataset Summary Report | TABLE | Dataset Attributes Uncataloged Summary |
| Uncataloged Dataset Allocation by Volume | BAR | Dataset Attributes Uncataloged Summary |
| Uncataloged Dataset Total | BAR | Dataset Attributes Catalog Summary |
| Uncataloged Dataset Allocation Total | BAR | Dataset Attributes Catalog Summary |

| Largest Uncataloged Datasets on Volume workspace | | |
|--|------------------|---------------------------------------|
| View | View type | Attribute group |
| Largest Uncataloged Datasets on Volume by Allocated Space | TABLE | Dataset Attributes Uncataloged Detail |
| Space Utilization for Largest Uncataloged Datasets on Volume | BAR | Dataset Attributes Uncataloged Detail |
| Uncataloged Dataset Space Allocation | PIE | Dataset Attributes Uncataloged Detail |

Chapter 17. Monitoring dataset groups

This topic describes the workspaces that monitor the status of Dataset Group Summary data. You access these workspaces through the Dataset Group Summary node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the Dataset Group Summary node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Dataset Group Summary node to navigate to the workspaces that you see in Table 16. You can view details regarding the following types of data set groups: Non-VSAM, VSAM, and GDG (generation data group). You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 16. Dataset Group Summary node: bulleted list showing the hierarchy of linked workspaces

| |
|--|
| Dataset Group Summary |
| • Dataset Group Details Primary |
| – Dataset Group Details Non-VSAM |
| – SMS Storage Class Details |
| – Dataset Performance Detail Report ¹ |
| – Dataset Group Details VSAM |
| – Dataset Group VSAM Attributes |
| • Catalog Details |
| • Catalog Dataset Attributes |
| • Last Catalog Backup |
| • SMS Storage Class Details |
| • Dataset Performance Detail Report ¹ |
| – Dataset Group Details GDG |
| – Dataset Group Details Non-VSAM |
| • SMS Storage Class Details |
| • Dataset Performance Detail Report ¹ |

Footnotes for Table 16:

1. This group of workspaces is organized in the following nested pattern:
 - Dataset Performance Detail Report
 - Dataset Details
 - Dataset Performance Detail Trend

Dataset Group Summary workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Dataset group space utilization graph

This graph displays total free space and group total used space information for each data set group.

Dataset group summary table view

The tabular portion covering the entire width of the workspace at the bottom of the screen displays one row of data for each defined data set group. The rows contain:

- The data set group name
- The total number of data sets
- The data set with the largest allocation
- The largest space allocation
- The data set with the largest free space
- The largest space free
- The data set with the oldest referenced date within the group
- The oldest referenced date for the identified data set within the group
- The data set with the worst MSR within the group, and the highest MSR of the data set
- The total used space for all of the data sets within the group
- The total number of I/Os for all the active data sets within the group
- The total average MSR for all the active data sets within the group
- The data set group used space percentage
- The total allocated space for all the data sets within the group
- The data set group percentage of free space
- The total free space for all the data sets within the group

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Dataset Group Summary node” on page 93.

Learn how to navigate to workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Dataset Group Summary node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “Dataset Group Summary workspace” on page 93 for more information on the primary workspace for this node.

| |
|--|
| Dataset Group Summary workspace |
|--|

| View | View type | Attribute group |
|---------------------------------|-----------|-----------------------|
| Dataset Group Summary | TABLE | Dataset Group Summary |
| Dataset Group Space Utilization | BAR | Dataset Group Summary |

| Dataset Group Details Primary workspace | | |
|---|-----------|-----------------------|
| View | View type | Attribute group |
| Total Allocations | BAR | Dataset Group Details |
| Response Time | BAR | Dataset Group Details |
| Dataset Group Primary Entries | TABLE | Dataset Group Details |
| I/O Activity | BAR | Dataset Group Details |

| Dataset Group Details VSAM workspace | | |
|--------------------------------------|-----------|-----------------------|
| View | View type | Attribute group |
| VSAM Cluster Report | TABLE | Dataset Group Details |
| Response Time | BAR | Dataset Group Details |
| CI/CA Splits | BAR | Dataset Group Details |
| Record Access | BAR | Dataset Group Details |

| Dataset Group Details Non-VSAM workspace | | |
|--|-----------|------------------------------|
| View | View type | Attribute group |
| Response Time | BAR | Dataset Group Details Volume |
| Dataset Information by Volume | TABLE | Dataset Group Details Volume |
| Space Utilization | BAR | Dataset Group Details Volume |
| Aggregate Dataset Information | TABLE | Dataset Group Details |

| Dataset Group Details GDG workspace | | |
|-------------------------------------|-----------|-----------------------|
| View | View type | Attribute group |
| Aggregate GDG Dataset Information | TABLE | Dataset Group Details |
| Response Time | BAR | Dataset Group Details |
| I/O Activity | BAR | Dataset Group Details |
| Total Allocations | BAR | Dataset Group Details |

| Dataset Group VSAM Attributes workspace | | |
|---|-----------|------------------------------|
| View | View type | Attribute group |
| Aggregate Dataset Information | TABLE | Dataset Group Details |
| Dataset Information by Volume | TABLE | Dataset Group Details Volume |
| Dataset Attributes | TABLE | Dataset Group Detail Attrs |

Chapter 18. Monitoring SMS configuration

This topic describes the workspaces that monitor the status of SMS configuration. You access these workspaces through the SMS Configuration node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the SMS Configuration node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the SMS Configuration node to navigate to the workspaces that you see in Table 17. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 17. SMS Configuration node: bulleted list showing the hierarchy of linked workspaces

| |
|----------------------------------|
| SMS Configuration |
| • SMS Base Configuration |
| – SMS Configuration Details |
| – SMS Tape Storage Group Details |
| – SMS Data Class Details |
| – SMS Management Class Details |
| – SMS Storage Class Details |

SMS Configuration workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help, including in the hover help.

SMS base configuration table view

The tabular report at the top of this workspace contains a single row of data describing SMS status in the reporting LPAR. Listed are:

- The SMS status and level
- The MVS™ level
- The TSO ID of the last user to update the SMS configuration
- The date and time of the last update
- The default management class
- The default number of tracks per cylinder, track size and unit type

This tabular report links to SMS Configuration Details workspace.

Storage groups table view

This tabular report contains one row for each storage group defined in the reporting LPAR. Each row lists:

- The storage group name and type
- The TSO ID of the last user to update the group
- The date and time of that update
- A description of the group

This tabular report links to SMS Storage Group Details and SMS Tape Storage Group Details workspaces.

Data classes table view

This tabular report contains one row for each defined SMS data class. Each row lists:

- The name of the data class
- The ID of the user who last updated the data class definition
- The date and time that the data class definition was updated
- A description of the data class

This tabular report links to SMS Data Class Details workspace.

Management classes table view

This tabular report contains one row for each defined SMS management class. Each row lists:

- The management class name
- The user ID that last updated this management class
- The date and time when this management class was last updated
- The storage administrator supplied description of the management class

This tabular report links to SMS Management Class Details workspace.

Storage classes table view

This tabular report contains one row for each defined SMS storage class. Each row lists:

- The name of the storage class
- The last user ID to update this storage class definition
- The date and time the storage class definition was last updated
- The storage administrator supplied description of the storage class

This tabular report links to Storage Class Details workspace.

Workspaces that are linked to this workspace

See “Navigating the workspaces of the SMS Configuration node” on page 97.

Learn how to navigate to workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the SMS Configuration node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of

workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “SMS Configuration workspace” on page 97 for more information on the primary workspace for this node.

| SMS Configuration workspace | | |
|------------------------------------|------------------|------------------------|
| View | View type | Attribute group |
| SMS Base Configuration | TABLE | SMS Configuration |
| Data Classes | TABLE | SMS Data Class |
| Management Classes | TABLE | SMS Management Class |
| Storage Classes | TABLE | SMS Storage Class |
| Storage Groups | TABLE | SMS Storage Group |

| SMS Storage Group Details workspace | | |
|--|------------------|--------------------------|
| View | View type | Attribute group |
| Storage Group Details Report | TABLE | SMS Storage Group |
| Storage Group Details Report | TABLE | SMS Storage Group |
| Storage Group Status | TABLE | SMS Storage Group Status |

| SMS Tape Storage Group Details workspace | | |
|---|------------------|--------------------------|
| View | View type | Attribute group |
| Storage Group Status | TABLE | SMS Storage Group Status |
| Tape Storage Group Details Report | TABLE | SMS Storage Group |
| Tape Storage Group Libraries | TABLE | SMS Storage Group |

| SMS Configuration Details workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| SMS Configuration Members | TABLE | SMS Systems |
| Cache Sets | TABLE | SMS Cache Sets |
| Base Configuration | TABLE | SMS Configuration |

| SMS Data Class Details workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Data Class Allocation | TABLE | SMS Data Class |
| Data Class VSAM Attributes | TABLE | SMS Data Class |
| Data Class Details | TABLE | SMS Data Class |

| SMS Management Class Details workspace | | |
|---|------------------|------------------------|
| View | View type | Attribute group |
| Management Class Details | TABLE | SMS Management Class |
| Management Class Details | TABLE | SMS Management Class |
| Management Class Details | TABLE | SMS Management Class |

| SMS Storage Class Details workspace | | |
|--|------------------|------------------------|
| View | View type | Attribute group |
| Storage Class Details | TABLE | SMS Storage Class |
| Millisecond Response Time | BAR | SMS Storage Class |

Chapter 19. Monitoring Storage Toolkit action requests

This topic describes the workspaces that monitor the status of Storage Toolkit action requests. You access these workspaces through the Storage Toolkit Action Requests node in the Tivoli Enterprise Portal.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

Navigating the workspaces of the Storage Toolkit node

This section summarizes the workspace links that are available through the Navigator node. Depending on your monitoring goal, you follow a specific series of links to get detailed information about storage resources.

You start in the Storage Toolkit Action Requests node to navigate to the workspaces that you see in Table 18. You can access the workspaces in this table using the standard path and method for navigating workspaces.

Table 18. Storage Toolkit Action Requests node: bulleted list showing the hierarchy of linked workspaces

| |
|-----------------------------------|
| Storage Toolkit Action Requests |
| • Storage Toolkit Action Requests |
| – Storage Toolkit Result Summary |
| – Storage Toolkit Result Detail |

Storage Toolkit Action Requests workspace

This topic describes the main components of the principal workspace for this node.

The following sections describe various components of this workspace. For a description of the data used as column headings in the workspace table view, see Chapter 23, “Storage Toolkit,” on page 251.

Views This workspace provides the following views. For a description of the attributes used as column headings in the workspace table view, review the attribute group definitions for this workspace in the online help. The help resource includes attribute descriptions that are displayed when you hover with the mouse pointer over a column heading.

Storage Toolkit Action Requests table view

The table view display one row of data for each action created by the user. The view contains the following information:

- Name of user who created action request
- Name of request as given by user
- Description of request as given by user
- The type of action.
- Number of times request has run
- Return code for most recently completed action
- Date/time request was created

- Date/time most recently completed action completed
- Command issued for execution
- Scheduled dates/time at which action request is to run

Workspaces that are linked to this workspace

See “Navigating the workspaces of the Storage Toolkit node” on page 101.

Learn how to navigate to workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9.

See Chapter 2, “Workspaces,” on page 7 for a list of all nodes provided in Tivoli OMEGAMON XE for Storage on z/OS.

Summary information for the workspaces in this node

This section lists the views and related attribute groups for the workspaces of the Storage Toolkit node.

See Chapter 2, “Workspaces,” on page 7 for general information on predefined workspaces and a list of all nodes that are provided by Tivoli OMEGAMON XE for Storage on z/OS. Each node corresponds to an area of concern for the monitoring of storage resources, and the node links to a unique set of workspaces. Learn how to navigate among workspaces here: “Accessing workspaces” on page 8. Also see “Integration with other products” on page 9 to learn about linking to the workspaces of other IBM Tivoli Monitoring products.

See “Storage Toolkit Action Requests workspace” on page 101 for more information on the primary workspace for this node.

| Storage Toolkit Action Requests workspace | | |
|---|-----------|-------------------------|
| View | View type | Attribute group |
| Action Requests | TABLE | Storage Toolkit Request |
| Requests with Last Return Code > 0 | BAR | Storage Toolkit Request |

| Storage Toolkit Result Summary workspace | | |
|--|-----------|--------------------------------|
| View | View type | Attribute group |
| Action Results Summary | TABLE | Storage Toolkit Result Summary |
| Return Code | BAR | Storage Toolkit Result Summary |
| Output Line Count | BAR | Storage Toolkit Result Summary |

| Storage Toolkit Result Detail workspace | | |
|---|-----------|---------------------------------|
| View | View type | Attribute group |
| Resource List | TABLE | Storage Toolkit Result Resource |
| Action Results Detail | TABLE | Storage Toolkit Result Detail |

Chapter 20. Directly monitoring situations

With direct situation analysis, you navigate directly from a situation event workspace to information regarding the original situation that triggered a situation event.

Problem solved by direct situation analysis:

Before direct situation analysis existed, there were many cases in which only the default situation event workspaces were available when a situation is triggered. To investigate the metrics that triggered a situation, the user traversed the Navigator view to access information on the true situation. Direct situation analysis gives you quick access to this information.

When direct situation analysis is available:

Direct situation analysis is available after a situation becomes true, a situation event is opened, and a situation event workspace is created.

When a situation resolves to true and is associated with a Navigator item, a virtual light or situation event indicator overlays the item and all related items above it in the Navigator. As you hover with the mouse over the indicator, a list is displayed. When you click one of the situations in this list, the situation event workspace opens. This workspace shows values of attributes for the situation that was triggered and expert advice, if any.

Direct situation analysis for the Application Summary workspace

The Application Summary workspace produces a situation event workspace when the following situation resolves to true: KS3_Applic_Resp_Time_Critical/Warning. You can navigate from the situation event workspace directly to the Application Dataset Response Time workspace to see basic data that is generated by the situation.

Note: The data might be incomplete or missing if data set performance data is not being collected for the volumes on which the data sets are located or for the application.

The data sets in the report are ordered by response time, in descending order. This workspace derives its data from the Application Monitoring attribute table.

Use this workspace to determine which data sets the application is utilizing that are not meeting performance objectives. The report lists all data sets that the application has accessed over the current RMF interval. The bar chart contains only those data sets that are not meeting their storage class performance objective.

Links

Right-clicking on the link to the left of each row of tabular data in the Application Dataset Response Time workspace produces a pop-up menu with links to the following workspaces:

- Dataset Details
- Dataset Performance Summary Trend

Direct situation analysis for the Channel Path workspace

The Channel Path workspace produces a situation event workspace when the following situation resolves to true: KS3_Channel_Busy_Pct_Critical/Warning. You can navigate from the situation event workspace directly to the Channel Volume Response Time workspace to see basic data that is generated by the situation.

The volumes in the report are ordered by response time, in descending order. This workspace derives its data from the DASD Volume Performance attribute table.

Use this workspace to determine which volumes, if any, are being impacted by the over utilization of the channel. The report lists all volumes associated with the channel. The chart immediately shows those volumes with an unacceptable response time.

Links

Right-clicking on the link to the left of each row of tabular data in the Channel Volume Response Time workspace produces a pop-up menu with links to the following workspaces:

- Dataset Performance Summary
- Device Performance Details
- Cross System Volume Performance
- Volume Performance Trend

Direct situation analysis for the Cache CU Performance workspace

The Cache CU Performance workspace produces situation event workspaces when the following situations resolve to true:

- KS3_Cachecu_DFW_Retry_Critical/Warning
- KS3_Cachecu_Read_HitP_Critical/Warning
- KS3_Cachecu_Write_HitP_Critical/Warning
- KS3_Cachecu_Inact_Vols_Critical/Warning

You can navigate from one of these situation event workspaces directly to the following workspaces to do direct situation analysis:

| Workspace for direct situation analysis | Comments |
|--|--|
| Situation: KS3_Cachecu_DFW_Retry_Critical/Warning | |
| Cache CU Volume DFW Retry | <p>Default navigation is to this workspace.</p> <p>The volumes in the report are ordered by DFW Retry Percent, in descending order. The graphs map DFW Retry Percent and DFW Hit Percent. This workspace derives its data from the Cache Devices attribute table.</p> <p>Use this workspace to determine which volumes, if any, are being impacted by poor DFW retry percent. The report lists all volumes on the control unit. The two bar charts display those devices that are not being treated well by the cache, in other words, poor hit rates.</p> |
| Cache CU DFW Retry Trend | <p>The plot chart maps DFW Retry Percent. This workspace derives its data from the Cache Control Unit attribute table.</p> <p>Use this workspace to determine if the poor DFW retry percent was a spike or a constant occurrence. The corrective action required varies depending on that information.</p> |

| Workspace for direct situation analysis | Comments |
|---|---|
| Cache CU Volume DFW Retry History | <p>The volumes in the report are ordered by DFW Retry Percent, in descending order. The graphs map DFW Retry Percent and DFW Hit Percent. This workspace derives its data from the Cache Devices attribute table.</p> <p>Use this workspace to determine which volumes, if any, were impacted by poor DFW retry percent. The report lists all volumes on the control unit at that point in time. The two bar charts display those devices that were not being treated well by the cache, in other words, poor hit rates, at that point in time.</p> |
| Situation: KS3_Cachecu_Read_HitP_Critical/Warning | |
| Cache CU Volume Read Hit Percent | <p>Default navigation is to this workspace. The report is sorted in ascending order by Read Hit Percent. This workspace derives its data from the Cache Devices attribute table.</p> <p>Use this workspace to determine which volumes were responsible for the poor control unit read hit percent. The report lists all volumes on the control unit. The two bar charts display those devices that are not being treated well by the cache, in other words, poor hit rates.</p> |
| Cache CU Read Hit Percent Trend | <p>The plot chart maps Read Hit Percent. This workspace derives its data from the Cache Control Unit attribute table.</p> <p>Use this workspace to determine if the poor read hit percent was a spike or a constant occurrence. The corrective action required varies depending on that information.</p> |
| Cache CU Volume Read Hit Percent History | <p>The volumes in the report are ordered by Read Hit Percent, in ascending order. The graphs map Read Hit Percent and Write Hit Percent. This workspace derives its data from the Cache Devices attribute table.</p> <p>Use this workspace to determine which volumes were responsible for the poor read hit percent. The report lists all volumes on the control unit at that point in time. The two bar charts display those devices that were not being treated well by the cache, in other words, poor hit rates, at that point in time.</p> |
| Situation: KS3_Cachecu_Write_HitP_Critical/Warning | |
| Cache CU Volume Write Hit Percent | <p>Default navigation is to this workspace. The report is sorted in ascending order by Write Hit Percent. This workspace derives its data from the Cache Devices attribute table.</p> <p>Use this workspace to determine which volumes were responsible for the poor control unit write hit percent. The report lists all volumes on the control unit. The two bar charts display those devices that are not being treated well by the cache, in other words, poor hit rates.</p> |
| Cache CU Write Hit Percent Trend | <p>The plot chart maps Write Hit Percent This workspace derives its data from the Cache Control Unit attribute table.</p> <p>Use this workspace to determine if the poor write hit percent was a spike or a constant occurrence. The corrective action required varies depending on that information.</p> |

| Workspace for direct situation analysis | Comments |
|---|--|
| Cache CU Volume Write Hit Percent History | <p>The volumes in the report are ordered by Write Hit Percent, in ascending order. The graphs map Read Hit Percent and Write Hit Percent. This workspace derives its data from the Cache Devices attribute table.</p> <p>Use this workspace to determine which volumes were responsible for the poor write hit percent. The report lists all volumes on the control unit at that point in time. The two bar charts display those devices that were not being treated well by the cache, in other words, poor hit rates, at that point in time.</p> |
| Situation: KS3_Cachecu_Inact_Vols_Critical/Warning | |
| Cache CU Volume Inactive Status | <p>The report contains a row for volumes with a cache status of inactive only. One bar chart shows the I/O count for the volumes. This workspace derives its data from the Cache Devices attribute table.</p> <p>Use this workspace to determine which volumes had an inactive status for cache.</p> |

Links

Right-clicking on the link to the left of each row of tabular data produces a pop-up menu with links to the following workspaces:

| Workspace for direct situation analysis | Links to other workspaces |
|---|--|
| Situation: KS3_Cachecu_DFW_Retry_Critical/Warning | |
| Cache CU Volume DFW Retry | <ul style="list-style-type: none"> • Dataset Performance Summary • Physical Volume Group Cache (if applicable, in other words, supported hardware) • Device Performance Details • Cross System Volume Performance Volume Cache Trend |
| Cache CU DFW Retry Trend | <ul style="list-style-type: none"> • Cache CU Volume Historical Performance • Cache CU Raid Rank Historical Data (if applicable) • New workspace: Cache CU Volume DFW Retry History |
| Cache CU Volume DFW Retry History | Dataset Historic Performance |
| Situation: KS3_Cachecu_Read_HitP_Critical/Warning | |
| Cache CU Volume Read Hit Percent | <ul style="list-style-type: none"> • Dataset Performance Summary • Physical Volume Group Cache (if applicable, in other words, supported hardware) • Device Performance Details • Cross System Volume Performance Volume Cache Trend |
| Cache CU Read Hit Percent Trend | <ul style="list-style-type: none"> • Cache CU Volume Historical Performance • Cache CU Raid Rank Historical Data (if applicable) • New workspace: Cache CU Volume Read Hit Percent History |
| Cache CU Volume Read Hit Percent History | Dataset Historic Performance |
| Situation: KS3_Cachecu_Write_HitP_Critical/Warning | |

| Workspace for direct situation analysis | Links to other workspaces |
|---|---|
| Cache CU Volume Write Hit Percent | <ul style="list-style-type: none"> • Dataset Performance Summary • Physical Volume Group Cache (if applicable, in other words, supported hardware) • Device Performance Details • Cross System Volume Performance • Volume Cache Trend |
| Cache CU Write Hit Percent Trend | <ul style="list-style-type: none"> • Cache CU Volume Historical Performance • Cache CU Raid Rank Historical Data (if applicable) • New workspace: Cache CU Volume Write Hit Percent History |
| Cache CU Volume Write Hit Percent History | Dataset Historic Performance |
| Situation: KS3_Cachecu_Inact_Vols_Critical/Warning | |
| Cache CU Volume Inactive Status' | Dataset Performance Summary |

Direct situation analysis for the Logical Control Unit workspace

The Logical Control Unit workspace produces situation event workspaces when the following situations resolve to true:

- KS3_LCU_Av_Delay_Q_Critical/Warning
- KS3_LCU_Cont_Rate_Critical/Warning
- KS3_LCU_IO_Rate_Sec_Critical/Warning

You can navigate from one of these situation event workspaces directly to the following workspaces to do direct situation analysis:

| Workspace for direct situation analysis | Comments |
|---|---|
| LCU Volume Event Performance | <p>Default navigation is to this workspace. The volumes in the report are sorted by response time in descending order. The entries in the bar chart are in descending order by response time. This workspace derives its data from the DASD Volume Performance attribute table.</p> <p>Use this workspace to determine which volumes had poor performance due to the situation. The report lists all volumes in the logical control unit. The bar chart immediately identifies those volumes with unacceptable response time.</p> |
| Logical Control Unit Trend | See "Logical Control Unit workspace" on page 39. |

Links

Right-clicking on the link to the left of each row of tabular data produces a pop-up menu with links to the following workspaces.

| Workspace for direct situation analysis | Links to other workspaces |
|---|---|
| LCU Volume Event Performance | <ul style="list-style-type: none"> • Dataset Performance Summary • Physical Volume Group Performance (if applicable, in other words, supported hardware) • Device Performance Details • Cross System Volume Performance • Volume Performance Trend |
| Logical Control Unit Trend | See “Logical Control Unit workspace” on page 39. |

Direct situation analysis for the SMS Storage Groups Performance workspace

The SMS Storage Groups Performance workspace produces situation event workspaces when the following situations resolve to true:

- KS3_Vol_Cache_DFW_Retry_Critical/Warning
- KS3_Vol_Cache_Read_HitP_Critical/Warning
- KS3_Vol_Cache_Write_HitP_Critical/Warning
- KS3_Vol_Perf_Resp_Time_Critical/Warning

You can navigate from one of these situation event workspaces directly to the following workspaces to do direct situation analysis:

| Workspace for direct situation analysis | Comments |
|---|--|
| Dataset Performance Event | <p>Default navigation is to this workspace. The rows in the report are sorted in descending order by Response Time. The entries in the bar chart are sorted in descending order by Response Time. This workspace derives its data from the Dataset Performance attribute table.</p> <p>Use this workspace to determine which data sets on the volume incurred poor performance. The rows in the report show all data sets by response time, from bad to good, accessed over the RMF interval. The bar chart graphically shows the response time of the data sets on the volume accessed over the RMF interval, making it obvious which data sets suffered.</p> |
| Volume Performance Trend | See “SMS Storage Groups Performance workspace” on page 53. |

Links

Right-clicking on the link to the left of each row of tabular data produces a pop-up menu with links to the following workspaces:

| Workspace for direct situation analysis | Links to other workspaces |
|---|--|
| Dataset Performance Event | <ul style="list-style-type: none"> • Dataset Performance Detail • Dataset Details • Dataset Performance Summary Trend |
| Volume Performance Trend | See “SMS Storage Groups Performance workspace” on page 53. |

Direct situation analysis for the SMS Storage Groups Space workspace

The SMS Storage Groups Space workspace produces situation event workspaces when the following situations resolve to true:

- KS3_Storage_Gr_Pct_Free_Critical/Warning
- KS3_Vol_Free_Space_Pct_Critical/Warning

You can navigate from one of these situation event workspaces directly to the following workspaces to do direct situation analysis:

| Workspace for direct situation analysis | Comments |
|--|---|
| Situation: KS3_Storage_Gr_Pct_Free_Critical/Warning | |
| Storage Group Percent Free | <p>Default navigation is to this workspace. These are new situations that provide alerts when the percentage of free space in a storage group falls below specified thresholds. They alert the user when there is a danger of allocation failure due to a lack of storage in the group.</p> <p>The rows in the report and entries in the bar charts are sorted in ascending order by Percent Free Space. This workspace derives its data from the DASD Volume Space attribute table.</p> <p>Use this workspace to determine which volumes in the group have the least amount of available space. The report lists all volumes in the group. The 'Volume Free Space' chart immediately identifies those volumes with a low amount of available space. The 'Fragmentation' chart immediately identifies the volumes in the group that are fragmented.</p> |
| Storage Group Percent Free Trend | <p>This workspace derives its data from the Volume Group Summary attribute table.</p> <p>Use this workspace to determine if the lack of free space is a spike or a trend. If it's a spike then migrating data might resolve the issue; however, if it's a trend then it might be time to add a volume or more to the group. The report shows complete space data for the group. The plot chart enables the user to immediately determine if the lack of available space is a trend or a spike.</p> |
| Situation: KS3_Vol_Free_Space_Pct_Critical/Warning | |
| Dataset Space Event | <p>Default navigation is to this workspace. Navigation to the new 'Volume Space Trend' workspace is also available from the existing 'SMS Storage Group Volume Space' and 'User DASD Group Volume Space' workspaces.</p> <p>This workspace includes the Tracks Unused attribute from the Dataset Space attribute group. The rows in the report are sorted in descending order by Tracks Unused. This workspace derives its data from the Dataset Space attribute table (report) and the DASD Volume Space attribute table (pie chart).</p> <p>Use this workspace to determine what data sets on the volume have the most unused space allocated. The chart displays the available space on the volume.</p> |

| Workspace for direct situation analysis | Comments |
|---|--|
| Volume Space Trend | <p>The plot chart plots the volume Percent Free Space attribute. The report consists of the following attributes from the DASD Volume Space attribute table:</p> <ul style="list-style-type: none"> • Total Capacity Megabytes • Free Space Megabytes • Percent Free Space • Fragmentation Index • Largest Free Extent MB • VTOC Index Status • SMS Status <p>This workspace derives its data from the DASD Volume Space attribute table.</p> <p>Use this workspace to determine if the shortage of available space is a spike or a trend. The action taken to correct this problem differs depending on that information. The report provides complete space information for the volume. The plot chart displays the free space percent over time.</p> |

Links

Right-clicking on the link to the left of each row of tabular data produces a pop-up menu with links to the following workspaces:

| Workspace for direct situation analysis | Links to other workspaces |
|--|--|
| Situation: KS3_Storage_Gr_Pct_Free_Critical/Warning | |
| Storage Group Percent Free | Dataset Space Summary Device Space Details |
| Storage Group Percent Free Trend | SMS Storage Group Volume Space History |
| Situation: KS3_Vol_Free_Space_Pct_Critical/Warning | |
| Dataset Space Event | Dataset Details |
| Volume Space Trend | None |

Direct situation analysis for the HSM Status workspace

The HSM Status workspace produces situation event workspaces when the following situations resolve to true:

- KS3_HSM_Backup_Queue_Critical/Warning
- KS3_HSM_Dump_Queue_Critical/Warning
- KS3_HSM_Migrate_Queue_Critical/Warning
- KS3_HSM_Recall_Queue_Critical/Warning
- KS3_HSM_Recovery_Queue_Critical/Warning

You can navigate from one of these situation event workspaces directly to the following workspace to do direct situation analysis: HSM Requests Details. See “DFSMSHsm Status workspace” on page 73.

Chapter 21. Situations

This chapter describes the predefined situations of the product.

Overview of situations

You can use the predefined situations shipped with Tivoli OMEGAMON XE for Storage on z/OS as-is or modify them to meet your requirements. If you choose to modify a predefined situation, first make a copy to ensure a fallback, if necessary. You can also create your own situations using the attributes provided by Tivoli OMEGAMON XE for Storage on z/OS.

Note: Do not modify the product-provided situations. If you want to modify a product-provided situation, copy the situation, modify the copy, and rename the copy.

Definition of a predefined situation

A situation is a logical expression involving one or more system conditions. Tivoli OMEGAMON XE for Storage on z/OS uses situations to monitor the systems in your network. To improve the speed with which you begin using Tivoli OMEGAMON XE for Storage on z/OS, the product provides situations that check for system conditions common to many enterprises. You can examine and if necessary, change the conditions or values being monitored to those best suited to your enterprise. Be sure to start the situations that you want to run in your environment.

Using situations

You manage situations from the Tivoli management portal using the Situation editor. Using the Situation editor you can perform the following tasks:

- Create a situation
- Save a situation
- Display a situation
- Edit a situation
- Start, stop, or delete a situation
- Investigate the situation event workspace for a situation

When you open the Situation editor, the left frame initially lists the situations associated with the Navigator item you selected. When you click a situation name or create a new situation, the right frame of the Situation editor opens to provide the following information about the situation and allow you to further define that situation:

Condition

View, add to, and edit the condition being tested.

Distribution

View the systems to which the situation is assigned and assign the situation to systems.

Expert advice

Write comments or instructions to be read in the situation event workspace.

Action

Specify a command to be sent to the system.

You can also enter Take Action commands, add a Take Action view to a workspace, select Take Action from the pop-up menu for an item in the Navigator's physical view, or create take action commands and save them for later use.

Until Reset a true situation when another situation becomes true or a specified time interval elapses.

Predefined situations descriptions

The following predefined situations are included in the Tivoli OMEGAMON XE for Storage on z/OS product.

KS3_Applic_Resp_Time_Critical

If VALUE S3_Application_Monitoring.High_Dataset_MSR GE 50

Monitors the response time components to determine the reason for a poor response time when an application is accessing a data set and the response time is greater than the critical threshold. Also examine the volume for over-utilization, cache settings, and the response time components at the volume level.

KS3_Applic_Resp_Time_Warning

If VALUE S3_Application_Monitoring.High_Dataset_MSR GE 40 and
VALUE S3_Application_Monitoring.High_Dataset_MSR LT 50

Monitors the response time components to determine the reason for a poor response time when an application is accessing a data set and the response time is greater than the warning threshold. Also examine the volume for over-utilization, cache settings, and the response time components at the volume level.

KS3_Cachecu_Cache_Stat_Critical

If VALUE S3_Cache_Control_Unit.Cache_Status NE Active

Monitors for the condition where caching is not active for the control unit. Use the SETCACHE command to activate caching, if appropriate.

KS3_Cachecu_DFW_Retry_Critical

If VALUE S3_Cache_Control_Unit.DFW_Retry_Percent GE 2

Monitors for the condition where the percent of DASD fast write attempts that cannot be satisfied because a shortage of available nonvolatile storage (NVS) space exceeds the critical threshold. Check for pinned NVS and correct the problem if NVS is pinned. Otherwise, if the impact on performance is not acceptable, you need to move a volume or data set to another cache control unit or add NVS to this control unit.

KS3_Cachecu_DFW_Retry_Warning

If VALUE S3_Cache_Control_Unit.DFW_Retry_Percent GE 1 and
VALUE S3_Cache_Control_Unit.DFW_Retry_Percent LT 2

Monitors for the condition where the percent of DASD fast write attempts that cannot be satisfied because a shortage of available nonvolatile storage (NVS) space has exceeded the warning threshold. Check for pinned NVS and correct the problem if NVS is pinned. Otherwise, if the impact on performance is not acceptable, move a volume or data set to another cache control unit or add NVS to this control unit.

KS3_Cachecu_Inact_Vols_Critical

If VALUE S3_Cache_Control_Unit.Deactivated_Volumes GE 15

Monitors for the condition where the number of deactivated volumes on the control unit exceeds the critical threshold. You can use the SETCACHE command to activate caching on the volumes, if necessary.

KS3_Cachecu_Inact_Vols_Warning

If VALUE S3_Cache_Control_Unit.Deactivated_Volumes GE 10 and
VALUE S3_Cache_Control_Unit.Deactivated_Volumes LT 15

Monitors for the condition where the number of deactivated volumes on the control unit exceeds the warning threshold. You can use the SETCACHE command to activate caching on the volumes, if necessary.

KS3_Cachecu_NVS_Stat_Critical

If Value S3_Cache_Control_Unit.NVS_Status NE Active

Monitors for the condition where nonvolatile storage is not active for the control unit. All writes to volumes on the control unit are written directly to the hard disk drive. Use the SETCACHE command to activate NVS (nonvolatile storage), if appropriate.

KS3_Cachecu_Read_HitP_Critical

If VALUE S3_Cache_Control_Unit.Read_Hit_Percent LE 50 and
VALUE S3_Cache_Control_Unit.Read_Hit_Percent GT 0

Monitors for the condition where the percent of read I/O requests resolved from cache has fallen below the critical threshold. If performance is a problem, look for volume with a low read hit percent and consider moving them to another control unit to balance the load. This condition can be caused by cache-unfriendly applications or a shortage of cache.

KS3_Cachecu_Read_HitP_Warning

If VALUE S3_Cache_Control_Unit.Read_Hit_Percent LE 60 and
VALUE S3_Cache_Control_Unit.Read_Hit_Percent GT 50

Monitors for the condition where the percent of read I/O requests resolved from cache has fallen below the warning threshold. If performance is a problem, look for volume with a low read hit percent and consider moving them to another control unit to balance the load. This condition can be caused by cache-unfriendly applications or a shortage of cache.

KS3_Cachecu_Trk_Dstg_Critical

If VALUE S3_Cache_Control_Unit.Track_Destaging_Rate GE 70

Monitors for the condition where the rate at which tracks are being removed from cache and written to DASD exceeds the critical threshold. If performance is being impacted, you need to migrate data sets or volumes to another cache control unit. An alternative is to increase the cache capacity.

KS3_Cachecu_Trk_Dstg_Warning

If VALUE S3_Cache_Control_Unit.Track_Destaging_Rate GE 50 and
VALUE S3_Cache_Control_Unit.Track_Destaging_Rate LT 70

Monitors for the condition where the rate at which tracks are being removed from cache and written to DASD exceeds the warning threshold. If performance is being impacted, you need to migrate data sets or volumes to another cache control unit. An alternative is to increase the cache capacity.

KS3_Cachecu_Trk_Stag_Critical

IF VALUE S3_Cache_Control_Unit.Track_Staging_Rate GE 70

Monitors for the condition where the movement of tracks from the physical device to cache has exceeded the critical threshold. If performance is impacted, you might need to move the logical volume that is causing the excessive activity or to move data sets on the logical volume.

KS3_Cachecu_Trk_Stag_Warning

IF VALUE S3_Cache_Control_Unit.Track_Staging_Rate GE 50 and
VALUE S3_Cache_Control_Unit.Track_Staging_Rate LT 70

Monitors for the condition where the movement of tracks from the physical device to cache has exceeded the warning threshold. If performance is impacted, you might need to move the logical volume that is causing the excessive activity or to move data sets on the logical volume.

KS3_Cachecu_Write_HitP_Critical

If VALUE S3_Cache_Control_Unit.Write_Hit_Percent LE 45 and
VALUE S3_Cache_Control_Unit.Write_Hit_Percent GE 0

Monitors for the condition where the percent of DASD/Cache fast write commands that were successfully processed without accessing the volume is below the critical threshold. If performance is impacted you might need to move a volume or data set to another control unit to balance the workload.

KS3_Cachecu_Write_HitP_Warning

If VALUE S3_Cache_Control_Unit.Write_Hit_Percent LE 50 and
VALUE S3_Cache_Control_Unit.Write_Hit_Percent GT 45

Monitors for the condition where the percent of DASD/Cache fast write commands that were successfully processed without accessing the volume is below the warning level. If performance is impacted you might need to move a volume or data set to another control unit to balance the workload.

KS3_Channel_Busy_Pct_Critical

If VALUE S3_Channel_Path.Complex_Percent_Utilized GE 85

Monitors high response time for I/O requests to volumes being serviced by the channel due to over utilization of that channel. You might need to balance the workload between channels by moving volumes or data sets.

KS3_Channel_Busy_Pct_Warning

If VALUE S3_Channel_Path.Complex_Percent_Utilized GE 70 and
VALUE S3_Channel_Path.Complex_Percent_Utilized LT 85

Monitors high response time for I/O requests to volumes being serviced by the channel due to over utilization of that channel. You might need to balance the workload between channels by moving volumes or data sets.

KS3_HSM_Backup_Held_Critical

If VALUE S3_HSM_Function_Summary.Function_Status EQ Held and
VALUE S3_HSM_Function_Summary.Function EQ Backup

Monitors the HSM backup function to see if it is being held. If the hold is inadvertent, issue the HSM RELEASE BACKUP command to allow the backup function to continue processing.

KS3_HSM_Backup_Queue_Critical

If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 50 and
VALUE S3_HSM_Function_Summary.Function EQ Backup

Monitors the HSM backup queue for a condition where the number of backup requests waiting exceeds the critical threshold. If the number of backup tasks is not at the maximum, use the HSM SETSYS MAXBACKUPTASKS command to increase the number of backup tasks, thus increasing the processing rate. Keep in mind that the number of available backup volumes serves as a constraint on the number of active backup tasks.

KS3_HSM_Backup_Queue_Warning

If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 15 and
VALUE S3_HSM_Function_Summary.Waiting_Requests LT 50 and
VALUE S3_HSM_Function_Summary.Function EQ Backup

Monitors the HSM backup queue for a condition where the number of backup requests waiting exceeds the warning threshold. If the number of backup tasks is not at the maximum, use the HSM SETSYS MAXBACKUPTASKS command to increase the number of backup tasks, thus increasing the processing rate. Keep in mind that the number of available backup volumes serves as a constraint on the number of active backup tasks.

KS3_HSM_Dump_Held_Critical

If VALUE S3_HSM_Function_Summary.Function_Status EQ Held and
VALUE S3_HSM_Function_Summary.Function EQ Dump

Monitors the HSM dump function to see if it is being held. If the hold is inadvertent, issue the HSM RELEASE DUMP command to allow dump processing to continue.

KS3_HSM_Dump_Queue_Critical

If VALUE S3_HSM_Function_Summary.Function_Status EQ Held and
VALUE S3_HSM_Function_Summary.Function EQ Dump

Monitors the HSM dump queue for a condition where the number of dump requests waiting exceeds the critical threshold. If the number of dump tasks is not at the maximum, use the HSM SETSYS MAXDUMPTASKS command to increase the number of dump tasks, thus increasing the processing rate. Keep in mind that the number of available tape drives serves as a constraint on the number of active dump tasks.

KS3_HSM_Dump_Queue_Warning

If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 15 and
VALUE S3_HSM_Function_Summary.Function EQ Dump and
VALUE S3_HSM_Function_Summary.Waiting_Requests LT 50

Monitors the HSM dump queue for a condition where the number of dump requests waiting exceeds the warning threshold. If the number of dump tasks is not at the maximum, use the HSM SETSYS MAXDUMPTASKS command to increase the number of dump tasks, thus increasing the processing rate. Keep in mind that the number of available tape drives serves as a constraint on the number of active dump tasks.

KS3_HSM_Inactive_Host_Warning

IF VALUE S3_HSM_Status.Inactive_HSM_Hosts GT 0

Monitors when an inactive HSM host has been detected. The event workspace for this situation has a link to the DFSMSHsm Host Details workspace.

KS3_HSM_Migrate_Held_Critical

If VALUE S3_HSM_Function_Summary.Function_Status EQ Held and
VALUE S3_HSM_Function_Summary.Function EQ Migration

Monitors the migrate function to see if it is being held. If the hold on the function is inadvertent, issue the HSM RELEASE MIGRATION command to allow migration to continue.

KS3_HSM_Migrate_Queue_Critical

If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 50 and
VALUE S3_HSM_Function_Summary.Function EQ Migration

Monitors the HSM migration queue for a condition where the number of migration requests waiting exceeds the critical threshold. If the number of migrate tasks is not at the maximum, use the HSM SETSYS MAXMIGRATIONTASKS command to increase the number of migration tasks, thus increasing the processing rate. Note that this affects only those migrations requested by automatic functions. Only one task is available to process command migration requests.

KS3_HSM_Migrate_Queue_Warning

If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 15 and
VALUE S3_HSM_Function_Summary.Waiting_Requests LT 50 and
VALUE S3_HSM_Function_Summary.Function EQ Migration

Monitors the HSM migration queue for a condition where the number of migration requests waiting exceeds the warning threshold. If the number of migrate tasks is not at the maximum, use the HSM SETSYS MAXMIGRATIONTASKS command to increase the number of migration tasks, thus increasing the processing rate. Note that this affects only those migrations requested by automatic functions. Only one task is available to process command migration requests.

KS3_HSM_Recall_Held_Critical

If VALUE S3_HSM_Function_Summary.Function_Status EQ Held and
VALUE S3_HSM_Function_Summary.Function EQ Recall

Monitors the recall function to see if it is being held. If the hold on the function is inadvertent, issue the HSM RELEASE RECALL command to allow recalls to resume.

KS3_HSM_Recall_Queue_Critical

If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 50 and
VALUE S3_HSM_Function_Summary.Function EQ Recall

Monitors the HSM recall queue for a condition where the number of recall requests waiting exceeds the critical threshold. If the number of recall tasks is not at the maximum, use the HSM SETSYS MAXRECALLTASKS command to increase the number of recall tasks, thus increasing the processing rate.

KS3_HSM_Recall_Queue_Warning

If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 15 and
VALUE S3_HSM_Function_Summary.Waiting_Requests LT 50 and
VALUE S3_HSM_Function_Summary.Function EQ Recall

Monitors the HSM recall queue for a condition where the number of recall requests waiting exceeds the critical threshold. If the number of recall tasks is not at the maximum, use the HSM SETSYS MAXRECALLTASKS command to increase the number of recall tasks, thus increasing the processing rate.

KS3_HSM_Recovery_Held_Critical

If VALUE S3_HSM_Function_Summary.Function_Status EQ Held and
VALUE S3_HSM_Function_Summary.Function EQ Recovery

Monitors the recovery function to see if it is being held. If the hold on the function is inadvertent, issue the HSM RELEASE RECOVER command to allow recovery function to resume.

KS3_HSM_Recovery_Queue_Critical

If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 50 and
VALUE S3_HSM_Function_Summary.Function EQ Recovery

Monitors the HSM recovery queue for a condition where the number of recover requests waiting exceeds the critical threshold. If the number of recovery tasks is not at the maximum, use the HSM SETSYS MAXDSRECOVERTASKS command to increase the number of recover tasks, thus increasing the processing rate. Keep in mind that the number of backup tape cartridges serves as a constraint on the number of active recovery tasks.

KS3_HSM_Recovery_Queue_Warning

If VALUE S3_HSM_Function_Summary.Waiting_Requests GE 15 and
VALUE S3_HSM_Function_Summary.Waiting_Requests LT 50 and
VALUE S3_HSM_Function_Summary.Function EQ Recovery

Monitors the HSM recovery queue for a condition where the number of recover tasks waiting exceeds the warning threshold. If the number of recovery tasks is not at the maximum, use the HSM SETSYS MAXDSRECOVERTASKS command to increase the number of recover tasks, thus increasing the processing rate. Keep in mind that the number of backup tape cartridges serves as a constraint on the number of active recovery tasks.

KS3_HSM_Status_Inactive_Crit

If VALUE S3_HSM_Status.HSM_Status EQ InActive

Monitors the status of the HSM. If status is not active, restart HSM.

KS3_LCU_Av_Delay_Q_Critical

If VALUE S3_Logical_Control_Unit.Average_Delay_Queue GE 0.500

Monitors for the condition where the average number of requests queued to devices assigned to a logical control unit due to busy conditions on physical paths has exceeded the critical threshold. If performance is impacted, you might be able to balance the workload across multiple LCUs by moving a volume or data set. Otherwise, you need to add physical paths to the LCU.

KS3_LCU_Av_Delay_Q_Warning

If VALUE S3_Logical_Control_Unit.Average_Delay_Queue GE 0.2 and
VALUE S3_Logical_Control_Unit.Average_Delay_Queue LT 0.500

Monitors for the condition where the average number of requests queued to devices assigned to a logical control unit due to busy conditions on physical paths has exceeded the warning threshold. If performance is impacted, you might be able to balance the workload across multiple LCUs by moving a volume or data set. Otherwise, you need to add physical paths to the LCU.

KS3_LCU_Cont_Rate_Critical

If VALUE S3_Logical_Control_Unit.Contention_Rate GE 1.001

Monitors for the condition where the rate at which I/O requests are being queued to devices on a logical control unit (LCU) due to busy conditions on physical paths has exceeded the critical threshold. If performance is impacted, you need to migrate volumes or data sets to another LCU, otherwise, you need to add physical paths to the LCU.

KS3_LCU_Cont_Rate_Warning

If VALUE S3_Logical_Control_Unit.Contention_Rate GE 0.2 and
VALUE S3_Logical_Control_Unit.Contention_Rate LT 1.001

Monitors for the condition where the rate at which I/O requests are being queued to devices on a logical control unit (LCU) due to busy conditions on physical paths has exceeded the warning threshold. If performance is impacted, you need to migrate volumes or data sets to another LCU, otherwise, you need to add physical paths to the LCU.

KS3_LCU_IO_Rate_Sec_Critical

If VALUE S3_Logical_Control_Unit.Channel_Path_I/O_Rate GE 600

Monitors for the condition where the I/O rate per second to volumes in the logical control unit (LCU) has exceeded the critical threshold. If performance is impacted, you need to balance the workload across multiple LCUs by moving volumes or data sets.

KS3_LCU_IO_Rate_Sec_Warning

If VALUE S3_Logical_Control_Unit.Channel_Path_I/O_Rate GE 200 and
VALUE S3_Logical_Control_Unit.Channel_Path_I/O_Rate LT 600

Monitors for the condition where the I/O rate per second to volumes in the logical control unit (LCU) has exceeded the warning threshold. If performance is impacted, you need to balance the workload across multiple LCUs by moving volumes or data sets.

KS3_RMM_CDS_Backup_Critical

IF VALUE S3_RMM_Control_Dataset.Days_Since_Last_Backup GT 3

The number of days since the last backup of the DFSMSrmm CDS or Journal exceeded the critical threshold.

KS3_RMM_CDS_Backup_Warning

IF VALUE S3_RMM_Control_Dataset.Days_Since_Last_Backup GT 1
AND VALUE S3_RMM_Control_Dataset.Days_Since_Last_Backup LE 3

The number of days since the last backup of the DFSMSrmm CDS or Journal exceeded the warning threshold.

KS3_RMM_CDS_Space_Critical

IF VALUE S3_RMM_Control_Dataset.RMM_Percent_Used GT 90

The percentage of space used by the DFSMSrmm CDS or Journal is greater than the critical threshold.

KS3_RMM_CDS_Space_Warning

IF VALUE S3_RMM_Control_Dataset.RMM_Percent_Used GE 80 AND
VALUE S3_RMM_Control_Dataset.RMM_Percent_Used LE 90

The percentage of space used by the DFSMSrmm CDS or Journal is greater than the warning threshold.

KS3_RMM_Exit_Status_Critical

IF ((VALUE S3_RMM_Config.EDGUX200_Status NE Enabled) OR
(VALUE S3_RMM_Config.EDGUX100_Status NE Enabled))

The DFSMSrmm EDGUX100 or EDGUX200 exit is not Enabled.

KS3_RMM_Journal_Status_Critical

IF VALUE S3_RMM_Config.Journal_Status NE Enabled

The DFSMSRmm Journal is either Disabled or Locked. DFSMSRmm does not allow further updates to the journal until BACKUP is run to back up the DFSMSRmm control data set and to clear the journal. If the Journal is Locked, DFSMSRmm fails any requests that result in an update to the DFSMSRmm control data set. Message EDG2103D might also have been issued to the DFSMSRmm operator console.

KS3_RMM_Operating_Mode_Warning

IF VALUE S3_RMM_Config.Operating_Mode NE Protect

DFSMSRmm is not operating in Protect mode. Certain actions that should be rejected are permitted if DFSMSRmm is not operating in protect mode, for example attempting to read a scratch tape volume.

KS3_RMM_Scratch_Tape_Critical

IF VALUE S3_RMM_Summary.Type EQ 0 AND
VALUE S3_RMM_Summary.Scratch_Volumes LT 100

The number of Scratch volumes is below the critical threshold.

KS3_RMM_Scratch_Tape_Warning

IF VALUE S3_RMM_Summary.Type EQ 0 AND
VALUE S3_RMM_Summary.Scratch_Volumes LT 200 AND
VALUE S3_RMM_Summary.Scratch_Volumes GE 100

The number of Scratch volumes is below the warning threshold.

KS3_RMM_Inactive_Critical

IF VALUE S3_RMM_Config.Subsystem_Status EQ Inactive

The DFSMSRmm subsystem is inactive.

KS3_Stg_Toolkit_Result_Critical

If VALUE S3_Storage_Toolkit_Result_Summary.Return_Code GT 4

A DFSMSHsm or DFSMSdss command issued by the Storage Toolkit returns a value greater than 4. The DFSMSHsm or DFSMSdss facilities typically return a value greater than 4 when a command fails to complete. Check the messages in the action log to determine whether the warning requires further attention.

KS3_Stg_Toolkit_Result_Warning

If VALUE S3_Storage_Toolkit_Result_Summary.Return_Code EQ 4

A DFSMSHsm or DFSMSdss command issued by the Storage Toolkit returns the value 4. The DFSMSHsm or DFSMSdss facilities typically return the value 4 as a warning. Check the messages in the action log to determine whether the warning requires further attention.

KS3_Storage_Gr_Pct_Free_Crit

IF VALUE S3_Volume_Group_Summary.Free_Space_Percent LT 5.0 AND
VALUE S3_Volume_Group_Summary.Group_Type EQ SMSGROUP AND
VALUE S3_Volume_Group_Summary.Free_Space_Percent GE 0.0

Monitors the percentage of free space available for allocation in the storage group and detects when free space has dropped below the critical threshold. To prevent allocation failures, you might have to either add one or more logical volumes to the storage group, or to move data sets off of the logical volumes in the storage group.

KS3_Storage_Gr_Pct_Free_Warning

```
IF VALUE S3_Volume_Group_Summary.Free_Space_Percent LT 10.0 AND  
VALUE S3_Volume_Group_Summary.Group_Type EQ SMSGROUP AND  
VALUE S3_Volume_Group_Summary.Free_Space_Percent GE 5.0",
```

Monitors the percentage of free space available for allocation in the storage group and detects when free space has dropped below the warning threshold. In order to prevent allocation failures, you might have to either add one or more logical volumes to the storage group, or to migrate data sets off of the logical volumes in the storage group.

KS3_Vol_Cache_DFW_Retry_Critical

```
If VALUE S3_Cache_Devices.DFW_Retry_Percent GE 2 and  
VALUE S3_Cache_Devices.I/O_Count GE 25
```

Monitors for the condition where the percentage of DASD fast write attempts for a volume that cannot be satisfied due to a shortage of available nonvolatile storage (NVS) space exceeded the critical threshold. Check for pinned NVS and correct the problem if NVS is pinned. Otherwise, if the impact on performance is not acceptable, move a volume or data set to another cache control unit or to add NVS to this control unit.

KS3_Vol_Cache_DFW_Retry_Warning

```
If VALUE S3_Cache_Devices.DFW_Retry_Percent GE 1 and  
VALUE S3_Cache_Devices.DFW_Retry_Percent LT 2 and  
VALUE S3_Cache_Devices.I/O_Count GE 25
```

Monitors for the condition where the percentage of DASD fast write attempts for a volume that cannot be satisfied due to a shortage of available nonvolatile storage (NVS) space exceeded the warning threshold. Check for pinned NVS and correct the problem if NVS is pinned. Otherwise, if the impact on performance is not acceptable, move a volume or data set to another cache control unit or to add NVS to this control unit.

KS3_Vol_Cache_Read_HitP_Critical

```
If VALUE S3_Cache_Devices.Read_Hit_Percent LE 45 and  
VALUE S3_Cache_Devices.Read_Hit_Percent GE 0 and  
VALUE S3_Cache_Devices.I/O_Count GE 25
```

Monitors for the condition where the cache read hit percent is below the critical threshold. If performance is impacted determine the reason for the low read hit percent. Common problems are cache-unfriendly applications and over-utilization of the control unit.

KS3_Vol_Cache_Read_HitP_Warning

```
If VALUE S3_Cache_Devices.Read_Hit_Percent LE 55 and  
VALUE S3_Cache_Devices.Read_Hit_Percent GT 45 and  
VALUE S3_Cache_Devices.I/O_Count GE 25
```

Monitors for the condition where the cache read hit percent is below the warning threshold. If performance is impacted determine the reason for the low read hit percent. Common problems are cache-unfriendly applications and over-utilization of the control unit.

KS3_Vol_Cache_Writ_HitP_Critical

If VALUE S3_Cache_Devices.Write_Hit_Percent LE 20 and
VALUE S3_Cache_Devices.Write_Hit_Percent GE 0 and
VALUE S3_Cache_Devices.I/O_Count GE 25

Monitors for the condition where the cache write hit percent for a volume is below the critical threshold. Check the status of the nonvolatile storage in the cache control unit. You can move volumes or data sets to balance the workload.

KS3_Vol_Cache_Writ_HitP_Warning

If VALUE S3_Cache_Devices.Write_Hit_Percent LE 30 and
VALUE S3_Cache_Devices.Write_Hit_Percent GT 20 and
VALUE S3_Cache_Devices.I/O_Count GE 25

Monitors for the condition where the cache write hit percent for a volume is below the warning threshold. Check the status of the nonvolatile storage in the cache control unit. You can move volumes or data sets to balance the workload.

KS3_Vol_Disabled_VTOC_Critical

If VALUE S3_DASD_Volume_Space.VTOC_Index_Status EQ Disabled

Monitors for the condition where a VTOC index has been disabled. This condition can degrade performance on the volume. Enable the VTOC index.

KS3_Vol_EAV_Fragment_Index_Crit

IF VALUE S3_DASD_Volume_Space.Extended_Address_Volume EQ Yes AND
VALUE S3_DASD_Volume_Space.Track_Managed_Fragmentation_Index GE 850

The fragmentation index in the track managed area of an Extended Address Volume exceeds the critical threshold.

KS3_Vol_EAV_Fragment_Index_Warn

IF VALUE S3_DASD_Volume_Space.Extended_Address_Volume EQ Yes AND
VALUE S3_DASD_Volume_Space.Track_Managed_Fragmentation_Index GE 650
AND VALUE S3_DASD_Volume_Space.Track_Managed_Fragmentation_Index LT 850

The fragmentation index in the track managed area of an Extended Address Volume exceeds the warning threshold.

KS3_Vol_EAV_Free_Space_Pct_Crit

IF VALUE S3_DASD_Volume_Space.Track_Managed_Percent_Free LE 5.0
AND VALUE S3_DASD_Volume_Space.Track_Managed_Percent_Free GE 0.0 AND
VALUE S3_DASD_Volume_Space.Extended_Address_Volume EQ Yes

The percentage of free space in the track managed area of an Extended Address Volume is below the critical threshold.

KS3_Vol_EAV_Free_Space_Pct_Warn

IF VALUE S3_DASD_Volume_Space.Track_Managed_Percent_Free LE 10.0
AND VALUE S3_DASD_Volume_Space.Track_Managed_Percent_Free GT 5.0 AND
VALUE S3_DASD_Volume_Space.Extended_Address_Volume EQ Yes

The percentage of free space in the track managed area of an Extended Address Volume is below the warning threshold.

KS3_Vol_Fragment_Index_Critical

If VALUE S3_DASD_Volume_Space.Fragmentation_Index GE 850

Monitors for the condition where a volume has a fragmentation index that exceeds the critical threshold. Defragment the volume so that free extents are combined to help prevent data set allocation failures.

KS3_Vol_Fragment_Index_Warning

If VALUE S3_DASD_Volume_Space.Fragmentation_Index GE 650 and
VALUE S3_DASD_Volume_Space.Fragmentation_Index LT 850

Monitors for the condition where a volume has a fragmentation index that exceeds the warning threshold. Defragment the volume so that free extents are combined to help prevent data set allocation failures.

KS3_Vol_Free_Space_Pct_Critical

If VALUE S3_DASD_Volume_Space.Percent_Free_Space LE 5 and
VALUE S3_DASD_Volume_Space.Percent_Free_Space GE 0

Monitors for the condition where the percentage of free space on a volume is below the critical threshold. If data sets on the volume require more space, then either migrate some data sets to another volume or release space from data sets that might be over-allocated.

KS3_Vol_Free_Space_Pct_Warning

If VALUE S3_DASD_Volume_Space.Percent_Free_Space LE 10 and
VALUE S3_DASD_Volume_Space.Percent_Free_Space GT 5

Monitors for the condition where the percentage of free space on a volume is below the warning threshold. If data sets on the volume require more space, then either migrate some data sets to another volume or release space from data sets that might be over-allocated.

KS3_Vol_Perf_Resp_Time_Critical

If VALUE S3_DASD_Volume_Performance.Response_Time GE 55 and
VALUE S3_DASD_Volume_Performance.I/O_Count GE 25

Monitors for the condition where response time for the volume exceeds the critical threshold. Look at the volume to see if high utilization is a problem. If so, it might be necessary to migrate data sets from the volume to reduce utilization. Also check the cache status of the volume. Look at the components of I/O to determine where the time is being spent and address the problem accordingly.

KS3_Vol_Perf_Resp_Time_Warning

If VALUE S3_DASD_Volume_Performance.Response_Time GE 35 and
VALUE S3_DASD_Volume_Performance.Response_Time LT 55 and
VALUE S3_DASD_Volume_Performance.I/O_Count GE 25

Monitors for the condition where response time for the volume exceeds the warning threshold. Look at the volume to see whether high utilization is a problem. If so, you can migrate data sets from the volume to reduce utilization. Also check the cache status of the volume. Look at the components of I/O to determine where the time is being spent and address the problem accordingly.

KS3_VTS_Disconnect_Time_Crit

If VALUE S3_VTS_Overview.Virtual_Disconnect_Time GE 500

Monitors for the condition where the logical control unit disconnect time for the virtual tape server exceeds the critical threshold. This condition is often an indication that the tape volume cache capacity is being exceeded.

KS3_VTS_Host_GB_Warning

If VALUE S3_VTS_Overview.Host_Channel_Activity_GB GE 18

Monitors for the condition where the activity between the MVS system and the virtual tape server on the host channels exceeds 19 GB over the hour interval. This condition can be an indication that the virtual tape server is being overloaded.

KS3_VTS_Pct_Copy_Throt_Warn

IF VTSTPVOLC.PCTCPT > 50

Monitors for the condition where copy is the predominant reason for throttling.

KS3_VTS_Pct_Wr_Over_Throt_Warn

IF VTSTPVOLC.PCTWROT > 50

Monitors for the condition where write overrun is the predominant reason for throttling.

KS3_VTS_Recall_Pct_Warning

If VALUE S3_VTS_Overview.Volume_Recall_Percent GE 20

Monitors for the condition where the percent of virtual tape mounts that required a physical tape mount to be satisfied exceeded the warning threshold. This condition can lead to unacceptably large virtual mount times. If so, then investigate the reason for the recalls. If rescheduling or removing the application workload is not possible, you need to increase the cache capacity of the VTS.

KS3_VTS_Virt_MtPend_Av_Warning

If VALUE S3_VTS_Overview.Average_Virtual_Mount_Pend_Time GE 300

Monitors for the condition where the average seconds required to satisfy a virtual mount in the virtual tape subsystem exceeded the warning threshold. If this condition persists, then further study is required to determine the cause for the elongated mount times. The condition might be due to VTS-hostile applications or to a shortage of VTS resources.

KS3_VTS_Virt_MtPend_Mx_Warning

VALUE S3_VTS_Overview.Maximum_Virtual_Mount_Pend_Time EQ 900

Monitors for the condition where the maximum seconds required to satisfy a virtual mount in the virtual tape subsystem exceeded the warning threshold. If this condition persists, then further study is required to determine the cause for the elongated mount times. The condition might be due to VTS-hostile applications or to a shortage of VTS resources.

Chapter 22. Attributes

Tivoli Enterprise Portal gathers data from remote agents located within your systems. The portal stores this data for OMEGAMON XE for Storage on z/OS in elements called attributes.

Each attribute is a characteristic of an object. For example, the All Channels Busy Percent attribute in the Logical Control Unit attribute group displays the percentage of the time that all channel paths to an LCU were busy.

Attribute groups

The OMEGAMON XE for Storage on z/OS attributes are organized into groups of related items. These groups of attributes comprise the attribute tables for the product.

Note: In the Tivoli Enterprise Portal, the names of attribute groups are preceded by "S3," which is the product code for Tivoli OMEGAMON XE for Storage on z/OS.

Attributes and workspaces

You can view the data reported for the attribute in the reports and charts displayed in the OMEGAMON XE for Storage on z/OS workspaces. The workspaces contain the table views you use to obtain information about the various aspects of your environment. The table views and graphs within each workspace report the storage attribute information you are monitoring. You can use the collected data to analyze and monitor performance.

Attributes and situations

Attributes are used in the predefined situations for the product. You can also create custom situations for to monitor the performance of your z/OS systems. These situations can monitor one aspect of your enterprise or analyze multiple conditions to alert you to problems that can occur when attribute values exceed their thresholds.

List of attribute groups and attributes

The following attribute groups are provided with the OMEGAMON XE for Storage on z/OS product:

- "Application Dataset Details" on page 127
- "Application Monitoring" on page 129
- "Application Volume Performance" on page 130
- "Application Volume Space" on page 132

- "Cache Control Unit" on page 133
- "Cache Devices" on page 134
- "Cache Raid Rank" on page 138

- "Channel Path" on page 138

- "Cross System Volume Summary" on page 140
- "Cross System Volumes" on page 141

- "DASD Physical Group" on page 142
- "DASD Volume Performance" on page 145
- "DASD Volume Space" on page 147

These attribute groups supply information to a database that tracks attributes related to space, DCB, and cataloging of data sets. This database collects data regarding all data sets in the z/OS environment so that you can create comprehensive reports or identify alert conditions throughout the system. See Chapter 25, “Dataset Attributes Database,” on page 311 for more information.

“Dataset Attributes Blocksize Summary” on page 149
“Dataset Attributes CA Split Summary” on page 149
“Dataset Attributes Catalog Summary” on page 150
“Dataset Attributes CI Split Summary” on page 151
“Dataset Attributes Creation Date Summary” on page 152
“Dataset Attributes Dataset Detail” on page 153
“Dataset Attributes DSORG Detail” on page 157
“Dataset Attributes DSORG Summary” on page 158
“Dataset Attributes Extent Summary” on page 159
“Dataset Attributes Inefficient Blocksize Summary” on page 160
“Dataset Attributes Mature Dataset Summary” on page 161
“Dataset Attributes Never Referenced Dataset Summary” on page 162
“Dataset Attributes Newborn Dataset Summary” on page 163
“Dataset Attributes Reference Date Summary” on page 163
“Dataset Attributes SMS Class Detail” on page 164
“Dataset Attributes SMS Class Summary” on page 165
“Dataset Attributes Space Allocated Summary” on page 166
“Dataset Attributes Space Unused Summary” on page 166
“Dataset Attributes System Summary” on page 167
“Dataset Attributes Uncataloged Detail” on page 169
“Dataset Attributes Uncataloged Summary” on page 170

“Dataset Group Details Volume” on page 175
“Dataset Group Details” on page 170
“Dataset Group Detail Attrs” on page 174
“Dataset Group Summary” on page 177
“Dataset Performance” on page 178
“Dataset SMS” on page 180
“Dataset Space” on page 180
“Dataset Users” on page 182

“HSM CDS” on page 182
“HSM Common Storage” on page 183
“HSM Dataset Recover Version” on page 184
“HSM Function Details” on page 184
“HSM Function Statistics” on page 185
“HSM Function Summary” on page 187
“HSM Host Function Status” on page 187
“HSM Host Request Summary” on page 189
“HSM Host Status” on page 189
“HSM Private Storage” on page 191
“HSM Requests” on page 191
“HSM Status” on page 192

“Logical Control Unit” on page 194
“Logical Control Unit Channel Paths” on page 195
“LCU Control Unit” on page 196

“RMM Configuration” on page 197
“RMM Control” on page 200
“RMM Control Data Set” on page 200
“RMM Dataset Details” on page 201
“RMM Location Details” on page 203
“RMM Owner Details” on page 204
“RMM Pending Actions” on page 204
“RMM System Summary” on page 205
“RMM Volume Details” on page 208
“RMM VRS Details” on page 214

“SMS Cache Sets” on page 215
“SMS Configuration” on page 215
“SMS Data Class” on page 216
“SMS Management Class” on page 217
“SMS Storage Class” on page 219
“SMS Storage Group” on page 220
“SMS Storage Group Status” on page 222
“SMS Systems” on page 222

“Symmetrix Configuration” on page 226
“Symmetrix Devices” on page 226
“Symmetrix Director Summary” on page 227
“Symmetrix Directors” on page 227
“Symmetrix SSIDs” on page 229

“Tape Device” on page 229
“Tape Group” on page 230

“Storage Toolkit Requests” on page 222
“Storage Toolkit Result Summary” on page 224
“Storage Toolkit Result Detail” on page 224

For more information related to this set of attributes, see Chapter 23, “Storage Toolkit,” on page 251.

“TotalStorageDS Array Configuration” on page 231
“TotalStorageDS Configuration” on page 232
“TotalStorageDS Extent Pool” on page 232
“TotalStorageDS Rank” on page 234
“TotalStorageDS SSIDs” on page 234

“Volume Group Summary” on page 235
“Volume Status” on page 238
“Volume Users” on page 238

“VTS Cache” on page 238
“VTS Capacity” on page 239
“VTS Composite Library” on page 241
“VTS Overview” on page 242
“VTS Physical Devices” on page 243
“VTS TS7700-series Cache Container” on page 244
“VTS TS7700-series Cache Partition Container” on page 245
“VTS TS7700-series Cache Preference Group” on page 246
“VTS TS7700-series Cluster” on page 247
“VTS Virtual Devices” on page 248
“VTS Virtual Devices Summary” on page 250

Application Dataset Details

The Application Dataset Details attributes provide space information on the data sets allocated to a specific application.

Block Size: Represents the physical length attribute of the records in this data set. For a DASD data set, block size is maintained in the DSCB. Because large blocks transfer more data, large block sizes result in

higher connect time. Larger blocks with fewer buffers might be appropriate for sequential data sets. Consider using DASD cache control for heavily used, read-only sequential data sets. Random access files need smaller block sizes with a larger number of PDS libraries, block sizes should usually reflect average member size.

Dataset Name: The MVS data set name.

Dataset Type: Represents the organization of this data set. This attribute can take the following decimal values:

- Direct_Access=11
- Direct_Access_Unmovable=12
- Extended_Seq_Unmovable=4
- Extended_Sequential=3
- Hierarchical_File=1
- ISAM=7
- ISAM_Unmovable=8
- Partitioned=13
- Partitioned_Unmovable=14
- PDS_Extended=2
- Physical_Seq_Unmovable=10
- Physical_Sequential=9
- Unknown=0
- Unknown=15
- VSAM=5
- VSAM_Unmovable=6

Job Name: The MVS JOBNAME, TSOUSERID, or STARTED TASK NAME associated with this application.

Logical Record Length: The length attribute associated with the logical records in this data set.

Number of Extents: The number of DASD extents allocated to this data set. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Managed System: The SMS name of this system.

Tracks Allocated: Represents the number of DASD tracks occupied by this data set. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Tracks Used: Represents the number of DASD tracks actually used by this data set. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Tracks Used Percent: Represents the value of Tracks Used divided by Tracks Allocated for this data set. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Volume: Represents the six-character identifier associated with this logical volume.

Application Dataset Monitor

This group of attributes provides information to other attribute groups. Users do not use these attributes. The attributes are documented for the purpose of background reference only.

Managed System: The SMS name of this system.

Application Monitoring

In addition to controlling the monitoring of applications with the Monitor Status and I/O Monitor Status columns, the Application Monitoring attributes return summary performance information for an application.

Application: The name of the batch job, TSO user ID, or started task associated with this application.

ASID Hex: The address space ID (in hex) associated with this application.

ASID: Indicates the number of the address space ID (in decimal) associated with this application.

Dataset Count: Represents the total number of data sets allocated to this application in this address space. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Dataset Name: Specifies the MVS data set name.

Dataset with High MSR: The MVS data set name with the longest response time in milliseconds associated with this application.

High Dataset MSR: Represents the longest response time in milliseconds for this application to access any data set. You can enable aggregation for this attribute in Tivoli Data Warehouse.

High Volume MSR: Represents the longest response time in milliseconds for any volume in use by this application. You can enable aggregation for this attribute in Tivoli Data Warehouse.

I/O Monitor Status: Used to toggle the status of I/O monitoring at the data set level for the application. This value is used for input only when defining a situation to start or stop monitoring of an application. This value is meaningless in a report, because it is always reported as an empty field. When specified as Y (Yes), it triggers MSR (millisecond response time) monitoring at the data set level for all data sets in use by the application.

Y = Turn ON data set-level I/O monitoring for this application.

N = Turn OFF data set-level I/O monitoring for this application.

I/O Second: Specifies the average number of I/O instructions per second issued by this application. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Managed System: The Storage Management Subsystem name of this system.

Monitor Status: Represents the current monitor status of this application. The values are: Y=The application is being monitored. N=The application is not being monitored.

SMF ID: The SYSTEM ID of the MVS system which is running this application.

Volume Count: Represents the total number of volumes allocated to this application in this address space. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Volume with High MSR Dataset: The volume label of the logical volume containing the data set with the longest response time in milliseconds for this application.

Volume with High MSR: The six-character identifier of the logical device with the longest response time in milliseconds in use by this application.

Volume: The six-character identifier associated with this logical volume.

Application Volume Performance

The Application Volume Performance attributes provide performance information for the logical volumes allocated to an application.

Average Command Response Delay: The average number of milliseconds that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device over the RMF interval.

Average HyperPAV Alias Count: The average number of aliased devices assigned to a base volume in a HyperPAV environment over the RMF interval.

Busy Percent: The ratio, expressed as a percentage, of the total busy time this interval divided by the interval length.

Bypass Cache Percent: The percentage of I/O that goes directly to a DASD device regardless of whether the data might be or currently is cached.

Cache Status: Indicates whether caching is enabled for this device. Values are Timeout, Active, Inactive, Pending Active, Pending Inactive, and n/a.

Cache to DASD Total Tracks: The number of tracks written from cache to physical disk during the cache interval. This value is captured asynchronously with respect to the transfer of data from the channel to the cache.

Cache to DASD Tracks per Sec: The rate per second of tracks written from cache to physical disk during the cache interval. This value is captured asynchronously with respect to the transfer of data from the channel to the cache.

CFW Read Percent: The percentage of I/O requests read from cache that were in cache due to a cache fast write.

CFW Write Percent: The percentage of I/Os that were cache fast write requests. Cache Fast Write (CFW) is a function of 3990 controllers that writes data to cache at channel speed. The data is later written to DASD. Cache fast write differs from DASD fast write in that the data is not written to NVS (nonvolatile storage).

Connect Time: The part of I/O during which data is located on the track and transferred between the DASD device and central storage.

Current PAV Exposures: The number of exposures for the device when the data was gathered.

DASD to Cache Non Promotes Second: The rate per second of operations that did not promote (copy) a track to cache because free cache segments were not available. Non promotes are tracks that the system does not copy to cache because free cache segments were not available.

DASD to Cache Non Promotes: The number of operations that did not promote (copy) a track to cache because free cache segments were not available. Non promotes are tracks that the system does not copy to cache because free cache segments were not available.

DASD to Cache Total Tracks Non-sequential: The number of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Total Tracks Sequential: The number of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

DASD to Cache Total Tracks: The number of tracks read from physical disk to cache over the cache interval for I/O operations.

DASD to Cache Tracks per Sec Non-sequential: The rate per second of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Tracks per Sec Sequential: The rate per second of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

DASD to Cache Tracks per Sec: The rate per second of tracks read from physical disk to cache over the cache interval for I/O operations.

DCBs Open: The count of open DCBs for data sets.

Device Address: The MVS device number, assigned by HCD, associated with this device.

Device MPL: The device Multi-Processing Level, which is the I/O density for a device. This metric is the product of the MSR (millisecond response time) and I/O per second for a device.

DFW Hit Percent: The ratio of I/O requests where data is written to both cache and nonvolatile storage that was satisfied by the storage control cache without requiring access to the hard disk.

DFW Retry Percent: The percentage of DASD fast write operations that did not complete due to NVS (nonvolatile storage) constraint.

DFW Status: Indicates whether DASD Fast Write is enabled on this device.

Disconnect Time: Time that I/O spends, freed from the channel, searching for or requesting data. This time includes moving the device to the requested cylinder and track, waiting for the record to rotate under the head (rotational delay), and waiting for the channel to become available.

I/O Count: The total number of I/O operations performed on this device during this interval.

I/O Per Second: The average number of input/output operations, per second, directed at this volume during this interval.

Inhibit Cache Percent: The percentage of I/O that used the MVS I/O mechanism to prevent a request from being cached, even if it is eligible. If the data is already in cache, the request is satisfied by cache.

IOSQ Delay: The average time that an I/O waits because the device is already busy (signified by the UCBBUSY bit for the device being on).

Job Name: The MVS JOBNAME, TSO user ID, or started task name associated with this application.

Managed System: The SMS name of this system.

Maximum PAV Exposures: The maximum number of exposures for a device over the current RMF interval.

MSR Connect Time Percent: The ratio of average connect time divided by average total response time.

PAV Exposure Changed: An indicator that the number of PAV exposures for the volume has changed over the interval.

Pend Time: The time that I/O is delayed in the path to the device. Pending time can be attributed to the channel or control unit path being busy.

Physical Device: The physical device type.

Read Hit Percent: The ratio of read requests that were satisfied by the storage control unit cache without requiring access to DASD.

Read I/O Count: The number of read requests in the current collection interval.

Read I/O Percent: The percentage of I/O requests that were read requests.

Reserved Percent: The percentage of time during the interval when a shared device was reserved by this system.

Response Time: The total hardware service time and the front end software queuing time in milliseconds involved for the average I/O request to the device.

Solid State Device: Indicates whether the logical volume is backed by a solid state physical device.

Storage Group Name: Indicates whether the number of PAV exposures for this device has changed during this interval.

Subsystem ID: The subsystem ID.

Time Stamp: The date and time this data was reported.

Volume: Represents the six-character identifier associated with this logical volume.

Write Hit Percent: The ratio of write hit requests that were satisfied by the storage control unit cache without requiring access to DASD.

Write I/O Count: The number of write requests in the current collection interval.

Application Volume Space

The Application Volume Space attributes provide space information for the logical volumes allocated to an application.

Device Address: The MVS device number, assigned by HCD, associated with this device.

Device Type: The MVS device type, for example 3380, 3390, or 2105.

Fragmentation Index: The qualitative measure of the scattered free space on a volume. The higher the number, the more severe the fragmentation.

Free Space Megabytes: Represents the total amount of free space currently available on this device, in megabytes.

Job Name: The JOBNAME, TSO ID, or started task name associated with this application.

Largest Free Extent MB: The largest contiguous free area available, in megabytes.

Managed System: The SMS name for this system.

Percent Free Space: The ratio of free space divided by capacity, expressed as a percentage.

SMS Conversion Status: The SMS conversion status of this volume. Values are: non-SMS, initial, converted, or excluded.

SMS Status: The SMS status of this volume. Values are enabled, quiesced_all, quiesced_new, disabled_all, or disabled_new.

Solid State Device: Indicates whether the logical volume is backed by a solid state physical device.

Storage Group Name: The 32-character SMS storage group name to which this volume belongs.

Total Capacity Megabytes: Represents the total data capacity of this device, in megabytes.

Volume: The six-character identifier associated with this logical device.

VTOC Index Status: The status of the VTOC index on this volume. Values are: Enabled, Disabled, Error, OS-VTOC, or Exclude List.

Cache Control Unit

The Cache Control Unit attributes provide performance and status information for cache control units.

Active Volumes: The number of volumes with ACTIVE caching status on the cache control unit.

Bypass Cache Percent: The percentage of I/Os that bypassed cache during this interval. The I/O request is sent directly to DASD. Tracks are not read from cache or promoted to cache. Also, the entries of tracks in the LRU list are not updated.

Cache MB Available: The total, in megabytes, of cache storage available on this Cache Control Unit.

Cache MB Configured: The total storage, in megabytes, of cache storage configured on this Cache Control Unit.

Cache Status: The caching status for this subsystem. Values are as follows: Active (in other words, caching for this subsystem is active), Inactive (in other words, caching for this subsystem is inactive), Pending Active, Pending Inactive, and Timeout.

CFW Read Percent: The percentage of all I/Os that were Read operations for data in cache that originated from a cache fast write during this interval. Cache fast write (CFW) is a form of fast write where the data is written directly to cache without using nonvolatile storage and is available for later destaging.

CFW Write Percent: The percentage or all I/Os that were Cache Fast Writes during this interval. Cache fast write (CFW) is a form of fast write where the data is written directly to cache without using nonvolatile storage and is available for later destaging.

Control Unit Type: The type of this cache control unit. This field is only populated for the physical devices that are supported with extended information. Examples of these devices include EMC Symmetrix, 2105, and 2107.

Deactivated Volumes: The number of volumes on this Cache Control for which the cache has been deactivated.

DFW Hit Percent: The percentage of all DASD I/Os to this control unit that were DASD Fast Writes during this interval. DASD Fast Write (DFW) is form of fast write to cache where the data is written concurrently to cache and nonvolatile storage and automatically scheduled for destaging to DASD.

DFW Retry Percent: The percentage of DASD I/O operations which required retry during this interval.

Inhibit Cache Percent: The percentage of I/O operations which inhibited the loading of data into cache.

Managed System: The SMS name of this system.

NVS KB Configured: The kilobytes of cache storage configured on this subsystem. NVS (nonvolatile storage) facilitates the DASD Fast Write feature which allows data to be written at electronic speeds directly into this nonvolatile memory in the controller, thus avoiding the many delays associated with accessing the actual disk.

NVS KB Pinned: The amount of storage that is unavailable because DASD failure is preventing the subsystem from destaging data. The data is pinned in cache.

NVS Status: The overall status of the nonvolatile storage (NVS). Values are:

Active=1 NVS is online and usable.

Inactive=2 NVS is not usable.

Pending_Active=3 Status is switching from inactive to active.

Pending_Inactive=4 Status is switching from active to inactive.

Timeout=0 Status is timed out.

Read Hit Percent: The percentage of read hits for each cache control unit.

Read I/O Percent: The percentage of I/O requests to the cache control unit that were read requests.

Serial Number Index: The serial number of the control unit. Used as an index into hardware specific displays.

Storage Facility ID: The Storage Facility Identifier for the device (Type-Model-Serial Number).

Storage Facility Series: The series of control unit, which defines the physical geometry of the subsystem. Can be one of: Symmetrix, Iceberg, RAMAC®, shark, DS6000™, or DS8000.

Subsystem ID Fullword: The subsystem ID (for Model 3990, 2105, and 2107 control units).

Subsystem ID Hex: The subsystem ID (for Model 3990, 2105, and 2107 control units).

Subsystem ID: The subsystem ID (for Model 3990, 2105, and 2107 control units).

Time Stamp: The date and time this data was reported. Primarily for use in historical reporting.

Track Destaging Rate: The number of tracks per second being removed from cache to physical disks over the cache interval.

Track Staging Rate: The number of tracks per second being brought into cache from the physical disks over the cache interval.

Write Hit Percent: The percentage of I/O operations that completed without accessing the hard disk drive.

Cache Devices

The Cache Devices attributes provide both cache status and cache performance information for a logical volume.

Avg Logical KB Read: The average number of kilobytes read from the logical volume per second over the cache interval.

Avg Logical KB Written: The average number of kilobytes written to the logical volume per second over the cache interval.

Avg Logical Read Time: The average time in milliseconds spent reading from the logical volume per I/O over the cache interval.

Avg Logical Write Time: The average time in milliseconds spent writing to the logical volume per I/O over the cache interval.

Avg Phys Read Response Time: The average time to read from the physical disk per operation in milliseconds over the cache interval.

Avg Phys Write Response Time: The average time in milliseconds to write to the physical disk per operation over the cache interval.

Bypass Cache Percent: The percentage of I/Os that bypassed cache during this interval. The I/O request is sent directly to DASD. Tracks are not read from cache and are not promoted to cache, and their entries in the LRU list are not updated.

Cache Status: The caching status for the subsystem to which this device is attached. Active means caching for this subsystem is active. Inactive means caching for this subsystem is disabled.

Cache to DASD Total Tracks: The number of tracks written from cache to physical disk during the cache interval. This value is captured asynchronously with respect to the transfer of data from the channel to the cache.

Cache to DASD Tracks per Second: The rate per second of tracks written from cache to physical disk during the cache interval. This value is captured asynchronously with respect to the transfer of data from the channel to the cache.

CFW Read Percent: The percentage of READ requests which were satisfied from cache due to a previous cache fast write.

CFW Write Percent: The percentage of all I/Os that were CACHE FAST WRITES during this interval. CACHE FAST WRITE allows data to be written to cache only. The data is later written to DASD. CFW differs from DASD Fast Write in that the data is not written to nonvolatile storage (NVS).

Channel Index: The identifier for channel path that identifies which channel path this device is attached to.

Control Unit Model: The model number of the control unit to which this device is attached.

DASD to Cache Non-promotes per Second: The rate per second of operations that did not promote (copy) a track to cache because free cache segments were not available.

DASD to Cache Non-promotes: The number of operations that did not promote (copy) a track to cache because free cache segments were not available.

DASD to Cache Total Tracks Non-sequential: The number of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Total Tracks Sequential: The number of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

DASD to Cache Total Tracks: The number of tracks read from physical disk to cache over the cache interval for I/O operations.

DASD to Cache Tracks per Sec: The rate per second of tracks read from physical disk to cache over the cache interval for I/O operations.

DASD to Cache Tracks per Second Non-sequential: The rate per second of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Tracks per Second Sequential: The rate per second of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

Device Address: The four hexadecimal digits that define the physical address of this device. Also known as unit or unit address.

Device Hex Address: The four-digit hexadecimal address of this device.

DFW Hit Percentage: The percentage of DASD Fast Writes that were satisfied from cache without requiring access to DASD. DASD Fast Writes are written to both cache and nonvolatile storage (NVS).

DFW Retry Percentage: The percentage of DASD Fast Write operations that did not complete due to nonvolatile Storage (NVS) constraints.

DFW Status: The status of DASD FAST WRITE function for this volume. Values are Active, Inactive, Unknown, or Timeout.

Extent Pool ID: The Extent Pool Identifier.

I/O Count: The number of cache eligible I/Os to the device over the cache collection interval.

Inhibit Cache Percentage: The percentage of I/O requests that set the Inhibit Cache Load mode in the channel program to prevent the data being cached. If the data is already in cache, the request is satisfied from cache.

Logical Control Unit Index: The number of LCU to which this device is attached. An LCU is the logical representation of either a single control unit, with or without attached devices, or a group of control units that share devices.

Logical KB Read: The total number of kilobytes read from the logical volume over the cache interval.

Logical KB Written Total: The total number of kilobytes written to the logical volume over the cache interval.

Logical Read Time: The amount of time in milliseconds spent reading from the logical volume over the cache interval.

Logical Write Time: The amount of time in milliseconds spent writing to the logical volume over the cache interval.

Managed System: The SMS name of this system.

Physical Device: The identifier of the physical device type of this device.

Physical KB Read Rate: The kilobytes read from the physical disk per second over the cache interval.

Physical KB Read Total: The total kilobytes read from the physical disk over the cache interval.

Physical KB Write Rate: The kilobytes written to the physical disk per second over the cache interval.

Physical KB Write Total: The total kilobytes written to the physical disk over the cache interval.

Physical Read Rate: The rate of I/O read operations from the physical disk per second over the cache interval.

Physical Read Total: The total number of I/O read operations from the physical disk over the cache interval.

Physical Write Rate: The rate of I/O write operations to the physical disk per second over the cache interval.

Physical Write Total: The total number of I/O write operations to the physical disk over the cache interval.

RAID Flag: Indicates the type of RAID device.

Read Hit Percent: The percentage of READ requests that were satisfied from cache.

Read I/O Count: The number of read requests for this device.

Read I/O Percent: The percentage of I/O requests that are read requests.

Solid State Device: Indicates whether the logical volume is backed by a solid state physical device.

SSID Index: The fullword representation of the SSID used for indexing the product.

Storage Facility ID: The storage facility identifier for the device (Type-Model-Serial Number).

Storage Group Index: The name of the storage group to which this volume belongs. The storage group is an SMS construct used to group DASD devices in order to separate the physical requirements of storing data from the logical requirements.

Subsystem ID Hex: The subsystem ID associated with this device.

Subsystem ID Selector: An index to limit the range of physical volumes for which data is requested.

Subsystem ID: The two-byte unique value used to identify a 3990 or equivalent storage control unit.

Time Stamp: Represents the date and time this data was collected.

Total Phys Read Response Time: The total time to read from the physical disk in milliseconds over the cache interval.

Total Phys Write Response Time: The total time to write to the physical disk in milliseconds over the cache interval.

User Group Index: A user DASD group is an OMEGAMON II® facility that enables you to define groups of DASD volumes by VOLSER and device address. You can also specify generic VOLSERS and devices (patterns) by using the wildcard character.

Volume: The volume serial number of the disk device.

Write Hit Percent: The percentage of I/O WRITE operations that were satisfied without having to access the physical drive. This occurs as the result of DASD Fast Write or Cache Fast Write features. DASD Fast Write is a default in controllers that support it. Cache Fast Write can be requested in the Define Extent command of the channel program.

Write I/O Count: The number of write requests for this device.

Cache Raid Rank

The Cache Raid Rank attributes provide performance and usage information for the ranks in an IBM 2105 storage facility.

Device Adapter ID: The two hexadecimal digit of the device adapter associated with this rank. The device adapter is the hardware component that provides the path control and data transfer functions.

Hard Disk Drives in Rank: The number of disk drives contained in this rank.

Managed System: The SMS name of this system.

Raid Rank ID: The four hexadecimal digit ID associated with this raid rank.

Read MB per Request: The average number of megabytes of data transferred per read request.

Read MB per Second: The average number of megabytes of data transferred per second for read requests.

Read Requests per Second: The average number of read operations per second.

Read Response Time MS: The average response time in milliseconds for read requests.

Read Sectors per Request: The average number of hard disk sections read per read request.

Sector Size: The size in bytes of the sectors on the hard disks associated with this rank.

Sectors Read per Second: The average number of hard disk sectors read per second for read requests.

Subsystem ID Fullword: The subsystem ID of the logical control unit to which this rank is attached.

Subsystem ID: The subsystem ID of the logical control unit to which this rank is attached.

Time Stamp: The date and time associated with this row of data.

Volume: The volume serial number of the first logical volume on this rank.

Write MB Per Request: The average number of megabytes written per request for this rank.

Write MB Per Second: The average number of megabytes of data written per second for this rank.

Write Requests Per Second: The average number of write operations per second for this rank.

Write Response Time MS: The average time in milliseconds required to complete a write operation.

Write Sectors per Request: The average number of physical sectors written per write request.

Channel Path

The Channel Path attributes contain channel configuration and performance information. Extended information is provided for FICON[®] channels.

Bus Percent Utilization: For FICON channels only, represents the percentage of the maximum available bus cycles being used in this collection interval.

Complex MB Second Read: For FICON channels only, represents the number of megabytes of data transferred for read requests by this channel path for all LPARs in this complex.

Complex MB Second Written: For FICON channels only, represents the average number of megabytes written by this channel path per second during this interval for the complex.

Complex Percent Read Units: For FICON channels only, represents the ratio of total read units for this channel, for the complex, divided by the maximum read units for this channel.

Complex Percent Utilized: The ratio of total channel busy time divided by elapsed time for this interval.

Complex Percent Write Units: For FICON channels only, represents the ratio of the number of write units for the complex divided by the maximum write units for this channel.

CPMF Group Number: The number of the Dynamic Channel Path Management Group to which this channel path is assigned.

DCM Status: The Dynamic Channel Path Management Status. DCM allows you to dynamically assign channels for controlling units to respond to peaks in demand.

Description: The type of channel connection. The values are as follows:

- Cluster_Bus_Receiver
- Cluster_Bus_Sender
- Coupling_Facility_Peer
- Coupling_Facility_Receiver
- Coupling_Facility_Sender
- CTC_Point_to_Point
- CTC_Switched_or_Point_to_Point
- CTC_Switched_Point_to_Point
- Direct_System_Device
- Emulated_I/O
- ESCON_Path_to_a_Block_Converter
- ESCON_Path_to_a_Byte_Converter
- ESCON_Point_to_Point
- ESCON_Switched_Point_to_Point
- FICON_Incomplete
- FICON_Point_to_Point
- FICON_Switched
- FICON_to_ESCON_Bridge
- Integrated_Cluster_Bus_Peer
- Internal_Coupling_Peer
- Internal_Coupling_Receiver
- Internal_Coupling_Sender
- Internal_System_Device
- Native_Interface
- Open_Systems_Adapter
- OSA_DIRECT_EXPRESS
- OSA_EXPRESS
- Parallel_Block_Multiplex
- Parallel_Byte_Multiplex
- Unknown

LPAR MB Second Read: For FICON channels only, represents the number of megabytes of data transferred for read requests by this channel path for this LPAR.

LPAR MB Second Written: For FICON channels only, is the average number of megabytes written by this channel path per second during this interval for this LPAR.

LPAR Percent Read Units: For FICON channels only, is the ratio of total read units for this LPAR, divided by the maximum read units for this channel.

LPAR Percent Utilized: The ratio of channel path busy time for this LPAR divided by elapsed time for this interval.

LPAR Percent Write Units: For FICON channels only, is the ratio of total write units for this channel, for this LPAR, divided by the maximum write units for this channel.

Managed System: The SMS name of this system.

NonSMS Volumes: Represents the number of DASD volumes accessible through this channel path that are not managed by SMS.

Path ID Hex: The two-digit hexadecimal number assigned by HCD identifying this channel path to the system.

Path ID: The two-digit hexadecimal identifier for this channel path assigned by HCD.

Shared Path: Denotes whether this path is shared with other LPARs (yes), or is dedicated to this LPAR (no).

SMS Volumes: The number of SMS-controlled DASD volumes accessible through this channel path.

Status: Denotes the availability state of this channel path.

Time Stamp: The date and time this data was reported. Primarily for use in historical reporting.

Total Volumes: Represents the total number of DASD volumes accessible by this channel path.

Cross System Volume Summary

The Cross System Volume attributes provide an overview of the performance of a logical device across multiple z/OS images, that is, it helps to identify problems with shared DASD.

Control Unit Model: The model number associated with this control unit.

Cumulative I/O Rate: Represents the sum of I/O rates for all sharing systems.

Device Contention Index: The percentage of time that I/Os to the device are delayed (pending), averaged for all sharing systems. This is a measure of the time that each system spends waiting for the device.

Device Model: The model of DASD. Device geometry for the most common models or their equivalents are as follows:

| Model | Tracks | Cylinders | Paths | Density |
|--------|--------|-----------|-------|---------|
| 3390-1 | 16685 | 1113 | 4 | TRIPLE |
| 3390-2 | 33390 | 2226 | 4 | TRIPLE |
| 3390-3 | 50055 | 3339 | 4 | TRIPLE |
| 3380-K | 39825 | 2655 | 4 | TRIPLE |

| Model | Tracks | Cylinders | Paths | Density |
|--------|--------|-----------|-------|---------|
| 3380-J | 13275 | 885 | 4 | SINGLE |
| 3380-E | 26550 | 1770 | 2 | DUAL |
| 3380-A | 13275 | 885 | 2 | SINGLE |
| 3375 | 11508 | 959 | 2 | SINGLE |
| 3350 | 16650 | 555 | 2 | SINGLE |

High Disconnect Time System: The SMFID of the sharing system reporting the highest disconnect time for this volume.

High Disconnect Time Value: The highest disconnect time reported by any sharing system.

High MSR System: The SMFID of the sharing system reporting the highest MSR (millisecond response time) value for this volume.

High MSR Value: Represents the highest MSR (millisecond response time) time reported by any sharing system.

Managed System: The SMS name of this system.

Storage Group Name: The 32-character name of the SMS storage group to which this volume belongs.

True % Busy: The percentage of time a device is busy as a result of activity from all sharing systems. This is a measure of the time that the device spends actually doing work.

User DASD Grp Name: The name of the user DASD group to which this volume belongs. User DASD group is an OMEGAMON II facility that enables you to define groups of DASD volumes by VOLSER and device address. You can also specify generic VOLSERS and devices (patterns) by using the wildcard character.

This feature gives you the flexibility of having OMEGAMON II report the performance and status of any combination of DASD volumes as a single entity. User DASD groups can include both SMS managed volumes and non-SMS managed volumes.

Volume: The six-character identifier assigned to a DASD or tape volume. Also known as the volume serial number or volume ID.

Cross System Volumes

The Cross System Volumes attributes provide performance information for a logical volume on each z/OS image on which it is shared.

Busy Percent: The percentage of channel path usage. For FICON channels, this is the percentage of the maximum work units being used.

Connect Time: A part of I/O during which data is located on the track and transferred between the DASD device and central storage.

Control Unit Model: The model number of the control unit for this hard disk drive.

Device Address: The MVS device number.

Disconnect Time: The time that I/O spends, free from the channel, searching for or requesting data. This time includes moving the device to the requested cylinder, waiting for the record to rotate under the head (rotational delay), and waiting for the channel to become available.

I/O Per Second: The average number of I/O operations per second.

IOSQ Delay: The average time that an I/O waits because the device is already busy (signified by the UCBBUSY BIT of the device being on).

Managed System: The SMS name of this system.

Pend Time: The time that I/O is delayed in the path to the device. Pending time can be attributable to the channel or control unit path being busy.

Response Time: The amount of time in milliseconds required for an I/O request to complete to a given device. MSR (millisecond response time) is divided into four components: IOS Queue, Pending, Disconnect, and Connect.

Storage Group Name: The 32-character name of the SMS storage group to which this volume belongs.

System: The systems management facility ID of the system that reported this row of data.

Volume: The six-character name assigned to DASD or tape volume. Also known as volume serial number of volume ID.

DASD Physical Group

The DASD Physical Group attributes provide performance information for a physical volume group. A physical volume group is the set of logical volumes that are related by a hardware specific construct.

Average Command Response Delay: The average number of milliseconds that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device over the RMF interval.

Average HyperPAV Alias Count: The average number of aliased devices assigned to a base volume in a HyperPAV environment over the RMF interval.

Busy Percent: The percentage of time a resource is busy processing I/O requests. For PAV volumes, this value is normalized by dividing the number of PAV exposures.

Bypass Cache Percentage: The percentage of I/Os that bypassed cache this interval. The I/O request is sent directly to DASD. Tracks are neither read from cache nor promoted to cache. Also, their entries in the LRU list are not updated.

Cache FW % Read: The percentage of all I/Os that were read operations for data in cache that originated from a cache fast write during this interval. Cache fast write is a form of fast write where the data is written directly to cache without using nonvolatile storage and is available for later destaging.

Cache FW % Write: The percentage of all I/Os that were cache fast writes during this interval. Cache fast write is a form of fast write where the data is written directly to cache without using nonvolatile storage and is available for later destaging.

Cache Status: The caching status for this volume. Active means caching for this volume is active and inactive means caching for this volume is inactive.

Cache to DASD Total Tracks: The number of tracks written from cache to physical disk during the cache interval. This value is captured asynchronously with respect to the transfer of data from the channel to the cache.

Cache to DASD Tracks per Sec: The rate per second of tracks written from cache to physical disk during the cache interval. This value is captured asynchronously with respect to the transfer of data from the channel to the cache.

Connect Time: Part of I/O during which data is located on the track and transferred between the DASD device and central storage.

Control Unit Model: The model number associated with the control unit to which this device is attached.

Current PAV Exposures: The current number of UCBs through which the device can be accessed.

DASD to Cache Non Promotes Second: The rate per second of operations that did not promote (copy) a track to cache because free cache segments were not available.

DASD to Cache Non-promotes: The number of operations that did not promote (copy) a track to cache because free cache segments were not available.

DASD to Cache Total Tracks Non-sequential: The number of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Total Tracks Sequential: The number of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

DASD to Cache Total Tracks: The number of tracks read from physical disk to cache over the cache interval for I/O operations.

DASD to Cache Tracks per Sec Non-sequential: The rate per second of tracks read from physical disk to cache over the cache interval for non-sequential I/O operations.

DASD to Cache Tracks per Sec Sequential: The rate per second of tracks read from physical disk to cache over the cache interval for sequential I/O operations.

DASD to Cache Tracks per Sec: The rate per second of tracks read from physical disk to cache over the cache interval for I/O operations.

DCBs Open: The count of open data sets.

Device Address: The MVS device number assigned by HCD and associated with this device.

Device MPL: The device multi-processing level, which is the I/O density for a device. This metric is the product of the MSR (millisecond response time) and I/O per second for a device.

DFW Hit Percent: The percentage of all DASD I/Os to this volume that were DASD fast writes during this interval. DASD fast write is a form of fast write to cache where the data is written concurrently to cache and nonvolatile storage and automatically scheduled for destaging to DASD.

DFW Retry Percent: The percentage of DASD I/O operations that required retry during this interval.

DFW Status: The status of the DASD fast write option for this volume. Value can be Active, Inactive, Unknown or Timeout.

Disconnect Time: The time that I/O spends, free from the channel searching for or requesting data.

I/O Count: The actual number of I/O requests issued for this device during this interval.

I/O per Second: The average number of I/O operations per second.

Inhibit Cache Percent: The percentage of I/O operations that inhibited the loading of data into cache. If the data is already in cache, the request is satisfied from cache.

IOSQ Delay: The average time that an I/O waits because the device is already busy. A device is determined to be busy if its UCBBUSY bit is on.

LCU Contention Rate: The average number of requests delayed due to LCU contention. LCU contention is the rate per second at which the SAP places delayed I/O requests on the CU-HDR for this LCU. This calculation is performed when all paths to the subchannel are busy and at least one path to the control unit is busy. For devices with only one path, or devices where multiple paths exist and the busy condition is immediately resolved, the IOP does not count the condition.

LCU Number: The three hexadecimal digit logical control unit number.

Managed System: The SMS name of this system.

Maximum PAV Exposures: The maximum number of UCBs through which the device can be accessed in this interval.

MSR Connect Time Percent: The percentage of MSR (millisecond response time) during which data is transferred (SEARCH + TRANSFER).

PAV Exposure Changed: Indicates that the number of PAV exposures to this device changed during this interval.

Pend Time: The time that I/O is delayed in the path to the device. Pending time can be attributable to the channel or control unit path being busy.

Read Hit Percent: The percentage of read hits for this volume.

Read I/O Count: The number of read operations directed to this device that completed during this interval.

Read I/O Percent: The percentage of I/O requests to this volume that were read requests.

Reserved Percent: The percentage of time during the interval when a shared device was reserved by this system.

Response Time: The amount of time in milliseconds it takes for an I/O request to complete to a given device. MSR (millisecond response time) is divided into four components: IOS Queue, Pending, Disconnect and Connect.

Selected Volser: The volume serial number associated with the physical volume selected for display.

SMS Status: The SMS volume status for the local system. Values are: Enabled, Quiesced_All, Quiesced_New, Disabled_All or Disabled_New.

Storage Group Name: The 32-character name of the SMS storage group to which this volume belongs.

Subsystem ID: The MVS identifier of the logical control unit to which this device is attached.

Volume: The VOLSER associated with this physical volume.

VTOC Index Status: The status of the VTOC index on this volume. Values are Enabled, Disabled, Error, OS_VTOC or Excluded.

Write Hit Percent: The percentage of I/O operations that completed without accessing the hard disk drive.

Write I/O Count: The number of write operations directed to this device that completed during this interval.

DASD Volume Performance

The DASD Volume Performance attributes provide status and performance information for logical volumes.

Average Command Response Delay: The average number of milliseconds that a successfully initiated start or resume function needs until the first command is indicated as accepted by the device over the RMF interval.

Average HyperPAV Alias Count: The average number of aliased devices assigned to a base volume in a HyperPAV environment over the RMF interval.

Busy Percent: The percentage of time a resource is busy processing I/O requests. For Parallel Access Volumes, this value is normalized by dividing by the number of PAV exposures.

Channels: The channel Path IDs (in hexadecimal) through which this device can be accessed.

CHPID SEL The CHPID selector is used as an index to limit the range of physical volumes for which data is requested.

Connect Time: The average time (in milliseconds) required for seek and transfer. This is the time it takes for data to be located on the track and transferred between the DASD device and central storage.

Control Unit Busy Delay Time: The average amount of I/O delay (in milliseconds) to this device caused by a control unit busy condition.

CU Model: The model indicator for this cache control unit.

Current PAV Exposures: The current number of UCBs through which the device can be accessed.

DCBs Open: The number of open data sets on this device.

Device Address: The MVS device number associated with this device.

Device Busy Delay Time: The average amount of I/O delay (in milliseconds) to this device caused by a device busy condition.

Device Hex Address: The device address in hexadecimal numbers.

Device Model: The MVS device model.

Device MPL: The device Multiprocessing Level, which is the I/O density for a device. This metric is the product of the MSR (millisecond response time) and I/O rate per second for the device. For a PAV volume, this value is normalized by dividing by the number of exposures.

Device Type: The MVS device type.

Director Busy Delay Time: The average amount of I/O delay (in milliseconds) to this device caused by a director port busy condition.

Disconnect Time: The average time in milliseconds that I/O waited because the device was not connected to the channel. This time includes moving the device to the requested cylinder, waiting for the record to rotate under the head (rotational delay), and waiting for the channel to become available.

I/O Count: The actual number of I/O requests issued to the device over the Resource Monitoring Facility (RMF) interval.

I/O Per Second: The average number of I/O operations per second.

IOSQ Delay: The average time (in milliseconds) that an I/O waits because the device is already busy. A device is determined to be busy if its UCBBUSY bit is on.

LCU Number Hex: The hexadecimal LCU number for this device.

LCU Number: The hexadecimal logical control unit number.

Managed System: The SMS name of this system.

Maximum PAV Exposures: The maximum number of UCBs through which the device can be accessed in this interval.

MSR Connect Time Percent: The percentage of the device response time during which the device was connected. This value is obtained by dividing the device pending time (in milliseconds) by the total millisecond response time.

MVS Status: The status of this device in the local MVS system. Values are Online, Offline, Pending Offline, Boxed, Not Ready, and Long Busy.

PAV Exposure Changed: Indicates whether the number of PAV exposures for the volume has changed over the interval.

Pend Time: The average time (in milliseconds) that I/O is delayed in the path to the device. Pending time can be attributable to the channel or control unit path being busy.

Physical Device: The type of physical device containing this logical device (supported for the following devices only: Shark 2105, Symmetrix, TDS). Values are 2105, Symmetrix, and TDS.

Raid Model: The raid type indicator. Values are Symmetrix, Iceberg, Ramac, or 2105.

Reserved Percent: The percentage of time during the interval when a shared device is reserved by this system.

Response Time: The total hardware service time and the front end software queuing time (in milliseconds) involved for the average I/O request to the device.

Solid State Device: Indicates whether the logical volume is backed by a solid state physical device.

SMFID: The SMF ID of the system from which this information was obtained.

SMS Conversion Status: The SMS conversion status of this volume. Values are Non-SMS, Initial, Converted, or Excluded.

SMS Status: The SMS volume status for a local system. Values are Not Connected, Enabled, Quiesced_All, Quiesced_New, Disabled_All, or Disabled_New.

Storage Group Name: The SMS storage group name to which this volume belongs.

Subsystem ID Index: The two-byte unique value used to identify a 3990 or equivalent storage control unit.

Time Stamp: The date and time associated with this row. Primarily for use in historical reporting.

User Group Index: The user-assigned name for a group of DASD volumes.

Volume: The volume serial number of a disk device currently in use by this application.

VTOC Index Status: The type of VTOC present on this device. Values are Enabled, Disabled, Error, OS VTOC, and Excluded.

DASD Volume Space

The DASD Volume Space attributes provide status and space information for logical volumes.

Allocated Space MegaBytes: The amount of space in megabytes on this volume which is allocated to data sets.

Cylinder Managed Space: Space on the volume, beginning at cylinder address 65520, is managed only in multicylinder units.

Device Address: The MVS device number assigned by HCD, associated with this device.

Device Model: The actual device model. Each device model has unique space and geometry attributes.

Device Type: The MVS device type. For example: 3380, 3390, or 2105.

Extended Address Volume: Volume has more than 65520 cylinders and is an Extended Address Volume.

Fragmentation Index: A number between 0 and 1000 that indicates the degree of scattered free space on the volume. The higher the number, the higher the amount of fragmentation.

Free DSCBs: The number of DSCBs in the VTOC of this volume that are unused.

Free Extents: The amount of free areas on the volume.

Free Space Megabytes: Represents the amount of free space on this volume in megabytes.

Free VTOC Index Records: The number of index records in the VTOC index that are available for use.

Largest Free Extent MB: The size in megabytes of the largest contiguous free area on the volume.

Managed System: The SMS name of this system.

Maximum Free Cylinders: The largest contiguous area of free cylinders on this device.

Maximum Free Tracks: The number of tracks in the largest free extent.

Percent Free Space: The ratio of free space divided by the total space for this volume expressed as a percentage.

Solid State Device: Indicates whether the logical volume is backed by a solid state physical device.

SMS Conversion Status: The SMS conversion status of this volume. Values are: Non-SMS, Initial, Converted, or Excluded.

SMS Status: The SMS volume status for the local system. Values are: Enabled, Quiesced_All, Quiesced_New, Disabled_All, or Disabled_New.

Storage Group Name: The 32-character name of the SMS storage group to which this volume belongs.

Time Stamp: Represents the date and time associated with this row.

Total Capacity Megabytes: The data capacity of this volume in megabytes.

Total Free Cylinders: Represents the total number of unallocated cylinders on this volume.

Total Free Tracks: The total number of separated areas of one track or larger on this volume that are available to allocate.

Track Managed Fragmentation Index: A measure of the degree to which free space in the track managed area on the volume is fragmented. The higher the number, the more severe the track managed area fragmentation.

Track Managed Free Cylinders: Total number of unallocated whole cylinders in the track managed area on this volume.

Track Managed Free Extents: Total number of free extents in the track managed area on this volume.

Track Managed Free_Space: Amount of free space in the track managed area on this volume in megabytes.

Track Managed Free Tracks: The total number of free tracks in the track managed area which are not contained in the total number of track managed area free whole cylinders.

Track Managed Largest Free Extent: The size, in megabytes, of the largest contiguous free area in the track managed area on the volume.

Track Managed Maximum Free Cylinders: Largest contiguous area of free cylinders in the track managed area on this device.

Track Managed Maximum Free Tracks: Number of additional free tracks in the extent which contains the Maximum Free Cylinders.

Track Managed Percent Free: Portion of the track managed area on this volume which is currently unused, expressed as a percentage.

Track Managed Total Capacity: Data capacity of the track managed area on this volume in megabytes.

User Group Index: The name of the user-defined group of the DASD volume to which this volume belongs.

Volume: The six-character identifier associated with this logical device.

VTOC Index Status: The status of the VTOC index on this volume. Values are: Enabled, Disabled, Error, OS-VTOC, or Excluded.

Dataset Attributes Blocksize Summary

This group of attributes monitors the count of the number of data sets that exist in each product-defined unique DCB BLKSIZE range. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Allocated GB: Total allocated GB for all data sets in the range.

Allocated MB: Total allocated MB for all data sets in the range.

Allocated Tracks: Total allocated tracks for all data sets in the range.

High Blocksize: The highest value in the blocksize range.

Low Blocksize: The lowest value in the blocksize range.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Datasets: The number of data sets in the range.

Percent of Allocated Space: The percentage of total allocated space that data sets in this blocksize range represent.

Percent of Total Datasets: The percentage of total processed data sets that data sets in this blocksize range represent.

Timestamp: The time that the row was created.

Unused GB: Total unused GB for all data sets in the range.

Unused MB: Total unused MB for all data sets in the range.

Unused Tracks: Total unused tracks for all data sets in the range.

Used GB: Total used GB for all data sets in the range.

Used MB: Total used MB for all data sets in the range.

Used Tracks: Total used tracks for all data sets in the range.

Dataset Attributes CA Split Summary

This group of attributes provides data on the Top N VSAM data sets in terms of total number of CA Splits. You define the value of n in the Dataset Attributes Data Collection dialog box. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Address: Address of detail information in the data space of the data set.

ALET: ALET (Address List Entry Token) of information in the data space of the data set.

Allocated GB: Total allocated GB for the data set.

Allocated MB: Total allocated MB for the data set.

Allocated Tracks: Total allocated tracks for the data set.

Associated DSN: The name of this component's associated data set.

CA Splits: Total number of Control Area splits.

CI Splits: Total number of Control Interval splits.

Dataset Name: The name of the data set.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Total unused GB for the data set.

Unused MB: Total unused MB for the data set.

Unused Tracks: Total unused tracks for the data set.

Used GB: Total used GB for the data set.

Used MB: Total used MB for the data set.

Used Tracks: Total used tracks for the data set.

Volser: The VOLSER on which the data set is located.

Dataset Attributes Catalog Summary

This group of attributes tracks the number of data sets cataloged in each system catalog. This data shows the distribution of data sets across catalogs and can allow rebalancing to improve performance.

When the Collect Catalog Data option is enabled, the catalog data is retrieved for each data set encountered in the VTOC of the volume. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Allocated GB: Total allocated GB represented by all of the data sets in the catalog.

Allocated MB: Total allocated MB represented by all of the data sets in the catalog.

Allocated Tracks: Total allocated tracks represented by all of the data sets in the catalog.

Catalog Name: Name of the catalog.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Total Datasets: The total number of processed data sets contained in the catalog.

Note: Data sets that are cataloged, but do not exist on any processed volume are not counted. The data sets are not counted because the data set name from the VTOC is the starting point for the catalog lookup.

Unused GB: Total unused GB represented by all of the data sets in the catalog.

Unused MB: Total unused MB represented by all of the data sets in the catalog.

Unused Tracks: Total unused tracks represented by all of the data sets in the catalog.

Used GB: Total used GB represented by all of the data sets in the catalog.

Used MB: Total used MB represented by all of the data sets in the catalog.

Used Tracks: Total used tracks represented by all of the data sets in the catalog.

Dataset Attributes CI Split Summary

This group of attributes provides data on the Top N VSAM data sets in terms of total number of CI Splits. You define the value of n in the Dataset Attributes Data Collection dialog box. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Address: Address of detail information in the data space of the data set.

ALET: ALET (Address List Entry Token) of information in the data space of the data set.

Allocated GB: Total allocated GB for the data set.

Allocated MB: Total allocated MB for the data set.

Allocated Tracks: Total allocated tracks for the data set.

Associated DSN: The name of this component's associated data set.

CA Splits: Total number of Control Area splits.

CI Splits: Total number of Control Interval splits.

Dataset Name: The name of the data set.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Total unused GB for the data set.

Unused MB: Total unused MB for the data set.

Unused Tracks: Total unused tracks for the data set.

Used GB: Total used GB for the data set.

Used MB: Total used MB for the data set.

Used Tracks: Total used tracks for the data set.

Volser: The VOLSER on which the data set is located.

Dataset Attributes Creation Date Summary

This group of attributes monitors the count of the number of data sets that exist in each product-defined unique Creation Date (in other words, number of days since creation) range. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Allocated GB: Total allocated GB for all data sets in the range.

Allocated MB: Total allocated MB for all data sets in the range.

Allocated Tracks: Total allocated tracks for all data sets in the range.

High Days Since Creation: The highest value in the creation date range.

Low Days Since Creation: The lowest value in the creation date range.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Datasets: The number of data sets in the range.

Percent of Total Allocated Space: The percentage of total allocated space that data sets in this creation date represent.

Percent of Total Datasets: The percentage of total processed data sets that data sets in this creation date range represent.

Timestamp: The time that the row was created.

Unused GB: Total unused GB for all data sets in the range.

Unused MB: Total unused MB for all data sets in the range.

Unused Tracks: Total unused tracks for all data sets in the range.

Used GB: Total used GB for all data sets in the range.

Used MB: Total used MB for all data sets in the range.

Used Tracks: Total used tracks for all data sets in the range.

Dataset Attributes Defs

This group of attributes provides information to other attribute groups. Users do not use these attributes. The attributes are documented for the purpose of background reference only. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Message: Message.

Dataset Attribute Definitions: Dataset Attribute bulk data.

Dataset Attributes Dataset Detail

This group of attributes provides data set information from the VTOC and catalog. The attributes in this group include details regarding control intervals (CIs) and control areas (CAs). See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Note: You can create queries and situations that generate more specific data than that provided in the default workspaces. Be aware that in a large enterprise the data set database can be tracking millions of data sets. Estimate the potential size of the responses to your queries and situations to prevent them from requiring excessive processing, network, and disk space resources. Also be aware that a maximum of 16384 rows are returned on any single query against the Dataset Attribute Detail table.

% Free Bytes per CI: The VSAM percentage of free bytes in all CIs.

% Free CIs per CA: The VSAM percentage of free CIs in all CAs.

Allocated GB: The gigabytes allocated to the data set on the volume.

Allocated MB: The megabytes allocated to the data set on the volume.

Allocated RBA: The VSAM high allocated RBA.

Allocated Tracks: The tracks allocated to the data set on the volume.

Associated Entry Name: The VSAM associated data set name.

Average LRECL: The VSAM average LRECL.

BLKSIZE: The block size.

CA Splits: The VSAM number of CA splits.

Catalog Entry Type: The catalog entry type. The values are as follows:

- Alternate Index
- Cluster
- Data Component
- Generation Data Group
- Generation Data Set
- Index Component
- Non-VSAM Data Set
- Unknown
- USERCAT

Catalog Name: The name of the catalog in which data set is cataloged.

Cataloged: The data set is cataloged.

CI Size: The VSAM control interval size.

CI Splits: The VSAM number of CI splits.

CIs per CA: The VSAM number of CIs per CA.

Collection Timestamp: Date and time of data collection for this row.

Component Free Space: Bytes of free space in a VSAM data set component.

Compressed User Data Size: The VSAM compressed user data size.

Compressible: This is a compressible data set.

Creation Date: The creation date.

DAEAddress: Address of detail information in the data space of the data set.

DAEALET: ALET of detail information in the data space of the data set.

Data Component: The data set is a VSAM data component.

Dataset Name: The name of the data set.

Days Since Creation: The number of days since creation.

Days Since Last Backup: The days since last backup.

Days Since Reference: The number of days since last reference.

Days Since/Until Expiration: The numbers days until expiration. This value is negative if expiration has passed.

Device Type: The device type.

DSORG: The data set organization type. The values are as follows:

VS: VSAM

IS: ISAM

PS: Physical Sequential

DA: Direct Access

PO: Partitioned

GS: Graphics

CX: BTAM OR QTAM Line Group

CQ: QTAM DIRECT ACC MSG QUE

MQ: QTAM PROB PROG MSG QUE

TX: TCAM Line Group

TQ: TCAM Message Queue

AM: VSAM Data Set/Data Space

TR: TCAM 3705

?: Unknown

Extended Address Space Eligible: Data set is eligible to use Extended Address Space of an Extended Address Volume.

Erase: The VSAM erase attribute.

EXCPS: The VSAM number of EXCPs.

Expiration Date: The expiration date.

Extended Format: This is an extended format data set.

Extents: The number of extents.

Free Bytes per CI: The VSAM number of free bytes in all CIs.

Free CIs per CA: The VSAM number of free CIs per CA.

GDG Attributes: GDG attributes.

GDG Last Altered Date: GDG last altered date.

GDG Level: GDG level.

Generation Dataset: This is a generation data set.

HFS: This is an HFS data set.

Imbed: The VSAM imbed attribute.

Index Component: The data set is a VSAM index component.

Is a Catalog: The data set itself is a catalog.

Key Length: The DASD hardware key length or, for keyed VSAM data sets, the keylen value.

Last Backup Date: The last backup date.

LRECL: The logical record length.

Managed System: Internal node name used to identify the system from which this row of data was obtained.

Maximum Buffer Size: The VSAM maximum buffer size.

Maximum Datasets in GDG: For generation data sets, the maximum number of data sets in the GDG.

Maximum Record Size: The VSAM maximum record size.

Owner: The VSAM data set owner.

Page Space: The cluster is a page space.

PDSE: This is a persistent data store extended (PDSE) data set.

Percent Free: The percent of allocated space free. For VSAM data sets, this value is calculated using: $(\text{High Allocated RBA} - \text{High Used RBA}) \setminus \text{High Allocated RBA}$.

Percent Used: The percent of allocated space used. For VSAM data sets, this is calculated using $(\text{High Used RBA} \setminus \text{High Allocated RBA})$.

Primary GB: The primary space allocation (in other words, the first extent) in GB.

Primary MB: The primary space allocation (in other words, the first extent) in MB.

Primary Tracks: The primary space allocation (in other words, the first extent) in tracks.

RACF®: This data set is RACF protected. This attribute is based on the VTOC RACF indicator, not on whether the data set might be RACF-protected by default.

RECFM: The record format. The values are as follows:

F: Fixed Length

V: Variable Length
U: Undefined Length
T: Track Overflow
B: Blocked
S: Spanned/Standard
A: ANSI Control Character
M: Machine Control Character

Records Deleted: The VSAM number of deleted records.

Records Inserted: The VSAM number of inserted records.

Records Retrieved: The VSAM number of retrieved records.

Records Total: The VSAM number of logical records.

Records Updated: The VSAM number of updated records.

Recovery: The VSAM recovery attribute.

Reference Date: The last reference date.

Replicate: VSAM Attribute: Replicate.

Reuse: The VSAM reuse attribute.

Secondary Allocation Quantity: The secondary allocation amount.

Secondary Allocation Record Units: Secondary space record allocation units.

Secondary Allocation Record Value: Secondary space allocation value when specified in Records.

Secondary Allocation Units: The space allocation units. The values are as follows: CYLINDER/CYL, TRACK/TRK, ABSTRACT/ABS, AVG BLOCK/AVB, AVG BLOCK ROUNDED/AVD, AVG RECORD/AVR

Share Option: The VSAM share options (cross-region, cross-system). The values are as follows:

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

SMS Data Class: The SMS data class.

SMS Managed: The data set is SMS managed.

SMS Management Class: The SMS management class.

SMS Storage Class: The SMS storage class.

SMS Storage Group: The SMS storage group.

Spanned: The VSAM spanned attribute.

Speed: The VSAM speed attribute.

Stripe Count: The stripe count.

Swap Space: The cluster is a swap space.

System Determined Blocksize: The block size as determined by the system.

Temporary Dataset: This is a temporary data set.

Timestamp: The time that the row was created.

Unique: The VSAM unique attribute.

Unmovable: The data set is unmovable.

Unused GB: The gigabytes not used by the data set on the volume.

Unused MB: The megabytes not used by the data set on the volume.

Unused Tracks: The tracks unused by the data set on the volume.

Used GB: The gigabytes used by the data set on the volume.

Used MB: The megabytes used by the data set on the volume.

Used RBA: The VSAM high used RBA.

Used Tracks: The tracks used by the data set on the volume.

Using Cylinder Managed Space: Dataset is using Cylinder Managed Space on an Extended Address Volume.

Volser: The volume on which the data set is located.

Volume Count: The number of volumes on which the data set is located.

Volume Sequence: The volume sequence number for this portion of the data set.

VSAM Dataset Organization: The VSAM data set organization types. The values are as follows: KSDS, ESDS, RRDS, LDS, and AIX®

Dataset Attributes DSORG Detail

This group of attributes monitors the Top N data sets in terms of allocated space for a specific DSORG. You define the value of n in the Dataset Attributes Data Collection dialog box. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Address: Address of detail information in the data space of the data set.

ALET: ALET (Address List Entry Token) of information in the data space of the data set.

Allocated GB: Allocated GB for the data set.

Allocated MB: Allocated MB for the data set.

Allocated Tracks: Allocated tracks for the data set.

BLKSIZE: The block size of the data set.

Creation Date: The date data set was created.

Dataset Name: The name of the data set.

DSORG: The data set organization.

Extents: The number of extents.

LRECL: The logical record length.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

RECFM: The record format.

Reference Date: The date data set was last referenced.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the data set.

Unused MB: Unused MB for the data set.

Unused Tracks: Unused tracks for the data set.

Used GB: Used GB for the data set.

Used MB: Used MB for the data set.

Used Tracks: Used tracks for the data set.

Volser: The VOLSER on which the data set is located.

Dataset Attributes DSORG Summary

This group of attributes provides summary counts for the number of data sets in each encountered DSORG. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Allocated GB: The allocated GB for all data sets with this DSORG.

Allocated MB: The allocated MB for all data sets with this DSORG.

Allocated Tracks: The allocated tracks for all data sets with this DSORG.

DSORG String: The data set organization type for this data set in character string format.

DSORG: The data set organization type. The values are as follows:

- Unknown
- Hierarchical_File
- PDS_Extended
- Extended_Sequential
- Extended_Seq_Unmovable
- VSAM
- VSAM_Unmovable

ISAM
ISAM_Unmovable
Physical_Sequential
Physical_Seq_Unmovable
Direct_Access
Direct_Access_Unmovable
Partitioned
Partitioned_Unmovable
Unknown

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Datasets: The number of processed data sets within the indicated DSORG.

Percent of Allocated Space: The percentage of total allocated space for all DSORGs that data sets with this DSORG represent.

Percent of Total Datasets: The percentage of total data sets for all DSORGs that data sets with this DSORG represent.

Timestamp: The time that the row was created.

Unused GB: Total unused GB for data sets with this DSORG.

Unused MB: Total unused MB for data sets with this DSORG.

Unused Tracks: Total unused tracks for data sets with this DSORG.

Used GB: Total used GB for data sets with this DSORG.

Used MB: Total used MB for data sets with this DSORG.

Used Tracks: Total used tracks for data sets with this DSORG.

Dataset Attributes Extent Summary

This group of attributes provides data on the Top N largest data sets in terms of total number of extents. You define the value of n in the Dataset Attributes Data Collection dialog box. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Address: Address of detail information in the data space of the data set.

ALET: ALET (Address List Entry Token) of information in the data space of the data set.

Allocated GB: Allocated GB for the data set.

Allocated MB: Allocated MB for the data set.

Allocated Tracks: Allocated tracks for the data set.

Dataset Name: The name of the data set.

DSORG: The data set organization.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the data set.

Unused MB: Unused MB for the data set.

Unused Tracks: Unused tracks for the data set.

Used GB: Used GB for the data set.

Used MB: Used MB for the data set.

Used Tracks: Used tracks for the data set.

Volser: The VOLSER on which the data set is located.

Dataset Attributes Inefficient Blocksize Summary

This group of attributes monitors the Top N data sets in terms of used space, where the BLKSIZE is less than or equal to 128 bytes (in other words, the 'inefficient' threshold that the product defines). You define the value of n in the Dataset Attributes Data Collection dialog box. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Address: Address of detail information in the data space of the data set.

ALET: ALET (Address List Entry Token) of information in the data space of the data set.

Allocated GB: Allocated GB for the data set.

Allocated MB: Allocated MB for the data set.

Allocated Tracks: Allocated tracks for the data set.

BLKSIZE: The block size of the data set.

Creation Date: The date that the data set was created.

Dataset Name: The name of the data set.

DSORG: The data set organization.

Extents: The number of extents.

LRECL: The logical record length.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

RECFM: The record format of the data set.

Reference Date: The date that the data set was last referenced.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the data set.

Unused MB: Unused MB for the data set.

Unused Tracks: Unused tracks for the data set.

Used GB: Used GB for the data set.

Used MB: Used MB for the data set.

Used Tracks: Used tracks for the data set.

Volser: The VOLSER on which the data set is located.

Dataset Attributes Mature Dataset Summary

This group of attributes monitors the Top N data sets in terms of allocated space, where the number of days since last reference is greater than 366, which is the threshold that the product uses to categorize a data set as mature. You define the value of n in the Dataset Attributes Data Collection dialog box. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Address: Address of detail information in the data space of the data set.

ALET: ALET (Address List Entry Token) of information in the data space of the data set.

Allocated GB: Allocated GB for the data set.

Allocated MB: Allocated MB for the data set.

Allocated Tracks: Allocated tracks for the data set.

Creation Date: The date the data set was created.

Dataset Name: The name of the data set.

Days Since Creation: The number of days since creation.

Days Since Reference: The number of days since last reference.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Reference Date: The date the data set was last referenced.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the data set.

Unused MB: Unused MB for the data set.

Unused Tracks: Unused tracks for the data set.

Used GB: Used GB for the data set.

Used MB: Used MB for the data set.

Used Tracks: Used tracks for the data set.

Volser: The VOLSER on which the data set is located.

Dataset Attributes Never Referenced Dataset Summary

This group of attributes monitors the Top N data sets, in terms of allocated space, where the last reference date has not been set. (In other words, the data set has never been opened, or it has been opened using an access method that does not maintain the last reference date field.) You define the value of n in the Dataset Attributes Data Collection dialog box. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Address: Address of detail information in the data space of the data set.

ALET: ALET (Address List Entry Token) of information in the data space of the data set.

Allocated GB: Allocated GB for the data set.

Allocated MB: Allocated MB for the data set.

Allocated Tracks: Allocated tracks for the data set.

Creation Date: The date the data set was created.

Dataset Name: The name of the data set.

Days Since Creation: The number of days since creation.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the data set.

Unused MB: Unused MB for the data set.

Unused Tracks: Unused tracks for the data set.

Used GB: Used GB for the data set.

Used MB: Used MB for the data set.

Used Tracks: Used tracks for the data set.

Volser: The VOLSER on which the data set is located.

Dataset Attributes Newborn Dataset Summary

This group of attributes monitors the Top N non-temporary data sets, where the creation date is within 1 day of the collection cycle start time. You define the value of n in the Dataset Attributes Data Collection dialog box. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Address: Address of detail information in the data space of the data set.

ALET: ALET (Address List Entry Token) of information in the data space of the data set.

Allocated GB: Allocated GB for the data set.

Allocated MB: Allocated MB for the data set.

Allocated Tracks: Allocated tracks for the data set.

Dataset Name: Name of the data set.

Extents: Number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the data set.

Unused MB: Unused MB for the data set.

Unused Tracks: Unused tracks for the data set.

Used GB: Used GB for the data set.

Used MB: Used MB for the data set.

Used Tracks: Used tracks for the data set.

Volser: VOLSER on which the data set is located.

Dataset Attributes Reference Date Summary

This group of attributes monitors the count of the number of data sets that exist in each product-defined unique Reference Date (in other words, number of days since last reference) range. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Allocated GB: Allocated GB for all data sets in the range.

Allocated MB: Allocated MB for all data sets in the range.

Allocated Tracks: Allocated tracks for all data sets in the range.

High Days Since Reference: The highest value in the reference date range.

Low Days Since Reference: The lowest value in the reference date range.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Datasets: The number of data sets in the range.

Percent of Allocated Space: The percentage of total allocated space that data sets in this reference date range represent.

Percent of Total Datasets: The percentage of total processed data sets that data sets in this reference date range represent.

Timestamp: The time that the row was created.

Unused GB: Total unused GB for all data sets in the range.

Unused MB: Total unused MB for all data sets in the range.

Unused Tracks: Total unused tracks for all data sets in the range.

Used GB: Total used GB for all data sets in the range.

Used MB: Total used MB for all data sets in the range.

Used Tracks: Total used tracks for all data sets in the range.

Dataset Attributes SMS Class Detail

This group of attributes provides data on the Top N data sets within an SMS Data/Storage/Management Class or Storage Group in terms of allocated space. You define the value of n in the Dataset Attributes Data Collection dialog box. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Address: Address of detail information in the data space of the data set.

ALET: ALET (Address List Entry Token) of information in the data space of the data set.

Allocated GB: Allocated GB for the data set.

Allocated MB: Allocated MB for the data set.

Allocated Tracks: Allocated tracks for the data set.

Dataset Name: The name of a data set in this class.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

SMS Class: Name of the SMS class.

Timestamp: The time that the row was created.

Type: The type of construct. The values are as follows: D (Data Class); S (Storage Class); M (management Class); G (Storage Group).

Unused GB: Unused GB for the data set.

Unused MB: Unused MB for the data set.

Unused Tracks: Unused tracks for the data set.

Used GB: Used GB for the data set.

Used MB: Used MB for the data set.

Used Tracks: Used tracks for the data set.

Volser: The VOLSER on which this data set is located.

Dataset Attributes SMS Class Summary

This group of attributes provides a summary of the number of SMS-managed data sets encountered in each of the DFSMS Data, Storage and Management classes and Storage Groups defined on the system. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Allocated GB: Total allocated GB represented by all of the data sets in the class.

Allocated MB: Total allocated MB represented by all of the data sets in the class.

Allocated Tracks: Total allocated tracks represented by all of the data sets in the class.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Datasets: The number of data sets encountered in the indicated class.

Percent of Total Allocated: The percentage of total allocated space represented by data sets in the class.

Percent of Total Datasets: The percentage of total processed data sets represented by data sets in the class.

SMS Class: The name of the SMS class. A blank class indicates that there is no construct of this type defined for the data set. A class of **Unknown** indicates that the class name is not defined in the active DFSMS configuration.

Timestamp: The time that the row was created.

Type: The type of construct. The values are as follows: D (Data class); S (Storage class); M (Management class); G (Storage Group)

Unused GB: Total unused GB for all data sets in this class.

Unused MB: Total unused MB for all data sets in this class.

Unused Tracks: Total unused tracks for all data sets in the class.

Used GB: Total used GB for all data sets in this class.

Used MB: Total used MB for all data sets in this class.

Used Tracks: Total used tracks for all data sets in this class.

Dataset Attributes Space Allocated Summary

This group of attributes provides data on the Top N data sets in terms of total allocated space. You define the value of n in the Dataset Attributes Data Collection dialog box. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Address: Address of detail information in the data space of the data set.

ALET: ALET (Address List Entry Token) of information in the data space of the data set.

Allocated GB: Allocated GB for the data set.

Allocated MB: Allocated MB for the data set.

Allocated Tracks: Allocated tracks for the data set.

Dataset Name: The name of the data set.

DSORG: The data set organization.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the data set.

Unused MB: Unused MB for the data set.

Unused Tracks: Unused tracks for the data set.

Used GB: Used GB for the data set.

Used MB: Used MB for the data set.

Used Tracks: Used tracks for the data set.

Volser: The VOLSER on which the data set is located.

Dataset Attributes Space Unused Summary

This group of attributes provides data on the Top N data sets in terms of total unused space. You define the value of n in the Dataset Attributes Data Collection dialog box. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Address: Address of detail information in the data space of the data set.

ALET: ALET (Address List Entry Token) of information in the data space of the data set.

Allocated GB: Allocated GB for the data set.

Allocated MB: Allocated MB for the data set.

Allocated Tracks: Allocated tracks for the data set.

Dataset Name: The name of the data set.

DSORG: The data set organization.

Extents: The number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the data set.

Unused MB: Unused MB for the data set.

Unused Tracks: Unused tracks for the data set.

Used GB: Used GB for the data set.

Used MB: Used MB for the data set.

Used Tracks: Used tracks for the data set.

Volser: The VOLSER on which the data set is located.

Dataset Attributes System Summary

This group of attributes provides summary data representing all volumes and data sets encountered in a collection cycle. One row exists in this table. The data in the row is refreshed at each collection cycle. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Allocated GB: The total GB allocated to all data sets on all volumes.

Allocated MB: The total MB allocated to all data sets on all volumes.

Allocated Tracks: The total tracks allocated to all data sets on all volumes.

Available GB: The total GB available on all volumes encountered.

Available MB: The total MB available on all volumes encountered.

Available Tracks: The total tracks available (in other words, device capacity) on all volumes encountered.

Collection End Time: The date and time of day when the last collection cycle ended.

Collection Start Time: The date and time of day when the last collection cycle began.

Datasets Processed: The total number of data sets on all volumes that were processed during the collection cycle.

Largest Dataset Allocated GB: The size of the data set with the largest amount of allocated space in GB.

Largest Dataset Allocated MB: The size of the data set with the largest amount of allocated space in MB.

Largest Dataset Allocated Tracks: The size of the data set with the largest amount of allocated space in Tracks.

Largest Dataset Name: The name of the data set with the largest amount of allocated space.

Largest Dataset Volume: Volume on which the data set with the largest amount of allocated space is located.

Largest Unused Space Dataset GB: The size of unused space in GB within the data set with the largest amount of allocated and unused space.

Largest Unused Space Dataset MB: The size of unused space in MB within the data set with the largest amount of allocated and unused space.

Largest Unused Space Dataset Name: The name of the data set with the largest amount of allocated and unused space.

Largest Unused Space Dataset Tracks: The size of unused space in Tracks within the data set with the largest amount of allocated and unused space.

Largest Unused Space Dataset Volume: Volume on which the data set with the largest amount of allocated and unused space is located.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Oldest Reference Dataset Age: The number of days that oldest data set has gone unreferenced. The reference date must be non-zero. In other words, data sets that have never been referenced are not counted.

Oldest Reference Dataset Name: The name of the data set with the largest number of days since last reference. The reference date must be non-zero. In other words, data sets that have never been referenced are not counted.

Oldest Reference Dataset Volser: The volume on which the data set with the largest number of days since last reference is located.

Percent Allocated: The percentage of all available space that is allocated.

Percent Free: The percentage of all allocated space that is allocated but unused.

Percent Used: The percentage of all allocated space that is in use.

Status Message: Status of the Dataset Attributes Database data collection process.

Timestamp: The time that the table row was created.

Uncataloged Datasets: The total number of uncataloged data sets encountered on all volumes. This attribute is valid only if 'Collect Catalog Data' option is on in collection control table).

Note: The SYS1.VVDS and SYS1.VTOCIX data sets are known to always be uncataloged. They are not included in the Uncataloged Dataset total.

Unused GB: The total GB allocated but unused by all data sets on all volumes.

Unused MB: The total MB allocated but unused by all data sets on all volumes.

Unused Tracks: The total tracks allocated but unused by all data sets on all volumes.

Used GB: The total GB used by all data sets on all volumes.

Used MB: The total MB used by all data sets on all volumes.

Used Tracks: The total tracks allocated and used by all data sets on all volumes.

Volumes Online: Total number of DASD volumes that were online at the time that collection ran on the system.

Volumes Processed: The total number of volumes processed.

Dataset Attributes Uncataloged Detail

This group of attributes provides data on the Top N uncataloged data sets on a volume in terms of allocated space.

Note: The SYS1.VVDS and SYS1.VTOCIX data sets are known to always be uncataloged. They are not included in the Uncataloged Detail for a volume.

When the Collect Catalog Data option is enabled, the catalog data is retrieved for each data set encountered in the VTOC of the volume. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Address: Address of detail information in the data space of the data set.

ALET: ALET (Address List Entry Token) of information in the data space of the data set.

Allocated GB: Allocated GB for the data set.

Allocated MB: Allocated MB for the data set.

Allocated Tracks: Allocated tracks for the data set.

Dataset Name: Name of the uncataloged data set.

Extents: Number of extents.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Unused GB: Unused GB for the data set.

Unused MB: Unused MB for the data set.

Unused Tracks: Unused tracks for the data set.

Used GB: Used GB for the data set.

Used MB: Used MB for the data set.

Used Tracks: Used tracks for the data set.

Volser: The VOLSER on which this data set is located.

Dataset Attributes Uncataloged Summary

This group of attributes tracks the number of uncataloged data sets by volume. If no catalog entry is found for a data set, the count of uncataloged data sets on that volume is incremented. The space totals are also incremented.

Note: The SYS1.VVDS and SYS1.VTOCIX data sets are known to always be uncataloged. They are not included in the Uncataloged Detail for a volume.

When the Collect Catalog Data option is enabled, the catalog data is retrieved for each data set encountered in the VTOC of the volume. See Chapter 25, "Dataset Attributes Database," on page 311 for more information.

Allocated GB: Total allocated GB for all uncataloged data sets on the volume.

Allocated MB: Total allocated MB for all uncataloged data sets on the volume.

Allocated Tracks: Total allocated tracks for all uncataloged data sets on the volume.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Timestamp: The time that the row was created.

Total Datasets: Total number of uncataloged data sets on this volume.

Unused GB: Total unused GB for all uncataloged data sets on the volume.

Unused MB: Total unused MB for all uncataloged data sets on the volume.

Unused Tracks: Total unused tracks for all uncataloged data sets on the volume.

Used GB: Total used GB for all uncataloged data sets on the volume.

Used MB: Total used MB for all uncataloged data sets on the volume.

Used Tracks: Total used tracks for all uncataloged data sets on the volume.

Volser: The name of the volume on which at least one uncataloged data set was found.

Dataset Group Defs

This group of attributes provides information to other attribute groups. Users do not use these attributes. The attributes are documented for the purpose of background reference only.

Dataset Mask Data: Group mask bulk data. This contains a mask name status and space collection flag. These fields are repeated for each Group Mask being maintained. Mask name: max 44 characters; Status : 0=inactive, 1=active, 2=stopped

Dataset Group Details

Use the Dataset Group Details attributes to view detailed information on your data set. The attributes in this group include details regarding control intervals (CIs) and control areas (CAs).

% Free Bytes In CI: The percentage of bytes free across all CIs. (VSAM only.) You can enable aggregation for this attribute in Tivoli Data Warehouse.

% Free CIS In CA: The percentage of free CIs across all CAs in the data set. (VSAM only.) You can enable aggregation for this attribute in Tivoli Data Warehouse.

Alternate Index Attributes: The alternate index attributes. Values are UPGRADE, and NO UPGRADE.

Associated Entry Name: The associated entry names.

Average LRECL: The VSAM data set average logical record size.

Average MSR: The average MSR (millisecond response time) for this data set.

Bytes Free In Compression: The number of bytes of free space in compression. (VSAM only.)

CA Splits: The number of CA splits. You can enable aggregation for this attribute in Tivoli Data Warehouse.

A CA split is the movement of half of the records in an existing CA to a new CA because a record add or update cannot be accommodated in the existing CA. This condition results in two approximately half-full CAs instead of one full and one empty CA. A CA split is inefficient because it requires a large number of I/O operations. (VSAM only.)

Catalog Name: The name of the catalog that the data set belongs.

CI Size: The control interval (CI) size of the data set. A CI is a unit of data that is transferred between auxiliary storage and virtual storage when an I/O request is made. It contains records, free space, and control information.

CI Splits: The number of control interval (CI) splits. You can enable aggregation for this attribute in Tivoli Data Warehouse.

A CI split is the movement of some records from an existing CI to another free CI in the same control area because a record add or update cannot be accommodated in the existing one. This condition results in two half-empty CIs instead of one full and one empty CI. (VSAM only.)

Compressed User Data Size: The size of compressed user data.

Compression Indicator: The compression of the data set. Values are EXTENDED FORMAT, COMPRESSIBLE, or EXTENDED FORMAT.

Creation Date: The date when the data set was created.

Dataset % Free: The percentage of space that is free in this data set. This value is calculated using: $(\text{Dataset High Allocated RBA} - \text{Dataset High Used RBA}) \backslash \text{Dataset High Allocated RBA}$. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Dataset % Used: The percentage of space that is used in this data set. This is calculated using: $\text{Dataset High Used RBA} \backslash \text{Dataset High Allocated RBA}$. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Dataset Flags: The data set entry type. Values are PRIMARY and NOT PRIMARY.

Dataset Group Name: A user-defined name that identifies a group of data set name masks that make up the Dataset Group.

Dataset High Allocated RBA: The data set high-allocated RBA.

Dataset High Used RBA: The data set high-used RBA.

Dataset Mask: The user-defined data set mask that matches the data set.

Dataset Name: The name of the data set for which the information was collected.

Dataset Type: The data set entry type. Values are: NON-VSAM DATA SET, GENERATION DATA GROUP, CLUSTER, DATA COMPONENT, ALTERNATE INDEX, GENERATION DATA SET, INDEX COMPONENT, PATH, ALIAS, USER CAT CONNECTOR ENTRY.

Days Since Creation: The number of days since the data set was created.

Days Unreferenced: The number of days since the data set was last referenced.

EXCPS: The number of EXCPS for the data set. (VSAM only.) You can enable aggregation for this attribute in Tivoli Data Warehouse.

Expiration Date: The date when the data set expires.

Extended Attributes: The extended attributes.

Extended Format Attributes (HEX): The extended format attributes expressed as hexadecimal values.

Extended Format Attributes: The extended format attributes.

First Extent Tracks Allocated: The primary space allocation.

Free CIS Per CA: The number of free control intervals (CIs) per control area (CA). (VSAM only.) You can enable aggregation for this attribute in Tivoli Data Warehouse.

GDG Dataset Attributes: The Dataset Generation Data Group Attributes. Values are as follows:

- Delete oldest GDS GATLIMIT exceeded.
- Delete all GDSs when GATLIMIT exceeded.
- Do not scratch data set when rolled off.
- Scratch data set when rolled off if volume mounted.

GDG Last Altered Date: The date the data set Generation Data Group (GDG) was last changed.

GDG Maximum Datasets: The maximum data sets in the Generation Data Group (GDG).

KSDS Key Length: The length of the KSDS key.

KSDS Key Position: The relative position of key within the data set.

Last Backup Date: The date of the most recent data set backup.

Log Parameter Set: The value of the log parameter set. Values are NEVER SET, LOG NONE, LOG UNDO, LOG ALL.

Managed System: The internal name used to identify the system from which this row of data was obtained.

Maximum Buffer Size: The maximum buffer size.

Maximum I/Os: The maximum I/Os for a volume for this data set.

Maximum MSR: The maximum MSR (millisecond response time) for this data set.

Maximum Record Size: The maximum record size for this data set.

NON-VSAM Attributes: The non-VSAM attributes information for this data set. Values are Active GDS, Deferred GDS, Rolled off GDS, PDSE, POSIX DS, NON VSAM.

Number Bytes per CI: The number of free bytes across all control intervals (CIs). (VSAM only.)

Number CIS per CA: Contains the number of control intervals (CIs) that are contained in a control area (CA) for this data set. A control area is a group of control intervals. (VSAM only.)

Owner: Identifies the data set owner.

Primary Volume: The volume serial number for primary space allocation.

Records Deleted: The number of deleted records for the data set. (VSAM only.) You can enable aggregation for this attribute in Tivoli Data Warehouse.

Records Inserted: The number of inserted records for the data set. (VSAM only.) You can enable aggregation for this attribute in Tivoli Data Warehouse.

Records Retrieved: The number of retrieved records for the data set. (VSAM only.) You can enable aggregation for this attribute in Tivoli Data Warehouse.

Records Updated: The number of updated records for the data set. (VSAM only.) You can enable aggregation for this attribute in Tivoli Data Warehouse.

Recovery Time: The recovery time, TOD value, local.

Reference Date: The date of the last reference for the data set.

Request Type: Request type.

Secondary Space Allocation: The secondary space allocation.

SMFID: The SMFID of the system where the data was collected.

SMS Data Class: The SMS data class for the data set.

SMS Management Class: The SMS management class for the data set.

SMS Storage Class: The SMS storage class for the data set.

Space Allocation Units: The space units used for allocation. Values are AVG BLOCK, AVG BLOCK ROUNDED, CYLINDER, TRACK, and n/a.

Striping Counts: The striping count for a striped data set.

Timestamp: The timestamp of when the data was collected. Displayed in UTC (Coordinated Universal Time).

Total I/Os: The total sum of I/Os across all volumes for this data set.

Total Records: The number of logical records for the data set. (VSAM only.) You can enable aggregation for this attribute in Tivoli Data Warehouse.

Total Track Allocation: The total sum of tracks allocated across all volumes for this data set.

Total Track Free: The total sum of tracks free across all volumes for this data set.

Total Track Used: The total sum of tracks used across all volumes for this data set.

Volume Count: The number of volumes for this data set.

VSAM Dataset Attributes (HEX): The VSAM data set attributes. The attributes are Speed, Unique, Reusable, Erase, Inhibit update, Temporary export, and Track overflow expressed as hexadecimal values.

VSAM Dataset Attributes: The VSAM data set attributes. The attributes are Speed, Unique, Reusable, Erase, Inhibit update, Temporary export, and Track overflow.

VSAM Dataspace Attributes (HEX): Defines the attributes of the VSAM data space for the data set, in hexadecimal. The values are Swap, No Swap, and Page Space (indicates the data set is a pagespace).

VSAM Dataspace Attributes: Defines the attributes of the VSAM data space for the data set. The values are Swap, No Swap, and Page Space (indicates the data set is a pagespace).

VSAM Reuse Attributes (HEX): Defines the VSAM reuse attributes for the data set in hexadecimal. Values (in hexadecimal) can be: Dataset has RACF discrete profile, Index component data set, Reusable data set, Erase specified (cluster only), Not defined, swap space (cluster only), page space (cluster only).

VSAM Reuse Attributes: Defines the VSAM reuse attributes for the data set. Values can be: Dataset has RACF discrete profile, Index component data set, Reusable data set, Erase specified (cluster only), Not defined, swap space (cluster only), page space (cluster only).

VSAM Share Attributes: The share attributes of a VSAM data set. The share attribute is used when a data or index component of a cluster, alternate index, or the data component of a catalog can be shared among users. Cross Region specifies the amount of sharing allowed among regions within the same system or within multiple systems using global resource serialization (GRS). Cross System specifies the amount of sharing allowed among systems.

VSAM Type Attributes (HEX): The type attributes of a VSAM data set in hexadecimal. Values (HEX) can be: KSDS or non-KSDS, Write check, Imbed, Replicate, Key-range data set, RRDS (Relative Record Dataset), Spanned records allowed, Non-unique or unique keys allowed, LDS, or VRRDS. A LDS is a VSAM data set that contains data but no control information. A VRRDS is a Relative Record Dataset and is similar to a non-VSAM relative file.

VSAM Type Attributes: The type attributes of a VSAM data set. Values can be: KSDS or non-KSDS, Write check, Imbed, Replicate, Key-range data set, RRDS (Relative Record Dataset), Spanned records allowed, Non-unique or unique keys allowed, LDS, or VRRDS. An LDS is a VSAM data set that contains data but no control information. A VRRDS is a Relative Record Dataset and is similar to a non-VSAM relative file.

Dataset Group Detail Attrs

Use the Dataset Group Details attributes to view detailed information on your data set.

Allocation values: The values of the data set allocations.

Allocations: The data set allocations. Values are Owner, Volume, Space Type, Primary Space, Secondary Space, High Allocated RBA, High Used RBA, Percent Free, and Percent Used.

Attribute values: The values of the data set attributes.

Attributes: The data set attributes. Values are as follows:

| | | |
|-----------------------|--------------------------|-------------------------------|
| Key Length | NOIMBED | Invalid COMUDSIZ and UDATASIZ |
| Relative Key Position | REUSE | Block level compression |
| Average LRECL | NOREUSE | NON-VSAM Data Set |
| Maximum LRECL | SPANNED | Generation data Group |
| Buffer Space | NONSPANNED | Cluster |
| EXCP Exit | WRITECHK | Data Component |
| Share Options | NOWRITECHK | Alternate Index |
| Association | REPLICATE | Generation Data Set |
| UPGRADE | NOREPLICATE | Index Component |
| NOUPGRADE | NUMBERED | Path |
| SPEED | LINEAR | Alias |
| RECOVERY | NONUNIQKEY | User Cat Connector Entry |
| UNIQUE | RACF discrete profile | NON VSAM |
| NON-UNIQUE | Index Component Data Set | Active GDS |
| ERASE | Reusable | PDSE |
| NOERASE | Erase Specified | Rolled off GDS |
| INDEXED | SWAP Space | Deferred GDS |
| NON-INDEXED | Page Space | POSIX DS |
| IMBED | | |

Dataset Group Name: A user-defined name that identifies a group of data set name masks that make up the Dataset Group.

Dataset Name: A name of the data set.

Managed System: The internal name used to identify the system from which this row of data was obtained.

Request Type: Request type.

Statistic Values: The values of the data set statistics.

Statistics: The data set statistics. Values are Total Records, Deleted Records, Inserted Records, Updated Records, Retrieved Records, CI Size, CI Splits, CA Splits, Number CIs per CA, Free CIs per CA, Free Space per CI, Percent Free CIs, Percent Free Bytes, EXCPS, Extents, and System Timestamp.

Dataset Group Details Volume

Use the Dataset Group Details Volume attributes to view detailed information on the data set of a volume.

Allocation Flag: The data set volume flags. The following are valid values:

CANDIDATE VOLUME
CANDIDATE VOLUME\CATALOG ENTRY ONLY
CANDIDATE VOLUME\OVERFLOW VOLUME
CANDIDATE VOLUME\OVERFLOW VOLUME\CATALOG ENTRY ONLY
CATALOG ENTRY ONLY
OVERFLOW VOLUME
OVERFLOW VOLUME\CATALOG ENTRY ONLY
PRIMARY VOLUME
PRIMARY VOLUME\CANDIDATE VOLUME
PRIMARY VOLUME\CANDIDATE VOLUME\CATALOG ENTRY ONLY
PRIMARY VOLUME\CATALOG ENTRY ONLY
PRIMARY VOLUME\INDIRECT
PRIMARY VOLUME\OVERFLOW VOLUME
PRIMARY VOLUME\OVERFLOW VOLUME\CATALOG ENTRY ONLY

PRIMARY VOLUME\ SYMBOLIC
n/a

Associated Entry Name: The associated entry name. (Example: VSAM Cluster component.)

Average MSR: The average MSR (millisecond response time) for the data set on this volume. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Dataset Group Name: A user-defined name that identifies a group of data set name masks that comprise the Dataset Group.

Dataset Name: The name of the data set for which the information was collected.

Extended Address Space Eligible: Data set is eligible to use Extended Address Space of an Extended Address Volume.

Extents: The number of extents for this data set.

GDG Generation Level: The Dataset Generation Group Level.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Percent Free: The percentage of tracks free for this data set on this volume. This is calculated using this formula: $(\text{Tracks Allocated} - \text{Tracks Free}) / \text{Tracks Free}$

Percent Used: The percentage of tracks used for this data set on this volume. This is calculated using this formula: $(\text{Tracks Allocated} - \text{Tracks used}) / \text{Tracks Used}$

Physical Block Size: The number of bytes per track.

Request Type: Request type.

Sequence Number: Identifies the file sequence number.

Sequence Order: The sequence order of the volume.

SMFID: The SMFID of the system where the data was collected.

Timestamp: The time stamp of when the data was collected.

Total I/Os: The total I/Os for this data set on this volume. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Tracks Allocated: The number of tracks allocated by this data set on this volume. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Tracks Free: The number of tracks free for this data set on this volume. This is calculated using: $\text{Tracks Allocated} - \text{Tracks Used}$.

Tracks Used: The number of tracks used by this data set on this volume. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Using Cylinder Managed Space: Dataset is using Cylinder Managed Space on an Extended Address Volume.

Volume: The Volume Serial Number for primary or candidate space allocation.

VSAM Associated Entry Name: The name of the VSAM Dataset Associated entry. For example, VSAM Cluster component.

VSAM High Allocation RBA: The VSAM high-allocated Relative Byte Address.

VSAM High CI RBA: The VSAM Relative Byte Address data control interval with high key.

VSAM High Used RBA: The VSAM high-used Relative Byte Address.

VSAM Percent Free: The percentage of space that is free in this VSAM data set. This is calculated using: $(\text{High Allocated RBA} - \text{High Used RBA}) \setminus \text{High Allocated RBA}$.

VSAM Percent Used: The percentage of space that is used in this VSAM data set. This is calculated using: $\text{High Used RBA} \setminus \text{High Allocated RBA}$.

Dataset Group Summary

Use the Dataset Group Summary attributes to view information on your data set.

Dataset Group Name: A user-defined name that identifies a group of data set name masks that make up the data set group.

Dataset with Highest MSR: The data set with the slowest MSR (millisecond response time) within the group.

Dataset with Largest Allocation: The data set with largest space allocation within the group.

Dataset with Largest Free Space: The data set with largest amount of unused space within the group.

Dataset with Largest Usage: The name of the data set with largest allocation, maximum highest space allocation, and oldest referenced date within the group, in tracks.

Dataset with Oldest Referenced: The data set with oldest referenced date (read, write, or opened) within the group.

Group Total Average MSR: The total average MSR (millisecond response time) for all the active data sets within the group, in milliseconds. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Group Total Free Space: The total free space for all the data sets within the group, in tracks. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Group Total I/Os: The total I/Os for all the active data sets within the group. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Group Total Space Allocation: The total allocated space for all the data sets within the group, in tracks. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Group Total Used Space: The total used space for all the data sets within the group, in tracks. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Highest MSR: The highest MSR (millisecond response time) for the identified data set within the group, in milliseconds. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Largest Space Allocation: The space allocation for the identified data set, in tracks. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Largest Space Free: The free space for the identified data set, in tracks. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Largest Space Usage: The largest allocation for the identified data set within the group.

Largest Usage Oldest Reference Date: The reference date for the identified data set within the group.

Managed System: An internal name used to identify the system from which this row of data was obtained.

Oldest Date Referenced: The last referenced date for the identified data set.

Percent Free: The percentage of the group total space allocation not in use. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Percent Used: The percentage of the group total space allocation in use. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Total Datasets: The total number of data sets in this data set group. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Dataset Performance

The Dataset Performance attributes provide performance information for a data set by application or summarized for all applications accessing the data set. They provide metrics at the I/O component level.

ASID: The address space ID of the job using this data set.

Cache Hit Percent: The sum of cache read hits plus cache write hits divided by the sum of cache-eligible reads plus cache-eligible writes.

Cache Read Hit Percent: The number of cache read hits divided by cache-eligible reads.

Cache Write Hit Percent: The number of cache write hits divided by the number of cache-eligible writes.

Connect Time Percent: The ratio of connect time divided by MSR (millisecond response time), expressed as a percentage.

Connect Time: Represents time in milliseconds spent locating data on the track and transferring between the DASD device and central storage.

Dataset Name: The MVS data set name.

Device Active Only Percent: The average device active only time per I/O in milliseconds.

Device Active Only Time: That portion in milliseconds of total MSR (millisecond response time) that occurs after 'CHANNEL-END' but before 'DEVICE-END'.

Device Address: The MVS device number, assigned by HCD, associated with the device on which the data set is located.

Disconnect Time Percent: The ratio of disconnect time divided by MSR (millisecond response time), expressed as a percentage.

Eligible Cache Percent: The percentage of total I/Os that were cache eligible.

Total I/O Count: The total number of I/Os (read and write) processed for the data set.

I/O Per Second: The average number of I/O operations per second for this data set during this interval.

IOSQ Time Percent: The ratio of IOSQ time divided by MSR (millisecond response time), expressed as a percentage.

IOSQ Time: The average time in milliseconds that an I/O waits because the device is already busy (signified by the UCBBUSY bit for a device being on).

JOB_START: Job Start value.

Jobname: The MVS JOBNAME that is using this data set.

Low SC Objective: The lower value of direct or sequential storage class MSR (millisecond response time) objective.

Managed System: The SMS name of this system.

NON_C_READ: Non C Read value.

NON_C_WRIT: Non C Writ value.

Pend Time Percent: The ratio of pending time divided by MSR (millisecond response time), expressed as a percentage.

Pend Time: The time in milliseconds that I/O is delayed in the path to the device. Pending time can be attributable to the channel or control unit path being busy.

Read Disconnect Time: The average time in milliseconds that read I/O waits because the device is not connected to the channel.

Read I/O Count: Total number of read I/Os processed for the data set.

Read Write Percent: The sum of cache-eligible reads plus non-cache reads divided by the total I/O.

Response Time Index: The total response time in milliseconds for this data set during this interval.

Response Time: The total millisecond response time for this data set during this interval.

SC Direct MSR Objective: The target direct response time specified for this storage class. This value is used during allocation to select a volume with acceptable performance characteristics. This value also influences caching.

SC Sequential MSR Objective: A construct in the SMS storage class that is used during allocation and indicates the desired response time for sequential accesses to this data set.

Storage Class Name: The name of the SMS storage class to which this data set belongs.

Time Stamp: The date and time used to control the scope of the query.

Total Disconnect Time: The average time in milliseconds that all I/O (read and write) waits because the device is not connected to the channel.

Volume: The six-character name of the DASD volume where this data set is located.

Write I/O Count: Total number of write I/Os processed for the data set.

Write Disconnect Time: The average time in milliseconds that write I/O waits because the device is not connected to the channel.

Dataset SMS

The Dataset SMS attributes provide information about the SMS constructs associated with a data set. Additionally, the number of stripes, if applicable, is provided.

Data Class: The SMS data class. This value describes attributes and space requirements for SMS managed data sets. Data Classes contain attributes that govern data set allocation. These attributes are set by the storage administrator and users cannot change them.

Dataset Name: The MVS data set name.

Managed System: The SMS name of this system.

Management Class: The SMS management class. Management class is a collection of management attributes, defined by the storage administrator, used to control the release of allocated but unused space; to control the retention, migration and backup of data sets; to control the retention and backup of aggregate groups; to control the retention, backup and class transition of objects.

Storage Class: The SMS storage class.

A collection of storage attributes that identify performance goals and availability requirements, defined by the storage administrator, used to select a device that can meet those goals and requirements.

Stripe Count: The number of stripes which comprise this data set.

A stripe is the portion of a striped data set that is located on one volume. The records in that portion are not always logically consecutive. The system distributes records among the stripes such that the volumes can be read from or written to simultaneously.

Dataset Space

The Dataset Space attributes provide information about the space allocation and usage for a data set.

Block Size: The physical record size for storage of data in a data set and transfer of that data during an I/O operation. For a DASD data set, block size is maintained in the data set control block (DSCB).

Because more data is being transferred, large block sizes result in higher connect time. Larger blocks with fewer buffers might be appropriate for sequential data sets. Consider using DASD cache control units for heavily used, read-only sequential data sets.

Random access files need smaller block sizes with a larger number of index buffers (when indexed). For PDS libraries, block sizes should usually reflect average number size.

Dataset Name: The 44-character MVS name of this data set.

Dataset Type: The data set organization type for this data set. Values are as follows:

Direct_Access

Direct_Access_Unmovable

Extended_Seq_Unmovable

Extended_Sequential
Hierarchical_File
ISAM
ISAM_Unmovable
Partitioned
Partitioned_Unmovable
PDS_Extended
Physical_Seq_Unmovable
Physical_Sequential
VSAM
VSAM_Unmovable
Unknown

Date Created: The date that the current DASD space for this data set was first allocated.

Date Referenced: The most recent date when this data set was opened.

Expiration Date: A date stamp kept on the DASD device where this data set is allocated that indicates when this data set is no longer required. This date stamp is written when the data set is first allocated. Attempting to delete the data set before this date results in a message to the MVS system operator.

Extended Address Space Eligible: Data set is eligible to use Extended Address Space on an Extended Address Volume.

Logical Record Length: The DCB parameter that specifies the length, in bytes, of each record in the data set (fixed format), or the maximum record length (variable or undefined format).

Managed System: The SMS name of this system. An internal node name used to identify the system from which this row of data was obtained.

Number of Extents: The number of separate areas of DASD space allocated to this data set.

Record Format: The DCB parameter that specifies the type and blocking characteristics of records in the data set. These values are:

- F Fixed Length Records
- FB Fixed Length Blocked Records
- V Variable Length Records
- VB Variable Length Blocked Records
- U Undefined Length Records
- D ISCI/ASCII Variable Length Records

See the IBM publication "MVS DFP 3.3 Using Data Sets" for more information on this and other DCB parameters.

Record Secondary Allocation Unit: The allocation unit for secondary space requests.

Record Secondary Allocation: The amount of space allocated to this data set in secondary allocations.

Secondary Allocation Quantity: The number of space units which are allocated for each data set extent beyond the primary allocation.

Secondary Allocation Unit: The unit in which space is allocated for each data set extent beyond the primary allocation.

Tracks Allocated: The number of tracks reserved for the use of this data set.

Tracks Used Percent: The ratio of tracks used divided by the number of tracks allocated. The value is obtained by dividing the value of the **Tracks Used** attribute by the value of the **Tracks Allocated** attribute.

Tracks Used: The number of tracks that actually contain data for this data set.

Unused Tracks: The number of tracks allocated to the data set that are not being used. For VSAM data sets, this is the number of tracks that are not being used at the current high RBA.

Using Cylinder Managed Space: Data set is using Cylinder Managed Space on an Extended Address Volume.

Volume Count: The number of DASD volumes used to contain this data set.

Volume Extents: The count of extents for this data set on this volume.

Volume Sequence: The sequence number of the volume in a multi-volume data set.

Volume: The six-character name assigned to this DASD volume associated with this row of data.

VSAM Key Displacement: The distance in bytes of the key from the beginning of the data record. See the IBM publication "MVS DFP 3.3 Using Data Sets" for more information on this parameter.

VSAM Key Length: The DCB parameter that specifies the length in bytes (0 - 255) of a record key. See the IBM publication "MVS DFP 3.3 Using Data Sets" for more information on this and other DCB parameters.

Dataset Users

The Dataset Users attributes provide the names of applications using or waiting for access to a data set. Also provided are the name of the application, the system ID, and the type of control (shared or exclusive) that the application requires.

Application: The application name associated with this row.

Application Type: Indicates the type of application.

Application Waiting: Indicates whether the application is waiting or not.

- No = The application is not waiting because it has either exclusive or shared control of the data set.
- Yes = The application is waiting because control of the data set is owned by another user.

Control: Indicates whether this application shares this data set with other applications (Shared) or not (Exclusive).

Dataset Name Index: The name of the data set that is enqueued.

Managed System: The SMS name of this system.

System: The SMFID (Systems Management Facility ID) of the system on which this application is running.

HSM CDS

The HSM CDS attributes provide space allocation and utilization information for each of the HSM control data sets.

DDNAME: The DDNAME of this data set or cluster. Specific DDNAMEs are associated with specific HSM functions:

DDNAMEs MIGCAT, MIGCAT2-4, BAKCAT, and BAKCAT2-4 are associated with migration.

DDNAME OFFCAT is associated with the OCDS, which contains information about tape backup volumes and tape migration level 2 volumes.

DDNAME Journal is associated with the HSM Journal.

The journal contains a duplicate of each record in the control data sets that has changed since the BCDS, and OCDS were last backed up.

Display Order: The order in which the data sets are displayed.

Managed System: The SMS name of this system.

Number of Extents Data Component: The number of extents in the data component of this cluster. A high number of extents can indicate a need for re-allocation or reorganization of this component. Use IDCAMS LISTCAT to review splits.

Number of Extents Index Component: The number of extents in the index component of this cluster. A high number of extents can indicate a need for re-allocation or reorganization of this component. Use IDCAMS LISTCAT to review splits.

Percent Available Space Data Component: The amount of available space in an HSM control data set, expressed as a percentage of the total available space. This value is based on the high allocated RBA and the available space in the component. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Percent Available Space Index Component: The amount of available space in an HSM control data set, expressed as a percentage of the total available space. This value is based on the high allocated RBA and the available space in the component. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Percent Free Space Data Component: The amount of unused space in an HSM control data set, expressed as a percentage of the total available space. This is based on the high allocated RBA and high used RBA. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Percent Free Space Index Component: The amount of unused space in an HSM control data set, expressed as a percentage of the total available space. This is based on the high allocated RBA and high used RBA. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Size (KB) Data Component: The total amount of space, in units of 1024 bytes, allocated to this component. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Size (KB) Index Component: The total amount of space, in units of 1024 bytes, allocated to this component. You can enable aggregation for this attribute in Tivoli Data Warehouse.

HSM Common Storage

This section describes the attributes in this group, including information regarding migration work elements (MWE).

Active Limit Percent: DFSMSHsm active limit for the common service area storage allocation. This limit is a percent of the specified maximum limit. The default is 90% of the specified maximum limit. After the specified percentage of common service area has been allocated and DFSMSHsm is active, only batch WAIT MWEs are added to the common service area queue.

Inactive Limit Percent: DFSMSHsm inactive limit for the common service area storage allocation. This limit is a percent of the specified maximum limit. After the percent of the specified maximum limit has been allocated and DFSMSHsm is inactive, no MWEs are added to the common service area queue. The default is 30% usage of the maximum limit common service area.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Maximum Common Storage: Maximum amount of common storage area (in bytes) to be allocated to all MWEs. After the limit has been reached, no MWEs are added to the common service area queue. The default is 102400 (100k) bytes.

Maximum MWE Address Space: Maximum number of MWEs to allow per address space. If any NOWAIT MWEs are added to the common service area queue for this address space after the maximum number has been reached, they are flagged to indicate the common service area storage allocated to the MWE is to be freed when the ARCCTL has copied the MWE into DFSMSHsm address space. The default is 4 MWEs.

Used Common Storage: The number of bytes of the common storage area that are currently being used.

HSM Dataset Recover Version

The HSM Dataset Recover Version attributes provide summary information regarding the backup versions of HSM data sets.

Backup Date Time: The date and time of the backup.

Catalog Status: The catalog status of the data set.

Generation: Generation number of the backup.

Version: Version number of the backup.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

HSM Function Details

This section describes the attributes in this group.

Automatic Requests: Number of DFHSM requests of this function type issued automatically. This metric is extracted from the DFHSM daily statistics record.

Average Allocate Time: The average allocation time (in seconds) for a request to be satisfied on this HSM host.

Average Completion Time: The average completion time (in seconds) for a request to be satisfied on this HSM host.

Average Process Time: The average processing time (in seconds) for a request to be satisfied on this HSM host.

Average Queue Time: The average queue time (in seconds) for a request to be satisfied on this HSM host.

Failed Requests: Count of failed requests of this function type extracted from the DFHSM daily statistics record.

Function Type: The function that was executed. Values are as follows:

- Primary_to_Level1
- Level1_to_Level2
- Primary_to_Level2
- Level1_to_Primary
- Level2_to_Primary
- Delete_Migrated_Data_Set
- Daily_Backup
- Spill_Backup
- Recovery
- Recycle_Backup_Volume
- Delete_Data_Set
- Recycle_L2_Tape_Volumes
- Delete_Backup_Versions

HSM Host Name: The name of the HSM host. This is the started task name.

HSM_Host_ASID: The ASID of the HSM host.

HSM_Host_ID: The ID of the HSM host.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Megabytes Read: Total megabytes read for this function type, extracted from the DFHSM daily statistics record. (_=-2; n/a=-1)

Megabytes Written: Total megabytes written for this function type, extracted from the DFHSM daily statistics record. (_=-2; n/a=-1)

Tracks Read: Total DASD tracks read for this function type, extracted from the DFHSM daily statistics record.

Tracks Written: Total DASD tracks written for this function type, extracted from the DFHSM daily statistics record.

User Requests: Number of requests initiated by user command for this function type as extracted from the DFHSM daily statistics record on this HSM host.

HSM Function Statistics

The HSM Function Statistics attributes provide summary statistics for the HSM activities, such as movement of data sets from level 1 to level 2 archives.

Automatic Requests: The number of DFHSM requests issued automatically by this function type, as extracted from the DFSMSHsm daily statistics record. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Average Allocate Time: The average allocation time (in seconds) for a request to be satisfied across all HSM hosts.

Average Completion Time: The average completion time (in seconds) for a request to be satisfied across all HSM hosts.

Average Process Time: The average processing time (in seconds) for a request to be satisfied across all HSM hosts.

Average Queue Time: The average queue time (in seconds) for a request to be satisfied across all HSM hosts.

Failed Requests: Represents the count of failed requests of this function type extracted from the DFSMSHsm daily statistics record. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Function Index: The function type, as extracted from the DFSMSHsm daily statistics record, to which this row of data applies.

Function Type: The type of function used. Valid types are as follows:

- Daily_Backup
- Delete_Backup_Versions
- Delete_Data_Set
- Delete_Migrated_Data_Set
- Level1_to_Level2
- Level1_to_Primary
- Level2_to_Primary
- Primary_to_Level1
- Primary_to_Level2
- Recovery
- Recycle_Backup_Volume
- Recycle_L2_Tape_Volumes
- Spill_Backup

Function Type String: The function type, displayed as a string.

Managed System: The SMS name of this system.

Megabytes Read: The total number of megabytes read for this function type, extracted from the DFSMSHsm daily statistics record. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Megabytes Written: The total number of megabytes written for this function type, extracted from the DFSMSHsm daily statistics record. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Tracks Read: The total number of DASD tracks read for this function type, extracted from the DFSMSHsm daily statistics record. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Tracks Written: The total number of DASD tracks written for this function type, extracted from the DFSMSHsm daily statistics record. You can enable aggregation for this attribute in Tivoli Data Warehouse.

User Requests: The number of requests initiated by user command for this function type as extracted from the DFSMSHsm daily statistics record. You can enable aggregation for this attribute in Tivoli Data Warehouse.

HSM Function Summary

The HSM Function Summary attributes provide status and queue information about each of the HSM functions.

Active Requests: The count of requests currently active for this function. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Dataset Requests: The count of requests either queued or active for this function for data sets. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Function: The status of this HSM function (Held or Not Held).

Function Str: The name of the function for this data row. Valid names are Migration, Recall, Backup, Recovery or Dump.

Function Status: The status of this HSM function. Values are Held, Not Held, or Held EOV.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Volume Requests: The count of requests either queued or active for this function for volumes. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Waiting Requests: The count of requests currently waiting for this function. Requests are queued either when the function is held, or maxtasks for this function is reached. You can enable aggregation for this attribute in Tivoli Data Warehouse.

HSM Host Function Status

This section describes the attributes in this group.

Aggregate Backup Status: The status of the HSM host to do aggregate backups. (Not_Held=0; Held=1)

Aggregate Recovery Status: The status of the HSM host to do aggregate recoveries. (Not_Held=0; Held=1)

Audit Activity: The HSM host is currently running audit (Active) or not running audit. (Inactive=0; Active=1)

Audit Status: The status of the audit function on this HSM host. (Not_Held=0; Held=1)

Automatic Backup Status: The status of the HSM host to do automatic backups. (Not_Held=0; Held=1; Held_EOV=2)

Automatic Dump Status: The status of the HSM host to do automatic dumps. (Not_Held=0; Held=1; Held_EOV=2)

Automatic Migration Status: The status of the HSM host to do automatic migrations. (Not_Held=0; Held=1; Held_EOV=2)

Backup Status: The status of the backup function on this HSM host. (Not_Held=0; Held=1; Held_EOV=2)

BCDS Backup Status: Indicates if the backup control data set is currently being backed up on this HSM host. (Inactive=0; Active=1)

Dataset Backup Status: The status of the HSM host to run data set backup functions. (Not_Held=0; Held=1; Held_DASD=2; Held_Tape=3)

Dump Status: The status of the dump function on this HSM host. (Not_Held=0; Held=1; Held_EOV=2)

Expire Backup Version Status: The status of the host to expire backup versions. (Not_Held=0; Held=1)

HSM Host ASID: The ASID of the HSM host.

HSM Host ID: The ID of the HSM host.

HSM Host Name: The name of the HSM host. This is the started task name.

HSM Host Type: The HSM host type. (Unknown=0; Main=1; Main(Primary)=2; Aux=5; and Aux(Primary)=6)

Journal CDS Backup Status: Indicates if the journal control data set is currently being backed up on this HSM host. (Inactive=0; Active=1)

List Activity: Indicates if the list function is active on this HSM host. (Inactive=0; Active=1)

List Status: The status of the HSM host to run list functions. (Not_Held=0; Held=1)

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Migration CDS Backup Status: Indicates if the migration control data set is currently being backed up on this HSM host. (Inactive=0; Active=1)

Migration Status: The status of the HSM host to run migration functions. (Not_Held=0; Held=1; Held_EOV=2)

Offline CDS Backup Status: Indicates if the offline control data set is currently being backed up on this HSM host. (Inactive=0; Active=1)

Recall Status: The status of the HSM host to run recall functions. (Not_Held=0; Held=1; Held_TSO=2)

Recover Status: The status of the HSM host to run recover functions. (Not_Held=0; Held=1; Held_EOV=2)

Recycle Activity: Indicates if the recycle function is active on this HSM host. (Inactive=0; Active=1)

Recycle Status: The status of the HSM host to run recycle functions. (Not_Held=0; Held=1)

Report Activity: Indicates if the report function is active on this HSM host. (Inactive=0; Active=1)

Report Status: The status of the HSM host to run report functions. (Not_Held=0; Held=1)

Tape Copy Status: The status of the HSM host to run tape copy functions. (Not_Held=0; Held=1)

Tape Recall Status: The status of the HSM host to run tape recall functions. (Not_Held=0; Held=1; Held_TSO=2)

Tape Recovery Status: The status of the HSM host to run tape recovery functions. (Not_Held=0; Held=1)

Tape Replace Status: The status of the HSM host to run tape replace functions. (Not_Held=0; Held=1)

HSM Host Request Summary

This section describes the attributes in this group.

Backup Requests Queued: The number of backup requests queued to this HSM host.

Delete Requests Queued: The number of delete requests queued to this HSM host.

Dump Requests Queued: The number of dump requests queued to this HSM host.

HSM Host ASID: The ASID of the HSM host.

HSM Host ID: The ID of the HSM host.

HSM Host Name: The name of the HSM host. This is the started task name.

HSM Host Type: The HSM host type. Unknown=0; Main=1; Main(Primary)=2; Aux=5; and Aux(Primary)=6

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Migrate Requests Queued: The number of migrate requests queued to this HSM host.

Recall Requests Queued: The number of recall requests queued to this HSM host.

Recovery Requests Queued: The number of recovery requests queued to this HSM host.

HSM Host Status

This section describes the attributes in this group.

Avg Dump Alloc Time: The average time (in seconds) spent waiting for allocation for dump requests on this host.

Avg Dump Elapsed Time : The average time (in seconds) it took to process dump requests on this host.

Avg Dump Process Time: The average time (in seconds) spent processing dump requests on this host.

Avg Dump Queue Time: The average time (in seconds) spent on the queue for dump requests on this host.

Failed Aggregate Backups: The number of failed aggregate backups on this host.

Failed Data Set Restores: The number of failed data set restores on this host.

Failed Fast Replic Recovers: The number of failed fast replicate recovers on this host.

Failed Fast Replicat Backups: The number of failed fast replicate backups on this host.

Failed Volume Dump Copies: The number of failed volume dump copies on this host.

Failed Volume Dumps: The number of failed volume dumps.

Failed Volume Restores: The number of volume restores that failed on this host.

HSM Host ASID: The ASID of the HSM host.

HSM Host ID: The ID of the HSM host.

HSM Host Name: The name of the HSM host. This is the started task name.

HSM Host Start Time: Start time and date for the HSM address space.

HSM Host Status: The status of the HSM host. This value is available only if the host has been previously started and another host has not been started in its place. (Active=0; Inactive=1)

HSM Host Type: The HSM host type. (Unknown=0; Main=1; Main(Primary)=2; Aux=5; and Aux(Primary)=6)

Interval Abends: Count of HSM requests that abended during this interval.

Interval CPU Usage: CPU time consumed by DFHSM to service requests during this interval.

Interval Requests: Total number of HSM requests during this interval.

Interval Start Time: Start time for the current HSM interval.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Recycled Back Up Volumes: Total number of BACKUP volumes recycled during this interval.

Recycled Migration Volumes: Total number of MIGRATION volumes recycled during this interval.

System Requested Dump Copies: The number of volume dump copies on this host requested by the system.

System Requested Dumps: The number of volume dumps on this host requested by the system.

Total Aggregate Backups: The number of aggregate backups done on this host.

Total Data Set Restores: The number of data set restores done on this host.

Total Fast Replic Recovers: The number of fast replicate recovers done on this host.

Total Fast Replicat Backups: The number of fast replication backups on this host.

Total Volume Dump Copies: The number of volume dump copies on this host.

Total Volume Dumps: The number of complete volume dumps attempted on this host.

User Requested Dump Copies: The number of volume dump copies on this host requested by a user.

User Requested Dumps: The number of volume dumps on this host requested by a user.

Volume Backups: Total number of requests for volume BACKUP during this interval.

Volume Migrations: Total number of requests for volume MIGRATIONs during this interval.

Volume Restores: Total number of requests for volume restores during this interval.

HSM Private Storage

This section describes the attributes in this group.

Central Storage Fixed Frames: The number of central storage fixed frames in use by the HSM host. (n/a=-1)

Central Storage Real Frames: The number of frames of central storage in use by the HSM host. (n/a=-1)

ESTOR Frames: The number of expanded storage frames in use by the HSM host. (n/a=-1)

HSM Host ASID: The ASID of the HSM host.

HSM Host ID: The ID of the HSM host.

HSM Host Name: The name of the HSM host. This is the started task name.

HSM Host Type: The HSM host type. (Unknown=0; Main=1; Main(Primary)=2; Aux=5; and Aux(Primary)=6)

Managed System: An internal node name used to identify the system from which this row of data was obtained.

HSM Requests

The HSM Requests attributes provide detailed information about the current and queued requests for HSM functions, such as data sets queued for recall.

Aggregate Group Name: The name of an HSM aggregate group to which this data set belongs.

An aggregate group is an SMS grouping of data sets where all data sets in the group have identical backup or recovery attributes (such as medium, retention period and destination).

You can define all data sets related to an application in an aggregate group so that they all can be treated alike and at the same time by SMS.

Dataset Name: The name of the data set for which a request has been made.

Function: The type of request. Each request type has a different queue. The types are as follows: Migration, Recall, Backup, Recovery, Delete, Vol_Recover, and Dump.

Function Index: The HSM function type to which this data row pertains.

HSM Host ID: The ID associated with this HSM host.

HSM Host Name: The name of HSM host associated with the request.

Managed System: The SMS name of this system.

Queue Time: The time that the request was queued.

Request Age: The number of minutes this request has waited to be serviced. A large number of queued requests or a large wait time can be caused by an insufficient number of tasks started of this function type, a high rate of requests or this function type being held.

Requesting Jobname: The name of the job/started task/TSOUSERID that submitted this request.

Request Number: The number assigned to this request. Requests are numbered in ascending sequence by DFHSM as they are received.

Request Status: The status of request. The request status defines the state of the request. The following states are valid:

Queued = Request has been received and is waiting for higher priority requests to complete.

Active = Request is currently being serviced.

Abended = The request has completed unsuccessfully.

Complete = The request has completed successfully.

Cancelled = The request was cancelled before it completed.

Request Time: The time that DFHSM received this request.

Request Type: Denotes whether this request is for a specific volume or specific data set.

Request Volume: The specific DASD volume related to this request.

Requesting User ID: The user ID associated with this request.

User Waiting: Specifies whether this is a synchronous (waiting) or asynchronous (not waiting) request.

HSM Status

The HSM Status attributes provide information on the status of HSM, including whether it is active, the start time, and some summary activity information.

Active HSM Hosts: The number of active HSM hosts on this z/OS image.

Avg Dump Alloc Time: The average time (in seconds) that a dump request spent waiting for allocation across all HSM hosts on the z/OS image.

Avg Dump Elapsed Time: The average time (in seconds) that a dump took to complete across all HSM hosts on the z/OS image.

Avg Dump Process Time: The average time (in seconds) that a dump request spent processing across all HSM hosts on the z/OS image.

Avg Dump Queue Time: The average time (in seconds) that a dump request spent waiting for execution across all HSM hosts on the z/OS image.

Failed Aggregate Backups: The number of failed fast aggregate backups across all HSM hosts on the z/OS image.

Failed Data Set Restores : The number of failed fast volume dump copies across all HSM hosts on the z/OS image.

Failed Fast Replic Recovers: The number of failed fast replicate recovers across all HSM hosts on the z/OS image.

Failed Fast Replicat Backups: The number of failed fast replicate backups across all HSM hosts on the z/OS image.

Failed Volume Dump Copies: The number of failed fast volume dump copies across all HSM hosts on the z/OS image.

Failed Volume Dumps: Failed volume dumps.

Failed Volume Restores: Failed volume restores.

HSM ASID: The address space in which the DFHSM task is currently running.

HSM Start Time: The start time and date for the HSM address space.

HSM Status: The current status of the HSM product. Status can be one of the following:

Inactive = The product is licensed, but not currently running.

Active = The product is licensed, installed, and ready to run.

Notinstalled = The product is not licensed or not defined as part of the system.

Inactive HSM Hosts: The number of inactive HSM hosts on this z/OS image.

Interval Abends: The count of HSM requests that abended during this interval. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Interval CPU Usage: The CPU time consumed by HSM to service requests during this interval. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Interval Requests: The total number of HSM requests during this interval. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Interval Start Time: The start time for the current HSM interval.

Managed System: The SMS name of this system.

Recycled Back Up Volumes: The total number of BACKUP volumes recycled during this interval. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Recycled Migration Volumes: The total number of MIGRATION volumes recycled during this interval. You can enable aggregation for this attribute in Tivoli Data Warehouse.

System Requested Dump Copies: The number of system requested fast volume dump copies taken by all HSM hosts on this z/OS image.

System Requested Dumps: The number of system requested dumps across all HSM hosts on the z/OS image.

Total Aggregate Backups: The number of aggregate backups taken by all HSM hosts on this z/OS image.

Total Data Set Restores: The number of data set restores taken by all HSM hosts on this z/OS image.

Total Fast Replic Recovers: The number of fast replicate recovers taken by all HSM hosts on this z/OS image.

Total Fast Replicat Backups : The number of fast replicate backups taken by all HSM hosts on this z/OS image.

Total Volume Dump Copies: The number of fast volume dump copies taken by all HSM hosts on this z/OS image.

Total Volume Dumps: Total volume dumps.

User Requested Dump Copies: The number of user requested fast volume dump copies taken by all HSM hosts on this z/OS image.

User Requested Dumps: The number of user requested dumps across all HSM hosts on the z/OS image.

Version: The version and release level of DFHSM on this system.

Volume Backups: The total number of requests for volume BACKUP during this interval. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Volume Migrations: The total number of requests for volume MIGRATIONS during this interval. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Volume Restores: Volume restores.

Logical Control Unit

The Logical Control Unit attributes provide performance information for a logical control unit.

All Channels Busy Percent: The percentage of the time all channel paths to this LCU were busy.

Attached Channels: The total number of channel paths that have defined connections to the LCU.

Average Delay Queue: The average number of delayed requests on the control unit header (CU-HDR). This field is calculated as the accumulated Q length minus the number of enqueued requests divided by the number of enqueued requests.

Channel Path I/O Rate: The Channel Path I/O rate for I/Os to this LCU. This field is calculated as the number of I/O operations accepted on this channel path divided by (sample time minus interval start time).

Contention Rate: The rate per second at which the SAP places delayed I/O requests on the CU-HDR for this LCU. This is done when all paths to the subchannel are busy and at least one path to the control unit is busy.

Control Unit Busy Percent: The percentage of requests deferred due to control unit busy.

DCM Channels Attached: The flag indicating that there are DCM channels attached to this LCU.

DCM Channels Channel Path I/O Rate: The Channel Path I/O rate for I/Os to this LCU. This field is computed as the number of I/O operations accepted on this channel path divided by (sample time minus interval start time).

DCM Channels Control Unit Busy Percent: Shows the relationship for each channel path of the LCU, between requests deferred due to control unit busy and total successful requests serviced by that path. Each CHPID of the LCU measures the distribution of control unit contention. This number is derived from: the number of deferred I/O requests due to **control unit busy** divided by the following combined values:

(the number of deferred I/O requests due to director port busy +
the number of deferred I/O requests due to control unit busy +
the number of successful I/O requests on that path) * 100

DCM Channels Defined: The number of DCM channels defined for this LCU.

DCM Channels Director Port Busy Percent: Indicates director port contention. This value represents the number of times an I/O request was deferred because the director port was busy during the measurement interval. This value is derived from the number of deferred I/O requests due to **director port busy** divided by the following combined values:

(the number of deferred I/O requests due to director port busy +

the number of deferred I/O requests due to control unit busy +
the number of successful I/O requests on that path) * 100

DCM Channels Maximum: The maximum number of DCM channels used for this LCU.

DCM Channels Minimum: The minimum number of DCM channels used for this LCU.

Director Port Busy Percent: The percentage the director port is busy. This field indicates director port contention. This value is the number of times an I/O request was deferred because the director port was busy during the measurement interval.

High DFW Retry Percent: The highest DFW retry percentage of any volume on this LCU.

High Volume MSR: The highest MSR (millisecond response time) of any volume on this LCU.

High Volume Utilization: The highest percentage busy of any volume on this LCU.

I/O Per Second: The channel Path I/O Rate for I/Os to this.

LCU Hex Number: The character representation to the hexadecimal number of this LCU.

LCU Number: The HCD-assigned number for this logical control unit. An LCU is a logical representation of either a single control unit, with or without attached devices, or a group of control units that share devices.

Low Read Hit Percent: The lowest read hit percentage of any volume on this LCU.

Low Write Hit Percent: The lowest write hit percentage of any volume on this LCU.

Managed System: The SMS name of this system.

Maximum HyperPAV Unavailable Alias Percent: The maximum percent of the number of I/Os that could not start for a control unit because a HyperPAV alias was not available.

Maximum I/Os Queued: The maximum number of I/Os queued for a control unit.

NonSMS Volumes: The total number of nonSMS-managed volumes on this LCU.

PDS Record: A single record of trend data retrieved from persistent data store.

SMS Volumes: The total number of SMS-managed volumes on this LCU.

Time Stamp: The date and time associated with this row.

Total Volumes: The count of the total number of volumes in this LCU.

Logical Control Unit Channel Paths

The Logical Control Unit Channel Paths attributes provide performance information for the channels associated with a logical control unit.

Channel Path IO Rate: The number of I/O operations per second accepted on this channel path.

CHPID: The channel path identifier. A channel path is the physical connection from the processor to I/O or peripheral devices. Some devices can operate with up to eight paths.

Complex Percent Utilized: The aggregate utilization of this CHPID by all processors in the complex.

Control Units Percent Busy: The control unit utilization is expressed as the percentage of time that all control units attached to the channel path are busy for I/O operations to this LCU.

Control Units: The ID of the physical control units that make up the logical control unit.

DCM Status: The dynamic Channel Path Management Status. Managed or unmanaged, dynamic channel path management provides the capability to dynamically assign channels to control units to respond to peaks in demand.

Description: The type of channel path identified by this CHPID.

Director Port Percent Busy: The percentage of initial selection attempts that were unsuccessful because the director port was busy.

LCU Number H: The system-assigned identifier of this logical control unit. An LCU is the logical representation of either a single control unit, with or without attached devices, or a group of control units that share devices.

LCU Number: The system-assigned identifier of this logical control unit. An LCU is the logical representation of either a single control unit, with or without attached devices, or a group of control units that share devices.

Managed System: The SMS name of this system.

Path ID: The Channel Path Identifier. A channel path is the physical connection from the processor to I/O or peripheral devices. Some devices can operate with up to eight paths.

Status: The status of this channel path on this LPAR.

LCU Control Unit

This group of attributes monitors I/O requests for a control unit that are related to the HyperPAV (Hyper parallel access volume) feature.

Control Unit: The physical control unit number.

HyperPAV Requests: The total number of HyperPAV I/O requests for the control unit.

I/Os Queued Alias Unavailable: The number of times an I/O could not start for the control unit because no HyperPAV aliases were available.

LCU Number: The HCD-assigned number for this logical control unit. An LCU is a logical representation of either a single control unit, with or without attached devices, or a group of control units that share devices.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Maximum HyperPAV Aliases Used: The maximum number of HyperPAV aliases in use concurrently.

Maximum HyperPAV Volume Aliases: The maximum number of HyperPAV aliases in use concurrently by a single logical volume.

Maximum I/Os Queued: The maximum number of I/Os concurrently queued for the control unit.

Percent I/Os Queued Alias Unavailable: The percent of I/Os that could not start for the control unit because no HyperPAV aliases were available.

Time Stamp: Date and time this data was reported. Primarily for use in historical reporting.

RMM Configuration

This table contains RMM configuration state and option attributes.

Accounting Source: Accounting is from the STEP or the JOB. (Job=0; Step=1)

Backup In Progress: Indicates whether backup is in progress. (No=0; Yes=1)

Backup Procedure: Specifies the name of the backup procedure.

BLP: Indicates whether RMM mode is active for Bypass Label Processing (BLP) mount requests. (RMM=0; NoRMM=1)

Cat Sync In Progress: Indicates whether catalog synchronization is in progress. (No=0; Yes=1)

Catalog Retention Period: Indicates the number of hours assigned as the catalog retention period.

Catalog SysID: Indicates whether the catalog SysID is set. (Set=0; NotSet=1)

CDS ID: Identifies the ID of the control data set.

Common Time: Indicates whether the Common Time feature is enabled. (Disabled=0; Enabled=1)

Date Format: Specifies the date format. (American=1; European=2; ISO=3; Julian=4)

Days Since Last Low on Scratch Processing: Specifies the days since the last low on the processing of scratch space.

Days Since Last Expiration Processing: Specifies the days since last expiration processing event.

Days Since Last Storage Location Proc: Specifies the days since last storage location processing event.

Days Since Last VRS Update: Specifies the days since the last VRS update.

Days Since CDS Extract: Specifies the days since a CDS extract.

Days Since Last Cat Sync: Specifies the days since last catalog synchronization.

Days Since Last Update: Specifies the days since the last update.

Default Retention Period: Specifies the number of days in the default retention period.

Distant Bins Free: Specifies the number of distant bins that are free.

Distant Bins Total: Specifies the total number of distant bins.

DSN Cmd Auth: Indicates whether DSN command authorization is enabled. (No=0; Yes=1)

EDGUX100 Status: Indicates the enablement status of the EDGUX100 installation exit. (None=0; Enabled=1; Disabled=2)

EDGUX200 Status: Indicates the enablement status of the EDGUX200 installation exit. (None=0; Enabled=1; Disabled=2)

Expiration In Progress: Indicates whether expiration is in progress. (No=0; Yes=1)

Extended Bin: Indicates the enablement status of extended bin support. (Disabled=0; Enabled=1)

Extract In Progress: Indicates whether an extract operation is in progress. (No=0; Yes=1)

IPL Date Check: Indicates the enablement status of the IPL date check. (No=0; Yes=1)

Journal Full Value: Specifies the value that defines the journal full condition.

Journal Status: Indicates the enablement status of the journal. (Disabled=0; Enabled=1; Locked=2)

Last Catalog Sync: Date and time of the last catalog synchronization.

Last CDS Extract: Specifies the last extraction performed on the CDS.

Last Expiration Processing: Days since last VRS update.

Last Low on Scratch Processing: Last low on the processing of scratch space.

Last Storage Location Processing: Last storage location processing.

Last Update Date: Specifies the date of the last update.

Last VRS Update: Specifies the last update to the vital record specification.

Library Racks Free: Specifies the number of free Library racks.

Library Racks Total: Specifies the total number of library racks.

Local Bins Free: Specifies the number of free local bins.

Local Bins Total: Specifies the total number of local bins.

Local Tasks: Specifies the number of local tasks.

Location In Progress: Indicates whether the Location In Progress option is enabled. (No=0; Yes=1)

Managed System: Specifies the name of the managed system.

Master Overwrite: Specifies the Master Overwrite setting that is in effect. (Add=1; Last=2; Match=3; User=4)

Max Hold: Specifies the Max Hold value that is in effect.

Maximum Retention Period: Maximum retention period in days. (Unlimited=-1)

Media Name: Specifies the name assigned to the media.

Move By: Move VRS volumes individual or by multi-volume set. (Volume=0; Set=1)

Operating Mode: Specifies the operating mode. (Manual=1; Record=2; Warning=3; Protect=4)

Owner Cmd Auth: Indicates whether owner command authorization is enabled. (No=0; Yes=1)

Parmlib Suffix: Specifies the suffix of the parameter library.

Pre ACS: Specifies whether Pre-ACS is enabled. (No=0; Yes=1)

RACF: Specifies the RACF option that is in use. (None=1; Predefined=2; Automatic=3; Cleanup=4)

Remote Bins Free: Specifies the number of free remote bins.

Remote Bins Total: Specifies the total number of remote bins.

Restore In Progress: Indicates whether a restore is in progress. (No=0; Yes=1)

Retain By: Specifies the retention policy for volumes. (Volume=0; Set=1)

Reuse Bin: Specifies the type of reuse allowed for a bin. (ConfirmMove=0; StartMove=1)

RMM System ID: Specifies the RMM system ID.

Scratch Procedure: Specifies the name of the scratch procedure.

SMF Audits Rec Num: Specifies the SMF audits rec num.

SMF Security Rec Num: Specifies the SMF security rec num.

SMSACS: Indicates whether the SMSACS option is enabled. (No=0; Yes=1)

SMSTape Command: Indicates whether the Command subparameter is enabled for the SMSTape operand. (No=0; Yes=1)

SMSTape Exits: Indicates whether the Exits subparameter is enabled for the SMSTape operand. (No=0; Yes=1)

SMSTape Purge: Indicates whether the Purge parameter is enabled for the SMSTape operand. (No=0; Yes=1; Asis=2)

SMSTape Scratch: Indicates whether the Scratch subparameter is enabled for the SMSTape operand. (No=0; Yes=1)

Stacked Volume Enable Status: Indicates the enablement status for stacked volumes. (None=0; Enabled=1; Disabled=2; Mixed=3)

Subsystem ID : Specifies the RMM subsystem ID.

Subsystem Status: Indicates the status of the RMM subsystem. (Inactive=0; Active=1)

Subsystem Type: Identifies the subsystem type. (Standard=0; Client=1; Server=2)

Tapevol Exit Purge: Specifies the Tapevol Exit Purge setting (Release=0; Expire=1; None=2).

Timestamp: Provides a timestamp.

UnCatalog: Indicates whether the UnCatalog specifier is enabled. (No=0; Yes=1; Scratch=2)

User Notification: Indicates whether user notification is enabled. (No=0; Yes=1)

Verify In Progress: Indicates that verification is in progress. (No=0; Yes=1)

VRS Change: Indicates the response to changes to the vital record specification. (Info=0; Verify=1)

VRS Name Precedence: Indicates name precedence for the vital record specification. (Job=1; Dataset=2)

VRS In Progress: Indicates whether vital record specification is in progress. (No=0; Yes=1)

VRSEL: Indicates the type of vital record processing. (Old=0; New=1)

VRSMIN Count: Specifies the minimum number of vital record specifications.

VRSMIN Action: Specifies the action that occurs when the minimum number of vital record specifications is reached. (Fail=0; Info=1; Warn=2)

RMM Control

This table contains basic RMM Control data.

Message: Information regarding the data that was obtained.

Controls: DFSMSrmm controls.

Timestamp: Time table row was created in CT timestamp format.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

RMM Control Data Set

This table contains RMM Control Data Set and Journal attributes and statistics.

Actual Percent Used : The percent of the data set that is used with reference to tracks. For control data sets, this percentage value is calculated by comparing the values of High Allocated and High Used RBA. For Journals, this percentage value is calculated by comparing Allocated Tracks and Last Track Used.

CA Splits: Number of component CA splits.

Catalog Entry Type: Catalog Entry Type for the data set. (Cluster=C Data_Component=D Index_Component=I NonVSAM=A)

CI Splits: Number of component CI Splits.

Creation Date: Dataset Creation Date.

Dataset Type: Type of CDS Dataset. (Master_CDS=0; Satellite_CDS=1; Journal=9)

Dataset Name: Name of the data set.

Days Since Creation: Number of days since the creation of the data set.

Days Since Last Backup: Number of days since the last data set backup.

Extents: Number of extents allocated to the data set.

GB Allocated: Number of gigabytes allocated.

GB Free: Number of gigabytes free.

GB Used: Number of gigabytes used.

Last Backup: Date of the last backup as reported by RMM (for the CDS) or as recorded in the catalog entry (for the Journal).

Managed System: Managed System.

MB Allocated: Number of megabytes allocated.

MB Free: Number of megabytes free.

MB Used: Number of megabytes used.

RMM Percent Used: Dataset percent used as reported by RMM.

Timestamp: Timestamp.

Tracks Allocated: Number of Tracks allocated to the data set.

Tracks Free: Number of tracks free.

Tracks Used: Number of tracks used by the data set.

Volser: Volume Serial on which data set is located.

RMM Dataset Details

This table contains attributes for a single RMM data set.

Abend While Open: Dataset was open at time of a system or application abend. (No=0; Yes=1)

Block Count: Number of blocks used by the data set on this volume.

Blocksize: The Maximum Dataset blocksize. Blocksize*Number of Blocks = Approximate data set size.

Cataloged: Data set is cataloged. (Unknown=0; No=1 ;Yes=2)

Creating DDname: DDname that created the data set.

Creating Device: Name of device on which the data set was created.

Creating Job: Name of job that created the data set.

Creating Program: Name of program that created the data set.

Creating Step: Name of step that created the data set.

Creation Date: Date on which the data set was created.

Data Class: DFSMS Data Class for the data set.

Dataset Name: Name of the data set.

Dataset Sequence: Relative position of the data set within a multiple data set group.

Date Last Read: Date when the data set was last accessed for Input.

Date Last Written: Date when the data set was last accessed for Output.

Days for Retention: Number of days until or since (negative number) the data set is no longer retained by the current VRS. Nonexpiring=-2147483648; WHILECATLG=-286331154

Days Since Creation: Number of days since the data set was created.

Days Since Last Read: Number of days since the data set was last accessed for input.

Days Since Last Write: Number of days since the data set was last accessed for output.

Days Since Original Expiration: Days Since/Until the Original Expiration date. Nonexpiring=-2147483648

Days Until Expiration: Number of days until or since (negative number) the data set expires. Nonexpiring=-2147483648

Expiration Date: Date on which the data set expires.

Last DDname: DDname that last accessed the data set.

Last Device: Number of the device that last accessed the data set.

Last Job: Name of the Job that last accessed the data set.

Last Program: Name of the Program that last accessed the data set.

Last Step: Name of the Step that last accessed the data set.

LRECL: Logical Record Length (LRECL).

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Management Class: DFSMS Management Class for the data set.

Original Expiration Date: The Expiration Date contained in the label when the data set was created.

Owner: ID of the owner of the volume where the data set is located.

Percent of Volume: How much of the physical tape media is occupied by the data set. **Note:** The accuracy of this number is based on information reported to the host at the end of a tape volume file. The number can only be provided if the hardware supports the provision of positioning information, such as with IBM 3590 Magstar® devices.

Physical Sequence: Position of the data set on the volume relative to other data sets.

Primary VRS Date: Primary VRS Date.

Primary VRS Jobname: Job name mask of the primary VRS for this data set.

Primary VRS Name: Data set name mask of the primary VRS for this data set.

Primary VRS Subchain Name: Name of subchain VRS.

Primary VRS Type: Type of the primary VRS selected for the data set.

- Undefined=0 Dataset=1 SMSMC=2 VRSMV=3 DSNMV=4 DSNMC=5

Record Format: Record Format (RECFM).

Retention Date: Date after which the data set is no longer retained by the current VRS.

Secondary VRS Date: Date on which the secondary VRS became active for the data set.

Secondary VRS Jobname: Job name mask of the secondary VRS for this data set.

Secondary VRS Mask: Mask that matches the VRS management value of SMS management class.

Secondary VRS Subchain Name: Name of the subchain VRS.

Security Class: Security Class name.

Storage Class: DFSMS Storage Class for the data set.

Storage Group: DFSMS Storage Group for the data set.

System ID: SMF ID of the system on which the data set was created.

Timestamp: Time table row was created in CT timestamp format.

Total Blocks: Number of blocks used by the data set on all volumes. This value includes all volumes up to and including the current volume.

Volser: Serial number of the volume on which the data set is located.

VRS Retained: Dataset is retained by a Vital Record Specification. (No=0; Yes=1)

RMM Location Details

This table contains attributes for a single RMM location.

Location: Location name.

Location Def Exists: Indication if the Location is defined with a LOCDEF command in the RMM PARMLIB. (No=0; Yes=1)

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Management Type: The shelf management technique in use for the storage location. (Unknown=0; Bin=1; NoBins=2)

Media: Media name for the location.

Priority: Priority of this location relative to other locations.

Timestamp: Time table row was created in CT timestamp format.

Type: Type of Location. (Shelf=0; Auto=1; Manual=2; Store=3)

RMM Owner Details

This table contains attributes for a single RMM owner.

Address1: Address Line 1.

Address2: Address Line 2.

Address3: Address Line 3.

Department: Department.

Email: Owner's email.

External Phone: External Phone number.

First Name: Owner First Name.

Internal Phone: Internal Phone number.

Last Name: Owner Last Name.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Name: Owner Name.

Number of Volumes: Total number of volumes owned.

Timestamp: Time table row was created in CT timestamp format.

User node: User node.

Userid: Owner's Userid.

RMM Pending Actions

This table contains attributes for a single RMM action.

Action: RMM Action Name. Values are as follows: (SCRATCH REPPLACE INIT ERASE RETURN NOTIFY)

Current Location: For Moves, the Current Location of the volumes.

Destination Location: For Moves, the Destination Location for the volumes.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Move Type: Return To Scratch move. (NoRTS=0; RTS=1)

Number of Volumes: Number of Volumes in specified action state.

Status: Status of the action. (Unknown=5; Complete=6; Confirmed=7; Pending=8)

Timestamp: Time table row was created in CT timestamp format.

RMM System Summary

This table contains attributes that summarize RMM at various levels.

% Volumes R/W Previous Day: Percent of total volumes read or written on the previous day.

Average Datasets Per Volume: Average number of data sets per volume.

Average Volume Use Count: Average volume use count for all volumes associated with this summarized object.

Average Compression Ratio: The average compression ratio for all volumes associated with this summarized object.

Average compression ratio measures the effectiveness of data compression and is expressed as the percentage of the original data written by the application that was actually recorded on tape. This metric is dependent on the capacity and percent used values being reported back to RMM by the hardware, which might not be the case for older tape devices.

The compression ratio is based on the average % used for all volumes, the total capacity of the volumes, and the total usage of all the volumes and is calculated as follows:

- Total usage is the $(\text{data-set blocksize}) * (\text{blockcount})$, which represents the amount of data written by the application for all data sets on the volume.
- Percent used (% used) is based on data actually written to the tape (assuming that the data is reported back by the hardware).
- Average capacity used is provided by this formula:
 $(\text{avg \% used for all vols}) * \text{total capacity}$
- Average compression ratio is provided by this formula:
 $100 - (\text{avg capacity used} / \text{total usage})$

Average Dataset Size: Average data set size, expressed in megabytes to 1 decimal place on all volumes associated with this summarized object.

Average Volume Utilization: Average volume utilization. Volume utilization (in other words, percent full) is based on where the last file on the tape ended. Only physical (not stacked or logical) volumes that are in Master or User status are included in the utilization calculation and distributions. Note that not all devices report this information back to RMM, so the value might be 0 for volumes that do in fact contain data.

Collection Status: Status of RMM Data collection process.

- Collection_completed=RMM_COLLECTION_COMPLETE
- Collection_is_waiting=RMM_COLLECTION_WAITING
- Collection_not_started=RMM_COLLECTION_INACTIVE
- Collection_parameter_errors=RMM_COLLECTION_PARAM_ERROR
- Collection_in_progress=RMM_COLLECTION_IN_PROGRESS
- Collection_halted_by_user=RMM_COLLECTION_HALTED
- Collection_halted_due_to_deadline=RMM_COLLECTION_DEADLINE
- Collection_terminated=RMM_COLLECTION_TERMINATED
- Collection_failed=RMM_COLLECTION_FAILED
- RMM_is_not_active=RMM_INACTIVE
- RMM_is_not_installed=RMM_NOT_INSTALLED

Data Created Previous Day: Data created on all volumes associated with this summarized object on the previous Julian day, expressed as MB to a single decimal place.

Datasets Created Previous Day: Total number of data sets created on the previous Julian day on all volumes associated with this summarized object.

Datasets R/W Previous Day: The number of data sets last read or written on the previous Julian day on all volumes associated with this summarized object (at the time that the last summary collection cycle ended).

Entry Volumes: Number of volumes in Entry status.

Erase Pending Volumes: Number of volumes in Erase Pending status.

Init Pending Volumes: Number of volumes in Init Pending status.

Init Volumes: Number of volumes in Init status.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Master Volumes: Number of volumes in Master status.

Name: Name of the object being summarized.

New Scratch Volumes: Number of volumes returned to scratch status during or since the last run of RMM expiration processing, as reported by the LISTCONTROL command.

Notify Pending Volumes: Number of volumes in Notify Pending status.

Permanent Read Errors Volumes: Total number of volumes associated with this summarized object with permanent read errors.

Permanent Write Errors Volumes: Total number of volumes associated with this summarized object with permanent write errors.

Permanent Read Errors: Total number of permanent read errors for all volumes associated with this summarized object.

Permanent Write Errors: Total number of permanent write errors for all volumes associated with this summarized object.

Physical Volumes: Number of physical (in other words, not LOGICAL or STACKED) volumes.

Replace Pending Volumes: Number of volumes in Replace Pending status.

Return Pending Volumes: Number of volumes in Return Pending status.

Scratch Pending Volumes: Number of volumes in Scratch Pending status.

Scratch Volumes: Number of volumes in Scratch status.

Summary Collection End Time: Time at which the last summary collection cycle ended.

Summary Collection Start Time: Time at which the last summary collection cycle started.

Temporary Read Errors Volumes: Total number of volumes associated with this summarized object with temporary read errors.

Temporary Write Errors Volumes: Total number of volumes associated with this summarized object with temporary write errors.

Temporary Read Errors : Total number of temporary read errors for all volumes associated with this summarized object.

Temporary Write Errors: Total number of Temporary write errors for all volumes associated with this summarized object.

Timestamp: Time table row was created in CT timestamp format.

Total Error Volumes: Total number of volumes associated with this summarized object with permanent or temporary read or write errors.

Total Errors: Total number of temporary or permanent read or write errors for all volumes associated with this summarized object.

Total Permanent Errors Volumes: Total number of volumes associated with this summarized object with permanent read or write errors.

Total Temporary Errors Volumes: Total number of volumes associated with this summarized object with temporary read or write errors.

Total Usage: Total Usage in GB. Usage is the amount of data written to the tape by the application, before any compression, expressed as gigabytes.

Total Capacity: Total Capacity of all volumes associated with this summarized object, in Gigabytes. Note that some tape devices do not report capacity info to RMM and this field might therefore be zero (0).

Total Datasets: Total number of data sets on all volumes associated with this summarized object.

Total Permanent Errors: Total number of permanent read or write errors for all volumes associated with this summarized object.

Total Temporary Errors: Total number of temporary read or write errors for all volumes associated with this summarized object.

Total Volumes: Total number of volumes in RMM for this object.

Total Write Mounts: Total number of Write Mounts for all volumes associated with this summarized object.

Type: Type of the object being summarized.

User Volumes: Number of volumes in User status.

Volume 0% Utilized: Number of Volumes 0% Utilized.

Volume 1-10% Utilized: Volumes 1-10% Utilized.

Volume 11-20% Utilized: Volumes 11-20% Utilized.

Volume 21-30% Utilized: Volumes 21-30% Utilized.

Volume 31-40% Utilized: Volumes 31-40% Utilized.

Volume 41-50% Utilized: Volumes 41-50% Utilized.

Volume 51-60% Utilized: Volumes 51-60% Utilized.

Volume 61-70% Utilized: Volumes 61-70% Utilized.

Volume 71-80% Utilized: Volumes 71-80% Utilized.

Volume 81-90% Utilized: Volumes 81-90% Utilized.

Volume 91-100% Utilized: Volumes 91-100% Utilized.

Volumes R/W Previous Day: The number of volumes associated with the summarized object that were last read or written on the previous Julian day (at the time that the last summary collection cycle ended).

RMM Volume Details

This table contains attributes for a single RMM volume.

Account Number: User accounting information.

Assigned Date: For a MASTER or USER volume the Assigned date field specifies the date the volume was assigned to an owner. That is the date when the volume was first added or when a SCRATCH volume became a MASTER volume as a result of a nonspecific mount request.

Availability: The Availability status of the volume. Valid values are as follows:

- Pending Release - The volume is not available.
- On Loan - The volume is not available because it is on loan, possibly at a different site.
- Open - The volume has been opened for a write operation and has not yet been closed.
- Vital Record - The volume is a vital record.
- Unknown - The availability status of the volume is not known.

Bin: The shelf location where the volume is currently stored.

Bin Media: The media type assigned to the shelf location where the volume is currently stored.

Container: Container where the volume is currently stored.

Creating Job Name: Name of the job that created the volume.

Creating System ID: ID of the system where the volume was defined to RMM.

Creation Date : Creation date of the RMM object.

Current Location: Location where the volume is currently stored.

Current Location Type: The type of location where the volume is currently stored.

- Shelf=0
- Store_Builtin=1
- Manual=2
- Auto=3

- Store_Bins=4
- Store_NoBins=5

Data Set Name: Name of the data set on the volume.

Dataset Record: RMM is recording details for all data sets on the volume (YES) or just the first data set (NO). (No=0; Yes=1)

Dataset Count: Number of data sets on the volume, or sequence number for a specific data set.

Date Last Read: Date the volume was last read from.

Date Last Written: Date the volume was last written to.

Days for Retention: Number of Days Until or since (negative number) the Retention date.

- Nonexpiring=-2147483648
- WHILECATLG=-286331154

Days Since Assigned: Number of Days Since Assigned date.

Days Since Creation: Number of days since the volume was created.

Days Since Last Read: Number of days since the volume was last read from.

Days Since Last Written: Number of days since the volume was last written to.

Days Since Movement Tracking: Number of days since the movement tracking date.

Days Until Expiration: Number of Days Until or since (negative number) the expiration date.
NonExpiring=-2147483648

Days Until Original Expiration Date: Number of Days Until or since (negative number) Original Expiration date. Nonexpiring=-2147483648

density: The recording density of the volume.

- Undefined=0
- 1600=1
- 6250=2
- 3480=3
- Compact=4

Description: Descriptive text about the volume.

Destination Type: The type of location where the moving volume will be stored.

- Shelf=0
- Store_builtin=1
- Manual=2
- Auto=3
- Store_Bins=4
- Store_NoBins=5

Destination Bin Media: The media type assigned to the shelf location where the moving volume will be stored.

Destination Bin Number: The shelf location where the moving volume will be stored.

Destination Location: The name of the target location where the moving volume will be stored.

Erase on Release: Erase and initialize the volume before performing the release action. (No=0; Yes=1)

Erase Pending: Volume will be erased upon release. (No=0; Yes=1)

Expiration Date: The date on which the volume will expire and its status set to Release Pending.

Expiration Date Ignore: The expiration date of the volume is set to the current date when removed from VRS control. (No=0; Yes=1)

Home Location: The library location for the volume.

Home Location Type: The type of library location for the volume.

- Shelf=0
- Store_builtIn=1
- Manual=2
- Auto=3
- Store_Bins=4
- Store_NoBins=5

In Transit: Volume is in transit to its destination location. (No=0; Yes=1)

Init on Release: Initialize the volume before performing the release action. (No=0; Yes=1)

Init Pending: Volume will be initialized upon release. (No=0; Yes=1)

Label Type: The label type of the volume. Values are as follows:

- AL - ANSI Label
- AUL - ASCII User Label
- BLP - Bypass Label Processing
- NL - No Label
- SL - Standard Label
- SUL - Standard User Label

Label Version: Version for the current label type.

- NoVersion=0
- Label_Version_1=1
- Label_Version_3=3
- Label_Version_4=4

Last Drive: Device address where the volume was last read or written.

Last Userid: Last Userid. The Last Changed by field specifies the ID of the user who last changed the details for this volume. If the volume was not last changed by a user, this field can contain one of these values:

- *OAM - volume last updated by OAM.

- *HKP - volume last updated by inventory management.
- *OCE - volume last updated by automatic volume recording when read or written to

Loan Location: Location of the volume if it is on loan.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Media Attributes: Media Attributes of the volume. Values are as follows:

- NONE - The volume has no special attributes.
- RDCOMPAT - This indicates that the volume was created using one recording format but might be mounted on a drive that supports that format for input processing, but another for output processing.

Media Compaction: A Compaction technique was used to record data to the volume. (Unknown=0; No=1; Yes=2)

Media Name: The Name of the volume's media.

Media Recording: The recording format of the volume.

- Non-Cartridge=0
- 18Trk=1
- 36Trk=2
- 128Trk=3
- 256Trk=4
- 384Trk=5
- EFMT1=6
- EFMT2=7

Media Type: The physical media type of the volume. Values are as follows:

- CST - The volume is a 3480 or 3490 cartridge.
- ECCST - Enhanced Capacity 3480 or 3490 cartridge.
- HPCT - High Performance Cartridge Tape.
- EHPCT - Extended High Performance Cartridge Tape
- ETC - IBM TotalStorage Enterprise Tape Cartridge.
- EWTC - IBM TotalStorage WORM Enterprise Tape Cartridge.
- EETC - IBM TotalStorage Enterprise Economy Tape Cartridge.
- EEWTC - IBM TotalStorage Enterprise Economy WORM Tape Cartridge.
- EXTTC - IBM TotalStorage Enterprise Extended Tape Cartridge.
- EXWTC - IBM TotalStorage Enterprise Extended WORM Tape Cartridge.

Move Mode: How the volume moves between locations. Auto=0; Manual=1

Movement Tracking Date: The date when the volume started to move between two locations or when the confirmation that the last movement had occurred was performed.

MVS Use: Volume can be used with the MVS operating system. (No=0; Yes=1)

Next Volume: Next Volume Serial number for multi-volume data sets.

Notify on Release: Notify the volume owner that the volume has expired. (No=0; Yes=1)

Notify Pending: Notify the owner of the volume that the volume will be released. (No=0; Yes=1)

OCE: Volume information was recorded at the last Open, Close, or EOVS event. (No=0; Yes=1)

Old Bin Media: The media type assigned to the shelf location where the volume was previously stored.

Old Bin Number: The shelf location where the volume was previously stored.

Old Location: The name of the location where the volume was previously stored.

Old Location Type: The type of location where the volume was previously stored.

- Shelf=0
- Store_builtin=1
- Manual=2
- Auto=3
- Store_Bins=4
- Store_NoBins=5

Original Expiration Date: Expiration Date in the volume label, which might have been changed in RMM using the CHANGEVOLUME command.

Owner: ID of the owner.

Owner Access: The level of access the owner has to the volume.

- Read=0
- Update=1
- Alter=2

Permanent Read Errors: Total number of permanent read errors on the volume.

Permanent Write Errors: Total number of permanent write errors on the volume.

Previous Volume: Previous Volume Serial number for multi-volume data sets.

Product Number: Number of the Product located on the volume.

Rack Number: Rack/shelf location number where the volume is located.

Release Actions: Action to take when RMM releases the volume.

- Unknown=0
- Scratch=8
- Replace=7
- Return=4

Replace Pending: Volume is to be replaced (for example, due to excessive I/O errors). (No=0; Yes=1)

Required Location: The intended target location for the volume.

Required Location Type: The intended target location type for the volume.

- Shelf=0
- Store_builtin=1
- Manual=2

- Auto=3
- Store_Bins=4
- Store_NoBins=5

Retain By Set: The volume is being retained because it is part of a multi-volume set. (No=0; Yes=1)

Retention Date: Highest date for all data sets on the volume after which the volume is no longer retained by the current VRS.

Return Pending: Volume will be returned to its owner upon release. (No=0; Yes=1)

Scratch Immediate: Return volume to scratch in a single inventory management run. (No=0; Yes=1)

Scratch Pending: Volume will return to scratch status after all other release actions are completed. (No=0; Yes=1)

Security Class: The Security Class of the volume.

Stacked Volume Count: Number of volumes contained in a stacked volume.

Storage Group: DFSMS Storage Group for the volume.

Temporary Read Errors: Total number of temporary read errors on the volume.

Temporary Write Errors: Total number of temporary write errors on the volume.

Timestamp: Time table row was created in CT timestamp format.

Use Count: Number of times the volume has been opened for either read or write.

Vendor: Manufacturer or supplier of the media.

VM Use: Volume can be used with the VM operating system. (No=0; Yes=1)

VOL1 Label: Serial number of the volume as recorded in the standard internal volume label.

Volume Capacity: The physical capacity of the tape volume in megabytes as reported by the hardware. Note: The value is displayed for those devices which provide capacity information, such as the IBM 3590 Magstar and later devices.

Volume Serial: Volume Serial number.

Volume Status: The Status of the volume. Values are as follows:

- MASTER - The volume contains valid user data and cannot be overwritten.
- USER - The volume is assigned to a user and is in use.
- SCRATCH - The volume is free and available for use.
- INIT - The volume is a scratch volume that is waiting to be initialized. After the volume is initialized, the status becomes SCRATCH.
- ENTRY - The volume is a scratch volume that is moving into an automatic library; once entered, the status becomes SCRATCH or INIT.

Volume Type: The type of the volume as defined to RMM.

- Physical=0
- Logical=1

- Stacked=2

Volume Usage: How much of the volume has been taken up by the data sets that are located on it, in kilobytes. This value is calculated as the sum of the usage of all the data sets on the volume. The usage for a data set is the product of the block count and the block size. Volume Usage is an indication of how much data has been written by an application prior to compression or compaction.

Volume Write Mount Count: Number of times the volume has been mounted for output and written to.

Volume Percent Full: How much of the volume has been taken up by the data written to it. The value is calculated using the reported position on the volume where the end of the last file was written.

Note: This value is only available for those devices which provide positioning information, such as the IBM 3590 Magstar and later devices. When data is written to a volume it is most probably compressed so that the volume usage, capacity and percent full values are only an indication of the remaining, available capacity.

Volume Sequence: Sequence number of the volume in a multi-volume set.

Worldwide ID: ID value set in the volume by the manufacturer and recorded by DFSMSrmm when the volume is first used. DFSMSrmm ensures that the correct volume is mounted by checking that the WWID still matches the recorded value.

RMM VRS Details

This table contains details for a Vital Record Specification (VRS) for RMM (removable media manager).

Count: Specifies how many days or how many cycles of a data set should be retained in the location specified in the Location field.

Delay Days: Specifies the number of days that a data set is retained in its current location before it is sent to a storage location.

Delete Date: Date when the VRS is deleted.

Description: Description string.

Expiration Date Ignore: Whether or not to honor the volume expiration date. (No=0; Yes=1)

Job Name: Specifies a job name mask used to select only those data sets which match the VRS and which were created by certain jobs.

Location: Where vital records managed by this VRS are retained.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Name: Name of the VRS. Depending on the VRS type could be a data set name, GDG base name, VOLSER or VRS name.

Next VRS: Name of the subsequent VRS in the chain.

Next VRS Type: Type of the subsequent VRS in the chain.

- Undefined=0
- Next=1

- And=2

Owner: Owner of the VRS.

Priority: Specifies a value to override the priority set for the location on a LOCDEF command.

Retain Until Expired: VRS will retain data sets until their expiration dates. (No=0; Yes=1)

Retain While Cataloged: Retain to the specified count value while the data set is cataloged. (No=0; Yes=1)

Retention Type: How data sets are retained by this VRS.

- Cycles=1
- Days=2
- RefDays=3
- Volumes=4
- ExtraDays=5
- ByDaysCycle=6

Scratch Immediate: Pending release volume goes to scratch after expiration processing. (No=0; Yes=1)

Timestamp: Time table row was created in CT timestamp format.

Type: The type of this VRS.

- Name=4
- Volume=5
- DSName=6
- PseudoGDG=7
- GDG=8

SMS Cache Sets

The Cache Sets attributes contain the name of the cache sets and the structures that they contain.

Cache Set Name: The name of this cache set. An SMS base configuration can contain a maximum of 256 Coupling Facility (CF) cache sets. Each set can contain up to eight Coupling Facility cache structures, which are defined by MVS policies.

Managed System: The SMS name of this system.

Structures: The structure names contained in this cache set.

SMS Configuration

The SMS Configuration attributes provide information about the base SMS installation.

Default Management Class: The name of the default management class. Management class is an optional SMS construct containing attributes that determine migration, backup and retention objectives for data sets.

Default Track Size: The tracks-size component of the default device geometry. SMS uses default device geometry when placing data sets, so that the amount of space is consistent and most efficiently used no matter what device type is chosen. Default device geometry isolates the user from the actual physical devices.

Default Tracks per Cylinder: The tracks-per-cylinder component of the default device geometry. SMS uses default device geometry when placing data sets, so that the amount of space is consistent and most efficiently used no matter what device type is chosen. Default device geometry isolates the user from the actual physical devices.

Default Unit Type: The unit-type component of the default device geometry.

Last Update User ID: The ID of the user that performed the last update.

Last Updated: The date and time of the last update to a storage group within the SMS configuration.

Managed System: The SMS name of this system.

MVS Level: The version number associated with the MVS system in use.

SCDS Name: The name of the Source Control Dataset. A linear VSAM data set containing an SMS configuration.

SMS Level: The version number associated with the DFP Product in use.

SMS Status: The status of the SMS subsystem (active or inactive) on the local system.

SMS Data Class

The SMS Data Class attributes provide information about the SMS data classes that have been defined.

Average Record Length: The average record length used for space allocation.

Compaction: Specifies whether the data sets allocated under this data class should be compressed when stored on tape or hard disk drive.

Control Area Free Space Percent: The percentage of free space to be left in the control area when it is loaded or after a CA split.

Control Interval Free Space Percent: The percent of free space to be left in the control interval when it is loaded or after a CI or CA split.

Control Interval Size: The size of the control interval.

Cross Region Share Option: The amount of sharing of clusters allowed by tasks in the same system.

Cross System Share Option: The level of sharing of the data sets allocated under this data class among multiple systems.

Data Class Name: The name of the data class. A data class is a collection of allocation and space attributes, defined by the storage administrator, that are used to create a data set.

Dataset Type: The format of the data sets allocated under this data class; either PDS, PDSE, HFS, or extended.

Description: A description of the data class.

Directory Blocks: The number of directory blocks in a persistent data store.

Expiration Date: The date the data sets allocated under this data class will be expired.

Extended Addressability: Allows a VSAM data set to grow beyond four GB when allocated in the extended format.

Index Options: The options for a VSAM index.

Key Length: The length of the key in bytes.

Key Offset: The displacement, in bytes, to the key field of a KSDS.

Last Update User ID: The ID of the user who last updated this data class definition.

Last Updated: The date and time that the data class definition was updated.

Logical Record Length: The logical record length of data sets allocated under this data class.

Managed System: The SMS name of this system.

Maximum Volumes: The maximum number of volumes data sets allocated under this data class can span.

Primary Space Allocation: The primary space allocation.

Record Bias: The system specifies whether to let VSAM determine how many and which type of buffers to use when accessing VSAM extended format data sets by batch processing.

Record Format: The record format of data sets allocated under this data class.

Record Organization: The organization of data sets allocated under this data class.

Record Scaling Factor: The scaling factor used to scale the average record length.

Reduce Space Percent: The percentage by which the allocation is reduced when the first attempt fails due to space being unavailable.

Retention Days: The number of days before data sets allocated under this data class are expired.

Secondary Space: The number of units to allocate for secondary extents.

SMS Management Class

The SMS Management Class attributes provide information about the SMS management classes that have been defined.

Automatic Backup: Specifies whether objects (data sets) of this class are automatically backed up by DFHSM.

Automatic Migration: The attributes which specify whether a data set is eligible for migration by command (command) automigration (auto) both, or none.

Backup Frequency: The backup frequency specifies how many days must elapse before DFHSM can back up data sets.

Backup Versions Dataset Deleted: Specifies the number of backup versions to retain for a data set that has been deleted from primary.

Backup Versions: Specifies the number of backup versions to retain for a data set.

Command Backup: Specifies the authority level necessary to issue backup commands for data sets in the management class. Values are as follows:

ADMIN: the storage administrator can issue backup commands

USER: user can issue backup commands

NONE: neither storage administrators nor users can issue backup commands

Copy Technique: Copy technique specifies the method to be used when data in this class is backed-up by DFSMSDss and DFSMSHsm.

- Standard
- Concurrent_Preferred
- Concurrent_Required
- Virtual_Preferred
- Virtual_Required
- Cache_Preferred
- Cache_Required

Description: A storage administrator supplied description of this management class.

Expiration NonUsage: Specifies the criteria to determine when an unreferenced data set is eligible for deletion during automatic space management.

Expiration: The number of days or date when the data set or object will be deleted.

GDS Elements on Primary: Specifies how many versions of a generation data set group use the normal primary days non-usage attribute for auto space management processing.

Last Update User ID: Specifies the user ID that last updated this management class.

Last Updated: The date and time when this management class was last updated.

Managed System: The SMS name of this system.

Management Class Name: The management class name. A named collection of management attributes describing the retention, backup and class transition characteristics for a group of objects in an object storage hierarchy.

Minimum Days Level 1: Specifies the minimum number of days since the data set was referenced on primary before a data set can be migrated from level1 to level2 migration.

Minimum Days Primary: Specifies the minimum number of days since a data set on primary was last referenced before it is eligible for migration.

Partial Release: The partial release specifies the conditions under which unused allocated space is released. Values are as follows:

NO: no release of unused space occurs.

YES: space is released at space management cycle time.

CONDITIONAL: if a non-zero secondary space allocation has been specified, release unused space at space management time.

Retain Extra Backup Version: Specifies the number of days to retain backup version in addition to the most recent one.

Retain Only Backup Version: Specifies the number of days to retain the most recent backup version of a deleted data set. This attribute applies only to data set no longer resident on primary.

Retention Limit: The retention limit value is a required value that limits the use of retention period (RETPD) and expiration date (EXPDT) values. These values are explicitly specified in JCL, are derived from data class definitions, or are explicitly specified in the OSREQ store macro. NO LIMIT effectively causes this value to be ignored.

Rolled off GDS: The rolled-off GDS action value indicates whether to expire rolled-off GDSs or make them eligible for migration.

SMS Storage Class

The SMS Storage Class attributes provide information about the SMS storage classes that have been defined.

Availability: Specifies whether data set access should continue in the event of a single device failure. This determines whether the data set is placed on RAID or duplexed volumes.

Cache Set Name: Specifies a name that is associated with a coupling facility structure that is used for data sets that are defined for VSAM record level sharing.

Description: The description of this storage class as entered by a storage administrator.

Direct Bias: Bias determines which volumes MSR (millisecond response time) performance numbers (READ, WRITE, or BOTH) to consider during volume selection. If you specify a READ (R) bias, a cache storage control should be available to allow caching. If you specify a WRITE (W) bias, the DASD fast write feather of an IBM 90 storage controller with cache should be available to allow the use of DASD fast write. If you do not specify a value for bias (BLANK), the MSR time determines whether caching or DASD fast write are used.

Direct MSR Objective: Specifies the desired millisecond response time for direct access I/O operations. This is used at allocation time to select a device capable of satisfying this objective.

Guaranteed Space: Specifies whether this storage class has the guaranteed space attribute. You can allocate space for single volume and multi-volume data sets by specifying a storage class with the guaranteed space attribute. The request fails when space is insufficient.

The guaranteed space attribute does not apply to objects, but it does apply to both VSAM and non-VSAM-managed data sets. For a multi-volume system-managed data set, primary space is preallocated on all volumes. The first volume becomes the primary volume. The remaining volumes become candidate volumes with preallocated space. When the primary extent on the current volume becomes full, VSAM attempts to create secondary extents on the current volume. If there is not enough space left on the current volume, VSAM uses the preallocated primary extent on the next volume. The next volume is converted from a candidate volume to a primary volume.

Guaranteed Synchronous Writes: Specifies whether guaranteed synchronous write is defined for this storage class. Ensures that data transfer occurs and is complete between the channel and the device when write operations are issued to a PDS.

Last Update User ID: The last user ID to update this storage class definition.

Last Updated: The date and time the storage class definition was last updated.

Managed System: The SMS name of this system.

Sequential Bias: Bias determines which volumes MSR (millisecond response time) performance numbers to consider during the volume selection. Values are READ WRITE, or BOTH.

Sequential MSR Objective: Specifies the desired millisecond response time for sequential access I/O operations. This is used at allocation time to select a device capable of satisfying this objective.

Storage Class Name: The name that identifies this storage class.

SMS Storage Group

The SMS Storage Group attributes provide information about the SMS storage groups that have been defined.

Automatic Backup System/Group: Specifies the name of the system or system group that is eligible to perform automatic backup processing.

Automatic Backup: Specifies whether the DASD volumes in this pool storage group are eligible for automatic backup processing.

Automatic Dump System/Group: Specifies the name of the system or system group that is eligible to perform automatic dump processing.

Automatic Dump: Specifies whether the DASD volume in this pool storage group are eligible for automatic dump processing.

Automatic Migration System/Group: Specifies the name of the system, or system group, that is eligible to perform an automatic migration processing.

Automatic Migration: Specifies whether this storage group is eligible for interval migration, primary space management migration, both, or no automatic migration.

Breakpoint Value: For Extended Address Volumes (EAV), when a disk space request is this size or more, expressed in cylinders, the system prefers to use the cylinder-managed space for that extent. This applies to each request for primary or secondary space for data sets that are eligible for the cylinder-managed space. If not enough cylinder-managed space is available, then the system uses the track-managed space or uses both areas. When the size of a disk space request is less than the breakpoint value, the system prefers to use the track-managed area and if enough space is not available there, then the system uses the cylinder-managed space or uses both areas. A value of 0 means that all EAS eligible data sets is allocated in the cylinder-managed space of a EAV by DADSM as long as space is available.

Description: The description of this storage group as entered by a storage administrator.

Dump Classes: Specifies up to five unique dump class names. When DFSMSHSM dumps DASD volumes that belong to the storage group it directs their contents to the dump classes.

Guaranteed Backup Frequency: Specifies the number days within the last backup period in which the backup process should have a copy of each of the data sets within the storage group. NOLIMIT specifies that data sets in the storage group are backed up according to management class specifications.

High Allocation Migration Threshold %: Specifies the upper limit, as a percentage of volume occupancy, above which new allocations to this volume is avoided by SMS, if possible. In addition, this threshold is used by HSM to determine whether data sets should be migrated off a DASD volume.

Last Update User ID: Specifies the user ID of the storage administrator who last updated this storage group definition.

Last Updated: Specifies the date and time this storage group definition was last updated by a storage administrator.

Low Allocation Migration Threshold %: The value that is used as the threshold goal in reducing the amount of space occupied on a DASD volume in the storage group during interval migration or daily space management. A value of zero causes all data sets in this storage group to be migrated if interval migration is also selected for this storage group.

Managed System: The SMS name of this system.

Storage Group Name: The name of this storage group as defined by a storage administrator.

Storage Group Type: Specifies the type of storage group defined by this row.

Tape Library1: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library2: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library3: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library4: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library5: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library6: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library7: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Library8: Specifies the name defined by a storage administrator for one of up to eight tape libraries that can own volumes within this storage group.

Tape Select: A tape library name used as an index value to find a storage group that can contain tape volumes owned by that library.

Track Managed Low Threshold %: Used as the threshold goal for the track managed area in reducing the amount of space occupied on a DASD volume in the storage group during interval migration or daily space management. This attribute specifies a lower limit, as a percentage of volume occupancy. A value of zero causes all data sets in this storage group to be migrated if interval migration is also selected for this storage group.

Track Managed High Threshold %: Specifies the upper limit, as a percentage of volume occupancy. When this value is exceeded, new allocations to this volume in the track managed area are avoided by SMS, if possible. In addition, this threshold is used by HSM to determine whether data sets should be migrated off a DASD volume.

VIO Maximum Size (KB): Specifies the maximum size, in kilobytes, of data sets to be allocated to VIO.

VIO Unit: Specifies the generic DASD device type that this storage group simulates.

Volser Select: A volume used as an index value to find the storage group that contains that volume.

SMS Storage Group Status

The SMS Storage Group Status attributes provide the SMS status of the storage group on each z/OS image in the SMSplex.

Managed System: The SMS name of this system.

Storage Group Name: Specifies the name of this storage group as defined in the base configuration.

Storage Group Status: The status for this storage group for this system or system group.

DisAll = Does not permit data sets in this storage group to be accessed.

DisNew = Does not permit the allocation of new data sets in this storage group.

Enable = Permits access to data sets in this storage group.

Notcon = The storage group is not defined.

QuiAll = Does not schedule more jobs that access data sets in the storage group.

QuiNew = Does not schedule any jobs that create new data sets in this storage group.

System ID: Specifies the name of the system or system group whose status is being displayed.

SMS Systems

The SMS Systems attributes provide information about the z/OS images that are in the same SMSplex.

Managed System: The SMS name of this system.

Name: Specifies the name of this system or system group as defined in the base configuration.

Type: Specifies whether this name is an individual system name or a system group name (name of a parallel sysplex).

Storage Toolkit Requests

This group of attributes provides details regarding actions that you have created and about the status of those actions.

Action Name: Name of the request as given by user.

Command: The command that the Storage Toolkit runs.

Command Property: The command properties to use for command resubmission.

Creator Name: Name of user who created action request.

Date Created: Date/time that the request was created.

Description: Description of the request as given by user.

Last Completed: Date/time that the most recent action completed.

Last Return Code: Return code for the most recently completed action.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Request ID: An ID number for the request.

Run Count: Number of times that the request has run.

Schedule: Dates and times when the action request runs. The value is a character string and specifies either a time or an interval. An **AT**[®], **IN**, or **EV** prefix in a schedule specification identifies the type of schedule for running action requests:

- The **AT** prefix is followed by a specific date and time, as in these examples. Each of the following example lines represents a separate example:

AT:hh:mm
AT:YY.DDD.hh:mm
AT:YY.MM.DD:hh: mm
AT:D1234567:hh: mm

(Go to the end of this bullet list for descriptions of the preceding variables.)

- The **IN** prefix is followed by an interval value, as in these examples. Each of the following example lines represents a separate example:

IN:hh:mm
IN:DDD:hh: mm

(Go to the end of this bullet list for descriptions of the preceding variables.)

- The **EV** (Every) prefix is followed by specific days and times of the week, as in this example:

*EV:D1234567 :hh:mm*where:

hh:mm is the hour and minute time or interval

YY is the two-digit year (00 to 99)

MM is the two-digit month (01 to 12)

DD is the two-digit day of the month (01 to 31)

DDD is the three-digit day of the year (001 to 366)

D1234567 indicates the days of the week (1=Monday to 7=Sunday). You can specify from 1 to 7 digits (days) after the letter D.

Type: Type of request. The request types are as follows:

- DFSMSHsm Volume Migration
- DFSMSHsm Volume Backup
- DFSMSHsm Dataset Migrate to Level 1
- DFSMSHsm Dataset Migrate to Level 2
- DFSMSHsm Dataset Backup
- DFSMSHsm Dataset Recover
- DFSMSHsm Dataset Recall
- DFSMSdss Volume Release
- DFSMSdss Volume Convert Test
- DFSMSdss Volume Convert Prepare
- DFSMSdss Volume Convert
- DFSMSdss Dataset Copy or Move
- DFSMSdss Dataset Release
- DFSMSdss Dataset Compress
- DFSMSdss Volume Defrag
- DFSMSdss Release
- DFSMSHsm Dataset Migrate
- DFSMSHsm Cancel Request
- DFSMSdss Compress
- DFSMSHsm
- DFSMSdss
- DFSMSrmm
- DFSMSrmm AddDataset
- DFSMSrmm ChangeDataset
- DFSMSrmm DeleteDataset
- DFSMSrmm AddOwner
- DFSMSrmm ChangeOwner
- DFSMSrmm DeleteOwner
- DFSMSrmm AddVolume
- DFSMSrmm ChangeVolume Confirm
- DFSMSrmm ChangeVolume
- DFSMSrmm DeleteVolume
- DFSMSrmm AddVRS
- DFSMSrmm ChangeVRS
- DFSMSrmm DeleteVRS
- IDCAMS
- ICKDSF
- TSO
- Mainframe
- Batch Job
- IDCAMS Dataset Delete
- IDCAMS Dataset Alter
- IDCAMS Dataset Catalog
- IDCAMS Dataset Allocate
- IDCAMS Dataset Define Cluster
- IDCAMS Dataset ListCat
- IDCAMS Dataset Print
- IDCAMS Dataset Repro
- IDCAMS Dataset Verify

Storage Toolkit Result Summary

This group of attributes provides a summary of the results from action requests that are either running or have completed.

Action Name: The name of the action.

Completed: Date/time that the action completed.

Execution Start: Date/time that the action request started.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Output Line Count: Number of lines of output.

Request ID: An ID number for the request.

Resource Count: Number of resources (VOLSERs, DSNs, and so on) that this action affected.

Result ID: An ID number for the result.

Return Code: Return code for the completed action.

Status: The status of the action. The status values are as follows:

- Submitted
- Executing
- Completed
- Cancelled
- Scheduled
- Held
- Pending
- BadToken
- InvalidJCL
- AuthorizationFailure
- DataSetError

Submitted: Date/time that the request was submitted.

Storage Toolkit Result Detail

This group of attributes provides details regarding the results from action requests that are complete.

The result for action requests can contain output from one or more steps or locations. For this reason, the **Step Name** and **Dataset or DD Name** attributes identify the source of the details.

- When you run a command from a command dialog or issue a command from the Issue Command dialog, the values for the **Step Name** and **Dataset or DD Name** are predefined by the Storage Toolkit.

| Command type | Step Name | Dataset or DD Name |
|--------------|-----------|--------------------|
| DFSMSdss | KS3DSS1 | SYSPRINT |
| DFSMSHsm | KS3HSM1 | KS3OUT |
| IDCAMS | KS3AMS1 | SYSPRINT |
| ICKDSF | KS3DSF1 | SYSPRINT |
| DFSMSrmm | KS3RMM1 | KS3OUT |

| Command type | Step Name | Dataset or DD Name |
|--------------|-----------|--------------------|
| TSO | KS3TSO1 | KS3OUT |

- When you create JCL using the Create Batch Job dialog box, the values for the **Step Name** and **Dataset or DD Name** reflect the values that you specified in the **Files whose contents should be copied for later viewing** fields on the **Options** tab.
- When you select **Copy JCL and JES logs for later viewing** on the **JCL** tab, the **Step Name** values are predefined by the Storage Toolkit. The **Dataset or DD Name** values are blank.

| Step Name | Description |
|-----------|----------------------------------|
| *JCL | The JCL that was submitted |
| *JESMSGLG | The contents of the JESMSGLG log |
| *JESJCL | The contents of the JESJCL log |
| *JESYSMSG | The contents of the JESYSMSG log |

No output from Mainframe commands is returned, because execution of the command is not synchronized with the execution of the batch job. Other reasons why you might not receive output include the following:

- The command that you are running might not have output that can be returned.
- A file or log to be copied for later viewing is empty.

Step Name: The name of a step in the JCL, the JCL log or JES log that identifies the source of this line of output.

Dataset or DD Name: The data set or DD name that identifies the source of this line of output.

File Name: The file name associated with this action result.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Output: The lines of output that the action produced.

Request ID: An ID number for the request.

Resource Name: Name of the resource (VOLSER, DSN, and so on) for which this output was generated.

Result ID: An ID number for the result.

Related information

“Using the Storage Toolkit” on page 254

Storage Toolkit Result Resources

This group of attributes tracks the identifiers for action names, requests, and results.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Resource Name: Name of the resource (VOLSER, DSN, and so on) for which this output was generated.

Request ID: An ID number for the request.

Result ID: An ID number for the result.

Symmetrix Configuration

The Symmetrix Configuration attributes provide status and configuration information about an EMC Symmetrix storage facility.

Cache Memory: The cache size in megabytes.

Cache Write Pending Ceiling Percent: The percentage of cache that can be dedicated to write I/O.

Code Check Sum: The microcode check sum.

Code Date: The date of the microcode (mdyy) in Hexadecimal format.

Code Version: The microcode version.

Logical Volumes: The number of logical units defined.

Managed System: The SMS name of this system.

Number of Directors: The number of directors.

Number of SSIDs: Specifies the number of SSIDs associated with this Symmetrix configuration.

Operating Hours: Specifies the number of hours since power-on.

Physical Volumes: Specifies the number of physical disks installed.

Serial Number Index: The eight low order characters of serial number which serve as the control unit index.

Serial Number: Specifies the serial number assigned.

SSID Index: One of up to 64 SSIDs that are associated with this Symmetrix configuration.

SSID_FWRD: The subsystem ID of the Symmetrix control unit that is controlling this device.

Time Stamp: Displays the date and time this data was collected from the virtual tape server (VTS).

Symmetrix Devices

The Symmetrix Devices attributes provide information about the physical disks in an EMC Symmetrix storage facility.

Buffer Ratio: The ratio of number of full buffers to number of empty buffers.

Cylinders: The number of cylinders mapped to the logical volume.

Director Number: The ID of the disk director managing this device.

Managed System: The SMS name of this system.

Mirror Number: The number of mirror volumes for this logical volume.

Not Ready Mirrors: The number of mirrors in a NOT-READY state.

RDF Device: The type of RDF device.

RDF Partner Symmetrix Device Number: The Remote Data Facility partner device number.

Remote Adapter Group: The remote adapter group.

SCSI Interface: The ID of the physical volume for which the logical volumes are being displayed.

Serial Number Index: The eight low order characters of serial number which serve as the control unit index.

SSID: The subsystem ID of the Symmetrix control unit that is controlling this device.

Symmetrix Device Number: The Symmetrix Logical device number.

Time Stamp: Displays the date and time this data was collected from the virtual tape server (VTS).

Valid Mirrors: The valid mirrors for this device.

Volser: The volume serial of a logical drive located on the physical drive.

Write Disabled Mirrors: The number of mirrors that are disabled for write.

Symmetrix Director Summary

The Symmetrix Director Summary attributes provide overview information about the directors in an EMC Symmetrix storage facility.

Director ID: The ID of the disk director managing these devices.

Director Index: The ID of the disk director managing these devices.

Disk IDX: The ID of the physical disk on which the logical devices are located.

Logical Volume Count: The number of logical volumes located on the physical drive.

MANAGED SYSTEM: The SMS name of this system.

SCSI Interface: The ID of the physical disk on which the logical devices are located.

Serial Number Index: The eight low order characters of serial number which serve as the control unit index.

SSID Index: The subsystem ID of the Symmetrix control unit that is controlling this device. To view the disk interfaces of a director controlled by another Symmetrix control unit, enter the four-character subsystem ID in this field to identify the control unit.

Subsystem ID: The subsystem ID of this Symmetrix disk director.

Time Stamp: Displays the date and time this data was collected from the virtual tape server (VTS).

Symmetrix Directors

The Symmetrix Directors attributes provide information about the disk directors on an EMC Symmetrix storage facility.

Cache Hit Pct Long Interval: The percentage of prefetched tracks found in cache during a long interval. A long interval is cache hits since IPL or STATUS CLEAR command.

Cache Hit Pct Short Interval: The percentage of prefetched tracks found in cache during a short interval. A short interval can be up to 10 minutes.

Cache Hits Long Interval: The count of cache hits since IPL or STATUS CLEAR command.

Cache Hits Short Interval: The count of cache hits during a short interval. A short interval can be up to 10 minutes.

Cache Misses Long: The number of cache misses during a long interval.

Cache Misses Short: The number of cache misses during a short interval.

Director Index: The ID of this director.

Director Number: Identifies the disk director.

Director Type: The type of director.

I/O Requests Long Interval: The number of I/O requests for a long interval.

I/O Requests Per Second Long Interval: The average number of I/O requests per second during a long interval. A long interval is since IPL or the STATUS CLEAR command.

I/O Requests Per Second Short Interval: The average number of I/O requests per second during a short interval. A short interval can be up to 10 minutes.

I/O Requests Short Interval: The number of I/O requests for a short interval.

I/Os Per Second Long Interval: The average number of I/O operations per second during a long interval. A long interval is since IPL or STATUS CLEAR command.

I/Os Per Second Short Interval: The average number of I/O operations per second during a short interval. A short interval can be up to 10 minutes.

Long Interval: A long measurement interval expressed as hours, minutes, and seconds.

Managed System: The SMS name of this system.

Prefetched Tracks Percent Used: The percentage of prefetched tracks used.

Serial Number Index: The eight low order characters of serial number which serve as the control unit index.

Short Interval: A short measurement interval expressed as hours, minutes, and seconds.

SSID Index: The SSID to which this director is attached.

SSID: The subsystem ID of the Symmetrix control unit controlling this disk director.

Time Stamp: Displays the date and time this data was collected from the virtual tape server (VTS).

Total CCW Chains: The total number of I/O requests handled by the director.

Total I/Os Long Interval: The total I/Os (reads and writes) for a long interval.

Total I/Os Short Interval: The total I/Os (reads and writes) for a short interval.

Total I/Os: The total number of I/Os.

Total Prefetch Cache Hits: The total number of prefetch cache hits.

Total Prefetch Tasks: The number of prefetch tasks.

Total Tracks Prefetched: The total number of prefetched tracks.

Total Write requests: Total write requests.

Unused Prefetch Tracks: The number of unused prefetch tracks.

Used Prefetch Tracks: The number of used prefetch tracks.

Write Pct Long Interval: The percentage of write commands during a long interval. The value is (Write channel commands) / (Reads and writes) A long interval is since IPL or STATUS CLEAR command.

Write Pct Short Interval: The percentage of write commands during a short interval. Write channel commands divided by reads and writes. A short interval can be up to 10 minutes.

Write Requests Long Interval: The count of write requests during a long interval.

Write Requests Short Interval: The count of write requests during a short interval.

Symmetrix SSIDs

The Symmetrix SSIDs attributes provide a list of the SSIDs associated with an EMC Symmetrix storage facility.

Managed System: The SMS name of this system.

SSID Index: One of up to 64 SSIDs associated with this Symmetrix configuration.

Subsystem IDs: One of up to 64 SSIDs associated with this Symmetrix configuration.

Tape Device

The Tape Device attributes provide status information about the tape devices in a tape group. A tape group is the set of devices in an IBM ATL, VTS, or the generic device type.

Device Address: The four hexadecimal digits that define a physical device. Also known as unit or unit address.

Device Status: The device status, including ONLINE, OFFLINE, ALLOCATED, or MOUNT PENDING.

Device Type: The generic tape drive type. Examples include 3480 and 3490.

Group Name: An OMEGAMON XE convention for identifying logical groupings of tapes by library of device. Tape subsystems that group volumes and devices using a library dataserwer are considered libraries. All other devices are grouped by device considered non-library or NO LIB. Each tape library is considered to be a tape group whose name is the same as the library name.

I/O Count: The count of I/O operations by this drive during this interval. You can enable aggregation for this attribute in Tivoli Data Warehouse.

I/O Per Second: The number of I/O operations by the device during this interval. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Jobname: The name of the job to which the tape drive is currently allocated.

Managed System: The SMS name of this system.

Mount Pending Time: The amount of time in minutes that a device has been in MOUNT PENDING status. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Permanent Errors: The number of permanent read or write errors that the device has incurred. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Storage Group: The name of the SMS tape storage group to which the tape volume belongs.

Temporary Errors: The number of temporary read and write errors that the device has incurred. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Volume: The volume serial number of the tape mounted or being mounted on the drive.

Tape Group

The Tape Group attributes provide status information for IBM ATLS, VTSs and generic groups of tape devices.

Available Drives: The number of tape drives that are online and have not been allocated to an application. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Available Slots: The number of slots in a tape library that are not in use. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Configured Drives: The number of tape drives configured to a tape group. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Console Name: The console associated with the tape library.

Drives Online: The number of tape drives that have been configured and are online. You can enable aggregation for this attribute in Tivoli Data Warehouse.

ID Number: The hardware identification of the tape library.

IO Per Second: The lowest number of I/Os, in I/Os per second, to a device currently active in the tape group. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Library Description: The description of the tape library.

Library Device: The tape library device model number.

Library Name: Shows the name of the tape library for which the information is displayed.

Library Type: The type of the tape library. Valid types are: automatic for automated tape libraries (ATLDS), and manual for manual tape libraries (MTLDS).

Managed System: The SMS name of this system.

Media1 Scratch Volumes: The number of scratch volumes of this type. Media1, for magnetic cartridges (3480 compatible). You can enable aggregation for this attribute in Tivoli Data Warehouse.

Media2 Scratch Volumes: The number of scratch volumes of this type. Media2, for enhanced capacity magnetic cartridges (3490E compatible). You can enable aggregation for this attribute in Tivoli Data Warehouse.

Media3 Scratch Volumes: The number of scratch volumes of this type. Media3, for magnetic cartridges (3590 compatible). You can enable aggregation for this attribute in Tivoli Data Warehouse.

Media4 Scratch Volumes: The number of scratch volumes of this type. Media4. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Media5 Scratch Volumes: The number of scratch volumes of this type. Media5. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Media6 Scratch Volumes: The number of scratch volumes of this type. Media6. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Media7 Scratch Volumes: The number of scratch volumes of this type. Media7. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Media8 Scratch Volumes: The number of scratch volumes of this type. Media8. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Mount Pend Time: The longest time in minutes that a device in the group has been waiting for a mount. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Mounts Pending: The number of tape drives waiting for mounts. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Permanent Errors: The highest number of permanent errors incurred by a tape drive in the group. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Temporary Errors: The highest number of temporary errors incurred by a tape drive in the group. You can enable aggregation for this attribute in Tivoli Data Warehouse.

Total Slots: The number of slots in a tape library. You can enable aggregation for this attribute in Tivoli Data Warehouse.

TotalStorageDS Array Configuration

The TotalStorage Array Configuration attributes provide configuration information about the arrays in a TotalStorageDS storage facility.

Array ID: The array Identifier.

Array Type: The RAID type of the array.

Array Width: The width of the array (number of Disk Drive Modules in the array).

DDM Capacity: The capacity (in gigabytes) of each Disk Drive Module in the array.

DDM Speed: The speed, in 1000 RPM, of the Disk Drive Modules in the array.

Extent Pool ID: The identifier of the extent pool.

Managed System: The internal node name used to identify the system from which this row of data was obtained.

Rank ID: The rank identifier.

SSD Array: Indicates if the array is a solid state device.

Storage Facility ID: The storage facility identifier of the device (Type-Model-Serial Number).

Time Stamp: The date and time associated with this row of data.

TotalStorageDS Configuration

The TotalStorage Configuration attributes provide configuration information about a TotalStorageDS storage facility.

Available Cache: The available TDS device cache in megabytes.

Configured NVS: The configured TDS device Non Volatile Storage in megabytes.

Logical Subsystems: The total number of logical subsystems defined.

Logical Volumes: The total number of logical volumes defined.

Managed System: The internal node name used to identify the system from which this row of data was obtained.

Pinned NVS: The pinned TDS device Non Volatile Storage in megabytes.

Storage Facility ID: The storage facility identifier of the device (Type-Model-Serial Number).

Storage Facility Serial Number: The storage facility serial number of the device.

Time Stamp: The date and time associated with this row of data.

Total Arrays: The total number of arrays defined.

Total Cache: The total TDS device cache in megabytes.

Total Extent Pools: The total number of extent pools defined.

Total Ranks: The total number of ranks defined.

Total SSD Arrays: The total number of solid state device arrays.

TotalStorageDS Extent Pool

The TotalStorage Extent Pool attributes provide information about the extent pools in a TotalStorageDS storage facility.

Allocated Extents: The total number of real allocated extents in the pool.

Avg Read Op Rate: The average read operations per second to all ranks in the extent pool.

Avg Read Resp Time: The average read response time (in milliseconds) to all ranks in the extent pool.

Avg Write Op Rate: The average number of write operations per second to all ranks in the extent pool.

Avg Write Resp Time: The average write response time (in milliseconds) to all ranks in the extent pool.

Dyn Reloc Src Extents: The total number of extents used as source for dynamic relocation.

Dyn Reloc Tgt Extents: The total number of extents used as the target for dynamic relocation.

Dyn Reloc Tgt Extents: The total number of extents used as the target for dynamic relocation.

Extent Pool ID: The identifier of the extent pool.

Managed System: The internal node name used to identify the system from which this row of data was obtained.

Max Read Op Rate: The maximum read operations per second to any rank in the extent pool.

Max Read Resp Time: The maximum read response time (in milliseconds) to any rank in the extent pool.

Max Write Op Rate: The maximum number of write operations per second to any rank in the extent pool.

Max Write Resp Time: The maximum write response time (in milliseconds) to any rank in the extent pool.

Pool Type: The extent pool type, either CKD (Count Key Data) or FB (Fixed Block).

Pool Utilization: The percentage of the total extents in the pool that are currently allocated.

Real Extents Converted: The total number of real extents converted in the extent pool.

Real Extents: The total number of real extents in the extent pool.

Real Pool Capacity: The total real capacity (in GB) of the extent pool.

SSD Ranks: The status of solid state devices in the extent pool, including All, None, and Mixed (both SSD and non-SSD ranks).

Storage Facility ID: The storage facility identifier for the device (Type-Model-Serial Number).

Storage Facility Serial Number: The storage facility serial number of the device.

Time Stamp: The date and time associated with this row of data.

Total Read Operations: The total read operations to all ranks in the extent pool.

Total Write Operations: The total number of write operations to all ranks in the extent pool.

Virt Extents Converted: The total number of virtual extents converted in the extent pool.

Virtual Extents: The total number of virtual extents in the extent pool.

Virtual Pool Capacity: The total virtual capacity (in gigabytes) of the extent pool.

TotalStorageDS Rank

The TotalStorage Ranks attributes provide performance information about the ranks in a TotalStorageDS storage facility.

Average Read Resp: The average read response time to the rank over the interval, in milliseconds.

Average Write Resp: The average write response time to the rank over the interval, in milliseconds.

Avg Bytes/ Read Op: The average number of bytes per read operation.

Avg Bytes/ Write Op: The average number of bytes per write operation.

Bytes Read Per Sec: The rate per second at which bytes were read from the rank over the interval.

Bytes Written Per Sec: The rate per second at which bytes were written to the rank over the interval.

Extent Pool ID: The extent pool identifier.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Number of Arrays: The total number of arrays that make up the rank.

Rank ID: The identification number of the rank.

Read Ops Per Sec: The rate per second that read operations were issued to the rank over the interval.

SSD Rank: Indicates that the rank is backed by SSD arrays.

Storage Facility ID: The storage facility identifier for the device (Type-Model-Serial Number).

Time Stamp: The date and time associated with this row of data.

Total Bytes Read: The total number of bytes read from the rank over the interval.

Total Bytes Written: The total number of bytes written to the rank over the interval.

Total Read Ops: The total number of read operations issued to the rank over the interval.

Total Read Resp: The total read response time to the rank over the interval, in milliseconds.

Total Write Ops: The total number of write operations issued to the rank over the interval.

Total Write Resp: The total write response time to the rank over the interval, in milliseconds.

Write Ops Per Sec: The rate per second that write operations were issued to the rank over the interval.

TotalStorageDS SSIDs

The TotalStorage SSIDs attributes provide a list of the SSIDs associated with a TotalStorageDS storage facility.

Managed System: The internal node name used to identify the system from which this row of data was obtained.

Storage Facility ID: The storage facility identifier for the device (Type-Model-Serial Number).

Subsystem IDs: A list of SubSystem IDs defined on this TotalStorage DS device.

Volume Group Summary

The Volume Group Summary attributes provide aggregate performance and aggregate space information for SMS storage groups and user DASD groups.

Busy: The percentage of time a resource is busy processing I/O requests. For PAV volumes, this value is normalized by dividing by the number of PAV exposures.

Description: This is a user-defined description used to identify an User DASD Group.

Device MPL: The highest MPL of all volumes associated with this storage group. This value is limited to those volumes that have met the minimum I/O requirements for exception processing.

Free Space Cylinders: The total amount of free space within the storage group, in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Free Space GB: The total amount of free space within the storage group, in gigabytes. This value does not include the space on volumes in an SMS storage group that have not been converted.

Free Space MB: The total amount of free space within the storage group, in megabytes. This value does not include the space on volumes in an SMS storage group that have not been converted.

Free Space Percent: The free space summed across all volumes within the group, expressed as a percentage. This value does not include volumes which are unavailable for new data set allocation.

Free Space Tracks: The total amount of free space within the storage group, in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Group Name: SMS storage group name or OMEGAMON User DASD Group name.

Group Type: The type of storage group described. Group type can be either of the following:

- **SMSGROUP:** Volumes owned by an SMS storage group of this name.
- **USERGROUP:** Volumes that are part of a user-defined group, which can contain SMS-controlled volumes and nonSMS-controlled volumes.

High Busy Percent: The highest percentage of time busy processing I/O requests. For PAV volumes, this value is normalized by dividing by the number of PAV exposures.

High DFW Retry Percent: The highest percentage of DFW Retries. A DASD Fast Write retry occurs when a DASD fast write operation cannot complete due to NVS (nonvolatile storage) constraints.

High Response Time: The highest MSR (millisecond response time) of any volume in the storage group.

High Volume Fragmentation Index: The value of the highest fragmentation index within the group. This value does not include volumes which are unavailable for new data set allocation.

Largest Free Extent Cylinders: The size of the largest free extent within the storage group, in cylinders. This value does not include volumes which are unavailable for new data set allocation.

Largest Free Extent GB: The size of the largest free extent within the storage group, in gigabytes. This value does not include volumes which are unavailable for new data set allocation.

Largest Free Extent MB: The size of the largest free extent within the storage group, in megabytes. This value does not include volumes which are unavailable for new data set allocation.

Largest Free Extent Tracks: The size of the largest free extent within the storage group, in tracks. This value does not include volumes which are unavailable for new data set allocation.

Low Read Hit Percent: The lowest read hit percentage of any volume in this group. A Read Hit occurs when a read request for data is made by the processor and resolved by cache. As a result, the request does not require DASD access.

Low Volume Free Space %: The lowest free space value for a volume within the group, expressed as a percentage. This value does not include volumes which are unavailable for new data set allocation.

Low Write Hit Percent: The lowest write hit percentage of any volume in this group. This value is limited to those volumes that have met the minimum I/O requirements for exception processing.

Managed System: The SMS name of this system.

Non-SMS Volumes: The number of non-SMS volumes in this storage group.

Non-Enabled Volumes: The indicator that this group contains SMS volumes with a status other than ENABLED.

SMS Volumes: Represents the number of SMS volumes in this storage group.

Storage Group Status: Represents the SMS storage group status for the system being monitored.

Storage Group Type: The SMS storage group type.

Time Stamp: The date and time this data was collected.

Total Space Cylinders: The total amount of space within the storage group, in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Total Space GB: The total amount of space within the storage group, in gigabytes. This value does not include the space on volumes in an SMS storage group that have not been converted.

Total Space MB: The total amount of space within the storage group, in megabytes. This value does not include the space on volumes in an SMS storage group that have not been converted.

Total Space Tracks: The total amount of space within the storage group, in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Total Volumes: The number of volumes in this storage group.

Track Managed Free Space MB: The total amount of track managed free space within the storage group, in megabytes. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Free Space Cylinders: The total amount of track managed free space within the storage group, in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Free Space GB: The total amount of track managed free space within the storage group, in gigabytes. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Free Space Percent: The percent of track managed free space in the storage group. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Free Space Tracks: The total amount of track managed free space within the storage group, in tracks. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed High Volume Frag Index: Value of the highest track managed area fragmentation index within the group. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Largest Free Extent MB: Size of the largest free track managed area extent within the storage group, in megabytes. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Largest Free Extent GB: Size of the largest free track managed area extent within the storage group, in gigabytes. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Largest Free Extent Cylinders: Size of the largest free track managed area extent within the storage group, in cylinders. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Largest Free Extent Tracks: Size of the largest free track managed area extent within the storage group, in tracks. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Low Volume Free Space %: Lowest track managed area free space value for a volume within the group, expressed as a percentage. This value does not include volumes which are unavailable for new dataset allocation.

Track Managed Total Space MB: The total amount of track managed space within the storage group, in megabytes. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Total Space Cylinders: The total amount of track managed space within the storage group, in cylinders. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Total Space Tracks: The total amount of track managed space within the storage group, in tracks. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Total Space GB: The total amount of track managed space within the storage group, in gigabytes. This value does not include the space on volumes in an SMS storage group that have not been converted.

Track Managed Used Space Percent: The percentage of track managed used space in the storage group or the User DASD group. This includes both the track and cylinder managed areas of the volumes.

User Group Name: This is a user-defined name used to identify a User DASD group.

Used Space Percent: The percentage of used space in the storage group or the User DASD group. This includes both the track and cylinder managed areas of the volumes.

VTOC Index Status: The worst status of any VTOC index in the group.

Volume Status

The Volume Status attributes provide the z/OS and SMS status of a volume across the z/OS images in an SMSplex.

Managed System: An internal node name used to identify the system from which this row of data was obtained.

MVS Status: The z/OS status of this device in the named system or group.

SMS Status: The SMS status of this device in the named system or group.

System ID: The SMF ID or group ID of the system from which this information was obtained.

Volser: The volume serial number of a disk device.

Volume Users

The Volume Users attributes allow a system administrator to identify the applications that are allocated to a logical volume and the data sets that are being used on the volume by each application.

Application: The name of the application (STC, TSO user, or batch job) using the volume.

ASID: ASID of the application using the volume, displayed as 4 hex digits.

Data Set Name: The name of the data set on the volume that is allocated.

DD Name: The DDNAME associated with the allocation.

Managed System: The name identifying the monitored system.

Open Count: Number of opens issued by the job against the dataset.

VTS Cache

The VTS Cache attributes provide performance and capacity information about the cache for a virtual tape server (VTS).

Average Recall Throttling Value: The average of the NON-ZERO throttling values over the last hour for the VTS where recall was the predominate reason.

Average Throttling Value: The average of both the ZERO and the NON-ZERO throttling values over the last hour for the VTS.

Average TVC Volume Age Minutes: The percentage of throttling over the last hour for the VTS where write overrun was the predominate reason.

Average Write Throttling Value: The average of the NON-ZERO throttling values over the last hour for the VTS where write overrun was the predominate reason.

Cache Hits: The count of cache hits a mount requests for a virtual tape volume in AVTS for which the virtual tape volume is already located in the TVC. In this case, there is no physical tape mount associated with the virtual mount request.

Fast Ready Mounts: The count of fast ready mounts. A mount attribute that allows a scratch mount request in a virtual tape server to be satisfied by the tape header information kept in the TVC for a previously-initialized virtual tape volume.

GB TVC Cache: The size of the cache in the TVC of the VTS in units of gigabytes.

Managed System: The SMS name of this system.

Maximum TVC Volume Age Minutes: The maximum age in minutes, of all virtual tape volumes in the TVC at the end of the reporting interval.

Pre-migration Mounts: The number of processes used to copy virtual volumes to logical volumes. For example, from the TVC to a physical tape, by the VTS. Normally, this number is relatively high after high write activity.

Recall Throttling Percent: The percentage of throttling over the last hour for the VTS where recall was the predominate reason.

Stage-In Mounts: The count of stage-in mounts. A stage-in mount is a mount request for a virtual tape volume in a virtual tape system for which the virtual tape volume is not residing in the TVC. In this case, there is a physical tape mount associated with the virtual mount request to access the logical volume from a physical tape.

Time Stamp: Displays the date and time this data was collected from the VTS.

TVC Virtual Volume Count: The number of virtual volumes in the TVC at the end of the reporting interval.

TVC Virtual Volume Size MB: The average size in megabytes of the virtual volumes in the TVC at the end of the reporting interval.

VTS Library Type: Specifies the library type of this VTS. Values for library type are:

- **STANDALONE VTS:** a VTS that is not part of a peer-to-peer VTS environment.
- **DISTRIBUTED VTS:** a VTS that is part of a peer-to-peer VTS environment, but has not been designated for user interface commands.
- **DISTRIBUTED USER INTERFACE VTS:** a VTS that is part of a peer-to-peer VTS environment, and has been designated for user interface commands.
- **COMPOSITE VTS:** the logical view of a peer-to-peer VTS environment.

VTS Name: Specifies the name of this virtual tape server.

Write Throttling Percent: The percentage of 30-second intervals where there was a non-zero throttling value and write overrun was the predominate reason.

VTS Capacity

The VTS Capacity attributes provide utilization information about the physical tapes utilized by a virtual tape server (VTS).

Active Data TB: The number of bytes, in terabytes, of data being managed by this VTS.

Active Logical Volumes(1000): The number of logical volumes managed by the VTS that contains active data. This data is reported in units of 1000.

Empty Physical Volumes: The number of empty physical cartridges available for use by the VTS.

Empty Volume Free Space GB: The estimated storage capacity in gigabytes, available on the empty 3590 cartridges managed by the VTS.

ESCON® Channels Connected: The number of ESCON channels connected to this VTS.

Host Channel Activity GB: The total, READ + WRITE, number of gigabytes (1,073,741,824) processed by the MVS host channels to the virtual tape volumes associated with the VTS.

Host Channel GB Read: The number of gigabytes (1,073,741,824) read by the MVS host channels from the virtual tape volumes associated with the VTS.

Host Channel GB Written: The number of gigabytes (1,073,741,824) written by the MVS host channels to the virtual tape volumes associated with the VTS.

Host Channel Read Percent: The percentage of all host channel byte movement that was the result of a read request by the MVS host channels from the virtual tape volumes associated with the VTS.

Managed System: The SMS name of this system.

Physical Tape GB Read: The number of gigabytes (1,073,741,824) read by the MVS host channels from the virtual tape volumes associated with the VTS.

Physical Tape GB Written: The number of gigabytes (1,073,741,824) written by the virtual tape server to the physical tape drives.

Physical Tape Read Percent: The percentage of all physical tape byte movement that was the result of a read request by the VTS.

Reclaim Threshold Percentage: The percentage identified when a physical cartridge is made available for reclamation. When the percentage of active data falls below the reclaim threshold, the physical volume is available to go through reclamation.

SCSI Channels Connected: The number of SCSI channels connected to this VTS.

Time Stamp: Displays the date and time this data was collected from the VTS.

Volumes 0-5% Utilized: The number of volumes containing 0 to 5% active data. This field is 0 for a composite library.

Volumes 10-15% Utilized: The number of volumes containing 10 to 15% active data. This field is 0 for a composite library.

Volumes 15-20% Utilized: The number of volumes containing 15 to 20% active data. This field is 0 for a composite library.

Volumes 20-25% Utilized: The number of volumes containing 20 to 25% active data. This field is 0 for a composite library.

Volumes 25-30% Utilized: The number of volumes containing 25 to 30% active data. This field is 0 for a composite library.

Volumes 30-35% Utilized: The number of volumes containing 30 to 35% active data. This field is 0 for a composite library.

Volumes 35-40% Utilized: The number of volumes containing 35 to 40% active data. This field is 0 for a composite library.

Volumes 40-45% Utilized: The number of volumes containing 40 to 45% active data. This field is 0 for a composite library.

Volumes 45-50% Utilized: The number of volumes containing 45 to 50% active data. This field is 0 for a composite library.

Volumes 50-55% Utilized: The number of volumes containing 50 to 55% active data. This field is 0 for a composite library.

Volumes 5-10% Utilized: The number of volumes containing 5 to 10% active data. This field is 0 for a composite library.

Volumes 55-60% Utilized: The number of volumes containing 55 to 60% active data. This field is 0 for a composite library.

Volumes 60-65% Utilized: The number of volumes containing 60 to 65% active data. This field is 0 for a composite library.

Volumes 65-70% Utilized: The number of volumes containing 65 to 70% active data. This field is 0 for a composite library.

Volumes 70-75% Utilized: The number of volumes containing 70 to 75% active data. This field is 0 for a composite library.

Volumes 75-80% Utilized: The number of volumes containing 75 to 80% active data. This field is 0 for a composite library.

Volumes 80-85% Utilized: The number of volumes containing 80 to 85% active data. This field is 0 for a composite library.

Volumes 85-90% Utilized: The number of volumes containing 85 to 90% active data. This field is 0 for a composite library.

Volumes 90-95% Utilized: The number of volumes containing 95 to 100% active data. This field is 0 for a composite library.

Volumes 95-100% Utilized: The number of volumes containing 95 to 100% active data. This field is 0 for a composite library.

VTS Library Type: Specifies the library type of this VTS. Values for library type are:

STANDALONE VTS: A VTS that is not part of a peer-to-peer VTS environment.

DISTRIBUTED VTS: A VTS that is part of a peer-to-peer VTS environment, but has not been designated for user interface commands.

DISTRIBUTED USER INTERFACE VTS: A VTS this is part of a peer-to-peer VTS environment, and has been designated for user interface commands.

COMPOSITE VTS : The logical view of a peer-to-peer VTS environment.

NOT APPLICABLE: The library type attribute does not apply to TS7700 resources.

VTS Name: The name of the VTS as assigned via the ISMF panels.

VTS Composite Library

The VTS Composite Library attributes provide performance statistics for a composite virtual tape server (VTS).

Cache Hit Mounts: The number of mount requests for a virtual tape volume in a VTS for which the virtual tape volume is already located in the TVC. In this case, there is no physical tape mount associated with the virtual mount request.

Category Mounts: The count of category or scratch mounts for this composite server.

Controller ID: The controller ID associated with this composite server.

GB Data Copied: The number of gigabytes of data copied by this composite server.

Managed System: The SMS name of this system.

Read Gigabytes: The amount, in gigabytes, of data read by the composite server.

Recall Mounts: The number of virtual mounts that result in a physical mount in the VTS.

Time Stamp: Displays the date and time this data was collected from the VTS.

Volumes Copied: The number of volumes copied by this composite server.

VTS Name: The name of the VTS as assigned through the ISMF panels.

Write Gigabytes: The amount, in gigabytes, of data written by this composite server.

VTS Overview

The VTS Overview attributes provide an overview of the performance and configuration of a virtual tape server (VTS).

Average Virtual Mount Pend Time: The average amount of time in seconds that an MVS system waited for a virtual tape volume to be mounted on a virtual tape server.

Grid Library Sequence Number This five-character field contains the Library Sequence Number of the Grid (Composite) library.

Host Channel Activity GB: The total number of gigabytes (read + write), processed by the MVS host channels to the virtual tape volumes associated with the VTS.

Library Serial Number: The manufacturer's serial number for the VTS.

Library Type: The physical tape library type.

Managed System: The SMS name of this system.

Maximum Virtual Mount Pend Time: The maximum amount of time in seconds that an MVS system waited for a virtual tape volume on a virtual tape drive.

Model number: The manufacturer's model number for the VTS.

Physical Drives Available: The number of physical tape drives available on the VTS.

Scratch Volumes: The number of empty physical cartridges available for use by the virtual tape server.

Time Stamp: Displays the date and time this data was collected from the VTS.

Virtual Devices Online: The number of virtual tape drives on line for the VTS.

Virtual Disconnect Time: The time that I/O spends, free from the channel, searching for or requesting data. This time includes the time spent waiting for the VTS to service the virtual device.

Volume Recall Percent: The percentage of virtual mounts that result in a physical mount in the VTS.

VTS Library Type: Specifies the library type of this VTS. Valid library types are:

- **STANDALONE VTS:** A VTS that is not part of a peer-to-peer VTS environment.
- **DISTRIBUTED VTS:** A VTS that is part of a peer-to-peer VTS environment, but has not been designated for user interface commands.
- **DISTRIBUTED USER INTERFACE VTS:** A VTS that is part of a peer-to-peer VTS environment, and has been designated for user interface commands.
- **COMPOSITE VTS:** The logical view of a peer-to-peer VTS environment.
- **NOT APPLICABLE:** The library type attribute does not apply to TS7700 resources.

VTS Name: The name of this virtual tape server.

VTS Physical Devices

The VTS Physical Devices attributes provide detailed performance and configuration information about the physical devices in a virtual tape server (VTS).

Average Concurrent Mounts: The average number of physical tape drives available to a VTS that were used at the same time.

Average Mount Pend Time: The average amount of time in seconds that this VTS waited for a tape volume to be mounted on a physical tape drive.

Managed System: The SMS name of this system.

Maximum Concurrent Mounts: The maximum number of physical tape drives available to VTS that were used at the same time.

Maximum Mount Pend Time: The maximum amount of time in seconds that this VTS waited for a tape volume to be mounted on a physical tape drive.

Migration Mounts: The number of physical tape drive mounts required to move new data from the TVC to the physical tape volumes associated with this VTS.

Minimum Concurrent Mounts: The minimum number of physical tape drives available to a VTS that were used at the same time.

Minimum Mount Pend Time: The minimum amount of time in seconds that this VTS waited for a tape volume to be mounted on a physical tape drive.

Physical Drives Available: The number of physical tape drives available to this VTS to move data between the TVC and the physical tape volumes.

Physical Drives Installed: The number of physical tape drives installed in this VTS.

Reclamation Mounts: The number of physical tape drive mounts required to remove obsolete data from the physical tape volumes and to consolidate active data on physical tape volumes associated with this VTS.

Stage Mounts: The number of mount requests in this VTS for a virtual tape volume for which the virtual tape volume is not located in the TVC. In this case, there is a physical tape mount associated with the virtual mount request to access the logical volume from a physical tape.

Time Stamp: Displays the date and time this data was collected from the VTS.

VTS LIB TY: Specifies the library type of this VTS.

Stand Alone: VTS that is not part of a peer-to-peer VTS environment.

Distributed: VTS that is part of a peer-to-peer VTS environment, but has not been designated for user interface commands.

Distributed UI: VTS that is part of a peer-to-peer VTS environment, and has been designated for user interface commands.

Composite: Logical view of a peer-to-peer VTS environment.

Not Applicable: The library type attribute does not apply to TS7700 resources.

VTS Name: Specifies the name of this virtual tape server.

VTS TS7700-series Cache Container

This attribute group provides basic information about a cache in a TS7700 tape virtualization system. All attributes in this table are derived from the "HSM - Cache Container" record of the "Hnode HSM Historical Record."

Active Hnode: Specifies whether this Hnode was the active node in charge of Tape Volume Cache (TVC) management at the end of the interval being reported.

Average Copy Throttle: Average copy throttle value during the interval. The value presented is the average of the non-zero throttling values where copy was the predominant reason for throttling. The value is reported in thousandths of a second.

Average Overall Throttle: Average of all throttling values during the interval. The calculation includes samples for periods where throttling was both zero and non-zero. The value is reported in thousandths of a second.

Average Write Overrun Throttle: Average write overrun throttle value during the interval. The value presented is the average of the non-zero throttling values where write overrun was the predominant reason for throttling. The value is reported in thousandths of a second.

Cluster ID: Hexadecimal value indicating the cluster ID. Values are 0 to 7.

Grid Library Sequence Number: Library Sequence Number of the Grid (Composite) library.

Machine Model: Machine model of the node. Initially this field is set to "V06".

Machine Serial Number: Serial number of the node. This field is left justified and padded with blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed. It should be noted that there can only be one machine per-cluster.

Machine Type: Machine type of the node. Initially this field is set to "3957".

Managed System: The SMS name of this system.

Node ID: A hexadecimal field that indicates the ID of the node.

Partition Container Number: An integer value 0-7 that indicates which cache partition you are viewing.

Percent Copy Throttle: Percentage of 30-second periods where there was at least one throttling value greater than zero and that copy was the predominant reason for throttling.

Percent Write Overrun Throttle: Percentage of 30-second periods where there was at least one throttling value greater than zero and that write overrun was the predominant reason for throttling.

Timestamp: Time table row was created in CT timestamp format.

Tape Volume Cache Size: Current size of the cluster's Tape Volume Cache (TVC) in increments of 1GB (1024 x 1024 x 1024). A TVC that is 1.7 TB in size is reported as x000006A4 (1700 decimal).

VE Code Level: Code level of the TS7700 Virtualization Engine (VE). The VE code level is expressed as Version.Release.Modification.Fix in a decimal form.

VTS TS7700-series Cache Partition Container

This attribute group provides basic information about a cache partition in a TS7700 tape virtualization system. All attributes in this table are derived from the "HSM - Cache - Partition Container" record of the "Hnode HSM Historical Record."

Average Cache Hit Mount Time: Average time, in milliseconds, taken to complete Cache Hit mounts during the interval. Mount time is accrued from the time the mount request is accepted by the system (PLF received, DE returned) until the x20 message is received from the Library Manager. The mount time is averaged into the interval time when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.

Average Cache Miss Mount Time: Average time, in milliseconds, taken to complete Cache Miss mounts during the interval. Mount time is accrued from the time the mount request is accepted by the system (PLF received, DE returned) until the x20 message is received from the Library Manager. The mount time is averaged into the interval time when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.

Average Fast Ready Mount Time: Average time, in milliseconds, taken to complete Fast-Ready mounts during the interval. Mount time is accrued from the time the mount request is accepted by the system (PLF received, DE returned) until the x20 message is received from the Library Manager. The mount time is averaged into the interval time when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.

Cache Hit Mounts: Number of mount requests completed that the data was resident in the Tape Volume Cache (TVC) during this interval. A mount is accredited to the interval when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.

Cache Miss Mounts: Number of mount requests completed that required recall from a stacked volume during this interval. A mount is accredited to the interval when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.

Cluster ID: Hexadecimal value indicating the cluster ID. Values are 0 to 7.

Fast Ready Mounts: Number of mount requests completed using the Fast Ready method during this interval. A mount is accredited to the interval when the x20 message is received from the Library Manager. This value is reset to 0 at the beginning of the interval.

Grid Library Sequence Number: Library Sequence Number of the Grid (Composite) library.

Machine Model: Machine model of the node. Initially this field is set to "V06".

Machine Serial Number: Serial number of the node. This field is left justified and padded with blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed. It should be noted that there can only be one machine per-cluster.

Machine Type: Machine type of the node. Initially this field is set to "3957".

Managed System: The SMS name of this system.

Node ID: A hexadecimal field that indicates the ID of the node.

Partition Container Number: An integer value 0-7 that indicates which cache partition you are viewing.

Partition Size: Amount of cache assigned to this partition. The value is reported in increments of 1 GB.

Timestamp: Time table row was created in CT timestamp format.

VE Code Level: Code level of the TS7700 Virtualization Engine (VE). The VE code level is expressed as Version.Release.Modification.Fix in a decimal form.

VTS TS7700-series Cache Preference Group

This attribute group provides basic information about a cache preference group in a TS7700 tape virtualization system. All attributes are derived from the HSM HSM - Cache - Partition - Preference Group Container record of the "Hnode HSM Historical Record."

35 Day Average Cache Age: 35-day rolling average of cache age, in minutes, of the virtual volumes migrated out of the cache partition that were assigned to the preference group this data is for. Cache age is measured from when a volume is created or recalled into cache until it has been migrated from cache. The cache age of each volume is rounded up to the nearest minute. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.

4 Hour Average Cache Age: Four-hour rolling average of cache age, in minutes, of the virtual volumes migrated out of the cache partition that were assigned to the preference group this data is for. Cache age is measured from when a volume is created or recalled into cache until it has been migrated from cache. The cache age of each volume is rounded up to the nearest minute. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.

48 Hour Average Cache Age: 48-hour rolling average of cache age, in minutes, of the virtual volumes migrated out of the cache partition that were assigned to the preference group this data is for. Cache age is measured from when a volume is created or recalled into cache until it has been migrated from cache. The cache age of each volume is rounded up to the nearest minute. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.

Cluster ID: Hexadecimal value indicating the cluster ID. Values are 0 to 7.

Data Resident in Cache: Amount of data in the TVC partition whose volumes are assigned to the preference this data is for. The value is reported in increments of 1 MB. Any residual data causes the value to be rounded up to the next higher value. This metric is the value at the end of the interval.

Grid Library Sequence Number: Library Sequence Number of the Grid (Composite) library.

Machine Model: Machine model of the node. Initially this field is set to "V06".

Machine Serial Number: Serial number of the node. This field is left justified and padded with blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed. It should be noted that there can only be one machine per-cluster.

Machine Type: Machine type of the node. Initially this field is set to "3957".

Managed System: The SMS name of this system.

Node ID: A hexadecimal field that indicates the ID of the node.

Partition Container Number: Integer value 0-7 that indicates which cache partition you are viewing.

Preference Group Number: Integer value (0 or 1) that denotes the preference group of this particular row.

Timestamp: Time table row was created in CT timestamp format.

VE Code Level: Code level of the TS7700 Virtualization Engine (VE). The VE code level is expressed as Version.Release.Modification.Fix in a decimal form.

Virtual Volumes in Cache: Number of virtual volumes in the Tape Volume Cache (TVC) partition that are assigned to the preference group this data is for. This is the value at the end of the interval.

Volumes Migrated Last 35 Days: Number of virtual volumes migrated from the cache partition over the past 35 days that are assigned to the preference group this data is for. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.

Volumes Migrated Last 4 Hours: Number of virtual volumes migrated from the cache partition over the past 4 hours that are assigned to the preference group this data is for. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.

Volumes Migrated Last 48 Hours: Number of virtual volumes migrated from the cache partition over the past 48 hours that are assigned to the preference group this data is for. This data is calculated once an hour, on the hour. The data for this field is calculated at the end of the interval.

VTS TS7700-series Cluster

This attribute group provides basic information about a cluster in a TS7700 tape virtualization system. The attributes include metrics that are derived from standard sources, such as the virtualization (vNode) or hierarchical data storage management (hNode). These nodes are part of the TS7700 Virtualization Engine distributed-node architecture.

Average Virtual Mount Pend Time: Average amount of time in seconds that an MVS system waited for a virtual tape volume to be mounted on this cluster of the selected TS7700 grid.

Cluster ID: Hexadecimal value indicating the cluster ID. Values are 0 to 7.

Grid Library Sequence Number: Library Sequence Number of the Grid (Composite) library.

Host Channel Activity GB: Total number of gigabytes (read + write), processed by the z/OS host channels to the virtual tape volumes associated this cluster of the selected TS7700 grid.

Machine Model: Machine model of the node. Initially this field is set to "V06".

Machine Serial Number: Serial number of the node. This field is left justified and padded with blanks. The format is XX-YYYYY where XX is the plant of manufacture and the YYYYY is the sequence number of the node's machine. The dash character (-) is fixed. It should be noted that there can only be one machine per-cluster.

Machine Type: Machine type of the node. Initially this field is set to "3957".

Managed System: An internal node name used to identify the system from which this row of data was obtained.

Node ID: This hexadecimal field indicates the Node ID. Valid values are 0-1.

Physical Drives Available: Number of physical tape drives available on this cluster of the selected TS7700 grid.

Scratch Volumes: Number of empty physical cartridges available for use in this cluster of the selected TS7700 grid.

Timestamp: Time table row was created in CT timestamp format.

VE Code Level: Code level of the TS7700 Virtualization Engine (VE). The VE code level is expressed as Version.Release.Modification.Fix in a decimal form.

Volume Recall Percent: Percentage of virtual mounts that result in a physical mount in this cluster of the selected TS7700 grid.

VTS Virtual Devices

The VTS Virtual Devices attributes provide detailed internal performance information about the virtual devices associated with a virtual tape server (VTS).

Average Cache Hit Mount Time: The average time required for a mount request for a virtual tape volume in a TVS for which the virtual tape volume is already located in the TVC. In this case, there is no physical tape mount associated with the virtual mount request.

Average Concurrent Mounts: The average number of virtual tape drives available to this VTS that were used at the same time.

Average Fast Ready Mount Time: The average fast-ready mount time for tape drives in this VTS. Requests in a virtual tape server are satisfied by the tape header information kept in the TVC for a previously-initialized virtual tape volume. This approach avoids a physical tape mount for volume verification and tape management system integrity checking.

Average Mount Pend Time: The average amount of time in seconds that an MVS system waited for a virtual tape volume to be mounted on a virtual tape drive.

Average Recall Mount Time: The average time required to perform a recall mount on this VTS. A recall mount is a virtual mount that results in a physical mount.

Average Time Mount Used: The average time in seconds that a virtual tape drive in this VTS was mounted. This time is accrued from the completion of the mount until a demount request is issued.

Data To Copy: The amount of data that to be copied in gigabytes.

LCU Disconnect Time: The average disconnect time in milliseconds for each I/O to each virtual drive in the logical control unit of this virtual tape server.

Managed System: The SMS name of this system.

Maximum Cache Hit Mount Time: The maximum time required for a mount request for a virtual tape volume in a TVS for which the virtual tape volume is already located in the TVC. In this case, there is no physical tape mount associated with the virtual mount request.

Maximum Concurrent Mounts: The maximum number of virtual tape drives available to this VTS that were used at the same time.

Maximum Fast Ready Mount Time: The maximum fast-ready mount time for tape drives in this VTS. Requests in a virtual tape server are satisfied by the tape header information kept in the TVC for a previously-initialized virtual tape volume. This approach avoids a physical tape mount for volume verification and tape management system integrity checking.

Maximum Mount Pend Time: The maximum amount of time in seconds that an MVS system waited for a virtual tape volume to be mounted on a virtual tape drive.

Maximum Recall Mount Time: The maximum time required to perform a recall mount on this VTS. A recall mount is a virtual mount that results in a physical mount.

Maximum Time Mount Used: The maximum time in seconds that a virtual tape drive in this VTS was mounted. This time is accrued from the completion of the mount until a demount request is issued.

Minimum Cache Hit Mount Time: The minimum time required for a mount request for a virtual tape volume in a TVS for which the virtual tape volume is already located in the TVC. In this case, there is no physical tape mount associated with the virtual mount request.

Minimum Concurrent Mounts: The minimum number of virtual tape drives available to this VTS that were used at the same time.

Minimum Fast Ready Mount Time: The minimum fast-ready mount time for tape drives in this VTS. Requests in a virtual tape server are satisfied by the tape header information kept in the TVC for a previously-initialized virtual tape volume. This approach avoids a physical tape mount for volume verification and tape management system integrity checking.

Minimum Mount Pend Time: The minimum amount of time in seconds that an MVS system waited for a virtual tape volume to be mounted on a virtual tape drive.

Minimum Recall Mount Time: The minimum time required to perform a recall mount on this VTS. A recall mount is a virtual mount that results in a physical mount.

Minimum Time Mount Used: The minimum time in seconds that a virtual tape drive in this VTS was mounted. This time is accrued from the completion of the mount until a demount request is issued.

Time Stamp: Displays the date and time this data was collected from the VTS.

Unit B180 Mounts: The mounts for unit B18-0.

Unit B181 Mounts: The mounts for unit B18-1.

Virtual Drives Configured: The number of virtual tape drives configured for this VTS.

Virtual Drives Online: The number of virtual tape drives online for this virtual tape server.

Volumes To Copy: The number of logical volumes that can be copied.

VTS LIB TY: Specifies the library type of this VTS.

Stand Alone: VTS that is not part of a peer-to-peer VTS environment.

Distributed: VTS that is part of a peer-to-peer VTS environment, but has not been designated for user interface commands.

Distributed UI: VTS that is part of a peer-to-peer VTS environment, and has been designated for user interface commands.

Composite: Logical view of a peer-to-peer VTS environment.

VTS Name: The name of this virtual tape server.

VTS Virtual Devices Summary

The VTS Virtual Devices Summary attributes provide response time information for the virtual devices assigned to a virtual tape server (VTS).

Busy Percent: The percentage of time this virtual drive was busy processing I/O requests.

Connect Time: The connect time for this virtual drive. Connect time is the time that I/O spends searching for or requesting data.

Device Address: The MVS address of this virtual drive.

Device Model: The generic tape drive type, such as 3590-1.

Device Status: The device status of this virtual drive including Online, Offline, Allocated, or Mount Pending.

Disconnect Time: The disconnect time for this virtual drive. Disconnect time is the time that I/O spends, free from the channel, searching for or requesting data.

I/O Per Second: The number of I/O operations completed per second for this virtual drive.

Managed System: The SMS name of this system.

Pend Time: The total pending for this virtual drive. Pending time is time that I/O is delayed in the path to the device. Pending time can be attributable to the channel or control unit path being busy.

Response Time: The average response time for this virtual drive in milliseconds. The amount of time (in milliseconds) it takes for an I/O request to complete.

Time Stamp: Displays the date and time this data was collected from the VTS.

VTS Name: The name of this virtual tape server.

Chapter 23. Storage Toolkit

This topic describes the Storage Toolkit feature, which enables you to use *actions* in the Tivoli Enterprise Portal. The Storage Toolkit actions provide convenient access to mainframe operations that let you target resources such as volumes and mainframe data sets. For example, you use the actions to migrate data sets or release free space.

Product-provided menus and dialog boxes enable you to initiate a broad range of commonly used mainframe commands. You can also issue custom commands. Using the Storage Toolkit, you can invoke commands from the following facilities in a z/OS system:

- DFSMSdss (Device Support Services)
- DFSMSHsm (Hierarchical Storage Manager)
- ICKDSF (Device Support Facilities)
- IDCAMS (Access Method Services)
- DFSMSRmm (DFSMS Removable Media Manager)
- Mainframe (z/OS console command)
- TSO (Time Sharing Option)

The Storage Toolkit also enables you to create and submit user-defined JCL streams.

Storage Toolkit menus and dialog boxes

This section describes the actions that are available in the context menus of the Storage Toolkit.

To access the actions, you select and right-click rows in the workspaces that have action menus. A context menu is displayed where you now select an item such as **Dataset Actions**. Then you choose the appropriate action from the submenu. When you select an action, a dialog box is displayed where you configure the action and submit the request.

The items on the context menu vary depending on the selected resources, however the following items are always available:

- Create Command
- Create Batch Job
- Submit Command or Job

These items do not have submenus. When you select them, you proceed directly to the corresponding dialog box.

Some actions can process multiple resources in a single request. In other words, you can run a single action against multiple resources. For example, to migrate multiple data sets you can proceed as follows:

1. Select the target resources in the Tivoli Enterprise Portal.
2. Right-click and select **Dataset Actions > Migrate**.
3. Configure the command.
4. Submit the request. The **Migrate** command is issued against each resource.

When you select an action, the Storage Toolkit launches a dialog box where you configure the action. The dialog boxes contain a set of standard tabs, which are described in “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258.

Custom commands and jobs

| Name of menu item | Corresponding dialog box |
|-----------------------|--|
| Create Batch Job | "Create Batch Job dialog box" on page 265 |
| Create Command | "Issue Command dialog box" on page 268 |
| Submit Command or Job | "Submit Command or Job dialog box" on page 269 |

Actions related to volumes that are displayed in workspaces of the Tivoli Enterprise Portal

| Name of menu item | Corresponding dialog box |
|-------------------|---|
| Backup | "Volume Backup dialog box" on page 271 |
| Convert | "Volume Conversion dialog box" on page 271 |
| Migrate | "Volume Migration dialog box" on page 271 |
| Release Space | "Volume Space Release dialog box" on page 271 |
| Compress | "Compress dialog box" on page 272 |
| Defrag | "Defrag dialog box" on page 272 |

Actions related to data sets that are displayed in workspaces of the Tivoli Enterprise Portal

| Name of menu item | Corresponding dialog box |
|-------------------|---|
| Backup | "Dataset Backup dialog box" on page 273 |
| Migrate | "Dataset Migration dialog box" on page 273 |
| Move & Copy | "Dataset Copy and Move dialog box" on page 274 |
| Recall | "Dataset Recall dialog box" on page 274 |
| Recover | "Dataset Recover dialog box" on page 274 |
| Release Space | "Dataset Space Release dialog box" on page 274 |
| Compress | "Compress dialog box" on page 275 |
| Print | "Print dialog box" on page 283 |
| Allocate | "Dataset Allocate dialog box" on page 275 |
| Rename | "Dataset Rename dialog box" on page 284 |
| Alter | "Dataset Alter dialog box" on page 276 |
| Delete | "Dataset Delete dialog box" on page 281 |
| Listcat | "Dataset Listcat dialog box" on page 281 |
| Repro | "Dataset Repro dialog box" on page 284 |
| Verify | "Dataset Verify dialog box" on page 286 |
| Catalog | "Dataset Catalog dialog box" on page 280 |
| Uncatalog | "Dataset Uncatalog dialog box" on page 286 |
| Define Cluster | "Dataset Define Cluster dialog box" on page 277 |

Actions related to DFSMSHsm

"Cancel HSM Request dialog box" on page 286. The menu item that corresponds to this dialog box is available when all the selected items are either queued or active. If the state of any selected item is **complete**, **abended**, or **cancelled**, the menu is not available.

Actions related to DFSMSrmm

| Name of menu item | Corresponding dialog box |
|--|--|
| Actions related to DFSMSrmm-related data sets | |
| Add Dataset | "Add Dataset dialog box" on page 288 |
| Change Dataset | "Change Dataset dialog box" on page 288 |
| Delete Dataset | "Delete Dataset dialog box" on page 288 |
| Actions related to DFSMSrmm-related owners | |
| Add Owner | "Add Owner dialog box" on page 289 |
| Change Owner | "Change Owner dialog box" on page 289 |
| Delete Owner | "Delete Owner dialog box" on page 289 |
| Actions related to DFSMSrmm-related volumes | |
| Add Volume | "Add Volume dialog box" on page 289 |
| Change Volume | "Change Volume dialog box" on page 290 |
| Confirm Actions | "Confirm Actions for All Volumes dialog box" on page 290 |
| Delete Volume | "Delete Volume dialog box" on page 290 |
| Actions related to DFSMSrmm-related VRS resources | |
| Add VRS | "Add VRS dialog box" on page 290 |
| Change VRS | "Change VRS dialog box" on page 291 |
| Delete VRS | "Delete VRS dialog box" on page 291 |

The documentation for the z/OS operating system includes reference information for the commands and parameters that the Storage Toolkit uses. You can access documentation for the version of z/OS that you use through the following Web page: <http://www-03.ibm.com/systems/z/os/zos/>

Related information

"Using the Storage Toolkit" on page 254

How the Storage Toolkit processes action requests

This section briefly describes how the Tivoli OMEGAMON XE for Storage on z/OS monitoring agent handles an action request.

Note: You can use the Storage Toolkit to target resources in a separate z/OS environment. However, this activity is valid only when that separate environment is located on DASD that is shared between the separate z/OS environments.

All actions, not only the actions that submit user-defined JCL streams, run as batch jobs on the z/OS system. Each type of command has a batch job interface which the Storage Toolkit uses to execute the action. The JCL for each command type is defined by the Configuration Tool and is subsequently used to run the command. The Configuration Tool is also used to define a default JOB card which is applied to each batch job, but can be overridden with a custom JOB card that you provide on the Storage Toolkit dialog box.

When you submit an action request, the Storage Toolkit starts a thread on z/OS to handle the execution of the request. The thread generates the JCL, substituting values for variables and expanding)D0T)ENDD0T clauses, submits the batch job for execution using the MVS Internal Reader, and waits for its completion. At the beginning of each JCL stream, the Storage Toolkit inserts a step that executes code to monitor the rest of the steps in the job. As each step completes, this code collects command output or, in

the case when your request is to run user-defined JCL, it collects the contents of corresponding data sets and DD names that you specified be copied for later viewing. DD names that are routed to a SYSOUT class are not collected until the final job step.

At the end of each JCL stream, the Storage Toolkit adds a final step that also executes Storage Toolkit code. The code in the final step collects SYSOUT from each step where it was requested, collects the return code of the batch job, and optionally, collects the JES output. Additionally, it notifies the Storage Toolkit thread that the batch job is complete. The thread then stores the results of the execution into the checkpoint database and terminates.

To accomplish this, the Storage Toolkit thread creates temporary data sets to contain the JCL, the results of the batch job and other files. The data sets are transient and are shared between the thread and the batch job. When the batch job completes, the data sets are deleted.

Activities required to process your request are executed under a security environment that the Storage Toolkit establishes using your Tivoli Enterprise Portal user ID. Because the security environment is established on the z/OS system where the monitoring agent runs, your portal user ID must match your mainframe ID on that system. If this ID is not a valid mainframe user ID, a security environment cannot be established and the action request fails.

The Storage Toolkit performs the following operations within the context of your security environment:

- Create and access all temporary data sets required by the batch job.
- Submit the batch job.

When your action request runs user-defined JCL, the Storage Toolkit also uses your security context to access the JCL data set and any other data sets you identify as containing substitution variables. If you are not authorized to the data sets, the action request fails.

Note: To ensure the security of Storage Toolkit commands, some type of security authentication must govern the users who log on to the Tivoli Enterprise Portal. You set up authentication on the hub monitoring server in the IBM Tivoli Monitoring environment. For further information, see “Authentication and authorization for running Storage Toolkit commands” on page 261.

Using the Storage Toolkit

This section describes the general process of using the Storage Toolkit and a typical scenario for use of the toolkit.

Scenario for using the Storage Toolkit

The Storage Toolkit enables the following scenario:

| | |
|-------------------|---|
| Phase One: | You create a situation that monitors available free space in your storage groups. Whenever the situation is triggered, you want to find the affected resources and take action. Specifically, you want to find and take action on the data sets with the most unused space. |
|-------------------|---|

| | |
|-------------------|--|
| Phase Two: | <p>You work in the Tivoli Enterprise Portal, instead of accessing the z/OS system and performing a manual process on each resource:</p> <ol style="list-style-type: none"> 1. Navigate to an appropriate workspace for the product, such as the Largest Datasets in SMS Class, and sort on Unused Tracks. 2. Select data sets that can be migrated. Note: In addition to the migration command, the Storage Toolkit offers a range of commands that is listed in the “Storage Toolkit menus and dialog boxes” on page 251 topic. You can also submit user-defined JCL streams, and you can issue mainframe commands that are not already available in the dialog boxes that the product provides. 3. Right-click the selected data sets and choose Dataset Actions > Migrate in the pop-up menu. 4. Use the Storage Toolkit dialog boxes (accessed through a right-click menu) to issue DFSMSHsm migrate commands for all selected data sets. The DFSMSHsm migrate command you selected from a single dialog box is issued against each data set that you selected. 5. View the results of the migration commands in one of the Storage Toolkit workspaces. |
|-------------------|--|

Process for using actions in the Storage Toolkit

The following steps describe the general process of using the actions provided by the Storage Toolkit. The options that are displayed in the context menus change, depending on what resources you select in a workspace view in the Tivoli Enterprise Portal. For example, if you have selected data sets in a workspace view, you see the **Dataset Actions** option in the context menu.

1. Select one or more rows from a view in a workspace. The rows correspond to resources such as volumes and data sets on z/OS systems.
2. Right-click within the rows. A context menu is displayed where you can select an appropriate option. The following example steps describe how to invoke the appropriate Storage Toolkit dialog box when you select *data sets* in a workspace view and right-click within the rows:
 - a. Select **Dataset Actions** from the context menu to access the submenu.
 - b. Select an appropriate data set action in the submenu, such as **Backup**. The corresponding dialog box is displayed.

The menu displays different options if you are working with other types of resources, such as volumes. In most cases, the action can simultaneously target multiple resources that you select in Step 1.

3. Make selections in the dialog box to configure the command that you want to issue. These selections are automatically transformed to command syntax that is valid for the mainframe environment. See “Storage Toolkit menus and dialog boxes” on page 251 for links to descriptions of all available dialog boxes.

Reusing action requests

The Storage Toolkit provides three methods for you to reuse an action. Consider whether you want to rerun the action using the same or different resources before you choose the method.

When you reuse an action, the new execution is processed under your security context on the z/OS system. In the following descriptions, it is assumed that you are authorized to the resources and command associated with the reuse action.

Method 1: Reusing an existing action using the same resources

1. Navigate to the Storage Toolkit Action Requests workspace (available from the Storage Toolkit Navigator node).
2. Right-click the action in the Action Requests view that you want to reuse, and select **Submit Request** from the context menu. The dialog box corresponding to the selected action is displayed.
3. Modify the settings that control how the request runs, as appropriate.

4. Click **OK**. The selected action request is submitted on the z/OS system to which you are connected and the same set of resources is affected by the action.

Method 2: Reusing an existing action using the same or different resources

1. Select one or more rows from a view in a workspace.
2. Right-click within the rows, and select **Submit Command or Job** from the context menu. From the Submit Command or Job dialog box, select the action you want to reuse. The dialog box corresponding to the selected action is displayed.
3. Modify settings as appropriate. If you want to reuse the action using the rows you selected in step 1, ensure that the **Use request's default data** option in the General tab is not selected. If you want to reuse the action using the resources that were used when the action initially ran, select the **Use request's default data** option.
4. Click **OK**. The action request is submitted on the z/OS system to which you are connected. The resources that are affected depends on the setting of **Use request's default data**.

Method 3: Reusing a specific execution of an action

When you rerun a specific execution of an action, the resources that are used are the same as those associated with the selected execution.

1. Navigate to the Storage Toolkit Action Requests workspace (available through the Storage Toolkit Navigator node).
2. Find the action that you want to reuse, and link to the Storage Toolkit Result Summary workspace.
3. Right-click the specific result in the Action Results Summary view that you want to reuse, and select **Submit Request** from the context menu. The dialog box that corresponds to the selected action is displayed.
4. Modify settings as appropriate.
5. Click **OK**. The selected action request is submitted on the z/OS system to which you are connected. The same set of resources used by the selected execution are affected by the action.

Action requests in workspaces that you create

If you create a workspace view, the Storage Toolkit action menus are available in the new view when the following conditions are true:

- Your workspace uses one of the queries that come with the product.
- AND
- The query originally had a Storage Toolkit action menu associated with it. These menu options are listed in "Storage Toolkit menus and dialog boxes" on page 251.

About results for action requests

The result for action requests can contain output from one or more steps or locations. For this reason, the **Step Name** and **Dataset or DD Name** attributes identify the source of the details.

- When you run a command from a command dialog or issue a command from the Issue Command dialog, the values for the **Step Name** and **Dataset or DD Name** are predefined by the Storage Toolkit.

| Command type | Step Name | Dataset or DD Name |
|--------------|-----------|--------------------|
| DFSMSdss | KS3DSS1 | SYSPRINT |
| DFSMSHsm | KS3HSM1 | KS3OUT |
| IDCAMS | KS3AMS1 | SYSPRINT |
| ICKDSF | KS3DSF1 | SYSPRINT |
| DFSMSrmm | KS3RMM1 | KS3OUT |
| TSO | KS3TSO1 | KS3OUT |

- When you create JCL using the Create Batch Job dialog box, the values for the **Step Name** and **Dataset or DD Name** reflect the values that you specified in the **Files whose contents should be copied for later viewing** fields on the **Options** tab.
- When you select **Copy JCL and JES logs for later viewing** on the **JCL** tab, the **Step Name** values are predefined by the Storage Toolkit. The **Dataset or DD Name** values are blank.

| Step Name | Description |
|------------|-----------------------------------|
| *JCL | The JCL that was submitted |
| *JESMSGLOG | The contents of the JESMSGLOG log |
| *JESJCL | The contents of the JESJCL log |
| *JESYSMSG | The contents of the JESYSMSG log |

No output from Mainframe commands is returned, because execution of the command is not synchronized with the execution of the batch job. Other reasons why you might not receive output include the following:

- The command that you are running might not return output.
- A file or log to be copied for later viewing is empty.
- You specified 0 as the value for the **Maximum output lines** option.
- The output is not available or the Storage Toolkit cannot capture it.

In some cases, you might find that command output can be captured using a different technique to run the command. For example, if you use the Issue Command dialog box to run the DFSMSHsm **LIST** command, the Storage Toolkit is not able to capture the output. However, you can write JCL that uses HSEND CMD WAIT to issue the **LIST** command and route the output to a data set. Using the Create Batch Job dialog box, you can then run this JCL and request the data set be copied for later viewing.

Checking for the completion of action requests

You might need to check the status of the toolkit action request between the time you click **OK** in the toolkit dialog box and the time that the job completes in the mainframe environment. In some cases, it can be unclear whether a job is taking a long time to run or has failed. The status of the job might be **EXECUTING** in both cases. You can respond to these possible conditions as described in this scenario:

| Conditions of the scenario, in sequential order | Your response options |
|---|---|
| You notice a request on the Tivoli Enterprise Portal that remains in EXECUTING state without completing. | Check the status of the batch job in the z/OS system. |
| Problem scenario 1: The batch job has failed due to a JCL error. | Cancel the request on the Tivoli Enterprise Portal. This action releases the thread and the resources that have been invoked. <i>(Optional)</i> Attempt to fix the JCL error and run the request again. |
| Problem scenario 2: The batch job requires more time to complete. | <ul style="list-style-type: none"> • Consider raising the priority of the job. • Consider adjusting the times when the job is allowed to run, if the job has been scheduled to run at off-peak hours. |

Related information

“Overview of common action commands” on page 291

“Storage Toolkit menus and dialog boxes” on page 251

This section describes the actions that are available in the context menus of the Storage Toolkit.

“Standard tabs in the dialog boxes of the Storage Toolkit”

The dialog boxes of the Storage Toolkit have similar layouts. This topic describes the tabs that are standard in most of the dialog boxes.

“Create Batch Job dialog box” on page 265

You use this dialog box to submit user-defined JCL. This action request is a more flexible equivalent to the toolkit commands that the product provides.

“Storage Toolkit limitations and hints” on page 292

“Storage Toolkit Result Detail” on page 224

Standard tabs in the dialog boxes of the Storage Toolkit

The dialog boxes of the Storage Toolkit have similar layouts. This topic describes the tabs that are standard in most of the dialog boxes.

To access the dialog boxes, you select and right-click one or more rows in the workspaces that have action menus. Then you choose one of the Storage Toolkit options in the pop-up menu to access the corresponding dialog box.

Related information

“Using the Storage Toolkit” on page 254

Options tab

Use this tab to choose configuration settings for the action that you want to run. The **Options** tab provides a subset of the available command parameters. Table 19 describes where to go for complete documentation of the available command parameters.

Table 19. Command reference guides

| |
|---|
| The documentation for the z/OS operating system includes reference information for the commands and parameters that the Storage Toolkit uses. You can access documentation for the version of z/OS that you use through the following Web page: http://www-03.ibm.com/systems/z/os/zos/ |
|---|

The options that you select on the **Options** tab are rendered as parameters on the command. You can view the raw text of the command including your options on the **Command** tab.

Note: If you make any modifications in the **Command** tab, further modifications in the **Options** tab are disabled. If you want to use the **Options** tab again, click **Cancel** to dismiss the dialog box, and begin constructing a new command in a new invocation of the dialog box.

Refer to “Setting substitution variables for toolkit commands” on page 262 for a scenario that describes how the **Options** tab and the **Command** tab interrelate.

The following exceptions exist regarding the **Options** tab of the Storage Toolkit:

- The **Options** tab for the Create Batch Job dialog box enables you to define substitution variables. You can click the **Show Data** button to access the Data View dialog box, where you can review the scope of the variables that you specify.
- The **Options** tab is not provided in the following dialog boxes, because they have no options to set: Dataset Space Release dialog box, Dataset Verify dialog box, Dataset Uncatalog dialog box, and Cancel HSM Request dialog box.
- The **Options** tab is not provided in the Issue Command dialog box, because you specify the command and its parameters as raw text in the **Command** tab of that dialog box.

General tab

Use the General tab to set name and description values for the action. Also use this tab to set the following options:

- **Run Options:**
 - **Run now, if no schedule:**
 - Select this option to cause the action request to run when you click **OK** in the dialog box. Keep in mind that any schedule that you specify in the Schedule tab takes precedence over the **Run now, if no schedule** option.
 - To save a Storage Toolkit action for later use, clear the **Run now, if no schedule** check box. The toolkit action request is saved, but it does not run. You can run the action later using the **Submit Request** or the **Submit Command or Job** context menu items. See “Reusing action requests” on page 255 for related information.
 - **Delete request and results after one run:** Select this option to cause deletion of the action request and results when the action request completes. If you do not need to review the results of a request and only intend to run the request one time, you might want to use this option. For example, you might want to delete the request and results for an MVS operator command that you need to run only one time. If you do not select this option, the request and results, persist until you delete them, as described in “Managing actions” on page 292.
 - **Use request’s default data:** When you reuse an existing action, select this option to resubmit it using the same resources that were used when the action was initially created. See “Reusing action requests” on page 255 for additional details.

- **Results:**

Note: Use the dialog box options that are listed here to manage the results of requests. You must manage the buildup of old results and unused requests. Otherwise, this data can build up and actually cause the Toolkit Checkpoint Database to run out of space.

- **Delete after (days):** This value determines the number of days after the toolkit action request completes that the results are automatically deleted. If the value is 0, no results are saved. If you do not specify a value, the results persist until you delete them.
- **Maximum output lines:** This value determines the number of lines that are saved from the execution of the action request. If the value is 0, no output lines are saved. If this field is blank, all of the output lines are saved, up to a maximum defined by configuration options that are chosen for your production environment. If you specify a value other than 0, the number of output lines that are saved may also be reduced based on the configuration options. See “Storage Toolkit limitations and hints” on page 292 for additional information.

Scheduling tab

Use this tab to schedule the action in one of the following ways:

Note: Time is in 12-hour format, so you must specify morning (**am**) or afternoon/night (**pm**).

- *At* a specific time and date OR
- *At* a specific time and days of the week. (Select the **Every** check box to enable the action to run every week on the schedule you specify.) OR
- *In* a specific time frame, when a specific amount of time has passed.

For details regarding the syntax of scheduling specifications, see the description of the **Schedule** attribute in the section “Storage Toolkit Requests” on page 222.

JCL tab

Use this tab to specify changes to JCL (job control language) characteristics as follows:

- **Replacement JCL JOB card:** Use this text box to define a new JCL JOB card specification to override the installation-specific JOB card. This job card can include substitution variables.
- **Extra JCL statements:** Use this text box to specify JCL statements to append to the JCL stream. For example, you can add DD statements to supplement a command in addition to control statements. This JCL can include substitution variables.

- **Dataset prefix for temp files:** Use this field to specify the prefix to use for the names of the temporary data sets that the Storage Toolkit creates to process the action request. The Storage Toolkit assigns unique names to the temporary data sets using this prefix as the high level qualifier. The Storage Toolkit accesses the temporary data sets under your security context. Therefore your user ID must have authority to create, read, write, and delete data sets whose names begin with the prefix. If you do not specify a value, the default is your user ID.
- **Copy JCL and JES logs for later viewing:** Select this option to save a copy of the submitted JCL as well as the various JES files (**JESMSG LG**, **JESJCL**, and **JESYSMSG**) that are produced when the batch job runs. You can view these files in the **Action Results Details** view of the Storage Toolkit Result Detail workspace.

Command tab

Use this tab to review and modify the command that you construct using the Storage Toolkit.

For most dialog boxes of the toolkit, this tab displays a command that summarizes the selections that you make in the **Options** tab. However, there are two exceptions from the standard behavior of the **Command** tab:

- When you use the “Issue Command dialog box” on page 268 (by selecting the **Create Command** option in the pop-up menu of a product workspace) the **Command** tab is blank, ready for you to manually enter all the specifications for a new command.
- When you use the “Create Batch Job dialog box” on page 265, the **Command** tab is not present. The functionality to run user-defined JCL is provided in the **Options** tab of that dialog box.

1. The **Command type** drop-down list displays the command type that corresponds to the type of toolkit action that you are constructing:

- DFSMSdss (Device Support Services)
- DFSMSHsm (Hierarchical Storage Manager)
- ICKDSF (Device Support Facilities)
- IDCAMS (Access Method Services)
- DFSMSrmm (Removable Media Manager)
- Mainframe (z/OS console command)
- TSO (Time Sharing Option)

The type of command is selected automatically while you are using any of the dialog boxes of the toolkit. (The exception is the Issue Command dialog box, where the command type field can be configured. By default, **DFSMSdss (Device Support Services)** is selected, but you can select any of the command types in the drop-down list.)

2. Experienced users can modify the command and parameters that are displayed in the **Command** text box.

Note: If you make any modifications in the Command tab, further modifications in the **Options** tab are disabled. If you want to use the **Options** tab again, click **Cancel** to dismiss the dialog box, and begin constructing a new command in a new invocation of the dialog box.

3. Define substitution variables, as needed, in the **Substitution variables and their run-time replacement values** table. The following guidelines apply to the use of substitution variables:

- Tivoli OMEGAMON XE for Storage on z/OS uses the percent sign (%) for the default substitution variables that it provides.
- The product provides generic attribute names, such as *DSN and *ROWNUMBER.

“Setting substitution variables for toolkit commands” on page 262 includes a description of how the **Options** tab and the **Command** tab interrelate.

Related information

“Using substitution variables in the Storage Toolkit” on page 262

This section explains how to use substitution variables with the Storage Toolkit and provides examples.

Authentication and authorization for running Storage Toolkit commands

The user ID for the Tivoli Enterprise Portal is fundamental to Storage Toolkit functionality. The following principles determine the success of Storage Toolkit operations:

- The user ID for the Tivoli Enterprise Portal enables toolkit activity in the mainframe environment.
- Authorization is achieved when the portal user ID matches a valid user ID in the z/OS system.
- You ensure the security of the z/OS environment by enabling user authentication on the hub monitoring server.

Keep in mind the following details regarding the portal user ID:

- The following list provides examples of the types of authentication that you could use to secure the use of the Storage Toolkit:
 - For a z/OS hub monitoring server, RACF, ACF2, Top-Secret, or equivalent hub-server authentication.
 - For a distributed hub monitoring server, user IDs can be authenticated either by the local system registry or by an external LDAP-enabled central registry.
- The Storage Toolkit does not have access to lists of valid user IDs in the mainframe environment. The toolkit depends on you to use a portal ID that matches a valid ID in the z/OS environment.
- For UNIX-based monitoring environments, you must ensure that the user ID exactly matches the user ID in the z/OS environment, including case sensitivity. For example, if the user ID in the z/OS system is USER1 (and not user1), you must log in to UNIX-based Tivoli Enterprise Portal with a USER1 user account. If you log in to the Tivoli Enterprise Portal with a user1 user account, Storage Toolkit actions would fail because user1 is not equivalent to USER1 on the z/OS system.
- You have the option to use various methods to resubmit actions that you defined previously. Authorization of a resubmission depends on the currently active user ID. (Resubmission does not depend on the ID used when the action was originally run.)
- To enable use of the temp files that are configured in the **JCL** tab, see the guidelines that are provided in the **JCL** tab section. The names of the target data sets start with the data set prefix that is specified for temp files. Keep in mind that the user ID must have the required authority to work with these data sets in the mainframe environment.

Data View dialog box

You use this dialog box to inspect the actual data that will substitute the substitution variables that you define in the **Command** or **Options** tabs.

You access this dialog box by clicking the **Show Data** button in the **Command** tab (or in the **Options** tab, when you are using the Create Batch Job dialog box). The **Command** tab is one of standard tabs that is described in “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. If you see inappropriate data in the Data View dialog box, respond as follows:

- If the substitution variables that you defined generate results that are not valid, click **Cancel** to return to the **Command** tab, and adjust your definitions.
- If resources are referenced in the dialog box that you did not intend to include, close the Data View dialog box, and click **Cancel** in the toolkit dialog box that you are configuring. Select the resources that you want to target, and access the toolkit dialog box again.

Note: If you select the **Use request’s default data** option in the **General** tab, the substitution uses the data from the original request. The data that is displayed in the Data View dialog box will reflect this fact.

The general process of setting substitution variables and evaluating their scope is described in “Setting substitution variables for toolkit commands.” Column headings in the Data View dialog box display the *ATTRIBUTE(%variable%)* format, where the attribute name is followed by the corresponding substitution variable name in parenthesis, as in the following example headings:

- ***DSN (%SUB01%)**
- **Managed System (%name%)**

In these example headings, ***DSN** and **Managed System** are the attribute names and **%SUB01%** and **%name%** are the names of the substitution variables.

Using substitution variables in the Storage Toolkit

This section explains how to use substitution variables with the Storage Toolkit and provides examples.

You can use substitution variables to indicate portions of a Storage Toolkit command or batch job that require substitution with volume names, dataset names, or other values. “Storage Toolkit limitations and hints” on page 292 provides helpful information about using substitution variables.

Related information

“Command tab” on page 260

“Create Batch Job dialog box” on page 265

You use this dialog box to submit user-defined JCL. This action request is a more flexible equivalent to the toolkit commands that the product provides.

Setting substitution variables for toolkit commands

Many of the action requests that the Storage Toolkit generates can include substitution variables. You can define these variables in the **Command** tab in Step 6 on page 263 as part of the following overall suggested procedure:

Note: A **Command** tab is not present in the Create Batch Job dialog box. In that dialog box, you set substitution variables in the **Options** tab.

1. Right-click rows in one of the views of a workspace to access the pop-up menu.
2. Select the Storage Toolkit command that you want to invoke. The **Options** tab is selected by default in the dialog box.

Note: An **Options** tab is not provided in the Issue Command dialog box, because you use this dialog box to create arbitrary free-form command and parameter strings. If you are using this dialog box, go to Step 6 on page 263

3. Make selections and enter values in the **Options** tab, as appropriate.
4. Select the **Command** tab. The **Command** field shows how the options that you selected in the **Options** tab are rendered in raw text within a command.
5. Before you make any changes in the **Command** tab, inspect the command parameters and syntax that are displayed to ensure that they match your expectation. If necessary, modify the command using either of the following approaches:
 - *Before you make any changes in the **Command** tab*, return to the **Options** tab and modify your selections.

Note: If you make changes in the **Command** tab, the **Options** tab is disabled, and you cannot modify the selections that you made there. To be able to work with the **Options** tab again, you can open a new instance of the toolkit dialog box. Click **Cancel** to dismiss the current instance of the dialog box, and access the dialog box again through the Tivoli Enterprise Portal.

- If you are familiar with command syntax, edit the raw text that you see in the **Command** field of the **Command** tab. Keep in mind that after you make changes in the **Command** tab, you cannot make further modifications to the **Options** tab.

- If you want to define substitution variables, access the **Command** tab and define them in the **Substitution variables and their runtime replacement values** area. In many cases, the substitution variables and their runtime replacement values table are primed with variables and values defined by the Storage Toolkit based on the action and resources that you selected.

Note: A **Command** tab is not present in the Create Batch Job dialog box. In that dialog box, you set substitution variables in the **Options** tab.

- Click the **Show Data** button to access the “Data View dialog box” on page 261 to confirm that the correct volumes or data sets are targeted by the substitution variables.

Standard substitution variables in the toolkit

The following general guidelines apply to the substitution variables that the toolkit uses:

- Tivoli OMEGAMON XE for Storage on z/OS uses a pair of percent signs (%) surrounding the *SUBnn* name to identify the substitution variables that it provides, such as %SUB01%.
- When you create variables for JCL files, it is recommended that you use the percent-sign (%) convention, as in this example: %MY_VAR%. This convention enables the Storage Toolkit to find these variables and list them automatically in the “Create Batch Job dialog box” on page 265. Variables that do not follow this convention are not listed automatically in the dialog box.
- Do not create a variable that could be a subset of some other variable. For example, do not name a variable **MY_VARNAME** and another variable **MY_VAR**.
- Do not create a variable that could match data in your command, JCL files, or generated JCL that should not be replaced. For example, in the following scenario a value might be replaced that should not be replaced:

- You define an **ADDR** variable and assigned to it a value of **01AE**.

- You use the variable in the following command:

```
ANALYZE UNITADDRESS(ADDR) SCAN
```

- Automatic substitution of the **ADDR** variable would generate the following command. The generated command includes a modification to the **UNITADDRESS** parameter that might have effects that you do not intend:

```
ANALYZE UNIT01AEESS(01AE) SCAN
```

DISP is another example of a variable name that is not recommended because that string is likely to appear on a JCL DD statement, where it should not be replaced.

- The product provides generic attribute names, such as ***DSN** and ***ROWNUMBER**.

To view the scope of substitution variables that you define in the **Options** tab, click the **Show Data** button.

How commands run when substitution variables are defined

You specify substitution variables and their mappings in the Substitution variables and their run-time replacement values area of the dialog box. Depending on your settings, the action request issues the command one or more times as follows:

| Number of rows selected in the workspace | Are substitution variables defined in the Command tab? | How the command runs |
|--|--|--|
| One | Yes or no | One time for the selected item. |
| Multiple | Yes | Multiple times; one time for each selected item. |

Examples of substitution variables

The following example table shows typical substitution variables for a command that is issued by the Storage Toolkit. A similar, two-column table is provided in the dialog boxes of the toolkit.

| Command Variable ¹ | Attribute or String ² |
|-------------------------------|----------------------------------|
| %dsname% | ENTRYNAME |
| my-volser | VOLUME |
| %job-name% | "USER01A" |
| %sysout-class% | "X" |
| %rownum% | *ROWNUMBER |
| %user-account% | *USERID |

1. The heading for this column is "**Variable**" when you are defining substitution variables for a batch job in the Create Batch Job dialog box. See "About the "Command Variable" (or "Variable") column" for more information.
2. You can enter raw text or use a drop-down menu to select the entries in the Attribute or string column. See "About the "Attribute or String" column" for more information.

Related information

"Create Batch Job dialog box" on page 265

You use this dialog box to submit user-defined JCL. This action request is a more flexible equivalent to the toolkit commands that the product provides.

About the "Command Variable" (or "Variable") column

An example of the **Command Variable** (or **Variable**) column for substitution variables is shown in "Examples of substitution variables" on page 263. These examples reflect names of substitution variables that might be present in user-defined JCL. The examples illustrate the following guidelines and characteristics for entries in the **Command Variable** column:

- Strings can consist of up to 64 characters from the following character collection: [a-zA-Z0-9_%.~]. The strings are not case sensitive when processed by the mainframe agent.
- The ampersand character (&) is not supported for a substitution variable. Otherwise, there could be confusion with other use of ampersands within the JCL.
- Do not allow variable names to be a substring of another substitution variable. For example, do not use **my_dsn** and **my_dsname** as substitution variables.
- Do not create a variable that could match data in your command, JCL files, or generated JCL that should not be replaced. For example, **DISP** is not recommended as the name of a substitution variable, because that string is likely to appear on a JCL DD statement, where it should not be replaced.

About the "Attribute or String" column

An example of the **Attribute or String** column for substitution variables is shown in "Examples of substitution variables" on page 263. In most cases, the values for substitution variables are taken from columns of data in the workspace where the toolkit action request originates. You use the drop-down list in the **Attribute or String** column to select a value. This value is used as the source of values for the corresponding substitution variable.

The toolkit detects the entries that are available for the row that you select in a workspace prior to invoking a Storage Toolkit dialog box. "Examples of substitution variables" on page 263 illustrates the following guidelines and characteristics for entries in the Attribute or string column:

- **ENTRYNAME:** In this example, the user selected a table attribute, **ENTRYNAME**, from a list box, for the value of the variable.
- **USER01A:** In this example, the user typed a literal name (**USER01A**) to correspond to the **%job-name%** variable.
- **%job-name%**, **%sysout-class%:** You can enter literal text strings in this column, instead of attribute names. The raw text can reference items in the JCL that are not related to the underlying table, such as

a `%job-name%`, `%sysout-class%`, and so on. The toolkit treats a value that you enclose in single or double quotes entered as a literal string rather than the name of a table column. A single occurrence of this string is used in the batch job or command.

- ***ROWNUMBER:** In this example, the user selected the ***ROWNUMBER** value from the drop-down menu that corresponds to the `%rownum%` variable. As a result of this setting a variable number of rows is generated for use in the JCL. The toolkit generates integer numbers ranging from 1 through the total number of data rows being passed from the underlying table. For example, you might select three rows from a workspace view when you create or resubmit an action request. In this case, the values 1, 2, and 3 would be generated.
- ***USERID:** In this example, the user selected the ***USERID** value from the drop-down menu that corresponds to the `%user-account%` variable. This value is filled in with a single occurrence of the name of the user account that is currently active in the Tivoli Enterprise Portal.
- The following points apply to setting up batch jobs in the Create Batch Job dialog:
 - When the toolkit targets a member of a partitioned data set or a sequential dataset that already exists, the toolkit scans that file for strings that meet the following naming convention: `%name%`, where *name* is the name of a substitution variable. The toolkit populates the dialog box with each of the corresponding variables for `%name%`.
 - The substitution of values for variables within the JCL might produce lines that are longer than 72 characters. If this occurs, this execution of the action request fails. A status of **Invalid JCL** is displayed in the Result Summary workspace.
 - The substitution of values for variables within other data sets that require variable substitution might produce lines that are longer than 80 characters. If this occurs, this execution of the action request fails. A status of **InvalidJCL** is displayed in the Results Summary workspace.
 - When configuration of the batch job is complete, you click **OK** in the Create Batch Job dialog box. At this point, the JCL file that you provide and any files outside that JCL that contain substitution variables are copied into temporary files. These temporary files are used when the batch job JCL is submitted. The toolkit does not modify the original files, because other users might be referencing the files at the same time.

Custom actions in the Storage Toolkit

This topic provides links to information about dialog boxes that are related to custom actions that you can create in the Storage Toolkit.

| Name of menu item | Corresponding dialog box |
|-----------------------|--|
| Create Batch Job | “Create Batch Job dialog box” |
| Create Command | “Issue Command dialog box” on page 268 |
| Submit Command or Job | “Submit Command or Job dialog box” on page 269 |

Create Batch Job dialog box

You use this dialog box to submit user-defined JCL. This action request is a more flexible equivalent to the toolkit commands that the product provides.

The “**Options** tab” on page 258 of this dialog box provides the fields that are described below. “Storage Toolkit limitations and hints” on page 292 provides helpful information about using this dialog box.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

See “Conventions for JCL files that you use with the Storage Toolkit” on page 267. Also see “Authentication and authorization for running Storage Toolkit commands” on page 261 regarding the authorization that is required for the action requests of the Storage Toolkit.

- **Mainframe fully qualified dataset containing the batch JCL:** Specifies the data set that contains the batch job. The data set must be a sequential data set or a member of a partitioned data set (PDS). The logical record length of the data set must be 80 and the data set must be a fixed length format. To edit or create the data set, click the **Edit JCL** button, which causes the “Edit Data Set dialog box” on page 268 to be displayed.

Note: The user ID for the Tivoli Enterprise Portal that you use must be authorized in the mainframe system to access the data set. When the user ID is not authorized, you see an error message when you click **Edit JCL**. For more information, see “Authentication and authorization for running Storage Toolkit commands” on page 261.

- **Substitution variables and their run-time replacement values:** Specifies replacement values to be used at run time. The replacement values allow the batch job to reference specific objects, such as data sets and volumes. For example, these objects might correspond to values in the table rows that you select prior to invoking this dialog box.

In the **Variable** column, you can enter a unique name for each substitution variable that exists in the data set. In the cells of the **Attribute or string** column, you can select from the available names. These names reflect attributes from the table rows that you selected in a workspace, prior to invoking the Create Batch Job dialog box. Alternatively, you can manually enter values in the cells. For example, you might want to enter values such as job name, **sysout** class, and other values that might not be present in the workspace table from which you are selecting rows.

Tip: Use the percent-sign convention for the substitution variables (*%name%*) that are present in your JCL data set. That way, this dialog box automatically finds and lists the JCL substitution variables.

Refer to “Standard substitution variables in the toolkit” on page 263 for guidelines that apply to the use of substitution variables.

- **Files whose contents should be copied for later viewing:** Specifies the DD names and data sets to be preserved as a record of the results of the batch job. The files are available for viewing in the Storage Toolkit Result Details workspace.

In the **Step name** column, you specify the name of the step whose data you want to save. In the **Fully Qualified Dataset or DD name** column, you specify the sequential data sets, members of partitioned data sets, or DD names that you want to copy. The following table shows some sample settings for this area of the dialog box:

Table 20. Sample settings for “Files whose contents should be copied for later viewing”

| Step name | Dataset or DD name |
|-----------|--------------------|
| STEPX | USER02.LOG.LISTING |
| STEPY | SYSPRINT |

The specification in row one preserves data from the **USER02.LOG.LISTING** data set.

The specification in row two preserves the **SYSPRINT** file in **STEPY**. This output is saved in the checkpoint file. The following restrictions apply:

- Only sequential data sets or members of a PDS are supported.
- VSAM data sets are not supported.
- Data sets or files that contain unprintable characters (for example hexadecimal code) might display improperly in the Storage Toolkit Result Details workspace.
- Output lines longer than 256 characters are truncated.

- **Fully Qualified Datasets needed by the job that also contain substitution variables:** Specifies data sets that are located outside the JCL job stream. Such data sets might be needed to dynamically supply substitution variables that correspond to objects (such as data sets and volumes) that are needed when

the batch job runs. These additional data sets can be sequential data sets or members of partitioned data sets. The logical record length of these data sets must be 80 and these data sets must be a fixed length format.

For example, an IEHLIST job might use a SYSIN DD statement that points to a DASD data set rather than in-stream control cards. That data set could contain substitution variables that need replacement at execution time.

If no data sets that contain substitution variables are referenced from the JCL, this area of the dialog box can remain empty.

You might encounter the following return codes when the attempt to access data sets fails:

| Return code | Description |
|----------------------------|---|
| 119 (DataSetError) | The User Data Set server (UDS) has failed. |
| 121 (AuthorizationFailure) | You are attempting to access a data set for which you do not have RACF authority. |
| 123 (DataSetError) | The data set does not exist or the data set is not a supported format or record length. |

You might also receive a return code (status) of 117 (**InvalidJCL**). This status typically indicates that the JCL exceeds 72 characters when the substitution variables are applied. It might also indicate other JCL-related errors, such as a missing JOB card, or that another data set requiring variable substitution exceeds 80 characters when the substitution variables are applied.

Be sure to consult the RKLVLLOG for additional messages regarding these results.

When configuration of the batch job is complete, you click **OK** in the Create Batch Job dialog box. At this point, the JCL file that you provide and any files outside that JCL that contain substitution variables are copied into temporary files. These temporary files are used when the batch job JCL is submitted.

Related information

“Using substitution variables in the Storage Toolkit” on page 262

This section explains how to use substitution variables with the Storage Toolkit and provides examples.

“Examples of substitution variables” on page 263

“Using the Storage Toolkit” on page 254

Conventions for JCL files that you use with the Storage Toolkit

The data set for the JCL that you create or edit must be a sequential data set or a member of a partitioned data set (PDS). Any JCL data set that you use with the Storage Toolkit must follow conventions to enable the Tivoli Enterprise Portal to substitute data values at run time. The following elements (marked in bold text in Figure 1) are substitution variables that will be altered by the Storage Toolkit at job run time:

- %job-name% (see line 000100)
- %sysout% (see line 000100)
- %my-volser% (see line 000800)
- %dsname% (see line 000800)

Figure 1. Variable elements for JCL files used with the Storage Toolkit

```

***** ***** Top of Data *****
000100 //%job-name% JOB (),'DHODS',REGION=4M,MSGCLASS=%sysout%,NOTIFY=&SYSUID,
000200 //          CLASS=A
000300 //* -----
000400 //LISTVTOC EXEC PGM=IEHLIST
000500 //SYSPRINT DD SYSOUT=*
000600 //SYSIN DD *
000700 )DOT
000800 LISTVTOC DUMP,VOL=3390=%my-volser%,DSNAME=(%dsname%)
000900 )ENDDOT
001000 /*
001100 //
***** ***** Bottom of Data *****

```

Tip: Use the percent-sign convention for the substitution variables (%name%) that are present in your JCL data set. That way, this dialog box automatically finds and lists the JCL substitution variables.

The)DOT and)ENDDOT lines in Figure 1 on page 267 indicate the start and end of table processing. The JCL lines between the)DOT and)ENDDOT lines are repeated for each table row that is passed into the mainframe agent at JCL execution time. Typically, these lines contain substitution variables that are replaced with new values at run time. For example, the output for these SYSIN JCL lines might look similar to this excerpt:

```

000600 //SYSIN DD *
000700 LISTVTOC DUMP,VOL=3390=TS0201,DSNAME=(SYS2.USER.CLIST)
000800 LISTVTOC DUMP,VOL=3390=TS0202,DSNAME=(SYS2.USER.PROCLIB)
000900 LISTVTOC DUMP,VOL=3390=MVSA12,DSNAME=(SYS2.PROD.LOADLIB)
001000 /*

```

Edit Data Set dialog box

You use this dialog box to create, edit, and save JCL, which is displayed in a scrollable ASCII text box in the dialog box. The user ID that you use in the TEP must be a valid mainframe user ID that provides access to the mainframe JCL data set.

Click the **Edit JCL** button in the Create Batch Job dialog box to access the Edit data set dialog box. The name of the data set being edited appears in the title of the dialog box. If the data set that you specify does not exist, the data set is created when you click **OK**. This dialog box can contain ASCII text only.

Line and column indicators displayed below the ASCII text box identify the current position of the cursor. A ruler is also available to assist with column identification. By default, the ruler is hidden. When displayed, it appears above the ASCII text box. Use **Ctrl+R** to toggle the display of the ruler.

The Edit data set dialog box also displays your editing mode. The editing mode appears below the ASCII text box. REP indicates that you are typing in replace mode (overwriting existing characters). INS indicates that you are typing in insert mode. Press the Insert key to toggle between insert and replace modes. Replace mode is the default.

A vertical gray bar in the ASCII text box is positioned at the logical record length of the data set. If the data set does not exist, a length of 80 is used. You cannot type beyond this bar, nor can characters shift past the bar when you are in insert mode. When the line on which you are typing reaches the bar, additional characters that you type are ignored. Use the **Enter** key to create new lines.

Issue Command dialog box

You use this dialog box to issue arbitrary commands, including IDCAMS, ICKDSF, DFSMSrmm, DFSMSHsm, DFSMSdss, TSO, and mainframe console commands in text format.

Note: An MVS command that you issue through the Issue Command dialog box is submitted through SDSF. A forward slash (/) must precede MVS commands, as in this example, which cancels a time-sharing user (*tso_user_ID*):

```
/C U=tso_user_ID
```

Also see “Authentication and authorization for running Storage Toolkit commands” on page 261 regarding the authorization that is required for the action requests of the Storage Toolkit.

Typically you use this feature when you are familiar with the syntax and parameters of a command. You construct and issue a mainframe command directly from the Tivoli Enterprise Portal, without having to directly access the mainframe environment. See “Authentication and authorization for running Storage Toolkit commands” on page 261 for a description of the relevant requirements for your portal user ID.

You supply the command details in the following fields of the dialog box:

- Command Type drop-down list
- Command text box
- Substitution variables and their run-time replacement values table

For additional information, see the description of the generic **Command** tab.

In the generic **Command** tab you enter *supplemental* parameters to a basic command that you specify in a graphical user interface. In contrast, the **Command** tab of the Issue Command dialog box is used to enter a complete command in text form. Because you specify all options for the command in the text box, the Issue Command dialog box does not have an **Options** tab.

Command tab

Use this tab to construct the raw text of the command that you want to issue.

1. Select one of the following types of command in the **Command type** drop-down list:

- DFSMSdss (Device Support Services)
- DFSMSHsm (Hierarchical Storage Manager)
- ICKDSF (Device Support Facilities)
- IDCAMS (Access Method Services)
- DFSMSrmm (Removable Media Manager)
- Mainframe (z/OS console command)
- TSO (Time Sharing Option)

You must select the type of command (corresponding to the type of facility in the mainframe environment) that matches the command type that you enter in the Command text box. Otherwise, the command is not processed successfully in the mainframe environment.

2. Type the command and its parameters in the **Command** text box. Use substitution variables as needed.
3. Define substitution variables, as needed, in the **Substitution variables and their run-time replacement values** table.
4. Click the **Show Data** button to access the “Data View dialog box” on page 261, where you can review the scope of the substitution variables that you specify.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Submit Command or Job dialog box

You use this dialog box to reuse a command that you previously issued through the Storage Toolkit. The relevant types of commands include the following:

- A standard command like the **Migrate** command for data sets that invokes the “Dataset Migration dialog box” on page 273.
- A batch job generated in the “Create Batch Job dialog box” on page 265.
- A command generated in the “Issue Command dialog box” on page 268

You select data rows in a workspace, such as data set names or volume names, and right-click and select **Submit Command or Job** to invoke the dialog box. The Submit Command or Job dialog box is displayed and includes the following features:

- A table of data regarding commands that have previously run, including the ID of the creator, name and type of command, and the syntax of the command.
- The **Show only actions applicable to the selected data row(s)** check box. When you select this option, the table updates to display only the toolkit actions that are relevant to the selected data rows or commands that do not require data from the underlying tables. (These rows are the ones that you select prior to right-clicking and selecting **Submit Command or Job** in the pop-up menu.)

When you select a toolkit action and click **OK** in the Submit Command or Job dialog box, the relevant toolkit dialog box is displayed. For example, if you click on an action that migrates data sets and click **OK**, the “Dataset Migration dialog box” on page 273 is displayed. At the same time, the relevant values of the data rows that you selected in the portal workspace are captured for use in the dialog box. When you submit the action, all the values are combined into a single batch file. The current Tivoli Enterprise Portal user ID is used for this transaction. For more information, see “Authentication and authorization for running Storage Toolkit commands” on page 261.

See “Reusing action requests” on page 255 to learn how to resubmit a toolkit action through a Storage Toolkit workspace.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Volume actions in the Storage Toolkit

This topic provides links to information about dialog boxes that are related to volume actions in the Storage Toolkit.

You access the following dialog boxes in the **Volume Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

| Name of menu item | Corresponding dialog box |
|-------------------|---|
| Backup | “Volume Backup dialog box” on page 271 |
| Convert | “Volume Conversion dialog box” on page 271 |
| Migrate | “Volume Migration dialog box” on page 271 |
| Release Space | “Volume Space Release dialog box” on page 271 |
| Compress | “Compress dialog box” on page 272 |
| Defrag | “Defrag dialog box” on page 272 |

You can also create a custom action. See “Custom actions in the Storage Toolkit” on page 265.

Volume Backup dialog box

You use this dialog box to issue the DFSMSShsm BACKVOL command to back up all eligible data sets on the selected volumes. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The action supplies the following information to the service:

- Device Type (default device category, Type or DASD)
- Backup Type (incremental or full)
- Backup data sets not backed up in the last n days. (Range of n : 0-999)

For data sets not managed by SMS, you can target data sets that have not been backed up in the last n days, where n ranges from 0 to 999.

Volume Conversion dialog box

You use this dialog box to issue the DFSMSdss CONVERTV command to convert the selected volumes to or from system-managed storage (SMS). To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The following parameters are available:

- **Action:** Prepare, Test, or Convert.
- Convert multi-volume data sets
- Redetermine classes for SMS data sets

Volume Migration dialog box

You use this dialog box to issue the DFSMSShsm MIGRATE VOLUME command to migrate all eligible data sets on the selected volumes. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

For data sets not managed by SMS, you can target data sets that have not been referenced in the last n days, where n ranges from 0 to 999.

Volume Space Release dialog box

You use this dialog box to issue the DFSMSdss RELEASE command to release allocated but unused space from all data sets on the selected volumes. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The following parameters are available:

- Minimum secondary allocation for release (n number of tracks, where n ranges from 0 to 99999999)
- Minimum unused space prior to release (n number of tracks, where n ranges from 0 to 99999999)
- Wait time in seconds between retry attempts (n seconds, where n ranges from 0 to 255)
- Number of attempts allowed to obtain volume access (n tries, where n ranges from 0 to 99)
- A fully-qualified data set name or pattern for data sets that you want to exclude from the space release operation. Use the **Fully Qualified Dataset Name or Pattern** table to manually enter data set names or DSN patterns.

Compress dialog box

You use this dialog box to issue the DFSMSDss COMPRESS command to compress all eligible data sets on the volumes that you selected in the Tivoli Enterprise Portal. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The dialog box provides the following options:

- **Exclude datasets qualified by:** You can add a qualifier to exclude specific data sets on the volume. Enter a fully-qualified name or pattern that identifies the data sets to exclude.
- **Include datasets qualified by:** You can add a qualifier to include specific data sets on the volume. Enter a fully-qualified name or pattern that identifies the data sets to include.
- **Run as DFSMSDss Administrator:** Select this option to cause the command to run with Administrator authorization in DFSMSDss. The user ID that you are using must be authorized as a storage administrator in DFSMSDss.
- **Dynamic Allocation:** Select this option to cause the command to use dynamic allocation for the affected data sets.

Defrag dialog box

You use this dialog box to issue the DFSMSDss DEFrag command to defragment the volumes that you selected in the Tivoli Enterprise Portal. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The dialog box provides the following options:

- **Exclude Datasets qualified by:** You can add a qualifier to exclude data sets. Enter a fully-qualified name or pattern that identifies the data sets to exclude.
- **Fragmentation Index:** Specify an integer n . When fragmentation drops below the value of n , the defragmentation operation ends. The DFSMSDss facility places a decimal point before the value that you specify for n , as in these examples:
 - If you specify 1, the value becomes .1.
 - If you specify 999, the value becomes .999.
 - If you specify 001, the value becomes .001.
- **Fast Replication:** Use the drop-down menu to set your preference for use of the DFSMSDss fast replication feature. If the **Fast Replication** option is blank (the default), fast replication is not enabled.
- **Run as DFSMSDss Administrator:** Select this option to cause the command to run with Administrator authorization in DFSMSDss. The user ID that you are using must be authorized as a storage administrator in DFSMSDss.
- **Dynamic Allocation:** Select this option to cause the command to use dynamic allocation for the affected data sets.
- **Peer-To-Peer Remote Copy for FlashCopy®:** Select this option to specify that a Peer-to-Peer Remote Copy (PPRC) primary volume is allowed to become a FlashCopy target volume. This option is ignored when FlashCopy is not used for defragmentation. RACF authorization might be required for this option.
- **Consolidate:** Select this option to enable the command to consolidate multiple extent data sets when possible.
- **Write Check:** Select this option to enable verification for the data moved by defragmentation. Verification increases the time required for the defragmentation operation.

Data set actions in the Storage Toolkit

This topic provides links to information about dialog boxes that are related to data set actions for the Storage Toolkit.

You access the following dialog boxes in the **Dataset Actions** submenu. This submenu is available in the pop-up menu that you access by right-clicking specific views in a workspace.

| Name of menu item | Corresponding dialog box |
|-------------------|---|
| Backup | "Dataset Backup dialog box" |
| Migrate | "Dataset Migration dialog box" |
| Move & Copy | "Dataset Copy and Move dialog box" on page 274 |
| Recall | "Dataset Recall dialog box" on page 274 |
| Recover | "Dataset Recover dialog box" on page 274 |
| Release Space | "Dataset Space Release dialog box" on page 274 |
| Compress | "Compress dialog box" on page 275 |
| Print | "Print dialog box" on page 283 |
| Allocate | "Dataset Allocate dialog box" on page 275 |
| Rename | "Dataset Rename dialog box" on page 284 |
| Alter | "Dataset Alter dialog box" on page 276 |
| Delete | "Dataset Delete dialog box" on page 281 |
| Listcat | "Dataset Listcat dialog box" on page 281 |
| Repro | "Dataset Repro dialog box" on page 284 |
| Verify | "Dataset Verify dialog box" on page 286 |
| Catalog | "Dataset Catalog dialog box" on page 280 |
| Uncatalog | "Dataset Uncatalog dialog box" on page 286 |
| Define Cluster | "Dataset Define Cluster dialog box" on page 277 |

You can also create a custom action. See "Custom actions in the Storage Toolkit" on page 265.

Dataset Backup dialog box

You use this dialog box to issue the DFSMSHsm BACKDS command to create a backup version of specific data sets. The action that you define targets the data sets that you select. To learn about the standard capabilities of a dialog box, see "Standard tabs in the dialog boxes of the Storage Toolkit" on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

This works with the following information:

- Target for the backup operation: Default, Tape, or DASD.
- Concurrent copying: Standard, Preferred, or Required

Dataset Migration dialog box

You use this dialog box to issue the DFSMSHsm MIGRATE DATASETNAME command to migrate the selected data sets. To learn about the standard capabilities of a dialog box, see "Standard tabs in the dialog boxes of the Storage Toolkit" on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Convert (recall immediately to a different volume) is allowed. Place a check mark in the Convert check box to recall data sets to a different volume immediately.

Dataset Copy and Move dialog box

You use this dialog box to issue the DFSMSdss COPY command to copy or move the selected data sets to new locations. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The dialog box sets the following COPY command parameters for an action that starts a DFDSS service:

- For source data set the dialog box can define the following behavior:
 - Process multi-volume data sets
 - Process VSAM sphere components
 - Force unmovable data set
 - Disposition of the source data set: keep, delete, or uncatalog (Uncat)
- For the destination of the move/copy operations the dialog box can set the following parameters:
 - Volser name
 - Fully Qualified Dataset name
 - Name of the storage class
 - Name of the management class
 - Other options:
 - Bypass the SMS ACS routine, Catalog (default), and Reblock
 - Conversion to PDS, PDSE, or no conversion

Dataset Recall dialog box

You use this dialog box to issue the DFSMSShsm RECALL command to recall the selected migrated data sets. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The dialog box defines the DA option for a DFSMSShsm service to the default, RELBLK, RELTRK or SAMETRK. You also have the option to specify the target volume.

Dataset Recover dialog box

You use this dialog box to issue the DFSMSShsm RECOVER command to restore the selected data sets from backed up versions of those data sets. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The action targets a single data set to be recovered from a backup operation:

- A new, fully qualified name for the recovered data set
- The volume to which the recovered data set is sent
- The option to replace the data set
- The backup version to be used in restoring the data set

Dataset Space Release dialog box

You use this dialog box to issue the DFSMSdss RELEASE command to release allocated but unused space from the selected data sets. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

This dialog box does not have options that you can set.

Compress dialog box

You use this dialog box to issue the DFSMSdss COMPRESS command to compress the data sets that you selected in the Tivoli Enterprise Portal. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The dialog box provides the following options:

- **Run as DFSMSdss Administrator:** Select this option to cause the command to run with Administrator authorization in DFSMSdss. The user ID that you are using must be authorized as a storage administrator in DFSMSdss.
- **Dynamic Allocation:** Select this option to cause the command to use dynamic allocation for the affected data sets.

Dataset Allocate dialog box

You use this dialog box to issue an IDCAMS ALLOCATE request to create a data set. You can optionally model it based on the selected data set. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Table 21. IDCAMS ALLOCATE options that are available in the Allocate dialog box of the Storage Toolkit

| Option | Description |
|-----------------------------------|---|
| Fully Qualified Dataset name | Type the fully-qualified name of the data set that you want to allocate. |
| Model name | Type the name of a model data set that you want to use. |
| Use model to allocate new dataset | Select this option when you want to use the value of Model name in the allocation of the new data set. |
| Disposition | You have the option to select a disposition value in this drop-down menu. |
| Record format | You have the option to select a value for the format of the record in this drop-down menu. Options include fixed-format records, variable-format records, and undefined-format records. |
| Dataset organization | You have the option to select the data-set organization type in this drop-down menu. Options include DA (direct access); DAU (direct access unmovable); PO (partitioned organization); POU (partitioned organization unmovable); PS (physical sequential); and PSU (physical sequential unmovable). |
| Dataset type | You have the option to select a data-set type in this drop-down menu. Options include partitioned data set (PDS) and the Hierarchical File System (HFS) for z/OS UNIX®. |
| Data class | Type the name of the data class for either SMS or non-SMS-managed data sets. |
| Record length | Specify the length, in bytes, of the largest logical record in the data set. |
| Management class | Type the name of the management class for a new data set. When possible, do not specify this value. Allow the value to default through the ACS routines. |
| Block size | Specify the block size of the data control block (DCB) for the data set. |
| Storage class | <i>For SMS-managed data sets</i> you can type the storage class name. When possible, allow this value to default through the ACS routines established by the system administrator. |

Table 21. IDCAMS ALLOCATE options that are available in the Allocate dialog box of the Storage Toolkit (continued)

| Option | Description |
|------------------|--|
| Volser(s) | Choose one of the following values in the list: <ul style="list-style-type: none"> The serial number of an eligible direct access volume on which a new data set is to be located. OR <ul style="list-style-type: none"> The serial number of a direct access volume on which an old data set is located. If you use VOLUME for an old data set, the data set must be on the specified volume for allocation to take place. If you do not include VOLUME, the system allocates new data sets to any eligible direct access volume. |
| Unit | Enter the name of the unit type to which a file or data set is to be allocated. You can list an installation-defined group name, a generic device type, or a specific device address. |
| Directory blocks | Specify the number of 256-byte records for the directory of a new partitioned data set. Use this parameter to allocate a new partitioned data set. |
| Space Units | Specify the units and amount of space to allocate to a new data set. Options include the following combinations units and values: <ul style="list-style-type: none"> BLOCK: An integer value to be multiplied by the Primary value that you specify in the Space area of this dialog box. AVBLOCK: An integer value to be multiplied by the Primary value that you specify in the Space area of this dialog box. CYLINDERS: The Primary value that you specify in the Space area of this dialog box is given in cylinders. (No Value is required for this option.) TRACKS: The Primary value that you specify in the Space area of this dialog box is given in tracks. (No Value is required for this option.) |
| Release space | Select this option when you want to delete unused space when the data set is closed. |
| Space | Allocates the amount of space for a new data set. <p>The Primary value specifies the initial number of units of space that is allocated for a data set. For a partitioned data set, a directory quantity is not necessary.</p> <p>The Secondary value specifies the number of units of space to be added to the data set each time the previously allocated space has been filled. You must provide the primary quantity along with the secondary value.</p> |

Dataset Alter dialog box

You use this dialog box to issue an IDCAMS ALTER request to alter the data sets and catalogs that you select in the Tivoli Enterprise Portal. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The dialog box provides the following options:

- **Fully Qualified Name:** Specifies a new, fully-qualified name for the data set that you have selected. If you have selected multiple data sets, this field is disabled. Other fields in the dialog box continue to be available.
- **Catalog:** Specifies a new name for the catalog that contains the entry.
- **Management class:** For SMS-managed data sets only, this option specifies a new name for the management class for a data set. The name must come from a list of names that the system administrator defines.
- **Storage class:** For SMS-managed data sets only, this option specifies a new name for the storage class. The name must come from a list of names that the system administrator defines.

Dataset Define Cluster dialog box

You use this dialog to issue an IDCAMS DEFINE CLUSTER command to define attributes for a cluster as a whole and for the components of a cluster.

The **Options** tab of this dialog box provides the following options in three tabbed subsections: the **Cluster** tab, the **Data** tab, and the **Index** tab. These tabs correspond to the cluster that you define and its data and index components. You specify many of the same attributes for each of the three entities (cluster, and its data and index components). The values that you specify are unique to each entity.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Table 22. IDCAMS DEFINE CLUSTER options that are available in the Dataset Define Cluster dialog box of the Storage Toolkit. Part 1 of 3, Cluster options

| Option | Description |
|-----------------------|---|
| Fully Qualified Name | Type the fully-qualified entryname for the cluster. This name is not propagated to the DATA and INDEX components of the cluster. For SMS and non-SMS-managed clusters, the component names must resolve to the same catalog as the cluster name of the data set. You can define a separate entryname for the cluster, its data component, and its index component. If no name is specified for the data and index component, a name is generated. When the cluster, data component, and index component are individually named, each can be addressed. |
| Fully Qualified Model | Type the fully-qualified name of the cluster that you want to use as a model. |
| Space Units | Select the unit of space that you want to allocate in the volume. Options include cylinders, kilobytes, megabytes, records, or tracks. Then specify a corresponding value for the unit that you selected. The Primary value that you specify allocates the initial amount of space to the cluster. The Secondary value that you specify allocates an amount of space for each time the cluster extends. |
| Volser(s) | Specifies the volumes on which a cluster is to have space. The maximum number of VOLSERs is 59. Choose one of the following values in the list: <ul style="list-style-type: none"> The serial number of an eligible direct access volume on which a new cluster is to be located OR <ul style="list-style-type: none"> The serial number of a direct access volume on which an old cluster is located. If you choose an old cluster for Volser , the cluster must be on the specified volume for allocation to take place. If you do specify a value for Volser , the system allocates new clusters to any eligible direct access volume. |
| Control interval size | The size of the control interval for the cluster. |
| Buffer space | Specifies the minimum space for buffers. If you do not specify a value, VSAM attempts to get enough space to contain two data component control intervals and, if the data is key-sequenced, one index component control interval. |
| Data class | Identifies the name (1-to-8 characters) of the data class for the data set. This metric provides the allocation attributes for new data sets. |
| Account | Specify up to 32-bytes of accounting information and user data for the data set. The information is supported for only SMS-managed VSAM and non-VSAM data sets. The information is used only for the data set level (not member level) of PDSE/PDS. |
| Management class | <i>For SMS-managed data sets</i> , you can type the name of the management class for a new data set. |
| Owner | Type the name of the owner of the cluster. |
| Storage class | <i>For SMS-managed data sets</i> , you can type the name of the storage class. Your storage administrator defines the range of names that you can use for storage classes. |

Table 22. IDCAMS DEFINE CLUSTER options that are available in the Dataset Define Cluster dialog box of the Storage Toolkit. Part 1 of 3, Cluster options (continued)

| Option | Description |
|------------------------|--|
| Catalog | Type the name of the catalog in which you want to define the entry. |
| Erase | Specifies that the cluster is overwritten with binary zeros when its entry in the catalog is deleted. |
| Recatalog | Select this option to cause the catalog entries for the cluster components to be re-created from information in the VVDS, if valid VVDS entries are found on the primary VVDS volume. |
| Reuse | Select this option to enable the cluster to be opened repeatedly as a reusable cluster. |
| Spanned | Select this option to enable a data record to cross the boundaries of a control interval. If the maximum length of a data record (as specified with RECORDSIZE) is larger than a control interval, the record is contained on more than one control interval. This option allows VSAM to select a control interval size that is optimal for the DASD. You cannot use this option in the definition of a linear data set cluster. |
| Speed | Select this option to prevent preformatting of the control areas of the data component during loading (creation) of a data set. |
| Data organization type | By default, the data organization type is set to Indexed. Use this option to specify the type of data organization for the cluster: Indexed, Linear, Non-indexed, or Numbered. |

Table 23. IDCAMS DEFINE CLUSTER options that are available in the Dataset Define Cluster dialog box of the Storage Toolkit. Part 2 of 3, (optional) Data component options

| Option | Description |
|-----------------------|--|
| Fully Qualified Name | Type the fully-qualified name of the data component. |
| Fully Qualified Model | Type the fully-qualified name of the data component that you want to use as a model. |
| Space Units | Select the unit of space that you want to allocate in the volume. Options include cylinders, kilobytes, megabytes, records, or tracks. Then specify a corresponding value for the unit that you selected. The Primary value that you specify allocates the initial amount of space to the data component. The Secondary value that you specify allocates an amount of space for each time the data component extends. |
| Volser(s) | Specifies the volumes on which a component is to have space. The maximum number of VOLSERs is 59. Choose one of the following values in the list: <ul style="list-style-type: none"> The serial number of an eligible direct access volume on which a new component is to be located OR <ul style="list-style-type: none"> The serial number of a direct access volume on which an old component is located. If you choose an old component for Volser, the component must be on the specified volume for allocation to take place. If you do specify a value for Volser, the system allocates new components to any eligible direct access volume. |
| Control interval size | The size of the control interval for the component. |
| Buffer space | Specifies the minimum space for buffers. If you do not specify a value, VSAM attempts to get enough space to contain two data component control intervals and, if the data is key-sequenced, one index component control interval |
| Record Size | Specifies the average and maximum lengths, in bytes, of the records in the data component. The minimum record size is 1 byte. <ul style="list-style-type: none"> Nonspanned records: For nonspanned records, the (maximum record size + 7) total cannot exceed the size of the control interval of the data component. In this case, the default values for Average and Maximum are 4089 and 4089, respectively. Spanned records: When you use a record size that is larger than one control interval, you must also select the Spanned option in the Cluster tab. In this case the default values for Average and Maximum are 4086 and 32600, respectively. |

Table 23. IDCAMS DEFINE CLUSTER options that are available in the Dataset Define Cluster dialog box of the Storage Toolkit. Part 2 of 3, (optional) Data component options (continued)

| Option | Description |
|---------------|--|
| Free Space | <p>Specifies the percentage of each control interval (CI) and control area (CA) to be set aside as free space when the cluster is initially loaded or when a mass insert is done. The value of CI% is the percentage of space to be preserved for adding new records and updating existing records with an increase in the length of the record.</p> <p>Because a CI is split when it becomes full, the CA might also need to be split when it is filled by CIs created by a CI split. The empty space in the control interval and control area is available for data records that are updated and inserted after the cluster is initially loaded.</p> <p>This parameter applies only to key-sequenced clusters, and variable-length relative records with variable-length records. CI-percent is the number of bytes that is equal to, or slightly less than, the percentage value of CI-percent. CA-percent is the number of control intervals equal to, or less than, the percentage of CA-percent.</p> |
| Share Options | <p>These settings determine how a component or cluster can be shared among users. To ensure integrity, you should be sure that share options specified at the DATA and INDEX levels are the same.</p> <ul style="list-style-type: none"> • Cross Region: Specifies the amount of sharing allowed among regions within the same system. • Cross System: Specifies the amount of sharing allowed among systems. Job steps of two or more operating systems can gain access to the same VSAM data set regardless of the disposition indicated in the DD statement for the data set in each step. <p>For information regarding the options, see the description of the IDCAMS DEFINE CLUSTER command consult <i>DFSMS Access Method Services for Catalogs</i>. A general link to this documentation is provided in Table 19 on page 258.</p> |
| Keys | <p>In this area, you can specify information about the prime key field of a key-sequence data records of a data set. This parameter applies only to key-sequenced clusters. The default is a key field of 64 bytes, beginning at the first byte (byte 0) of each data record. The key field of the cluster's index is called the prime key to distinguish it from other keys, called alternate keys.</p> <ul style="list-style-type: none"> • Length: Specifies the length of the key. • Offset: Specifies the displacement of the key (in bytes) from the beginning of the record. The sum of length plus offset cannot exceed the length of the shortest record. The length of the key can be 1 to 255 bytes. |
| Erase | Select this option to cause components to be overwritten with binary zeros when their entries in the catalog are deleted. |
| Reuse | Select this option to enable the component to be opened repeatedly as a reusable cluster. |
| Spanned | Select this option to enable a data record to cross the boundaries of a control interval. If the maximum length of a data record (as specified with Record Size) is larger than a control interval, the record is contained on more than one control interval. This option allows VSAM to select a control interval size that is optimum for the DASD. |
| Speed | Select this option to prevent preformatting of the control areas of the data component during loading (creation) of a data set. |

Table 24. IDCAMS DEFINE CLUSTER options that are available in the Dataset Define Cluster dialog box of the Storage Toolkit. Part 3 of 3, (optional) Index component options

| Option | Description |
|-----------------------|---|
| Fully Qualified Name | Type the fully-qualified name of the index component. |
| Fully Qualified Model | Type the fully-qualified name of the index component that you want to use as a model. |

Table 24. IDCAMS DEFINE CLUSTER options that are available in the Dataset Define Cluster dialog box of the Storage Toolkit. Part 3 of 3, (optional) Index component options (continued)

| Option | Description |
|-----------------------|--|
| Space Units | Select the unit of space that you want to allocate in the volume. Options include cylinders, kilobytes, megabytes, records, or tracks. Then specify a corresponding value for the unit that you selected. The Primary value that you specify allocates the initial amount of space to the data component. The Secondary value that you specify allocates an amount of space for each time the index component extends. |
| Volser(s) | Specifies the volumes on which a component is to have space. The maximum number of VOLSERs is 59. Choose one of the following values in the list: <ul style="list-style-type: none"> The serial number of an eligible direct access volume on which a new component is to be located OR <ul style="list-style-type: none"> The serial number of a direct access volume on which an old component is located. If you choose an old component for Volser, the component must be on the specified volume for allocation to take place. If you do specify a value for Volser, the system allocates new components to any eligible direct access volume. |
| Control interval size | The size of the control interval for the component. |
| Share Options | These settings determine how a component or cluster can be shared among users. To ensure integrity, you should be sure that share options specified at the DATA and INDEX levels are the same. <ul style="list-style-type: none"> Cross Region: Specifies the amount of sharing allowed among regions within the same system. Cross System: Specifies the amount of sharing allowed among systems. Job steps of two or more operating systems can gain access to the same VSAM data set regardless of the disposition indicated the DD statement for the data set in each step. For information regarding the options, see the description of the IDCAMS DEFINE CLUSTER command consult <i>DFSMS Access Method Services for Catalogs</i> . A general link to this documentation is provided in Table 19 on page 258. |
| Reuse | Select this option to enable the component to be opened repeatedly as a reusable cluster. |

Dataset Catalog dialog box

You use this dialog box to define a catalog entry for a single non-VSAM data set you have selected in the Tivoli Enterprise Portal.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The action request defines a catalog entry for a non-VSAM data sets or a collection of objects. The dialog box provides the following options:

- **Catalog:** Identifies the catalog in which the non-VSAM data set or OAM object is to be defined.
- **Device types:** Specifies the device types of the volumes containing the non-VSAM data set being defined. If the non-VSAM data set is located on different device types, the device types must be specified in the same order as the volume serial numbers listed in the **Volser(s)** parameter.

You can specify a generic device name that is supported by your system, using up to six characters. For example, 3380, 3390, 3400-2, and 3400-5 are valid device types.

If you expect to change the device type of the system residence volume, you can enter four zeroes (0000) in one of the rows of the Device Types table. At LOCATE and DELETE time, this field is resolved to the device type. This method allows you to use the non-VSAM data sets without having to

recatalog them to point to the new volume. A Device Type value of 0000 causes the actual device type to be determined from the current system residence volume (or its logical extension) at the time the catalog entry is retrieved.

To avoid an erroneous result, you must use this wildcard convention (0000) with a **Volser(s)** parameter of the following types:

- An indirect volume serial (*****) identifier.
- An extended indirect volume serial (a system symbol) identifier in the form &xxxxx, where &xxxxx is a symbol contained in some system library. The volume must be mounted and on line at the time the catalog entry is retrieved from the catalog. If it is not mounted and on line, the catalog request is terminated with a return and reason code.
- **Volser(s)**: Specifies the volumes that contain the non-VSAM data set. A **Volser(s)** value is required when you define a non-OAM non-VSAM data set. As mentioned in the previous field description, you can also specify indirect volume serial identifiers and extended indirect volume serial identifiers. When you use the wildcard specification of six asterisks (*****), the system dynamically resolves this specification to the system residence volume serial number, whenever the catalog entry is used. This value is not resolved when the DEFINE NONVSAM command is processed. This behavior allows you to later change the volume serial number of system residence volume without also having to recatalog the non-VSAM data sets on that volume.
- **Recatalog**: Specifies that the catalog entries are recreated if valid VVDS entries are found on the primary VVDS volume. RECATALOG can be specified only for an SMS-managed data set.

Dataset Delete dialog box

You use this dialog box to issue an IDCAMS DELETE request to delete the data sets that you select in the Tivoli Enterprise Portal. To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The dialog box provides the following options:

- **Entry type**: Select an option in this drop-down list to identify the type of object or entry to be deleted. If the object to be deleted is a catalog, truenamename entry, or VSAM volume record, you must specify its entry type as USERCATALOG, TRUENAME, NVR, or VVR. If you delete a migrated data set without specifying the entry type, DFSMSHsm deletes the data set without recalling it.
- **Catalog**: Type the name of the catalog that contains the entries to be deleted.
- **Erase**: Select this option to cause the deleted entries to be overwritten with binary zeros.
- **Force**: Select this option to enable deletion of entries, even if they are not empty.
- **Purge**: Select this option to enable deletion of entries, even when a retention date has been specified.
- **Recovery**: Select this option to enable retention of information regarding the deleted records, in anticipation of a recovery operation. When you select **Recovery**, the **Erase**, **Force**, **Purge**, and **Noscratch** options are disabled.
- **Noscratch**: Select this option to enable deletion of the catalog entry from the catalog without mounting the volume that contains the object defined by the entry.

Dataset Listcat dialog box

You use this dialog box to issue an IDCAMS LISTCAT request for the data sets that you select in the Tivoli Enterprise Portal. The dialog box provides the options shown in Table 25 on page 282.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Table 25. IDCAMS LISTCAT options that are available in the Dataset Listcat dialog box of the Storage Toolkit

| Option | Description |
|-------------------|--|
| Target Entries | <ul style="list-style-type: none"> • Use selected row entries: Select this option to run the LISTCAT command against the data sets that you selected in the Tivoli Enterprise Portal. • Level: Select this option to run the LISTCAT command against data sets that match the level of qualification that you specify in the text field. |
| Entry Type: | <ul style="list-style-type: none"> • Alias: Select this option to cause alias data sets to be listed. • Alternate index: Select this option to specify that entries for alternate indexes are to be listed. If Alternate index is specified and Data and Index are not also specified, entries for the data and index components of the alternate index are not listed. • Cluster: Select this option to specify that cluster entries are to be listed. If Cluster is specified and Data and Index are not also specified, entries for the data and index components of the cluster are not listed. • Data: Select this option to specify that entries for data components of clusters and alternate indexes are to be listed. • Generation group: Select this option to specify that entries for generation data groups are to be listed. GDSs in the active state, existing at the time the LISTCAT command is entered, are identified as such when ALL is specified • Index: Select this option to specify that entries for index components of key-sequenced clusters and alternate indexes are to be listed. If a VSAM name of the object is specified and INDEX is coded, only the index component entry of the object is listed. When INDEX is the only entry type parameter coded, the index component of the catalog is not listed. • Non-VSAM: Select this option to specify that entries for non-VSAM data sets are to be listed. • Page space: Select this option to specify that entries for page spaces are to be listed. • Path: Select this option to specify that entries for paths are to be listed. • User catalog: Select this option to specify that catalog connectors are to be listed. The user catalog connector entries are in the master catalog. (User catalog connector entries can also be in a user catalog, but the operating system does not recognize them when searching for a user catalog.) |
| Creation (days) | <p>Type a days value to cause entries of the indicated type (CLUSTER, DATA, and so on) to be listed only if they were created in the specified number of days ago or earlier. The maximum number that can be specified is 9999. Zero indicates that all entries are to be listed.</p> |
| Expiration (days) | <p>Type a days value to cause entries of the indicated type (CLUSTER, DATA, and so on) to be listed only if they expire in the specified number of days or earlier. The maximum number that can be specified is 9999 and it indicates that all entries are to be listed. Zero indicates that only entries that have already expired are to be listed.</p> |
| Listing type | <p>Select this option to specify the fields to be included for each entry listed. The following list shows the options that are available:</p> <ul style="list-style-type: none"> • Name specifies that the name and entry type of the entries are to be listed. Some entry types are listed with their associated entries. The entry type and name of the associated entry follow the name of the listed entry. • History specifies that only the following information is to be listed for each entry: name, entry type, ownerid, creation date, expiration date, and release. • Volume specifies that the information provided by specifying HISTORY, plus the volume serial numbers and device types allocated to the entries, are to be listed. • Allocation specifies that the information provided by specifying VOLUME plus detailed information about the allocation are to be listed. • All specifies that all fields are to be listed. |

Table 25. IDCAMS LISTCAT options that are available in the Dataset Listcat dialog box of the Storage Toolkit (continued)

| Option | Description |
|---------|--|
| Catalog | Specifies the name of the catalog that contains the entries that are to be listed. If Catalog is coded, only the entries from that catalog are listed. See the "Catalog Search Order for LISTCAT" topic in the IDCAMS documentation for information about the order in which catalogs are searched. |

Print dialog box

You use this dialog box to issue an IDCAMS PRINT request to print the data sets that you select in the Tivoli Enterprise Portal. To learn about the standard capabilities of a dialog box, see "Standard tabs in the dialog boxes of the Storage Toolkit" on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

The dialog box provides the following options:

Print tab

Access the **Print** tab in the **Options** tab of this dialog box to set attributes for format and range of printing:

- **Format of Output:** Select **Character**, **Dump**, or **Hex**.
- **Print Range Attributes:** Choose a **From/To** range or a **Skip/Count** range.
 - **From/To:** When you choose the **From/To** range type, specify the following values:
 - **Range Type:** This drop-down menu offers the following options:
 - **Key:** To specify the keys of the first and the last record you want to print.
 - **Address:** To specify the relative byte address (RBA) of the first and the last record to be printed.
 - **Number:** To specify the relative record number of the first and the last record to be printed.
 - **From** and **To** fields: Enter values that correspond to the **Range Type** that you specified:
 - To target multiple items that have similar key names, you can specify an initial character string followed by a wildcard character (*).
 - Other valid range values include decimal, hexadecimal, and binary numeric.
 - **Skip/Count:** When you choose the **Skip/Count** range type, specify the following values:
 - **Number of records to skip:** Zero (0) is the default.
 - **Count of records to print:** The number of subsequent records to print. All records, except those skipped, are printed by default

To print all the records choose either of the **From/To** and **Skip/Count** radio buttons, but leave the other fields blank.

Note: When you are working in the Print dialog box and you select the **From/To** or **Skip/Count** radio buttons, the radio buttons remain selected. If you decide to deselect one of the radio buttons, you must click **Cancel** to dismiss the dialog box. Access the dialog box again and make the revised selections that you want.

DBCS tab

Access the **DBCS** tab to configure printing of dual-byte character sets (DBCS):

- **Use DBCS handling:** Select this check box to enable the printing of dual-byte characters.
 - **Skip Options:** Specify skip options in this area. Select the **Skip DBCS check** option to specify a table of **Offset1** and **Offset2** pairs, which are byte offsets in the logical record that define the beginning and the end of an area where DBCS checking is not performed. Select the **No DBCS check** option to disable validity checking for dual-byte characters.

To omit any **Skip Options** in the IDCAMS PRINT command select the **Skip DBCS check** option and leave the offset table blank.

- **Shift Options:** Specify shift options in this area. These options serve the scenario in which SO (shift out) and SI (shift in) characters are not present already in the record that you are

printing. Select the **Insert shift** option to specify a table of **Offset1** and **Offset2** pairs, which are byte offsets in the logical record that indicate where the DBCS characters begin (**Offset1**) and end (**Offset2**).

Select **Insert all** when the record to be printed consists of only dual-byte characters. This option causes a single SO character to be inserted at the beginning of each record and SI to be inserted at the end of the record.

To omit any **Shift Options** in the IDCAMS PRINT command select the **Insert shift** option and leave the offset table blank.

Dataset Rename dialog box

You use this dialog to issue an IDCAMS ALTER request to rename a specific dataset that you select in the Tivoli Enterprise Portal. The dialog box provides the following option: **Fully Qualified New name**. Type a new, fully-qualified name for the data set in this field. This dialog box is a simplified version of the “Dataset Alter dialog box” on page 276.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Dataset Repro dialog box

You use this dialog box to issue an IDCAMS REPRO request for the data set that you select in the Tivoli Enterprise Portal. The dialog box displays options in the Repro and Merge/**DBCS** tabs that are located in the **Options** tab.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Table 26. IDCAMS REPRO options that are available in the Dataset Repro dialog box of the Storage Toolkit, Part 1 of 2, Repro tab

| Option | Description |
|--|--|
| Target Dataset section | Specify the fully qualified name of the target data set in the Fully Qualified Dataset field. |
| Source Dataset Copy Options section | |
| Volume entries | Use this option to copy/merge a tape volume catalog to another tape volume catalog. Specify the entries of the source tape volume catalogs to be merged or copied in this field. |
| Error limit | Type a failure limit in this field. Use this parameter to set a limit to the number of errors that REPRO copy tolerates. |
| Replace | Select this option to indicate that a record in the source cluster is to replace a record in the target cluster when the source cluster is copied into the target cluster. |
| Reuse | Select this option if the target data set is to be opened as a reusable data set. This parameter is valid only for VSAM data sets. |

Table 26. IDCAMS REPRO options that are available in the Dataset Repro dialog box of the Storage Toolkit, Part 1 of 2, **Repro** tab (continued)

| Option | Description |
|------------------------|---|
| Repro Range Attributes | <p>Choose a From/To range or a Skip/Count range.</p> <ul style="list-style-type: none"> • From/To: When you choose the From/To range type, specify the following values: <ul style="list-style-type: none"> – Range Type: This drop-down list offers the following options: <ul style="list-style-type: none"> - Key: To specify the keys of the first and the last record you want to copy. - Address: To specify the relative byte address (RBA) of the first and the last record to be copied. - Number: To specify the relative record number of the first and the last record to be copied. – From and To fields: Enter values that correspond to the Range Type that you specified: <ul style="list-style-type: none"> - To target multiple items that have similar key names, you can specify an initial character string followed by a wildcard character (*). - Other valid range values include decimal, hexadecimal, and binary numeric. • Skip/Count: When you choose the Skip/Count range type, specify the following values: <ul style="list-style-type: none"> – Number of records to skip: Zero (0) is the default. – Count of records to copy: The number of subsequent records to copy. All records, except those skipped, are copied by default. <p>To copy all the records choose either of the From/To and Skip/Count radio buttons, and leave the other fields blank.</p> |

Table 27. IDCAMS REPRO options that are available in the Dataset Repro dialog box of the Storage Toolkit, Part 2 of 2, **Merge/DBCS** tab

| Option | Description |
|-----------------------|---|
| Merge Entries section | <p>Mergecat: You cannot select this option when Use DBCS support is selected. Select Mergecat to specify that the source catalog entries are to be merged with the target catalog entries and that the merged entries are to be deleted from the source catalog upon a successful merge operation.</p> <p>You can specify the entries to be merged either as the list of Entries (entry names or generic names) or as a single Level of qualification.</p> <ul style="list-style-type: none"> • You can specify a list of names and generic names in the Entries field to match all the entries to be merged. When using a generic name with the ENTRIES parameter, entries must have one qualifier in addition to those specified in the command. • You can specify a Level value in order that all entries matching the level of qualification you indicated are to be merged. |
| Use DBCS support | <p>Use this option to specify that bytes in the logical record contain DBCS characters. The Repro command checks to ensure the DBCS data meets DBCS criteria. This parameter cannot be specified with MERGECAT.</p> |
| Skip section | <p>Use the Skip DBCS check options in this section to indicate that characters between Offset1 and Offset2 are not to be checked for DBCS criteria during Repro command processing. You can specify several Offset1-Offset2 pairs here.</p> <ul style="list-style-type: none"> • Offset1 indicates the byte offset in the logical record to be copied at which checking is to cease until offset2 is reached. • Offset2 indicates the byte offset in the logical record after which checking is to resume. Offset2 must be greater than Offset1. <p>No DBCS check: Select this option to disable DBCS verification checking.</p> |

Table 27. IDCAMS REPRO options that are available in the Dataset Repro dialog box of the Storage Toolkit, Part 2 of 2, Merge/DBCS tab (continued)

| Option | Description |
|---------------|---|
| Shift section | <p>Specify insertion of a DBCS shift in this section. Select Insert Shift to specify offsets. If SO and SI characters are not present, you must specify shift values. (By default, the logical record is assumed to already contain SO and SI characters, and the Repro command checks during copying to ensure that the DBCS data meets DBCS criteria.)</p> <ul style="list-style-type: none"> • Offset1 indicates the byte offset in the logical record to be copied before which a SO character is to be inserted. • Offset2 Indicates the byte offset in the logical record to be copied after which a SI character is to be inserted. Offset2 must be greater than Offset1 and the difference must be an odd number. <p>Select Insert all to specify that the entire logical record is assumed to contain only DBCS characters.</p> |

Dataset Uncatalog dialog box

You use this dialog to issue an IDCAMS DELETE NOSCRATCH request to uncatalog the datasets that you select in the Tivoli Enterprise Portal. An **Options** tab is not provided, because there are no options to set.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Dataset Verify dialog box

You use this dialog box to issue an IDCAMS VERIFY request for the data sets that you select in the Tivoli Enterprise Portal. The VERIFY command causes a catalog to correctly reflect the end of a VSAM data set after an error occurs while closing a VSAM data set. This dialog box does not have the **Options** tab, because there are no options to set.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

DFSMSHsm actions in the Storage Toolkit

This topic provides links to information about dialog boxes that are related to DFSMSHsm actions in the Storage Toolkit.

You access the Cancel HSM Request dialog box in the pop-up menu that you access by right-clicking specific views in a workspace.

You can also create a custom action. See “Custom actions in the Storage Toolkit” on page 265.

Cancel HSM Request dialog box

You use this dialog box to issue the Cancel HSM Request command to cancel active or queued DFSMSHsm (“HSM”) requests. Unlike other dialog boxes for the Storage Toolkit, the **Options** tab is not displayed in the Cancel HSM Request command dialog box. This action has no HSM command parameter options to set.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Accessing the Cancel HSM Request feature

The Cancel HSM Request feature is available through the HSM Status workspace for Tivoli OMEGAMON XE for Storage on z/OS. You access and use the feature as follows:

1. Navigate to the HSM Status workspace in the Navigator. The workspace includes the HSM Functions Status Report table.
2. Click one of the blue link icons on the left side of the rows in that HSM Functions Status Report table. The HSM Request Details workspace is displayed. This workspace includes the HSM Requests view that lists HSM requests.
3. In the HSM Requests view select one or more rows (HSM commands) to cancel. Press the Control (**Ctrl**) key and click multiple rows, as needed.
4. Right-click on the selected rows to access the context menu.
5. Select **Cancel HSM Request**. The Cancel HSM Request dialog box is displayed.

Note: The menu item is available when all the items that you have selected are either queued or active. If the state of any selected item is **complete**, **abended**, or **cancelled**, the menu is not available.

6. As needed, enter options in the fields of the **General** tab, **Schedule** tab, and **JCL** tab of the dialog box. “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258 describes these tabs. Use the **Replacement JCL JOB card** field in the **JCL** tab to assign a unique job ID to this job. If you do not do this, the batch job that is submitted to process the HSM cancel request might not run until any job that is currently executing with that job ID completes.
7. Click **OK** in the dialog box to confirm your settings. A confirmation dialog box is displayed.
8. Read the warnings in the confirmation dialog and then click **OK**, as appropriate.

See “Warning about cancellation of HSM requests” for further information.

Warning about cancellation of HSM requests

Cancelling HSM requests can have negative consequences. This section introduces some issues that you should consider before cancelling an HSM request. See *z/OS DFSMSHsm Storage Administration Reference* (<http://publibz.boulder.ibm.com/epubs/pdf/dgt2s440.pdf>) for detailed information about these issues.

- **Active versus stopped:** The Cancel HSM Request feature is intended for use when *active* HSM tasks are suspended. Cancelling an HSM task that is actually moving toward completion might have unintended consequences.
- **Required followup activity:** When an **HSM RECOVER REPLACE** command is cancelled, it might be necessary to reallocate and recatalog the target data set before you issue the **HSM RECOVER** command again.
- **When duplexing output tapes:** There is a limitation in **HSM BACKDS** processing when output tapes are duplexed. Under some circumstances, processing stops. Cancelling DFSMSHsm might be required to release the original tape and tape drive.
- **Backup tasks:** Before you cancel a suspended HSM journal backup task, allow any active **HSM CS** backup tasks to run to completion. If an **HSM CDS** backup task is also suspended, cancel it before you cancel the HSM journal backup task that is suspended.

DFSMSrmm actions in the Storage Toolkit

This topic provides links to information about dialog boxes that are related to DFSMSrmm actions in the Storage Toolkit.

The DFSMSrmm TSO command is documented in the *DFSMSrmm Guide and Reference* (SC26-7404-07) available at <http://publibz.boulder.ibm.com/epubs/pdf/dgt2r360.pdf>. Consult this manual in order to identify required parameters.

You access the following dialog boxes in the pop-up menu that you access by right-clicking specific views in a workspace.

You can also create a custom action. See “Custom actions in the Storage Toolkit” on page 265.

Add Dataset dialog box

You use this dialog box to issue the DFSMSrmm ADDDATASET command from RMM data-set-oriented workspaces. You must specify a fully-qualified name for the data set.

You can specify the following details for the data set that you add:

- Fully Qualified DSN
- Volser
- Retention area:
 - Retention days
 - Expires on
- File sequence
- Block size
- Logical record length
- Record format

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Change Dataset dialog box

You use this dialog box to issue the DFSMSrmm CHANGEDATASET command from RMM data-set-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Delete Dataset dialog box

You use this dialog box to issue the DFSMSrmm DELETEDATASET command from RMM data-set-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Add Owner dialog box

You use this dialog box to issue the DFSMSrmm ADDOWNER command from RMM owner-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Change Owner dialog box

You use this dialog box to issue the DFSMSrmm CHANGEOWNER command from RMM owner-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Delete Owner dialog box

You use this dialog box to issue the DFSMSrmm DELETEOWNER command from RMM owner-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Add Volume dialog box

You use this dialog box to issue the DFSMSrmm ADDVOLUME command from RMM volume-oriented workspaces. The dialog box presents the most frequently used parameters and options for the command.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Change Volume dialog box

You use this dialog box to issue the DFSMSrmm CHANGEVOLUME command from RMM volume-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Confirm Actions for All Volumes dialog box

You use this dialog box to issue the DFSMSrmm CHANGEVOLUME command from RMM volume-oriented workspaces. This dialog box exploits the CONFIRM... parameters of this DFSMSrmm command.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Delete Volume dialog box

You use this dialog box to issue the DFSMSrmm DELETEVOLUME command from RMM volume-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Add VRS dialog box

You use this dialog box to issue the DFSMSrmm ADDVRS command from RMM VRS-oriented workspaces. You must specify a fully-qualified name for the data set.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Change VRS dialog box

You use this dialog box to issue the DFSMSrmm CHANGEVRS command from RMM VRS-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Delete VRS dialog box

You use this dialog box to issue the DFSMSrmm DELETEVRS command from RMM VRS-oriented workspaces.

The **Options** tab contains the most commonly used keywords of this command. Not all keywords are presented in the **Options** tab. If a command keyword that you need is not presented, access the **Command** tab for this dialog box and enter the appropriate keyword and parameters in the **Command** input area.

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Overview of common action commands

The Storage Toolkit provides actions that start commands on target resources in the z/OS system. Table 28 shows typical goals for some of the action commands in the Storage Toolkit.

Table 28. Common commands that are enabled by the Storage Toolkit

| Command type | Goal |
|----------------------------|---|
| HSM Volume Services | Migrate data sets on volume not referenced for a certain number of days Perform HSM incremental backup of eligible data sets on volume |
| HSM Dataset Services | Migrate data set to level 1 Migrate data set to level 2 Backup data set Recover data set from backup version Recall a migrated data set |
| DFDSS Volume Services | Release unused space in data sets on volume TEST volume convert to SMS format Prepare volume for conversion to SMS format Convert volume to SMS format |
| DFDSS Dataset Services | Copy/move data set Release unused space within data set |
| Cancel HSM Request Service | Cancel active or queued DFSMSshm (“HSM”) requests |

Related information

“Using the Storage Toolkit” on page 254

Managing actions

The following workspaces help you view and manage actions:

- **Storage Toolkit Action Requests workspace:** The Storage Toolkit Action Requests workspace of the Storage Toolkit enables you to view and maintain reports on actions. You select and right-click items in the workspace to access context menus that allow you to create, submit, delete all, and delete the items.
Action summaries and the data from the running of actions persist until you delete them. Select **Delete All Results** to delete all of the results associated with the selected action requests. Select **Delete Request** to delete the selected action requests and all of the results associated with them.
- **Storage Toolkit Result Summary workspace:** The Storage Toolkit Action Requests Summary workspace provides a link to the Storage Toolkit Result Summary workspace. As with the principal workspace, you select and right-click items to select management operations, such as delete. Action summaries and the data from the running of actions persist until you delete them. Select **Delete Result Details** to delete the output data associated with the selected results. Select **Delete** to delete the selected results and the output data associated with the selected results.
- **DFSMSHsm Request Details workspace:** The DFSMSHsm Request Details workspace provides context menus that enable you to select and cancel active or queued DFSMSHsm requests. See “Cancel HSM Request dialog box” on page 286 for more information.

Storage Toolkit limitations and hints

This topic describes limitations, hints and other considerations for the Storage Toolkit. You can avoid problems by remaining aware of these issues. Also remain aware of other issues documented throughout this chapter.

- **Data sets not locked:** The Storage Toolkit enables you to click the **Edit JCL** button in the Create Batch Job dialog box, and edit a data set that is located in the mainframe environment. However, you must ensure that no one accesses the data set at the same time, because the data set is not locked during your editing session. Another user could edit the data set, for example, in TSO using ISPF, during your editing session. That user would not see a “data set in use” message; if he saves his changes before you save yours, your changes overlay his. If that user is still editing the data set when you attempt to save your changes, the attempt to save your changes fails. The results of each editing session might be unpredictable depending on the scenarios and editing tools that the other user uses.
- **EBCDIC code page 037 only:** The JCL editor provided by Tivoli OMEGAMON XE for Storage on z/OS supports EBCDIC code page 037 only. The editing or authoring of JCL in other code pages such as EBCDIC code page 930 (in other words, Japanese EBCDIC) is not supported.
- **Storage Toolkit first and last steps:** In every batch job, the Storage Toolkit inserts a Toolkit-specific first step that sets up a monitor over the steps in the job. The monitor collects targeted output from the steps in the job. The Storage Toolkit also appends a last step to the end of the job. The last step collects targeted SYSOUT from the previous steps along with the JES output and the return code. It also notifies the monitoring agent that the batch job is complete. The Toolkit-specific first step is inserted immediately before the first EXEC, PROC, or INCLUDE statement that it locates in the JCL.

Note: If your JCL has an INCLUDE statement before an EXEC or PROC statement, the INCLUDE member must not contain JCL statements, such as JCLLIB, that must precede the first job step. Because the Toolkit first step is inserted before the INCLUDE, the batch job fails in this case.

- **Using the null statement to mark the end of a job:** When you use the Create Batch Job dialog box to execute user-defined JCL, the Storage Toolkit generates a copy of your JCL to which it adds a Toolkit-specific first step and a Toolkit-specific last step. If your JCL ends with the null JCL statement denoting end of job,

```
//
```

that null statement is removed from the generated JCL because the job does not end until after the Toolkit-specific last step runs.

- **Conditional processing:** Do not specify a COND parameter on the JOB card or on an EXEC statement that might cause the Storage Toolkit steps that were inserted into the JCL to not run. If you use the COND parameter or you use IF/THEN/ELSE/ENDIF statements, you must ensure that the Storage Toolkit first and last steps run.
- **No support for multiple jobs in user-defined JCL:** When you use the Create Batch Job dialog box to execute user-defined JCL, the JCL must not include multiple jobs. The Storage Toolkit does not support this type of JCL. Results are unpredictable.
- **JOB card:**
 - When you use the Create Batch Job dialog box to execute user-defined JCL, the batch job is submitted using the Replacement JCL JOB card which you specify on the JCL tab. This overrides a JOB card that might be present in the JCL. If you do not specify a Replacement JCL JOB card, your installation-specific JOB card is used.
 - Do not specify a CLASS or a TYPRUN option on your JOB card that just copies or scans the job. Because the batch job does not execute, your action request remains in the EXECUTING state. You must cancel the request to release the thread and resources associated with it and to remove it from EXECUTING state.
 - Do not specify a COND parameter on the JOB card that might cause the Storage Toolkit first and last steps inserted into the JCL to not run.
 - When you request the JES output be copied for later viewing, make sure the MSGLEVEL on your JOB card is set to the level of output that you desire.
 - When you specify your JOB card, consider assigning it a unique job name. If the name matches a batch job that is executing on your z/OS system, your job might be delayed until the executing job completes. To prevent this, assign a unique name to your job.
- **Using a PROC in user-defined JCL:** When you use the Create Batch Job dialog box, you can run JCL that executes a procedure, however the Storage Toolkit might not be able to properly copy the contents of files associated with steps in that procedure:
 - The procedure can be instream or in a system or private procedure library. If you use a private procedure library, you must ensure the JCLLIB statement precedes the Storage Toolkit first step.
 - You can request files that are referenced in the procedure be copied for later viewing, but with certain limitations:
 - The step name is the step name in your JCL that executes the procedure. You cannot specify the step names that are in the procedure itself.
 - If the procedure consists of a single step, the contents of the requested files is returned.
 - If there are multiple steps in the procedure, the contents of a requested data set or DD name that references a data set is returned for each procedure step (in other words, multiple times). The contents of a DD name that is routed to SYSOUT is returned for each procedure step in which the SYSOUT DD name is defined (ie., one or more times).
- **Use step names in user-defined JCL:** Do not include steps in your user-defined JCL without step names, if you intend to copy files for later viewing that are associated with those steps. The Storage Toolkit requires a step name.
- **/*XMIT:** Do not use the /*XMIT statement in any of your JCL. The Storage Toolkit does not support this. Results are unpredictable.

- **DYNAMNBR:** If you submit user-defined JCL that allocates data sets, be aware that the Storage Toolkit allocates data sets in each step, too. If necessary, you might need to use the DYNAMNBR parameter on your EXEC statement to allow for your data sets and 3 Storage Toolkit data sets.
- **JCL errors:** The Storage Toolkit last step appended to the end of each batch job notifies the monitoring agent when the batch job completes. When this occurs the action request that is pending completion of the job is updated with the results of the batch job. If the last step in the batch job does not run, for example, the batch job failed with a JCL error or conditional processing bypassed the last step, the action request remains in an EXECUTING state. If you determine that an action request is in EXECUTING state longer than you anticipate, you must check the status of the batch job on the z/OS system. If the job failed such that the last step did not run, you must cancel the execution of the action request in the Tivoli Enterprise Portal. This releases the thread and resources associated with the request and removes it from EXECUTING state. You can then determine why the job failed, correct the error and resubmit the request.
- **Return codes:** Certain return codes, which are generally paired with a status, are set by the Storage Toolkit when it detects an error processing an action request. The following table lists common return codes and their corresponding status:

| Return code | Status |
|-------------|--|
| 117 | This status typically indicates that the JCL exceeds 72 characters when the substitution variables are applied. It might also indicate other JCL-related errors, such as a missing JOB card, or another data set requiring variable substitution that exceeds 80 characters when the substitution variables are applied. |
| 119 | The User Data Server has ended abnormally or the batch job has ended, but the Storage Toolkit is unable to determine the return code. |
| 121 | Authorization has failed. |
| 123 | A data set error has occurred, such as the data set containing your JCL does not exist. Messages in the RKLVLLOG of the Tivoli Enterprise Monitoring Server might help you analyze this result. |

If you see a return code that does not look familiar, you might want to convert it to its hexadecimal equivalent, because the code might indicate an abend in the batch job. For example, a return code of 193 is the same as x'0C1'.

- **Variable substitution and line limits:** Substitution variables that are defined through the Storage Toolkit are replaced when the action request runs. The Storage Toolkit creates temporary data sets to contain the updated statements.

There are two basic categories of data sets (JCL and Other) that are updated with substitution variables. The Storage Toolkit processes them as follows:

- **JCL data set:** Variable substitution is applied to all components of the batch job, including the JOB card, extra JCL, the Toolkit-specific steps, and the body of your JCL along with any instream data sets. The Storage Toolkit interprets column 72 as a continuation character and preserves its contents. The data between columns 2 and 71 might shift left or right depending on the size of the variable name and its substitution value. If the data shifts beyond column 71, the request fails. The return code for the request is set to 117 and the status for this execution is **InvalidJCL**. You must perform these actions:
 - Verify the substitution variables and values are correct and do not have unintentional consequences to the components of the batch job
 - Correct the JCL to ensure that no line exceeds the limit.

The action request may also fail with return code 117 if any lines in the original components of the batch job exceed 80 characters.

- **Other data set:** Variable substitution is applied to all records in other data sets that you specify as needed by the job that also contains substitution variables. The toolkit makes no assumptions about the contents of the data set and considers each line, from column 1 to column 80, as a line of data. Variable substitution may cause the data in columns 2 through 80 to shift left or right depending on

the size of variables names and their values. If the data shifts beyond column 80 (excluding trailing blanks), the request fails. The return code for the request is set to 117 and the status for this execution is InvalidJCL. You must perform these actions:

- Verify that the substitution variables and values are correct and do not have unintentional consequences to the contents of the data set.
- Correct the contents of the data set to ensure that no line exceeds the limit.
- **Validating JCL:** When you write JCL for use in the Create Batch Job dialog box, always check the validity of the statements before you submit the batch job. For example, when you edit JCL in the Edit JCL dialog box, consider whether line lengths will exceed the 72-byte limit after variable substitution is performed. When substitution variables are replaced in the JCL at execution time, resultant JCL lines that contain more than 72 bytes causes the JCL to not be submitted. A status of InvalidJCL is displayed in the Result Summary workspace for the action request.
- **Reserved variable names:** The Storage Toolkit reserves the following variable names. You must not use these names for your own variables:

```
%%KS3TK_CMD_DSN%%  
%%KS3TK_HSM%%  
%%KS3TK_DYNAMNBR%%
```

- **Fully Qualified Datasets needed by the job that also contain substitution variables:** When you use the Create Batch Job dialog box, you can specify additional data sets that contain substitution variables. The Storage Toolkit creates a temporary data set with the updates and replaces the name of the original data set with the temporary one in its copy of your JCL. In order for the names to be replaced, the data sets must be referenced in your JCL; they cannot be in cataloged procedures or INCLUDE members that your JCL might use.
- **JES output:**
 - The techniques that the Storage Toolkit uses to collect JES logs and system-handled output data sets (SYSOUT) require your z/OS operating system use JES2.
 - Because the Storage Toolkit last step collects JES output just before the batch job ends, some of the messages you normally see in the JES logs such as the job start (IEF375I) and job end (IEF376I) messages are not included in the JES output.
- **Mainframe commands:**
 - Mainframe console commands are submitted through an SDSF batch job interface. A forward slash (/) must precede the command, as in this example, which cancels a time-sharing user (*tso_user_ID*):

```
/C U=tso_user_ID
```
 - Command output is not returned for Mainframe console commands because execution of the command is not synchronized with execution of the batch job.
 - Because execution of the command is not synchronized with execution of the batch job, the return code associated with the action request reflects the submission of the command to the SDSF batch job interface. It does not reflect the execution of the command itself.
 - Because the Storage Toolkit uses SDSF, your z/OS operating system must use JES2.
- **Shared DASD:** The temporary data sets that the Storage Toolkit creates to contain the generated JCL, the results data set, and other files are shared between the Toolkit and the batch job. Because the batch job can run on a z/OS system in your SYSPLEX different from the one where the monitoring agent runs, the temporary data sets must be created on DASD shared across the systems. Your installation can control the location of the temporary data sets using options in the Configuration Tool. These options also control the location of data sets created using the Edit JCL option in the Create Batch Job dialog box.

In addition, when you use the Create Batch Job dialog box, you specify the data set containing the JCL you wish to submit and, optionally, specify data sets needed by the job that also contain substitution variables. These data sets must be cataloged and located on on-line DASD that is accessible to the z/OS system where the monitoring agent runs.

- **APF-authorized load library on remote systems:** The Storage Toolkit inserts a first step and last step into every batch job. These steps run Toolkit code that is located in the TKANMODL load library for

your installation's run time environment. The load library must be APF-authorized. If the batch job runs on the same z/OS system as the monitoring agent, the load library is normally already APF-authorized. If your batch job runs on another z/OS system in your SYSPLEX, you must ensure the load library is APF-authorized on that system as well. The load library must also be located on DASD that is shared across the systems.

- **Unprintable characters:** You must ensure that the files you specify as Files whose contents should be copied for later viewing on the Create Batch Job dialog box or the output from a command on the Issue Command dialog box will contain character data only. If the files or command output contain unprintable characters (for example, hexadecimal data), these characters might not display properly in the Storage Toolkit Result Detail workspace.
- **Checkpoint dataset storage exhausted:** When you submit an action request, information about the request is stored in the checkpoint database. When the request completes, results of the execution are also stored. The information stored in the checkpoint database includes elements such as:
 - The name and description of the request
 - The time the request was submitted and the time it completed
 - The resources associated with the request
 - The return code from the execution of the request
 - The output produced by the execution of the request, which might include:
 - Command output
 - Files copied for later viewing
 - The submitted JCL
 - The JES files produced by the batch job.

The Storage Toolkit checkpoint database is allocated and initialized when the runtime environment is created. This database is never extended. You must ensure that sufficient space exists for the activities of the database. Otherwise, the database can run out of space. You can remove the requests and the results when you no longer need them. You can also use options on the **General** tab of the Storage Toolkit dialog boxes to help you manage the results.

Note: If the results from the execution of an action request exceed the free space available in the checkpoint database, the output is lost entirely. The error message KS3T830E SERVICE CHECKPOINT DATASET STORAGE EXHAUSTED in the RKLVLLOG of the Tivoli Enterprise Monitoring Server indicates this condition. The *IBM Tivoli OMEGAMON XE for Storage on z/OS: Troubleshooting Guide* provides further information about this issue.

Related information

“Using the Storage Toolkit” on page 254

Chapter 24. Dataset Group feature

Tivoli OMEGAMON XE for Storage on z/OS offers several GUI-based features to help you identify and solve storage-related application problems. The graphical user interface enables you to navigate quickly to problem areas and automatically apply solutions, instead of using more time-consuming methods in the z/OS operating system:

- This chapter describes how to use the Dataset Group feature and expands on the information in Section 2 of *OMEGAMON II for SMS: User's Guide, GC32-9402*. The following topics are covered in this chapter:
 - How to add, edit, delete, and set the collection interval for data set groups
 - Data set masks and the rules for constructing them
 - Describes a number of ways that you can use data set groups to resolve storage problems.
- Another problem-solving tool is described in Chapter 25, “Dataset Attributes Database,” on page 311.
- Another problem-solving tool is described in Chapter 23, “Storage Toolkit,” on page 251.

The Dataset Group feature versus the Dataset Attributes feature

OMEGAMON XE for Storage on z/OS gives you two powerful facilities for managing and monitoring data sets: the Dataset Group feature and the Dataset Attribute functions. Because both of these functions provide similar (although not identical) information, you need to be aware of when one approach is better than the other:

- Construct and use Dataset Groups when you have specific sets of similarly named data sets that need to be monitored, either for space or performance monitoring.

Note: To see data-set-level I/O performance metrics, the volumes on which those data sets are located must have been defined for data-set-level I/O monitoring. For details, see the configuration procedure for the DASD Device Monitoring panel in the *IBM Tivoli OMEGAMON XE for Storage on z/OS: Planning and Configuration Guide*.

- Enable Dataset Attribute collection when you want to monitor overall space utilization, exceptions, statistics, trends, and so on with respect to data sets.
- Define Situations for Dataset Group tables that raise alerts if specific groups of data sets are in danger of running out of space or are performing poorly, as described in this section.
- Define Situations against Dataset Attribute tables to raise alerts when any data set exceeds a defined threshold (for example, number of extents) or when overall conditions (for example, total system free space) warrant that a warning be raised.

Working with data set groups

The following procedures show you how to add, edit, and delete data set groups using the Dataset Group feature. The procedure for setting the collection interval for a data set group is also described.

Note: The menu selections in the following procedures might not be available to the user depending on their security permissions.

Dataset Group Properties dialog box

You use this dialog box to define a group of data sets based on filtering rules (masks) and related properties.

You access this dialog box through the Dataset Group Summary workspace, which is described in Chapter 17, “Monitoring dataset groups,” on page 93. The settings to modify are as follows:

- **Mask.** You type a mask definition following the rules used by DFDSS data set name masking. See “Data set group masking rules” on page 299.
- **Active.** Select this checkbox to cause data collection from the catalog to be performed for the data sets that fall under a mask.
- **Space Data.** Select this checkbox to collect space data for the data sets that fall under a mask.

To learn how to add and configure a data set group see “Adding a data set group.”

Dataset Group Data Collection dialog box

You use this dialog box to set the interval at which data collection occurs.

You access this dialog box through the Dataset Group Summary workspace, which is described in Chapter 17, “Monitoring dataset groups,” on page 93. The settings to modify are as follows:

- **RMF interval.** Data is collected at the predefined time interval of the resource measurement facility (RMF).
- **On Demand.** Data is collected when the user requests collection.
- **mmn Mins (5 - 360).** Data is collected at the interval of minutes that you define, where *mmn* is an integer from 5 to 360.

Adding a data set group

Use the following procedure to add a new data set group.

1. In the navigator, click **Dataset Group Summary**.
2. Right-click on a row within the Dataset Group Summary table, and then click **Add Group**.
3. Enter a group name in the **Add New Group** window, click **OK**.
4. In the **Dataset Group Properties** window, type a mask value, press Enter.
5. Edit the group properties as desired. You can specify the following:
 - **Mask.** A mask against which catalogued data set names are processed. The masks follow the rules used by DFDSS data set name masking. For additional information on mask rules see “Data set group masking rules” on page 299. To add a new mask click the **Add a new mask...** button located in the upper left corner of the dialog box.
 - **Active.** When selected, data collection from the catalog is performed for the data sets that fall under the mask.
 - **Space Data.** When selected, space data is collected for the data sets that fall under the mask.
6. Click **OK** to finish adding the data set group.

Editing and deleting data set groups

You can edit a group definition in the Dataset Summary table by right-clicking in the group’s row, and then click **Edit Group**. You can then edit the group properties.

To delete a group in the Dataset Summary table. Right-click in the group’s row, and then click **Delete Group**.

Setting the collection interval

Use the following procedure to specify the collection interval.

1. In the Dataset Summary table, right-click in the row for any group, and then click **Collection Interval**. You set the interval for all groups in one operation.
2. Specify the group data collection interval with one of the three options:
 - **RMF Interval.** When selected, the collection interval is synchronized to the RMF interval.

- **On Demand.** When selected, the data is collected in real-time. However, data is collected no more often than 5 minutes.
- **Mins.** When selected, the collection interval can be specified to take place automatically every *N* minutes between 5 and 360 minutes.

3. Click **OK** to set the collection interval.

Data set group masking rules

A data set group mask is a character string that describes the data set group names for which you want information returned. Asterisks (*) and percent signs (%) are used to represent qualifiers or parts of qualifiers. Table 29 explains the use of asterisks and percent signs in masks.

Table 29. Mask characters

| Mask Characters | Description |
|-----------------|--|
| * | A single asterisk by itself indicates that either a qualifier or one or more characters within a qualifier can occupy that position. An asterisk can precede or follow a set of characters. |
| ** | A double asterisk indicates that zero or more qualifiers can occupy that position. A double asterisk cannot precede or follow any characters; it must be preceded or followed by either a period or a blank. |
| % | A single percent sign by itself indicates that exactly one alphanumeric or national character can occupy that position. |
| %% | One to eight percent signs can be specified in each qualifier. |

Mask scenarios

Consider a scenario in which you have the following data set groups:

- MYGROUP.DATA.SET1
- MYGROUP.DATA.SET2
- MYGROUP.DATA.SET30
- VSAM.MYGROUP.DATA.SET

Table 30 shows several masks and the information that each mask returns.

Table 30. Sample masks

| Mask | Information returned | Information excluded |
|--------------------|--|---|
| MYGROUP.DATA.* | MYGROUP.DATA.SET1 MYGROUP.DATA.SET2 MYGROUP.DATA.SET30 | VSAM.MYGROUP.DATA.SET |
| ** .DATA. ** | MYGROUP.DATA.SET1, MYGROUP.DATA.SET2, MYGROUP.DATA.SET30, and VSAM.MYGROUP.DATA.SET | — |
| MYGROUP.DATA.SET% | MYGROUP.DATA.SET1 and MYGROUP.DATA.SET2 | MYGROUP.DATA.SET30 and VSAM.MYGROUP.DATA.SET |
| MYGROUP.DATA.SET%% | MYGROUP.DATA.SET30 | MYGROUP.DATA.SET1, MYGROUP.DATA.SET2, and VSAM.MYGROUP.DATA.SET |

Using data set groups

The Dataset Group function provides a powerful facility for you to quickly examine and monitor the space and performance attributes of a collection of data sets. By defining a series of data set groups and masks, any collection of data sets can be displayed as a report within the portal, and Situations can be written to monitor the space or performance attributes of those data sets. This section presents examples and ideas of how this function can add value for storage management and administration.

Dataset Groups can be thought of as a more powerful version of the ISPF DSLIST utility (that is, 3.4). In addition to creating lists of data sets, OMEGAMON allows you to name and save those lists in groups, and provides much more powerful name masking capabilities. In addition, all of the attributes associated with the data sets in the generated lists can be used in reports, alerts and Storage Toolkit actions. The OMEGAMON Dataset Group function, in its current form, is not intended as a direct substitute for ISPF 3.4. While ISPF 3.4 is geared towards producing interactive, user-generated ad-hoc lists of data sets and running simple actions against those data sets, one at a time, the OMEGAMON Dataset Groups are more static in nature, and are better used for viewing large sets of related files and their attributes (individually and as a group), generating alerts based on those attributes, and taking actions.

Note: The new Dataset Group function differs from the User DASD Group and Application functions in that Dataset Groups start at the system catalog by using a name mask to create a list of matching data sets; User DASD Groups start at the volume or storage group level, and Applications start at the job level, to create their volume/data set lists. The important difference here is that using Dataset Groups you do not need to know where a data set is located or which application accesses it. Using this data set orientation can be more useful for end-users and storage administrators who prefer to perform certain tasks from this perspective.

The following examples describe how to use Tivoli OMEGAMON XE for Storage on z/OS to monitor a number of events that might occur in your system. They are:

- “Obtaining an overview of space usage”
- “Identifying data sets with excess space” on page 301
- “Create an alert when response time exceeds specific thresholds” on page 302
- “Identify VSAM data sets with excessive CA/CI splits” on page 304
- “Cataloged and not-allocated data sets” on page 306
- “Notify the Production Scheduling group when a critical data set exceeds a specified number of extents” on page 308
- “Automatically migrate infrequently used data sets” on page 308

Obtaining an overview of space usage

Follow this procedure if you want to determine the amount of space that is being used by your applications. Initially you can create a group of your critical data sets, possibly dividing them into development, test, and production data sets.

1. Use the processes described in “Adding a data set group” on page 298 to add a new data set group. In the case of the example shown in Figure 2 on page 301, there are four groups, for development, test, production, and history data sets.

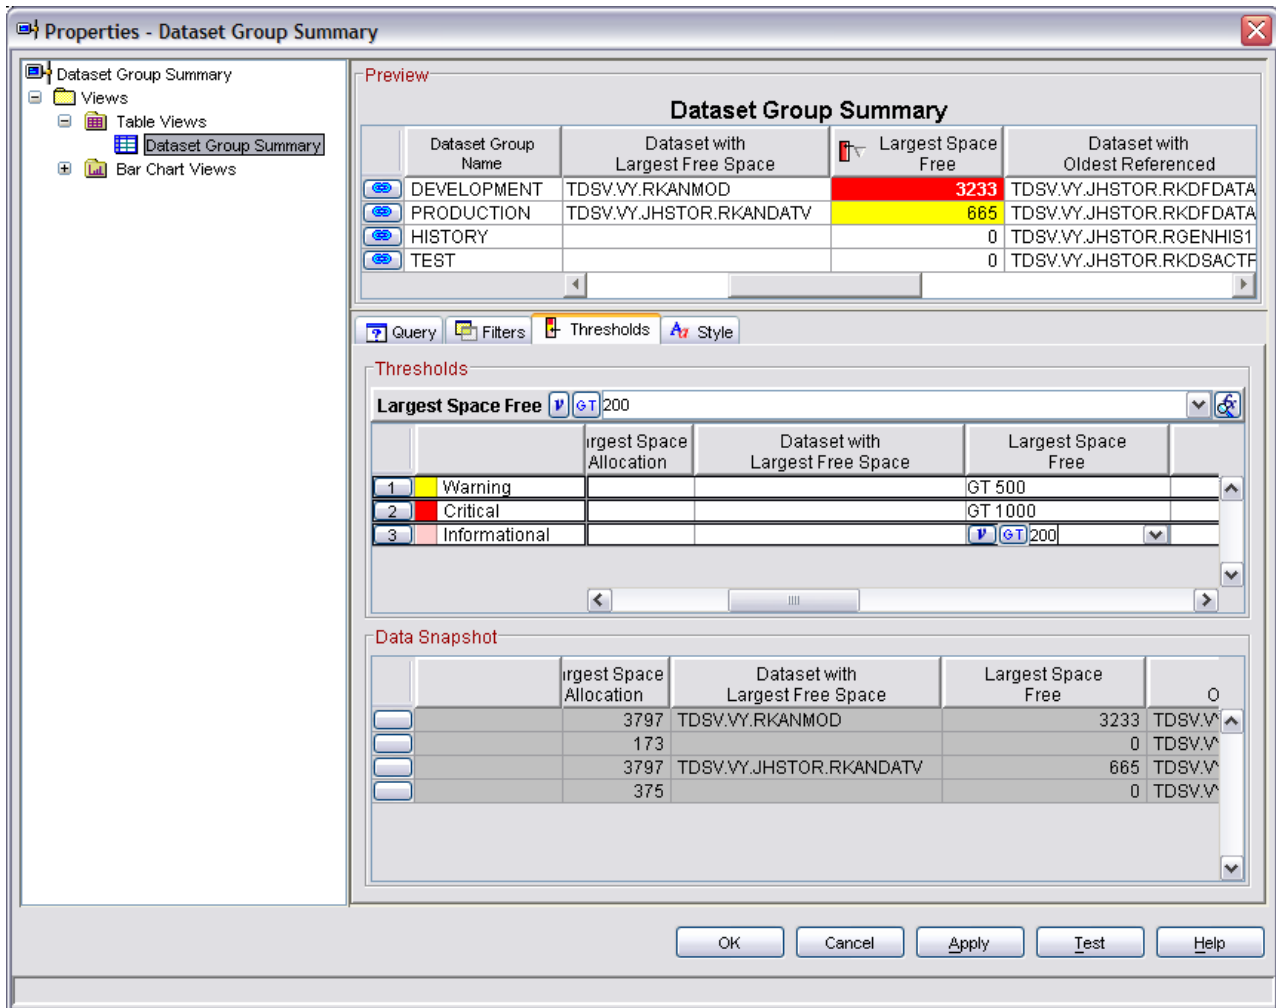


Figure 2. Dataset Group Summary workspace showing four new data set groups

Identifying data sets with excess space

Data set groups can help you identify data sets for which you have allocated too much space. If DFSMS Data and Storage Class ACS routines are not configured to enforce space restrictions or release unused primary space, the aggregate amount of allocated and unused space across the enterprise can be large. Your company might waste resources on unnecessary storage upgrades.

To monitor for this condition, perform the following steps:

- Create a group of data sets.
- Create a threshold to monitor the Largest Free Space for those data sets.
- After you have isolated a group of data sets, it is possible to analyze individual data sets.

Note: You can also use the Dataset Attribute Summary workspace to monitor space usage on a system-wide basis if the group orientation does not provide sufficient scope for your storage management activities.

To set up this threshold do this:

1. In Figure 2, right-click anywhere within the Dataset Group Summary table, click **Properties** to open the Properties editor.

2. Click the **Thresholds** tab.
3. Select the Largest Space Free attribute and change the operator in the first row to **Greater than (>)** and the value to **1000** or a reasonable value for your system. Check that the indicator is set to critical.
4. Repeat the process on the second and third line to set a threshold that indicates an acceptable warning and informational level.
5. Click **Apply** and **OK**.

Whenever the Largest Free Space value for the data set group exceeds the specified thresholds, that group is highlighted in the table view. You can analyze the resulting list of data sets and take appropriate action.

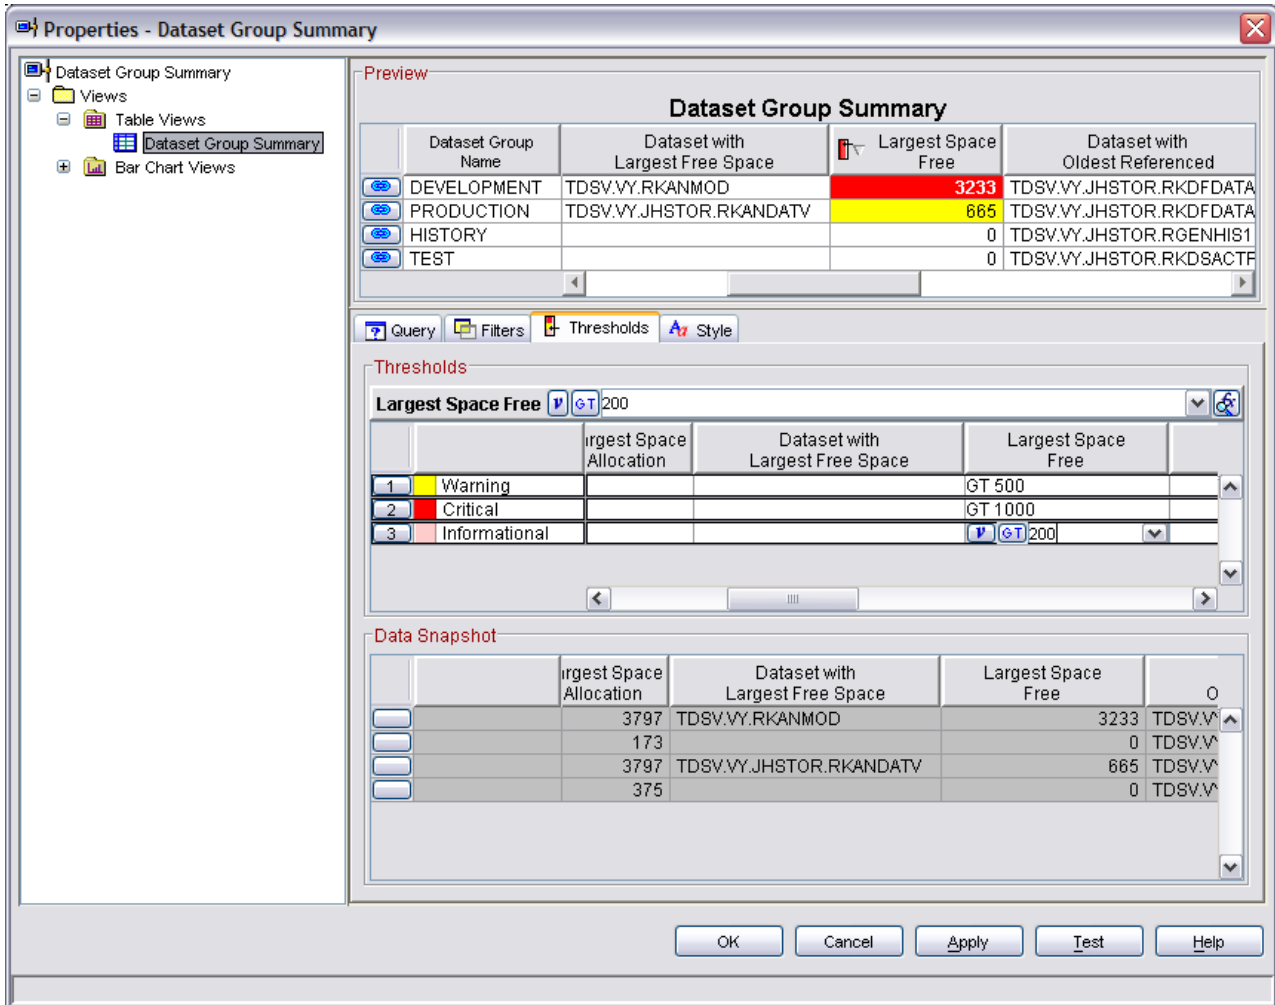


Figure 3. Properties panel of Dataset Group Summary workspace showing a new threshold

Create an alert when response time exceeds specific thresholds


Response times are critical to a business. The Application Monitoring feature of OMEGAMON is useful when you know the name of the JOB, STC, or TSO user whose data sets need to be monitored. There are other cases, however, when it is a group of data sets you need to monitor, no matter how or where they are accessed.

In this case, you can perform the following steps:

- Create a group called, for example Critical, encompassing the critical data sets.

- Define both Group Summary thresholds and situations to generate an alert when response time (MSR) or the number of I/O events exceeds the expected limits.

Create a situation as follows:

1. Right-click the Dataset Group Summary node in the navigation pane.
 2. Click **Situations** in the pop-up menu. The Situations for - Dataset Group Summary dialog box is displayed.
 3. Click the **Create new Situation** icon () to access the Create Situation dialog box.
 4. Enter the name and description of your new situation.
 5. Click **OK** to access the Select condition dialog box.
 6. Select **S3_Dataset_Group_Details** in the Attributes Group column.
 7. While pressing the **Ctrl** key, select both **Dataset Group Name** and **Maximum MSR** in the Attribute Item column.
 8. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
 9. In the **Formula** tab, enter appropriate values in the fields. For example, enter an operator of **Equal (==)** and a value of "Production" for the Dataset Group Name, and an operator of **Greater than (>)** and a value of 20 for the Maximum MSR, see Figure 4 on page 304.
 10. Select the **Distribution** tab.
 11. Add the appropriate managed systems to the **Assigned systems** list for this situation.
 12. Click **Apply** and **OK** to complete the definition your new situation and close the dialog box.
- Whenever a response time exceeds specific thresholds, this situation is triggered and an alert is sent.

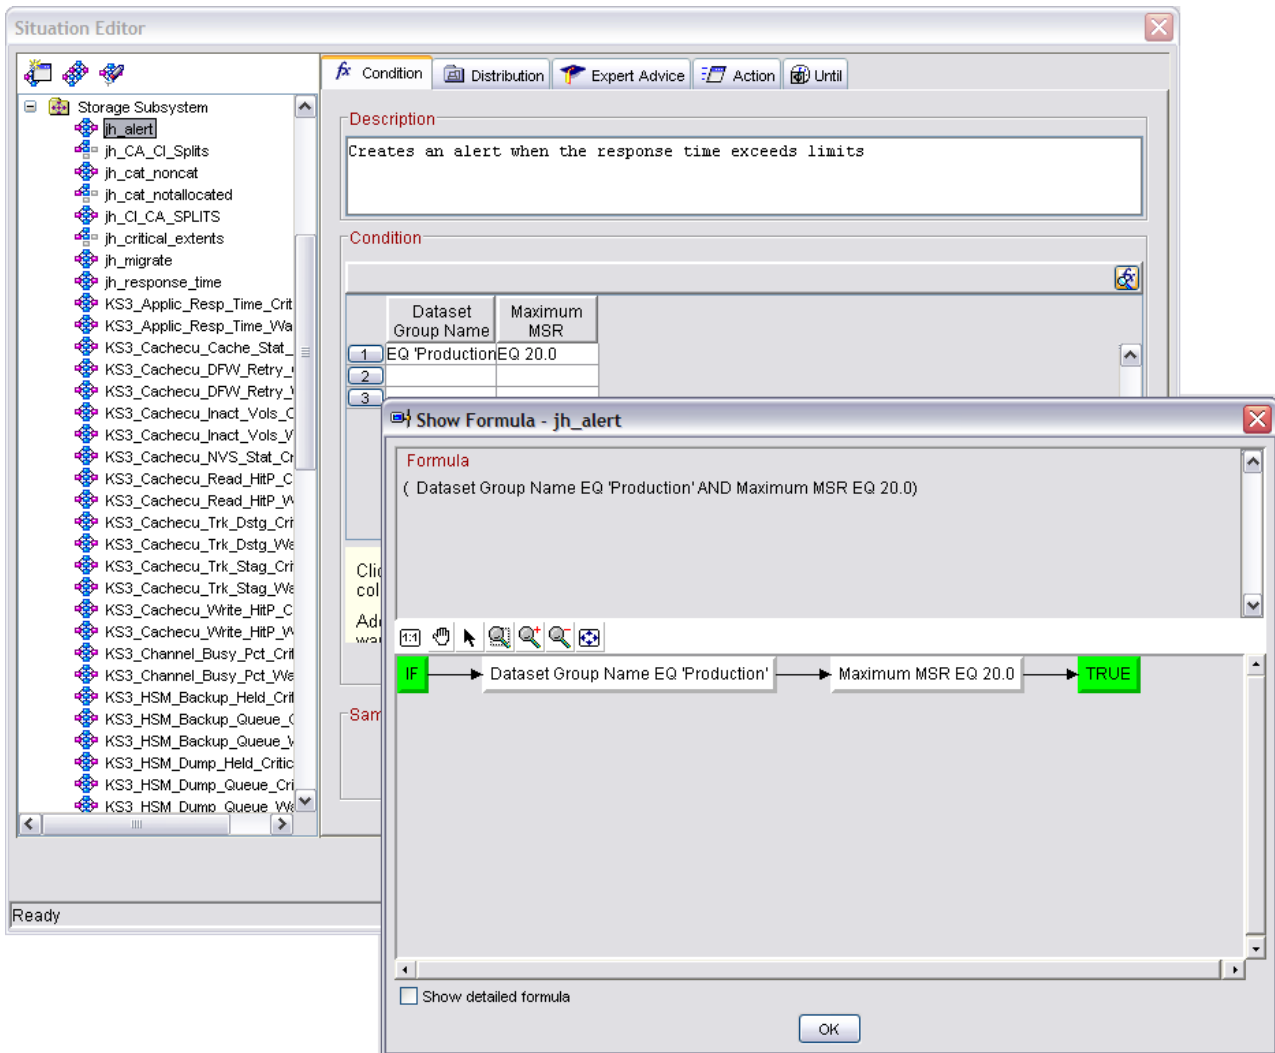


Figure 4. Creation of an alert in the Dataset Group Summary workspace

Identify VSAM data sets with excessive CA/CI splits



VSAM data sets undergo CI and CA splits as they are updated and records are added or inserted. This behavior is not always a bad thing, however for some data set update patterns it can cause excessive processing and I/O.

Note: You can also use the CA/CI Split Summaries from the Dataset Attribute Summary workspace to see the data sets with the largest number of splits across the entire enterprise, rather than limited to defined groups.

To monitor for this condition, perform the following steps:

- Define a group encompassing the VSAM data sets that you know do not experience high split rates.
- Create a Situation to generate a warning alert when any data set in that group exceeds the threshold.
- Using Action or Workflow you can also automatically issue a command that would be trapped by, for example, AOC to cause the data set in question to be REPROed and reDEFINED to eliminate the splits.

Assuming that you have identified and grouped the VSAM data sets that you want to monitor, create a situation as follows:

1. Right-click the Dataset Group Summary node in the navigation pane.
2. Click **Situations** in the pop-up menu. The Situations for - Dataset Group Summary dialog box is displayed.
3. Click **Create new Situation** icon () to access the Create Situation dialog box.
4. Enter the name and description of your new situation.
5. Click **OK** to access the Select condition dialog box.
6. Select **S3_Dataset_Group_Details** in the Attributes Group column.
7. While pressing the **Ctrl** key, select both **CI Splits** and **CI Splits** in the Attribute Item column.
8. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
9. In the **Formula** tab, enter appropriate values in the fields. Change the operators to **Greater than (>)** and enter the value **8** for **CA Splits** and **4** for **CI Splits**.
10. Click **Apply**.
11. To view the formula, click  in the right-hand side of the Situation editor panel.
12. Select the **Distribution** tab.
13. Add the appropriate managed systems to the **Assigned systems** list for this situation.
14. Click **Apply** and **OK** to complete the definition your new situation and close the dialog box.
Whenever a excessive CA/CI splits exist, this situation is triggered and correction action can occur automatically.

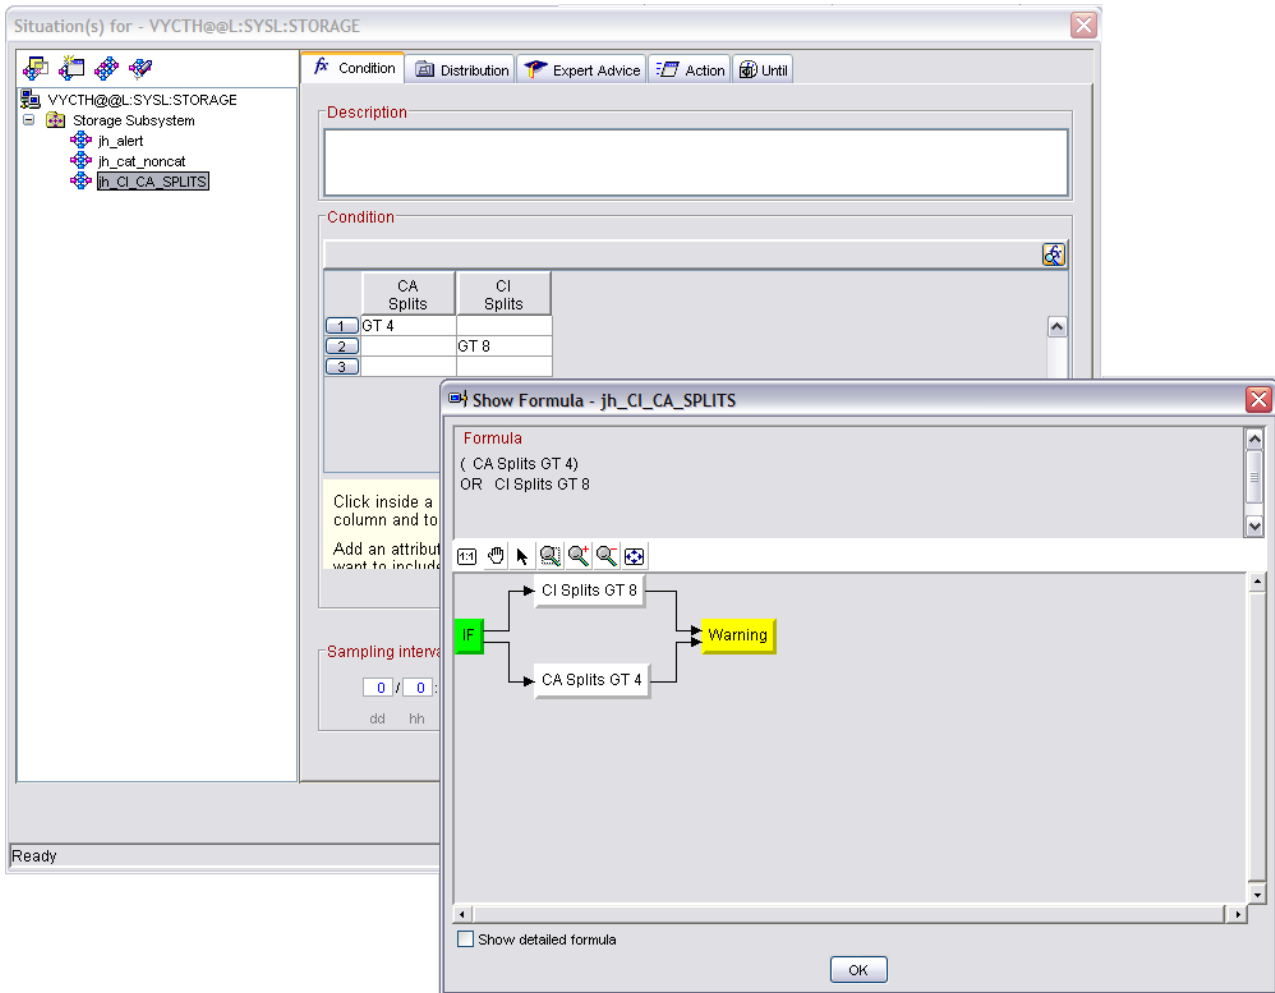


Figure 5. Creating a situation for a data set group for VSAM attributes workspace

Cataloged and not-allocated data sets

Catalog all data sets in a DFSMS environment. Although you follow this policy, there can still be cases where a data set exists in the catalog, but has been moved or deleted from the specified volume.

Note: As a converse to this scenario, you can use the Uncataloged Dataset Summary report of the Dataset Attribute Summary to identify data sets that are allocated, but uncataloged.


To monitor for this condition, perform the following steps:

- Create a data set group.
- Create a situation using the Allocation Flag attribute from the Dataset Group Details Volume attribute group in the Situation editor.

Create a situation as follows:

1. Right-click the Dataset Group Summary node in the navigation pane.
2. Click **Situations** in the pop-up menu. The Situations for - Dataset Group Summary dialog box is displayed.

3. Click the **Create new Situation** icon () to access the Create Situation dialog box.

4. Enter the name and description of your new situation.
5. Click **OK** to access the Select condition dialog box.
6. Select **S3_Dataset_Group_Details_Volume** in the Attribute Group column.
7. Select **Allocation Flag** in the Attribute Item column.
8. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
9. In the **Formula** tab, enter values for Allocation Flag as follows:
 - a. On the first line of the Allocation flag, open the drop-down list and select **CANDIDATE VOLUME/CATALOG ENTRY ONLY**.
 - b. On the second line of the Allocation flag, open the drop-down list and select **OVERFLOW VOLUME/CATALOG ENTRY ONLY**.
10. To view the formula, click  in the right-hand side of the Situation editor panel.
11. Select the **Distribution** tab.
12. Add the appropriate managed systems to the **Assigned systems** list for this situation.
13. Click **Apply** and **OK** to complete the definition your new situation and close the dialog box.

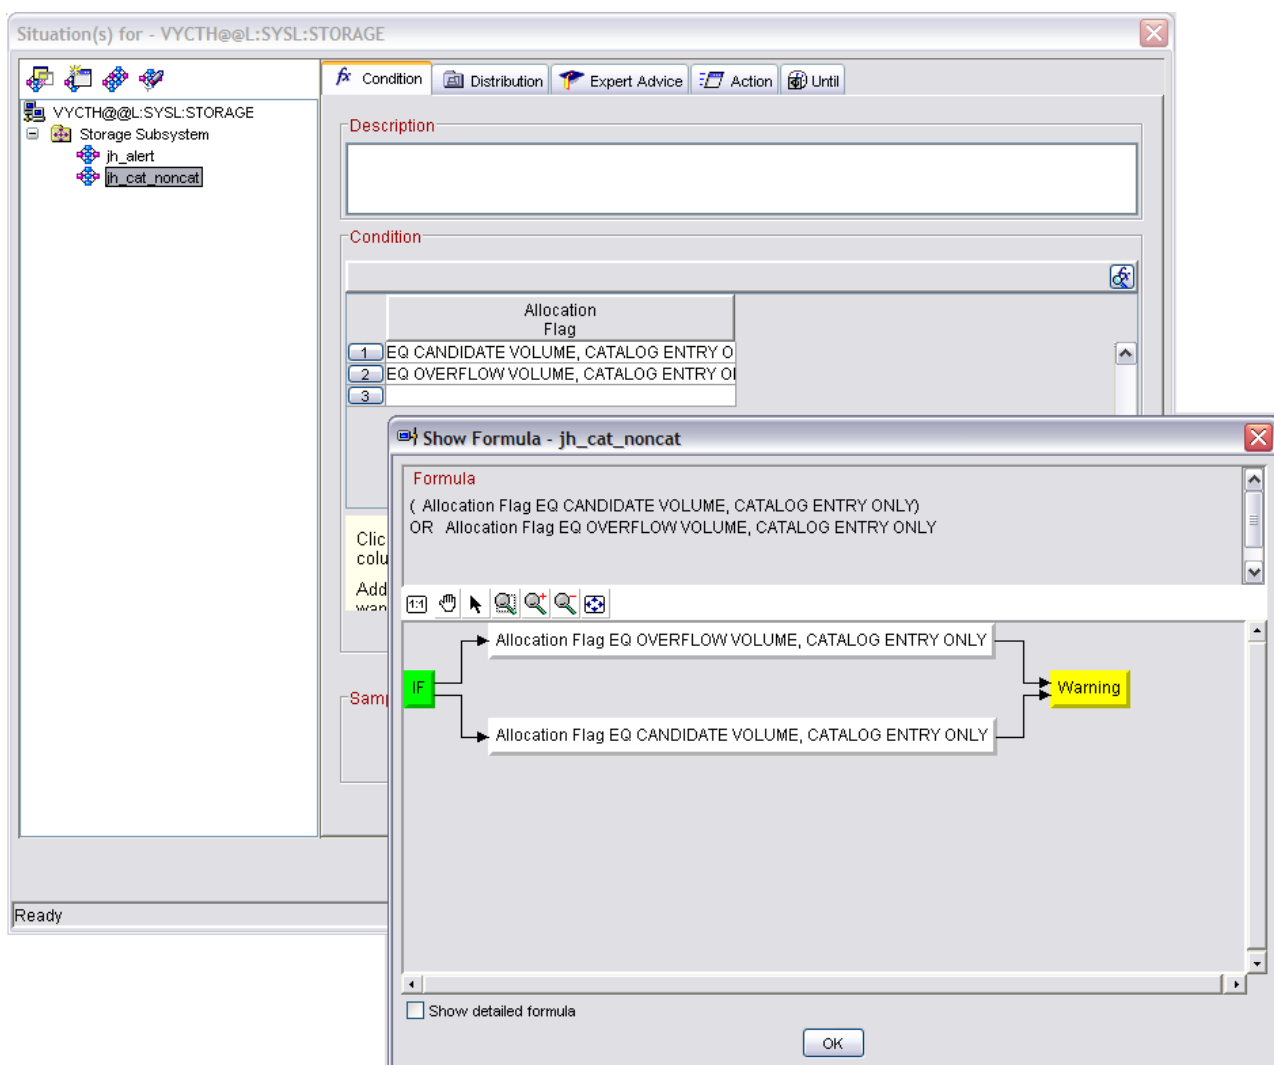


Figure 6. Creating a situation in the Dataset Group Details workspace to monitor cataloged and non-cataloged data sets

Notify the Production Scheduling group when a critical data set exceeds a specified number of extents

Overnight batch jobs have a small window in which they must complete, leaving very little, if any, room for ABENDs and reruns. If a production data set exceeds its available space and causes the batch update job to terminate with an x37 ABEND, online applications might not be able to start in time, and business might be affected.

To monitor for this condition, perform the following steps:


- Create a Dataset Group containing the critical production data sets.
- Define a Situation to raise an alert if any of those data sets approaches the 16 extent limit.
- Add an action to the Situation to send a pager message to the relevant person.

Note: You can use the name of the data set and the number of its extents as substitution variable in the action.

Create the following situation:

1. Right-click the Dataset Group Summary node in the navigation pane.
2. Click **Situations** in the pop-up menu. The Situations for - Dataset Group Summary dialog box is displayed.



3. Click the **Create new Situation** icon () to access the Create Situation dialog box.
4. Enter the name and description of your new situation.
5. Click **OK** to access the Select condition dialog box.
6. Select **S3_Dataset_Group_Details_Volume** in the Attribute Group column.
7. Select **Extents** in the Attribute Item column.
8. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
9. In the **Formula** tab, enter values for **Extents**. Change the operator to **Greater than (>)** and enter the value **12**.
10. Select the **Distribution** tab.
11. Add the appropriate managed systems to the Assigned systems list for this situation.
12. Select the **Action** tab.
13. Click **Universal message** and enter send message &dsname.
14. Click **Apply** and **OK** to complete the definition of your new situation and close the dialog box.

Whenever the number of extents for your critical data set exceeds your threshold, this situation is triggered and notification is sent.

The paging command that you use depends on the underlying product in the Windows® or z/OS environment.

Automatically migrate infrequently used data sets

During the course of a day a large number of permanent data sets are created, possibly accessed once, and then never referenced again. If those data sets are not automatically migrated by DFHSM from disk to near-line storage, they waste valuable space.



To monitor for this condition, perform the following steps:

- Create a Situation using the **Dataset Group Details** attribute group.

- When this situation is true, an HMIGRATE command is issued at the host to cause DFHSM to migrate the data set in question.

Note: This is not a process for replacing Management class and DFSMSHsm processing options but rather a way of augmenting those functions using OMEGAMON intelligent functions.

Create a Dataset Group for permanent data sets that you know are created, accessed once, and should then be migrated. Then create the following situation:

1. Right-click the Dataset Group Summary node in the navigation pane.
2. Click Situations in the pop-up menu. The Situations for - Dataset Group Summary dialog box is displayed.
3. Click the **Create new Situation** icon () to access the Create Situation dialog box.
4. Enter the name and description of your new situation.
5. Click **OK** to access the Select condition dialog box.
6. Select **S3_Dataset_Group_Details** in the Attribute Group column.
7. Select **Primary Volume** in the Attribute Item column.
8. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
9. In the **Formula** tab, enter values for Primary Volume. Change the operator to **Not equal (!=)** and enter the value **MIGRAT**.
10. Click **Add Conditions** to access the Select condition dialog box.
11. Choose **Days Unreferenced** in the Attribute Item column.
12. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
13. In the **Formula** tab, enter values for **Days Unreferenced**. Change the operator to **Greater than (>)** and enter the value **1**.
14. Click **Add Conditions** to access the Select condition dialog box.
15. Choose **Dataset Group Name** in the Attribute Item column.
16. Click **OK** to return to the Situations for - Dataset Group Summary dialog box.
17. In the **Formula** tab, enter values for Dataset Group Name. Change the operator to **Equal (==)** and enter the value **groupname**. For example, the group name could be **PAYROLL_TEMP**. This group can be a specific set of data sets that you know are created and then referenced only once.
18. Select the **Distribution** tab.
19. Add the appropriate managed systems to the Assigned systems list for this situation.
20. Select the **Action** tab.
21. Enter the command you want to use in the System command entry box. For example, you could enter **HMIGRATE &S3_Dataset_Group_Details:Dataset_Name**. To add the variable click **Attribute Substitution** and select **Dataset name**.
22. Select the **Formula** tab and click  which should be similar to Figure 7 on page 310.
23. Click **Apply** and **OK** to complete the definition your new situation and close the dialog box.

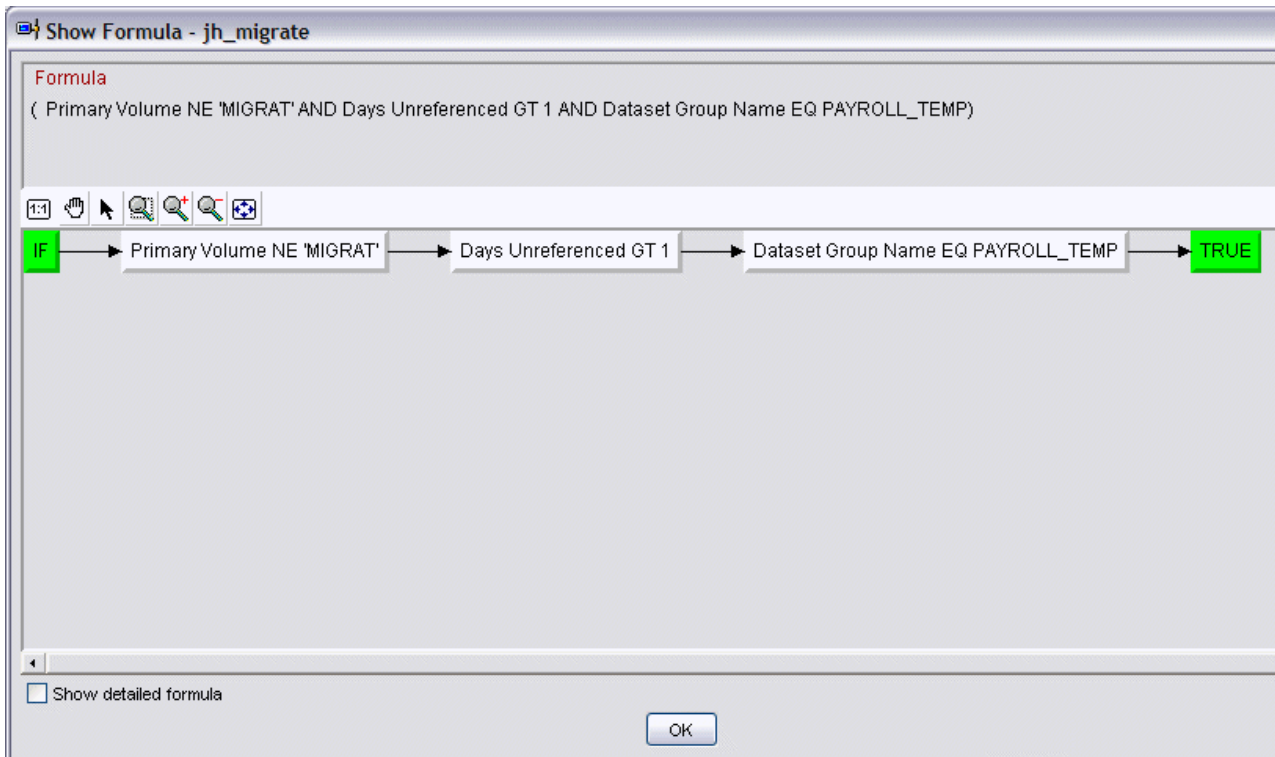


Figure 7. Partial view of a formula that monitors the migration of infrequently used data sets

Chapter 25. Dataset Attributes Database

Tivoli OMEGAMON XE for Storage on z/OS, version 4.2.0, provides the Dataset Attributes Database to help you manage data sets. This database maintains attributes related to data set space, DCB, and cataloging. The “Dataset Attributes System Summary workspace” on page 88 provides the principal view of the database. The principal workspace links to related workspaces that track data for common aspects of data set management and administration. Also see “Configuring the Dataset Attribute Database feature” on page 313.

Benefits

The database collects data regarding all data sets on all volumes in the z/OS environment (except for volumes that you have excluded) so that you can perform the following operations:

- View comprehensive summary reports regarding general trends or specific low-level details.
- Identify exception conditions regarding data sets throughout the environment.
- Select data sets from a workspace view and initiate Storage Toolkit action commands against them.
- Obtain an installation-wide view of data set space utilization, exceptional conditions, and summary statistics.
- Quickly identify resources that require attention, such as data sets that have excessive unused space, extents, or CA/CI splits.
- Avoid the use of utility programs that have limited ability to aggregate and display data, such as IEHLIST.

“The Dataset Group feature versus the Dataset Attributes feature” on page 297 clarifies the strengths of this feature.

Optimization

You must configure the Dataset Attributes Database feature to optimize performance. Otherwise, the amount of data that you collect might consume excessive system resources. See “Configuring the Dataset Attribute Database feature” on page 313 and also “Optimizing data collection for the Dataset Attributes Database feature” on page 314.

For example, you can optimize configuration of the following behaviors:

- Exclusion of specific SMS Storage Groups and Volumes
- Collection interval
- Number of items listed in the Summary reports
- Inclusion of catalog data

Note: Data collected for the Dataset Attributes Database is maintained in an in-memory data space for high-performance access. The data is not saved to disk, except for those Summary tables that have been configured for historical collection. For this reason, following a restart of the Tivoli Enterprise Monitoring Server address space, the database must be rebuilt by initiation of a collection cycle, either manually or at a scheduled time.

Attribute derivation and definition

This section clarifies the concepts and values associated with specific attributes.

Percent used and free

- Percent Used and Free for PDS, PDSE, and HFS data sets is calculated based on the allocated and used tracks from the VTOC, and does not take into account internal usage or wastage within the data set.
- Tracks Used is set equal to Tracks Allocated for PDSE and HFS data sets.
- For VSAM data sets Percent Used and Free is based on the volume-level High Used and High Allocated RBAs (relative byte addresses).

Space values

The following details govern the collection and display of space values in the product:

- Within the Dataset Attributes Database reports and queries, all space values are displayed as Tracks, Megabytes or Gigabytes. By default, the space unit for all delivered Views is Tracks. The default space value columns can be changed as follows:

1. Modify the **Filters** tab of the Properties dialog box for the workspace.
2. Save the workspace under a new name.

Note: Refer to the *IBM Tivoli Monitoring: User's Guide* to learn how to display alternate columns in a predefined workspace.

- Megabytes are defined as 1024*1024. Gigabytes are defined as 1024*1024*1024.
- All calculated megabyte and gigabyte values are calculated from the underlying byte totals, and are rounded up.
- Space values in summary reports for multivolume data sets are for the portion of the data set located on the indicated volume.
- The Dataset Detail workspace lists all individual pieces of a multivolume data set (assuming that all volumes on which it is located were processed during collection and selected by the Query).
- Tracks allocated/used
 - **Tracks Allocated** is always the sum of all tracks in all extents on the volume from the DSCB.
 - **Tracks used for non-VSAM data sets** comes from DS1LSTAR in the DSCB.
 - **Tracks used for VSAM data sets** is calculated as follows:

$$\text{High Used RBA on the volumes} / (\text{VSAM physical record size}) * (\text{VSAM physical records per track})$$

That value is then used to calculate Percent Used and Percent Free and to calculate the total megabytes and gigabytes used.

Note: Other utilities such as DCOLLECT might assume that VSAM data sets are either 100% Used or 100% Free. If you are comparing aggregate results between utilities (for example, with enterprise-wide space values) be aware that the OMEGAMON calculation is intended to give a more accurate picture of space usage within VSAM data sets.

- The SYS1.HASPACE data set (also called the JES2 spool) is always assumed to be 100% Used (Space Used=Space Allocated).

Catalog Values

The following details pertain to catalog values in the product:

- The VTOC Index (SYS1.VTOCIX.Vnnnnnn) and VVDS (SYS1.VVDS.nnnnnn) are not included in the Uncataloged Summary.
- Because the Dataset Attribute database is built starting at the Volume level, there might be entries in a Catalog that are not represented in the Catalog Detail display. The entries that are not displayed can include the following:
 - Data sets that are in the catalog but do not physically exist on any volume
 - Tape data sets

- VSAM Clusters that do not physically exist on disk
- Datasets that are cataloged in more than one catalog

Use cases

The Dataset Attributes Database feature enables scenarios for resource management, such as the following:

- Identify data sets in Unmovable status or with excessive CA/CI splits that might affect application performance
- Find large data sets that have a small Blocksize value, and reblock them to reduce elapsed time. (The referencing application must support reblocking.)
- Define a situation that raises an alert when the number of extents used by a catalog exceeds a specified value.
- Use the Tivoli Enterprise Portal to immediately fix a data set problem that you have found, as described in Table 31.

Table 31. Using the Tivoli Enterprise Portal to manage data sets

| | |
|--|---|
| Phase One: become aware of a problem and decide how to respond | <ol style="list-style-type: none"> 1. You create a situation to monitor available free space in your storage groups. 2. When the situation fires, you want to find and take action. |
| Phase Two: identify targets and take corrective action * | <ol style="list-style-type: none"> 1. Access any row within any TotalStorage workspace. 2. Click on the Unused column to resort the table by most unused space. 3. Select the candidate data sets, right click, and issue the Release Space action from the Dataset Actions popup menu item. This action initiates a DFDSS RELEASE command for each selected data set. |
| <p>* The Dataset Attributes Database and Storage Toolkit features that you use in Phase Two replace procedures that require more effort:</p> <ul style="list-style-type: none"> • If you did not have the Dataset Attributes Database feature, you would have to look at a workspace for each volume in the storage group, one at a time. Also, you would have to compile and sort the list of data sets for each volume to determine which ones have the most unallocated space. • If you did not have the Storage Toolkit feature, you would have to issue DFDSS RELEASE for each data set individually. See Chapter 23, "Storage Toolkit," on page 251. | |

Configuring the Dataset Attribute Database feature

The Dataset Attributes Database feature is turned off by default. It is recommended that you first run a collection on a subset of volumes to test how the feature affects performance. Use the information from this test session to estimate the impact of collection on the address space of the Tivoli Enterprise Monitoring Server and on the overall z/OS system. This estimate can help you choose optimal configuration settings, such as, the volumes and storage groups to exclude from collection, what time of day to start collection of data and what deadline (if any) to impose for completion of collection.

You use this dialog box to Use the Dataset Attribute Collection Control dialog box to configure data collection for the Dataset Attribute Database feature. Right-click in a row of the Dataset Attribute Summary workspace and select **Dataset Attribute Collection Control** to launch the dialog box.

Table 32. Dataset Attribute Data Collection dialog box

| Section in the dialog box | Comments |
|---|---|
| Data Collection settings: | |
| <ul style="list-style-type: none"> • Start data collection at (<i>hh:mm</i>) | When you select this check box, most of the other fields in the dialog box become active for you to configure. <i>hh</i> is an integer from 1 to 12 and <i>mm</i> is an integer from 0 to 60. |

Table 32. Dataset Attribute Data Collection dialog box (continued)

| Section in the dialog box | Comments |
|---|---|
| <ul style="list-style-type: none"> Stop data collection at (hh:mm) | <p>Use this setting to prevent data collection from continuing into a time period when critical jobs run that require significant amounts of system resources. <i>hh</i> is an integer from 1 to 12 and <i>mm</i> is an integer from 0 to 60.</p> <p>If you specify a stop time and collection is halted because that time is reached, attributes and reports for every data set processed up to the point that collection was halted are available for viewing from the Tivoli Enterprise Portal.</p> |
| <ul style="list-style-type: none"> Collect catalog data | <p>When you deselect this check box no catalog information is obtained for any data set.</p> <p>Note: By deselecting this option, you reduce resource consumption and elapsed time. However, several summary reports are empty when you do not collect this data, for example, CA/CI Split Summaries and Catalog/Uncataloged Summaries.</p> |
| <ul style="list-style-type: none"> Number of data rows to collect for Top N reports | <p>The default is 20 rows. The range is 1 to 100 rows. Many of the reports for the Dataset Attribute Database feature involve Top N data. You can lower or raise this setting to adjust the amount of impact that this feature has on system resources versus the number of rows returned in each Summary report.</p> <p>For data sets that have the same key value as the smallest key in the table (for example, lowest number of allocated tracks in the Top N Space Allocated table), the first data set encountered with that value is the one that is displayed in the report.</p> <p>Defaults for Top N data collection are as follows:</p> <ul style="list-style-type: none"> An inefficient block size is defined as blksize < 128. A mature data set is defined as days since last reference > 366. |
| Exclusions: | |
| <ul style="list-style-type: none"> Volumes Storage Groups | <p>You can specify VOLSER or storage group names to exclude volumes from data collection. A volume that matches either exclusion specifications is excluded. You can specify up to 32 exclusions of each type.</p> <p>You can specify a partial name pattern and place the wildcard character (*) at the end. Volumes and Storage Groups whose names match the pattern are excluded.</p> |
| Data Collection Status: | |
| <ul style="list-style-type: none"> Status message Force Start button or Force Stop button | <p>Status messages include, <i>Collection not active</i>, <i>Collection in progress</i>, <i>Collection waiting</i>, <i>Collection failed</i>, <i>Collection halted due to deadline</i>, and <i>Collection halted by user or operator</i>.</p> <p>Click Force Start to manually force data collection to start. A Force Stop button is available while data collection is running to manually force data collection to stop. Results are displayed, including all data sets processed up to the point that collection was halted.</p> |

Optimizing data collection for the Dataset Attributes Database feature

Follow these guidelines to optimize use of the Dataset Attribute Database feature:

- In a multi-system, shared hard-disk-drive environment, only enable Dataset Attribute collection on one of the connected systems, ideally a non-production system. This approach avoids collection of the same data set space information on multiple systems, and places the work for collection on the most lightly loaded system.
- Use product-provided queries and workspaces rather than user-created, as they have been crafted and tuned for maximum performance. User-defined queries can incur significant overhead if a very large number of data sets needs to be examined to resolve the predicates.
- User-created situations written against the Dataset Detail table run against the complete data set database. Do not set a Situation Sampling Interval smaller than the Dataset Attribute Database collection cycle, because attributes in the database are not updated more often than that.

Command-line option

If you have the authority to issue z/OS operator commands, you can use the command-line interface to manage data collection for the Dataset Attributes Database feature as follows:

F temsname,S3DA command_function

where *temsname* is the name of the Tivoli Enterprise Monitoring Server started task and *command_function* is one of the following:

| Command function | Purpose |
|------------------|---|
| START or COLLECT | Immediately start a collection cycle. Previously defined volume and storage group exclusion lists are used, if any exist. |
| STATUS | Display the current status of a Dataset Attributes Database collection cycle that is running. |
| HALT or STOP | Immediately stop an collection cycle. Attributes and reports for every data set are processed up to the point that the STOP command was issued are available for viewing from the Tivoli Enterprise Portal. |
| TERM | Terminate all processing for the Dataset Attributes Database feature. |

You find messages and output from any of these commands in the RKLVLLOG output from the Tivoli Enterprise Monitoring Server.

Custom queries

This section provides tips and suggestions regarding queries that you create for the Dataset Attributes Database feature:

- Queries that you create for the Data Set Details table (a query that involves the entire data set attribute database) must have at least one predicate specified. If you use only one attribute for selection, you cannot use the attribute in conjunction with STR or COUNT.
- You must use uppercase characters for the character string values in your queries, such as data set name or volume serial number. Otherwise, the data is not displayed in the Tivoli Enterprise Portal.
- A single custom query can return a maximum of 16000 rows. This behavior prevents storage shortages and potential abends in the Tivoli Enterprise Monitoring Server. A query (for example, DSORG=PO) might match a large number of data set detail rows. This processing load can negatively affect all other IBM Tivoli Monitoring applications.
- When you specify the predicate for a custom query, position the most specific attribute as the first column so that it functions as the primary index. The primary index limits the search for the additional attributes and, if specified with sufficient granularity, can greatly reduce response time.

For example, if you want to see all data sets that are greater than 1000 cylinders on volume PRI000, specify the VOLUME attribute as primary index, so that the query initially searches for data sets from volume PRI000 only. (If you specify the number of cylinders (1000) as the primary index, the query

finds every data set in the z/OS environment that is larger than 1000 cylinders and then selects those that are located on PRI000. This query would be much more time consuming.)

- The string 00:00:00 functions as a wildcard specification in queries that use Creation Date, Last Reference Date or Expiration date. Specify the Time value of such queries to be 00:00:00 because that value does not have a time associated with it and the returned converted times must not contradict the query specification.
- The following examples show the potential value of queries:
 - To see all components associated with a specific VSAM Cluster, enter the cluster name in the Associated Dataset attribute column.
 - Find VSAM data sets that have REPLICATE or IMBED.
 - Locate Uncataloged VSAM data sets.
 - Identify catalogs that have exceeded a threshold number of extents.

Chapter 26. Launching of IBM TotalStorage Productivity Center

IBM Tivoli Monitoring provides an Application Launch feature. Tivoli OMEGAMON XE for Storage on z/OS uses this feature in selected workspaces to provide launching of the IBM TotalStorage Productivity Center.

About the IBM TotalStorage Productivity Center

The IBM TotalStorage Productivity Center is a software tool that helps you manage complex, heterogeneous open storage environments to increase storage capacity utilization, and to improve administrative efficiency. Typically this software is integrated with a disk subsystem, such as a DS8000 storage server.

About the launch feature in Tivoli OMEGAMON XE for Storage on z/OS

In Tivoli OMEGAMON XE for Storage on z/OS, the launch feature provides convenient access to the IBM TotalStorage Productivity Center. The feature enables the following example scenario:

| | |
|--|---|
| Phase One: become aware of a problem and decide how to respond | <ol style="list-style-type: none">1. You suspect that a DS8000 storage facility has slow response time.2. You decide to investigate hardware issues in IBM TotalStorage Productivity Center. |
| Phase Two: identify targets and take appropriate corrective action * | <ol style="list-style-type: none">1. You right-click the workspace to access the context menu.2. You select Launch and immediately access the sign-on screen for the instance of IBM TotalStorage Productivity Center for the storage resource that you want to investigate.3. You discover a need to allocate space to the extent pool that generated the alert.4. You allocate some unallocated storage space to the extent pool. |
| * If you did not have the launch feature, you must complete the following manual process: log on to the mainframe computer, find the affected resource, and find the HTML page for IBM TotalStorage Productivity Center and launch it. | |

Launching the IBM TotalStorage Productivity Center

IBM TotalStorage Productivity Center is a browser-based software environment. Tivoli OMEGAMON XE for Storage on z/OS calls the URL of the browser page to make a connection. You run the **KS3LaunchConfig** command to configure the launch feature of Tivoli OMEGAMON XE for Storage on z/OS. The command identifies the location of the browser software and the URL of the IBM TotalStorage Productivity Center. See “Configuring the launch feature” on page 318.

You launch the IBM TotalStorage Productivity Center as follows:

1. Right-click on a resource that is displayed in a workspaces. See “Workspaces that provide the launch feature” on page 318.
2. Select **Launch** in the context menu. The Create or Edit Launch Definitions dialog box is displayed.
3. Select the instance of the IBM TotalStorage Productivity Center that you want to launch.
4. Click **Launch**. The sign-on screen for IBM TotalStorage Productivity Center is displayed in a Web browser page. After you sign in, you can use the features of IBM TotalStorage Productivity Center to manage the hardware resource that you have selected in the workspace.

Workspaces that provide the launch feature

This section lists the Tivoli OMEGAMON XE for Storage on z/OS workspaces that provide the launch feature. You access these linked (secondary) workspaces through links in either of two primary workspaces: “Cache CU Performance workspace” on page 23 or “Cache CU Status workspace” on page 34. The secondary workspace links are available when you have selected a DS8000 resource in a workspace view.

- TotalStorage Configuration
 - TotalStorage Ranks
 - TotalStorage Array Configuration
 - TotalStorage Rank Trend
 - TotalStorage Extent Pool Volumes
 - TotalStorage Volume Trend
 - TotalStorage Extent Pool Trend
 - TotalStorage Rank History
 - TotalStorage Volume History
- TotalStorage CU Volumes
 - TotalStorage Volume Trend

Configuring the launch feature

You use the command-line interface to maintain the launch points (insert, update, and remove operations). The available launch points are predefined by Tivoli OMEGAMON XE for Storage on z/OS. Perform the following steps to configure the ability to launch IBM TotalStorage Productivity Center in Tivoli OMEGAMON XE for Storage on z/OS:

1. Access a command-line interface.
2. Navigate to the directory where the command is stored. By default the command is stored in the following path: C:\IBM\ITM\CNP\
C:\IBM\ITM\CNP\
C:\IBM\ITM\CNP\
3. Run the **KS3LaunchConfig** command to insert, change, or remove target points for the launch feature.

Note: Before you run the **KS3LaunchConfig** command you must navigate in the command line interface to the directory where the command is stored. By default the command is stored in the following path: C:\IBM\ITM\CNP\
C:\IBM\ITM\CNP\
C:\IBM\ITM\CNP\

- To add a new launch target use the following flags:

```
KS3LaunchConfig -insert -loc=LOCATION -url=URL
```

where,

- *LOCATION* is the local path of the browser application that you use to access the IBM TotalStorage Productivity Center. For example, the entry for Microsoft® Internet Explorer might be "C:\Program Files\Internet Explorer\iexplore.exe".
- *URL* is the URL of the browser page for IBM TotalStorage Productivity Center.

Note: If the *LOCATION* parameter or the *URL* parameter contains embedded blank spaces, you must enclose the parameter in quotation marks.

- To make a change or add a launch target use the following flags:

```
KS3LaunchConfig -change -loc=LOCATION -url=URL
```

where,

- *LOCATION* is the local path of the browser application that you use to access the IBM TotalStorage Productivity Center. For example, the entry Microsoft Internet Explorer might be "C:\Program Files\Internet Explorer\iexplore.exe".

- *URL* is the URL of the browser page for IBM TotalStorage Productivity Center.

Note: If the *LOCATION* parameter or the *URL* parameter contains embedded blank spaces, you must enclose the parameter in quotation marks.

- To remove all of the launch target points that were created by this command use the following flags:

```
KS3LaunchConfig -remove
```

The following command illustrates the use of the command:

Note: Before you run the **KS3LaunchConfig** command you must navigate in the command line interface to the directory where the command is stored. By default the command is stored in the following path:
C:\IBM\ITM\CNP\.

```
KS3LaunchConfig -insert -loc="c:\Program Files\Internet Explorer\iexplore.exe" -url=http://your.TPC.url
```

where *your.TPC.url* represents the URL of the browser page for IBM TotalStorage Productivity Center.

Chapter 27. Reports available for IBM Tivoli OMEGAMON XE for Storage on z/OS

IBM Tivoli OMEGAMON XE for Storage on z/OS version 4.2 includes reports that run under Tivoli Common Reporting, a reporting tool and strategy common across Tivoli products. Tivoli Common Reporting provides a consistent approach to viewing and administering reports. This reporting environment runs on Windows, Linux, and UNIX. For more information about Tivoli Common Reporting platforms, refer to the *Tivoli Common Reporting: User's Guide*. To learn more about how the OMEGAMON XE monitoring agents on zSeries use Tivoli Common Reporting, refer to *IBM Tivoli OMEGAMON XE and IBM Tivoli Management Services on z/OS: Reports for Tivoli Common Reporting*.

The reports provided with IBM Tivoli OMEGAMON XE for Storage on z/OS are historical reports run against data collected on a DB2-based Tivoli Data Warehouse 6.2 Fix Pack 1 or later running on Windows. Reports for IBM Tivoli OMEGAMON XE for Storage on z/OS are provided as *report packages*, zipped archives containing reports, documentation, graphics, and dynamic link libraries. The IBM Tivoli OMEGAMON XE for Storage on z/OS report package was shipped as a zip file on the Application CD in the REPORTS directory, and the REPORTS directory is divided into subdirectories named with the three-character prefix that identifies the product (the prefix for IBM Tivoli OMEGAMON XE for Storage on z/OS is ks3). For example, on a Windows machine, if the DVD/CD drive is labelled D:, reports would be in directories such as: D:\REPORTS\ks3.

These reports are useful for the following IT professionals:

- The network systems programmer who troubleshoots TCP/IP issues
- The application analyst or ID manager
- The IT manager or service level advisor who validates service level agreements
- The capacity planner
- The IMS™ service manager
- The z/OS system administrator
- The CICS® system administrator
- The storage administrator

Using these reports

Before you can use the reports in this report package, do the following:

- Install Tivoli Common Reporting, using the information found in the *Tivoli Common Reporting: User's Guide*.
- Ensure that your environment meets the requirements described in the "Prerequisites" section of the *IBM Tivoli OMEGAMON XE and IBM Tivoli Management Services on z/OS: Reports for Tivoli Common Reporting*.
- Perform the setup functions described in the *IBM Tivoli OMEGAMON XE and IBM Tivoli Management Services on z/OS: Reports for Tivoli Common Reporting*:
 - Ensuring that historical reporting is enabled
 - Importing the report package
 - Configuring the data source
 - Generating a sample report

Refer to the "Troubleshooting" section of the *IBM Tivoli OMEGAMON XE and IBM Tivoli Management Services on z/OS: Reports for Tivoli Common Reporting* book if you have any problems performing these tasks.

To develop reports of your own or edit the reports provided with this product, you need an Eclipse BIRT Report Designer, which is a non-charge, open-source tool. This tool is not included with Tivoli Common Reporting, but can be downloaded from <http://www.eclipse.org/birt/phoenix/> or from the Tivoli Common Reporting page at IBM developerWorks® (<http://www.ibm.com/developerworks/spaces/tcr>). You also need the *Tivoli Common Reporting: Development and Style Guide* on the IBM developerWorks Web site: <http://www.ibm.com/developerworks/spaces/tcr>.

Typical types of reports available with Tivoli Common Reporting

A report can either be run on demand or a snapshot can be created at any time for viewing later. An *on-demand* report is a formatted version of a report run with the currently available data and viewed immediately. When you run an on-demand report, you can specify the parameter values to use; the report is then formatted using the current data and displayed in the report viewer.

In addition to creating on-demand reports, you can also create, save, and access *snapshot* reports, saved versions of a report using data from a particular time. Snapshots are useful for generating unattended reports in advance. Snapshots are of particular value in avoiding the wait for a big report to run. Reports may take a long time to run when you request a huge amount of data, such as the past month of real time measurements. Refer to the *IBM Tivoli Common Reporting User's Guide* for information about this report type or look in the online help for Tivoli Common Reporting.

Because Tivoli Management Services supports the summarization and pruning of data, many OMEGAMON XE packaged reports can also generate *summarized* reports. If a packaged report supports summarized data and the Summarization Agent has been configured to collect data for the attribute group required (at the Tivoli Enterprise Portal console), then selected reports will provide an option to specify a summarization period (Hourly, Daily, Weekly, Monthly, Quarterly, or Yearly). The resulting reports reflect data summarized at that level. You can change the summarization period dynamically from a generated report without completing the parameters definition dialog again by clicking a different summarization period in the Available Summarization Time Periods area of a summarized report. For more information about summarization and pruning, see the *IBM Tivoli Monitoring: Administrator's Guide*.

The reports for IBM Tivoli OMEGAMON XE for Storage on z/OS are described in tables that follow. The report description tables include the following information:

- A description of the report
- The most logical users of this report
- The default attribute graphed in the report
- Other attributes from the same attribute group and workspace that you could graph instead
- Resource identifiers that form the matrix for the report and are not selectable
- Other resource identifiers you can specify that act as filters for the data shown in the report. The choices in this cell usually correspond to the choices in the Report Parameters dialog.
- The name of the workspace that displays the same data as this report
- The attribute group or table in Tivoli Data Warehouse that is the source of the attributes
- The types of reports available

OMEGAMON XE for Storage on z/OS reports

This section lists the Tivoli Common Reporting reports available for OMEGAMON XE for Storage on z/OS:

- Storage HSM Migrate and Recall Function Requests - Summarized report
- Storage System Volume and Dataset DASD Usage - Summarized report
- Storage Dataset Group DASD Usage - Summarized report

Chapter 28. Working with the DFSMSrmm Status feature

This topic describes key concepts and capabilities of the DFSMSrmm Status feature.

Tivoli OMEGAMON XE for Storage on z/OS, Version 4.2.0, provides a DFSMSrmm monitoring and analysis function to help you manage the critical DFSMSrmm subsystem. The DFSMSrmm Status workspace provides a high-level overview of DFSMSrmm statistics and links to related workspaces that track data for common aspects of DFSMSrmm management and administration.

Related information

“RMM Collection Control dialog box” on page 326

You use this dialog box to configure the data collection settings for the DFSMSrmm feature.

Chapter 15, “Monitoring DFSMSrmm status,” on page 77

This topic describes the workspaces that monitor the status of DFSMSrmm resources. You access these workspaces through the DFSMSrmm Status node in the Tivoli Enterprise Portal.

Benefits

This topic describes the benefits of the DFSMSrmm Status feature.

The DFSMSrmm Status feature analyzes and reports on entries in the DFSMSrmm Control Data Set (CDS) regarding all data sets on all tape volumes that DFSMSrmm is managing. This feature allows you to perform the following operations:

- View comprehensive summary reports regarding general trends or specific low-level details:
 - Locations Owners
 - Accounting codes
 - Creating Jobs or Programs
 - Vital Record Specifications (VRS)
 - Volumes and Datasets
- Identify exception conditions regarding DFSMSrmm CDS space and backup, scratch tape counts, configuration state, and so on.
- Select tape volumes, data sets, owners, or VRSs from a workspace view and initiate Storage Toolkit action commands against them.
- Obtain an installation-wide view of DFSMSrmm CDS space utilization, exceptional conditions, and summary statistics.
- Quickly identify resources that require attention, such as Pending Action or Pending Move volumes, and initiate that action using the Storage Toolkit.
- Avoid the use of the standard DFSMSrmm reporting utility program that requires user customization.

Optimization

This topic describes how to optimize performance of the DFSMSrmm Status feature.

You should configure the DFSMSrmm Status feature to optimize performance. Otherwise, the amount of data that you collect can consume excessive system resources. For example, you can optimize configuration of the following behaviors:

- Collection interval
- Inclusion of dataset-level attribute data

Also see “RMM Collection Control dialog box” on page 326 and also “Optimizing data collection for the DFSMSrmm Status feature” on page 325.

Note: Data collected for the DFSMSrmm Status feature is maintained in an in-memory data space for high-performance access. The data is not saved to disk, except for those Summary tables that have been configured for historical collection. For this reason, after you restart the Tivoli Enterprise Monitoring Server address space, the data space must be rebuilt by initiation of a collection cycle, either manually or at a scheduled time.

Attribute derivation and definition

This topic clarifies the concepts and values associated with specific attributes of the DFSMSrmm Status feature.

- Only volumes with a type of PHYSICAL and a status of MASTER or USER are included in when calculating the following statistics:
 - Capacity
 - Utilization
 - Compression
- Compression, which is meant to represent the percentage reduction of data written from the host that is actually recorded on the tape volume, is calculated as follows:
$$1 - ((\text{Volume Capacity} * \text{Volume \% Full}) / (\text{App Data written to volume}))$$
- The following attributes include PHYSICAL, LOGICAL, and STACKED volume types:
 - Total volumes
 - Total datasets
 - Average volume use count
 - Average datasets per volume
 - Percent of volumes read/written previous day
- The following dates are shown as NONEXPIRING in the Days To/Days From attributes that are calculated by OMEGAMON:
 - 99000
 - 99365
 - 99366
 - 98000
 - 98365
- The **Days Until Expiration** attribute displays a positive value to represent a date in the future, while a negative the number represents the days since expiration. Similarly, a positive Days for Retention value signifies the remaining number of days that the dataset or volume is retained, while a negative value indicates the number of days since the retention date.
- History can be enabled for the DFSMSrmm Summary, Configuration, and CDS tables only. CDS and Configuration data is collected at the interval specified in the History Collection Configuration dialog box of the Tivoli Enterprise Portal. Summary attributes are written to the history files only once during any collection cycle, because the summary data does not change until the next collection cycle is run.
 - Locations Owners
 - Accounting codes
 - Creating Jobs or Programs
 - Vital Record Specifications (VRS)
 - Volumes and Datasets

Scenarios for using the DFSMSrmm feature

This topic describes typical scenarios for the DFSMSrmm feature.

The DFSMSrmm feature enables scenarios for powerful tape environment monitoring and management, such as the following:

- Using the existing OMEGAMON Tape Devices workspace (under the navigation tree Tape Group node), you can link to the new DFSMSrmm Volume Details workspace to see all DFSMSrmm attributes for a mounted and in-use tape volume.

- When an individual leaves the organization, you can use the Owner Summary workspace and link to the Owner Realtime Volumes workspace to see all volumes owned by that individual. Using the **Storage Toolkit Change Volume** command you could then change the owner name to a current employee.

Table 33. Using the Tivoli Enterprise Portal to manage DFSMSrmm tape volumes and data sets

| | |
|---|---|
| Phase One: Become aware of a problem and decide how to respond | Create a copy of the product-provided situation <code>KS3_RMM_Scratch_Tape_Warning</code> , and set the low and high scratch volume counts to values appropriate for your installation. Start the situation. |
| Phase Two: Identify targets and take corrective action | When the situation is triggered, navigate to the DFSMSrmm Pending Actions workspace, and use the Storage Toolkit Confirm Actions for all Volumes command for any volumes that are in pending scratch state. You could also use the Storage Toolkit Add Volume command to add new volumes to DFSMSrmm. |

Optimizing data collection for the DFSMSrmm Status feature

This topic describes ways to optimize data collection for the DFSMSrmm Status feature.

Follow these guidelines to optimize use of the DFSMSrmm Status feature:

- In a multi-system, shared hard-disk-drive environment where the DFSMSrmm subsystem is running on multiple systems and sharing a single CDS, only enable OMEGAMON DFSMSrmm collection on one of the connected systems, ideally a non-production system. This behavior avoids collection of the same DFSMSrmm information on multiple systems, and places the work for collection on the system that is most lightly loaded.
- Use product-provided queries and workspaces rather than user-created, because they have been crafted and tuned for maximum performance. User-defined queries can incur significant overhead when a very large number of data sets or volumes must to be examined to resolve the predicates.
- User-created situations written against the DFSMSrmm Summary table should not have their Sampling Interval smaller than the DFSMSrmm collection cycle, because attributes are not updated more often than that.

Command-line option

This topic describes the command-line option for the DFSMSrmm Status feature.

If you have the authority to issue z/OS operator commands, you can use the command-line interface to manage data collection for the DFSMSrmm Status feature as follows:

```
F temsname,S3RM command_function
```

where *temsname* is the name of the Tivoli Enterprise Monitoring Server started task and *command_function* is one of the following:

START or COLLECT

Immediately start a collection cycle.

STATUS

Display the current status of a running DFSMSrmm collection cycle.

HALT or STOP

Immediately stop a collection cycle. Attributes and reports for every volume and data set are processed up to the point that the **STOP** command was issued are available for viewing from the Tivoli Enterprise Portal.

TERM Terminate all processing for the DFSMSrmm Status feature.

You find messages and output from any of these commands in the RKLVLOG output from the Tivoli Enterprise Monitoring Server

Custom queries

This topic provides tips and suggestions regarding queries that you create for the DFSMSrmm Status feature.

You must use uppercase characters for the character string values in your queries, such as data set name or volume serial number, otherwise the data is not displayed in the Tivoli Enterprise Portal. A single DFSMSrmm table query can return a maximum of 4000 rows. This behavior prevents storage shortages and potential abends in the Tivoli Enterprise Monitoring Server. A query (for example using the predicate **JOB=***) might match an extremely large number of data set detail rows. This processing load can negatively affect all other IBM Tivoli Monitoring applications.

RMM Collection Control dialog box

You use this dialog box to configure the data collection settings for the DFSMSrmm feature.

Account for the following issues when you install Version 4.2.0 of OMEGAMON XE for Storage to ensure that the DFSMSrmm feature operates correctly:

- You must delete, reallocate, and initialize the RKS3DSA x persistent data stores. Otherwise, the new DFSMSrmm tables cannot be correctly defined for historical collection.
- You must establish security authorization for the Tivoli Enterprise Monitoring Server, as defined in the DFSMSrmm product documentation. Otherwise, commands cannot be issued from the monitoring server started task. Several commands, for example LISTCONTROL, require specific levels of security system authority.

The DFSMSrmm Status feature is turned off by default. It is recommended that you first run a collection cycle off-hours to test how the feature may affect performance. Use the information from this test session to estimate the impact of collection on the address space of the Tivoli Enterprise Monitoring Server and on the overall z/OS system. This estimate can help you choose optimal configuration settings, such as what time of day to start collection of data, what deadline (if any) to impose for completion of collection, and whether to include data-set-level attributes in the collection.

See Chapter 28, “Working with the DFSMSrmm Status feature,” on page 323 for additional information.

Related information

Chapter 28, “Working with the DFSMSrmm Status feature,” on page 323

This topic describes key concepts and capabilities of the DFSMSrmm Status feature.

Chapter 15, “Monitoring DFSMSrmm status,” on page 77

This topic describes the workspaces that monitor the status of DFSMSrmm resources. You access these workspaces through the DFSMSrmm Status node in the Tivoli Enterprise Portal.

Details regarding the Collection Control dialog box

Use the DFSMSrmm Collection Control dialog box to configure data collection for the DFSMSrmm feature. Right-click in a row of the DFSMSrmm Status workspace and select **RMM Collection Control** to launch the dialog box.

Table 34. RMM Collection Control dialog box

| Section/option | Comments |
|-----------------------|----------|
| Collection Scheduling | |

Table 34. RMM Collection Control dialog box (continued)

| Section/option | Comments |
|---|---|
| Start collection at (hh:mm) | <p>Select this check box to activate most of the other fields in the dialog box. <i>hh</i> is an integer from 1 to 12, and <i>mm</i> is an integer from 0 to 60.</p> <p>By default, collection is set to run daily. You deselect specific days of the week in the Every area to prevent collection on specific days.</p> <p>Start collection at causes the backend controller to schedule Summary Collection to begin at the specified time on the specified days of the week. If unchecked, Summary collection does not automatically run.</p> |
| Stop collection at (hh:mm) | <p>Use this setting to prevent data collection from continuing into a time period when critical jobs run that require significant amounts of system resources. You can also use this setting to prevent collection when the DFSMSrmm subsystem itself is in heavy use and requires exclusive access to its CDS. <i>hh</i> is an integer from 1 to 12, and <i>mm</i> is an integer from 0 to 60.</p> <p>If you specify a stop time and collection is halted because that time is reached, attributes and reports for every volume and data set processed up to the point that collection was halted are available for viewing from the Tivoli Enterprise Portal.</p> <p>Stop collection at causes the backend controller to halt the RMM Summary Collection process at the indicated time of day. This response can be useful in environments where the RMM CDS is extremely large and the summary process interferes with normal RMM processing, or where a limit should be placed on the amount of elapsed time/resources consumed by the summary collection process.</p> |
| Collection Options | |
| Summarize RMM data set attributes | <p>When you deselect this check box, no DFSMSrmm information is obtained for any tape data set. You continue to see a total data set count, however, because this metric is obtained by accumulating the total number of data sets on each processed volume.</p> <p>Note: By deselecting this option, you reduce resource consumption and elapsed time. However, several summary report workspaces are empty when you do not collect this data, specifically, Job, Program, and VRS summary. In addition, several attributes that appear on various summaries can also not be calculated:</p> <ul style="list-style-type: none"> • Data sets created previous day • Data created previous day • Data sets read/written previous day |
| Collection Status | |
| Status message: Force Start button or Force Stop button | <p>Status messages include, Collection not active, Collection in progress, Collection waiting, Collection failed, Collection completed, Collection halted due to deadline, and Collection halted by user or operator.</p> <p>Click Force Start to manually force data collection to start. A Force Stop button is available while data collection is running to manually force data collection to stop. Results are displayed that include all volumes and data sets processed up to the point that collection was halted. And the status shown is Collection completed.</p> <p>If you stop the collection by clicking on Force Stop button, the state in the Collection Status area is Collection halted by user or operator. The status is not Stopped.</p> |

To learn about the standard capabilities of a dialog box, see “Standard tabs in the dialog boxes of the Storage Toolkit” on page 258. The topic includes a link to the Reference guides for commands that are used in the Storage Toolkit.

Appendix A. Documentation library

This appendix contains information about the publications in the Tivoli OMEGAMON XE for Storage on z/OS library and about other publications related to Tivoli OMEGAMON XE for Storage on z/OS.

See *IBM Tivoli Monitoring and OMEGAMON XE Products: Documentation Guide*, SC23-8816, for information about accessing and using the publications. You can find the *IBM Tivoli Monitoring and OMEGAMON XE Products: Documentation Guide* in the IBM Tivoli Monitoring and OMEGAMON XE Information Center at <http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/>.

To find a list of new and changed publications, click **What's new** on the Welcome page of the IBM Tivoli Monitoring and OMEGAMON XE Information Center. To find publications for the previous version of a product, click **Previous information centers** on the Welcome page for the product.

Tivoli OMEGAMON XE for Storage on z/OS library

The following documents are available for Tivoli OMEGAMON XE for Storage on z/OS:

- *IBM Tivoli OMEGAMON XE for Storage on z/OS: Planning and Configuration Guide*, SC23-9702
Documents the installation and configuration tasks necessary for the implementation of IBM Tivoli OMEGAMON XE for Storage on z/OS. This document is written for z/OS systems programmers and others who are responsible for installing and configuring IBM Tivoli OMEGAMON XE for Storage on z/OS.
- *IBM Tivoli OMEGAMON XE for Storage on z/OS: User's Guide*, SC23-9703
Introduces the features, workspaces, attributes, and predefined situations for the IBM Tivoli OMEGAMON XE for Storage on z/OS product and supplements the user assistance provided with this product.
- *IBM Tivoli OMEGAMON XE for Storage on z/OS: Tuning Guide*, SC23-9704
Provides information on enhancing the performance of Tivoli OMEGAMON XE for Storage on z/OS so that you can get the maximum benefit with the least amount of resources.
- *IBM Tivoli OMEGAMON XE for Storage on z/OS: Troubleshooting Guide*, GC23-9705
Contains messages for the IBM Tivoli OMEGAMON XE for Storage on z/OS product and OMEGAMON II for SMS component and information to help solve problems with the IBM Tivoli OMEGAMON XE for Storage on z/OS product.
- *Tivoli OMEGAMON XE for Storage on z/OS user assistance*
Helps operators understand and use the provided data, attributes, commands, and situation to monitor performance and availability in the context of the product.

Shared OMEGAMON XE publications

The following publications provide information common to the OMEGAMON XE products:

- *Quick Start Guide*, GI11-8918
Summarizes the installation and setup of an OMEGAMON XE monitoring agent on z/OS.
- *Common Planning and Configuration Guide*, SC23-9734
Gives instructions for planning and configuration tasks common to the components of Tivoli Management Services on z/OS and the OMEGAMON XE monitoring agents on z/OS.
- *Upgrade Guide*, SC23-9745
Gives instructions for complete and staged upgrades to V4.2.0 of the OMEGAMON XE products.
- *End-to-End Response Time Feature Reference*, SC27-2303

Provides instructions and reference information for the End-to-End Response Time Feature, which supplies response time data to several OMEGAMON XE products.

- *Reports for Tivoli Common Reporting, SC27-2304*

Explains how to use the Tivoli Common Reporting tool to create reports from data displayed in the Tivoli Enterprise Portal and stored in the Tivoli Data Warehouse database.

IBM Tivoli Monitoring library

The following publications provide information about IBM Tivoli Monitoring V6.2 and about the commonly shared components of Tivoli Management Services:

- *Quick Start Guide, GI11-8058*

Introduces the components of IBM Tivoli Monitoring.

- *Exploring IBM Tivoli Monitoring, SC32-1803*

Provides a series of exercises to help you explore IBM Tivoli Monitoring. By completing the activities in this workbook, you will install and configure your monitoring environment, explore both the graphical and command-line interfaces of the product, use some of the new features, and work with several monitoring agents.

- *Installation and Setup Guide, GC32-9407*

Provides instructions for installing and configuring IBM Tivoli Monitoring components on Windows, Linux®, and UNIX systems.

- *Program Directory for IBM Tivoli Management Services on z/OS, GI11-4105*

Gives instructions for the SMP/E installation of the Tivoli Management Services components on z/OS.

- *IBM Tivoli Monitoring: Configuring the Tivoli Enterprise Monitoring Server on z/OS, SC27-2313*

Gives detailed instructions for using the Configuration Tool to configure Tivoli Enterprise Monitoring Server on z/OS systems. Includes scenarios for using batch mode to replicate monitoring environments across the z/OS enterprise. Also provides instructions for setting up security and for adding application support to a Tivoli Enterprise Monitoring Server on z/OS.

- *Administrator's Guide, SC32-9408*

Describes the support tasks and functions required for the Tivoli Enterprise Portal Server and clients, including Tivoli Enterprise Portal user administration.

- Tivoli Enterprise Portal online help

Provides context-sensitive reference information about all features and customization options of the Tivoli Enterprise Portal. Also gives instructions for using and administering the Tivoli Enterprise Portal.

- *User's Guide, SC32-9409*

Complements the Tivoli Enterprise Portal online help. The guide provides hands-on lessons and detailed instructions for all Tivoli Enterprise Portal features.

- *Command Reference, SC32-6045*

Provides detailed syntax and parameter information and examples of the commands you can use in IBM Tivoli Monitoring.

- *Troubleshooting Guide, GC32-9458*

Provides information to help you troubleshoot problems with the software.

- *Deployment Guide, SC23-8873*

Outlines key activities such as planning, maintenance, and tuning considerations that are needed to ensure a successful deployment of your IBM Tivoli Monitoring environment.

- *Messages, SC23-7969*

Lists and explains messages generated by all IBM Tivoli Monitoring components and by z/OS-based Tivoli Management Services components (such as Tivoli Enterprise Monitoring Server on z/OS and TMS:Engine).

- *Upgrading from V5.1.2, GC32-1976*
Gives instructions for upgrading custom resource models from IBM Tivoli Monitoring V5.1.2 to IBM Tivoli Monitoring V6.2.
- *IBM Tivoli Universal Agent User's Guide, SC32-9459*
Introduces you to the IBM Tivoli Universal Agent, an agent of IBM Tivoli Monitoring. The IBM Tivoli Universal Agent enables you to use the monitoring and automation capabilities of IBM Tivoli Monitoring to monitor any type of data you collect.
- *IBM Tivoli Universal Agent API and Command Programming Reference Guide, SC32-9461*
Explains the procedures for implementing the IBM Tivoli Universal Agent APIs and provides descriptions, syntax, and return status codes for the API calls and command-line interface commands.
- *Agent Builder User's Guide, SC32-1921*
Explains how to use the Agent Builder for creating monitoring agents and their installation packages, and for adding functions to existing agents.
- *Fix Pack Readme and Documentation Addendum*
Describes the installation process for a fix pack and provides updated information on problems and workarounds associated with the fix pack. A new version of this document is created for each fix pack.

Related publications

You can find useful information about the OMEGAMON XE monitoring agents in the *IBM Tivoli Monitoring and OMEGAMON XE Information Center* at the following Web address: <http://publib.boulder.ibm.com/infocenter/tivihelp/v15r1/>.

This set of documents include the following items, which are related to configuration of the OMEGAMON II for SMS component, which is optional in an Tivoli OMEGAMON XE for Storage on z/OS monitoring environment.

- *OMEGAMON II for SMS Configuration and Customization Guide*
- *OMEGAMON II for SMS User's Guide*
- *OMEGAMON II for SMS Administrator's Guide*
- *OMEGAMON II for SMS Tuning Guide*

Other sources of documentation

You can also obtain technical documentation about IBM Tivoli Monitoring and OMEGAMON XE products from the following sources:

- IBM Tivoli Open Process Automation Library (OPAL)
<http://www.ibm.com/software/tivoli/opal>
OPAL is an online catalog that contains integration documentation and other downloadable product extensions. This library is updated daily.
- IBM Tivoli Distributed Monitoring and Application Management wiki
<http://www.ibm.com/developerworks/wikis/display/tivolimonitoring/Home>
The IBM Tivoli Distributed Monitoring and Application Management wiki is a source of best practices, white papers, discussion forums, and more. The wiki includes integration documentation as well as downloadable product extensions.
- Redbooks
<http://www.redbooks.ibm.com/>
IBM Redbooks, Redpapers, and Redbooks Technotes provide information about products from platform and solution perspectives.
- Technotes

You can find Technotes through the IBM Software Support Web site at <http://www.ibm.com/software/support/probsub.html>, or more directly through your product Web site, which contains a link to Technotes (under **Solve a problem**).

Technotes provide the latest information about known product limitations and workarounds.

Appendix B. Support information

If you have a problem with your IBM software, you want to resolve it quickly. This section describes the following options for obtaining support for IBM software products:

- “Using IBM Support Assistant”
- “Obtaining fixes”
- “Receiving weekly support updates” on page 334
- “Contacting IBM Software Support” on page 335

Using IBM Support Assistant

The IBM Support Assistant is a free, stand-alone application that you can install on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products you use.

The IBM Support Assistant saves you the time it takes to search the product, support, and educational resources. The IBM Support Assistant helps you gather support information when you need to open a problem management record (PMR), which you can then use to track the problem.

The product-specific plug-in modules provide you with the following resources:

- Support links
- Education links
- Ability to submit problem management reports

For more information, and to download the IBM Support Assistant, see <http://www.ibm.com/software/support/isa>. After you download and install the IBM Support Assistant, follow these steps to install the plug-in for your Tivoli product:

1. Start the IBM Support Assistant application.
2. Select **Updater** on the Welcome page.
3. Select **New Properties and Tools** or select the **New Plug-ins** tab (depending on the version of IBM Support Assistant installed).
4. Under **Tivoli**, select your product, and then click **Install**. Be sure to read the license and description. If your product is not included on the list under **Tivoli**, no plug-in is available yet for the product.
5. Read the license and description, and click **I agree**.
6. Restart the IBM Support Assistant.

Obtaining fixes

A product fix might be available to resolve your problem. To determine which fixes are available for your Tivoli software product, follow these steps:

1. Go to the IBM Software Support Web site at <http://www.ibm.com/software/support>.
2. Select **Tivoli** in the **Select a brand and/or product** drop-down list to access the Select a product drop-down list.
3. Select **IBM Tivoli OMEGAMON XE for Storage on z/OS** in the **Select a product** drop-down list.
4. Click the **Go** arrow that is located to the right of the drop-down list. The product support Web page is displayed, including the blue **IBM Tivoli OMEGAMON XE for Storage on z/OS** support list box on the right.

Note: The links in this list box go to dedicated Web pages for the product regarding topics such as downloads and troubleshooting. In many cases, using these links leads you to product-specific information more quickly than entering search terms.

5. To obtain information on APARs, perform the following steps:
 - a. Click **Troubleshoot** in the **IBM Tivoli OMEGAMON XE for Storage on z/OS** support list box to access the **Troubleshoot** support page.
 - b. Click **APARs** in the **Browse by document type** area to access a list of APARs, with the most recent APARs first.
 - c. Enter search terms to find specific types of APARs, as needed.
6. To obtain information on fixes, fix packs, and other service updates for IBM Tivoli OMEGAMON XE for Storage on z/OS, perform the following steps:
 - a. Click **Download** in the **IBM Tivoli OMEGAMON XE for Storage on z/OS** support list box to access the **Download** support page.
 - b. Click the **Recommended fixes** link to access the **Recommended Maintenance Service Levels** page. Information for IBM Tivoli OMEGAMON XE for Storage on z/OS is provided in several places on this page.

For more information about the types of fixes that are available, see the *IBM Software Support Handbook* at <http://techsupport.services.ibm.com/guides/handbook.html>.

Receiving weekly support updates

To receive weekly e-mail notifications about fixes and other software support news, follow these steps:

1. Go to the IBM Software Support Web site at <http://www.ibm.com/software/support>.
2. Click **My support** in the far upper-right corner of the page under **Personalized support**.
3. If you have already registered for **My support**, sign in and skip to the next step. If you have not registered, click **register now**. Complete the registration form using your e-mail address as your IBM ID and click **Submit**.
4. The **Edit profile** tab is displayed.
5. In the first list under **Products**, select **Software**. In the second list, select a product category (for example, **Systems and Asset Management**). In the third list, select a product sub-category (for example, **Application Performance & Availability** or **Systems Performance**). A list of applicable products is displayed.
6. Select the products for which you want to receive updates.
7. Click **Add products**.
8. After selecting all products that are of interest to you, click **Subscribe to email** on the **Edit profile** tab.
9. In the **Documents** list, select **Software**.
10. Select **Please send these documents by weekly email**.
11. Update your e-mail address as needed.
12. Select the types of documents you want to receive.
13. Click **Update**.

If you experience problems with the **My support** feature, you can obtain help in one of the following ways:

Online

Send an e-mail message to erchelp@ca.ibm.com, describing your problem.

By phone

Call 1-800-IBM-4You (1-800-426-4968).

Contacting IBM Software Support

IBM Software Support provides assistance with product defects. The easiest way to obtain that assistance is to open a PMR or ETR directly from the IBM Support Assistant (see “Using IBM Support Assistant” on page 333).

Before contacting IBM Software Support, your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM. The type of software maintenance contract that you need depends on the type of product you have:

- For IBM distributed software products (including, but not limited to, Tivoli, Lotus®, and Rational® products, as well as DB2® and WebSphere® products that run on Windows or UNIX operating systems), enroll in Passport Advantage® in one of the following ways:

Online

Go to the Passport Advantage Web site at http://www-306.ibm.com/software/howtobuy/passportadvantage/pao_customers.htm.

By phone

For the phone number to call in your country, go to the IBM Software Support Web site at <http://techsupport.services.ibm.com/guides/contacts.html> and click the name of your geographic region.

- For customers with Subscription and Support (S & S) contracts, go to the Software Service Request Web site at <https://techsupport.services.ibm.com/ssr/login>.
- For customers with IBMLink™, CATIA, Linux, OS/390®, iSeries®, pSeries®, zSeries®, and other support agreements, go to the IBM Support Line Web site at <http://www.ibm.com/services/us/index.wss/so/its/a1000030/dt006>.
- For IBM eServer™ software products (including, but not limited to, DB2 and WebSphere products that run in zSeries, pSeries, and iSeries environments), you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for eServer software products, go to the IBM Technical Support Advantage Web site at <http://www.ibm.com/servers/eserver/techsupport.html>.

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. From other countries, go to the contacts page of the *IBM Software Support Handbook* on the Web at <http://techsupport.services.ibm.com/guides/contacts.html> and click the name of your geographic region for phone numbers of people who provide support for your location.

To contact IBM Software support, follow these steps:

1. “Determining the business impact”
2. “Describing problems and gathering information” on page 336
3. “Submitting problems” on page 336

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Use the following criteria to understand and assess the business impact of the problem that you are reporting:

Severity 1

The problem has a *critical* business impact. You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.

Severity 2

The problem has a *significant* business impact. The program is usable, but it is severely limited.

Severity 3

The problem has *some* business impact. The program is usable, but less significant features (not critical to operations) are unavailable.

Severity 4

The problem has *minimal* business impact. The problem causes little impact on operations, or a reasonable circumvention to the problem was implemented.

Describing problems and gathering information

When describing a problem to IBM, be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently. To save time, know the answers to these questions:

- Which software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can you re-create the problem? If so, what steps were performed to re-create the problem?
- Did you make any changes to the system? For example, did you make changes to the hardware, operating system, networking software, and so on.
- Are you currently using a workaround for the problem? If so, be prepared to explain the workaround when you report the problem.

Submitting problems

You can submit your problem to IBM Software Support in one of two ways:

Online

Click **Submit and track problems** on the IBM Software Support site at <http://www.ibm.com/software/support/probsub.html>. Type your information into the appropriate problem submission form.

By phone

For the phone number to call in your country, go to the contacts page of the *IBM Software Support Handbook* at <http://techsupport.services.ibm.com/guides/contacts.html> and click the name of your geographic region.

If the problem you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the Software Support Web site daily, so that other users who experience the same problem can benefit from the same resolution.

Notices

This information was developed for products and services offered in the U.S.A. IBM 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 be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right 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 have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give 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 (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply in the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "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 not apply to you.

This publication might include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM make improvements or changes in the products or the programs described in this publication at any time without notice.

Licensees of this program who want to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 U.S.A.

Such information be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Programming License Agreement, or any equivalent agreement between us.

Trademarks

IBM, the IBM logo, and [ibm.com](http://www.ibm.com) are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at <http://www.ibm.com/legal/copytrade.shtml>

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo 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, and service names may be trademarks or service marks of others.

Index

A

- about attributes 125
- about situations 111
- about this document xiii
- about Tivoli OMEGAMON XE for Storage on z/OS 1
- accessing workspaces 8
- action commands in Storage Toolkit 291
- actions 251, 258, 292
 - custom 265
 - data-set related 273
 - DFSMSHsm-related 286
 - DFSMSrmm-related 287
 - volume-related 270
- application dataset details 127
- application monitoring 129
- application summary workspace 13
- application volume performance 130
- application volume space 132
- associating reports with attributes 3
- attribute definitions for DFSMSrmm
 - Status feature 324
- attribute group
 - attribute item 1
- attribute groups
 - application dataset details 127
 - application monitoring 129
 - application volume performance 130
 - application volume space 132
 - cache control unit 133
 - cache devices 134
 - cache raid rank 138
 - channel path 138
 - cross system volume summary 140
 - cross system volumes 141
 - dasd physical group 142
 - dasd volume performance 145
 - dasd volume space 147
 - dataset attributes blocksize
 - summary 149
 - dataset attributes CA split
 - summary 149
 - dataset attributes catalog
 - summary 150
 - dataset attributes CI split
 - summary 151
 - dataset attributes creation date
 - summary 152
 - dataset attributes dataset detail 153
 - dataset attributes DSORG detail 157
 - dataset attributes DSORG
 - summary 158
 - dataset attributes extent
 - summary 159
 - dataset attributes inefficient blocksize
 - summary 160
 - dataset attributes mature dataset
 - summary 161
 - dataset attributes never referenced
 - dataset summary 162

- attribute groups (*continued*)
 - dataset attributes newborn dataset
 - summary 163
 - dataset attributes reference date
 - summary 163
 - dataset attributes SMS class
 - detail 164
 - dataset attributes SMS class
 - summary 165
 - dataset attributes space allocated
 - summary 166
 - dataset attributes space unused
 - summary 166
 - dataset attributes system
 - summary 167
 - dataset attributes uncataloged
 - detail 169
 - dataset attributes uncataloged
 - summary 170
 - dataset group detail attrs 174
 - dataset group details 170
 - dataset group details volume 175
 - dataset group summary 177
 - dataset performance 178
 - dataset SMS 180
 - dataset space 180
 - dataset users 182
 - HSM CDS 182
 - HSM common storage 183
 - HSM data set recover version 184
 - HSM function details 184
 - HSM function statistics 185
 - HSM function summary 187
 - HSM host function status 187
 - HSM host request summary 189
 - HSM host status 189
 - HSM private storage 191
 - HSM requests 191
 - HSM status 192
 - logical control unit 194
 - logical control unit channel
 - paths 195
 - RMM configuration 197
 - RMM control 200
 - RMM control data set 200
 - RMM data set details 201
 - RMM location details 203
 - RMM owner details 204
 - RMM pending actions 204
 - RMM system summary 205
 - RMM volume details 208
 - RMM VRS details 214
 - SMS cache sets 215
 - SMS configuration 215
 - SMS data class 216
 - SMS management class 217
 - SMS storage class 219
 - SMS storage group 220
 - SMS storage group status 222
 - SMS systems 222
 - storage toolkit request s 222

- attribute groups (*continued*)
 - storage toolkit result detail 224
 - storage toolkit result resources 225
 - storage toolkit result summary 224
 - Symmetrix configuration 226
 - Symmetrix devices 226
 - Symmetrix Director summary 227
 - Symmetrix Directors 227
 - Symmetrix SSIDs 229
 - tape device 229
 - tape group 229
 - TotalStorageDS array
 - configuration 231
 - TotalStorageDS configuration 232
 - TotalStorageDS extent pool 232
 - TotalStorageDS rank 234
 - TotalStorageDS SSIDs 234
 - volume group summary 235
 - volume status 238
 - volume users 238
 - VTs cache 238
 - VTs capacity 239
 - VTs composite library 241
 - VTs overview 242
 - VTs physical devices 243
 - VTs TS7700-series Cache 244
 - VTs TS7700-series Cache
 - Partition 245
 - VTs TS7700-series Cache Preference
 - Group 246
 - VTs TS7700-series Cluster 247
 - VTs virtual devices 248
 - VTs virtual devices summary 250
- attribute groups listed 125
- attributes
 - composition of 1
 - organization of 1
 - structure of 1
 - understanding 1
 - using 1
- authentication of Storage Toolkit
 - commands 261
- authorization of Storage Toolkit
 - commands 261

B

- benefits of DFSMSrmm Status feature 323
- building situations 1

C

- cache control unit 133
- cache cu performance workspace 23
- cache cu status workspace 35
- cache devices 134
- cache raid rank 138
- Cancel HSM Request dialog box 286
- cancelling DFSMSHsm requests 287

- cancelling DFSMSShm requests, warning 287
- channel path 138
- channel path workspace 18
- command line option
 - IBM TotalStorage Productivity Center 318
- Command tab, Storage Toolkit 260
- command-line option
 - Dataset Attributes Database 315
- command-line option for DFSMSrmm Status feature 325
- commands
 - running from Tivoli Enterprise Portal 251, 317
 - Storage Toolkit 251
- commands, Take Action 4
- composition of an attribute 1
- cross system volume summary 140
- cross system volumes 141
- cross-product workspace links 9
- custom queries
 - Dataset Attributes Database 315
- custom queries for DFSMSrmm Status feature 326
- customer support 335

D

- dasd physical group 142
- dasd volume performance 145
- dasd volume space 147
- data set group space utilization graph 94
- data set group summary report
 - adding a data set group 298
 - editing and deleting data set groups 298
 - setting the collection interval for a data set group 298
- data set group summary workspace 94
- Dataset Attributes Database, benefits 311
- dataset attributes blocksize summary 149
- dataset attributes CA split summary 149
- dataset attributes catalog summary 150
- dataset attributes CI split summary 151
- dataset attributes creation date summary 152
- Dataset Attributes Database 311
 - benefits 311
 - calculation of values 311
 - command-line option 315
 - configuration 313
 - custom queries 315
 - optimization 311
 - optimizing 314
 - scenarios 313
- Dataset Attributes Database, optimization 311
- dataset attributes dataset detail 153
- dataset attributes DSORG detail 157
- dataset attributes DSORG summary 158
- dataset attributes extent summary 159
- Dataset Attributes feature
 - compared to Dataset Group feature 297
- Dataset Attributes feature compared to Dataset Group feature 297
- dataset attributes inefficient blocksize summary 160
- dataset attributes mature dataset summary 161
- dataset attributes never referenced dataset summary 162
- dataset attributes newborn dataset summary 163
- dataset attributes reference date summary 163
- dataset attributes SMS class detail 164
- dataset attributes SMS class summary 165
- dataset attributes space allocated summary 166
- dataset attributes space unused summary 166
- dataset attributes system summary 167
- dataset attributes system summary workspace 88
- dataset attributes uncataloged detail 169
- dataset attributes uncataloged summary 170
- Dataset Group Data Collection dialog box 298
- dataset group detail attrs 174
- dataset group details 170
- dataset group details volume 175
- Dataset Group feature 297
 - compared to Dataset Attributes feature 297
 - find data sets with excess space 301
 - finding cataloged, not-allocated data sets 306
 - finding excessive CA/CI splits, VSAM data sets 304
 - masking rules 299
 - migrate infrequently used data sets 308
 - tracking number of extents 308
 - tracking response times 302
 - tracking space usage 300
 - using 300
 - working with 297
- Dataset Group feature compared to Dataset Attributes feature 297
- Dataset Group Properties dialog box 297
- dataset group summary 177
- dataset performance 178
- dataset SMS 180
- dataset space 180
- dataset users 182
- defining an attribute 1
- deleting 292
- description of Tivoli OMEGAMON XE for Storage on z/OS 1
- DFDSS facility
 - accessing from Tivoli Enterprise Portal 251
- DFSMSShsm status workspace 73
- DFSMSShsm Status workspace 73
- DFSMSrmm Status feature 323
 - Add Dataset dialog box 288
 - Add Owner dialog box 289
 - Add Volume dialog box 289

- DFSMSrmm Status feature (*continued*)
 - Add VRS dialog box 290
 - attribute definitions 324
 - benefits 323
 - Change Dataset dialog box 288
 - Change Owner dialog box 289
 - Change Volume dialog box 290
 - Change VRS dialog box 291
 - command-line option 325
 - Confirm Actions for All Volumes dialog box 290
 - custom queries 326
 - Delete Dataset dialog box 288
 - Delete Owner dialog box 289
 - Delete Volume dialog box 290
 - Delete VRS dialog box 291
 - optimization 323
 - optimizing data collection 325
 - scenarios 324
 - Summary Collection Control dialog box 326
 - use cases 324
- DFSMSrmm status workspace 79
- DFSMSrmm System Summary workspace 79
- dialog boxes
 - Dataset Group Data Collection 298
 - Dataset Group Properties 297
 - Storage Toolkit 251
- direct situation analysis 103
 - for application summary workspace 103
 - for cache CU performance workspace 104
 - for channel path workspace 103
 - for HSM status workspace 110
 - for logical control unit workspace 107
 - for SMS storage groups performance workspace 108
 - for SMS storage groups space workspace 109
- documentation 329
 - Tivoli OMEGAMON XE for Storage on z/OS 329
- dynamic linking, about 11
- dynamic workspace links 9

E

- enabled workspaces
 - IBM TotalStorage Productivity Center 318
- evaluating
 - tuning decision 3
- event
 - investigating 3
- event indicator 3
- events analysis 103

F

- favorites, adding a workspace to 4
- filtering 3
- finding cataloged, not-allocated data sets 306

finding data sets with excess space 301
finding excessive CA/CI splits, VSAM
data sets 304
fixes, obtaining 333
formats for information 2

G

General tab, Storage Toolkit 258
grouping data sets 297
GUI launch
IBM TotalStorage Productivity
Center 317

H

hardware settings
accessing the IBM TotalStorage
Productivity Center 317
HSM CDS 182
HSM common storage 183
HSM data set recover version 184
HSM facility
accessing from Tivoli Enterprise
Portal 251
HSM function details 184
HSM function statistics 185
HSM function summary 187
HSM host function status 187
HSM host request summary 189
HSM host status 189
HSM private storage 191
HSM requests 191
HSM requests, cancelling 286
HSM status 192

I

IBM Redbooks 333
IBM Support Assistant 333
IBM Tivoli Distributed Monitoring and
Application Management wiki 331
IBM Tivoli Open Process Automation
Library (OPAL) 331
IBM TotalStorage Productivity
Center 317
command line option 318
enabled workspaces 318
GUI launch 317
scenario 317
identifying
system bottlenecks 3
integration with other products 9
investigating an event 3

J

JCL tab, Storage Toolkit 259

K

KS3_Applic_Resp_Time_Critical 112
KS3_Applic_Resp_Time_Warning
situation 112

KS3_Cachecu_Cache_Stat_Critical
situation 112
KS3_Cachecu_DFW_Retry_Critical
situation 112
KS3_Cachecu_DFW_Retry_Warning
situation 112
KS3_Cachecu_Inact_Vols_Critical
situation 112
KS3_Cachecu_Inact_Vols_Warning
situation 113
KS3_Cachecu_NV5_Stat_Critical
situation 113
KS3_Cachecu_Read_HitP_Critical
situation 113
KS3_Cachecu_Read_HitP_Warning
situation 113
KS3_Cachecu_Trk_Dstg_Critical
situation 113
KS3_Cachecu_Trk_Dstg_Warning
situation 113
KS3_Cachecu_Trk_Stag_Critical
situation 113
KS3_Cachecu_Trk_Stag_Warning
situation 114
KS3_Cachecu_Write_HitP_Critical
situation 114
KS3_Cachecu_Write_HitP_Warning
situation 114
KS3_Channel_Busy_Pct_Critical
situation 114
KS3_Channel_Busy_Pct_Warning
situation 114
KS3_HSM_Backup_Held_Critical
situation 114
KS3_HSM_Backup_Queue_Critical
situation 114
KS3_HSM_Backup_Queue_Warning
situation 115
KS3_HSM_Dump_Held_Critical
situation 115
KS3_HSM_Dump_Queue_Critical
situation 115
KS3_HSM_Dump_Queue_Warning
situation 115
KS3_HSM_Inactive_Host_Warning 115
KS3_HSM_Migrate_Held_Critical
situation 115
KS3_HSM_Migrate_Queue_Critical
situation 116
KS3_HSM_Migrate_Queue_Warning
situation 116
KS3_HSM_Recall_Held_Critical
situation 116
KS3_HSM_Recall_Queue_Critical
situation 116
KS3_HSM_Recall_Queue_Warning
situation 116
KS3_HSM_Recovery_Held_Critical
situation 116
KS3_HSM_Recovery_Queue_Critical
situation 117
KS3_HSM_Recovery_Queue_Warning
situation 117
KS3_HSM_Status_Inactive_Crit
situation 117
KS3_LCU_Av_Delay_Q_Critical
situation 117

KS3_LCU_Av_Delay_Q_Warning
situation 117
KS3_LCU_Cont_Rate_Critical
situation 117
KS3_LCU_Cont_Rate_Warning
situation 118
KS3_LCU_IO_Rate_Sec_Critical
situation 118
KS3_LCU_IO_Rate_Sec_Warning
situation 118
KS3_RMM_CDS_Backup_Critical
Situation 118
KS3_RMM_CDS_Backup_Warning
Situation 118
KS3_RMM_CDS_Space_Critical
Situation 118
KS3_RMM_CDS_Space_Warning
Situation 118
KS3_RMM_Exit_Status_Critical
Situation 118
KS3_RMM_Inactive_Critical
Situation 119
KS3_RMM_Journal_Status_Critical
Situation 119
KS3_RMM_Operating_Mode_Warning
Situation 119
KS3_RMM_Scratch_Tape_Critical
Situation 119
KS3_RMM_Scratch_Tape_Warning
Situation 119
KS3_Stg_Toolkit_Result_Critical
situation 119
KS3_Stg_Toolkit_Result_Warning
situation 119
KS3_Storage_Gr_Pct_Free_Crit
situation 119
KS3_Storage_Gr_Pct_Free_Warning
situation 120
KS3_Vol_Cache_DFW_Retry_Critical
situation 120
KS3_Vol_Cache_DFW_Retry_Warning
situation 120
KS3_Vol_Cache_Read_HitP_Critical
situation 120
KS3_Vol_Cache_Read_HitP_Warning
situation 120
KS3_Vol_Cache_Writ_HitP_Critical
situation 121
KS3_Vol_Cache_Writ_HitP_Warning
situation 121
KS3_Vol_Disabled_VTOC_Critical
situation 121
KS3_Vol_EAV_Fragment_Index_Crit 121
KS3_Vol_EAV_Fragment_Index_Warn 121
KS3_Vol_EAV_Free_Space_Pct_Crit 121
KS3_Vol_EAV_Free_Space_Pct_Warn 121
KS3_Vol_Fragment_Index_Critical
situation 122
KS3_Vol_Fragment_Index_Warning
situation 122
KS3_Vol_Free_Space_Pct_Critical
situation 122
KS3_Vol_Free_Space_Pct_Warning
situation 122
KS3_Vol_Perf_Resp_Time_Critical
situation 122

- KS3_Vol_Perf_Resp_Time_Warning situation 122
- KS3_VTS_Disconnect_Time_Crit situation 122
- KS3_VTS_Host_GB_Warning situation 123
- KS3_VTS_Pct_Copy_Throt_Warn situation 123
- KS3_VTS_Pct_Wr_Over_Throt_Warn situation 123
- KS3_VTS_Recall_Pct_Warning situation 123
- KS3_VTS_Virt_MtPend_Av_Warning situation 123
- KS3_VTS_Virt_MtPend_Mx_Warning situation 123

L

- libraries
 - IBM Tivoli Monitoring 330
 - OMEGAMON XE shared publications 329
- library, Tivoli OMEGAMON XE for Storage on z/OS 329
- limitations and hints
 - Storage Toolkit 292
- links to workspaces
 - in Application Summary node 13
 - in Cache CU Performance node 21
 - in cache cu performance workspace 24
 - in Cache CU Status node 33
 - in cache cu status workspace 35
 - in Channel Path node 17
 - in channel path workspace 18
 - in Dataset Attributes System Summary node 87
 - in dataset attributes system summary workspace 89
 - in Dataset Group Summary node 93
 - in dataset group summary workspace 94
 - in DFSMSHsm Status node 73
 - in DFSMSHsm Status workspace 73
 - in DFSMSrmm Status node 77
 - in Logical Control Unit node 37
 - in logical control unit workspace 40
 - in SMS Configuration node 97
 - in SMS configuration workspace 98
 - in SMS Storage Groups Performance node 51
 - in SMS storage groups performance workspace 54
 - in SMS Storage Groups Space node 59
 - in SMS storage groups space workspace 60
 - in Storage Toolkit node 101
 - in Tape Group node 43
 - in tape group workspace 44
 - in User DASD Groups Performance node 65
 - in user DASD groups performance workspace 67
 - in User DASD Groups Space node 69

- links to workspaces (*continued*)
 - in user DASD groups space workspace 70
 - in Virtual Tape Subsystems node 45
 - in virtual tape subsystems workspace 13, 17, 21, 33, 37, 43, 45, 46, 51, 59, 65, 69, 87, 93, 97, 101
- list of attribute groups 125
- logical control unit 194
- logical control unit channel paths 195
- logical control unit workspace 39

M

- mainframe commands
 - running from Tivoli Enterprise Portal 251
- manage hardware settings
 - accessing the IBM TotalStorage Productivity Center 317
- managed system 1, 3
- managing 258, 292
- managing actions 258, 292
- manually deleting 292
- masking rules for data set grouping 299
- masks for data set grouping 299
- menus
 - Storage Toolkit 251
- metrics available for monitoring
 - application dataset details group 127
 - application monitoring 129
 - application volume performance 130
 - application volume space 132
 - cache control unit 133
 - cache devices 134
 - cache raid rank 138
 - channel path 138
 - cross system volume summary 140
 - cross system volumes 141
 - dasd physical group 142
 - dasd volume performance 145
 - dasd volume space 147
 - dataset attributes blocksize summary 149
 - dataset attributes CA split summary 149
 - dataset attributes catalog summary 150
 - dataset attributes CI split summary 151
 - dataset attributes creation date summary 152
 - dataset attributes dataset detail 153
 - dataset attributes DSORG detail 157
 - dataset attributes DSORG summary 158
 - dataset attributes extent summary 159
 - dataset attributes inefficient blocksize summary 160
 - dataset attributes mature dataset summary 161
 - dataset attributes never referenced dataset summary 162
 - dataset attributes newborn dataset summary 163
- metrics available for monitoring (*continued*)
 - dataset attributes reference date summary 163
 - dataset attributes SMS class detail 164
 - dataset attributes SMS class summary 165
 - dataset attributes space allocated summary 166
 - dataset attributes space unused summary 166
 - dataset attributes system summary 167
 - dataset attributes uncataloged detail 169
 - dataset attributes uncataloged summary 170
 - dataset group detail attrs 174
 - dataset group details 170
 - dataset group details volume 175
 - dataset group summary 177
 - dataset performance 178
 - dataset SMS 180
 - dataset space 180
 - dataset users 182
 - HSM CDS 182
 - HSM common storage 183
 - HSM data set recover version 184
 - HSM function details 184
 - HSM function statistics 185
 - HSM function summary 187
 - HSM host function status 187
 - HSM host request summary 189
 - HSM host status 189
 - HSM private storage 191
 - HSM requests 191
 - HSM status 192
 - logical control unit 194
 - logical control unit channel paths 195
 - RMM configuration 197
 - RMM control 200
 - RMM control data set 200
 - RMM data set details 201
 - RMM location details 203
 - RMM owner details 204
 - RMM pending actions 204
 - RMM system summary 205
 - RMM volume details 208
 - RMM VRS details 214
 - SMS cache sets 215
 - SMS configuration 215
 - SMS data class 216
 - SMS management class 217
 - SMS storage class 219
 - SMS storage group 220
 - SMS storage group status 222
 - SMS systems 222
 - storage toolkit requests 222
 - storage toolkit result detail 224
 - storage toolkit result resources 225
 - storage toolkit result summary 224
 - Symmetrix configuration 226
 - Symmetrix devices 226
 - Symmetrix Director summary 227
 - Symmetrix Directors 227

- metrics available for monitoring (continued)
 - Symmetrix SSIDs 229
 - tape device 229
 - tape group 230
 - TotalStorageDS array configuration 231
 - TotalStorageDS configuration 232
 - TotalStorageDS extent pool 232
 - TotalStorageDS rank 234
 - TotalStorageDS SSIDs 234
 - volume group summary 235
 - volume status 238
 - volume users 238
 - VTS cache 238
 - VTS capacity 239
 - VTS composite library 241
 - VTS overview 242
 - VTS physical devices 243
 - VTS TS7700-series Cache 244
 - VTS TS7700-series Cache Partition 245
 - VTS TS7700-series Cache Preference Group 246
 - VTS TS7700-series Cluster 247
 - VTS virtual devices 248
 - VTS virtual devices summary 250
- migrating infrequently used data sets 308
- monitoring
 - Applications 13
 - Cache CU Performance 21
 - Cache CU Status 33
 - Channel Path 17
 - Dataset Attributes System Summary 87
 - Dataset Group Summary 93
 - DFSMSHsm status 73
 - DFSMSRmm status 77
 - DFSMSRmm Status 77
 - Logical Control Unit 37
 - performance 1, 3
 - SMS Configuration 97
 - SMS Storage Groups Performance 51
 - SMS Storage Groups Space 59
 - Storage Toolkit Action Requests 101
 - Tape Group 43
 - User DASD Groups Performance 65
 - User DASD Groups Space 69
 - virtual tape subsystems 45

N

- new features 4
- nodes in the Navigator
 - Application Summary 13
 - Cache CU Performance 21
 - Cache CU Status 33
 - Channel Path 17
 - Dataset Attributes System Summary 87
 - Dataset Group Summary 93
 - DFSMSHsm status 73
 - DFSMSHsm Status 73
 - DFSMSRmm status 77
 - DFSMSRmm Status 77
 - DFSMSRmm Status node 77

- nodes in the Navigator (continued)
 - Logical Control Unit 37
 - SMS Configuration 97
 - SMS Storage Groups Performance 51
 - SMS Storage Groups Space 59
 - Storage Toolkit 101
 - Tape Group 43
 - User DASD Groups Performance 65
 - User DASD Groups Space 69
 - virtual tape subsystems 45
- notices from IBM 337

O

- OMEGAMON XE for Storage on z/OS
 - reports 322
- on-demand reports 322
- optimization of DFSMSRmm Status feature 323
- optimizing
 - Dataset Attributes Database 314
- Optimizing data collection for DFSMSRmm Status feature 325
- options
 - Storage Toolkit 251
- Options tab, Storage Toolkit 258
- overview
 - attributes 1
 - features 1
 - new features 4
 - Take Actions 4
 - workspaces 2
- overview of options
 - Storage Toolkit 251

P

- predefined situations 111
- problem determination
 - describing problems 336
 - determining business impact 335
 - submitting problems 336
- problem resolution 333
- process,
 - Storage Toolkit 254
- properties
 - defining for a workspace 3
- publications
 - IBM Tivoli Distributed Monitoring and Application Management wiki 331
 - IBM Tivoli Monitoring 330
 - IBM Tivoli Open Process Automation Library (OPAL) 331
 - OMEGAMON XE 329
 - Redbooks 331
 - related 331
 - shared 329
 - Technotes 331
 - Tivoli Management Services on z/OS 329, 330
 - Tivoli OMEGAMON XE for Storage on z/OS 329
 - types 329

R

- Redbooks 331, 333
- relationship between attributes and reports 3
- remote agents 1
- report description tables 322
- report types
 - on-demand 322
 - snapshot 322
 - summarized 322
- reports 2
 - application summary 13
 - cache cu performance 23
 - cache cu status 35
 - channel path 18
 - dataset attributes system summary 88
 - dataset group summary 93
 - DFSMSHsm status 73
 - DFSMSHsm Status 73
 - DFSMSRmm status 79
 - DFSMSRmm System Summary 79
 - logical control unit 39
 - OMEGAMON XE for Storage on z/OS 322
 - SMS configuration 97
 - SMS storage groups performance 53
 - SMS storage groups space 60
 - storage toolkit action requests 101
 - tape group 43
 - user DASD groups performance 67
 - user DASD groups space 69
 - virtual tape subsystems 45
- RMM configuration 197
- RMM control 200
- RMM control data set 200
- RMM data set details 201
- RMM location details 203
- RMM owner details 204
- RMM pending actions 204
- RMM system summary 205
- RMM volume details 208
- RMM VRS details 214

S

- scenario
 - IBM TotalStorage Productivity Center 317
- scenarios
 - Dataset Attributes Database 313
 - Storage Toolkit 254
- scenarios for DFSMSRmm Status feature 324
- scenarios,
 - Storage Toolkit 254
 - Storage Toolkit limitations and hints 292
- Scheduling tab, Storage Toolkit 259
- shared publications 329
- Show Data button 258
- situation analysis 103
- situation event
 - investigating 3
- situation event workspace 3
- situation event workspace, opening 3

- situations
 - see the KS3... prefix 111
 - using 111
- SMS base configuration table 97
- SMS cache sets 215
- SMS configuration 215
- SMS configuration workspace 97
- SMS data class 216
- SMS management class 217
- SMS storage class 219
- SMS storage group 220
- SMS storage group status 222
- SMS storage groups performance 53
- SMS storage groups space workspace 60
- SMS systems 222
- snapshot reports 322
- Software Support
 - contacting 335
 - describing problems 336
 - determining business impact 335
 - overview 333
 - receiving weekly updates 334
 - submitting problems 336
- sorting 4
- space usage for data set groups 300
- state change 3
- status of DFSMSrmm 323
- Storage Toolkit 251
 - action commands 291
 - actions
 - custom 265
 - data-set related 273
 - DFSMSHsm-related 286
 - DFSMSrmm-related 287
 - volume-related 270
 - Allocate dialog box 275
 - Alter dialog box 276
 - authentication and authorization 261
 - Cancel HSM Request dialog box 286
 - Catalog dialog box 280
 - Compress (Dataset) dialog box 275
 - Compress (Volume) dialog box 272
 - Create Batch Job dialog box 265
 - Data View dialog box 261
 - Dataset Backup dialog box 273
 - Dataset Copy and Move dialog box 274
 - Dataset Migration dialog box 273
 - Dataset Recall dialog box 274
 - Dataset Recover dialog box 274
 - Dataset Space Release dialog box 274
 - Define Cluster dialog box 277
 - Defrag dialog box 272
 - Delete dialog box 281
 - dialog boxes 251
 - Edit JCL dialog box 268
 - HSM requests, cancelling 286
 - Issue Command dialog box 269
 - limitations and hints 292
 - Listcat dialog box 281
 - menus and options 251
 - options 251
 - Print dialog box 283
 - process 254
 - Rename dialog box 284
 - Repro dialog box 284
 - scenarios 254

- Storage Toolkit (*continued*)
 - Submit Command or Job dialog box 269
 - tabs 258
 - Uncatalog dialog box 286
 - Verify dialog box 286
 - Volume Backup dialog box 271
 - Volume Conversion dialog box 271
 - Volume Migration dialog box 271
 - Volume Space Release dialog box 271
 - storage toolkit action requests workspace 101
 - Storage Toolkit dialog boxes 251
 - storage toolkit requests 222
 - storage toolkit result detail 224
 - storage toolkit result resources 225
 - storage toolkit result summary 224
 - structure of an attribute 1
 - summarized reports 322
 - support assistant 333
 - Symmetrix configuration 226
 - Symmetrix devices 226
 - Symmetrix Director s 227
 - Symmetrix Director summary 227
 - Symmetrix SSIDs 229

T

- tabs
 - Storage Toolkit
 - Command 260
 - General 258
 - JCL 259
 - Options 258
 - Scheduling 259
 - Take Action 4
 - take action commands
 - running a saved command 4
 - tape device 229
 - tape group 230
 - tape group workspace 43
 - Technotes 331
 - threshold values, setting 3
 - Tivoli Common Reporting
 - types of reports 322
 - Tivoli OMEGAMON XE for Storage on z/OS
 - description 1
 - Tivoli OMEGAMON XE for Storage workspaces 7
 - TotalStorage Productivity Center 317
 - TotalStorageDS array configuration 231
 - TotalStorageDS configuration 232
 - TotalStorageDS extent pool 232
 - TotalStorageDS rank 234
 - TotalStorageDS SSIDs 234
 - tracking number of extents in a data set group 308
 - tracking response times in a data set 302
 - Trademarks 338

U

- use cases
 - Dataset Attributes Database 313

- use cases for DFSMSrmm Status feature 324
- user DASD groups performance workspace 67
- user DASD groups space workspace 69
- using attributes 1
- using dataset grouping 300
- using reports 2
- using workspaces 2

V

- viewing
 - attribute information 3
- virtual tape subsystems workspace 45
- volume group summary 235
- volume status 238
- volume users 238
- VTS cache 238
- VTS capacity 239
- VTS composite library 241
- VTS overview 242
- VTS physical devices 243
- VTS TS7700-series Cache 244
- VTS TS7700-series Cache Partition 245
- VTS TS7700-series Cache Preference Group 246
- VTS TS7700-series Cluster 247
- VTS virtual devices 248
- VTS virtual devices summary 250

W

- working with data set groups 297
- workspace
 - adding to favorites 4
 - Event 3
- workspace links
 - in Application Summary node 13
 - in Cache CU Performance node 21
 - in cache cu performance workspace 24
 - in Cache CU Status node 33
 - in cache cu status workspace 35
 - in Channel Path node 17
 - in channel path workspace 18
 - in Dataset Attributes System Summary node 87
 - in dataset attributes system summary workspace 89
 - in Dataset Group Summary node 93
 - in dataset group summary workspace 94
 - in DFSMSHsm Status node 73
 - in DFSMSHsm Status workspace 73
 - in Logical Control Unit node 37
 - in logical control unit workspace 40
 - in SMS Configuration node 97
 - in SMS configuration workspace 98
 - in SMS Storage Groups Performance node 51
 - in SMS storage groups performance workspace 54
 - in SMS Storage Groups Space node 59

- workspace links (*continued*)
 - in SMS storage groups space
 - workspace 60
 - in Storage Toolkit node 101
 - in Tape Group node 43
 - in tape group workspace 44
 - in User DASD Groups Performance
 - node 65
 - in user DASD groups performance
 - workspace 67
 - in User DASD Groups Space
 - node 69
 - in user DASD groups space
 - workspace 70
 - in Virtual Tape Subsystems node 45
 - in virtual tape subsystems
 - workspace 13, 17, 21, 33, 37, 43, 45, 46, 51, 59, 65, 69, 87, 93, 97, 101
- workspace, event
 - opening 3
- workspaces 2
 - accessing 8
 - application summary 13
 - cache cu performance 23
 - cache cu status 35
 - channel path 18
 - data set group primary data sets 94
 - dataset attributes system
 - summary 88
 - dataset group summary 93
 - defining properties 3
 - DFSMSHsm status 73
 - DFSMSHsm Status 73
 - DFSMSrmm status 79
 - DFSMSrmm System Summary 79
 - direct situation analysis 103
 - formats for information 2
 - logical control unit 39
 - overview 7
 - SMS base configuration 97
 - SMS configuration 97
 - SMS storage groups performance 53
 - SMS storage groups space 60
 - storage toolkit action requests 101
 - tape group 43
 - user DASD groups performance 67
 - user DASD groups space 69
 - virtual tape subsystems 45
- workspaces overview 2
- workspaces with attributes 3

Z

- z/OS commands
 - running from Tivoli Enterprise
 - Portal 251



Printed in USA

SC23-9703-01

