


Integral Single Document Storage Procedures and Guidelines



Integral Single Document Storage Procedures and Guidelines

Note

Before using this information and the product it supports, read the information in "Notices" on page 157.

This edition applies to version 4.1.1 of IBM FileNet Image Services (product number 5724-R95) and to all subsequent releases and modifications until otherwise indicated in new editions.

© **Copyright International Business Machines Corporation** 2002, 2008. **All rights reserved.**

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

Contents

About This Manual 12

Audience 12

Document revision history 12

Accessing IBM FileNet Documentation 13

Conventions Used in this Manual 13

Typing Instructions 13

Screen Displays 14

Console Displays 14

Observe Cautions, Important Notes, Notes, and Tips 14

Command Syntax 15

 Optional Parameters 15

 Required Parameters 15

Comments and Suggestions 16

1 Single Document Storage (SDS) Overview 17

Introduction 17

Availability 21

Compatibility 21

Prerequisites 21

Image Services software 21

Server memory 22

Special consideration for Windows Servers 22

- Major SDS features 22**
 - SDS repository information 22
 - Repository connector libraries 23
 - Standard IS tools manage the configuration information 23
 - Database Maintenance Document Family 23
 - SDS-only document families 24
 - Support for multiple SDS units 24
 - Backward compatibility with NLS-SDS 25
 - Object retention 26
 - Retention types 27
 - The SDS Update Retention Tool 28
 - Indefinite hold 29
 - Content deletion 30
 - deldocs 30
 - Document services 30
 - Client applications 31
 - Checksumming 31
 - SDS Enable and Disable capability 32
 - Support for all CFS-IS document types 33
 - Integration into the IS committal/migration process 33
 - Committal/migration auto-restart 33
 - Retrieval redirection for individual documents 33
 - Page cache usage 34
 - Error handling 34

2 How to set up your SDS system 35

- Performance Recommendations and Best Practices 36**
- Multiple storage library servers with SDS 36**

Preferred library settings	37
How to configure Single Document Storage devices	38
Adjust Performance Tuning Parameters	38
Rebuild the IS configuration files and update the MKF databases	40
Set up the repository device	41
Determine the repository location string	41
Define the SDS configuration in Image Services	42
View the SDS details in the System Configuration Editor	48
Delete the SDS configuration from Image Services	49
Define SDS Families and Document Classes	51
How to define SDS Families	51
Supported family configurations	52
Configure the document class	52
Multi-Committal systems with SDS	52
Compatible systems	52
Incompatible systems	53

3 How to manage your SDS system 55

SDS system monitoring	56
Monitor RSVP and INFO messages in Storage Library Control (SLC)	56
RSVP scripting	57
IS Syslog and Subsystem Logging	57
IS Performance Monitoring (perf_mon)	58
Simple Network Management Protocol	59
SDS System States	59
About SDS background job tools	60

Convert and move data to integral SDS 61

Convert an NLS-SDS unit to integral SDS 61

Syntax 62

Examples 63

Convert an MSAR/optical/cache system to integral SDS 64

Syntax 65

Examples 66

Update the retention on specific documents 67

Syntax 69

Place SDS documents on indefinite hold 71

Syntax 72

Examples 73

Copy between two SDS units 73

Syntax 75

Examples 76

Restore or import data to an SDS unit 77

Syntax 77

Examples 79

Using sds_import on Centera 79

The DOC_tool utility 81

SDS States 81

DOC_Tool commands for managing SDS devices 82

ENable Command 83

DISable Command 84

DEMAND Queue Command 84

SDS Command 85

SDS command prompt and options 86

SDS List option (l) 86

SDS Summary option (s) 87
 SDS Detailed info option (d) 87
 SDS Worker option (w) 97
 SDS All option (a) 100
 SDS Find object option (f) 102
 SDS Delete option (DELETE) 104

Appendix A – Log Files and Messages 106

Error Logs 106

SDS Error Message Tuples 107

The fn_msg Tool 107

SDS (238) Messages 108

SDS general messages 108

SDSw Repository messages 111

Background SDS migration job messages 113

SDS Fastbatch committal messages 114

Appendix B – EMC Centera 115

Support policy 115

Performance tuning suggestions for Centera 116

Configuration steps 117

Set up Centera Profile and Pool Entry Authorization (.pea) file 117

 Create the home pool 117

 Create the pool entry authorization profile 118

 Verify the .pea profile file 120

 Copy the .pea file to the IS server 122

 Create new links for HP-UX 9000 only 122

- Modify the .profile and .cshrc files 123
- Add the Centera device to the IS configuration 123

Appendix C – IBM DR550 124

- What is IBM DR550? 124**
- What is IBM Tivoli Storage Manager? 124**
- TSM Configuration Steps 124**
 - Set up the IS software for IBM DR550 125
 - Restart the IS Software 129
 - Verify the TSM API Environment Variables 129
 - First Time Installation of TSM 129
 - Pre-existing installation of TSM 130
 - Set the MaxSessions Parameter 131
 - Add the DR550 device to the IS configuration 133
 - Restart the IS Software 133
 - Update the SDS family 133
- Set API Environment Variables for Tracing (Optional) 133**
 - Set TSM Tracing Method 134
 - Tracing on UNIX Systems 135
 - Tracing On Windows Systems 135
 - Restart the IS Software 136

Appendix D – Windows network security settings 137

- How to create a Domain Account for Image Services 137**

How to set security properties for remote network storage devices 138

How to set security properties for remote windows server storage devices 138

How to stop IS software and IS ControlService 140

How to modify the IS ControlService to use the domain account 140

How to add the domain account to the following groups: Administrators, FNADMIN, FNOP, FNUSR 141

How to modify the local security policy for Domain Account 142

How to configure a storage library to use the MSAR directory 143

Start the IS Software 143

Appendix E – Centera License Agreement 144

Glossary 146

Notices 157

Trademarks 161

U.S. Patents Disclosure 161

About This Manual

This document provides procedures and guidelines for Single Document Storage (SDS) on IBM® FileNet® Image Services (IS) systems for both Windows® Server and UNIX® platforms.

Important

Before performing any procedures to setup and configure your system, read this document (especially Chapters 1 and 2) to familiarize yourself with the concepts and features of an SDS system.

Audience

This manual is written for IS System Administrators and support personnel. We assume that you are familiar with the IS Application Executive (Xapex), which includes Storage Library Control, Background Job Control, and Database Maintenance, as well as the IS System Configuration Editor. We also assume that you are familiar with your operating system environment and workstation operations.

Document revision history

IS version	Date	Comment
4.1	June 2008	Documentation refresh.
4.1	Feb. 2008	Extensive revisions to include information on HCAP, sds_import tool, multiple server support, and so on.
4.1	Nov. 2007	Initial release.

Accessing IBM FileNet Documentation

To access documentation for IBM FileNet products:

- 1 Navigate to the Information Management support page (www.ibm.com/software/data/support).
- 2 Select the appropriate IBM FileNet product from the “Select a category” list.
- 3 From the Product Support page, click **Product Documentation** under Learn.
- 4 From the Product Documentation page:
 - a If necessary, click the Doc Link for the appropriate component product to display the document list.
 - b Click the icon in the appropriate release column to access the document you need.

Conventions Used in this Manual

The following paragraphs discuss the ways in which we call your attention to information throughout this document.

Typing Instructions

To indicate commands, values, or other information you enter at your keyboard, we use the following indentation and typeface:

help [CSM_exim]

Screen Displays

Windows and screen examples shown in the procedures in this document are from Windows servers. These windows and screens could look different on UNIX servers.

Console Displays

Information you see displayed at your console is shown in this document in the following manner:

```
Surface '3176' : 1 document processed
Local doc_id = '2235007' Original doc_id = '2235007'
Original ssn = '8502'
Primary copy. No tranlog copy exists.
* document successfully deleted from databases. *
* Purging pages from disk... *
* This document has been successfully purged. *
```

Observe Cautions, Important Notes, Notes, and Tips

Important information and warnings appear in cautions, important notes, notes, and tips. Read these items carefully:

CAUTION Signals possible damaging consequences of an action, such as loss of data or time.

Important Gives added emphasis to notes that contain particularly vital information that must not be skipped over. While all the information in each section is important, these notes are especially important.

Note Draws your attention to essential information you should read.

Tip Introduces an idea that might make your work easier.

Command Syntax

Command syntax definitions are indented:

```
ddexim -e > <filename>
```

Optional Parameters

Optional parameters and keywords are within square brackets:

```
ddexim [-e] > <filename>
```

Required Parameters

Parameters that require you to provide information are shown within angle brackets (< >).

For example, for the following command:

```
ddexim -e > <filename>
```

you must substitute the name of a command for the parameter in angle brackets, such as:

```
ddexim -e > myfile
```

Comments and Suggestions

Send your comments by e-mail to comments@us.ibm.com. Be sure to include the name of the product, the version number of the product, and the name and part number of the book (if applicable). If you are commenting on specific text, include the location of the text (for example, a chapter and section title, a table number, a page number, or a help topic title).

Single Document Storage (SDS) Overview

Single Document Storage (SDS) is an integral function of IBM FileNet Image Services that provides the ability to store documents on magnetic, “Write Once, Read Many” (WORM) protected storage devices, which enforce retention protection at the storage level.

Introduction

The development of SDS has been motivated by the desire from customers to find high-performance, low-maintenance alternatives to optical technology, and the introduction in the market of affordable, very large protected magnetic storage solutions that typically offer the following benefits:

- **Address Regulatory Requirements:** provide advanced, storage-level retention and disposition functionality typically required for compliance regulations such as SEC Rule 17a-4.
- **Virtually Unlimited Scalability:** offers network storage with a wide range of settings and configuration options that match the high scalability of Image Services systems.
- **Transparent provisioning:** as the volume of information continues to grow or as technology upgrades are needed, adding storage is as easy as plugging in another storage unit—with no downtime, planned or unplanned—transparently providing a large storage continuum to Image Services systems.
- **Tiered Storage:** some devices permit the off-load or migration of data that is accessed less often to slower, less expensive media

such as tape, which also provide infinite storage capacity and portable disaster recovery capabilities.

- **Built-in High Availability and Disaster Recovery:** provide retail hardened hardware with built-in storage redundancy using RAID (redundant array of independent disks) or RAIN (redundant array of independent nodes) technology depending on device, often coupled with replication.
- **Reduce Total Cost of Ownership:** provide fast, affordable on-line access to fixed content assets by managing more content with fewer staff, by automating key data and storage management functions with self-management, self-configuration, and self-healing capabilities.

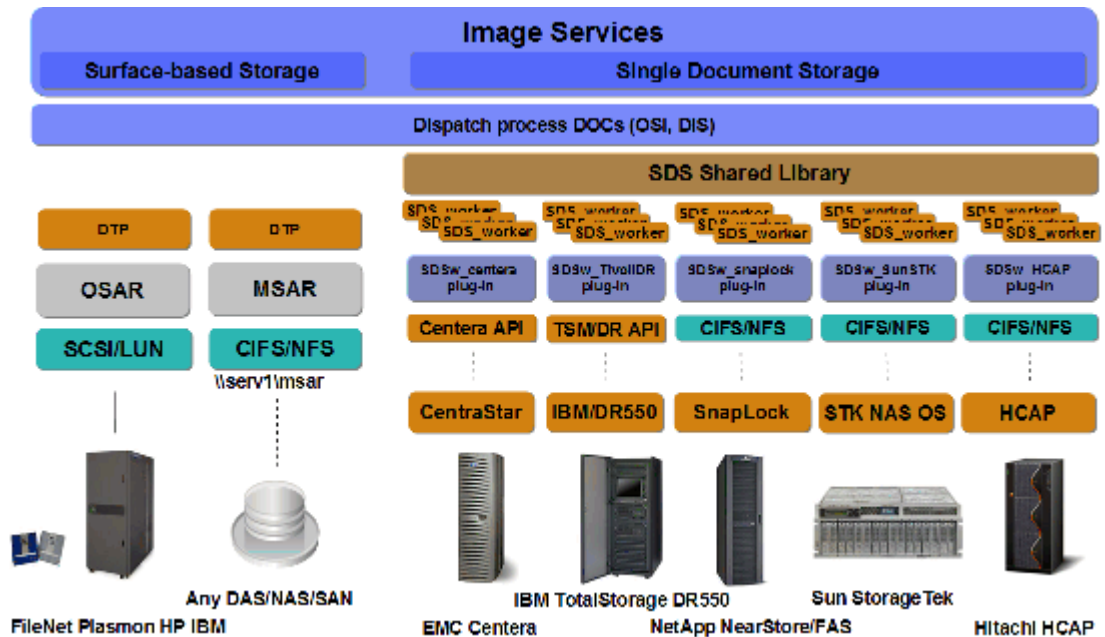
“Single Document Storage” describes the process of storing and retrieving documents as individual content objects through an API or a file system interface. Unlike surface-based devices where document pages are stored on physical or virtual “surface containers” that are internally controlled by IS, documents are stored in and retrieved from SDS devices as individual content objects through an API or a file system interface.

Also, while physical retention of documents is enforced at the platter level in surface-based storage, retention is enforced at the document object level in SDS devices, which gives Image Services the ability to manage and propagate retention and deletion of individual documents at the storage level.

Integral SDS archival and retrieval processes take advantage of the robust service architecture of Image Services on a level equivalent with OSAR/MSAR technology.

Connectivity with each SDS device is implemented through a specific pluggable connector library that is part of IS. This release of IS supports the following devices:

- **EMC Centera**
with the SDSw_centera connector library
- **IBM TotalStorage® DR550**
with the SDSw_tivoliDR connector library
- **NetApp SnapLock**
with the SDSw_snaplock connector library
- **Sun StorageTek 5320**
with the SDSw_SunSAR connector library
- **Hitachi Content Archive Platform (HCAP)**
with the SDSw_HCAP connector library



The device-specific pluggable libraries are loaded on demand by a configurable number of "worker" processes that handle all the low-level operations of writing and reading content to and from the device. Each worker establishes a channel of communication with the device. The more the workers, the higher the throughput and the higher the demand on system resources. While still being subject to the usual hardware or other environment variables that might limit performance, this model allows for a tight, componentized, efficient, and scalable implementation that supports up to 256 independent SDS devices, each with up to 99 worker processes.

Integral SDS workers use the IS Document Buffer Pool (DBP) and shared memory to read from and write to the SDS device. This greatly reduces the footprint of the worker processes and helps manage resources effectively.

SDS object references are maintained in the document locator (DOCS) table in the MKF permanent database in the same way as OSAR/MSAR document references.

Since Integral SDS is at the heart of IS, it is also part of the document entry and migration transactional path. If IS fails during the document capture process, the transaction is either undone or recovered depending at which stage of the process the failure occurs, in exactly the same manner as OSAR/MSAR.

And finally, Integral SDS comes with an efficient SDS object format, in which data is stored internally in a tight layout borrowed from MSAR. This design makes it especially fast to convert from OSAR/MSAR to SDS.

Availability

Integral SDS is available in IBM FileNet **Image Services 4.1.1** and higher.

Compatibility

Integral SDS can retrieve documents from storage devices configured with both NLS-SDS and Integral SDS.

NLS-SDS, the former Single Document Storage product for IS, requires a separate Near Line Storage (NLS) Archive software component and supports only one SDS device.

See the *Image Services Single Document Storage and Retrieval Procedures and Guidelines* for IS 4.1 for more information about NLS-SDS. To download this manual from the IBM support page, see [**“Accessing IBM FileNet Documentation” on page 13.**](#)

Prerequisites

The following prerequisites are required to be in place before you configure your SDS system.

Image Services software

The following releases of IS support integral SDS:

- **IS 4.1.1** and higher

Server memory

Each IS server must have **1 GB** of RAM for use by IS. If the server is also running other applications, this requirement is in addition to the RAM required for those applications.

Special consideration for Windows Servers

If you are configuring a file system repository such as SnapLock or StorageTek in a Windows Server environment, the Windows Server must be a member of a Windows Domain. See [“Appendix D – Windows network security settings” on page 137](#) for more information about configuring integral SDS on a Windows server

Major SDS features

This section provides a functional description of Integral SDS software.

SDS repository information

The Integral SDS repository information is stored in the Document locator (DOCS) table in the MKF permanent database. The DOCS table also contains the content location for surface based storage (that is, the surface_id and offset). Integral SDS uses the following amounts of additional space in the DOCS table (perm_db):

- Each document stored to IS without Integral SDS, takes up 24 bytes*
- Each document stored to IS with SDS and stored to snaplock/HCAP/Sun_SAR, takes up 36 bytes*.

- Each document stored to IS with SDS and centera, takes up 96 bytes*. (The clip_id length is about 53 bytes.)
- Each document stored in IS with SDS and DR550 takes up 72 bytes*.

*Might vary somewhat.

Repository connector libraries

Based on the SDS configuration settings, IS dynamically loads the connector libraries (such as SDSw_centera) that interact with specific repositories.

Standard IS tools manage the configuration information

The **fn_edit** and **fn_build** tools store configuration information in the IS Configuration Database (CDB). The IS System Configuration Editor, **fn_edit**, contains new procedures to add, view, and modify the SDS unit, and the CDB schema has been extended to store the SDS information. The new SDS Storage Library tab enables you to view and modify SDS information.

Database Maintenance Document Family

The Database Maintenance Document Family configuration in Xapex includes support for integral SDS units. Database Maintenance also lists the configured SDS units that are associated with a primary document family, and supports an SDS-only setting for document families that are stored only on an SDS device.

The **ddexim** program also supports the new document family SDS properties. The ddexim tool saves (exports) and restores (imports) document class and index information in the index database. The ddexim program uses the SDS unit name for export and import, rather than the SDS ID.

SDS-only document families

Integral SDS supports SDS-only documents. A document family can be configured so that documents are only associated with an SDS unit. Rather than having a transaction log on other media as a backup, an SDS-only document family depends on the built-in high availability capabilities of the SDS device for redundancy.

Support for multiple SDS units

Integral SDS supports up to 255 SDS units per system. You can still have only one NLS-SDS unit, but you can add up to 255 new integral SDS units to your IS system.

Some advantages of supporting multiple SDS units are:

- Multiple Centera Storage pools can be used on one IS system
- A different SDS unit can be specified for each document family. That is, one SDS unit can be associated with one document family/class, allowing for aggregation of content.
- Any combination of repository types can be supported on an IS system. For example, SDS NetApps SnapLock, DR550, Centera, and Sun StorageTek repository types can co-exist on one IS system.

- Migration of data between different repository types can be done easily with the **sds_copy** tool.
- Each SDS unit can be independently configured with:
 - Different retrieval priorities (low, medium, or high).
 - Different retention settings (none, indefinite, enabled).
 - Different content deletion settings (delete or do not delete).
 - As many as 99 dedicated SDS_worker processes.

Backward compatibility with NLS-SDS

Integral SDS is backward compatible with NLS-SDS. Integral SDS operates in a dual mode where documents can be retrieved from storage devices configured in either in NLS-SDS and Integral SDS.

- You can continue to store NLS-SDS information in the DOCTABA table.
- You can use the **sds_convert_nls** tool to move the NLS-SDS information into the DOCS table.
- You can use the **sds_copy** tool to copy the content itself from an NLS-SDS unit to a integral SDS unit.

All of the above modes can co-exist on a system.

The NLS-SDS readers (SDS_CSAR_reader, SDS_FS_reader and SDS_ISAR_reader) are still used (when configured) to read documents stored in an NLS-SDS repository. The NLS-SDS readers only support read requests.

SDS_tool can still be used with documents that are stored in the NLS-SDS repository unit, but the NLS-SDS information must also be stored in the DOCTABA table.

SDS_tool gives you the following capabilities for managing the NLS-SDS objects:

- You can display and verify NLS-SDS information about stored documents.
- You can recover documents from an NLS-SDS configured device.
- You can modify some of the document's SDS attributes on the SDS device.

The **sds_convert_nls** tool supports moving or copying of NLS-SDS information to the DOCS table. (The move option removes DOCTABA NLS-SDS information, while the copy option leaves the DOCTABA NLS-SDS information in place.)

With the copy option you can continue using SDS_tool on documents stored in an NLS-SDS unit. SDS_tool does not support documents created in Integral SDS. The other Integral SDS tools support similar features as SDS_tool.

Object retention

Object retention is a way of protecting content for a specified period of time. After the retention date for a document has been reached, the document can be deleted.

Image Services retention has always been stored in the DOCTABA table in the index database. With Integral SDS, retention value is

propagated to the SDS object itself and is also stored in the DOCS table in the MKF permanent database.

- Retention of a new SDS object is inherited from the document class when the SDS object is created as a result of committal.
- Retention of an SDS object is inherited from DOCTABA when the SDS object is created as a result of migration, such as when documents in MSAR storage are migrated to an SDS device.

The **sds_update_retention** tool synchronizes the SDS object retention value with the retention value stored in DOCTABA.

Retention types

Retention and disposition for a document class is set in the Disposition portion of the Define/Update Document Classes screen in Xapex. The following illustration shows the Archive and Delete radio buttons, the Months From: field, and the Date Filed and Date Closed radio buttons:

Disposition:

Archive

Delete

Months From:

Date Filed

Date Closed

- **Chronological retention** – the retention value is based on the date a document is created or entered into the repository. The document class disposition “Date Filed” corresponds to chronological retention.
- **Event-based retention (EBR)** – Retention is based on an event that triggers the disposition. In Image Services, closing the

document is the triggering event. Until this event occurs, the retention setting is indefinite. After a document is marked as closed, the retention is set to the current date plus the number of months specified in the “Months From” field. The document class disposition “Date Closed” corresponds to event-based retention.

Note Only certain SDS devices, such as Centera, support EBR.

- **Virtual event-based retention** – for those SDS devices that do not support EBR, such as SnapLock and StorageTek, IS provides a way to emulate event-based retention. The document class disposition “Date Closed” also corresponds to virtual event-based retention.

The SDS Update Retention Tool

The retention settings stored in the SDS device generally mirror the values in DOCTABA in the IS index database. A newly committed document uses the default retention values specified in the document class definition.

Because the triggering event does not occur automatically, both event-based and virtual event-based retention types require that you run the **sds_update_retention** tool periodically to synchronize the retention settings in the SDS object and the MKF permanent database with the settings in DOCTABA.

- For EBR documents on Centera devices, `sds_update_retention` sets the trigger event.
- For virtual EBR documents on SnapLock and StorageTek, `sds_update_retention` sets the trigger event and can optionally extend the retention if the triggering event has not occurred. The `extend` option lets you prolong the retention indefinitely.

- When a retention setting expires, a new retention setting can be applied.
- Retention settings can be extended, but they cannot be shortened.
- If the specified retention is shorter than the existing retention, the existing retention remains.
- The retention of virtual EBR documents can be extended by a specified offset (in number of days) from the current date, or to a specific date (yyyy-mm-dd).
- If the document is already marked closed (triggered) in DOCTABA, then the trigger date takes precedence over the extension.

See [“Update the retention on specific documents” on page 67](#) for more information about the `sds_update_retention` tool.

Indefinite hold

You can use the `sds_hold` tool to place an indefinite hold on any document during litigation or for some other purpose to prevent the document from being deleted if its retention expires. A hold overrides the current retention value, whether chronological or event-based, and must be explicitly released when the hold is no longer needed. When the hold is released, the retention value is returned to its original setting. The document can only be deleted when the retention has expired **and** the hold has been released.

Note Holds are not supported on SnapLock and StorageTek devices.

See [“Place SDS documents on indefinite hold” on page 71](#) for more information about the `sds_hold` tool.

Content deletion

A configurable option in `fn_edit` lets you specify whether the content stored in the SDS unit should be deleted from the SDS unit when reference to the document is deleted from the DOCS table in the MKF permanent database.

Note You cannot delete a document stored in an SDS device within 24 hours of its committal, even if the SDS device does not have a retention setting.

deldocs

The preferred way to delete document content from an SDS device is with the **deldocs** tool. The **-i** option enables you to manually delete content from the SDS device, provided the document has reached the end of its retention and has not been placed on hold. If the content cannot be deleted, the `deldocs` tool does not delete the entries from DOCTABA or the DOCS table.

Document services

When a document is deleted with the Xapex Document Services “Delete Doc/Folder” option, if the retention has not been reached, the document properties in the DOCS table and in DOCTABA are deleted, but the content is **not** deleted.

When a deletion fails, document services writes detailed information about the object in the system log: `doc_id`, retention expiration date, hold, retention type, `sds_obj_id`, and so on. Later, you can delete the content with the `DOC_tool DELETE` option.

Client applications

If you have the appropriate access rights, you can also delete a document after you have displayed it in a client application, such as IDM desktop, or an ISTK application.

This method deletes the document from the DOCS table, DOCTABA, cache, and the SDS device without checking the “delete date” in DOCTABA first. If the delete content option is set, and if the retention has not expired, the content is **not** deleted, and an error message is logged. As with a failed Document Services deletion, you can delete the content with the DOC_tool DELETE option.

Checksumming

Image Services uses checksumming to verify document integrity. Normally, a checksum is created for each page when a document is captured. The checksum is then stored with the document when the document is committed. A configurable option in the IS System Configuration editor allows you to select one of three checksumming options: none, minimum (the default), or maximum.

For documents stored on MSAR and optical media, if a checksum is not calculated when the document is captured, it is not stored with the document, and IS cannot validate the checksum later, even though the option for checksumming is set to maximum.

Integral SDS will use the checksum that is stored during capture in the SDS document BLOb header. However, if the document was captured with no checksum, IS calculates the checksum while writing the document to SDS and appends it to the SDS document BLOb. Upon retrieval, Integral SDS will verify the checksum of every document.

SDS Enable and Disable capability

The integral SDS “enable” and “disable” capability is consistent with surface based enable and disable of each surface. SDS units and surfaces can be read disabled, write disabled, or both.

For SDS units, only DOC_tool **Enable** and **Disable** commands are supported. (Xapex Storage Library Control does not support SDS disable or enable functions.) The SDS unit DOC_tool setting remains the same between restarts of IS.

MSAR/optical surfaces can also be disabled because of an error condition. For example, an MSAR or optical surface might be write-disabled because the MSAR surface file is out of resources. This is different from a surface that has been write-disabled by a user. Surfaces that are disabled because of error conditions cannot be enabled with the Storage Library Control “enable surface” window, but must be addressed through an RSVP window. SDS units are also managed this way.

Note This type of disabled setting of an SDS unit or a surface is **not** persistent after the IS software is restarted, but if the error condition is persistent the SDS unit or surface will be disabled again and an RSVP or INFO message will be regenerated.

When an SDS unit is disabled for reads (for example, if there is a sudden network connection issue), queued read requests are redirected to MSAR or optical devices. If redirection is not possible and notification is enabled, SDS error status is sent through **ds_notify**. If an SDS unit is disabled with DOC_tool, all queued read requests for that SDS unit are redirected to MSAR or optical devices, if possible. Display applications will not hang when an SDS unit is disabled or

goes down if the document is stored on another permanent storage device such as MSAR or optical.

Support for all CFS-IS document types

Integral SDS uses the DOCS table in the MKF permanent database and can support documents federated through CFS-IS from a Content Engine server.

Integration into the IS committal/migration process

Integral SDS is fully integrated into the IS committal/migration process. Every supported committal type has been extended to include SDS: fast batch committal (FBC), regular committal, and cache-only.

The IS media family infrastructure has also been expanded to include SDS units.

Committal/migration auto-restart

Part of being fully integrated into the committal/migration infrastructure of IS includes automatic restart of SDS writes when IS software is restarted, as it is with MSAR and optical media.

Retrieval redirection for individual documents

Integral SDS supports retrieval redirection between SDS and MSAR/Optical. Integral SDS also supports retrieval redirection from MSAR/Optical to SDS. Redirection is done automatically based on the DOCS table information.

Page cache usage

Integral SDS makes extensive use of the native IS page cache instead of using a working directory. Integral SDS write and read requests go directly between the SDS repository and the page cache to increase performance.

Integral SDS takes advantage of the Cache Services (CSM) resource management infrastructure. Writes are done in Document Buffer pool chunk sizes. The Document Buffer Pool buffer size is defined using the IS System Configuration Editor, `fn_edit`.

Error handling

Integral SDS error handling is built into the conventional IS error handling. For example, if the SDS repository runs out of resources to store document objects, the SDS repository is write-disabled and an RSVP message is generated in the Storage Library Control RSVP window.

If an SDS repository is down or cannot be accessed, an INFO message appears in the Storage Library Control window to indicate the state of the SDS unit. The RSVP message triggering can be used with both SDS RSVP and SDS INFO messages.

2

How to set up your SDS system

This chapter provides procedures and recommendations for planning and setting up an SDS system. Follow the procedures in this chapter to setup and configure your SDS system. The main topics provided in this chapter include:

- [“Performance Recommendations and Best Practices” on page 36](#)
- [“Multiple storage library servers with SDS” on page 36](#)
- [“How to configure Single Document Storage devices” on page 38](#)
- [“View the SDS details in the System Configuration Editor” on page 48](#)
- [“Delete the SDS configuration from Image Services” on page 49](#)
- [“Define SDS Families and Document Classes” on page 51](#)
- [“Multi-Committal systems with SDS” on page 52](#)

Performance Recommendations and Best Practices

Many factors can affect the performance of your SDS system. The following list contains some recommendations that could help to improve the performance of your system.

Recommendations that can improve performance

- On systems using local storage, when possible, install all databases (index, permanent, and transient) on separate physical devices.
- If you are using a high throughput attached storage device, put all databases on those devices, thereby gaining flexibility with regard to redundancy and high availability, without sacrificing performance.
- Increase your network connection speed to a **minimum** of 100Mbps.
- Upgrade to an NAS device with a higher rated I/O speed.
- Run in a full duplex data transmission environment, when possible, to increase network throughput speed. Verify that the network settings match on all servers and storage devices in the environment.
- Isolate SDS traffic from other network traffic.

Tip ECM Lab Services provides planning and tuning services.

Multiple storage library servers with SDS

When you configure an SDS device in an IS system that has multiple storage servers, the SDS device must be accessible by every storage

library server in the system. Each storage server contains the descriptions of all the SDS devices, so when you configure an SDS unit with the IS Configuration Editor (fn_edit), the configured repository location string will be used on all storage servers.

Tip For SnapLock, use the same mount point name on all storage servers.

Preferred library settings

The Database Maintenance application in Xapex allows you to select a preferred server and library under certain conditions:

- If the family has MSAR and SDS, or MSAR-only, you can set a preferred server and library. The committal path is through the preferred storage server.
- If the family has MSAR and SDS, or MSAR-only, but you choose not to select a preferred server and library, the committal path alternates between the storage servers.
- However, if the family is defined as SDS-only, you cannot set a preferred server and library. The committal path is through the primary (document locator) storage server.
- If the IS system is a cache-only system, whether the family is SDS-only or MSAR-only, you cannot set a preferred server and library. The committal path alternates between the storage servers.

Tip If you are using **sds_migr** tool to copy or move MSAR surfaces to SDS in a multi-storage library server environment, you can improve system performance if you generate multiple files of document IDs and group them so that all the IDs whose MSAR surfaces belong to the

same storage server are in the same file. Then transfer each file to its associated storage server and run **sds_migr** on each server with the **-doc_file** option.

How to configure Single Document Storage devices

Adjust Performance Tuning Parameters

Use the IS System Configuration Editor to configure the Single Document Storage devices that you need.

Before you configure the SDS device, adjust the performance tuning parameters.

- 1 As a FileNet software user, such as **fnsw**, launch the IS System Configuration Editor by entering:

fn_edit &

- 2 When the Configuration Editor window displays, select the Performance Tuning tab.
- 3 Select the System Processes subtab, and locate the Number of dtp_tran processes field.

If your server has more than one CPU, increase the number of dtp tran processes to half the number of CPUs.

The dtp tran process is responsible for cataloging documents in the Index database using Document Service's fast batch committal option. Performance can be improved if more than one dtp_tran process is running on a system. The maximum number is 16.

Note Increasing the number of dtp tran processes should only be done on systems that have multiple CPUs. The number of dtp_tran processes should never exceed the total number of CPUs on your system.

- 4 Select the Server Memory subtab, and set the following parameters to their maximum values:

Parameter	Maximum Value
Document Buffer Count	256
Document Buffer Size (KB)	1024 KB
Directory Buffer Count	256
Directory Buffer Size (KB)	256 KB

Parameter	Optimum Value
Max. Shared Memory (MB)	1024 MB

- 5 Optionally, increase the Permanent Buffer Pool Size and the Transient Buffer Pool Size.
- The Permanent Buffer Pool Size is 5000 KB (5 MB) by default.
 - The Transient Buffer Pool Size is 100000 KB (100 MB) by default.

You can increase system performance by having a small database entirely resident in memory. The size of these buffers should be almost twice the size (in KB) of the data and index (but not the description) in the database. (The minimum size of the permanent database is 100 MB; the minimum size of the transient database is 320 MB.)

- 6 Exit from fn_edit and save your changes.

Rebuild the IS configuration files and update the MKF databases

After you have increased the Buffer Pool Size (KB) field for one or more of the MKF permanent or transient datasets and/or the Max Shared Memory (MB) field, follow these steps:

- 1 Rebuild the system configuration files:

```
fn_build -a
```

Check that the command runs successfully and that no errors have occurred.

- 2 If you changed only the Max Shared Memory (MB) field, skip to the next step.

If you changed any of the Buffer Pool Size fields, enter the following command to update the appropriate databases on your server:

```
fn_util update<xxxx>
```

where <xxxx> is the MKF database (for example, trans or perm) whose buffer pool size you modified.

Repeat this step for each MKF database whose buffer pool size you increased.

- 3 Stop all Image Services processes by entering the following command:

```
killfnsw -DAy
```

- 4 Start the Image Services software:

```
initfnsw start
```


Make a note of any errors that occur when restarting the Image Services software after a configuration change. In many cases, the size of the recovery logs might also need to be increased in proportion to the new MKF permanent or transient buffer pool sizes.

Set up the repository device

- **Centera:** Create Profile (.pea) file. Skip to **“Appendix B – EMC Centera” on page 115** to create the Pool Entry Authorization file.
- **DR550:** Run the TSM (Tivoli Storage Manager) scripts, which are provided with Image Services 4.1.1, in the /fnsw/etc directory.
- **Snaplock, SunStorageTek, HCAP:** Create a mount point to the device, such as /SDS_device/dept.

Note The FileNet software users (such as **fnsw**) and groups (such as **fnadmin** and **fnop**) must have the appropriate access rights to these devices.

Determine the repository location string

When you add an SDS device to the Image Services configuration, you will be prompted to enter a repository location string. The format of the repository string varies from one SDS device to another.

- **Centera:** the repository location string is in the format <domain name>?<absolute path to the Pool Entry Authorization (.pea) file>.
- **DR550:** the repository location string is the domain name.
- **Snaplock, SunStorageTek, HCAP:** the repository location string is the mount point of the device.

The repository location string for your SDS device might look similar to one of these:

Centera: centera1.yourcorp.com?/fnsw/local/sd/1/merged.pea

DR550: hq-TivoliSM.yourcorp.com

SnapLock: /SDS_snap2/finance

StorageTek: /SDS_sunsar2/sales

Hitachi HCAP /SDS_HCAP/acctng

Define the SDS configuration in Image Services

To configure an SDS device you will need the following information:

- SDS unit type
- SDS repository path or domain name
- Number of SDS workers (1-99, the default is 3)
- SDS unit retrieval priority (high, medium, low)
- SDS unit debug level (off, min, max)
- SDS retention type (none, infinite, enabled)
- SDS retention default (only needed if retention is enabled)
- SDS content deletion enabled? (yes or no)
- SDS foreign unit ssn and name

- 1 As a FileNet software user, such as **fnsw**, launch the IS System Configuration Editor by entering:

fn_edit &

- 2 From the Procedures tab, select, “Configure a Single Document Storage unit” and click **Run**.
- 3 From the list of servers that appears, select the server to add the SDS unit to.
- 4 Select the SDS connector library from the list of supported SDS library types:

- SDSw_snaplock
- SDSw_SunSAR
- SDSw_tivoliDR
- SDSw_centera
- SDSw_HCAP

The SDS connector library provides the software connection between the IS server and the SDS device. If the SDS connector library for the device you are configuring is not listed, but has been qualified as compatible with the current release of IS, select **Other**, and specify the name:

- SDSw_newdevice

where “newdevice” is the specific device type. The new connector library must be in the /fnsw/lib/shobj directory.

- 5 Enter a unique object name for your SDS unit, such as Snaplock1. The maximum length is 18 characters.
- 6 Enter the repository location string for the SDS unit.
- 7 Enter the number of SDS workers.

The main function of the SDS worker process is to handle read and write requests. For read requests, the SDS_worker reads SDS data and writes it to page cache. For write requests, the SDS_worker reads the document content from cache and forms the SDS data object and writes it to the SDS unit. The SDS_worker also handles copy requests, which are similar to write requests.

Three SDS workers are assigned by default, but you can specify as many as ninety-nine (99). You can adjust this number up or down in the future as you monitor system performance using the DOC_tool sds_info utility.

Tip Start with ten workers and then use the **DOC_tool sds** command to see if they're being fully utilized. Add more workers as needed. For more information about the DOC_tool sds command, see [“The DOC_tool utility” on page 81](#).

Too many workers might cause memory allocation issues with the connections to the SDS devices, especially Centera. A 238,3,1 error tuple would be a symptom of this problem. See [“SDSw Repository messages” on page 111](#) for more information.

Enter an appropriate number of SDS workers.

8 Select the SDS retrieval priority.

Each SDS unit can have a different retrieval priority. For example, on an IS system with optical, MSAR, and SDS storage devices, setting high SDS priority means that IS will send retrieval requests to the SDS unit first, MSAR second, and optical media last. The priority settings correspond to the following retrieval sequences:

- **high** – SDS, MSAR, optical

- **medium** – MSAR, SDS, optical
- **low** – MSAR, optical, SDS

Choose the appropriate priority for this SDS unit. (For best SDS performance, select **high**.)

9 Select the SDS debug level.

The debug setting you choose (off, min, or max) determines the amount of information that is reported in the system log. Use the **vl** tool to view the system log.

- **off** – when debug is set to off, no SDS activity is reported in the system log.
- **min** – integral SDS logs a single message for each retrieval, write, or copy operation.
- **max** – integral SDS continues to log the same information as the “min” setting, but additional information (doc_id, request priority, target cache id, cache ssn, first_page, number of pages) is logged when the request is complete. The SDS_worker also logs SDS object information such as the ClipID for Centera devices.

Tip On a production system, setting the debug level to min or max could have a negative performance impact.

10 Select the document retention type.

- **none** - no retention policy is established.
- **infinite** - documents are retained permanently. (This retention type is the equivalent of WORM–Write Once, Read Many.)

- **enabled with default offset** - enables document retention.
- **enabled with default date** - enables document retention.

If you select none, the retention is actually set to the current time plus one day. As a result, you cannot remove a document immediately after committal. You must wait one day before removing it. Document content is always protected against modification.

If you select enabled with default offset or default date, virtual EBR is enabled, and the value will be set to be consistent with the retention setting in DOCTABA. (See [“Object retention” on page 26](#) for more information).

- 11** If you selected **enabled with default offset** in the previous window, you are prompted: Enter the default retention offset (days).

Enter the number of days (up to a maximum of 12,000) that you want to retain the documents stored in this SDS unit. This value is only used for virtual EBR.

- 12** If you selected **enabled with default date**, in the retention type window, enter the default retention date [yyyy-mm-dd].

Enter the future date when the documents stored in this SDS unit will expire. This value is only used for virtual EBR. For example, you might enter **2027-12-31**.

- 13** Content deletion enabled?

When a document is deleted from an Image Services system, whether by Document Services or by the deldocs tool, the document properties are deleted from the MKF permanent database (DOCs table) and the index database (DOCTABA) if the retention has expired and the

document is not on hold. The document content remains in storage unless you enable content deletion. See [“Content deletion” on page 30](#) in Chapter 1 for more information.

Select **Yes** to make the document content eligible for deletion from the SDS unit when the document properties are deleted.

14 Is this a foreign SDS unit?

If this SDS unit was previously configured on another IS system and contains data from that system, it is considered a “foreign” SDS unit. Foreign devices are read-only on the current system; no new documents can be committed to this device. Documents from the “foreign” SDS unit can be imported to the current SDS device with the `sds_import` tool. See [“Restore or import data to an SDS unit” on page 77](#) in the next chapter.

Select **Yes** if this SDS unit is “foreign” to your current IS system. Otherwise, select **No**.

15 If you indicated in the previous window that the SDS unit is “foreign,” you are prompted to enter the original system serial number.

Enter the System Serial Number of the IS system to which this “foreign” SDS unit was originally attached. The System Serial Number is a unique ten-digit number (provided by IBM) from 0000000001 to 4294967295. For example, 0000100728.

16 If you indicated that the SDS unit is “foreign,” enter the original SDS name, such as `snaplock5`, that was configured for this “foreign” SDS unit on the original IS system.

17 When the procedure has completed successfully, save your changes and exit from the System Configuration Editor.

- 18 Stop the IS Software and run **fn_build -a** to regenerate the system configuration files. Your SDS unit is ready to put into production.

View the SDS details in the System Configuration Editor

After you have configured the SDS unit, you can view the details in the System Configuration Editor on the Single Document Storage tab.

- 1 Launch the System Configuration Editor.
fn_edit &
- 2 Select the Single Document Storage tab.
- 3 Verify that the information displayed is correct. Make any needed changes. Some fields, such as SDS unit number, SDS name, and SDS connector library, cannot be modified.

Important

Use extreme caution if you change the **repository location string**. You might not be able to access documents that have already been committed.

The system looks for the most recent successful configuration settings in `/fnsw/local/logs/sds`. If you enter the wrong repository location string, then restart IS, and later correct it, you will continue to receive errors.

When you change the repository location string, you must also move or delete the sds device's log files in `/fnsw/local/logs/sds`.

– For new system configurations, if the initial setting is incorrect, then you must remove the logs for a successful change to occur.

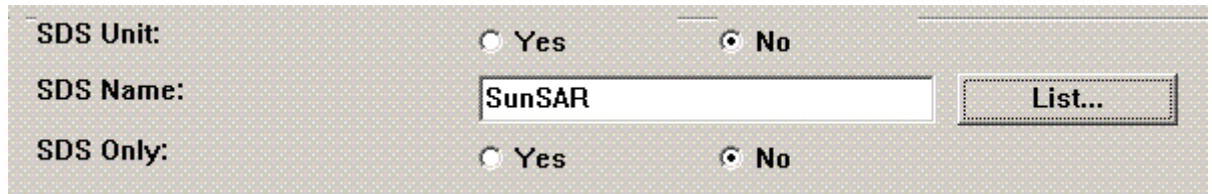
– For existing system configurations, incorrect changes will not take effect if prior successful logs exist.

- 4 Save your changes and exit from the System Configuration Editor.
- 5 Stop the IS Software and run **fn_build -a** to regenerate the system configuration files.

Delete the SDS configuration from Image Services

In the event that you need to remove an SDS device from your IS configuration, follow these steps. If there are multiple SDS units configured on your IS system, you will need to know the name of the unit you want to remove.

- 1 If any document families are associated with the SDS unit you plan to delete, use the Database Maintenance Define/Update Families window in the Application Executive (Xapex) to remove all document family associations with that SDS unit. Click the SDS Unit: **No** radio button, as shown in the following example.



The screenshot shows a dialog box with the following fields and controls:

- SDS Unit:** Radio buttons for **Yes** and **No**. The **No** button is selected.
- SDS Name:** A text input field containing the text "SunSAR". To the right of the field is a button labeled "List..".
- SDS Only:** Radio buttons for **Yes** and **No**. The **No** button is selected.

- 2 Launch the IS System Configuration editor.
fn_edit &
- 3 On the Procedures tab, select the procedure to “Delete a Single Document Storage unit” and click **Run**.

- 4 If more than one SDS unit has been configured on your IS system, you are prompted, “Select the SDS unit to delete.”

A list of the Single Document Storage units that are currently configured on your IS system appears in the window. Select the unit you want to delete, and click **Next**.

- 5 Then you are prompted, “Is this SDS unit associated with any document families?”

If no families are associated with this SDS unit, click the **No** radio button.

Note If you click the **Yes** radio button, you will see an error message when you try to continue.

Error: Cannot delete an SDS unit that is associated with a document family

Use the Database Maintenance application in the Application Executive (Xapex) to remove any document family associations with the SDS unit you are planning to delete. When all families have been removed, rerun this procedure to delete the SDS unit.

- 6 When the procedure to delete the SDS unit is finished, select the Single Document Storage tab and verify that the configuration values for the SDS unit you deleted no longer appear.

Define SDS Families and Document Classes

Use this information to define new SDS families and classes.

How to define SDS Families

Define SDS Families using the Database Maintenance application of the IS Application Executive (Xapex).

Important

Select Primary as the Family type. If you select Tranlog as the Family type, an error message will display when you try to save the SDS family settings.

- To associate the documents in this family with an SDS unit (not MSAR or optical), click the SDS Unit: **Yes** radio button at the bottom of the Define/Update Families window.
- To assign the media family to an SDS unit, choose the appropriate SDS Name from the list of currently configured SDS units.
- If the documents in this family will be associated with only an SDS unit (not MSAR or optical), click the SDS only: **Yes** radio button at the bottom of the Define/Update Families window.

The screenshot shows a configuration window with three rows of controls:

- SDS Unit:** Two radio buttons, "Yes" (selected) and "No".
- SDS Name:** A text input field containing "SunSAR" and a "List..." button to its right.
- SDS Only:** Two radio buttons, "Yes" and "No" (selected).

When configuring an SDS-only family, you must select a media type, but it is not used.

Supported family configurations

- SDS-only media
- MSAR/optical primary + SDS
- MSAR/optical primary family + MSAR or optical transaction log + SDS

Note IS does not support a transaction log family on an SDS device.

If no libraries such as MSAR or optical are configured on the system, the Migration Delay and the No Migrate options in the document class are ignored and documents will migrate to SDS right away. Otherwise, the Migration Delay and the No Migrate options are used with SDS as usual.

Configure the document class

Configure the document class to use with the new SDS family. There is no difference from the usual procedure.

Multi-Committal systems with SDS

Integral SDS is supported with multiple system committal. Both the target and the source families can be associated with SDS units.

Compatible systems

In a multi-committal environment where both IS systems are compatible (the ranges of document ID numbers do not overlap each

other), you can define a family on one system that points to a remote SDS family.

Each compatible IS system has a unique range of document ID numbers. For example, if source IS system A has a range of document ID numbers from 100,000 to 5,000,000, and target IS system B has a range of document ID numbers from 6,000,000 to 9,000,000, the two systems are compatible.

Incompatible systems

If the two IS systems are incompatible (the ranges of document ID numbers overlap), do not define a family that points to a remote SDS family.

For example, when you commit an MSAR document to a remote family, the MSAR document will be successful because the IS software writes the document to a unique MSAR file and assigns a local document ID number that does not conflict, but the SDS document will not be stored on the SDS device. Instead, you will see an error like this in the system log:

```
2007/11/13 15:37:02.340 79,0,23 <fnsw> dtp b 1 (29226) ...[WARNING]
Fastbatch object 4160000000 is from an incompatible system ssn=101410110. Skip
writing all documents with doc_ids (100723-100724) to SDS device id 5.
```

If the target family on the incompatible system is SDS-only, an SDS_worker on the target system will remove the cache object and the write request, and will report the following error in the system log:

```
2007/11/13 15:37:02.340 79,0,23 <fnsw> SDS_worker 5 2 (29226) ...[WARNING]
Fastbatch object 4160000000 is from an incompatible system ssn=101410110. Skip
writing all documents with doc_ids (100725-100726) to SDS device id 5.
```

3

How to manage your SDS system

This chapter contains administrative procedures for managing and maintaining your SDS system. The following procedures are included:

- [“SDS system monitoring” on page 56](#)
- [“Convert and move data to integral SDS” on page 61](#)
- [“Update the retention on specific documents” on page 67](#)
- [“Place SDS documents on indefinite hold” on page 71](#)
- [“Copy between two SDS units” on page 73](#)
- [“Restore or import data to an SDS unit” on page 77](#)
- [“The DOC_tool utility” on page 81](#)

Important

The **docchk** tool cannot be used with SDS documents. The **docchk** tool only works with surface-based storage, such as optical or MSAR.

The **stdoccpy** tool cannot be used to copy documents from an SDS-only family to MSAR or optical media.

SDS system monitoring

To monitor the health and performance of your IS system, you can use the following tools:

Monitor RSVP and INFO messages in Storage Library Control (SLC)

Log on to the Image Services Application Executive (Xapex), and select Storage Library Control from the Applications menu.

The title of the SLC main window shows the name of the local storage library server. The SLC main window has four sections:

- Informational Messages section displays status information.
- RSVP Messages section displays operator instructions.
- The RSVP/INFO Trigger field indicates whether or not the RSVP/INFO triggering script is enabled.
- Storage Libraries section displays a line for each library or ODU attached to the server.

The SLC RSVP and INFO messages have been expanded to include SDS information.

See the *Image Services System Administrator's Handbook* for more information on Storage Library Control. To download this manual from the IBM support page, see [**“Accessing IBM FileNet Documentation” on page 13.**](#)

RSVP scripting

Two new string-type arguments have been added to the list of fixed arguments that are passed to the messaging script you create. Like the other arguments, the `sds_id` and `sds_name` arguments apply to both RSVP and INFO messages.

Argument	Field	Type	Valid Range	Comments
18	<code>sds_id</code>	integer	2-255	The id number assigned to the SDS unit when it was configured with the IS Configuration Editor.
19	<code>sds_name</code>	Variable string[250]		The name you assigned to the SDS unit in the IS Configuration Editor. For example, <code>DR550_acct</code> .

See Appendix E, “Message Triggering” in the *Image Services System Administrator’s Handbook* for more information on creating customized scripts. To download this manual from the IBM support page, see [“Accessing IBM FileNet Documentation” on page 13](#).

IS Syslog and Subsystem Logging

Errors, warnings, and informational messages are recorded in the standard system log. Use the `vl` command to view the log file.

The SDSw connector libraries return device-specific errors from the SDS unit to IS, as in the following examples.

SnapLock error message:

```
SDS_open_object: Error from SDSw_open_object  err_msg='SDSw_open_object:
open() failed errno=2 docfile='/SDS_sunsar/SDS_test2/101410035/SunSAR/
00/0F/42/1000000' Does not exist' doc_id=1000000
```

Centera error message:

```
SDS_open_object: Error from SDSw_open_object  err_msg='FPCLip_
Open[Flat](Pool=2306434306920808704, Clip=0)Error ==-10006, SysErr=0
Mess  =Wrong parameter detected Trace
=FPObjectGuid::SetID(xxxxx)<FPCLip_Open(-,xxxxx,2)' doc_id=10000000
```

DR550 error message:

```
SDS_open_object: Error from SDSw_open_object  err_msg='Can't find object,
ObjName=/DR550/00/98/96/10000000 ANS0272I (RC121)  The operation is
finished.' doc_id=10000000
```

IS Performance Monitoring (perf_mon)

The **perf_mon** tool gathers statistical data from the following areas: cpu usage for kernel and non-kernel activities; overall file system activities; kernel file system activities; network I/O statistics; database server counts; document services counts, RPC statistics, and security statistics. The perf_mon tool also reports read and write counts of SDS documents.

See the *Image Services System Tools Guide* for more information on using the **perf_mon** tool. To download this manual from the IBM support page, see [“Accessing IBM FileNet Documentation” on page 13.](#)

Simple Network Management Protocol

Simple Network Management Protocol (SNMP) is an industry-standard protocol for network management. SNMP enables network administrators to manage network performance, find and solve network problems, and plan for network growth.

You can use your SNMP-compliant network management software to access a wide variety of information in your IS system through the IS Management Information Base (MIB), which has been enhanced to support SDS.

Although the SNMP protocol is a standard, there are many variations in specific implementations. Always refer to the manuals that came with your network management software for details.

See the *Image Services SNMP Reference Manual* for more information on configuring and using the SNMP. To download this manual from the IBM support page, see [“Accessing IBM FileNet Documentation” on page 13.](#)

SDS System States

When an error occurs, integral SDS might change the system state.

- The SDS System states are **enabled**, **disconnected**, **read-only**, or **write-only**.

- When the System state changes, an SDS_worker issues an RSVP or INFO message. For example, if the SDS_worker does not have write permission to a repository, the SDS device will be set to read-only and an RSVP message will be generated.
- If an SDS unit is disconnected, it can be reconnected automatically and set to enabled System state.
- These states are changed automatically as a result of system conditions, and they do not remain in effect when the IS software is restarted.

The **DOC_tool sds** command can be used to view the SDS sytem state.

About SDS background job tools

A suite of command line tools help you manage documents in your SDS environment:

sds_convert_nls

converts documents from NLS-SDS to integral SDS

sds_copy

copies documents between two SDS units

sds_hold

places a document on indefinite hold or releases a hold

sds_import

imports SDS documents from another IS system

sds_migr

migrates documents from MSAR or optical media to SDS

sds_update_retention

updates the retention for EBR documents

All **sds_*** command line tools run in the background, and you can monitor their progress with Background Job Control (BJC) in Xapex. To view the usage statement for each tool, run the command with no options.

These tools generate the standard BJC progress and completion logs. They also have a **-writelist** option that generates a list of documents processed.

Each **sds_*** tool has the standard BJC capabilities: Auto restart, suspend, abort, and log files. See Chapter 6, “Background Job Control” in the *Image Services System Administrators Handbook* for a complete description of BJC. To download this manual from the IBM support page, see [“Accessing IBM FileNet Documentation” on page 13](#).

Convert and move data to integral SDS

The tools described in this section will help you copy, move, and convert data to integral SDS. They include:

- [“Convert an NLS-SDS unit to integral SDS” on page 61](#)
- [“Convert an MSAR/optical/cache system to integral SDS” on page 64](#)

Convert an NLS-SDS unit to integral SDS

This procedure describes how to convert an existing NLS-SDS unit to use integral SDS.

The **sds_convert_nls** tool moves the NLS-SDS references from the index database to the MKF permanent database.

The SDS information in DOCTABA can be moved or copied to the DOCS table. The document content remains in place and unchanged in the SDS repository.

These documents can still be retrieved by either the SDS_CSAR_reader, the SDS_ISAR_reader, or the SDS_FS_reader process. In addition, use the `-copy` option to preserve the index database entry. This allows for the continued support of the NLS-SDS **SDS_tool** which relies on the index database entry to operate.

In a multiple storage library server environment, run this tool on the primary storage library server.

Log on as a member of the **fnadmin** group to run **sds_convert_nls**.

Syntax

```
sds_convert_nls  { -doc_file <filename> |
                  -first <doc_id> -last <doc_id> |
                  -all }
                  [ -copy ]
                  [ -overwrite ]
                  [ -writelist ]
```

Required parameters

Only one of the document specifiers (`doc_file`, `range`, or `all`) is required.

- doc_file <filename>** Specifies the name of a text file of document ids, one id per line.
- first <doc_id>** Specifies the first document in a range.
- last <doc_id>** Specifies the last document in a range.

-all Migrate all documents to integral SDS.

Optional parameters

-copy Copies, not moves, the archive and ClipID index values to the MKF permanent database.

-overwrite Allows migration again even if the documents have already been migrated.

-writelist Generates a list of migrated documents in a file named **ConvertNLSDocs.nnn**. The file is in the `/fnsw/local/logs/bkglog/` directory.

Examples

To copy all SDS-related database entries from DOCTABA to the DOCS table, and create a list of the copied documents, enter:

```
sds_convert_nls -copy -all -writelist
```

(With the `-copy` option, the entries in DOCTABA are preserved so they can be used by `SDS_tool`.)

To move a range of SDS-related database entries from DOCTABA to the DOCS table, enter:

```
sds_convert_nls -first 24680 -last 86420
```

(Without the `-copy` option, the DOCTABA entries will be removed.)

Convert an MSAR/optical/cache system to integral SDS

The **sds_migr** tool copies or moves documents from conventional storage (MSAR, optical, or cache-only) to SDS.

Note When you use the `-sds_only` option, which indicates move mode, the references to surface locations are removed from the Document Locator table in the MKF permanent database. Future retrievals will only be done from the SDS unit, not MSAR, optical, or cache.

When a document is copied or moved with the **sds_migr** tool, the retention is set in the target object.

If you are copying or moving MSAR surfaces to SDS in a multiple storage library server environment, you can run this tool on any storage server, even if the MSAR surfaces are on another storage server.

Note IS does not support the use of the **stdoccopy** tool to copy documents to or from an SDS unit.

Use the **stdoccopy** tool only for surface-based copying.

Log on as a member of the **fnadmin** group to run **sds_migr**.

Syntax

```
sds_migr  -sds <SDS unit name>
          { -doc_file <filename> |
            -first <doc_id> -last <doc_id> |
            -all }
          [ -overwrite ]
          [ -sds_only ]
          [ -use_page0 ]
          [ -writelist ]
```

Required parameters

(The SDS unit name and only one of the document specifiers (doc_file, range, or all) are required.)

- sds <SDS unit name>** The configured SDS unit name. NLS/CSAR (SDS 1) cannot be used.
- doc_file <file-name>** Specifies the name of a text file of document ids, one id per line.
If you are copying or moving MSAR surfaces to SDS in a multi-storage library server environment, you can improve system performance if you generate multiple files of document IDs and group them so that all the IDs whose MSAR surfaces belong to the same storage server are in the same file. Then transfer each file to its associated storage server and run sds_migr on each server with this (-doc_file) option.
- first <doc_id>** Specifies the first document in a range.
- last <doc_id>** Specifies the last document in a range.
- all** Migrate all documents to SDS.

Optional parameters

- overwrite** If the document is already associated with an SDS unit and the -overwrite flag is selected, the document will be rewritten to the target SDS unit. By default, sds_migr will not migrate a document if it is already associated with an SDS unit. (If the target and the current SDS unit are the same, this option is ignored.)
- sds_only** Deletes references to the migrated documents on MSAR or optical media from the DOCS table in the MKF permanent database.
- use_page0** By default, the metadata stored with the SDS object is extracted from DOCTABA when possible. If the -use_page0 option is selected, the metadata is taken from page0 of the document. Select this option to speed up the copy job because it will not need to extract this information from DOCTABA. Generally, it is preferable to use the metadata that is stored in DOCTABA because the indexing information might have been modified.

Note This option takes the **original** index values from the document on the optical or MSAR surface. If there are updated short descriptors (created by the stsurfupdate tool) with newer index values on the surface, they are ignored.

- writelist** Generates a list of migrated documents in a file named **SDSMigrDocs.nnn**. The file is in the /fnsw/local/logs/bkglog/ directory.

Examples

To copy all documents to an SDS unit named snaplock1, enter:

```
sds_migr -sds snaplock1 -all
```

To copy a range of documents to an SDS unit named centera42, enter:

```
sds_migr -sds centera42 -first 534654 -last 999973
```

Tip If you have a large number of documents to migrate, you can optimize the process by running multiple instances of `sds_migr` using different document number ranges.

If you are migrating from optical media, do not run more instances of `sds_migr` than there are drives in the optical library to prevent unnecessary disk swapping.

To move a group of documents whose document ids are listed in a text file to the SDS unit named `SDS_solaris`, enter:

```
sds_migr -sds SDS_solaris -doc_file doclist.txt -sds_only
```

Update the retention on specific documents

The `sds_update_retention` tool updates the retention of specific documents based on their retention settings in DOCTABA in the index database. Use this tool to update the virtual event-based retention on repositories such as SnapLock that do not support EBR.

- If DOCTABA indicates that the document has non-event-based retention, `sds_update_retention` will update the SDS object retention if it is not consistent with DOCTABA.
- For event-based documents (whose retention is determined only when the document is closed):

- **If the SDS device supports EBR:**

If the event is triggered in DOCTABA (a close event has occurred and a delete date is set), `sds_update_retention` will

set the SDS object retention to be triggered with the retention expiration date in DOCTABA.

If the event is not triggered, `sds_update_retention` will not change the SDS object retention.

- **If the SDS device does not support EBR:**
(also known as virtual EBR)

Run `sds_update_retention` with the `-extend` on a regular schedule to protect virtual EBR objects before the triggering event occurs.

If the object has not been triggered, the `-extend` value is applied. For example, if you run `sds_update_retention` weekly, you could extend retention by seven days each time. If the object has been triggered, the retention setting in DOCTABA is applied.

See [“Object retention” on page 26](#) for more information.

Important

The `sds_update_retention` tool operates **only** on SDS documents that are cataloged in DOCTABA in the IS index database. Documents captured and cataloged on Content Engine and stored in IS cannot be updated with this tool because they are not cataloged in the IS index database.

In a multiple storage library server environment, run this tool on the primary storage library server.

Log on as a member of the `fnadmin` group to run `sds_update_retention`.

Limitations

- The sds_update_retention tool is not supported on DR550 devices because DR550 uses predefined document management classes.
- Retention cannot be decreased.
- Documents set to infinite retention cannot be modified.
- Chronological retention cannot be changed to EBR.
- On Centera devices, retention type “none” cannot be changed to EBR using the -override option.

Syntax

```
sds_update _retention  -sds <SDS unit name>
                        { -doc_file <filename> |
                          -first <doc_id> -last <doc_id> |
                          -all }
                        [ -extend [yyyy-mm-dd | <offset> ] ]
                        [ -overwrite ]
                        [ -writelist ]
```

Required parameters

The SDS unit name and one of the document specifiers (-doc_file, -first/-last, or -all) are required.

-sds
<SDS unit name>

The configured SDS unit name. NLS SDS (SDS 1) cannot be used.

-doc_file
<filename>

Specifies the name of a text file of document ids, one id per line.

- first <doc_id>** Specifies the first document in a range.
- last <doc_id>** Specifies the last document in a range.
- all** Updates the retention of all documents on the specified SDS unit.

Optional parameters

- extend** (Only applies to documents with virtual EBR retention.)

This option extends the retention of the specified documents either to a specific future date (yyyy-mm-dd), or by a certain number of days from the current date. (For example, if the current retention expires 89 days from today, and you specify -extend 90, you only add one day to the retention value.)

Also, you cannot shorten a retention setting. If the retention value you specify is shorter than the existing retention, the existing retention remains.

For virtual EBR documents, the trigger overrides the -extend option.

- override** For documents with a retention setting of “none,” this option forces the retention update to be based on their current settings in DOCTABA. However, on Centera devices, retention cannot be changed from “none” to EBR using this option.
- writelist** Generates a list of updated documents in a file named **SDSUpdateDocs.nnn**. The file is in the /fnsw/local/logs/bkglog/ directory.

Examples

To update the retention to a specific date on all documents in an SDS repository named SunStorTek_Accts, enter:

```
sds_update_retention -sds SunStorTek_Accts -all -extend 2017-12-31
```

To extend the virtual EBR retention of a certain range of documents for sixty days from today, and create a list, enter:

```
sds_update_retention -sds SunStorTek_Accts -first 123123 -last 234234 -extend 60 -writelist
```

To update retention for a group of documents listed in a file named uprent.txt based on the current retention setting in DOCTABA, enter:

```
sds_update_retention -sds SunStorTek_Accts -doc_list uprent.txt
```

Place SDS documents on indefinite hold

The **sds_hold** tool sets the retention of documents to an indefinite hold or releases an existing hold.

Placing a hold on a group of documents overrides the standard retention values of those documents, which might be needed for litigation or some other purpose. The standard retention values are maintained, and even though the retention of a particular document expires, the hold prevents the document content from being deleted from the system. The document properties in the DOCS table and DOCTABA can still be removed. Only one hold can be placed on a document. Additional hold requests are ignored.

Note The **sds_hold** tool is not supported with SnapLock or StorageTek devices.

In a multiple storage library environment, run this tool on any storage server. Log on as a member of the **fnadmin** group to run **sds_hold**.

Syntax

```
sds_hold    -sds <SDS unit name>
             { -doc_file <filename> |
               -first <doc_id> -last <doc_id> |
               -all }
             [ -release ]
             [ -writelist ]
```

Required parameters

The SDS unit name and one of the document specifiers (-doc_file, -first/-last, or -all) are required.

- sds**
<SDS unit name> The configured SDS unit name. NLS/CSAR (SDS 1) cannot be used.
- doc_file**
<filename> Specifies the name of a text file of document ids, one id per line.
- first <doc_id>** Specifies the first document in a range.
- last <doc_id>** Specifies the last document in a range.
- all** Holds (or releases) all documents to the specified SDS unit.

Optional parameters

- release** Releases the hold on the specified documents.
- writelist** Generates a list of held or released documents in a file named **SDSHoldDocs.nnn**. The file is in the `/fnsw/local/logs/bkglog/` directory.

Examples

To place an indefinite hold on a range of documents on an SDS unit named `Cent_archive`, enter:

```
sds_hold -sds Cent_archive -first 139383 -last 949392
```

To release the hold on a group of documents listed in a file named `release.txt` and create a list, enter:

```
sds_hold -sds Cent_archive -doc_file release.txt -writelist
```

Copy between two SDS units

The **sds_copy** tool copies data between SDS units. The source SDS unit can be either NLS-SDS unit type or integral SDS unit type. The target SDS **cannot** be an NLS-SDS unit since writes are not supported.

Note IS does not support the use of the **stdoccopy** tool to copy documents to or from an SDS unit.

Use the **stdoccopy** tool only for surface-based copying.

You can copy content from any SDS unit type to any other SDS unit type (for example, from Centera to DR550). When a document is copied between newer SDS unit types, the content will be removed from the source SDS unit if the content deletion setting in `fn_edit` is Yes. If the retention has not expired, then the source SDS object will not be removed, but the copy will be successful and a message will be logged.

When content is also available on MSAR, it will be read from MSAR because copying from MSAR has been optimized.

When a document is copied or moved with the **sds_copy** tool, the retention is set in the target object to be consistent with retention setting in DOCTABA (if retention is enabled on the target SDS unit).

If a document is on hold in the source SDS unit, it will not be copied. A message will be logged to indicate that the held document must be released before it can be copied.

In a multiple storage library server environment, run this tool on any storage library server.

Log on as a member of the **fnadmin** group to run **sds_copy**.

Syntax

```
sds_copy  -source <SDS unit name>
          -target <SDS unit name>
          { -doc_file <filename> |
            -first <doc_id> -last <doc_id> |
            -all }
          [ -sds_only ]
          [ -use_page0 ]
          [ -writelist ]
```

Required parameters

The `-source` option, the `-target` option, and one of the document specifiers (`-doc_file`, `-first/-last`, or `-all`) are required.

-source
<SDS unit name>

The configured source SDS unit name (required).

-target
<SDS unit name>

The configured target SDS unit name (required).

-doc_file
<filename>

Specifies the name of a text file of document ids, one id per line.

-first <doc_id>

Specifies the first document in a range.

-last <doc_id>

Specifies the last document in a range.

-all

Migrates all documents to the target SDS unit.

Optional parameters

-sds_only

Deletes MSAR or optical references to the copied documents from the MKF permanent database.

-use_page0 By default, the metadata stored with the SDS object is extracted from DOCTABA when possible. If the `-use_page0` option is selected, the metadata is taken from page0 of the document. Select this option to speed up the copy job because it will not need to extract this information from DOCTABA. Generally, it is preferable to use the metadata that is stored in DOCTABA because the indexing information might have been modified.

Note This option takes the **original** index values from the document on the optical or MSAR surface. If there are updated short descriptors (created by the `stsurfupdate` tool) with newer index values on the surface, they are ignored.

-writelist Generates a list of migrated documents in a file named **SDSCopyDocs.nnn**. The file is in the `/fnsw/local/logs/bkglog/` directory.

Examples

To copy a group of documents whose document id numbers are listed in a file named `copydocs.txt`, from an SDS device named `OldCentera` to an SDS device named `NewSnaplock`, enter:

```
sds_copy -source OldCentera -target NewSnaplock -doc_file copydocs.txt
```

To copy a range of documents from an SDS device named `DR550_finance` to an SDS device named `DR550_acctg`, enter:

```
sds_copy -source DR550_finance -target DR550_acctg -first 123456 -last 234567
```

To copy all documents from an SDS device named `Snaplock_A` to an SDS device named `Snaplock_B`, and create a list of copied documents, enter:

```
sds_copy -source Snaplock_A -target Snaplock_B -all -writelist
```

Restore or import data to an SDS unit

The **sds_import** tool imports or restores data to an SDS repository, and works much like the **stdocimp** tool, which is used with data on MSAR and optical media. The DOCS and DOCTABA tables can be recovered using this tool. Like **stdocimp**, **sds_import** assumes that the INX data dictionary information is restored using a tool such as **ddexim**. The **sds_import** tool can also be used to import an SDS repository to another IS system.

When a document is imported or recovered, the retention setting of the object remains unchanged. The SDS object retention setting is propagated to DOCTABA.

In a multiple storage library server environment, run this tool on the primary storage library server.

Log on as a member of the **fnadmin** group to run **sds_import**.

Syntax

```
sds_import  -sds <SDS unit name>
            { -doc_file <filename> |
              -first <doc_id> -last <doc_id> |
              -all }
            [ -nonexactclass ]
            [ -noinsertdoctaba ]
            [ -security { none | doc } ]
            [ -updatesnt ]
            [ -noworkfloqueue ]
            [ -writelist ]
```

Required parameters

The SDS unit name and one of the document specifiers (-doc_file, -first/-last, or -all) are required.

- sds**
<SDS unit name> The name of a configured SDS unit. NLS/CSAR (SDS 1) cannot be used.
- doc_file**
<filename> Specifies the name of a text file of document ids, one id per line.
- first <doc_id>** Specifies the first document in a range.
- last <doc_id>** Specifies the last document in a range.
- all** Migrates all documents to the target SDS unit.

Optional parameters

- nonexactclass** This option allows you to import documents even if the source and target document classes do not match exactly.
- noinsertdoctaba** Does not insert documents into DOCTABA. You must also specify the -security none option if documents from non-existent classes are being imported.
- security**
{ none | doc } Sets the document security. If you specify -security none, the security of the imported document is set to “anyone.” If you specify -security doc, the security is set to the value from the document on optical disk.

The default is to use the security from the document class.
- updatesnt** Updates the scalar_number table in the MKF permanent database if the document id to be imported is greater than the next available document id.

-noworkfloqueue Does not insert the document into a WorkFlo queue if a queue is defined for the document class.

-writelist Generates a list of imported documents in a file named **SDSImportDocs.nnn**. This file is in the `/fnsw/local/logs/bkglog/` directory.

Examples

To import a group of documents whose document numbers are listed in a file named `importdocs.txt`, to an SDS device named `NewCentera`, enter:

```
sds_import -sds NewCentera -doc_file importdocs.txt
```

To import a range of documents to an SDS device named `DR550_acctg`, enter:

```
sds_import -sds DR550_acctg -first 123456 -last 234567
```

To restore all documents to an SDS device named `Snaplock_A`, enter:

```
sds_import -sds Snaplock_A -all
```

Using `sds_import` on Centera

On Centera devices, when extending the retention of an existing object, a separate object with the same document id is created on the device. As a result, there can be duplicate objects with the same document id.

Consequently, if you run `sds_import` to import the document, all the objects with the same document id will be processed.

This results in multiple counts in the sds_import log file and the processed id list from the -writelist option.

For example, if you enter the following command:

```
sds_import -sds Centera -first 122240 -last 122243 -writelist
```

you would expect four documents to be imported.

However, the resulting **SDSLog.000181** file might look like this:

```
Information for bkg_sds job number 181 started at 02/06/2008 13:52:13
Job parameters:
Import SDS Unit: centera_new (15)
-writelist: yes
Documents from document range (First doc=122240, Last doc=122243).

End of job information:
Number of non fatal errors encountered: 0
Number of documents processed: 5
Job completed at 02/06/2008 14:07:00
```

And the resulting **SDSImportDocs.000181** file would look like this:

```
DOC ID
-----
122240 7DS19L96ITM1Fe7U82R0BI6AQUGG412V6A4K4G0S4AVKEUTHRSR4K
122241 EP5QTTR3HPI5Le36BIH7Q6CPDL1G412V6ALEEH0R75IHL4A7NI5MO
122242 2KPL267SHMQQTTe2IGNTTOF6AJPJG412V7UH25L029UP32HMKOG2A6
122243 5B52BO08FKVO5e9D7LFMAEBK3D8G412V7VG8A80UDEP3N62M2D49J
122243 8B8ICHLHFKHOFe9H4D7IQ60Q6E2G412V7VK75M0UDEP3N62M2D49J
```

Although the last two Doc IDs are the same, the ClipIDs are unique.

The DOC_tool utility

Use **DOC_tool** to see information about your SDS device, such as the device state, and to enable or disable the SDS device.

For full information on using DOC_tool, refer to the *Image Services System Tools Reference Manual*. To download this manual from the IBM support page, see [“Accessing IBM FileNet Documentation” on page 13](#).

SDS States

Use the DOC_tool **SDS** command to view the state of an SDS unit. The following are valid SDS states:

- The SDS User states are **enabled**, **read-only**, **write-only**, or **disabled**.
 - These states can be changed by the user with the DOC_tool DISable and ENable commands, and remain in effect when the IS software is restarted.
- The SDS System states are **enabled**, **disconnected**, **read-only**, or **write-only**.
 - These states are changed automatically as a result of system conditions, and they do not remain in effect when the IS software is restarted.
 - When the System state changes, an SDS_worker issues an RSVP or INFO message. For example, if the SDS_worker does not have write permission to a repository, the SDS device will be set to read-only and an RSVP message will be generated.
 - If an SDS unit is disconnected, it can be reconnected automatically and set to enabled System state.

Note When the SDS device is disconnected, the SDS worker retries the open device command indefinitely. After it has successfully opened the SDS device, the SDS worker sets the SDS unit to an enabled state. At this point, any new requests can now be dispatched to the enabled SDS device.

The User and System states are managed independently of each other.

DOC_Tool commands for managing SDS devices

Use DOC_tool to monitor and manage the SDS units in your IS system. The SDS command with the list “l” option displays a list of SDS units and their SDS unit IDs. DOC_tool uses the SDS unit IDs to identify the SDS unit for various operations.

- The **SDS** commands allow you to gather information about the SDS unit and the SDS worker processes.
- The **ENable** and **DISable** commands allow you to enable and disable the SDS unit.
- The **DEMAND queue** command displays a list of pending read and write requests for the SDS unit you specify.

Note In a multi-storage server environment, an SDS device can be accessed by each storage server in the system. So, when enabling or disabling an SDS device with DOC_tool, you must run DOC_tool on each storage server in the system to completely enable or disable all paths to the device.

These commands are also available for cache-only servers. For more information on DOC_tool, see the *Image Services System Tools*

Reference Manual. To download this manual from the IBM support page, see [“Accessing IBM FileNet Documentation” on page 13](#).

ENable Command

The DOC_tool **ENable** command allows you to enable SDS units that have been set to DISABLED by DOC_tool DISable. This command is only available in the allowupdate DOC_tool mode.

Note If an SDS unit has been disabled by a network problem or hardware failure, it cannot be enabled with this command.

Depending on the type of storage library selected, the ENable command displays the following screens.

MSAR Library

```
drive, Platter(surface), SDS? ('d', 'p', 'SDS')
```

Optical Library

```
Slot, drive, OpticalLibrary, Platter(surface), SDS?  
('s', 'd', 'o', 'p', 'SDS')
```

Cache-only Servers

When a server is a Cache-only server, there can be no MSAR or optical library selected.

- If the SDS unit is in a DISABLED state, the following prompt displays after you enter the ENable command:

```
SDS unit? <y=yes>  
SDS id: 2
```

DISable Command

The DOC_tool **DISable** command allows you to disable SDS units that have been set to ENABLED. Like the ENABLE command, this command is only available in the allowupdates mode.

The screen prompts displayed are identical to those displayed for the ENABLE command.

```
<DOC_tool:lib a>allowupdates
password? *****
<DOC_tool:lib a>dis
Drive, Platter(surface), SDS? ('d', 'p', 'SDS'): sds
SDS id: 2
Disable reads and writes [no]:? (y/n):
Disable read-only (no => disable just writes)[no]:? (y/n):
DISABLED unit =2 for writes only (old state=USER ENABLED)
```

You can disable the SDS unit for writes only, reads only, or both. (The ENABLE command fully enables the device.) The system state settings are not affected by this command.

DEMAND Queue Command

The DOC_tool **DEMAND queue** command displays a list of pending read and write requests for the SDS unit you specify. In the following example, there are no pending read or copy requests, but there are four pending write requests:

```
<DOC_tool:lib a>dem
SDS unit? <y=SDS/n=surf> [y]:
SDS unit ID: 2

Read Q:      : count = 0

WRITE Q:     : count = 4
  Request Type: WRITE_REQUEST
    doc_id=101592, doc_index=101592
    numpages=1, fbc=FALSE, fam_id=3
  Request Type: WRITE_REQUEST
    doc_id=101593, doc_index=101593
    numpages=1, fbc=FALSE, fam_id=3
  Request Type: WRITE_REQUEST
    doc_id=101594, doc_index=101594
    numpages=1, fbc=FALSE, fam_id=3
  Request Type: WRITE_REQUEST
    doc_id=101595, doc_index=101595
    numpages=1, fbc=FALSE, fam_id=3

COPY Q:      : count = 0
```

SDS Command

The DOC_tool SDS command allows you to view the current state of the SDS unit and its associated SDS workers, as well as other performance information.

SDS command prompt and options

The **SDS** command has several display options:

- s** – Summary report
- d** – Detailed report
- w** – Worker report
- a** – All information report
- f** – Find object report
- l** – List report

Important

Several of these options, such as the Detailed report, the Worker report, and the All information report, use a large amount of system processing resources to gather the requested information.

If you plan to run any of these options repeatedly in a background script, keep in mind that doing so could slow system performance.

SDS List option (l)

The list option provides a list of currently configured SDS devices and their SDS ID numbers:

Sample output - List option

```
<DOC_tool:lib a>sds
Summary info, Detailed, Worker info, All info, Find object, or list?
('s', 'd', 'w', 'a', 'f', 'l') : l

The current time is Tue Oct 30 12:37:13 2007

SDS info: list  sds_id      Name
                2         NetApps2
                4         SunSAR
                5         DR550
                6         Centera
```

SDS Summary option (s)

The summary option provides a brief display of SDS device characteristics. This option prompts you to select all SDS units or one specific SDS unit.

Sample output - Summary option

```
<DOC_tool:lib a>sds
Summary info, Detailed, Worker info, All info, Find object, or list?
('s', 'd', 'w', 'a', 'f', 'l') : s

The current time is Tue Oct 30 12:37:20 2007

SDS info: SUMMARY option
All SDS units mode (y/n) [y]: : n
SDS unit ID: 2
***** SDS unit = NetApps2 (2)
      SYSTEM state = SYSTEM ENABLED (0x0)
      USER state = USER ENABLED (0x0)
      Worker = 'SDS_worker' Number Instances = 3
      info = '/SDS_snaplock/SDS_test2'
      SDS priority = high
      DEBUG Setting = MAX
      dynamic repository lib = 'SDSw_snaplock'
      retention default offset (10 days)
      SDS content delete setting=YES
```

SDS Detailed info option (d)

When you select the **Detailed** option, the report displays detailed information about the performance of the SDS device in several sections:

- SDS Configuration and state
- Total Accumulated counters
- Read Request Performance
- Write/Copy Request Performance
- Cumulative time spend in CSM cache operations and SDS device library calls.
- Average queue wait time. How long request are queued before an SDS worker starts to process it.

Sample output - Detailed option

```
<DOC_tool:lib a>sds
Summary info, Detailed, Worker info, All info, Find object, or list?
('s', 'd', 'w', 'a', 'f', 'l') : d

The current time is Tue Oct 30 12:37:33 2007

SDS info: Detailed option
All SDS units mode (y/n) [y]: : n
SDS unit ID: 2

***** SDS unit = NetApps2 (2)
        SYSTEM state = SYSTEM ENABLED (0x0)
        USER state = USER ENABLED (0x0)
        Worker = 'SDS_worker' Number Instances = 3
        info = '/SDS_snaplock/SDS_test2'
        SDS priority = high
        DEBUG Setting = MAX
        dynamic repository lib = 'SDSw_snaplock'
        retention default offset (10 days)
        SDS content delete setting=YES

(continued on next page)
```


Sample output - Detailed option, continued

```
Total Accumulated counters from all workers(3):
** Configured workers =3 active workers=2
TOTAL WORKER COUNTERS (sds_id=2):
  Read Requests processed: 1
  Write Requests processed: 64
  Copy Requests processed: 0
  Errors: 0
  Requests processed = 65
  Successful requests processed = 65
  Errors = 0
AVERAGE ACCUMULATED ELAPSE TIMES:
  Up time:      1670.561958 secs/workers  (27.842699 mins)
  Idle time:   1584.622689 secs/workers  (26.410378 mins) (94.86%)
  Total processing time: 85.914369 secs/workers
                        (1.321760 secs/reqs)
                        (0.167149 secs/image page)
                        (0.012215 secs/KB)
  AVG requests queue wait time: 34.442542 secs/reqs

*****Total READ REQUEST PERFORMANCE (sds_id=2)
  Total retrieval requests = 1
  Images retrieved from SDS = 1
  Data retrieved = 0.013184MB
  AVG Image Size = 13.500000K
  Number of read requests where the whole blob fits
    into the internal image_buffer (256K): 1
  Number of read requests where the whole blob does not fits
    into the internal image_buffer (256K): 0
  Cache hits: 0
  Number of redirection: 0
  Number of redirection errors: 0

(continued on next page)
```

Sample output - Detailed option, continued

```

Total Time to process read requests: 0.079858 secs (0.001331 mins)
Time in SDS device open and read: 0.001743 secs (0.000029 mins)
                                   (0.001743 secs/reqs)
                                   (0.001743 secs/image page)
                                   (0.000129 secs/KB)
Time in SDS device read only: 0.000079 secs (0.000001 mins)
                                   (0.000079 secs/reqs)
                                   (0.000079 secs/image page)
                                   (0.000006 secs/KB)
Time in cache(CSM) to process read: 0.077418 secs (0.001290 mins)
                                   (0.077418 secs/reqs)
                                   (0.077418 secs/page)
                                   (0.005735 secs/KB)

*****Total WRITE/COPY REQUEST PERFORMANCE (sds_id=2)
Total write requests = 64
Total copy requests = 0
Documents written = 113 (FBC=50, MSAR reads=50, Cache=13, Copy=0)
Images written = 513
Data written = 6.855503MB
AVG Image Size = 13.684279K
Cache hits in copy: 0
Total Time to process write requests:      171.748132 secs (2.862469
mins)
Total Time to process copy requests:      0.000000 secs (0.000000
mins)
Time in SDS device create and write object: 168.531598 secs (2.808860
mins)
                                   (2.633306 secs/reqs)
                                   (0.328522 secs/image page)
                                   (0.024007 secs/KB)

(continued on next page)

```

Sample output - Detailed option, continued

```

Time in SDS device write only: 168.287921 secs (2.804799 mins)
                               (2.629499 secs/reqs)
                               (0.328047 secs/image page)
                               (0.023973 secs/KB)
Time in cache(CSM) to process write/copy: 0.480883 secs (0.008015
mins)
                                               (0.007514 secs/reqs)
                                               (0.000937 secs/page)
                                               (0.000069 secs/KB)
Time in MSAR read to process write/copy: 0.068860 secs (0.001148
mins)

```

SDS Configuration and state

The following table describes the **SDS Configuration Settings** section of the summary:

SDS Configuration Settings and state

Configuration Settings	Description
SDS unit	The name of the SDS unit, such as NetApps2, which is assigned during configuration in fn_edit.
SYSTEM state	Shows the current SDS unit state: Enabled, Read-only, Write-only, or Disconnected.
USER state	Shows the current SDS user state: Enabled, Read-only, Write-only, or Disabled.
Worker name and number of instances	The number of SDS_worker instances can be from 1 to 99.
info	This field contains the repository connection string.

SDS Configuration Settings and state, continued

Configuration Settings	Description
SDS priority	Can be set to High, Medium or Low. This setting determines if read request will go to SDS first, or to MSAR or optical.
DEBUG setting	Settings are MAX, MIN, and OFF. If set to MAX, more information will be logged as SDS requests are processed.
dynamic repository lib	The associated SDS connector library name, such as SDSw_snaplock.
retention setting	Retention default offset <number of days> - or - Retention default date <future date>
SDS content deletion setting	Yes, content will be deleted from the SDS device when the corresponding entries in the DOCS table are deleted from the IS system. - or - No, content will remain intact in the SDS device when the corresponding entries in the DOCS table are deleted from the IS system.

Total Accumulated counters

Note Counters start incrementing when Image Services starts and cannot be reset unless the IS software is recycled.

The following tables describe the **Total Accumulated Counter** sections of the summary:

Total Accumulated Counters Information

Counter	Description
Read/Write/Copy Requests processed	Total number of read/write/copy requests that have been processed by all of SDS_worker processes serving a single SDS unit.
Successful requests processed	Total number of requests that have been successfully processed.
Errors	Total number of errors encountered. Only errors that were not corrected after retries and did not result in termination of SDS worker processes. This class of errors can be characterized as errors that resulted in an error completion status being sent back to the client with the ds_notify queue.

Average Accumulated Elapse Times

Counter	Description
Up time	Accumulated average up time of all SDS workers (seconds/workers). This total accumulated time of all SDS workers divided by the number of workers. The elapsed time since the last request was done for each is used to calculate the up time.
Idle time	Accumulated average idle time (seconds/workers). This is time that workers are idle waiting for request. The elapsed time of the last request is added to the idle time.

Average Accumulated Elapse Times, continued

Counter	Description
Total processing time	<p>Total elapsed processing time per worker. The time between the dequeue of an SDS request and the completion of the request process.</p> <p>The average time elapsed per request is also displayed (seconds/requests). The average time elapsed per image page (seconds/image page) and elapsed per K-byte (seconds/K-byte) is also displayed.</p>
AVG requests queue wait time	<p>The average wait time in seconds per request. The wait time is the elapsed time between the time a request is enqueued by the dispatch process and the time the SDS_worker dequeues the request. This means that a submission time is calculated and stored in SDS QMA structure when the request is enqueued in the SDS queue.</p> <p>After the SDS_worker dequeues the request, it will get the current time and subtract the submission time to calculate the elapsed or wait time. The wait time is accumulated and an average is displayed here.</p>

Read Request Performance

The following table describes the **Read Request Performance** section of the summary:

Total Read Request Performance

Accumulated Counter	Description
Total retrieval requests	Total number of retrieval requests submitted.
Images retrieved	Total number of images retrieved.
Data retrieved	Total amount of data transferred in K-bytes to CSM cache from the SDS repository. This does not include document pages that were requested and might have been read into the internal buffers and subsequently the page objects were already in cache.
AVG Image Size	Average size of an image page transferred.

Total Read Request Performance, continued

Accumulated Counter	Description
Number of read requests where the whole BLOB fits into the internal image_buffer	Number of requests where the whole document BLOB fits into the internal buffer. The size of the configurable internal buffer is displayed.
Number of read requests where the requested images do not fit into the internal image_buffer	Number of requests where the whole document BLOB does not fit into the internal buffer. The size of the configurable internal buffer is displayed.
Cache hits	Cache hits encountered by the SDS_worker process.
Number of redirections	This counter indicates the number of read requests that have been re-directed to MSAR or optical. These re-directions might have occurred due to errors or because the SDS unit has been disabled with DOC_tool. This counter includes successful and unsuccessful re-directions. The number of successful re-directions can be determined by subtracting the number of redirection errors.
Number of redirection errors	Indicates how many redirections have failed.
Total Time to process read request	Total elapsed time to process all read requests. This is the total time between the dequeue of an SDS request and the completion of the read request.
Time in SDS device open and read	Indicates the time spent in opening and reading calls to the SDS entry point only.
Time in SDS device read only	Indicates the time spent reading from the SDS entry point only.
Time in cache(CSM) to process read	Indicates the time spent processing the read request in page cache.

Write/Copy Request Performance

The following table describes the **Write/Copy Request Performance** section of the summary:

Total Write/Copy Request Performance

Accumulated Counter	Description
Total write requests	Total number of write requests processed.
Total copy requests	Total number of copy requests processed.
Documents written	Total number of documents written. A single document can be composed of several pages, which are stored as individual images.
Images written	Total number of images written. This is the number of document pages, not documents.
Data written	Total amount of data transferred in K-bytes from CSM cache to the SDS repository.
AVG Image Size	Average size of an image page transferred.
Cache hits in copy	Indicates the number of cache hits while processing copy requests.
Total Time to process write requests	Total elapsed time for processing write requests. This is the total time between dequeue of an SDS request and the completion of the write request.
Total Time to process copy requests	Total elapsed time for processing copy requests. This is the total time between dequeue of an SDS request and the completion of the copy request.
Time in SDS device create and write object	Indicates the time spent processing create and write calls to the SDS entry point only.
Time in SDS device write only	Indicates the time spent writing to the SDS entry point only.
Time in cache(CSM) to process write/copy	Indicates the time spent processing the write/copy request in page cache.
Time in MSAR read to process write/copy	Indicates the time spent reading from MSAR to process the write/copy request.

SDS Worker option (w)

When you select the **Worker** option, basic SDS information is displayed, along with information about each active worker and a detail worker report. The information displayed is similar to worker detail reports, but the counters and time accumulated are specific to each worker. The fields displayed in the detail report and Worker report are the same except there is an additional field, the “Last request done” field, which is a time stamp of the last completed request.

Sample output - Worker Option

```
<DOC_tool:lib a>sds
Summary info, Detailed, Worker info, All info, Find object, or list?
('s', 'd', 'w', 'a', 'f', 'l') : w

SDS info: WORKER option

***** SDS unit = 1 state = ENABLED SDS priority = medium *****
  DEBUG Setting = MAX
  Worker = 'SDS_worker' Number Instances = 3
  info = 'WorkingDirectory=/NLS/tivoli domain=10.10.8.35'
  Current number of outstanding requests in the SDS queue: 0

Worker instance 1:
WORKER COUNTERS:
  Requests processed: 3
  Requests satisfied via SDS device (IBM DR550): 2
  Successful requests processed: 2
  Errors: 1 (last error = ca64000a)
  Retrieved image pages count: 9
  Retrieved data: 0.719210MB

(continued on next page)
```

Sample output - Worker Option, continued

```
AVG Image Size = 81.830078K
AVG request size = 245.490234K
Number of SDS device read (API calls): 11
Number of SDS device query calls: 0
Number of SDS device requests where the whole document fits
  into the internal image_buffer (128K): 0
Number of SDS device requests where the requested images fits
  into the internal image_buffer (128K): 1
Number of SDS device requests where the requested images do not
  fit into the internal image_buffer (128K): 1
Cache hits: 0
```

WORKER ACCUMULATED TIMES:

```
Up time =                76481.966902 secs (1274.699448 mins)
Total Idle time =       76477.196948 secs (1274.619949 mins) (99.99%)
Last request done =     76355.779166 secs (1272.596319 mins)
Total processing time:  4.752298 secs
                        (1.584099 secs/reqs)
                        (0.528033 secs/image page)
                        (0.006453 secs/KB)
Time in get_doc_from_filesystem():  0.000790 secs
Time in get_doc_from_device():       4.076062 secs
Time in SDS device open and read (API calls):  2.991772 secs
                                                (0.997257 secs/reqs)
                                                (0.332419 secs/image page)
                                                (0.004062 secs/KB)
Time in SDS device read only (API calls): 2.133966 secs
                                                (0.711322 secs/reqs)
                                                (0.237107 secs/image page)
                                                (0.002898 secs/KB)
Time in SDS device querying database: 0.000000 secs
```

(continued on next page)

Sample output - Worker Option, continued

```
Time in writing to page cache(CSM): 0.438663 secs
                                   (0.146221 secs/reqs)
                                   (0.048740 secs/page)
                                   (0.000596 secs/KB)
Total SDS requests queue wait: 0.741531 secs
                                   (0.247177 secs/reqs)
                                   (0.082392 secs/page)
                                   (0.001007 secs/KB)
```

```
.
.
.
```

SDS All option (a)

When you select the **All info** option, all of the previous information mentioned above is displayed.

Sample output - All Option

```
<DOC_tool:lib a>sds
Summary info, Detailed, Worker info, All info, Find object, or list?
('s', 'd', 'w', 'a', 'f', 'l') : a

SDS info: ALL option

***** SDS unit = 1 state = ENABLED   SDS priority = medium *****
      DEBUG Setting = MAX
      Worker = 'SDS_ISAR_reader' Number Instances = 1
      info = 'workingdirectory=/NLS/tivoli domain=10.10.8.35'
      Current number of outstanding requests in the SDS queue: 0

SDS (unit=1) queue is currently empty

Worker instance 1:
WORKER COUNTERS:
  Requests processed: 3
  Requests satisfied via SDS device (IBM DR550): 2
  Successful requests processed: 2
  Errors: 1 (last error = ca64000a)
  Retrieved image pages count: 9
  Retrieved data: 0.719210MB
  AVG Image Size = 81.830078K
  AVG request size = 245.490234K
  Number of SDS device read (API calls): 11
  Number of SDS device query calls: 0
  Number of SDS device requests where the whole document fits
  into the internal image_buffer (128K): 0

(continued on next page)
```

Sample output - All Option, continued

```
Number of SDS device requests where the requested images fits
into the internal image_buffer (128K): 1
Number of SDS device requests where the requested images do not
fit into the internal image_buffer (128K): 1
Cache hits: 0
WORKER ACCUMULATED TIMES:
Up time =                76594.523156 secs (1276.575386 mins)
Total Idle time =       76589.753202 secs (1276.495887 mins) (99.99%)
Last request done =    76468.335420 secs (1274.472257 mins)
Total processing time: 4.752298 secs
                        (1.584099 secs/reqs)
                        (0.528033 secs/image page)
                        (0.006453 secs/KB)
Time in get_doc_from_filesystem(): 0.000790 secs
Time in get_doc_from_device():    4.076062 secs
Time in SDS device open and read (API calls): 2.991772 secs
                                           (0.997257 secs/reqs)
                                           (0.332419 secs/image page)
                                           (0.004062 secs/KB)
Time in SDS device read only (API calls): 2.133966 secs
                                           (0.711322 secs/reqs)
                                           (0.237107 secs/image page)
                                           (0.002898 secs/KB)
Time in SDS device querying database: 0.000000 secs
Time in writing to page cache(CSM): 0.438663 secs
                                           (0.146221 secs/reqs)
                                           (0.048740 secs/page)
                                           (0.000596 secs/KB)
Total SDS requests queue wait: 0.741531 secs
                                (0.247177 secs/reqs)
                                (0.082392 secs/page)
                                (0.001007 secs/KB)
(continued on next page)
```

Sample output - All Option, continued

```
Total Accumulated counters from all workers(1):
TOTAL WORKER COUNTERS:
  Requests processed = 3
  Successful requests processed = 2
  Requests satisfied via SDS device (IBM DR550): 2
  Errors = 1
  last error = ca64000a
.
.
.
```

SDS Find object option (f)

The **Find object** option enables you to display characteristics of a specific document in a specific SDS unit.

This option allows you to see the contents of the descriptor stored in the SDS object header and the retention setting of the SDS object. The tool reads the information directly from the record, not from the MKF permanent database.

Note The entry for SDS OBJ ID is blank for file system storage devices such as SnapLock and StorageTek. For Centera devices, the SDS OBJ ID contains the ClipID.

For foreign documents, you must specify the original document ID.

:

Sample output - Find object option

```

<DOC_tool:lib a>sds
Summary info, Detailed, Worker info, All info, Find object, or list?
('s', 'd', 'w', 'a', 'f', 'l') : f

The current time is Tue Oct 30 12:38:24 2007

SDS ID (2-255) : 2
Doc ID : 101589

SDS OBJ ID :
Doc ID=101589 found in SDS=NetApps2 (2)
Dump SDS object header? <y=yes/n=no>[y]:
SDS header info:
version.....1
sds_id.....2
local_doc_id.....101589
local_ssn.....101410035
obj_size.....14396

Document descriptor:
Doc hdr length: 456
check_word....0xad0cde5c  format.....2
num_pages.....1  total_sectors.....0
doc_id.....101589  system_id.....101410035
doc_location.....0  alt_surf_ssn.....0
annot_base_doc.....0  alt_doc_loc.....0
annot_page_num.....0  alt_surf_id.....0
doc_type.....0x00  key_info_len.....310
class_name_len.....18  sec_rd_nm_len.....9
sec_wrt_nm_len.....9  sec_exe_nm_len.....9
contig_sectors.....0  back_contig.....0

(continued on next page)

```

Sample output - Find object option, continued

```
(continued from previous page)

checksums.....1  character_set.....11
deleted.....0
cluster_id.....0x0000000000000000
class_name.....'Sds_snaplock_only'
sec_rd_name.....'(ANYONE) '
sec_wrt_name.....'(ANYONE) '
sec_exe_name.....'(ANYONE) '

Key type=2(str)  Key name='F_DOCFORMAT' len=11
                  Key value='application/excel;name="Simple excel.xls"'
len=41
Key type=2(str)  Key name='F_DOCLOCATION' len=13
                  Key value='nnnnnnnnnn' len=10
Pg...1  Sector#.....0  Byte_offset..572  #Bytes...13824

Pg...0  Checksum..0x081badff  Pg...1  Checksum..0xb305e8c7

Get SDS retention settings? <y=yes/n=no>[y]:
event_base=FALSE held=FALSE retention=11/09/2008 11:34:09
```

SDS Delete option (DELETE)

(Requires ALLOWUPDATES.)

The **DELETE** option enables you to delete the content of a specific document from a specific SDS unit.

This option is only available in allowupdates mode.

Sample output - DELETE option

```
<DOC_tool:lib a>allowupdates
password? *****
<DOC_tool:lib a>sds
Summary info, Detailed, Worker info, All info, Find object, List, DELETE?
('s', 'd', 'w', 'a', 'f', 'l', 'DELETE') : DELETE
SDS ID (2-255) : 10
Doc ID : 3945937
SDS OBJ ID :
Doc ID=3945937 found in SDS=QA_hcap (10)
Deleted Doc ID=3945937 SDS=QA_hcap (10) sds_obj='' (length=0)
<DOC_tool:lib a>
```

Appendix A – Log Files and Messages

This appendix contains information about errors you might encounter with integral Single Document Storage.

Error Logs

When an Image Services related error is encountered, the FileNet error tuple and error text, if available, are reported. If a utility or file system error is encountered, a utility error code is reported. The utility error codes are generally related to an incorrect configuration, bad input file format, or a problem with the SDS file system. File system errors (Code 1024) are almost always the responsibility of the System Administrator to correct.

The following are examples of error types and locations related to all utilities.

IS error logs – SDS errors and all IS related messages and errors (Storage Library, SDS_FS_reader, database, security, and so on) are logged in the IS elog. Use the **vl** command to view error logs.


Default UNIX - /fnsw/local/logs/elogs

Default Windows - \fnsw_loc\logs\elogs

To view more information about an IS or SDS error, you can use the **fn_msg tool** directly on the error tuple. See [“The fn_msg Tool”](#) below.

SDS Error Message Tuples

Error messages include a three-part numerical identifier called a tuple (for example, 238,3,1). SDS error tuples (generated by the SDS abstract or SDS_reader programs) use the same format as other IS error tuples. The following diagram identifies each part of an **IS** error tuple:



For example, if the first number of the error tuple is **NOT** zero, then the error tuple is either an SDS error or an IS error.

Note An SDS or IS error **never** has a 0 as the first number of the identifier.

The fn_msg Tool

Use the **fn_msg** tool to retrieve additional information for SDS or IS error tuples. Enter the following at the command line:

```
fn_msg <tuple>
```

where <tuple> is the three-part number preceding the message text.

For example, if you receive this message:

```
238,3,1 Optical disk drive not ready.
```

Enter this command to display additional information:

fn_msg 238,3,1

The following text displays:

```
<SDS,3,1> SDSw Repository: connection error. Typically caused by repository down or network down. The SDS is read and write disabled. SDS_worker will redirect read requests and attempt to re-connect. INFO Message will be generated in SLC window
```

For example, the hexadecimal equivalent of the error tuple <238,3,1> is 0xEE003001. In this case, running any of the following commands produces the same text:

```
fn_msg 238,3,1  
fn_msg 0xee003001  
fn_msg ee003001
```

Note The **fn_msg** tool accepts hexadecimal values in either upper or lower case hexadecimal numbers and does not require leading **0x**'s.

SDS (238) Messages

The following is a general list of SDS errors that you might encounter. To view more information about these errors, use the **fn_msg** tool directly on the error tuple. See [“The fn_msg Tool” on page 107](#)

SDS general messages

238,0,1 SDS informational message.
(SDS_err_info)

- 238,0,2** Invalid session handle. This is a programmatic problem.
(SDS_err_invalid_sess)
- 238,0,3** A particular document is not stored in the SDS.
(SDS_err_no_document)
- 238,0,4** An unsupported request type was encountered by a SDS worker.
(SDS_err_not_type_supported)
- 238,0,5** The SDS unit is disabled.
(SDS_err_disabled)
- 238,0,6** A particular SDS is not configured.
(SDS_err_not_configured)
- 238,0,7** A particular SDS unit is not valid.
(SDS_err_invalid_sds_unit)
- 238,0,8** The checksum stored with the SDS document page is not consistent with the retrieved document. This may indicate data corruption.
(SDS_err_chksum)
- 238,0,9** A critical error was encountered by a SDS worker and as result the worker aborted. See sys_log for the reason.
(SDS_err_worker_abort)
- 238,0,10** The Doc_id or ssn stored in the document header and read request does not match. See sys_log.
(SDS_err_header_doc_or_ssn_mismatch)
- 238,0,11** Invalid worker instance number. See sys_log.
(SDS_err_invalid_instance)
- 238,0,12** The version number of SDS Worker Info structure does not match. This would indicate that there is mixture of old and new modules (SDS, worker, DOC_tool, etc). See sys_log.
(SDS_err_info_struct_ver_mismatch)

238,0,13 Failed to access file or directory.
(SDS_err_file_access)

238,0,14 A particular SDS is no longer configured.
(SDS_err_no_longer_configured)

238,0,15 For old NLS SDS ID=1 only reads are supported.
(SDS_err_old_sds_id_invalid_operation)

238,0,16 A particular SDS is no longer configured in Document Family.
(SDS_err_no_longer_configured_in_fdt)

238,0,17 SDS ID may not be configured for tranlog document family. It must be associated with primary document family.
(SDS_238_tranlog)

238,0,18 Unexpected bytes read from CSM or MSAR surface file.
(SDS_badreadlen)

238,0,19 Cannot set the persistent & non-persistent state at the same time.
(SDS_err_invalid_state)

238,0,20 SDSw pluggable library load error.
(SDS_load_SDSw_error)

238,0,21 SDS unit is write disabled.
(SDS_err_write_disabled)

238,0,22 SDS unit is read disabled.
(SDS_err_read_disabled)

238,0,23 Object is not found in SDS unit.
(SDS_err_obj_not_found)

SDSw Repository messages

238,3,1 SDSw Repository: connection error. Typically caused by repository down or network down. The SDS is read and write disabled. An SDS_worker will redirect read requests and attempt to re-connect. INFO Message will be generated in SLC window.

(SDS_err_repository_connect)

For example, if the IS system cannot connect to the SDS unit, you might see a message similar to this:

```
2008/02/04 07:46:28.177 238,3,1 <fnsw> SDS_worker 7 74 (8876) ... [SERIOUS]
Error from SDSw_open_repository repository_id=10.10.8.35
err_msg='Can't Init TSM/ Client Node=FNP8IM_15197
      ANS1075E (RC102) *(*)The operating system refused a TSM request for memory
allocation.'
```

In this case, you might have too many SDS_workers. Use `fn_edit` to lower the number of SDS_workers for the SDS device.

238,3,2 SDSw Repository: Resource error. The SDS is write disabled. Generate RSVP. Respond to RSVP will enable SDS for writes.

(SDS_err_repository_no_resource)

238,3,3 SDSw Repository: permission error. The SDS is read and write disabled.

(SDS_err_repository_permission_err)

238,3,4 SDSw Repository: write permission error. The SDS is write disabled. Read can still processed.

(SDS_err_repository_write_permission_err)

238,3,5 SDSw Repository: read permission error. The SDS is read disabled. Write can still processed.

(SDS_err_repository_read_permission_err)

238,3,6 SDSw Repository: No object could be located from the repository.

(SDS_err_no_object)

238,3,7 SDSw Repository: Corrupt object.
(SDS_err_corrupt_object)

238,3,8 SDSw Repository: Retryable error condition. Close and open the connection and start over. This is useful for transactional repositories. DR550.
(SDS_err_retryable_err)

238,3,9 SDSw Repository: Read Single object failure. Not retryable.
(SDS_err_nonretryable_err)

238,3,10 SDSw Repository: Write Single object failure. Not retryable.
(SDS_err_nonretryable_write_err)

238,3,11 SDSw Repository: Could not allocate memory.
(SDS_err_no_memory_err)

238,3,12 SDSw Repository: Get repository info error.
(SDS_err_get_info_err)

238,3,13 SDSw Repository: Set retention error.
(SDS_err_set_retention_err)

238,3,14 SDSw Repository: Delete content error.
(SDS_err_delete)

238,3,15 SDSw Repository: Unexpected error see elog for more info.
(SDS_err_unexpected_system_err)

238,3,16 SDSw Repository: Duplicate object. That means the document object was already created. This may occur when migration to SDSw is interrupted after write out to the SDSw was complete but the DOCS table was not updated.
(SDS_err_duplicate_obj)

238,3,17 SDSw Repository: Connected.
(SDS_connected)

238,3,18 SDSw Repository: Target repository software version is not supported.

(SDS_err_target_repository_version_not_support)

238,3,19 SDSw Repository: Invalid Parameter where SDSw client and SDSw libraries do not match up.

(SDS_err_invalid_parameter)

238,3,20 SDSw Repository: Pending delete.

(SDS_err_pending_delete)

Background SDS migration job messages

238,4,1 Failed to write to background task log file.

(4_logwrt_err)

238,4,2 SDS background job logic error

(4_logic_err)

238,4,3 Background job terminated due to operator abort command.

(4_operabort_err)

238,4,4 Archive index not found in the database.

(SDS_archive_index_not_found)

238,4,5 This document has already had its ARCHIVE and or CLIP ID migrated from the index database to the docs table. If you want to migrate the entry again, run sds_convert_nls and specify the -override option.

(SDS_already_converted)

238,4,6 SDS background job syntax error.

(4_syntax_err)

SDS Fastbatch committal messages

238,5,1 Invalid offsets of cache and od info in client attrs of batch
(SDS_err_badbatchoffset)

238,5,2 Invalid number of pages in batch
(SDS_err_badbatchpages)

238,5,3 Unexpected number of bytes read
(SDS_err_badreadlen)

238,5,4 Page zero not found
(SDS_err_pagezeronotfound)

Appendix B – EMC Centera

Support policy

Image Services 4.1.1 and higher includes the EMC Centera Software Developers Kit (SDK). The CentraStar version must be 3.1.1 or higher. Support is limited to the functional scope of IS integration with Centera, as described in the IS documentation. New features appearing in a release of CentraStar are not automatically supported.

Note According to EMC, all current versions of CentraStar are backward compatible with all versions of the Centera SDK, provided that the Centera device is configured to be compatible with the environment that the Centera SDK supports. Also, specific versions of the SDK are backward compatible with CentraStar from a specific release up. Refer to the Centera documentation for SDK/CentraStar compatibility.

Centera is not supported with HP Integrity servers.

Since Centera SDK/CentraStar version 3.1, Centera supports both Event-Based Retention (EBR) and Chronological retention. With EBR, client applications can store a document with an "indefinite" retention, waiting for a business event to define and activate it a later time, as would be required for example for protecting a document from deletion for three years after close date. Also Centera provides the capability to place "retention holds" on documents subject to legal inquiries. The C-Clips associated with these documents cannot be deleted until they are released from retention hold. These capabilities require EMC's Advance Retention Management License for Governance and Compliance Plus editions.

Note To take advantage of Event-Based Retention in Centera Governance and Compliance Plus editions, customers should make sure that they are appropriately licensed for it.

Performance tuning suggestions for Centera

As you configure the Centera device with your Image Services (IS) system, keep the following suggestions and best practices in mind.

- If the Centera device is set to 1000BaseT FD (full duplex), ensure that your Layer 2 switch settings match exactly.
- Ensure your Ethernet 1000BaseT cables are Category 6 and nothing less (such as Category 5 or 5e).
- The IS server and the Centera device should always be on the same subnet / VLAN.
- Consider using embedded BLObs for all object sizes 100 KB or less.
- Verify that the Centera device is set to Storage Strategy Performance = FULL.
- Ensure that all IP addresses or fully qualified DNS names are properly coded on the IS server .pea profile file or input string where the connection IP addresses to the Centera device are coded.
- Ensure that the IS index database is on the faster SAN attached DMX storage and that your placement of this database is highly optimized for performance.
- When you use the Image Services System Configuration Editor (fn_edit) in Chapter 2 to configure the Centera device with IS, ensure that the SDS retrieval priority is set to HIGH.

- Also in `fn_edit`, verify that the Document Buffer Size is equal to the BLOB size wherever possible. When the whole BLOB fits into the internal buffer, all the pages of the document starting with the first requested page will be written to page cache.
- For best performance under heavy workloads, any IS application that writes to, or reads from, a storage device, such as Centera, must use multiple process threads for any I/O operation.

Configuration steps

In addition to the general configuration settings discussed in the main body of this document, some specific configuration settings apply to EMC Centera. These settings are discussed in the following sections.

Note To improve SDS import and recovery performance, each Centera SDS unit should have its own Centera Pool.

Set up Centera Profile and Pool Entry Authorization (.pea) file

Create a Pool Entry Authorization file to activate the security for accessing the Centera device. Create one .pea file for each Centera device. (Refer to Centera documentation for more information about creating profile and .pea files.)

Create the home pool

- 1 Log on to the Centera device and launch the Centera interactive shell.
- 2 Enter the following command at the `Config#` prompt:

```
Config# create pool IMPool
```

where `IMPool` is the name of the pool file you want to create.

- 3 When you are prompted for the Pool Mask, enter:

rdqeDcwh

- 4 Accept the Pool Quota of 1 GB.
- 5 Centera asks if you want to replicate this pool and if you want to issue the command now. Enter **yes** to both prompts.

The sequence of commands looks similar to this:

```
Config# create pool IMPool
Pool Mask [rdqeDcw-]: rdqeDcwh
Pool Quota: [1 GB]:
Do you want to replicate this pool? (yes, no) [yes]:
Issue the command?
  (yes, no) [no]: yes
Created pool IMPool with ID 6086fa3a-1dd2-11b2-bdd1-867da7cfd572-9.
```

Create the pool entry authorization profile

- 1 At the Config# prompt, enter:

Config# **create profile IMDevel**

where IMDevel is the name of the .pea profile file you want to create for this Centera device.

- 2 Accept the default values for all of the prompts, except these:

Enable Profile? **yes**

Home Pool: **IMPool** (or whatever name you chose)

Granted Rights: **rdqeDcwh**

Profile-Metadata Capability? **yes**

Compliance Role? **yes**
 Issue the command? **yes**

The sequence of commands looks similar to this:

```
Config# create profile IMDevel
Profile Secret [generate]:
Enable Profile? (yes, no) [no]: yes
Profile Type (access, cluster) [access]:
Home Pool [default]: IMPool
Granted Rights for the Profile in the Home Pool [rdqeDcw-]: rdqeDcwh
Profile-Metadata Capability? (yes, no) [no] yes
Accesscontrol role? (yes, no) [no]:
Audit role? (yes, no) [no]:
Compliance role? (yes, no) [no]: yes
Configuration role? (yes, no) [no]:
Monitor role? (yes, no) [no]:
(continued on next page)
```

```
(continued from previous page)
Replication role? (yes, no) [no]:
Issue the command?
  (yes, no) [no]: yes
Establish a Pool Entry Authorization for application use? (yes, no) [no]: yes
Please enter Pool Authorization creation information: C:\IMDevel.pea

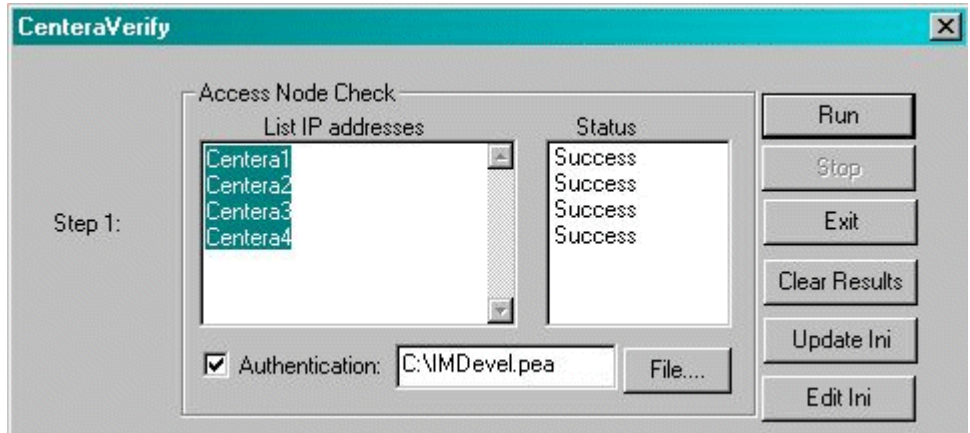
Config#
```

- 3 Repeat the steps in this section on each Centera unit to create a separate profile for each unit. (The profiles can be shared, but sharing makes sds_import slower.)

Note If you answer **yes** to the Replication role prompt, you must create a merged profile. See the Centera documentation for details.

Verify the .pea profile file

On the Centera device, run the CenteraVerify program to make sure the .pea profile file has been configured successfully.



The first step in the CenteraVerify program verifies that it can access each Centera device successfully.

Step 2:

Cluster Information			
Cluster ID:	6086fa3a-1dd2-11b2-bdd1-867da7cfd572		
Available Capacity:	21,416	of	31,097 GB
Version Information:	Centra Star	SDK	Centra Verify
	3.1.2-1835-831-20	3.1.473	3.1.22
	Write:	Read:	
Capabilities:	Allowed	Allowed	
	Delete:	Purge:	
	Allowed	Not allowed	
	Query:	Privileged delete	
	Allowed	Allowed	
	Exist:	Monitor:	
Allowed	Not allowed		
Replica Address:	9.39.55.104:3218,9.39.55.105:3		Test...

The second step verifies Cluster Information and displays the capabilities that are allowed or not allowed in the .pea profile file.

The third step verifies the read, write, and delete capabilities.

Step 3:

Write Read Test		
	Progress	Status
Write:	3 files have been written	Success
Read:	3 clips have been read	Success
Delete:	3 clips have been deleted	Success

Copy the .pea file to the IS server

Copy the .pea file from the Centera unit to every IS server in the system. The .pea file must be in the same directory path on each server, such as **/fnsw/local/sd/1/merged.pea**

Create new links for HP-UX 9000 only

On HP-UX servers (HP 9000 only), you must link the Centera libraries that are under **/fnsw/lib/shobj** to a new **/fnsw/lib/shobj/3rdparty/centera** directory and then add that path to the environment variable **SHLIB_PATH** to support EMC Centera SDK 3.1.

On HP-UX 9000 systems only, create new links for the client shared libraries for Centera:

- 1 Stop the FileNet software and kill all remaining processes by entering:

```
initfnsw -y stop  
killfnsw -ADy
```

- 2 As the FileNet software user, such as **fnsw**, create a new directory for the Centera shared library links:

```
mkdir /fnsw/lib/shobj/3rdparty  
chown fnsw:fnusr /fnsw/lib/shobj/3rdparty  
mkdir /fnsw/lib/shobj/3rdparty/centera  
chown fnsw:fnusr /fnsw/lib/shobj/3rdparty/centera
```

- 3 Change to the newly created directory:

```
cd /fnsw/lib/shobj/3rdparty/centera
```

- 4 Create new links for Centera by entering:

```
In -s /fnsw/lib/shobj/libFPLibrary.sl libFPLibrary.sl
In -s /fnsw/lib/shobj/libFPLibrary.sl libFPLibrary32.sl
In -s /fnsw/lib/shobj/libFPParser32.sl libFPParser32.sl
In -s /fnsw/lib/shobj/libPAI_module32.sl libPAI_module32.sl
```

- 5 Set the ownership for the new links by entering:

```
chown -h fnsw:fnusr *
```

Modify the .profile and .cshrc files

Modify the .profile and .cshrc files for the FileNet software user, such as fnsw, and any other users in the fnadmin and fnop groups who are responsible for starting and stopping the Image Services software.

- 1 Log in and use your preferred text editor to edit the .profile and .cshrc files.
 - In the Bourne or Korn shell, add the following line to the .profile file:

```
export SHLIB_PATH=$SHLIB_PATH:/fnsw/lib/shobj/3rdparty/centera
```
 - In the C shell, add the following line to the .cshrc file:

```
setenv SHLIB_PATH $SHLIB_PATH:/fnsw/lib/shobj/3rdparty/centera
```
- 2 Repeat the previous step for each user.
- 3 Then start the Image Services software.

Add the Centera device to the IS configuration

Return to [Chapter 2, “Define the SDS configuration in Image Services,” on page 42](#) to add this SDS unit to the Image Services configuration.

Appendix C – IBM DR550

What is IBM DR550?

The IBM TotalStorage DR550 packages storage, server, and software retention components into a lockable cabinet. For installation information, refer to the *IBM TotalStorage DR550 Version 2.0 Installation, Setup, Operations, and Problem Determination Guide* at the IBM Web site.

Image Services 4.1.1 includes the IBM Tivoli® Storage Manager (TSM) client libraries.

Note HP Integrity servers do not support IBM DR550.

What is IBM Tivoli Storage Manager?

TSM is a software product that enables you to protect your organization's data through features like backup and recovery and archive management. (See the IBM Web site for details.)

TSM Configuration Steps

In addition to the general configuration settings discussed in the main body of this document, there are some specific configuration settings that apply to the IBM DR550 discussed in the following sections.

- Create the SDS family and document classes (without specifying the SDS unit) as described in **Chapter 2, “Define SDS Families and Document Classes,” on page 51.**

- Run the TSM scripts as described in this appendix, beginning with the section, **“Set up the IS software for IBM DR550” on page 125.**
- Use the IS System Configuration Editor (fn_edit) to add the DR550 device to the IS configuration as described in **Chapter 2, “Define the SDS configuration in Image Services,” on page 42.**
- Restart the Image Services software.
- Update the SDS family with the SDS unit information as described in **Chapter 2, “Define SDS Families and Document Classes,” on page 51.**

Set up the IS software for IBM DR550

The IS software contains the following two scripts: TSM_SetupSvr and TSM_SetupSvr.awk. To set up an instance of IBM DR550, complete the following steps:

- 1 After you have created the family and document classes for your SDS documents, run the ddexim tool on the IS server you want to communicate with the IBM DR550 device:

```
ddexim -e > inxdb.<ssn>
```

where <ssn> is the system serial number of the IS system.

- 2 Copy the following files from the IS server to the DR550 device:

```
inxdb.<ssn> TSM_SetupSvr TSM_SetupSvr.awk.
```

- 3 On the IBM DR550 console, run the TSM_SetupSvr script with the <ssn> argument as shown below to generate three files: FN<ssn>, FN<ssn>.scr, and FN<ssn>.rpt.

TSM_SetupSvr <ssn>

- 4 On the IBM DR550 console, change the permissions of the .cmd file:

chmod +x FN<ssn>.cmd

- 5 On the IBM DR550 console, run the following command to set up the domain, policyset, and management class, and pipe the output to a log file:

FN<ssn>.cmd | tee FN<ssn>.out

where <ssn> is the ssn of the IS server.

- 6 Review the output to make sure no errors occur. Below is an example of the beginning, middle, and ending of this file without any errors.

```
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 2, Level 0.0
(c) Copyright by IBM Corporation and others 1990, 2003. All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 3, Level 0.0
  Server date/time: 11/22/07  17:47:16  Last access: 11/21/07  19:30:40

ANS8000I Server command: 'Delete Script FN16164_Setup'
ANR1455E DELETE SCRIPT: Command script FN16164_SETUP does not exist.
ANS8001I Return code 11.

ANS8002I Highest return code was 11.

...

ANR1514I Policy set FN16164SET activated in policy domain FN16164.
ANR1462I RUN: Command script FN16164_SETUP completed successfully.

ANS8002I Highest return code was 0.

...

ANS8000I Server command: 'Query Node FNP8IM_16164 Format=Detailed'
ANR2034E QUERY NODE: No match found using this criteria.
ANS8001I Return code 11.

ANS8002I Highest return code was 11.
```

In this example, the line in bold text indicates instance FN16164 was successfully set up.

- 7 Review the FN<ssn>.cmd file. On the IBM DR550 console, run the line in the file that has been “commented out” (bold in the example below)

to create the IBM DR550 instance. You can copy the line without the comment character from the file to the command line.

```
dsmadm -id=admin -password=admin Delete Script FN16164_Setup
dsmadm -id=admin -password=admin Define Script FN16164_Setup File=`pwd`/
FN16164.scr Desc=\'FileNet - SSN 16164 Domain Setup Script\'
dsmadm -id=admin -password=admin Run FN16164_Setup
#dsmadm -id=admin -password=admin Register Node FNP8IM_16164 SysAdmin
Domain=FN16164 Comp=Client
dsmadm -id=admin -password=admin Query Domain FN16164 Format=Detailed
dsmadm -id=admin -password=admin Query Policy FN16164 Format=Detailed
dsmadm -id=admin -password=admin Query MgmtClass FN16164 Format=Detailed
dsmadm -id=admin -password=admin Query Node FNP8IM_16164 Format=Detailed
```

- 8** After you run the line from the above file, you receive output on the console screen similar to the example below, and toward the end you will see output that indicates **FNP8IM_<ssn>** is registered in policy domain **FN<ssn>** (bold below).

```
IBM Tivoli Storage Manager
Command Line Administrative Interface - Version 5, Release 2, Level 0.0
(c) Copyright by IBM Corporation and others 1990, 2003. All Rights Reserved.

Session established with server TSM: AIX-RS/6000
  Server Version 5, Release 3, Level 0.0
  Server date/time: 11/22/05  17:47:16  Last access: 11/21/05  19:30:40

ANS8000I Server command: 'Register Node FNP8IM_95780 SysAdmin Domain=FN95780 Comp=Client'
ANR2060I Node FNP8IM_95780 registered in policy domain FN95780
ANR2099I Administrative userid FNP8IM_95780 defined for OWNER access to node FNP8IM_95780
```

Note If a new document class is added to the system, repeat the steps in this section to rerun the **TSM_SetupSvr** script and add the new document classes to the IBM DR550 environment.

Restart the IS Software

Restart the IS software on all servers to put the new configuration into effect.

Verify the TSM API Environment Variables

The dsm files must be set up for the TSM software to communicate with the IBM DR550 device.

First Time Installation of TSM

If this is a new installation of TSM and you are using TSM only with IS, you do not need to do much. By default, IS creates the dsm.sys (UNIX only) and dsm.opt (UNIX and Windows) configuration files in the following locations when the IS software is brought up, depending upon your operating system:

- /fnsw/local/logs/sds/SDSw_tivoliDR/<SDS unit name>/<instance number>/dsm.sys (UNIX)
- /fnsw/local/logs/sds/SDSw_tivoliDR/<SDS unit name>/<instance number>/dsm.opt (UNIX – placeholder file)
- <drive>:\fnsw_loc\logs\sds\SDSw_tivoliDR\<SDS unit name>\<instance number>\dsm.opt (Windows)

Note One dsm.sys and one dsm.opt file are created for each instance of TSM on a UNIX system, but the dsm.opt on UNIX is a placeholder file for trace file information or other options.

On a Windows system, only the dsm.opt file exists and it holds all of the information that both the dsm.sys and dsm.opt files hold on UNIX.

For information on setting up API tracing, go to [“Set API Environment Variables for Tracing \(Optional\)” on page 133](#).

These files are expected by the initial TSM API call (dsmSetup) for the APIs to work correctly. API library error messages are logged to one of the following files:

UNIX

/fnsw/local/logs/IBM_DR550/rdr_TSM_err_yyyymmdd

WIN

<drive>:\fnsw_loc\logs\IBM_DR550\rdr_TSM_err_yyyymmdd

Pre-existing installation of TSM

If this is a pre-existing installation of TSM and you are using TSM with other applications in addition to IS, you probably already have versions of the **dsm.sys** and **dsm.opt** files (UNIX) or a **dsm.opt** file (Windows) in a pre-existing location. If the environment variables already point to this location, the IS software does not touch these files.

In the dsm.sys file (UNIX) or the dsm.opt file (Windows), add the lines shown in the following example. Make sure the settings are appropriate for your system, and that the Tcpserveraddress line matches the IP address of the IBM DR550 as configured in the repository location string in fn_edit.

Servername	FileNet_SDS
Commmethod	tcpid
Tcpserveraddress	212.3.248.36
EnableArchiveRetentionProtection	yes

Note The **dsm.opt** file is used for optional API tracing or other options. For information on setting up API tracing, go to [“Set API Environment Variables for Tracing \(Optional\)” on page 133](#).

Set the MaxSessions Parameter

It is important to consider how many systems are going to be in communication with the IBM DR550 and to make sure the MaxSessions parameter is set to accommodate that number and more. To do this, complete the following steps:

- 1 Login the DSM Admin Console (IBM DR550).
- 2 Enter the following command at the TSM> prompt to see the current setting of the Maximum Sessions parameter:

query system

You will see output similar to the following, where the Maximum Sessions setting is shown in bold text:

```
tsm: TSM>query system

*****
*** ---> Q STATUS
*****

Storage Management Server for AIX-RS/6000 - Version 5, Release 3, Level 0.0

                Server Name: TSM
    Server host name or IP address:
                Server TCP/IP port number: 1500

(continued on next page)
```

(continued from previous page)

```
Server URL:
Crossdefine: Off
Server Password Set: No
Server Installation Date/Time: 02/26/04 17:40:30
Server Restart Date/Time: 11/09/07 16:01:25
Authentication: On
Password Expiration Period: 90 Days
Invalid Sign-on Attempt Limit: 0
Minimum Password Length: 0
WEB Admin Authentication Time-out (minutes): 5
Registration: Closed
Subfile Backup: No
Availability: Enabled
Accounting: Off
Activity Log Retention: 7 Days
Activity Log Number of Records: 5804
Activity Log Size: <1 M
Activity Summary Retention Period: 30 Days
License Audit Period: 30 Days
Last License Audit: 12/04/06 16:01:46
Server License Compliance: Valid
Central Scheduler: Active
Maximum Sessions: 256
Maximum Scheduled Sessions: 128
Event Record Retention Period: 10 Days}
```

The example above shows Maximum Sessions set to 256, which would mean no more than 256 users could concurrently be accessing a particular DR550 device.

- 3 Determine the setting for your site (suggest 100 per server) that is going to communicate with the DR550 device and enter a command similar to the following at the TSM> prompt:

SETOPT MAXSESSIONS 200

where you are changing the setting to 200.

Add the DR550 device to the IS configuration

Use the IS System Configuration Editor (fn_edit) to add the DR550 device to the IS configuration as described in [Chapter 2, “Define the SDS configuration in Image Services,” on page 42](#). When you have finished step 18 on page 47, return to this appendix.

Restart the IS Software

Restart the IS software on all servers to put the DR550 configuration changes into effect.

Update the SDS family

- Return to [Chapter 2, “Define SDS Families and Document Classes,” on page 51](#) to update the SDS family with the appropriate information for DR550.

Set API Environment Variables for Tracing (Optional)

Use the following tracing method for troubleshooting the IBM DR550 device.

- You can optionally set the DSMI_CONFIG environment variable to take advantage of the API tracing provided by Tivoli Storage Manager.

Note Be aware that tracing might have a significant impact on system performance.

Set TSM Tracing Method

TSM tracing provides you with very detailed tracing information. To enable TSM tracing, add two lines of text to the dsm.opt file (for both UNIX and Windows) and point the DSMI_CONFIG environment variable to the full path name <path> of your dsm.opt file.

To set the environment variable enter the following commands:

On UNIX platforms:

- For Bourne and Korn shell users, enter the following command:

```
export DSMI_CONFIG=/<path>/dsm.opt
```

- For C shell users, enter the following command:

```
setenv DSMI_CONFIG /<path>/dsm.opt
```

On Windows platforms, you must use the graphical user interface to set the environment variable:

- a Right-click My Computer > select Properties > select the Advanced tab > click Environment Variables.
- b In the Environment Variables window, click **New** in the System variables pane.
- c In the New System Variable pop-up window, enter **DSMI_CONFIG** in the Variable Name field and **<drive>:\<path>\dsm.opt** in the Variable Value field (where <path> is the location of your dsm.opt file) and click **OK**.

Tracing on UNIX Systems

After setting the environment variables and either creating new or using the default **dsm.opt** file, make sure they have content similar to that shown in the following example:

In the **dsm.opt** file, enter trace information as shown.

```
Traceflags api perform
Tracefile <api trace file name>
```

All the extensive API trace file information provided by TSM is in the **<api trace file name>** you specify in the **dsm.opt** file.

Tracing On Windows Systems

After setting the environment variables and either creating a new file or using the default **dsm.opt** file, make sure the contents are similar to the following example. Make sure the settings are appropriate for your system, and that the **Tcpserveraddress** line matches the IP address of the IBM DR550 from the **as_conf.g** file.

```
Servername                FileNet_SDS
Commmethod                tcpid
Tcpserveraddress          212.3.248.36
EnableArchiveRetentionProtection yes
Traceflags api perform
Tracefile <api trace file name>
```

All the extensive API trace file information provided by TSM is in the **<api trace file name>** you specify in the **dsm.opt** file.

Restart the IS Software

Stop the IS software with **killfnsw** and restart the IS software on all servers to put the TSM Tracing into effect.

Return to [Chapter 2, “Define SDS Families and Document Classes,” on page 51](#) for information on adding the DR550 to the SDS family.

Appendix D – Windows network security settings

To access file system repositories (the directories where the repository strings point) on remote network storage devices such as Network-Attached Storage (NAS) or Storage Area Network (SAN) devices, or on remote Windows devices such as a Windows shared drive or disk array in a Windows environment, you must configure your IS system appropriately.

Important

Use this procedure to create a folder for the MSAR creation directory before committing any documents.

How to create a Domain Account for Image Services

Create, or verify, a domain group account in a domain where both the IS server and the remote storage device are members.

For example, <domain>\domain fnadmin

Use a domain account that includes the owner of the IS Control Service and any other users that would need to administer MSAR surfaces.

Note

The examples in this appendix use **domain1** for the domain name and **domain fnadmin** for the user name.

How to set security properties for remote network storage devices

Use this procedure if you will be accessing a remote network storage device, such as NAS or SAN.

Note If you will be accessing a remote Windows storage device, skip this section and continue to, [“How to set security properties for remote windows server storage devices”](#).

Perform the following steps on the **remote network storage device**.

- 1 Create the MSAR directory with ample space allocated for your MSAR surfaces.
- 2 Set up permissions for the MSAR directory and give <domain>\domain fnadmin full control.
- 3 Continue to the section, [“How to stop IS software and IS ControlService”](#) on page 140.

How to set security properties for remote windows server storage devices

Use this procedure if you will be accessing a remote Windows storage device such as a Windows shared drive or disk array in a Windows environment. The security properties must give full control to the domain user account that the IMSService is using.

Note The remote Windows server, containing the shared device, does not need to have IS installed.

Perform the steps in this procedure on the **remote Windows Server** with the Storage Device.

- 1 In Windows Explorer, create a directory folder for MSAR.

Note In this procedure the directory folder is named MSAR_Share.

- 2 Right-click on the MSAR folder (MSAR_Share) and select *Properties*. The properties window for the folder opens.
- 3 In the properties window, select Security tab to display the Security tab dialog box.
- 4 Click the **Add** button to open the Select Users, Computers, or Groups window.
- 5 From the drop-down list in the Look in: box, select the domain.
- 6 Select the domain user you created in the section, **“How to create a Domain Account for Image Services” on page 137**, click **Add**, and then click **OK**.
- 7 In the Security tab dialog box, select the domain user and check the **Full Control** checkbox, under the **Allow** column, to grant full control permissions for the domain user selected. Leave the properties window open.


How to stop IS software and IS ControlService

Perform the steps in this procedure on the **IS server**.

- 1 Open a Command Prompt and stop IS:

initfnsw -y stop

- 2 From the Control Panel, open Administrative Tools and then open Services.
- 3 In the Services window, right-click IS ControlService and select **Stop**.

Note You can also click the  button on the tool bar at the top of the Services window to stop the IS Control Service.

How to modify the IS ControlService to use the domain account

Perform the steps in this procedure on the **IS server**.

- 1 In the IS ControlService Properties window, select the Log On tab.
- 2 From the Log On tab, click the **Browse** button to open the Select User window.
- 3 From the drop-down list in the Look in: box, select your domain.
- 4 Select your domain user name and click **OK**.
- 5 In the IS ControlService window, enter the domain user password, confirm the password entry, and click **OK**.
- 6 In the Services window, right-click IS ControlService and select **Start**.

Note You can also click the  button on the Services tool bar at the top of the window to start the IS Control Service.

- 7 After verifying that the IS ControlService starts successfully, close the Services window.

How to add the domain account to the following groups: Administrators, FNADMIN, FNOP, FNUSR

Perform the steps in this procedure on the **IS server**.

- 1 From the Taskbar, click Start > Programs > **Administrative Tools** and click **Computer Management**.
- 2 Expand the Local Users and Groups folder and select **Groups**.
- 3 Double-click the Administrators group to open the Administrators Properties window.
- 4 Click **Add** to open the Select Users or Groups window.
- 5 From the drop-down list in the Look in: box, select your domain.
- 6 Select the domain user name you created in the section, [“How to create a Domain Account for Image Services” on page 137](#), click **Add** and then **OK**.
- 7 Click **OK** to close the Administrators Properties window.
- 8 Repeat Steps 3 through 7 for the remaining groups: FNADMIN, FNOP, and FNUSR.

How to modify the local security policy for Domain Account

Perform the steps in this procedure on the **IS server**.

Modify the local security policy to give the domain account permissions for the following policies:

- Act as part of the operating System
- Log on as a service
- Increase quotas
- Replace a process token

- 1** From the Taskbar, click Start, point to Programs, point to the **Administrative Tools**, and click **Local Security Policy**.
- 2** Expand the Local Policies folder and select the **User Rights Assignment** folder.
- 3** Right-click on the policy selection you want to add and select **Security**.
- 4** Click the **Add** button to open the Select Users or Groups window.
- 5** From the drop-down list in the Look in: box, select your domain name.
- 6** Select the user name you created in the section, [**“How to create a Domain Account for Image Services” on page 137**](#), click the **Add** button, and click **OK**.
- 7** Click **OK** to close the Local Security Policy Setting window.
- 8** Repeat Steps 3 through 7 for the remaining Policy selections. After all Policy selections have been modified, continue to the next section.

How to configure a storage library to use the MSAR directory

Perform the steps in this section on the **IS server**.

Use the System Configuration Editor to configure a storage library for the remote server.

- 1** Open the System Configuration Editor.
- 2** Select the Configure an MSAR Storage Library procedure and click **Run**.
- 3** Respond to the prompts in the procedure as appropriate for your site. When you get to the step where you enter the MSAR surface directory for the MSAR Library, enter the UNC (Unified Naming Convention) path to the share name. For example: \\<servername>\sharename
- 4** Complete the remaining steps of the procedure, exit the System Configuration Editor and save your changes.

Start the IS Software

Open a Command Prompt and start the IS software:

```
initfnsw start
```

Appendix E – Centera License Agreement

Copyright © 2003 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

Copyright © 1991-2, RSA Data Security, Inc. Created 1991. All rights reserved.

License to copy and use this software is granted provided that it is identified as the "RSA Data Security, Inc.

MD5 Message-Digest Algorithm" in all material mentioning or referencing this software or this function. RSA Data Security, Inc. makes no representations concerning either the merchantability of this software or the suitability of this software for any particular purpose. It is provided "as is" without express or implied warranty of any kind.

These notices must be retained in any copies of any part of this documentation and/or software.

Copyright (c) 1995-2002 International Business Machines Corporation and others. All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Glossary

Ageable cache object

Cache objects that can be aged out of page cache. Typically a cache object is made ageable after it has been migrated to permanent storage and the write request is removed.

Background Job Control

Background Job Control is an Image Services application used to start, control, and monitor background jobs.

background jobs

Background jobs are time-intensive IS software functions designed to run at lower priority and without user attention (that is, in the background). As such, they do not interfere with other tasks. Copying media, importing documents from storage media, and consolidating media are examples of functions that are performed as background jobs. Background job progress is checkpointed and will automatically restart from where it stopped if the software is restarted.

blob

A collection of binary data stored as a single entity in a database management system. The term blob was originally used by James Starkey at DEC to refer to an amorphous chunk of data. Blob was later interpreted as a “backronym” for basic large object or binary large object.

cache

Cache is the magnetic disk space used to store documents on the way to and from storage media. Cache can also act as permanent storage if you do not use optical or MSAR storage media. Portions of cache (referred to as logical caches) are allocated for storage of different document types. IS uses multiple cache types.

Cache-only

Some IS customers choose not to use optical or MSAR storage libraries for their systems. All of their committed documents stay locked in page cache. This type of configuration is called a Cache-only system.

Cache-only Storage Server

Some FileNet customers do not use MSAR or optical storage libraries with their systems. All of their committed documents stay locked in page cache.

checksum

The arithmetic sum of the binary data in an object. Checksumming is used to verify image integrity. The checksum is normally calculated when a document is captured and stored with the document when it is committed. A checksum can be recalculated and compared with the original checksum at any point in the lifecycle of the document.

committal

Committal is the stage at which scanned images become retrievable documents. Index values are written to the index database and an entry in the document locator table of the permanent database is inserted. For systems equipped with a

storage library or SDS unit, requests to migrate the documents to the storage library or SDS unit are written to the write_request table.

CSM Page cache

CSM Page cache, also known as retrieval cache, is a cache containing all documents being committed to or retrieved from storage media. Page cache is an ageable cache, meaning that documents not locked can be aged out to make room for newer documents.

database

A database is a collection of logically related records or files managed by a database management subsystem. The IS system uses two types of databases: a third-party relational database for index data and WorkFlo queues and **“MKF (multi-keyed file)”** for document addresses, work in progress, and other purposes.

ddexim

The ddexim tool saves (exports) and restores (imports) document class and index information in the index database.

Docs table

Also known as the Document locator table, this Permanent MKF database table contains information about each document such as the associated primary and tranlog surfaces and offsets to the document (Long descriptor or TOC). There is only one docs table per system and it is on the primary Storage Library Server. Documents stored on all Storage Library Servers will have an entry in the docs table.

Doctaba table

The relational database table where all user Index information about documents is stored.

Document class

A logical grouping of similar documents. A document's scanning, indexing, and security characteristics are defined by the document class. All documents assigned to that class share the same characteristics. A document class is associated with a media family.

Document family

This is also known as media family. There are two types of families: primary and tranlog families. A family is used to migrate documents to a configured media type. New SDS ID and SDS only properties have been added to the family paradigm. These properties are only valid for the primary families. The database maintenance family configuration UI (Xdodprog) supports these new SDS properties.

Document locator (DOCS) table

The document locator table is part of the MKF permanent database and resides on the primary storage server. The DOCS table contains a record for each document id number and its location whether on optical media, MSAR, or integral SDS.

EBR

Event-Based Retention.

EMC Centera domain name

Used to connect to the appropriate Centera object repository.

EMC Centera Object Repository

EMC Centera Object Repository stores data based on the actual content value. This is a content based addressing system.

Fast Batch Committal

Fast batch committal is a quick way of moving documents from batch cache to page cache. In a fast batch committal environment, all documents and all pages of the documents for the batch are in one cache object, so the entire batch is committed in a single operation. Unlike normal committal, fast batch committed documents are not retrievable from the cache. They must be migrated to and then from a storage library before a retrievable copy of the document is in the page cache.

If the "Fast Batch Breakup" option is selected, the fastbatch object remains in the cache after it has been written to a surface. This option also causes the fastbatch object to be broken up so that individual documents are retrieved.

Image Services

Image Services is a set of servers and services providing a single document image database. The database includes a single index database, a single document locator database, and the collection of document images on storage media.

In box surface

A surface that is physically available in a library.

index database

The index database contains document and folder information and is implemented by relational database management system such as Oracle, DB2®, or Microsoft® SQL Server™.

Locked objects

Images not yet written to storage media are "locked" in cache and cannot be aged out (that is, they cannot be deleted or moved). In a system without a storage library, all documents are locked objects in page cache.

MKF (multi-keyed file)

The MKF stands for multi-keyed file databases. They include the transient database, the permanent database, the NCH database, and the security database.

Migration

Migration is the process of writing documents from batch cache to the storage library. It can also be used to describe the retrieval of documents from the storage library into page cache.

MSAR

Magnetic Storage and Retrieval library. This is a FileNet term for using available magnetic storage for IS documents and treating this space as a storage library.

Multiple storage library server system

IS systems can have multiple storage library servers, each of which can manage multiple libraries. If you have multiple storage library servers, one serves as the master server (called a document locator server) that keeps track of the contents of all storage libraries. Some

diagnostic tools and commands must be run from the document locator server.

The document locator server includes the complete permanent database that maps each document number into one or two media locations. Other storage library servers contain small permanent databases that store this information until enough accumulates to transfer to the document locator server.

NAS (Network-Attached Storage)

File-level computer data storage connected to a computer network providing data access to network clients.

NLS-SDS

Near Line Storage - Single Document Storage. The previous implementation of Single Document Storage and Retrieval used an add-on software component known as Near Line Storage. The system was limited to one SDS device, and configuration files had to be updated manually.

OSAR

OSAR is an acronym for the proprietary Optical Storage and Retrieval unit. An OSAR is a type of storage library.

out-of-box surface

A surface that is not physically available in a library. An RSVP message is issued if a document on that surface has been requested.

permanent database

The permanent database stores the media location of each document entered into the system and contains tables for media surfaces, media families, and annotations.

primary family

A collection of media that stores documents from one or more document classes. Before you can complete the definition of a document class, you must assign it to a media family so the system knows where to store the documents in that class. You can also create a primary family to be used exclusively for archiving the index and permanent databases.

primary storage server

In IS systems that are configured with multiple storage servers, the primary storage server is the storage server where the centralized document locator table is maintained. In addition, the surf_locator and family_locator tables also reside on the primary storage server.

QMA queue

The QMA queue is a part of the Queue Manager Abstract, which manages shared memory queues. The IS software has many well known or permanent queues that are created during software startup. Among these are a scheduler queue for a dsched process, a dtp queue for each dtp process, a dtp_tran queue for each dtp_tran process, a remote committal queue for a rmt_commit process, and so on.

These permanent queues are created or found by the entry points QMA_find_queue. Permanent queues cannot be deleted without restarting the IS software.

The first time `QMA_find_queue` is called for a permanent queue it is created, and subsequent calls to `QMA_find_queue` find the queue instance. The `QMA_find_queue` creates or finds a permanent queue with a queue class ID and associated size of each queue entry. In addition, `QMA_find_queue` can be called with up to 256 instance indicators. The instance indicator is used to support multiple SDS units by mapping the SDS unit number to QMA instance number.

Each instance of a QMA queue class is a separate queue and is managed independently. To send a request, the sender process must create a queue entry (with `QMA_create_entry`) and insert it into the tail or head of the queue instance (with `QMA_insert_last` or `QMA_insert_first`). The receiver process can dequeue a request from the top of the queue by calling `QMA_wait_remove_first`.

QMA has a locking mechanism to prevent two different processes from accessing or retrieving from the same queue at the same time. SDS uses the QMA structure to manage each SDS unit.

For SDS, the IS dispatch process calls the `SDS_submit_read_request` entry, which in turn calls the `QMA_create_entry` entry to queue a read request to the SDS QMA queue instance.

Each SDS worker task calls the `SDS_get_next_request` entry, which will in turn call the `QMA_wait_remove_first` entry to dequeue read requests from the SDS QMA queue instance. The SDS enhancement will take advantage of the existing locking and separately managed queue instances infrastructure in QMA.

Read Enabled Surface

A database setting (surf_info table) that enables retrieval requests for a particular surface. If "read enable" is set to "not available," read requests will not be submitted to the particular surface.

Remote System Committal

Documents entered on one IS system can be transmitted to and committed on another IS system at a remote location.

RSVP message

An RSVP message requires the operator to respond and perform some action, such as loading storage media. RSVP messages are displayed by the **“SLC (Storage Library Control)”** UI program.

SAN (Storage Area Network)

A mass storage device independently managed for access by many different users across a network.

Secondary Storage Server

For Multiple Storage Configured systems, the Secondary Storage Server is any Storage Library Server that is not the Primary Storage Server.

Server

A single computer that provides services to other computers (clients) in the system.

SLC (Storage Library Control)

SLC is an IS application that monitors the status of storage libraries on the server where SLC is running and provides messages about the media. SLC also provides status information, reports, and a user

interface to initiate such actions as enabling or disabling drives, pre-formatting media, and so on.

SSN (System Serial Number)

The system serial number is a unique ten-digit number provided to you by IBM. For UNIX systems, this number is based on the unique processor ID of the server. For Windows systems, this number is based on the network adapter address.

system - Image Services

An Image Services server, or multiple servers, in its own domain is often referred to as a system.

transaction log (tranlog)

The IS system is configured to write one or more backup copies of committed documents to transaction log media. A media family can be configured to be a tranlog family. A primary family can be configured to use a tranlog family. Tranlogs, when they exist, are always written before the primary.

virtual EBR

Event-Based Retention for SDS devices without native EBR support. See [“EBR” on page 149](#).

Write_request table

This is a transient MKF table used to insure that writes to permanent storage occurs. When a document is committed to an IS system a write_request entry is created and write request is submitted. If the write is not complete and IS software is recycled, the IS software automatically regenerates the write request to the device from write_request entries.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (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 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS 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 may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish 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
J46A/G4
555 Bailey Avenue
San Jose, CA 95141-1003
U.S.A.

Such information may 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 Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of

non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml.

FileNet is a registered trademark of FileNet Corporation, in the United States, other countries, or both.

Microsoft and Windows 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.

U.S. Patents Disclosure

This product incorporates technology covered by one or more of the following patents: U.S. Patent Numbers: 6,094,505; 5,768,416; 5,625,465; 5,369,508; 5,258,855.



Program Number: 5724-R95

Printed in USA

GC31-5535-01

