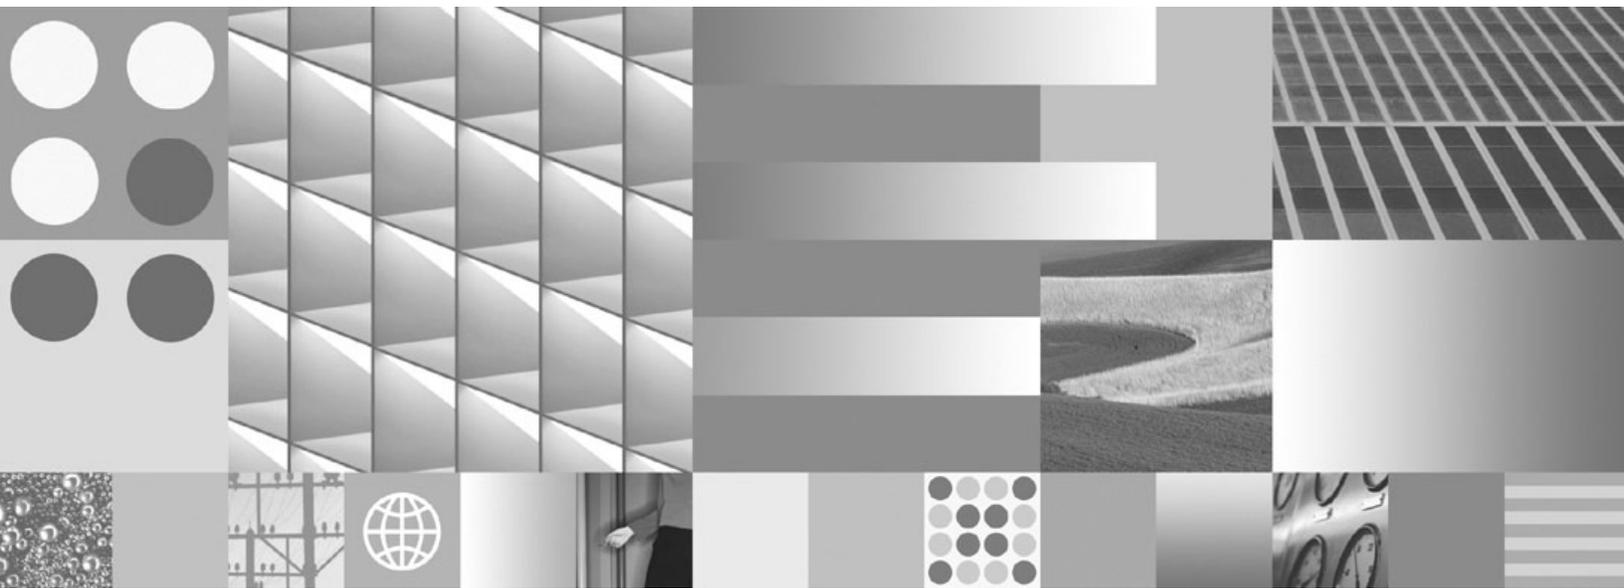
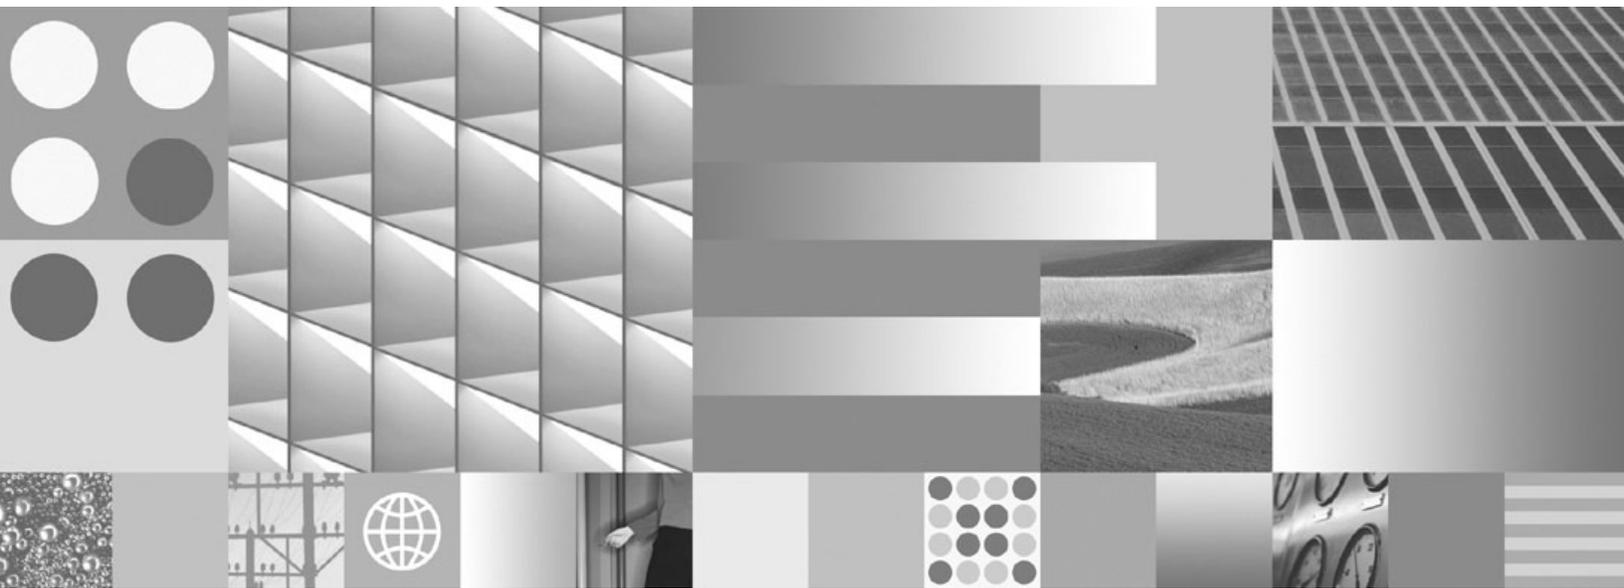


4.1.2



Single Document Storage and Retrieval Guidelines (NLS-based)

4.1.2



Single Document Storage and Retrieval Guidelines (NLS-based)

Note

Before using this information and the product it supports, read the information in "Notices" on page 194.

This edition applies to version 4.1.2 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 2005, 2008.

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

Contents

About This Manual 14

Audience 14

Document revision history 15

Accessing IBM FileNet Documentation 15

Conventions Used in this Manual 15

File Paths 15

Console Displays 16

Cautions, Notes, and Tips 16

Command Syntax 17

 Optional Parameters 17

 Required Parameters 17

IBM FileNet Education 17

Feedback 18

Documentation feedback 18

Product consumability feedback 18

1 Getting Started

Introduction 19

NLS Archive 22

 The File System Working Directory 22

 The Archive Process 24

SDS Retrieval 25

The Retrieval Process 26

Prerequisites 28

Image Services Software 28

AIX/6000 29

HP-UX (HP 9000 and HP Integrity) 29

Sun/Solaris 29

Windows Server 29

Image Services Toolkit Software 29

NLS Archive Software 30

NLS Archive Version 30

NLS Archive Disk Space 30

2 Installing NLS Software

Install NLS Software 31

Install NLS Software on a UNIX Platform 31

Install NLS Software on a Windows Server 34

3 Configuring NLS Software

NLS Archive Prerequisites 38

Update the System Hosts File 38

Environment Path 39

IS Toolkit File Ownership and Permissions on UNIX 40

Creating the NLS Configuration File 40

NLS Configuration Sample File (NLS_cfg.sample) 41

Configuring the NLS.cfg file 44

Logon Attributes 45

Using the optional Encrypted Password Functionality (PS_Password Program) 46
Configuring UserName and PassWord in NLS.cfg File 48
Program Attributes 49
 WorkingDirectory 49
 LogDirectory 50
Other Attributes 50
Saving the NLS.cfg file 59
Device-specific Settings 60

4 Configuring SDS File System Retrieval

Creating the SDS Configuration File 61

sds_conf File Format 63
 SDS Name (sds_name) 63
 Worker Program Name (worker) 63
 SDS_FS_reader Information Field (info) 63
 Retrieval Priority (sds_priority) 65
 Low Priority 65
 Medium Priority 66
 High Priority 66
 Activity Reporting (debug) 67
Multiple IS Storage Library Servers 68

Configuring Shared Resources 68

NLS.cfg and sds_conf Settings 68

Device-specific Settings 69

5 Administering and Running NLS

Running the NLS_Archive Function 70

Autostart/Autostop NLS Method 70

Command Line Startup/Shutdown Method 72

Command Line Startup 72

Command Line Shutdown 73

NLS_FLAGS File 74

Windows Service Setup Method 75

Removing the FileNet NLS Archive Service 78

Archiving Documents 79

Database Mode 79

Document Range Override 80

File List Mode 81

Surface Mode 82

6 Administering and Running SDS File System Retrieval

Using the DOC_tool Utility 84

SDS States 84

ENable Command 86

DISable Command 87

SDSinfo Command 88

SDSinfo command prompt and options 88

SDSinfo Basic Information 89

SDSinfo Summary Option (s) 90

SDSinfo Queue Option (q) 97

SDSinfo Worker Option (w) 99

SDSinfo All Option (a) 102

Retrieving Documents 104

7 Adding a Physical Volume

Installing the New Physical Volume 105

Updating the NLS.cfg File 106

Updating the sds_conf File 106

Appendix A – Troubleshooting NLS Archive

Reporting Problems 108

**Utility/Operating System and IS Toolkit
Version Relationship 111**

Error Logging Files 112

NLS Archive Logs 112

Viewing Logs in a Windows Environment 112

Viewing Logs in a UNIX Environment 113

Other Logs 113

NLS Error Message Tuples 114

NLS Error Types 114

NLS Specific 115

IS Errors 115

NLS Error Codes 115

Error Codes Known to Need Specific Actions 118

When to Use wal_purge (UNIX Only) 120

Operating System Path Issue 121

Shell Problems Causing NLS Archive Failures (UNIX Only) 121

Appendix B – Troubleshooting SDS Retrieval

Error Logs 122

SDS Error Message Tuples 123

Using the fn_msg Tool 124

TPI (192) Messages 125

SDS (238) Messages 126

Appendix C – Centera Storage and Retrieval

Support Policy 129

What is Centera? 130

FNP_CLIPID 131

FileNet Specific Centera Object Format 131

Additional Internal Buffer Counter Information 133

 Whole Document BLOB Fits into the Internal Buffer 133

 Requested images Fit into the Internal Buffer 134

 Requested Pages Do Not Fit into the Internal Buffer 135

Configuration Steps 136

Set up the NLS software for EMC Centera 136

Verify the NLS.cfg File 136

 PgmAttribute 136

 SDS-Keywords 136

Verify the sds_conf file 146

Set up Centera Pool Entry Authorization (PEA) File 147
Restart the Image Services Software 148
Define a Document Class 148

Appendix D – SnapLock Storage and Retrieval

What is SnapLock? 151

Configuration Steps 151

Set up the NLS software for NetApp SnapLock 151

Verify the NLS.cfg File 151

 PgmAttribute 152

 SDS-Keywords 152

Verify the sds_conf file 157

Restart the Image Services Software 158

Define a Document Class 158

Appendix E – IBM DR550

What is IBM DR550? 160

What is IBM Tivoli Storage Manager? 160

Install the Tivoli API Client Package 160

Install the gsk* modules (AIX servers only) 161

Set Up TSM Library and Error Message Files 162

Define a Document Class 164

TSM Configuration Steps 165

Set up the NLS software for IBM DR550 165

Verify the NLS.cfg File 169

Verify the sds_conf file 174
Restart the Image Services Software 175
Verify the TSM API Environment Variables 175
 First Time Installation of TSM 176
 Pre-existing installation of TSM 176
Set the MaxSessions Parameter 177
Restart the Image Services Software 179
Setup NLS_Archive 179
Set API Environment Variables for Tracing (Optional) 180
Set TSM Tracing Method 181
 Tracing on UNIX Systems 182
 Tracing on Windows Systems 182
 Restart the Image Services Software 183
Set FileNet ISAR Tracing Method 183

Appendix F – Sun StorageTek 5320 Storage and Retrieval

What is Sun StorageTek™ 5320? 185

Configuration Steps 186

Set up the NLS software for Sun StorageTek 5320 186
Verify the NLS.cfg File 186
 PgmAttribute 186
 SDS-Keywords 186
Verify the sds_conf file 191
Restart the Image Services Software 192
Define a Document Class 192

Notices 194

Trademarks 198

U.S. Patents Disclosure 198

About This Manual

This document, *Image Services Single Document Storage and Retrieval Procedures and Guidelines*, contains information for installing, configuring, and administering the single document file system storage and retrieval product in conjunction with EMC Centera, NetApp SnapLock, IBM® TotalStorage® DR550, and Sun StorageTek 5320 storage solutions. It is intended to be read by people who are familiar with the UNIX® and Windows® Server operating systems, as well as with the FileNet® Image Services (IS) software.

Note

Throughout this document, the Single Document Storage (SDS) File System Retrieval product might also be referred to as the SDS Retrieval product and the IBM Tivoli® Storage Manager software might also be referred to as TSM.

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
4.1.2	Nov. 2008	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 On the Product Support page, click **Documentation** and then click **Product Documentation**.
- 4 On the Product Documentation page, locate the document you need, then click the icon in the appropriate release column to access the document.

Conventions Used in this Manual

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

File Paths

Since this manual is used for all platforms, examples of file path designations, where used, are given for both UNIX and Windows platforms. For example:

UNIX

`/fnsw/local/sd`

WIN

`<drive>:\fnsw_loc\sd`

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. *
```

Cautions, Notes, and Tips

Three message types call your attention to important information:

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

Note Draws your attention to essential information you should be sure to read.

Tip Introduces an idea that might make your work easier.

Command Syntax

Command syntax definitions are indented and in bold text:

```
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 file for the parameter in angle brackets, such as:

```
ddexim -e > myfile
```

IBM FileNet Education

IBM provides various forms of education. Please visit the IBM Information Management support page at (www.ibm.com/software/data/support).

Feedback

Documentation feedback

Send comments on this publication or other IBM FileNet Image Services documentation 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 help topic title, a chapter and section title, a table number, or a page number).

Product consumability feedback

Help us identify product enhancements by taking a [Consumability Survey](#). The results of this comprehensive survey are used by product development teams when planning future releases. Although we are especially interested in survey responses regarding the most recent product releases, we welcome your feedback on any of our products.

The survey takes approximately 30 minutes to complete and must be completed in a single session; there is no option to save a partially completed response.

Getting Started

Single document file system storage and retrieval enables FileNet Image Services to store documents in a file system on protected storage devices such as EMC Centera, Network Appliance Filer and NearStore, and IBM TotalStorage DR550.

Introduction

Single document file system storage and retrieval is composed of two components:

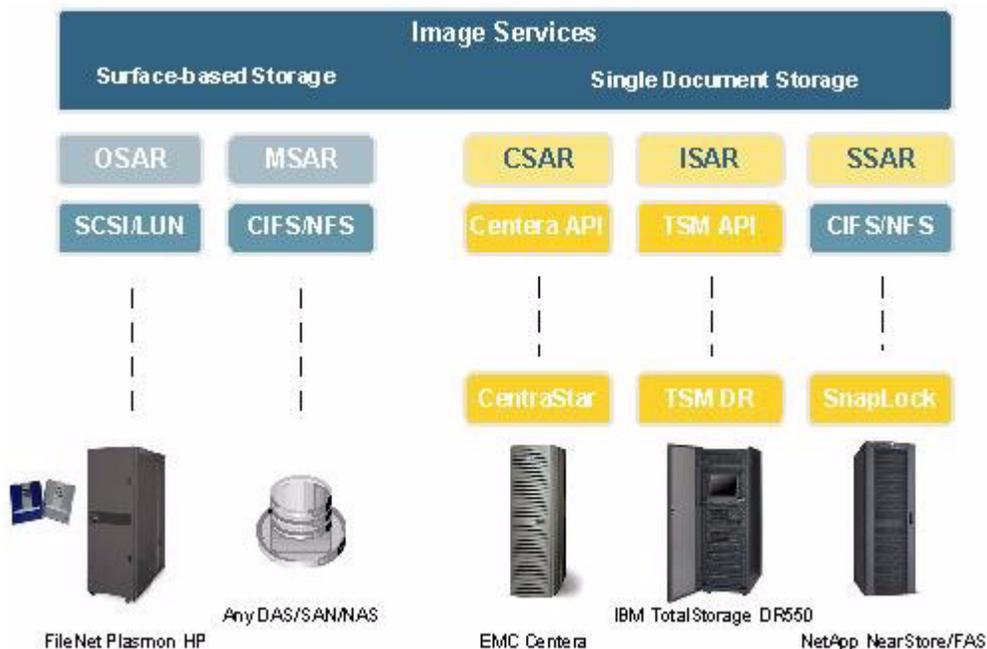
- **Near Line Storage (NLS)** stores Image Services documents on a the protected storage device.
- **Single Document Storage (SDS)** retrieves the documents from the protected storage device and returns them to Image Services.

The NLS component builds archive objects from Image Services documents, stores them in the appropriate file system, and applies retention parameters.

The SDS component is an integral part of Image Services, which supports these protected storage products:

- **CSAR** - EMC Centera Storage and Retrieval
- **ISAR** - IBM TotalStorage DR550 Storage and Retrieval
- **SSAR** - NetApp SnapLock Storage and Retrieval
- **SunSAR** - Sun StorageTek 5329 Storage and Retrieval

Surface-based storage technology, such as an OSAR and MSAR, writes a collection of pages (a document) into the container that represents an optical disk or magnetic surface file. Single document storage technology stores a collection of pages (a document) as individual objects on the protected storage device as shown in the following diagram:



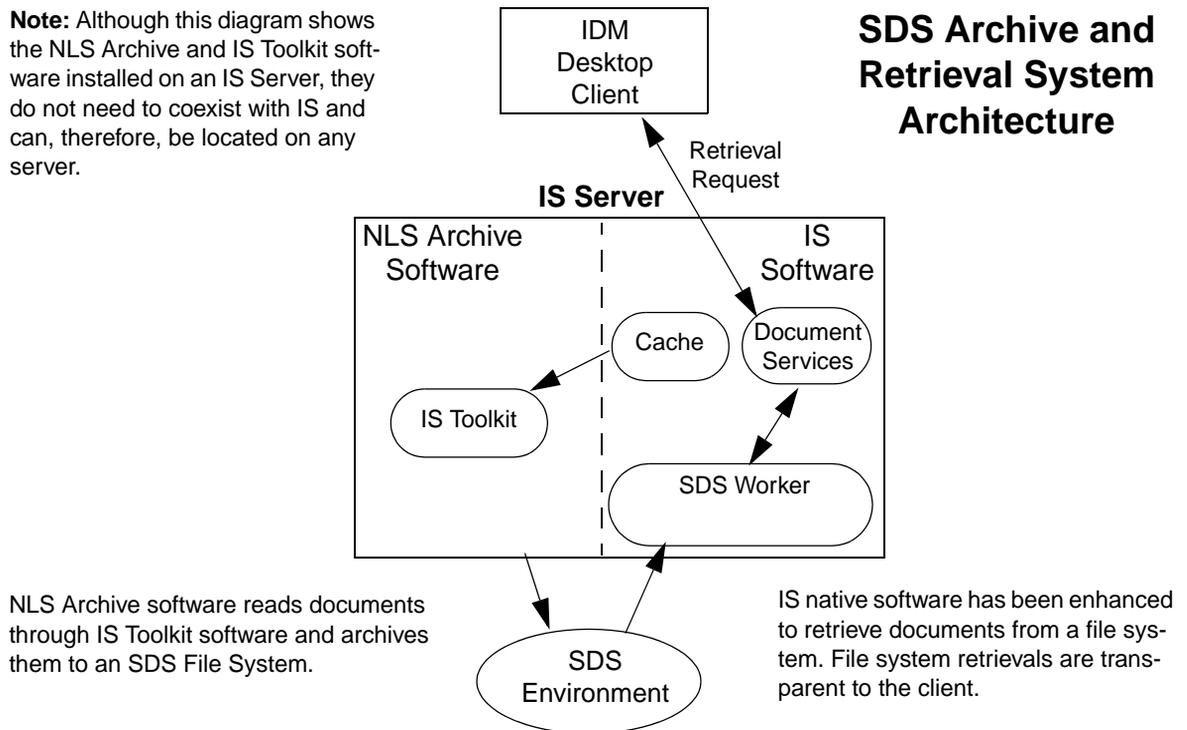
The FileNet xSAR Product Family

Note The SDS software includes an SDS worker process used to retrieve documents from the SDS device. Each supported SDS device has a unique and different worker. In the above example, CSAR = SDS_CSAR_reader, ISAR = SDS_ISAR_reader, and SSAR = SDS_FS_reader. See the appropriate appendix for your storage device. Only

one SDS unit worker per storage library server is supported in this release. (The unit number in the sds_conf file must be set to 1.)

The SDS Retrieval system architecture overview is illustrated in the following diagram.

Note: Although this diagram shows the NLS Archive and IS Toolkit software installed on an IS Server, they do not need to coexist with IS and can, therefore, be located on any server.



SDS Archive and Retrieval System Architecture

NLS Archive

The NLS Archive storage function is the process that identifies documents to be archived, requests storage of the documents, and updates the system to indicate that the documents have been archived.

Note Any type of file system can be used.

NLS Archive uses the FileNet IS Toolkit software to extract documents from Image Services and store them in the SDS File System.

The File System Working Directory

The SDS File System is specified as the WorkingDirectory. The WorkingDirectory specifies the top-level path where documents are stored.

The WorkingDirectory file system name is user-defined by editing the WorkingDirectory field in the NLS.cfg file so it is similar to the following:

```
; General information
PgmAttribute {
;   WorkingDirectory=/NLS
;   LogDirectory={Defaults HOME}
```

A sub directory is automatically created using the System Serial Number for the Domain, and all document files are stored in this sub directory, such as /NLS/<SSN>/.

Important Choose the WorkingDirectory file system name carefully, because once your system is installed and configured, it cannot easily be changed. In addition, multiple working directories are not supported.

Only one WorkingDirectory can be configured for archiving on any one system.

Keep in mind that the WorkingDirectory must be large enough to contain the documents that you anticipate will be stored here. Proper planning is important, to ensure that you have allocated enough disk space for this. However, this space is expandable as needed.

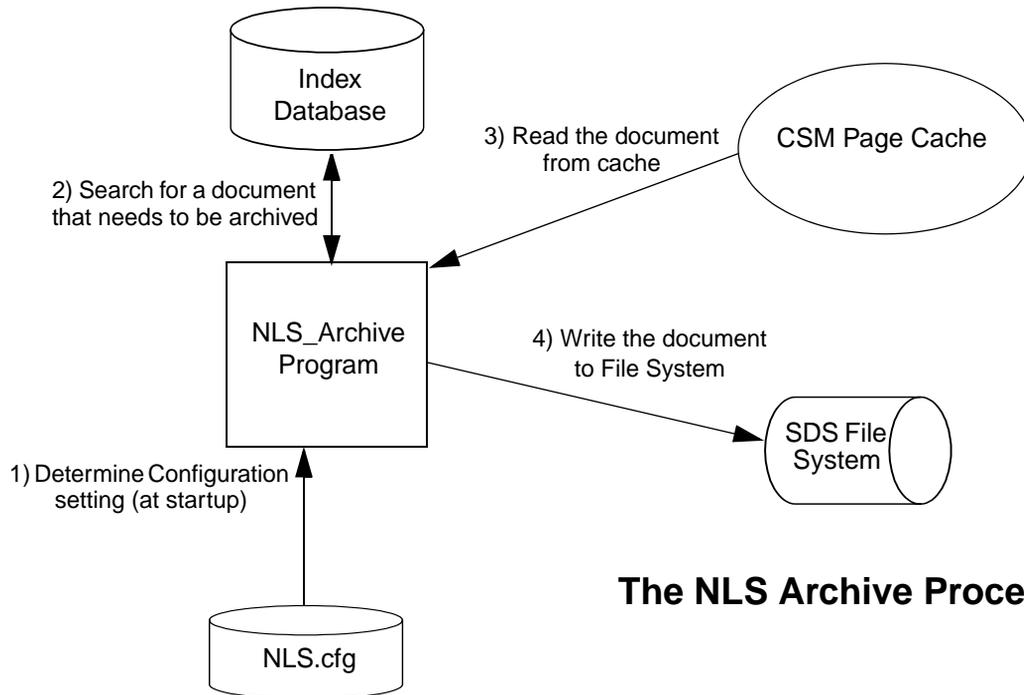
Be sure to also monitor your disk space carefully during production mode, because if your file system gets full, the NLS process stops archiving documents and shuts down.

Note If your WorkingDirectory approaches the maximum size for a physical volume on your SDS device, you might need to add another volume. (For example, the current maximum raw physical volume size for some storage devices is 16 TB. The effective storage is about 10 TB.)

When you add another physical volume, you must update the NLS.cfg file to reference the WorkingDirectory on the new volume. Although you can specify only one WorkingDirectory in the NLS.cfg file for NLS archiving, you can specify several Working Directories in the sds_conf file for SDS retrievals.

The following diagram illustrates the relationships between the major components of the NLS process for storing documents.

The NLS Archive Process

**The Archive Process**

Before archiving can begin, there are two requirements:

- Images must be entered into the system, indexed, and committed in the normal manner.
- NLS software must be installed and then configured for your specific storage device in the NLS.cfg file.

Refer to the numeric order in the diagram above for component interaction that the NLS Archive process normally follows to archive a document.

- 1 NLS reads configuration settings from the NLS.cfg file at system startup.
- 2 NLS searches for documents that are associated with document classes that have the FNP_ARCHIVE INX user index defined with a NULL value. (The Null value indicates that the document has not been archived yet.)
- 3 If the document is not already in cache, the NLS_Archive program submits a pre-fetch request for the document to be read into CSM page cache.

If the document is already in cache, the NLS_Archive program continues with the next step.

- 4 The NLS_Archive program reads the document from cache, forms a single document (object) file, and writes it to the SDS File System.

The single document (object) file includes a header portion and all pages in the document. After the document file has been written, the FNP_ARCHIVE index field is updated with the current date.

SDS Retrieval

Document retrievals are done directly from a file system using the SDS Worker, the retrieval component of the SDS software.

Note Only one SDS unit is supported in this SDS release, so the unit number in the sds_conf file **must** always be set to 1. For example, you

cannot configure both an SDS_FS_reader and SDS_CSAR_reader or SDS_ISAR_reader on the same system.

With SDS Retrieval, it is not necessary to integrate with a special client side library to access documents stored in an SDS File System. In addition, no changes to the client retrieval and display programs are required to access documents stored in the file system.

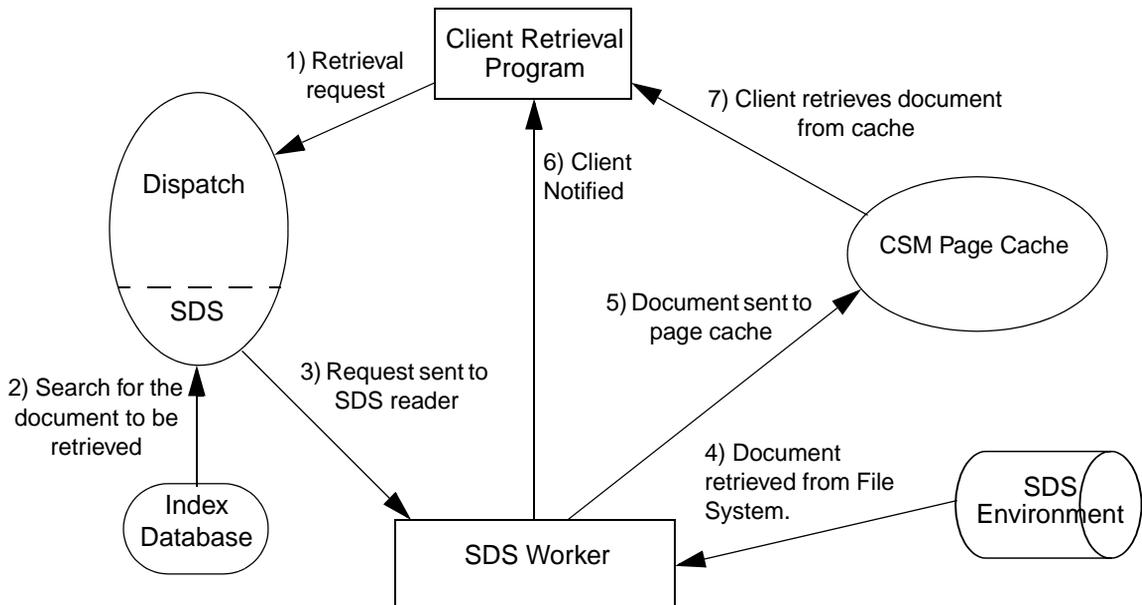
The Retrieval Process

The dispatch software component of the IS software has been enhanced to know how to process retrieval requests for documents located in an SDS File System, as well as documents located in CSM page cache, or on an MSAR or optical surface. Retrieval requests are processed based on the retrieval priority setting in the sds_conf file. For more information, see [**“Retrieval Priority \(sds_priority\)” on page 65.**](#)

When processing a retrieval request from an SDS File System, the request is sent to the SDS Worker (SDS_FS_reader) process. Since the SDS Worker process is integrated into the IS 4.0 (Service Pack 4 and later) software, the retrieval process is transparent to the user. Therefore, as long as the IS System server(s) are configured correctly, retrieval requests are made normally using the native Image Services software.

The following diagram illustrates the relationships between the major components of the SDS File System Retrieval Process.

The SDS File System Retrieval Process



The SDS File System Retrieval Process

Refer to the numeric labels in the diagram for the order of component interaction that the SDS Retrieval process normally follows to retrieve a document from a file system.

- 1** The retrieval request is sent to the Dispatch component. If the document is already in cache, there is no need to dispatch the request. The IS software indicates to the client that the document is in cache and retrieves the document from there.
- 2** If the document is not in cache, the dispatch process queries the Index Database's doctaba (FNP_ARCHIVE User Index) to determine if the document has been written to the SDS File System. Depending on the

SDS priority setting, and the state of the SDS, the dispatch process might dispatch the request to an MSAR or optical surface.

Note The SDS must be enabled to dispatch a request to the file system. If SDS is not enabled, the request is sent to an MSAR or optical surface.

- 3 Once the dispatch process determines that the request should be processed by the SDS Worker, the retrieval request is sent to the SDS Worker.
- 4 SDS Worker retrieves the document from the SDS File System.
- 5 SDS Worker sends the document to CSM Page Cache.
- 6 SDS Worker notifies the client that the document is in Page Cache.
- 7 The client retrieves the document from Page Cache.

Prerequisites

The following prerequisites are required to be in place before you install the NLS Archive software and configure your SDS system.

Image Services Software

The following release of Image Services is supported:

- **IS 4.1** or higher

Refer to the appropriate Installation or Update procedure for your platform:

AIX/6000

- [*Image Services Installation and Configuration Procedures for AIX/6000*](#)
- [*Image Services Update Procedure for AIX/6000*](#)

HP-UX (HP 9000 and HP Integrity)

- [*Image Services Installation and Configuration Procedures for HP-UX*](#)
- [*Image Services Update Procedure for HP-UX*](#)

Sun/Solaris

- [*Image Services Installation and Configuration Procedures for Solaris*](#)
- [*Image Services Update Procedure for Solaris*](#)

Windows Server

- [*Image Services Installation and Configuration Procedures for Windows Server*](#)
- [*Image Services Update Procedure for Windows Server*](#)

Image Services Toolkit Software

The following version of ISTK is supported:

- **ISTK 4.1** or higher

Image Services Toolkit (ISTK) must be installed before you install the NLS Archive software. For instructions on installing ISTK, refer to the Readme file located on the ISTK software package.

To check your ISTK version level, enter the following command for your platform:

UNIX`/fnsw/client/bin/fn_rel`**WIN**`<ISTK directory>\fnsw\client\bin\fn_rel`

NLS Archive Software

NLS Archive must be installed on the same server as the ISTK software.

However, while it might be common to install the NLS Archive and IS Toolkit software on the same server as Image Services, they do not need to coexist with IS and can, therefore, be located on a different server from Image Services.

NLS Archive Version

The following version of NLS Archive is supported:

- **NLS Archive 2.0.0** or higher

NLS Archive Disk Space

The file system disk space must be large enough to contain the documents that you anticipate will be stored, so plan accordingly. Directory disk space usage depends on the size of document files, volume of documents, and other factors such as file retention schedules.

Installing NLS Software

This chapter describes how to install the NLS software.

Installation by a FileNet Certified Professional (FCP) is recommended. For more information on the FCP program, refer to the IBM Web site (<http://www.ibm.com/us>), Products > Software > Information Management > Training and certification > Professional Certification > FileNet Certified Professional Program (FCP).

Important

Prior to installing the NLS software, ensure that Image Services Toolkit 4.0 (IS 4.0 SP5 edition) or later has been installed on the IS server. Image Services contains the SDS retrieval software.

Install NLS Software

To install the software on a **UNIX** platform, continue with the section **[“Install NLS Software on a UNIX Platform” on page 31.](#)**

To install the software on a **Windows** platform, skip to the section **[“Install NLS Software on a Windows Server” on page 34.](#)**

Install NLS Software on a UNIX Platform

This section describes the installation of NLS software for a UNIX platform.

The following procedure installs the NLS modules into the directory:
/fnsw/local/bin.

- 1 On the server where you are installing the NLS modules, logon as a user with **root** privileges.
- 2 Insert the NLS software CD into the CD-ROM drive.
- 3 Execute the appropriate mount command for your UNIX Operating System.

Note HP operating systems require a special mount command to read an ISO9660 CD type. If you have an HP system, continue with **Step 4**. For AIX or Sun systems, skip to **Step 5 on page 33**.

- 4 HP operating systems require a special mount command to read an ISO9660 CD type. Perform the following steps:
 - a To determine the CD-ROM device file name, enter the command:

ioscan -fnC disk
 - b Locate the CD-ROM device file name on the **ioscan** display. For example:

Class	I	H/W Path	Driver	S/W State	H/W Type	Description
=====						
.						
.						
disk	5	0/0/2/0.1.0	sdisk	CLAIMED	DEVICE	HP DVD-ROM 305
			/dev/dsk/c3t2d0		/dev/rdisk/c3t2d0	
.						

- c Mount the CD-ROM device on the /cdrom directory by entering a command similar to the following:

```
>mount -o cdcase /dev/dsk/c3t2d0 /cdrom
```

where **/dev/dsk/c3t2d0** is the CD-ROM device file name shown on the ioscan display.

- 5 Change directory to the CD-ROM mount point. This should place you at the root of the CD where the “ps_install” executable is located.
- 6 Run the installation program for the NLS Archive modules using the following command:

```
>csh ./ps_install
```

Regarding the above command, note the following:

- The “csh” shell specification in the command is needed if the user is not using the csh shell. If the user is already using the csh shell, including “csh” in the command does not cause any problems for the install.
 - There can be more than one csh (c shell) available on a UNIX system. The use of the csh (c shell) when executing the ps_install program can fail depending upon the default csh. Should there be a failure during installation it might become necessary to issue the ./ps_install command without specifying a shell. Another alternative would be to issue the command specifying the path to an alternate csh on the system.
- 7 After the installation is complete, the NLS Archive modules in /fns/local/bin can be moved to another directory, if desired.

- 8 You are now ready to configure the NLS software. Skip to **Chapter 3, “Configuring NLS Software,” on page 38.**

Install NLS Software on a Windows Server

This section describes how to install NLS Archive Modules on a Windows platform.

- 1 On the server where you are installing the NLS Archive modules, log on as a user with Administrator privileges.
- 2 Insert the NLS Archive software CD into the CD-ROM drive.

The Autorun program should automatically begin the installation. If Autorun is turned off, do one of the following to start the installer.

- a Open Windows Explorer, locate the PS_Install.exe file in the root directory and double-click the file to start it manually.
- b Select Start, Run on the Start Menu and enter the following command:

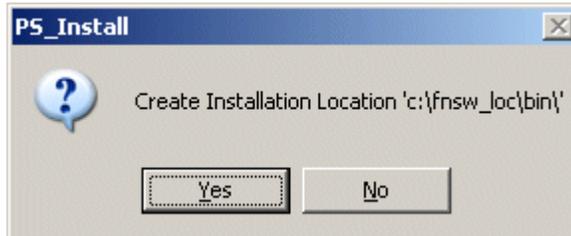
<CD-ROM drive letter>:PS_Install.exe

- 3 When the Application Package Installer dialog box appears, click the **Continue** button to proceed.

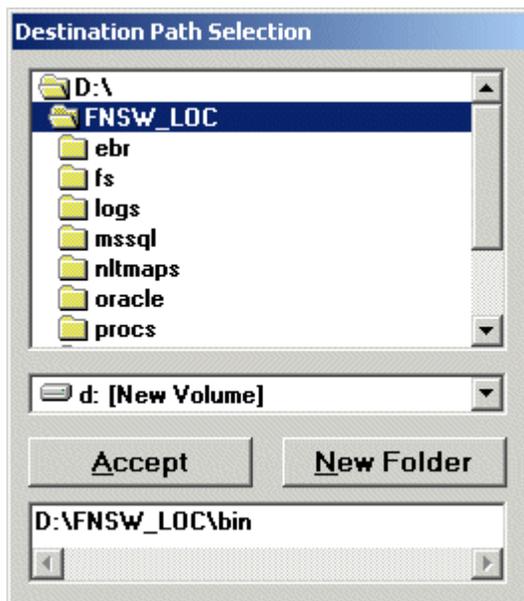
The Windows Installation Utility screen appears.

- 4 Select the WAL (IS Toolkit) version installed on your system from the Available Releases field, and click the **Install NLS** button.

The Create Installation Location dialog box appears.



- 5 Verify the install directory.
 - a If you want to create the installation location in the directory shown is this dialog, click **Yes** and continue to [Step 7 on page 37](#) .
 - b To specify a different installation location, click **No** and continue to [Step 6](#) below.
- 6 If you clicked the **No** button in the above step, the Destination Path Selection dialog box will open.



- a Select or create the destination path to the new installation location and click **Accept**.

Note

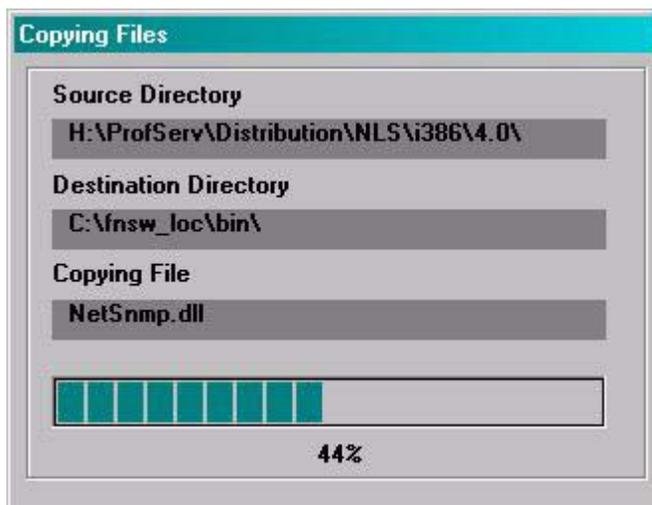
A standard installation directory for the NLS Archive modules is:
<drive>: \fnsw_loc\bin

Important

DO NOT SELECT a path with spaces in the path name, for example "\Program Files". Path names with spaces will not work with NLS Archive modules.

- b Click **Yes** to confirm the new installation directory location for the NLS files.

The Copying Files dialog box appears.



- 7 The Copying Files dialog box displays the installation progress. When the installation has completed click **OK**.
- 8 After the installation is complete, the NLS Archive modules in <drive>:\fnsw_loc\bin can be moved to another directory, if desired.
- 9 Click **Quit** to close the Windows Installation Utility window.

3

Configuring NLS Software

This chapter describes how to configure the NLS software.

NLS Archive Prerequisites

The NLS Archive component uses the IS Toolkit software and needs to know how to access the IS system it is archiving documents from. Although it can run on another server, for performance reasons, it typically runs on the IS server.

Complete the following sections to configure the NLS Archive software.

Update the System Hosts File

The hosts file can be found in the following locations on your server:

UNIX

/etc/hosts

WIN

<drive>:\WINNT\system32\drivers\etc\hosts

IS Toolkit applications require specific entries in the “hosts” file in order to correctly logon through IS Toolkit. This entry is specific to IS Toolkit applications and must be added to the “hosts” file before running a FileNet utility.

The FileNet domain name entry is case-sensitive, and must appear exactly as it appears in the FileNet Application Executive.

The TCP/IP address is the first entry and is followed by a tab. The server name is next, followed by a space. The FileNet domain name is the third entry, followed by another space. The last entry is the four-part FileNet nch server name, which contains the domain name, the organization name, and the string “-nch-server”.

For example:

```
192.48.11.12 server domain domain-orgname-nch-server
```

If your system uses IPv6 addresses, the entry might look like this:

```
fe80::fefe:0927:2638 server domain domain-orgname-nch-server
```

Environment Path

The PATH in the user environment (usually fns user) must be setup so that the IS Toolkit shared libraries are accessed by an IS Toolkit call, and not the IS shared libraries. This is explained in the IS Toolkit installation readme file, and is also covered in this document in an effort to emphasize and include the utility PATH location.

The configuration of the PATH is different for UNIX and Windows Servers, and it is important not to confuse the two.

The examples below show the IS Toolkit vs. IS PATH hierarchy.

UNIX

To configure the PATH in UNIX:

```
/fns/client/bin:/fns/local/bin:/fns/bin
```

WIN

To configure the PATH in Windows:

```
<drive>:\fnsw\client\bin;<drive>:\fnsw\client\shobj;  
<drive>:\fnsw_loc\bin;<drive>:\fnsw\bin;<drive>:\fnsw\lib\shobj;
```

The PATH should not be limited to the examples above, as there are other software and OS requirements found in each environment PATH that must be retained. In addition, if the NLS software has been moved from /fnsw/local/bin, or \fnsw_loc\bin, to another directory, this directory must be included in the PATH.

IS Toolkit File Ownership and Permissions on UNIX

FileNet Image Services Toolkit provides installation instructions, which outlines the UNIX environment permissions and ownership settings on the IS Toolkit modules. These instructions **must** be followed to ensure correct access of IS Toolkit modules.

However, the IS Toolkit modules **must** be owned by fnsw:fnusr. There are two modules that are owned by root:fnusr. Refer to the *IS Toolkit Installation Instructions* for complete instructions.

Creating the NLS Configuration File

The NLS.cfg file is a text file used to configure the NLS Archive product. This file includes logon attributes such as the FileNet user name and password, and the Image Services domain name and organization.

Also included in the NLS.cfg file are program attributes such as the NLS working directory where documents are stored. In addition, there are tuning parameters and other specific configuration parameters.

NLS Configuration Sample File (NLS_cfg.sample)

Below is an example of the NLS_cfg.sample file. This is the file you will edit in the next section to create the NLS.cfg file for your system. Since only NLS is used for archiving, some of the settings shown in the sample file are not used.

Note The example shows directory paths for a UNIX directory. For a Windows system, the directory path entry would follow the correct Windows directory convention.

Also, blank lines are not valid in the NLS.cfg file, and a line beginning with a semi-colon (;) is a comment line that is ignored by the program.

NLS_cfg.sample file

```
*****
; This module contains the configuration information for the XXX.
*
*****
; $Author: giffj $
; $Date: 10/.0/.2 .1:.2:.4 $
; $Header: NLS_cfg.sample,v 1.4 10/.0/.2 .1:.2:.4 giffj Exp $
; $Locker: $
; $Revision: 1.4 $
; $Source: /usr/ProfServ/ISR_200/NLS/src/RCS/NLS_cfg.sample,v $
; $State: Exp $
*****
; INSTALL HISTORY:
;          INSTALLED BY          DATE          DESCRIPTION
;-----
;          ----- -MM/DD/YYYY- -----
;
; (continued on next page)
```

```
(continued from previous page)
;*****
;   Logon Information
;*****
LogonAttribute {
  UserName="SysAdmin"
  PassWord="SysAdmin"
  Domain="serverA"
  Organization="YourCo"
}
;   General information
PgmAttribute {
;   WorkingDirectory=/NLS
;   LogDirectory={Defaults HOME}
;   Control for set size
;   MaxFetchDocs={Default=1000}
;   MaxQueryDocs={Default=1000}
;   MaxFetchChildren controls the maximum number of concurrent Add2Q
processes
;   MaxFetchChildren={Default=3}
;
;   MinDocId=(Default=100000)
;   MaxDocID=(Default=399999999)
;   Walkback=(Default=100000)
;   AppMaxDocID=(Default=399999999)
;   MinDate=(Default=980101)
;
;   CleanCache=(Default=False)
;   Timing=(Default=False)
;
(continued on next page)
```

```
(continued from previous page)
; The following time based keywords all have a minimum of 1 seconds
; and a maximum of 3600 seconds. If set to zero or a negative number
; it will default to 5 seconds.
;   FetchSleep=(in sec, default=600)
;   ArchSleep=(in sec, default=600)
;   FetchTimeout=(in sec, default=3600)
;
;   NOTE: Values that are out of range take on max value for that item
;
; Display format for the year %Y (default) is 'xxxx', %y is 'xx'.
; This effects displaying of a date only and extends the length of
; file names.
;   YearFmt="%Y"
;   Format is Year, month, day. Year can be 2(19xx) or 4 digits
;   NOTE: Values that are out of range take on max value for that item
;
; ; SDS-Keywords
;   Provides the type of interface for SDS
;   Centera = EMC/Centera Protected Storage
;   Tivoli   = IBM/Tivoli Storage Manager
;   Snaplock = NetApps-style Protect Storage File System
;             (SnapLock,StorageTek-Trinity,PermaBit, etc...)
;   Filesystem = Non-Protected Storage File System
;             (StorageTek-ASM)
;   SdsDevice = {Default="Filesystem"}
;
;   Provides the DNS and/or IP information for SDS device.
;   SdsDomain="central,212.3.248.36"
;   Sets the buffer size(k) for IO to SDS device.
;   SdsBufSize={Default=10240}
(continued on next page)
```

```
(continued from previous page)
;   Sets the network timeout value to SDS device, in seconds.
;   SdsTimeout={Default=120}
;   Sets the number of retries to SDS device.
;   SdsRetry={Default=6}
;   Sets the length of sleep between retries, in seconds.
;   SdsSleep={Default=120}
;   Sets the Retention policy for FileNet objects where 0 or No
Retention
;   is the Default.
;   0 = No Retention mode (actually set to 5 mins)
;   -1 = Infinite Retention mode
;   1 or greater than 1 = set retention using input value in hours
;   NOTE: 875520 = 100 yrs (100 yrs * 12 mths * 30.4 dys * 24 hrs)
;   See Documentation for details
;   SdsRetention={Default=0}
;   OR a date is valid in the form of "yyyy/mm/dd" or "mm/dd/yyyy"
;   SdsRetention="2005/08/23"
;   Sets the disposition of the FileNet object (a FileNet
;   document with index values) in the Working directory
;   as well as the location of the clip id information
;   SdsDisposition={Default="FS"}
}
;END OF CONFIGURATION FILE
```

Configuring the NLS.cfg file

- 1 Locate the sample NLS configuration file (NLS_cfg.sample) in the directory where the NLS software is installed.

UNIX

/fnsw/local/bin

WIN

<drive>:\fnsw_loc\bin

- 2 Make a copy of the NLS_cfg.sample file by saving it as NLS.cfg in the directory where you are running the NLS Archive program.
- 3 Edit the new NLS.cfg file using any text editor, making configuration changes appropriate for your system. See the appropriate appendix in this document for additional information. You can accept most of the defaults in NLS.cfg file, but at a minimum you must set at least the following parameters:
 - **LogonAttribute** (UserName, Password, Domain, Organization)
 - **PgmAttribute** (WorkingDirectory, LogDirectory)

Logon Attributes

Several logon attributes in the LogonAttribute section of the NLS.cfg file must be configured for the FileNet IS user who is running the NLS_ archive program. These attributes must match the user name and password attributes that have already been configured on the IS server for this user.

For more information, see [“Device-specific Settings” on page 60](#) in the “Configuring Shared Resources” section.

```
LogonAttribute {
  UserName="SysAdmin"
  PassWord="SysAdmin"
  Domain="serverA"
  Organization="YourCo"
}
```

There are two ways to enter logon attribute information.

- The System Administrator can enter the values for the UserName and PassWord via ascii text form. If this method is used, the password for the user is visible to anyone with read access to the NLS.cfg file.
- The login attributes can be configured using the optional Encrypted Password functionality. To use the Encrypted Password method, follow the instructions below.

Using the optional Encrypted Password Functionality (PS_Password Program)

The PS_Password program can be used to create the “.ps_passwd” file which contains the encrypted records. The first time the PS_Password program is run, you need to provide a password for the creation of the “.ps_passwd” file.

Important

You **must** remember the .ps_passwd file password or you will not be able to run the PS_Password program a second time to make changes or additions to the .ps_passwd file.

If you lose or forget the password, the .ps_passwd file has to be deleted and recreated via the PS_Password program Initialization Mode.

- 1 Login as **fns**w and run the **PS_Password** program:
 - a For UNIX systems enter:

```
cd /fnsw/local/bin
```

```
PS_Password
```

- b For Windows 2000 Server, enter the following at a Command Prompt:

```
cd <drive>:\fnsw_loc\bin
```

PS_Password

The first time you run the PS_Password program, you are notified that it is running in “Initialization Mode”.

- 2 When prompted, enter and confirm the password for the Password Application Key.

CAUTION The character limit for this password is 8 characters and **cannot** be changed once it has been set.

- 3 When you are prompted to enter an application name, enter NLS.
- 4 Enter a Password for the login being used, and the FileNet user name for that login.

Note You must provide a valid FileNet user name that has SysAdmin privileges within the IS domain. See the section, **“Device-specific Settings” on page 60** for more information.

Once the application name, password, and user name have been entered, the program displays the application and user name with validation that the password has been set.

- 5 When the “Enter PS Application Name” prompt appears again, select **Enter** to discontinue the PS_Password program. A list of all applications and user names entered is displayed.

The following is an example on a UNIX system:

```

/fnsw/local/bin> PS_Password
Creating PS Password File - Initialization Mode
Enter Password Application Key      :
ReEnter Password Application Key    :
Enter PS Application Name           : NLS
Application (39) 'NLS'
Enter New Password for Application NLS      :
ReEnter Password for Application NLS      :
Enter User Name                     : SysAdmin
PS Program 'NLS': User=SysAdmin, Password Set
Enter PS Application Name           :
**** Current Applications Set ****
** Application (39) 'NLS' [User='SysAdmin']
*****

```

Configuring UserName and PassWord in NLS.cfg File

After the PS_Password program has been run and the .ps_passwd file has been created, you can configure the UserName and PassWord values in the NLS.cfg file.

The UserName value in the configuration file should be set to the application name entered in the .ps_passwd file. In this case, UserName="NLS".

The PassWord value in the LoginAttribute section of the NLS.cfg file should be removed or commented out. If the application sees the PassWord parameter in the NLS.cfg file during start-up, it does **not** use the encrypted password that was entered when the PS_Password program was run, and the NLS Archive program does not run.

The following is an example of the NLS.cfg LogonAttribute section when using the encrypted password functionality:

```

;*****
;
;          LOGON INFORMATION
;*****
;NLS.cfg.sample
LogonAttribute {
    UserName="NLS"
;    PassWord="" <this parameter has been commented out>
; For a source system the Domain should be the source domain
; For any other system (target or 3rd party) the Domain
; should be the target target domain
    Domain="serverA"
    Organization="YourCo"
}

```

Program Attributes

Directory attributes in the PgmAttribute section of the NLS.cfg file must be defined.

WorkingDirectory

The SDS File System (**WorkingDirectory**) must be replaced with a directory structure that fits your environment. This is done by editing the **WorkingDirectory** field in the NLS.cfg file.

The WorkingDirectory is the directory that specifies the top-level path where documents are stored in the SDS File System. A directory is created under this path using the System Serial Number for the

Domain, and all document files are stored in this directory. For more information on the WorkingDirectory, see [“The File System Working Directory” on page 22.](#)

LogDirectory

The LogDirectory, specified in the NLS.cfg file, is where journal entries are made which contain all the log files created each time the NLS Archive modules are run.

Note If not specified in the NLS.cfg file, the LogDirectory is created under the current directory (default **/fnsw/local/bin/journals** for UNIX, or **<drive>:\fnsw_loc\bin\journals** for Windows).

The file format is:

arc<date> for NLS_Archive logs
ads<date> for SDS_writer logs

Note that only if NLS.cfg/Max ArchChildren>0 are SDS journal files created.

These are ordinary text files that can be read when the program is operating in a UNIX environment.

Windows 2000 might, or might not, allow access to the files depending upon whether the log is locked by the utility.

Other Attributes

The following attributes can be left at their default values or adjusted according to your system requirements.

AppMaxDocId

This configuration technique should only be attempted with the assistance of your service representative.

The AppMaxDocID is a setting used when the environment calls for two different instances of NLS_Archive to be run against two different document id ranges. This environment can occur when peer systems are running Image Services Multi-System Committal (MSC) with two different document id ranges being saved to the same service. NLS_Archive archives documents in the low range according to this setting and then looks for more work as it moves forward through the specified range.

NLS_Archive continues to run until it receives an NLS_stop command. As NLS_Archive runs, it archives documents until it has no more to archive and then it sleeps. When NLS_Archive awakens, it begins looking for more documents from the last document ID saved. It is important to note that NLS_Archive never goes back to the beginning of the range while it continues to run. The program only moves forward.

When running two instances of NLS_Archive, set this value to be the highest document id in the low document id set range.

The value can also be overridden by specifying it on the command line:

-app <doc id>

ArchSleep - default = 600 sec. (10 min.)

The minimum period between the start of a set and the start of the next set for archiving documents.

CleanCache - default = False

During archiving, if CleanCache is set to True and the document in cache is ageable, the document is deleted after the NLS Archive software has archived it.

Note

Do not set CleanCache to True if you are using migration delay on the document class you are archiving to a file system repository, such as SnapLock. When CleanCache is set to true for file system repositories, the NLS Archive software does not take migration delay into account.

DDEXim - default = True

When set to True, this attribute creates a text description of the Data Dictionary for backup. The default is True. For more details on DDEXim see the DDEXim section of the *IS System Tools Reference Manual*.

Important

The DDEXim feature has been known to cause the NLS modules to hang on some systems when NLS modules are first started. Setting the value to False solves this problem.

FetchSleep

This attribute allows you to bypass the default sleep periods in an effort to better tune system performance. The default sleep time is automatically adjusted to 1 second if the set size is set at 1 in the MaxFetchDocs keyword value. The sleep value automatically adjusts for sets of 10 items or less with the maximum sleep value to be automatically set at 10 seconds.

The standard set size for most production systems is 999 (assuming the system has resources to handle such a load). If the FetchSleep keyword is not set then the default sleep value between sets is automatically adjusted to a maximum of 600 seconds (10 minutes). Systems with extremely large cache space might not need to wait 10

minutes and systems with extremely small cache space might need to wait more than 10 minutes. When setting FetchSleep, the minimum setting is 0 (zero) and the maximum is 21,600 (6 hours).

The default is 600 seconds (10 minutes) when the set size is not set to a figure between 1 and 10.

FetchTimeout - default = 3600 sec. (60 min.)

This attribute is the maximum time that the NLS_Add2Q waits for a document to be available. It is only active in a UNIX Operating System environment.

IgnoreChksum - default = False

The ability to bypass checksumming should only be used under the direction of FileNet Certified Profession (FCP). The omission of this keyword value could cause corrupted documents to be imported.

MaxDocId

This attribute is used to tell NLS_Archive to exit when it reaches this value or to exit when it reaches the last doc_id on the system that is under this value. The use of this setting is just like AppMaxDocid but instead of sleeping and awakening, the NLS_Archive program stops when it reaches the last document id which is still under the set value.

Should both MaxDocId and AppMaxDocId be set, the Archiver ignores MaxDocId and performs the task using AppMaxDocId functionality.

The value can also be overridden by specifying it on the command line:

-max <doc id>

MaxArchChildren - default = 3/range 0-64

This attribute allows the SDS code to do multi-processing, potentially increasing the throughput of NLS Archiving. The MaxArchChildren

value is the number of concurrent SDS_Write (child) processes that NLS starts up. Note that only if MaxArchChildren>0 are sds<date> journal files created.

MaxFetchChildren

This attribute specifies the maximum number of simultaneous document reads from the WorkingDirectory. When a Tape Silo is used, this should be set no higher than the number of tape drives that are configured for reading. If a simple file system is used, this should initially be set to eight (8), the maximum value. If after running, it is found to be placing too much load on the system, this value should be reduced until system performance is acceptable and this number is as high as possible.

MaxFetchDocs - default and max. = 999

This attribute is the maximum number of documents processed as a set during retrieval and archiving.

Note

If both MaxFetchDocs and MaxQueryDocs are set, the lower of the two values is used for the set size.

MaxQueryDocs

The maximum number of documents processed as a set during archive in Database Mode.

MaxQueryDocs affects the Archive set size and the NLS Archive sleep period in the following manner:

- If the Set Size in the MaxQueryDocs value is set to 30 or more, the sleep value is internally set to a maximum of 30 seconds.
- If the MaxQueryDocs value is set to be less than 30 seconds, this default internal parameter is adjusted to the same value as the

MaxQueryDocs value. Therefore, if the MaxQueryDocs value is set to 16, the internal sleep value is adjusted to 16 seconds.

To illustrate this functionality, if the **ArchSleep** value is set to 10 seconds and the **MaxQueryDocs** value contains a set size set to 15, then the sleep time between sets is a minimum of 25 seconds, depending upon when the previous work was completed. If the previous work time exceeds the sleep interval, then the overlap is subtracted from the sleep value, and the next set could begin processing in less than the 25 seconds of default and configured sleep.

Note The NLS Archive modules contain an additional internal default sleep function that runs in addition to any other sleep value configured. This internal sleep provides for a 30 second or less sleep interval that is necessary for internal housekeeping. There are three different set size values that can be configured in the NLS.cfg file.

- MaxQueryDocs affects the Archive set size
 - MaxFetchChildren affects the Dispatch set size
 - MaxFetchDocs affects the Retrieval and Archive set size
-

DocReportFrequency - default = 999/min. = 2/suggest = 250

This attribute is the frequency with which entries are made into the journal files. The value of this attribute represents the number of documents between reports. This can be used as a progress indicator. The journal log entry enters a batch status, number of documents in the batch, and the memory used to process the batch.

For NLS Archive, if DocReportFrequency is greater than MaxQueryDocs, then DocReportFrequency is set to MaxQueryDocs. In other words, DocReportFrequency cannot be greater than MaxQueryDocs.

MinDate

The minimum date of an Image Services document processed during archiving in Database Mode.

MaxErrorsAllowed - default and max. = 65535/min. = 1

This attribute sets the number of errors that are allowed within the period specified by ErrorDelay. If this number is exceeded, the program will stop. The maximum and default is 65535; the minimum is 1.

ErrorDelay - default = 5/max. = 10

This attribute sets the window of time (in seconds) for which MaxErrorsAllowed measures its maximum errors. The minimum is 5 seconds. A setting of zero defaults to 5 seconds.

MinDocId

The minimum Image Services document number processed during archiving when in Database Mode.

The value can also be overridden by specifying it on the command line:

-min <doc id>

OSNice

This is the minimum period between sets. This causes the NLS modules to sleep and allows other processes to get processing time. This parameter is useful when the NLS modules are run during normal production time.

Primary - default = True

This setting only affects the directory/folder structure created in the WorkingDirectory.

When set to FALSE, it allows for the backup structure to be changed from the Primary Surface ID (surface_id_1) to the Secondary Surface ID (surface_id_2) (tranlog).

Primary=True causes a WorkingDirectory structure to be created starting with the FileNet Domain ssn. The surface_id_1 number is changed into a hexadecimal and split up into a four folder structured path. The balance of the folder structure consists of the actual document ID number being changed into a hexadecimal, and split up into a two folder structure.

Primary=False causes a WorkingDirectory structure to be created starting with the FileNet Domain ssn. The surface_id_2 number is used in decimal format as the next folder under the ssn folder. The following folder structuring is created by splitting up the hexadecimal form of the actual document ID into 2 more folders.

If the Secondary Surface ID (surface_id_2) is zero (0), there is no tranlog configured on the system. In this situation, the folder structure of the backup is still in hexadecimal just as the Primary=True (default), but the structure is not as deep because the starting point is zero (0). The last two folders in the structure are still a hexadecimal form of the document ID.

It is important to keep the Primary setting consistent, as the setting MUST be the same in the NLS.cfg file when retrieving from an Archive. Changing this configuration keyword after initial installation causes retrievals to fail.

Timing - default = false

When set to true, the NLS Archive modules write timing information to the log files. Normally Timing is set to false during production.

Timing logging output to the journal logs is used in conjunction with the `timer.awk` script for performance analysis. The command below is used with `timer.awk` and the journal file to display various timing figures to the screen.

- **For UNIX:** `awk -f timer.awk journals/arc20010328`
- **For Windows:** `awk -f timer.awk journals\arc20010328`
(must have `awk` for Windows)

WalkBack - default = 100,000

The NLS Archive modules use an automatic walk back algorithm of 100,000 to ensure that the starting point for the day's document backup does not miss any documents. In specialized environments there can be a break in new document IDs that exceeds this default setting of 100,000. This can occur when there is more than one document entry device, and they each have their own unique document ID numbering scheme. It can also occur when throughput is very large, and there is a Storage Library migration delay that could present a situation where documents are not migrated in a timely manner.

There might also be situations where no walk back is necessary and you have determined there is no risk of missing images. In this situation, the `WalkBack` value can be set to a number less than 100,000 to save `NLS_Archive` processing time. Your FCP can assist you in a review of your processes in an effort to determine if it is necessary to configure a custom `WalkBack` value for your site.

The default or configured `WalkBack` uses the `.nlsrc` file for a logical starting point at the startup of `NLS_Archive`. This `.nlsrc` file is updated with new surface and document id information at the end of each successful set as `NLS_Archive` is running.

Example:

Setting the **WalkBack** equal to 1 million causes the NLS Archive modules to begin at the document ID that is 1 million back from the last entry in the **.nlsrc** file.

```
WalkBack=1,000,000
```

If the resulting Document ID returned is less than the valid minimum document ID on the system, the NLS modules function as if this is running for the first time.

YearFmt - default = %Y

This is the display format for the year portion of the date that is appended to the NLS working files. This parameter can change the length of a file name. The default is '%Y' which displays in the 4-digit format "yyyy". A 2-digit "yy" format can be displayed by changing the default to '%y'.

Saving the NLS.cfg file

Save and close the NLS.cfg file to keep all of the changes you have made.

Device-specific Settings

For NLS.cfg parameters for specific devices, refer to the appropriate appendix:

- [“Appendix C – Centera Storage and Retrieval” on page 129](#)
- [“Appendix D – SnapLock Storage and Retrieval” on page 151](#)
- [“Appendix E – IBM DR550” on page 160](#)
- [“Appendix F – Sun StorageTek 5320 Storage and Retrieval” on page 185](#)

Configuring SDS File System Retrieval

This chapter describes how to configure the SDS retrieval software, as well as how to configure shared resources.

Creating the SDS Configuration File

The SDS_FS_reader (SDS Worker) can be configured to suit the specific needs of your system.

To configure SDS Worker, you need to manually create an sds_conf file. Using your preferred text editor (such as **vi**), create the sds_conf file in the following directory.

UNIX

`/fnsw/local/sd/1`

WIN

`<drive>:\fnsw_loc\sd\1`

Note

After creating, or editing, the sds_conf file, the **IS software must be recycled** for changes in the sds_conf file to take effect.

The `sds_conf` file definition format is as follows:

```
SDS <SDS Unit ID> {  
  
    sds_name    "<SDS label name>"  
    worker      <SDS worker executable> <# of worker processes>  
    info        "<SDS storage specific information>"  
    sds_priority [ low | medium | high ]  
    debug [ off | min | max ]  
}
```

Note

Only one SDS unit is supported in this SDS release, so the unit number in the `sds_conf` file **must** always be set to 1. For example, you cannot configure both an `SDS_FS_reader` and `SDS_CSAR_reader` or `SDS_ISAR_reader` on the same system.

Below is an example of how the `sds_conf` file might look for a NetApp SnapLock device:

```
SDS 1 {  
    sds_name      NLS SDS  
    worker        SDS_FS_reader 20  
    info          "WorkingDirectory=/fs"  
    sds_priority  medium  
    debug         max  
}
```

sds_conf File Format

SDS Name (sds_name)

This field is not currently used. The "sds_name" field might be used in the future to denote the SDS type when the SDS Control UI window is implemented.

Worker Program Name (worker)

This field indicates the worker program name and the number of worker instances to be started by the ds_init process.

Choose the appropriate worker name for your storage device:

- **SDS_FS_reader** - NetApp Snaplock
- **SDS_ISAR_reader** - IBM DR550
- **SDS_CSAR_reader** - EMC Centera
- **SDS_SunSAR_reader** - Sun StorageTek 5329

One (1) to ninety-nine (99) workers are supported. Twenty (20) is an arbitrary number of workers to start with. You can adjust this number up or down in the future as you monitor system performance.

SDS_FS_reader Information Field (info)

The **info** field indicates the storage specific information for the reader. It includes the working directory path and the IP address of the domain.

The **WorkingDirectory** path contains the following information.

```
"WorkingDirectory=<root directory name>/fs"
```

If the SDS Worker is an SDS_CSAR_reader or SDS_ISAR_reader, the field contains the following information:

```
"WorkingDirectory=<root directory name> Domain=[<domain name>,<IP>]"
```

Note Be sure to use double-quotes in the **info** field entry.

WorkingDirectory is an absolute path to the SDS File System.

Note You can specify more than one file system for SDS retrievals. For example, the current maximum raw physical volume size for some storage devices is 16 TB, so you might need to add more volumes to create several 16 TB file systems to accommodate all your data. The info statement in your sds_conf file might look like this:

```
info "workingdirectory=/fs4;/fs3;/fs2;/fs1 domain=..."
```

The file systems are searched in the order you list them in the info statement, so the active NLS WorkingDirectory should be listed first, then the others in order by age. The entire info line can be 4096 characters long. There should be no spaces between the file system names and no carriage returns—just semi-colons.

Domain is a comma-separated string containing one or more addresses of the available Access Nodes of the pool.

Note There are no spaces in the Domain string (such as, Domain=[<IP-address>]).

Retrieval Priority (`sds_priority`)

The `sds_priority` field indicates retrieval precedence. SDS priority can be set to, Low, Medium, or High.

- **Low** - requests first go to MSAR, then optical, and finally SDS.
- **Medium** - requests first go to MSAR, then SDS, and finally optical.
- **High** - requests first go to SDS, then MSAR, and finally optical.

Note Retrieval priority only takes effect if the requested document is not in CSM page cache. If the document is already in cache, there is no need to migrate it from the SDS file system, MSAR, or optical.

Low Priority

If priority is set to low, the program determines if the document is on an “in box” MSAR that is read enabled. If this is the case, the request is directed to the “in box” MSAR surface.

If the associated MSAR surface is not read enabled or the surface is “out of box”, the program determines if there is an “in box” optical surface that can be used. If there is, the request is directed to the optical surface.

If the associated optical surface is not read enabled or the surface is “out of box”, the program determines if the document has been stored in an SDS File System by looking at the `FNP_ARCHIVE INX` user index value. If SDS has not been configured on the system, or the document has not been stored in an SDS File System, the `SDS_not_configured` or `SDS_document_not_archived` error tuple is returned.

If the document is not found in the SDS File System, or "in box" surfaces, an RSVP is generated and Document Services returns a ASE_INTERVENTION_REQ or ASE_NOT_IN_OSAR status to the client retrieval program.

Medium Priority

If priority is set to medium, the program determines if the document is on an "in box" **MSAR** surface that is read enabled. If this is the case, the request is directed to the "in box" MSAR surface.

If the associated MSAR surface is not read enabled, or the surface is "out of box", the program determines if the document is stored in an SDS File System by looking at the FNP_ARCHIVE INX user index value. If it is, the request is directed to the SDS File System.

If the document is not found in the SDS File System, the program determines if there is an "in box" **optical** surface that can be used. If there is, the request is directed to the optical surface.

If there is no MSAR, SDS File System, or optical surface available, an RSVP is generated and Document Services returns a ASE_INTERVENTION_REQ or ASE_NOT_IN_OSAR status to the client retrieval program.

High Priority

If priority is set to high, the program determines if the document is stored in an SDS File System by looking at the FNP_ARCHIVE INX user index value. If it is, the request is directed to the SDS File System.

If the document is not found, the program determines if an “in box” MSAR or optical surface can be used. If this is the case, the request is directed to the MSAR or optical surface.

Important

Centera Users only: With priority set to high, if the document cannot be found because of an invalid index value (FNP_CLIPID) or a malfunctioning Centera device, you get a message similar to “Retrieval to cache failed” and the retrieval ceases before attempting to retrieve documents from MSAR or optical as expected. In this scenario, the only way to get the retrieval to continue past the Centera device is to disable the SDS device through the **DOC_tool** command.

If there is no SDS File System, MSAR, or optical surface available, an RSVP is generated and Document Services returns a ASE_INTERVENTION_REQ or ASE_NOT_IN_OSAR status to the client retrieval program.

Activity Reporting (debug)

There are three debug settings: “off” “min” and “max.” Activity is reported in the view log.

- **off** - when debug is set to off, no SDS_FS_retrieval_worker activity is reported in the view log.
- **min** - the SDS_FS_retrieval_worker logs a sys_log message indicating that a document has been retrieved from the SDS file system.
- **max** - the SDS_FS_retrieval_worker 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 by the SDS abstract (SDS_request_done) when

the request is complete. For a production system, setting the debug flag to min or max might have a negative performance impact.

Multiple IS Storage Library Servers

If you have multiple Storage Library servers, you can use the same `sds_conf` file on each storage library server. Make sure the `sds_conf` file is copied to the following directory on each storage library server:

For UNIX: The `sds_conf` file in `/fnsw/local/sd/1`

For Windows: The `sds_conf` file in `<drive>:\fnsw_loc\sd\1`

In some cases, you must modify the **WorkingDirectory** field in the `sds_conf` file so all servers point to the same target SDS File System.

Note The shared SDS File System can be mapped to different local names on each server.

Configuring Shared Resources

This section describes configuration settings that must be consistent for both NLS Archive and Image Services SDS.

NLS.cfg and `sds_conf` Settings

The `WorkingDirectory` (NLS File System) name must be the same in both `NLS.cfg` and `sds_conf` files. The NLS Archive process also must have read and write access to this directory. In addition, the Image Services **fnsw** user must have, at a minimum, read access rights to the SDS file system.

Note When specifying the path for the shared WorkingDirectory, only use UNC format. **Do Not** use a drive letter in the path. Since this is a remote shared drive, you **must** use the UNC path.
For example: \\<servername>\sharename

Device-specific Settings

For sds_conf parameters for specific devices, refer to the appropriate appendix:

- [“Appendix C – Centera Storage and Retrieval” on page 129](#)
- [“Appendix D – SnapLock Storage and Retrieval” on page 151](#)
- [“Appendix E – IBM DR550” on page 160](#)
- [“Appendix F – Sun StorageTek 5320 Storage and Retrieval” on page 185](#)

Administering and Running NLS

This chapter describes procedures for administering and running the NLS_Archive program.

Running the NLS_Archive Function

The NLS Archive program is started with **NLS_start**, and shut down with **NLS_stop**. NLS_Archive can be started/stopped using one of three methods:

- Automatically when Image Services software starts, or stops, see [“Autostart/Autostop NLS Method” on page 70](#).
- Manually from a command line, see [“Command Line Startup/Shutdown Method” on page 72](#).
- Using Windows Service, see [“Windows Service Setup Method” on page 75](#).

Autostart/Autostop NLS Method

Use this procedure to allow the NLS Archive software to automatically start and stop, in database mode, when the Image Services software starts and stops. If the NLS Archive software is not running directly on the Image Services Server, this autostart procedure cannot be used.

- 1 Create an `ims_start` and `ims_stop` file for UNIX, or `ims_start.bat` and `ims_stop.bat` for Windows, in the following directories:

UNIX`/fnsw/local/sd`**WIN**`<drive>:\fnsw_loc\sd`

Note For UNIX, the `ims_start` and `ims_stop` file names must be created in “lower case” and have permissions set to execute.

- 2 To cause NLS Archiving to automatically begin, in database mode, when Image Services is started, edit the `ims_start`, or `ims_start.bat` file, by adding the following to the file:

Note This causes Image Services to look for a file named `ims_start` after the IS comes up. This file can be implemented to contain scripting to start other applications dependent upon Image Services.

UNIX

For UNIX:

```
cd /fnsw/local/bin
NLS_start
```

WIN

For Windows:

```
cd <drive>:\fnsw_loc\bin
NLS_start
```

Note NLS_Archive will not start successfully if there are any problems logging on to the IS. This can happen if Image Services is slow to start. If this happens, a “sleep” for a few seconds can be added to facilitate proper timing.

- 3 To cause NLS Archiving to automatically stop when Image Services is stopped, edit the `ims_stop`, or `ims_stop.bat` file by adding the following to the file:

UNIX

For UNIX:

```
cd /fnsw/local/bin
NLS_stop
```

WIN

For Windows:

```
cd <drive>:\fnsw_loc\bin
NLS_stop
```

Note NLS_Archive does not shut down until it has completed processing its current set. A sleep should be entered into the `ims_stop` file to accommodate the longest time period it takes for NLS_Archive to process a set in your environment. If the IS software goes down during the processing of a set, NLS_Archive shuts down with an error because it has lost its IS logon.

Command Line Startup/Shutdown Method

Command Line Startup

If you are running NLS Archive on a separate server from Image Services, or if you want to manually start NLS Archive, use the following commands to start the NLS Archive program.

There are two ways to start the archive function from a command line: **NLS_start** or **NLS_Archive**. In each case, you must first change directory to the location of the NLS install point.

When using either the NLS_start or NLS_Archive programs, the directory where the NLS.cfg file is located can be specified by using the **-h** option.

For example:

```
UNIX:      cd /fnsw/local/bin
           NLS_start -h <path to NLS.cfg>
           or
           NLS_Archive -h <path to NLS.cfg>
```

```
Windows:  cd <drive>:\fnsw_loc\bin
           NLS_start -h <path to NLS.cfg>
           or
           NLS_Archive -h <path to NLS.cfg>
```

where:

-h <path> optionally specifies the full directory path to the NLS.cfg file if it is not located in the default directory

As an alternative, you can run these programs directly if you change directory to the install point, and have the NLS.cfg file located there.

Command Line Shutdown

Use the following commands to stop the NLS Archive software. Whenever the stop command is invoked, the NLS Archive modules attempt

to shut down gracefully by completing the processing of the current set in progress prior to shutdown.

To stop NLS Archive modules, use the following commands:

```
UNIX:      cd /fnsw/local/bin
           NLS_stop -h <path to NLS.cfg>

Windows:   cd <drive>:\fnsw_loc\bin
           NLS_stop -h <path to NLS.cfg>
```

where:

-h <path> optionally specifies the full directory path to the NLS.cfg file if it is not located in the default directory

As an alternative, you can stop these programs directly if you change directory to the install point, and have the NLS.cfg file located there.

Note When NLS_start.exe has been issued on a Windows command line, the NLS Archive modules must be stopped via the command line option using the NLS_stop program.

NLS_FLAGS File

NLS_start creates an NLS_FLAGS file upon startup. NLS_stop enters a stop sequence into the NLS_FLAGS file. The NLS_FLAGS file is read by each of the above programs while running. When an NLS program receives the correct stop sequence through the NLS_FLAGS file, the program stops upon completion of the set it is currently processing. If the program is not working on a set, or in a sleep cycle, it stops immediately upon receiving the stop sequence. It is important to note that when a program is run on the command line without the use of

NLS_start that the presence of the NLS_FLAGS file with the proper stop sequence in it causes the program to stop immediately without completing any work.

If you elect to run an NLS program without the use of NLS_start, simply delete the NLS_FLAGS file.

Windows Service Setup Method

Starting and stopping the NLS Archive modules can be managed using Windows Services when running NLS Archive on a Windows 2000 Server platform. This feature can be implemented after the NLS Archive modules have been installed.

- 1 From a Command Prompt, change to the directory where the NLS Archive modules are installed. The default install directory is:
<drive>:\fnsw_loc\bin
- 2 Enter the following command to create the Windows Service:

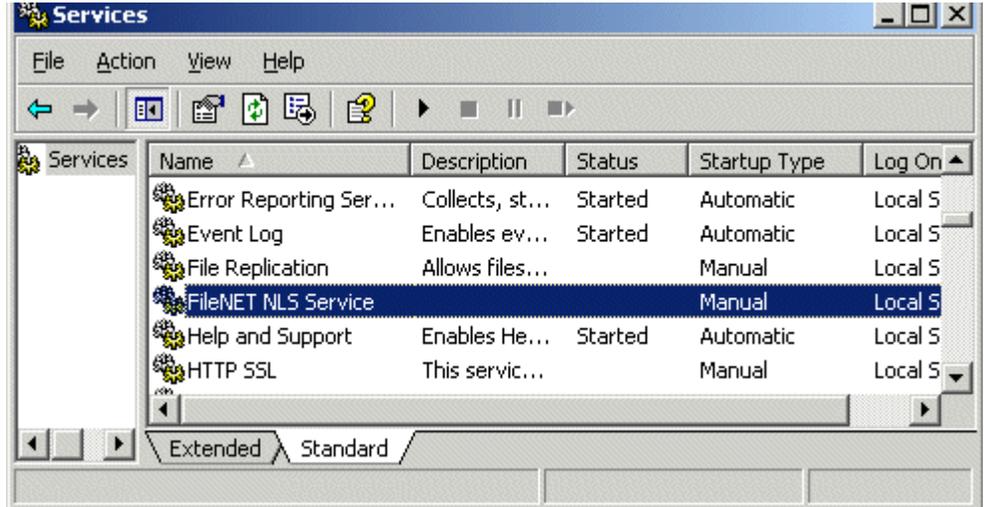
NLS_start -service

Note This command does NOT start the NLS Archive modules. It only creates the NLS Archive Service.

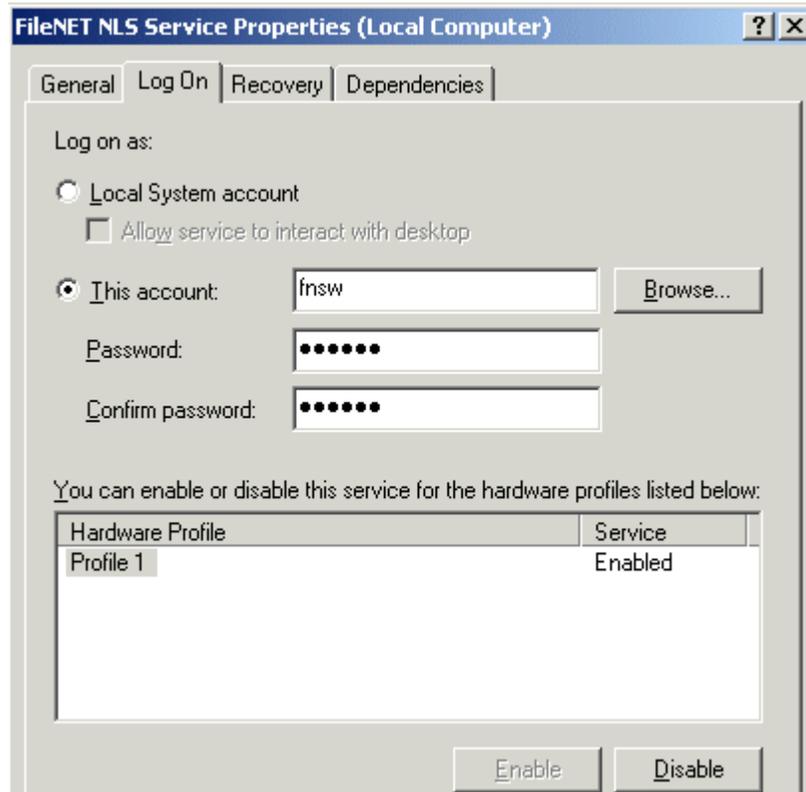
After the Service has been installed, the following message displays.

Successfully installed service NLS

- 3 On the Windows server, open Administrative Tools and double-click the **Services** icon. The Services dialog displays.



- 4 In the Services dialog box, right-click the FileNet NLS Service and select Properties. The Properties dialog box displays.



- 5 Select the Log On tab, and select the “This account:” radio button”
- 6 In the “This Account” dialog box, enter the “fnsw” user or a user in the fnusr group specifically created for running the NLS Archive modules.
- 7 Enter and Confirm the password for the account.
- 8 After the login properties have been set and all configuration items have been completed, start the NLS Archive modules by selecting the

“FileNet NLS Service” in the Windows server window and then clicking **Start**.

Note The Service Logon Account **MUST** be configured before the Service is used to start the NLS Archive modules or the startup will fail. The notification of a failure could take more than a minute to report back to the user.

Important When using the Service Applet to stop the NLS Archive modules, the applet might receive premature notice that the NLS Archive modules have stopped. However, some processes could still be shutting down and immediate restart through the applet can cause corrupted multiple processes to be run. Check the Windows Task Manager to ensure all NLS Archive processes have stopped.

Removing the FileNet NLS Archive Service

The following procedure removes the FileNet NLS Archive Service.

Check the Service through the Service applet window and ensure the service has been stopped.

Enter the following at a Command Prompt:

```
cd <drive>:\fnsw_loc\bin  
NLS_start -noservice
```

Archiving Documents

The document archive program, NLS_Archive, is designed to operate in a number of modes. These modes include:

- **Database Mode** - This mode automatically starts, and runs in the background, when NLS is started. It archives all new documents or any document that has not been previously archived.
- **File List Mode** and **Surface Mode** - These modes are invoked from a command line and run in the foreground. They archive specific sets of documents and are useful for archiving entire platters of documents.

Database Mode

When the NLS_Archive program is run without parameters, or invoked through the NLS_start program, it is run in Database Mode. In this mode the program performs these steps:

- Reads the file .nlsrc and the MinDocId attribute to determine a starting document id for the archive operation (uses the maximum of .nlsrc and MinDocId). If .nlsrc and MinDocId are not present in the NLS Archive modules home directory then the archive program begins at the lowest numbered document in the index database and works forward.
- At the starting point in the database, the NLS Archive modules search for any documents that have FNP_ARCHIVE as an index field, where FNP_ARCHIVE is null.
- The NLS_Archive program reads the document from cache. if the document is not in cache already, a pre-fetch operation is executed to put it in cache.

- NLS_Archive performs a conversion operation to take the individual cache pages and pack them all into a single file. Note that this file does not conform to any TIFF standard and is used only for purposes of archiving and retrieving FileNet documents.
- NLS_Archive writes the file (now a single file instead of a collection of pages) to the WorkingDirectory (SDS File System). From the perspective of the archive program, it sees only a UNIX file system and performs a standard UNIX write operation. The Image Services document number is used as the file name and it is placed in a subdirectory below the WorkingDirectory using a hashing algorithm based on the document id.
- After NLS_Archive successfully writes the file to the WorkingDirectory, the program updates the FNP_ARCHIVE index field with the current date. This signifies that the document has been archived to the file system repository.

Document Range Override

In Database Mode, the NLS_Archive program uses the MinDocID, MaxDocID, and AppMaxDocId values in the NLS.cfg file to determine the range of documents that it can process within.

In this mode, these values within the NLS.cfg file can be overwritten by specifying them on the command line.

-min <doc id> Sets the MinDocID value to the document ID specified.

-max <doc id> Sets the MaxDocID value to the document ID specified.

-app <doc id> Sets the AppDocID value to the document ID specified.

Example:

NLS_Archive -min <doc ID> -max <doc ID>

or

NLS_start -min <doc ID> -app <doc ID>

It is not required to specify two of these items, only the ones that need changing. A min document ID can be specified without specifying a max or appmax document ID. You also can specify a max or appmax document ID without specifying a min document ID.

File List Mode

The NLS_Archive program can be started and run independently of the other server processes when using File List Mode. This allows the user to archive a pre-defined list of documents. In file list mode, NLS_Archive does not check to determine if the documents have already been archived. It copies any document in the list to the WorkingDirectory regardless of the value in FNP_ARCHIVE. However, the Document Class must have the FNP_ARCHIVE index assigned.

Perform these steps to archive a list of documents:

- 1 Create a text file containing a list of documents.
- 2 Change directory:

cd /fnsr/local/bin

- 3 Enter one of the following commands to run the NLS_Archive program in background:

```
NLS_Archive -f <filelist> &
```

or

```
NLS_start -f <filelist>
```

Surface Mode

Similar to File List Mode, the NLS_Archive program can run independently in Surface Mode. In Surface mode NLS_Archive queries the database and create a list of document IDs to Archive where the document ID list is of the specified surface. Like File List Mode the documents found in the list are archived regardless if they have been previously archived. Should there be a copy of the document in the WorkingDirectory, it is overwritten by the newly Archived version.

Perform these steps to archive a surface of documents:

- 1 Change directory:

```
cd /fnsw/local/bin
```

- 2 Enter the one of the following commands to run the NLS_Archive program in background:

```
NLS_Archive -surf <surface ID> &
```

or

```
NLS_start -surf <surface ID>
```

Using the `-both` flag will tell `NLS_Archive` to backup both sides of the surface. For example:

`NLS_Archive -surf <surface ID> -both &`

or

`NLS_start -surf <surface ID> -both`

6

Administering and Running SDS File System Retrieval

This chapter contains information for administering and running the SDS software.

Using the DOC_tool Utility

DOC_tool provides information on your SDS device such as the device state. It also allows you to enable and disable an SDS device. For full information on using DOC_tool, refer to the [*Image Services System Tools Reference Manual*](#).

SDS States

The state of an SDS can be viewed using the DOC_tool “SDSinfo” command. The following are valid SDS states:

- **SDS not configured** - This state indicates that the SDS unit has not been configured.
- **SDS enabled** - This state indicates that the SDS unit has been configured and is enabled. Read requests can only be submitted to an SDS that is in an enabled state.
- **SDS disabled** - This state indicates that the SDS unit has been disabled. The only way that an SDS unit can be disabled is through the DOC_tool "DISable" command. The SDS disabled state is persistent across re-cycles, because a disabled SDS file is created.

When the Image Service software is re-cycled and the disabled SDS file exists, the SDS unit is marked as disabled. To enable the SDS unit, the DOC_tool "ENable" command must be used.

When an SDS unit is disabled, the dispatch process does not dispatch any new requests to the SDS worker (SDS_FS_reader). No new requests are added to the SDS worker queue until the SDS is re-enabled. The SDS worker process continues to process any outstanding SDS requests, but outstanding requests are not re-directed to surface based storage.

- **SDS disconnected** - This state indicates that the SDS unit is not available because of a connection problem encountered by the SDS worker. The disconnected state is **not** persistent across Image Services recycles. The SDS worker sets the SDS to a disconnected state if the WorkingDirectory/ssn cannot be accessed.

Once an SDS is set to a disconnected state, the dispatch software does not submit additional requests to the SDS unit. Therefore, if a document is available on MSAR or optical, a read request is re-directed to the surface based storage even if the SDS high or SDS medium priority has been set.

In addition, an insertion RSVP could also be generated if the SDS unit is in a disconnected state and associated surface(s) are "out of box". When the SDS is disconnected, the SDS_FS_reader tries to access the WorkingDirectory/ssn indefinitely.

Note

On Centera or DR550, when the SDS device is disconnected, the SDS worker retries the Pool Open command indefinitely. Once 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 DOC_tool utility program is used to manage SDS. The DOC_tool utility program supports two general SDS functions.

- The **ENable** and **DISable** commands have been enhanced to enable and disable the SDS unit.
- A **SDSinfo** command allows you to gather information about the SDS unit and the SDS worker processes (SDS_FS_reader).

These commands are also available for cache-only servers. For more information on DOC_tool, see the [***Image Services System Tools Reference Manual***](#).

ENable Command

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

Note It is not possible to change an SDS unit from a "DISCONNECTED" state to "ENABLED" state using the DOC_tool program.

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 library selected. If the SDS unit is in a DISABLED state, the following prompt displays after the ENable command is entered:

```
SDS unit? <y=yes>
```

If the ENable command is entered on a Cache-only server that **does not** have an SDS configured, the following message displays:

```
Function only available on systems with storage  
libraries configured.
```

If the ENable command is entered on a Cache-only server that **does** have an SDS configured, but not in a DISABLED or DISCONNECTED state, the following message displays:

```
No Storage libraries configured and no SDS units are DISABLED  
Use the SDSinfo command to get more info on the SDS units.
```

DISable Command

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

The screen prompt displayed are identical to those displayed for the ENable command.

SDSinfo Command

The DOC_tool SDSinfo command allows you to view the current state of the SDS unit and its associated SDS worker (SDS_FS_reader), as well as other performance information.

SDSinfo command prompt and options

The **SDSinfo** command allows you to select four display options as shown below. Options include: **(s)** a Summary report, **(q)** a Queue report, **(w)** a Detail worker report, or **(a)** an All information report.

```
<DOC_tool:lib a>sdsinfo
Summary info, Queued request info, Worker info, or All info?
('s','q','w','a')
```

If the SDSinfo command is invoked and there are no SDS units configured on the server, the following message is displayed:

```
<DOC_tool:lib a>sdsinfo
No SDS have been configured on this server
```

SDSinfo Basic Information

The following table describes basic information that is displayed for all reports, regardless of which **SDSinfo** option is selected.

Basic SDS Information

Field Name	Description
SDS unit number	The SDS unit always set to 1 in this release.
SDS state	The SDS unit can be set to a DISABLED state (via the DISable Command option of DOC_tool, to an ENABLED state (via the ENable Command option of the DOC_tool, and a DISCONNECTED state (via the SDS_worker [SDS_FS_reader] because the SDS_worker has encountered a problem making it impossible to process requests).
SDS Priority	This can be set to high, medium, or low and this setting is extracted from the sds_conf file. This setting is for when a document is not found in cache. High priority means the document is first looked for on the SDS unit, then in “in box” MSAR, and finally on optical. Medium priority means the document is first looked for in “in box” MSAR, then on the SDS unit, and last on optical. Low priority means the document is first looked for in “in box” MSAR, then on optical and finally on the SDS unit.
SDS Worker name	The SDS Worker name uses the format SDS_<device>_reader NetApp SnapLock - SDS_FS_reader IBM DR550 - SDS_ISAR_reader EMC Centera - SDS_CSAR_reader
Number SDS worker instances	This indicates how many SDS worker instances are running, The number of worker instances is configurable through the sds_conf file and the valid number instances are 1 to 99.
SDS info	The SDS Worker expects the following label: workingdirectory = <directory name> This is a maximum of 1024 characters.
SDS queue length	This displays the current number (0-n) of SDS queued requests. A high number here might indicate a problem.

SDSinfo Summary Option (s)

When the **Summary** option is selected, two summary reports are displayed; the **Counter Summary** and the **Accumulated Average Time Summary** reports. Each of these reports is displayed along with the basic SDS information.

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

The following table describes the **Counter Summary Worker** report:

SDS Counter Summary Worker Information

Counter Name	Description
Requests Processed	Total number of requests that have been processed by all of the SDS Worker processes.
Successful requests processed	Total number of requests that have been successfully processed.
Requests satisfied via SDS device (IBM DR550)	Total number of requests that have been processed using SDS. By subtracting the "Requests satisfied via SDS" from the "Requests Processed," the number of requests processed using the NLS working directory can be determined.
Errors	Total number of errors encountered. These are only errors that were not corrected after retries and did not result in termination of SDS worker processes.
Last_error	Last error encountered by the last SDS worker process. If no errors are encountered by SDS workers, this field is not displayed.
Images retrieved	Total number of images retrieved.
Data retrieved	Total amount of data, in K-bytes, transferred to CSM cache from the NLS working directory.
AVG request size	Average size of data transfer per requests. A standard deviation for this field is also displayed, and this can be used to gauge the other Standard Deviations to follow.

SDS Counter Summary Worker Information, continued

Counter Name	Description
AVG image size	Average size of image pages transferred.
Number of SDS device reads (API calls)	Centera: Total number of calls to the FPTag_BlobReadPartial API entry point. IBM DR550: Total number of this sequence of API calls to read a DR550 Tivoli object.
Number of SDS device query calls	IBM DR550 Only: Number of queries for an internal object ID based on the object name.
Number of SDS device requests where the whole document fits into the internal image_buffer (128K)	Centera and IBM DR550: Requests where the whole document fits into the internal buffer. The size of configurable internal buffer is displayed. For more information, see “Whole Document BLOB Fits into the Internal Buffer” on page 133.
Number of SDS device requests where the requested images fit into the internal image_buffer (128K)	Centera and IBM DR550: Requests where all of the images fit into the internal buffer. The total size of these images is displayed. For more information, see “Requested images Fit into the Internal Buffer” on page 134.
Number of SDS device requests where the requested images do not fit into the internal image_buffer (128K)	Centera and IBM DR550: Requests where all of the images do not fit into the internal buffer. For more information, see “Requested Pages Do Not Fit into the Internal Buffer” on page 135.
Cache hits	Cache hits encountered by the SDS_FS_reader process.

The following table describes the **Accumulated Average Time Summary Worker** Report information, which is also displayed using the Summary option.

SDS Accumulated Average Time Summary Worker Report

Average Accumulated Times	Description
Up time	Accumulated average up time of all the SDS workers in seconds/workers.
Idle time	Accumulated average up time of all the SDS workers in seconds/worker. This is the time that workers are idle waiting for requests to be processed.
Total processing time	Total elapsed processing time per worker. The average time elapsed per request is also displayed in seconds/workers with its standard deviation. The average time elapsed per image page (seconds/image page) and elapsed per K-byte (seconds/K-byte) is also displayed.
Time in get_doc_from_file_system()	Accumulated average time in the get_doc_from_file_system() subroutine in seconds/workers. This is the accumulated time taken to execute the get_doc_from_file_system subroutine which includes extracting the requested document pages from the file system and writing them to page cache.
Time in get_doc_from_device()	Centera and IBM DR550: Accumulated average time in the get_doc_from_device() subroutine in seconds/workers. This is the accumulated time taken to execute the get_doc_from_device subroutine which includes extracting the requested document pages from the SDS device and writing them to page cache.

SDS Accumulated Average Time Summary Worker Report, continued

Average Accumulated Times	Description
AVG time in SDS device open and read (API calls)	<p>Accumulated average time doing API calls (Seconds/workers).</p> <p>Centera: This includes operations such as Open Pool, Open Clip and BLOB reads. It is a subset of the operations performed by the <code>get_doc_from_centera</code> subroutine. The average time elapsed per request is also displayed (seconds/requests) with the Standard Deviation. A summed squared of each elapsed time to do the Centera API calls per requests is calculated and stored. This is done so the Standard Deviation for Centera API calls time per requests can be calculated. The average time elapsed per image page (seconds/image page) and elapsed per K-byte (seconds/K-byte) is also displayed.</p> <p>IBM DR550: Includes the set of API calls to connect to the DR550, query session information and API version number, and the sequence of API calls to read a DR550 object.</p>
AVG time in SDS device read only (API calls)	<p>Centera: Accumulated average time doing FileNet BLOB reads.</p> <p>IBM DR550: Includes the set of API calls to read a DR550 object.</p> <p>For both, this is a subset of the accumulated time doing API calls. The average time elapsed per request is also displayed (seconds/requests) with the Standard Deviation. The average time elapsed per image page (seconds/image page) and elapsed per K-byte (seconds/K-byte) is also displayed.</p>
AVG time in SDS device query calls	<p>IBM DR550 Only: The average in doing a query call to obtain an internal object ID based on the object name.</p>

SDS Accumulated Average Time Summary Worker Report, continued

Average Accumulated Times	Description
AVG time in writing to page cache (CSM)	Accumulated average time doing CSM cache calls. The accumulated time includes time spent writing to cache. The average time elapsed per request is also displayed (seconds/requests) with the Standard Deviation. The average time elapsed per image page (seconds/image page) and elapsed per K-byte (seconds/K-byte) is also displayed.
AVG requests queue wait time	The average wait time in seconds per request. The wait time is the elapsed time of when a request is enqueued by the dispatch process, and the time the SDS_FS_reader dequeues the request. The wait time is accumulated and an average is displayed here. In addition, a Standard Deviation of wait time in seconds per request is also calculated and displayed.

Sample output - Summary Option

```
<DOC_tool:lib a>sdsinfo
Summary info, Queued request info, Worker info, or All info?
('s','q','w','a'): : s

SDS info: SUMMARY 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

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
  Images retrieved = 9
  Data retrieved = 0.719210MB
  AVG request size = 245.490234K (STDEV=477699.203356)
  AVG Image Size = 81.830078K
  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
AVERAGE ACCUMULATED ELAPSE TIMES:
  Up time: 76152.546358 secs/workers (1269.209106 mins)
  Idle time: 76147.776404 secs/workers (1269.129607 mins) (99.99%)
(continued on next page)
```

Sample output - Summary Option, continued

```
Total processing time: 4.752298 secs/workers
                        (1.584099 secs/reqs)
                        (0.848008 STDEV secs/reqs)
                        (0.528033 secs/image page)
                        (0.006453 secs/KB)
Time in get_doc_from_filesystem(): 0.000790 secs/workers
Time in get_doc_from_device():      4.076062 secs/workers
AVG time in SDS device open and read (API calls): 2.991772 secs/workers
                                                (0.997257 secs/reqs)
                                                (0.840033 STDEV secs/reqs)
                                                (0.332419 secs/image page)
                                                (0.004062 secs/KB)
AVG time in SDS device read only (API calls): 2.133966 secs/worker
                                                (0.711322 secs/reqs)
                                                (0.237107 secs/image page)
                                                (0.002898 secs/KB)
AVG time in SDS device query calls: 0.000000 secs/worker
AVG time in writing to page cache (CSM): 0.438663 secs/workers
                                                (0.146221 secs/reqs)
                                                (0.179214 STDEV secs/reqs)
                                                (0.048740 secs/image page)
                                                (0.000596 secs/KB)
AVG requests queue wait time: 0.247177 secs/reqs
                                (0.426779 STDEV secs/reqs)
<DOC_tool:lib a>
```

SDSinfo Queue Option (q)

When the **Queue** option is selected, outstanding SDS queued request information is displayed along with the basic SDS information. For each outstanding request the following information is displayed:

Queue Option information

Queue Option	Description
Doc_id	source doc_id
Doc_index	target document ID
Request type	For this enhancement the request type is always read request type.
Ssn	SSN of the source
Cache_id	Target cache_id
Cache_ssn	Cache or target SSN
Firstpage	First page to retrieve
Numpages	Number of pages to retrieve
Prim_surf_id	Primary surface ID
Tran_surf_id	Tranlog surface ID
Notify_ptr	Notification pointer
Has_clipid	Indicates the CLIP ID was retrieved from doctaba when the request was dispatched

If there are requests in the queue when the queue option is selected, the screen looks similar to the sample output screen below:

Sample output - Queue Option

```
<DOC_tool:lib a>sdsinfo
Summary info, Queued request info, Worker info, or All info?
('s','q','w','a')
<DOC_tool>q

SDS info: QUEUED option

***** SDS unit = 1 state = ENABLED SDS priority = high *****
  DEBUG Setting = MIN
  Worker = 'SDS_FS_reader' Number Instances = 1
  info = 'workingdirectory=/surf/HSM_DISK
  Current number of outstanding requests in the SDS queue: 1
Request Type: READ_REQUEST
  doc_id = 1004259, doc_index = 1004259, ssn = 10347,
  cache_id = 1, cache_ssn = 10347, firstpage = 1, numpages = 1,
  prim_surf_id = 33002, tran_surf_id = 0
  notify_ptr =b0691b88, has_clipid = TRUE
<DOC_tool:lib a>SDS
```

If there are no requests in the SDS queue, a message similar to the following displays indicating that the queue is empty:

Sample output - Empty Queue Option

```
<DOC_tool:lib a>sdsinfo
Summary info, Queued request info, Worker info, or All info?
('s','q','w','a')
<DOC_tool>q

SDS info: QUEUED option

***** SDS unit = 1 state = ENABLED   SDS priority = high *****
        DEBUG Setting = MAX
        Worker = 'SDS_FS_reader' Number Instances = 4
        info = 'workingdirectory=/surf/HSM_DISK
        Current number of outstanding requests in the SDS queue: 0

SDS (unit=1) queue is currently empty
```

SDSinfo Worker Option (w)

When the **Worker** option is selected, basic SDS information is displayed, along with information about each active worker and a summary worker report. The information displayed is similar to worker summary reports, but the counters and time accumulated are specific to each worker. The fields displayed in the Summary 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 request done.

Sample output - Worker Option

```
<DOC_tool:lib a>sdsinfo
Summary info, Queued request info, Worker info, or All info?
('s','q','w','a'): : w

SDS info: WORKER option

***** SDS unit = 1 state = ENABLED SDS priority = medium *****
  DEBUG Setting = MAX
  Worker = 'SDS_ISAR_reader' 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
  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
(continued on next page)
```

Sample output - Worker Option, continued

```
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
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)
.
.
.
```

SDSinfo All Option (a)

When the all option is selected, all the information mentioned above is displayed.

Sample output - All Option

```
<DOC_tool>sdsinfo
Summary info, Queued request info, Worker info, or All info?
('s','q','w','a')::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

```
(continued from previous page)
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

```
(continued from previous page)

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
.
.
.
```

Retrieving Documents

The dispatch process component of the SDS software knows how to process all retrieval requests for documents whether they are located in CSM page cache, in an SDS File System, or on an MSAR or optical surface.

Note The priority setting in the `sds_conf` file determines where a retrieval request is directed.

When processing an SDS retrieval request, the request is sent to the SDS Worker (such as `SDS_FS_reader`, `SDS_CSAR_reader`, or `SDS_ISAR_reader`). Since the SDS Worker is integrated into the Image Services software, the retrieval process is transparent. Therefore, as long as the IS system is configured correctly, retrieval requests are made as usual using regular Image Services utilities.

Adding a Physical Volume

This chapter describes procedures for adding a physical volume to your SDS configuration and for modifying the `NLS.cfg` and `sds_conf` files to recognize it.

If your WorkingDirectory approaches the maximum size for a physical volume on your SDS device, you might need to add another volume. (For example, the current maximum raw physical volume size for some storage devices is 16 TB. The effective storage capacity is about 10 TB.)

When you add another physical volume, you must update the **NLS.cfg** file to reference the WorkingDirectory on the new volume. And you must update the **sds_conf** file to add the file system on the new volume to the WorkingDirectory search path.

Although you can specify only one WorkingDirectory in the `NLS.cfg` file for NLS archiving, you can specify several Working Directories in the `sds_conf` file for SDS retrievals.

Installing the New Physical Volume

Follow the manufacturer's instructions for installing the new physical volume in your SDS device. Create a file system for NLS using the same naming convention used for previously configured NLS file systems.

For example, if your first NLS file system is `/NLS/SnapLock`, you might name the file system on the new volume `/NLS2/Snaplock`.

Updating the NLS.cfg File

After you have successfully installed the new volume and created the new file system, you need to edit the NLS.cfg file.

- 1 Locate the NLS.cfg file.

UNIX`/fnsw/local/bin`**WIN**`<drive>:\fnsw_loc\bin`

- 2 Use your preferred text editor to modify the WorkingDirectory line. For example, if the WorkingDirectory line for your SnapLock device currently looks like this:

```
WorkingDirectory=/NLS/Snaplock
```

You would change it to the new directory path on the new volume: :

```
WorkingDirectory=/NLS2/Snaplock
```

- 3 Save your changes and close the file.

The changes will take effect the next time the NLS software is restarted.

Updating the sds_conf File

The NLS.cfg file can specify only one WorkingDirectory because content is archived to only one volume at a time. But the sds_conf file can specify several WorkingDirectory paths because content can be retrieved from multiple volumes.

The info statement in your sds_conf file might look like this:

```
info "WorkingDirectory=/fs4;/fs3;/fs2;/fs1 Domain=..."
```

- 1 Locate the sds_conf file.

UNIX

/fnsw/local/sd/1

WIN

<drive>:\fnsw_loc\sd\1

- 2 Use your preferred text editor to modify the WorkingDirectory section of the info line. For example, if the info line for your SnapLock device currently looks like this:

```
info "WorkingDirectory=/NLS/Snaplock"
```

You would add the new directory path of the file system on the new volume ahead of the previous file system path: :

```
info "WorkingDirectory=/NLS2/Snaplock;/NLS/Snaplock"
```

The file systems are searched in the order you list them on the info line, so the active NLS WorkingDirectory should be listed first, then the others in order by age. The entire info line can be 4096 characters long. There should be no spaces between the file system names and no carriage returns—just semi-colons.

- 3 Save your changes and close the file.

The changes will take effect the next time the Image Services software is restarted.

Appendix A – Troubleshooting NLS Archive

This appendix describes how to troubleshoot problems you might encounter when using NLS Archive. These include:

- [“Reporting Problems”](#)
- [“Utility/Operating System and IS Toolkit Version Relationship”](#)
- [“Error Logging Files”](#)
- [“NLS Error Message Tuples”](#)
- [“NLS Error Codes”](#)
- [“Error Codes Known to Need Specific Actions”](#)
- [“When to Use wal_purge \(UNIX Only\)”](#)
- [“Operating System Path Issue”](#)
- [“Shell Problems Causing NLS Archive Failures \(UNIX Only\)”](#)

Reporting Problems

When a problem with NLS Archive is perceived, a sample run with the minimum amount of data needs to be executed. The debug option '-D' and the spy option '-S' can be used to assist you in identifying a problem, or when requested by an FCP.

Note Do not use -M unless specifically requested by an FCP.

Debugging messages are added to the journal log file. This file, along with a description of the event that presented itself should be sent to

the appropriate support individuals. Failure to provide all output files delays the ability to diagnose the problem.

Note SDS debugging output does not appear on Windows if NLS.cfg's setting for MaxArchChildren is greater than 0.

When reporting a problem, send the following information as applicable for your system.

- Site name or Case Number.
- System Configuration
 - Operating System platform and version
 - Database platform and version
 - Run **fn_util whichfn** and supply output to support
 - Image Services version and stamp /fnsw/lib/shobj/SysV
 - IS Toolkit version
 - IS Toolkit stamp /fnsw/client/shobj/SysV
 - Directory listing of IS Toolkit installation point to verify permissions when in a UNIX environment
 - (ls -alR /fnsw/client --> <filename>)
 - Is IS Toolkit/Utility installed directly on IS or on a remote system? If remote, supply the Operating System for remote environment.
 - PS Utility version (for example, **NLS_start -v <domainname>**)

- PS Utility stamp on UNIX systems only (for example, **stamp NLS_Archive**) (Stamp not available in Windows environment)
- Directory listing of PS Utility installation point to verify permissions when in a UNIX environment
(ls -al /fnsw/local/bin --> <filename>)
- CheckSSN output from display in DOS prompt or Command window
- State problem
- Outline steps on how to reproduce problem
- Please give exact command used to start/stop the PS Utility
- Output from display in DOS prompt or Command window of Utility runtime
- Nohup files - for example, /tmp/<Utility_Name.datetime>
- Utility Journal logs (/fnsw/local/bin/journals/*)
- Utility Configuration file NLS.cfg
- Input files used (DocID list files, eob files, transact.dat files, images, etc.)
- IS Toolkit Error logs (/fnsw/client/tmp/walyyyyymmdd)
- IS Error Logs (elog#####)
- Stack trace of core (if a core file is produced)
DO NOT send the core file.

Note See the **“Error Logging Files” on page 112** for an explanation of the log file and default paths.

Utility/Operating System and IS Toolkit Version Relationship

The utility build version selected is dependent upon the IS Toolkit version in use on the IS Server. The utilities have been compiled to run with IS Toolkit 3.6 and 4.0. The installation software package detects the version of IS Toolkit and installs the correct compile of the selected Utility when installing on a UNIX system. When installing on a Windows server, the user is responsible for knowing the IS Toolkit version and selecting the correct Utility compile when asked during the setup procedure. Should the user need to download a version of a Utility off the support download site, it is necessary to verify the IS Toolkit version by running a stamp command on a IS Toolkit module in the IS Toolkit install directory.

To verify the OS version level, enter:

UNIX system:

```
>fn_util whichfn  
>stamp /fnsw/lib/shobj/*SysV*
```

Windows system:

```
fn_util whichfn  
stamp <drive>:\fnsw_loc\lib\shobj\SysV.dll
```

To stamp an IS Toolkit module, enter:

UNIX system:

```
stamp /fnsw/client/shobj/*SysV*
```

Windows system:

```
stamp <drive>:\fnsw\client\bin\SysV.dll
```

The “system” number and “rel_type” are both important informational items to take into account when selecting the correct utility version.

Error Logging Files

Errors that occur during the operation of this utility can be reported in several places, depending on the error type. 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 rectify.

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

NLS Archive Logs

Viewing Logs in a Windows Environment

On Windows systems, log files are located in the default directory: **<drive>:\fnsw_loc\bin\journals** or in the directory set in the NLS.cfg file.

Viewing log files when the NLS Archive modules are running in a Windows environment can be limited because of file access conflicts. It is necessary to stop the NLS Archive modules or use the Windows **more** command in a Command prompt.

To view a log file while NLS_Archive is running:

- 1 Open a Command prompt.

- 2 Change to the following directory:

```
cd <drive>:\fnsw_loc\bin\journals
```

- 3 Enter the command:

```
more < arc YYYYMMDD
```

Viewing Logs in a UNIX Environment

Journal log files are found in the default journals directory, **/fnsw/local/bin/journals** or in the directory set in the NLS.cfg file.

Any batch progress and general error messages are reported in the journal files.

Journal files are named as follows:

- arcYYYYMMDD
- sdsYYYYMMDD

Note that only if NLS.cfg/MaxArchChildren>0 are sdsYYYYMMDD journal files created.

Other Logs

IS error logs – all IS related messages and errors (Storage Library, SDS_FS_reader, database, security, etc.) Applicable to all utilities. Use the vl command to view error logs.

Default UNIX - /fnsw/local/logs/elogs

Default Windows - \fnsw_loc\logs\elogs

IS Toolkit error logs – all IS Toolkit messages and errors (an IS error log specifically for the IS Toolkit environment) Applicable to all utilities.

Default UNIX - /fnsw/client/logs/walYYYYMMDD

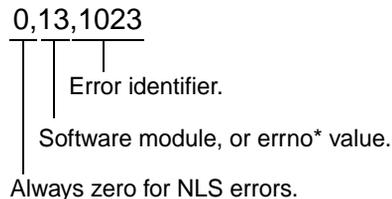
Default Windows - \fnsw\client\logs\walYYYYMMDD

Nohup logs – all standard out / standard error messages (system errors, preamble and postamble logging from utility, output generally captured via standard out that would be suppressed when using `-q` for quiet mode, etc). Applicable to utilities on UNIX ONLY.

Default UNIX - /tmp/{product}.YYYYMMDD

NLS Error Message Tuples

NLS Error messages include a three-part numerical identifier called a tuple (for example, 0,13,1023). The following diagram identifies each part of an **NLS** error tuple:



* The most common time you see an 'errno' value is for IO errors.

NLS errors are also logged into the NLS applications journal log. See [“Error Logging Files” on page 112](#) for more information.

NLS Error Types

NLS can display two types of errors; NLS specific errors and IS errors.

NLS Specific

If the first number of the tuple identifier **is** zero (for example, 0, 13, 1023) **AND** the last number **is** greater than 1000, then the error is an NLS error. The last number, in this case 1023, is the specific error number listed in the NLS applications documentation.

The second number in this tuple, in this case 13, is the '**errno**' generated from the native operating system. The most common time you see an '**errno**' value is for IO errors. Using the `fn_msg` tool does **not** work for these types of errors.

IS Errors

IS errors can also be displayed by NLS. IS errors can be identified as being any error that does not conform to the NLS specific parameters described above. For example, the first number of the tuple identifier **is not** zero and the last number **is not** greater than 1000.

Using the `fn_msg` tool **does** work for these types of errors. See [**“Using the fn_msg Tool” on page 124.**](#)

NLS Error Codes

The following table is a general list of errors you might encounter.

Note These error code numbers can change in subsequent builds and might not be reflected in this document.

ERROR NUMBER	ERROR TEXT
999	lowest error – empty
1000	Incorrect number of cfg argument
1001	cfg bad key word
1002	cfg exceed max
1003	cfg expect brace
1004	cfg file is empty
1005	cfg get failed
1006	cfg invalid index type
1007	cfg invalid keyword
1008	cfg invalid number
1009	cfg too many token
1010	cfg no 2 nd quote
1011	cfg Required Directory missing from Configuration File
1012	cfg not digit
1013	Value for Keyword is not valid
1014	cfg open fail
1015	cfg too many index
1016	cfg unexpected eof
1017	end of file
1018	EOB file contains bad information
1019	cannot convert

ERROR NUMBER	ERROR TEXT
1020	Data Error. Checksum of file is bad
1021	DocClass info not in configuration file
1022	Number of Documents in EOB and transact.dat mismatch
1023	Document ID List File is required (-f option)
1024	Can't exec command
1025	Index is not defined in document class
1026	Invalid document number
1027	Invalid object id
1028	Invalid object name
1029	Invalid ISR type
1030	IO error
1031	Critical IO error – Program Abort
1032	Can't kill a child process
1033	Line too long
1034	malloc failed
1035	Miscellaneous error
1036	Source domain must be supplied for this application
1037	Missing required index in configure
1038	Required command-line option missing
1039	No annotations were found for this document
1040	Document does not contain any page
1041	Multiple Classes are not supported by this application
1042	One or more critical resources are not available
1043	Not a valid FileNet format. See SC_convert.out for details

ERROR NUMBER	ERROR TEXT
1044	Not enough space
1045	open file fail
1046	Number of Pages in EOB and transact.dat mismatch
1047	Value for required index is missing
1048	Can't spawn the external report program
1049	System Error, see Errno
1050	Program Terminal ID is Required
1051	Image file did not validate
1052	Too many docs in subbatch. MaxDocPerSubBatch too small
1053	Too many docs returned by query. Must query on unique index
1054	Batch contains more than one family
1055	Too many pages in doc
1056	Too many sub batches. MaxBatchSize too small
1057	Transact file contains incorrect format
1058	Required input file was not provided
1059	PS Shared Library Mismatch-Contact an FCP
1062	ISR Validation Failure
1063	Unsupported SSN
1064	Preliminaries Complete

Error Codes Known to Need Specific Actions

,#### – folder/file path does not exist

0,13,#### – folder/file permissions are incorrect

15,16,17 – This is an IS network error and needs to be resolved by the System Administrator.

156,0,24 - <NCH,0,24> “A network related error was encountered.”
The most common reasons for this error are:

- They do not have the correct “nch-server” entry in the /etc/hosts file.
- They provided an incorrect Domain or Organization name in the NLS.cfg file.
- The FileNet IS software is stopped.
- There is a network problem.

202,6,3 - ./CheckSSN (10662) WARNING: Current process is not registered

This error usually means that the IS Toolkit libraries are not installed properly and the PS Utility is accessing the IS libraries instead of the IS Toolkit libraries.

202,0,9 – SysV segmentation violation: Problem between IS and IS Toolkit memory management modules. Any error of this type is difficult to detect. The following is a list of known problems that can receive this error:

- 202,0,9 – IS Toolkit install incomplete (Modules not copied down entirely or configuration incomplete.)

- 202,0,9 – IS Toolkit and Utility compile mismatch (Utility is compiled for each IS Toolkit version. See section on Hardware and Software requirements.)
- 202,0,9 – Utility installation incomplete (Not all modules were copied down or permissions on NLS Archive modules are incorrect on UNIX.)
- 202,0,9 – SysV module from IS is being accessed instead of the SysV module in the IS Toolkit installed directory. Check your PATH to see that the IS Toolkit path is listed before the IS path. Test has found problems where the IS shobj path has been input into another PATH value that causes the IS SysV to load when a IS Toolkit call is made thus causing a segmentation violation.

Error output that can be found in the nohup log (/tmp/NLS_importyyyyymmdd) file when there is an incorrect path present in the environment:

```
sys_log: fnc_get_process_name error: program is not
registered. pid = 15701
Process aborting due to segmentation violation...
```

When to Use wal_purge (UNIX Only)

The IS Toolkit program /fnsw/client/bin/wal_purge can clean up numerous problems but it can also be used incorrectly when IS Toolkit programs are running and cause application/program failures. Only run wal_purge when all IS Toolkit applications are stopped. This can be done via a script that runs immediately after bringing up the FileNet software but before starting any IS Toolkit applications/programs.

The following cases are times when wal_purge can be used to clean up an application problem.

- If the FileNet software is shut down before the utility is stopped.
- If any changes are made to the NLS.cfg file.
- If a newly created doc class or index is not recognized by the utility.
- Various security problems with the utility.

Operating System Path Issue

An Operating System contains an environment file that can load paths where origins can be difficult to detect. If your site experiences odd behavior when your utility is running, it is important to check all PATH variables. The **set** command can be used to view the environment variables. If a PATH is displayed through **set** but you are having problems locating the source of that PATH setting, check the /etc/profile file for incorrect entries.

Shell Problems Causing NLS Archive Failures (UNIX Only)

UNIX Operating System environments differ greatly between shells and environment variables configured within a shell. Some UNIX Operating System shells can cause the NLS Archive modules to fail during various execution and communication areas. If this occurs, it is recommended that you configure the environment for the User that is running the NLS Archive modules to use the C shell (csh).

Appendix B – Troubleshooting SDS Retrieval

This appendix describes contains information about errors you might encounter with SDS retrieval.

Error Logs

When an Image Services related error is encountered, the FileNet error tuple and error text, if available, is 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, etc.) are logged in the IS elog. Applicable to all utilities. Use the **vl** command to view error logs.

Default UNIX - /fnsw/local/logs/elogs

Default Windows - \fnsw_loc\logs\elogs

IS Toolkit error logs – all IS Toolkit messages and errors (an IS error log specifically for the IS Toolkit environment); applicable to all utilities.

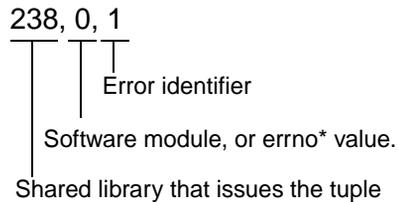
Default UNIX - /fnsw/client/logs/walYYYYMMDD

Default Windows - \fnsw\client\logs\walYYYYMMDD

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

SDS Error Message Tuples

Error messages include a three-part numerical identifier called a tuple (for example, 238,0,1). SDS error tuples (generated by the SDS abstract or SDS_reader programs) use the same format that is used for Image Services error tuples. The following diagram identifies each part of an **Image Services** error tuple:



* The most common time you see an 'errno' value is for IO errors.

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. The SDS error is usually sent directly from the IS Toolkit call.

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

Using the `fn_msg` Tool

You can 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:

```
30,0,1 Optical disk drive not ready.
```

Enter this command to display additional information:

```
fn_msg 30,0,1
```

The following text displays:

```
<ODX,0,1> Drive not ready.  
This typically is an intermittent drive problem which can be recovered from by either  
retrying the I/O, or removing/reinserting the media in the drive and then retrying the I/O.
```

In this example, ODX indicates the issuing shared library is the Optical Disk Transfer, followed by additional information for resolving the error.

The system sometimes displays an error code as a single hexadecimal number, instead of displaying the three-part error tuple. You can use the `fn_msg` tool to interpret the hexadecimal version of the error tuple as well.

For example, the hexadecimal equivalent of the error tuple <30,0,1> is 0x1E00001. In this case, running any of the following commands produces the same text:

```
fn_msg 30,0,1
fn_msg 0x1e000001
fn_msg 1e000001
```

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

TPI (192) Messages

If an entry similar to the one below appears in the log file, there might be an insufficient number of Document Services DS_NOTIFY processes:

```
192,1,6 <fnsw> ds_notify (3975) ... [WARNING]
Error return from TPI_SendPEPReq IP=112.128.39.75 socket = 2520,
requestid=58728052 doc_id = 111414 notification"
```

The DS_NOTIFY process in Image Services notifies a client application when the requested document has been retrieved into cache. On very active servers it would be beneficial to tune the number of DS_NOTIFY processes.

Tip Use the FileNet Image Services System Configuration Editor (`fn_edit`) to change the number of DS_NOTIFY processes on the Performance Tuning tab, System Processes subtab.

For example, the usual rule for storage libraries is to have at least one DS_NOTIFY process per drive. In a configuration with three optical storage libraries with two drives each, the recommended number of DS_NOTIFY processes is 6.

The same approach is recommended for SDS devices as well, that is, one DS_NOTIFY process per SDS Reader.

In most cases, however, you do not need to modify this setting unless you see a warning message in the log file.

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 [“Using the fn_msg Tool” on page 124](#).

<p>238,0,1 SDS Informational message. (SDS_err_info)</p> <p>238,0,2 Invalid session handle. (SDS_err_invalid_sess)</p> <p>238,0,3 Indicates that a particular document is not stored in the SDS. (SDS_err_no_document)</p> <p>238,0,4 Indicates that an unsupported request type was encountered by an SDS worker. (SDS_err_not_type_supported)</p> <p>238,0,5 Indicates a particular SDS unit is disabled. (SDS_err_disabled)</p> <p>238,0,6 Indicates a particular SDS unit is not configured. (SDS_err_not_configured)</p> <p>238,0,7 Indicates a particular SDS unit is not valid. (SDS_err_invalid_sds_unit)</p> <p>238,0,8 The checksum stored with the SDS document page is not consistent with the retrieved document. (SDS_err_chksum)</p> <p>(continued on next page)</p>

(continued from previous page)

238,0,9 A critical error was encountered by an SDS worker and as a result the worker aborted.

(SDS_err_worker_abort)

238,0,10 Doc_id or ssn stored in the document header and read request do not match.

(SDS_err_header_doc_or_ssn_mismatch)

238,0,10 Doc_id or ssn stored in the document header and read request do not match.

(SDS_err_header_doc_or_ssn_mismatch)

238,0,11 Invalid worker instance number.

(SDS_err_invalid_instance)

238,0,12 The version number of the SDS Worker Info structure does not match.

(SDS_err_info_struct_ver_mismatch)

238,1,4 The NLS directory has not been configured.

(SDS_err_nls_root_not_config)

238,1,10 The NLS indexes are not defined correctly.

(SDS_err_bad_nls_index_def)

238,1,11 The FNP_ARCHIVE user index field is not defined.

(SDS_err_FNP_ARCHIVE_not_def)

238,2,1 IBM DR550 error that is fatal and as a result SDS worker is terminated. See sys_log.

(SDS_err_fatal_dr550)

238,2,2 IBM DR550 error. See message in sys_log.

(SDS_err_dr550)

238,2,3 Cannot open IBM DR550 repository. The repository may be down, the network may be down or the IP address may be incorrectly entered in the sds_conf or dsm.sys or dsm.opt file.

(SDS_err_dr550_conf)

238,2,4 The FNP_CLIPID index contains an invalid object descriptor specified by the writer task for an IBM DR550 device.

(SDS_err_dr550_bad_obj_desc)

(continued on next page)

(continued from previous page)

238,2,5 This IBM DR550 session has been terminated.

(SDS_err_dr550_sess_disconnect)

238,2,6 Experiencing network error and cannot talk to IBM DR550 server. Check TCP/IP installation and LAN connection.

(SDS_err_dr550_network_err)

238,2,7 No object is found on the IBM DR550 based on this object ID or query name.

(SDS_err_dr550_objid_notfound)

238,2,8 Cannot connect to IBM DR550 because node name of this server is not registered on IBM DR550. Please follow documentation to set up node name, management class and policy domain on DR550 before connecting to it.

(SDS_err_dr550_node_notfound)

Appendix C – Centera Storage and Retrieval

This appendix contains information specific to the Centera storage and retrieval (CSAR) option for FileNet Image Services.

Support Policy

Image Services 4.1.2 is qualified with EMC Centera Software Developers Kit (SDK) version 3.1.1 and is supported in connection with all versions of CentraStar that are compatible with this version of the SDK.

Support is limited to the functional scope of IS integration with Centera, as described in the FileNet documentation, and that it should not be assumed that new features appearing in a release of CentraStar are automatically supported.

Furthermore, support of Event-Based Retention (EBR) is only available in connection with CentraStar 3.1.x and later. To take advantage of Event-Based Retention in Centera Governance and Compliance Plus editions, customers must acquire the Advanced Retention Management (ARM) license from EMC and have EMC activate the event-based retention on the licensed system."

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.

What is Centera?

Centera is a storage repository device manufactured by EMC Corporation. It uses a category of storage called Content Addressed Storage (CAS). Rather than address a data object by its file name at a physical location, a CAS device uses a content address that is based on attributes of the content itself. This results in a unique identifier (or Content Address) that is used to store and retrieve the data object.

This data access paradigm is known as Content Addressing. Instead of requiring an application to store and retrieve static content in the traditional address location-based approach, Centera presents a flat address scheme to the archive program.

When a reference data object is initially stored in the Centera repository, the application is given a “claim check” that is uniquely derived from the object’s content. Subsequent retrievals of that data object are made by simply giving the claim check, that uniquely identifies the object, back to the repository, and the data object is returned.

The “claim check” referred to above, is called a CLIP ID. The CLIP ID is a package containing the user’s data (BLOB) and associated C-CLIP Descriptor File (CDF), or Meta-data stored in a tag form.

After the document file has been written, and the FNP_ARCHIVE index field has been updated with the current date, the NLS Archive program copies the document from the NLS File System to the Centera repository.

Once the document has been copied to the Centera repository, a CLIP ID is returned from the repository. The NLS Archive program stores the CLIP ID in a CLIP ID file (<doc_id>.clipid), and/or in the FNP_CLIPID INX User Index.

Note The SdsDisposition settings in the NLS.cfg file determine whether the FileNet Object BLOB and CLIP ID file are retained in the NLS File System.

The system uses the CLIP ID to locate the document when a retrieval is requested.

FNP_CLIPID

The FNP_CLIPID is an index setting that can be stored either in the Index Database or the NLS File System. Where it is stored depends on the SdsDisposition setting in the NLS.cfg file.

If the Index Database is used, you must configure the Document Class to contain the Index field FNP_CLIPID to access this attribute.

FNP_CLIPID is a 65 character 'string' type hard coded name. Default attributes must be set for this index value.

Note To define user indexes (FNP_ARCHIVE and FNP_CLIPID) and create document classes that use these index fields, use the Database Maintenance application in the FileNet Application Executive. Refer to the *System Administrator's Handbook* for more information.

FileNet Specific Centera Object Format

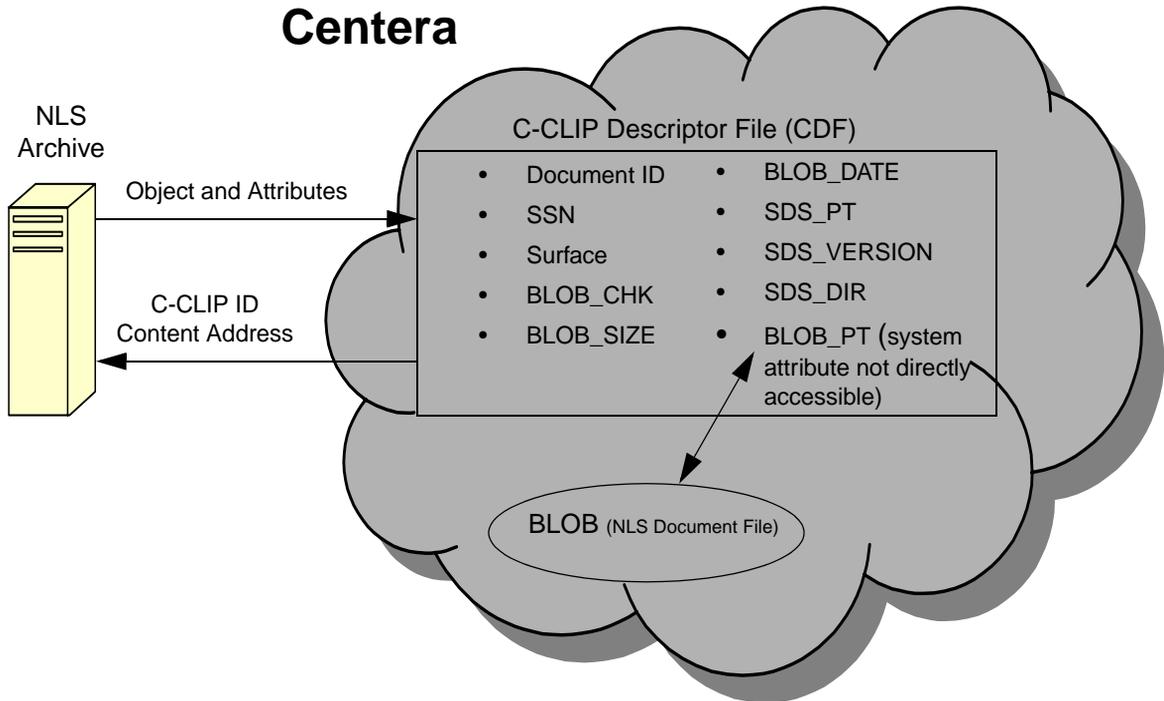
The NLS Archive storage function also converts the documents into a FileNet BLOB (Binary Large Object) file, and requests storage of the object. During this process, the NLS Archive program temporarily stores IS document images in an NLS File System, and then later migrates them to a Centera object repository using Centera API calls.

FileNet stores specific meta-data with each document in the C-CLIP Descriptor File (CDF) and BLOB which contains the document. The BLOB includes header info and the TOC to each document image page.

A Centera top tag with the FileNet name is created for each document with the following tag attributes:

Attribute Name	Example	Description
F_DOCNUMBER	123456789	Original Doc Id
SSN	120000002	Original System Serial Number
SURFACE	5002	Original Primary Surface Number
BLOB_CHK	3245643	BLOB Checksum Generated by NLS
BLOB_SIZE	32768	BLOB Size
BLOB_DATE (Not Yet Implemented)		UTC (Coordinated Universal Time) of BLOB write
SDS_PT (Not Yet Implemented)		<ul style="list-style-type: none"> • BLOB Page Table, fixed array of 1000 unsigned longs containing page offsets (this allows direct access to each page in the document) • First value is page count; first page offset is always 0. • Page count = 0 for Version 0, NLS format since the page table is in the BLOB already. • Max. 4K embedded BLOB
SDS_VERSION	0	<ul style="list-style-type: none"> • 0 = ISCC (NLS Format) • 1 = SDS (IMS Format)
SDS_DIR (Not Yet Implemented)		Document Index Record (DIR), embedded BLOB. This TAG attribute stores the DIR in a BLOB format.

When the top tag is created, a single BLOB object with the document file is written for each document as illustrated in the diagram below.



Additional Internal Buffer Counter Information

This section describes the following three Internal Buffer counters.

Whole Document BLOB Fits into the Internal Buffer

The BLOB size is part of the CLIP meta-data. If the BLOB size is less than or equal to the internal buffer size, the whole BLOB is read into the internal buffer with single BLOB Read call. Using a single BLOB

read maximizes system performance, so it is important to have DBP buffer size equal to the BLOB size whenever possible. A counter is maintained for the number of requests where the whole BLOB fits into the internal buffer. This counter is displayed through the DOC_tool program.

When the whole BLOB fits into the internal buffer, all the pages of the document starting with the first requested page are written to page cache. For example, if we have a 100 page document and the entire document fits into the internal buffer, and page 5 (starting page = 5 and requested pages = 1) is requested, page 5 to 100 are automatically written to cache. This is different from how the IS software retrieves documents from MSAR or optical.

This scenario was created for the following three reasons:

- The internal buffer cannot exceed 1 MB, so reading the whole document into page cache should not create a cache fullness problem. If this is an issue, then the cache is probably too small or the DBP size needs to be reduced.
- A retrieval of a multiple page document that is relatively small almost always results in retrieving more than one page. The IS software automatically does mid-priority retrievals.
- The overhead of doing BLOB reads is high, so if multiple BLOB reads can be combined it improves performance.

Requested images Fit into the Internal Buffer

If the whole document can not fit into the internal buffer, the header and Table Of Contents (TOC) information is read into the internal buffer with one BLOB Read. The TOC includes the offset and size of each page of the document.

The SDS_CSAR_reader process determines if the requested pages fit in the internal buffer. If the requested pages fit into the buffer, one BLOB read is done to read all the requested pages. Only the requested pages are written to page cache. For example, if pages 2 to 9 have been requested from a 100 page document and the whole document does not fit into the internal buffer, pages 2 to 9 are read in a single BLOB read and pages 2 to 9 are written to page cache. This results in two BLOB reads. The first BLOB read gets the header and TOC information to find out if the requested pages fit into the internal buffer, and the second BLOB read determines the offset of the first requested page.

Requested Pages Do Not Fit into the Internal Buffer

If the requested pages do not fit in the internal buffer, the header and TOC information are read into the internal buffer with one BLOB read, and each requested page requires at least one BLOB read.

In the case where the individual page exceeds the internal buffer, multiple BLOB reads are required per page. Only requested pages are read to the internal buffer and written to cache.

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 EMC Centera discussed in the following sections.

Set up the NLS software for EMC Centera

In [Chapter 2, “Installing NLS Software,” on page 31](#) you installed the NLS software on the Image Services Root/Index server.

Verify the NLS.cfg File

In [Chapter 3, “Configuring NLS Software,” on page 38](#) you created the NLS.cfg file for your Centera repository.

PgmAttribute

In the PgmAttribute section, make sure that the appropriate keyword values have been set.

- Make sure the **WorkingDirectory** and **LogDirectory** are set appropriately.
- Make sure the **SdsDevice** is set to “**Centera**”

SDS-Keywords

In the SDS-Keywords section, make sure that appropriate keyword values have been set

SdsDevice

Provides the type of interface for Single Document Storage.

Example:

```
SdsDevice="Centera"
```

SdsDomain

Provides a single or list of DNS and/or IP Address information to the EMC Centera system.

Note The SdsDomain name in the NLS.cfg file must be the same as the domain name in the sds_conf file.

Example:

```
SdsDomain="212.3.248.36"
```

SdsBufSize

Sets the buffer size in K-Bytes (KB) for IO to Centera (Default=10240, Minimum=1, Maximum=10240)

Example:

```
SdsBufSize=10240
```

SdsTimeout

Sets the network timeout value to Centera in seconds (Default=120, Minimum=0, Maximum=600)

Example:

```
SdsTimeout=120
```

SdsRetry

Sets the number of retries to Centera (Default=3, Minimum=0, Maximum=99)

Example:

```
SdsRetry=3
```

SdsSleep

Sets the length of sleep between retries, in seconds (Default=1, Minimum=1, Maximum=100)

Example:

```
SdsSleep=1
```

SdsRetention

Sets the document retention to adhere to a “Compliance” or “Non Compliance” model. The Image Services software controls this document setting. The **SdsRetention** keyword value in the NLS.cfg file overrides all other default values for retention on the Centera repository unless the FileNet IS **F_DELETEDATE** and/or **F_ARCHIVEDATE** values are set. (Specifics to this exception are defined below.)

Note

If the FileNet IS software is not configured correctly for your site’s environment model on a Centera CE or CE+ model repository, the document objects could end up with non-compliance settings.

0 = set no retention

When setting this value to zero, FileNet sets the document retention to 5 minutes before the “no retention” value setting.

-1 = set to infinite retention (WORM)

All documents entered with **SdsRetention=-1** are WORM compliant for an infinite length of time.

number greater than 0 (zero)

Sets retention using input value in hours. This minimum setting is used only when the IS **F_DELETEDATE** and/or **F_ARCHIVEDATE** are not present and the levels of CentraStar (v3.1.1 and above) and the Centera SDK (v3.1.1 and above) cannot support EBR (Event Based Retention).

If either **F_DELETEDATE** or **F_ARCHIVEDATE** are set, this date takes precedence over the **SdsRetention**=(number greater than 0) NLS.cfg keyword setting. If neither dates are set, and the environment supports EBR, the object will be assigned a retention status of EBR. Otherwise, the input value is used.

If both dates are set, the **F_DELETEDATE** takes precedence over the **F_ARCHIVEDATE** and the **SdsRetention**=(number greater than 0) NLS.cfg keyword setting. A date can also be set instead of an hour value. This provides for all incoming documents to have the same retention date, allowing for extension of the retention or deletion of the object by SDS_tool in a scheduled maintenance event.

Valid date formats are either "**yyyy/mm/dd**" or "**mm/dd/yyyy**". If a date that has passed is used, the retention will default to the date of archiving.

Examples:

If **SdsRetention**=48 and **F_DELETEDATE** and **F_ARCHIVEDATE** are **not** set. The document has retention set for 2 days.

If **SdsRetention**=48 and **F_DELETEDATE** is set, the 48 hours is ignored and retention is set for the amount of hours based on **F_DELETEDATE**.

If SdsRetention="12/31/2009" (or SdsRetention="2009/12/31") and F_DELETEDATE and F_ARCHIVEDATE are NOT set, the document has retention set for the number of days/hours/minutes/seconds until 12/31/2009.

SdsDisposition

Sets the disposition of the FileNet Object (a FileNet document with index values) in the Working Directory as well as the location of the Clip ID information.

FS – (Default) This causes the CLIP ID meta file to be written to the same file system location as the FileNet object within the Working Directory. Its name is {docid}.clipid. It contains the CLIPID returned from Centera for the object's Clip.

IDX – This causes the CLIPID to be entered into the Index database, in the Index **FNP_CLIPID** if it exists for this document class.

DEL – This causes the FileNet object file in the Working Directory to be deleted after the CLIPID has been saved where specified by the FS or IDX setting.

All three of these dispositions can be specified. Either **FS** or **IDX** must be specified to allow **DEL** to be used. If **IDX** and **DEL** are specified, the Index **FNP_CLIPID** must exist for the class and the index update must succeed before the delete of the FileNet object takes place. Likewise, a successful write and verify of the meta-file are required if **FS** and **DEL** are specified. When all three are requested, both the Index restrictions and the NLS File System restrictions are enforced before the delete of the FileNet object takes place.

Table of combinations:

- ‘meta file’ refers to the {docid}.clipid file containing a 65 byte string which is the Centera CLIPID for this object.
- ‘object file’ refers to the {docid} FileNet object file in the WorkingDirectory path which is the BLOB stored on Centera.
- Updating the IndexDB causes the 65-byte string CLIPID to be entered into the FNP_CLIPID index field for this document id.

Combination	Explanation
FS	<ul style="list-style-type: none"> • Retain ‘meta file’ - CLIP ID file (<docid>.clipid) • Do not update IndexDB - No FNP_CLIPID Index value • Do not delete ‘object file’ - The FileNet Object document file on the NLS File System is not deleted after it has been stored to Centera.
IDX	<ul style="list-style-type: none"> • Update IndexDB - FNP_CLIPID User Index • Do not retain ‘meta file’ - CLIP ID file • Do not delete ‘object file’ - The FileNet Object document file on the NLS File System is not deleted after it has been stored to Centera.
IDX, FS	<ul style="list-style-type: none"> • Update IndexDB - FNP_CLIPID User Index • Retain ‘meta file’ - CLIP ID file • Do not delete ‘object file’ - The FileNet Object document file on the NLS File System is not deleted after it has been stored to Centera.
IDX,DEL	<ul style="list-style-type: none"> • Update IndexDB - FNP_CLIPID User Index • Delete ‘object file’ - The FileNet Object document file on the NLS File System is deleted after it has been stored to Centera. • Do not retain ‘meta file’ - CLIP ID file

Combination	Explanation
FS, DEL	<ul style="list-style-type: none">• Retain 'meta file' - CLIP ID file• Delete 'object file' - The FileNet Object document file on the NLS File System is deleted after it has been stored to Centera.• Do not update IndexDB - FNP_CLIPID User Index
FS, IDX, DEL	<ul style="list-style-type: none">• Retain 'meta file' - CLIP ID file• Delete 'object file' - The FileNet Object document file on the NLS File System is deleted after it has been stored to Centera.• Update IndexDB - FNP_CLIPID User Index

For example:

SdsDisposition="FS"

Your NLS.cfg file might look similar to the following:

```

;*****
; This module contains the configuration information for EMC Centera. *
;*****
; $Author: giffj $
; $Date: 10/.0/.2 .1:.2:.4 $
; $Header: NLS_cfg,v 1.4 10/.0/.2 .1:.2:.4 giffj Exp $
; $Locker: $
; $Revision: 1.4 $
; $Source: /usr/ProfServ/ISR_200/NLS/src/RCS/NLS_cfg.sample,v $
; $State: Exp $
;*****
; INSTALL HISTORY:
;          INSTALLED BY          DATE          DESCRIPTION
;-----
;*****
; Logon Information
;*****
LogonAttribute {
    UserName="SysAdmin"
    PassWord="SysAdmin"
    Domain="serverA"
    Organization="YourCo"
}
; General information
PgmAttribute {
    WorkingDirectory=/NLS/EMC_Centera
    LogDirectory=/fnsw/local/logs/EMC_Centera
; Control for set size
; MaxFetchDocs={Default=1000}
; MaxQueryDocs={Default=1000}
; MaxFetchChildren controls the maximum number of concurrent Add2Q processes
; MaxFetchChildren={Default=3}
;
; Control for set size
; MaxFetchDocs={Default=1000}
(continued on next page)

```

```

(continued from previous page)
;   MaxQueryDocs={Default=1000}
;   MaxFetchChildren controls the maximum number of concurrent Add2Q processes
;   MaxFetchChildren={Default=3}
;
;   MinDocId=(Default=100000)
;   MaxDocID=(Default=3999999999)
;   Walkback=(Default=100000)
;   AppMaxDocID=(Default=3999999999)
;   MinDate=(Default=980101)
;
;   CleanCache=(Default=False)
;   Timing=(Default=False)
;
;   The following time based keywords all have a minimum of 1 seconds
;   and a maximum of 3600 seconds. If set to zero or a negative number
;   it will default to 5 seconds.
;   FetchSleep=(in sec, default=600)
;   ArchSleep=(in sec, default=600)
;   FetchTimeout=(in sec, default=3600)
;
;   Display format for the year %Y (default) is 'xxxx', %y is 'xx'.
;   This effects displaying of a date only and extends the length of
;   file names.
;   YearFmt="%Y"
;   Format is Year, month, day. Year can be 2(19xx) or 4 digits
;   NOTE: Values that are out of range take on max value for that item
;
;   SDS-Keywords
;   Provides the type of interface for SDS
;   Centera = EMC/Centera Protected Storage
;   Tivoli = IBM/Tivoli Storage Manager
;   Snaplock = NetApps-style Protect Storage File System
;               (SnapLock,StorageTek-Trinity,PermaBit, etc...)
;   Filesystem= Non-Protected Storage File System
;               (StorageTek-ASM)
;   SdsDevice={Default="Filesystem"}
;   SdsDevice="Centera"
(continued on next page)

```

```
(continued from previous page)
; Provides the DNS and/or IP information for SDS device.
SdsDomain="212.3.248.36"
; Sets the buffer size(k) for IO to SDS device.
; SdsBufSize={Default=10240}
; Sets the network timeout value to SDS device, in seconds.
; SdsTimeout={Default=120}
; Sets the number of retries to SDS device.
; SdsRetry={Default=6}
; Sets the length of sleep between retries, in seconds.
; SdsSleep={Default=120}
; Sets the Retention policy for FileNet objects where 0 or No Retention
; is the Default.
; 0 = No Retention mode (actually set to 5 mins)
; -1 = Infinite Retention mode
; 1 or greater than 1 = set retention using input value in hours
; NOTE: 875520 = 100 yrs (100 yrs * 12 mths * 30.4 dys * 24 hrs)
; See Documentation for details
SdsRetention={Default=0}
; OR a date is valid in the form of "yyyy/mm/dd" or "mm/dd/yyyy"
; SdsRetention="2005/08/23"
; Sets the disposition of the FileNet object (a FileNet
; document with index values) in the Working directory
; as well as the location of the clip id information
; SdsDisposition={Default="FS"}
SdsDisposition="IDX,FS,DEL"
;END OF CONFIGURATION FILE
```

Return to [Chapter 4, “Configuring SDS File System Retrieval,” on page 61](#) to create and configure the sds_conf file.

Verify the sds_conf file

In [Chapter 4, “Configuring SDS File System Retrieval,” on page 61](#) you created the sds_conf file for your EMC Centera device. Make sure that it looks similar to the following example:

```
SDS 1 {
  sds_name      NLS SDS
  worker        SDS_CSAR_reader 20
  info          "WorkingDirectory=/NLS/EMC_Centera domain=212.3.248.36"
  sds_priority  medium
  debug         off
}
```

If your EMC Centera device has more than one physical volume, the sds_conf file might look like this:

```
SDS 1 {
  sds_name      NLS SDS
  worker        SDS_CSAR_reader 20
  info          "WorkingDirectory=/NLS2/EMC_Centera;/NLS/EMC_Centera domain=212.3.248.36"
  sds_priority  medium
  debug         off
}
```

Whenever you add a new physical volume, add its file system to the beginning of the WorkingDirectory list. The SDS reader searches for content to retrieve from the first volume listed in the WorkingDirectory series. If it doesn't find it there, it searches the second volume listed, and so on. The new volume will contain the newest information, which will most likely be retrieved more frequently than from older volumes.

Note The domain in the sds_conf file must be the same as the SdsDomain in the NLS.cfg file.

Set up Centera Pool Entry Authorization (PEA) File

You might need to set up Centera's PEA file to activate the security for accessing the Centera device. To do this, complete the following steps:

- 1 Go to the location where you copied the Centera PEA file (normally called `rwe.pea`) on the C drive in the `centera` directory: for example, **C:\centera\rwe.pea**.
- 2 in the `rwe.pea` file, make sure the following parameters are set as shown below, depending on your site's needs.
 - If the security file is going to be used, the following parameters **must** be set:

Read: **ON**
Write: **ON**
Exist: **ON**
 - To allow the Image Services software to remove expired objects, the following parameter **must** be set:

Delete: **ON**
- 3 In the `NLS.cfg` file, augment the `SdsDomain` line to include multiple IP addresses, the path, and the question mark (?) to the PEA file, as shown below:

```
...  
; Provides the DNS and/or IP information for SDS device.  
SdsDomain="212.3.248.36,10.6.7.8?c:\centera\rwe.pea"  
; Sets the buffer size(k) for IO to SDS device.  
...
```

For more information on PAI modules, refer to the *Centera Programmer's Guide*, P/N 0690011217 available through EMC.

- 4 In the `sds_conf` file, augment the `info` line to include multiple IP addresses, the path, and the question mark (?) to the PEA file, as shown below:

```
SDS 1 {
  sds_name      NLS SDS
  worker        SDS_CSAR_reader 20
  info          "workingdirectory=/NLS/EMC_Centera
domain=212.3.248.36,10.6.7.8?c:\centera\rwe.pea"
  sds_priority  medium
  debug         off
}
```

Restart the Image Services Software

Restart the Image Services software on all servers to put the `sds_conf` file changes into effect.

Define a Document Class

NLS Archive modules update an index field called `FPN_ARCHIVE`. Therefore, any Document Class that is to be archived must have this index field included in its definition.

- `FPN_ARCHIVE` is a user defined Index of 'date' type. It must have default attributes.

When the NLS Archive modules are run using an EMC Centera storage unit, you can store the CLIPIDs in the Index database, or in the NLS File System.

If the Index database is used, you must configure the Document Class to contain the Index field FNP_CLIPID to access this attribute.

- FNP_CLIPID is a 65 character 'string' type with default values.

Note To define and create document classes that use the above index fields, use the Database Maintenance application in the FileNet Application Executive. Refer to the *System Administrator's Handbook* for more information.

Complete the following steps on the Image Services Root/Index server:

- 1 Through the Database Maintenance function of the Application Executive (Xapex), create a family that uses MSAR or optical disk. Create two indices:
 - **FNP_CLIPID**, type string, 65 character length
 - **FNP_ARCHIVE**, type date
- 2 Create a documentation class that uses the two indices.
- 3 Commit documents to the document class.
- 4 Run “NLS_Archive -f <docidfile>” or “NLS_start” so the documents are archived to the Centera repository.
- 5 Verify that the documents are archived by checking that the FNP_ARCHIVE field is not NULL. Verify that the FNP_CLIPID is not null. It actually has the following format:

```
strain(fnsw)/usr/ProfServ/TestEnv/NLStest/11001120/00/00/52/0C/00/1A  
1B. cat 1711103.clipid
```

```
8D4AHD5N6GLHMeCHKGOVUB5UAS6 11001120:21004:1711103
```

where:

8D4AHD5N6GLHMeCHKGOVUB5UAS6 is the CLIPID

11001120 is the ssn

21004 is the surf_id

1711103 is the doc_id

This is an example of a CLIPID returned from Centera (viewed from the NLS file system or alternatively via IDM Desktop (FNP_CLIPID field)).

- 6 Return to [Chapter 5, “Administering and Running NLS,” on page 70](#) for information on running the NLS_Archive function and archiving documents using the EMC Centera device.

Appendix D – SnapLock Storage and Retrieval

This appendix contains information specific to the SnapLock storage and retrieval (SSAR) option for FileNet Image Services.

What is SnapLock?

SnapLock software runs on NetApp NearStore, fabric-attached storage (FAS) and network-attached storage (NAS) devices. All or a portion of any device can be used for SnapLock magnetic WORM volumes, which permits the conversion of standard NetApp NAS storage to an application-based storage solution.

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 SnapLock discussed in the following sections.

Set up the NLS software for NetApp SnapLock

In [Chapter 2, “Installing NLS Software,” on page 31](#) you installed the NLS software on the Image Services Root/Index server.

Verify the NLS.cfg File

In [Chapter 3, “Configuring NLS Software,” on page 38](#) you created the NLS.cfg file for your SnapLock device.

PgmAttribute

In the PgmAttribute section, make sure that the appropriate keyword values have been set.

- Make sure the **WorkingDirectory** and **LogDirectory** are set appropriately.
- Make sure the **SdsDevice** is set to “**Snaplock**”

SDS-Keywords

In the SDS-Keywords section, make sure that appropriate keyword values have been set.

SdsDevice

Provides the type of interface for Single Document Storage.

Example:

```
SdsDevice="Snaplock"
```

SdsRetention

Sets the Document retention to adhere to a “Compliance” or “Non Compliance” model. The FileNet IS software controls this document setting. The **SdsRetention** keyword value in the NLS.cfg file overrides all other default values for retention on the file system repository unless the FileNet IS **F_DELETE** and/or **F_ARCHIVE** values are set. (Specifics to this exception are defined below.)

Note If the FileNet IS software is not configured correctly for your site’s environment model, the document objects could end up with non-compliance settings.

0 = set no retention

When setting this value to zero, FileNet sets the document retention to 5 minutes before the “no retention” value setting.

-1 = set to infinite retention (WORM)

All documents entered with **SdsRetention=-1** are WORM compliant for an infinite length of time.

number greater than 0 (zero)

Sets retention using input value in hours. This minimum setting is only used when the IS **F_DELETEDATE** and/or **F_ARCHIVEDATE** are **not** present. If either **F_DELETEDATE** or **F_ARCHIVEDATE** are set, this date takes precedence over the **SdsRetention=(number greater than 0)** NLS.cfg keyword setting. If both dates are set, the **F_DELETEDATE** takes precedence over the **F_ARCHIVEDATE** and the **SdsRetention= (number greater than 0)** NLS.cfg keyword setting.

Valid date formats are either "**yyyy/mm/dd**" or "**mm/dd/yyyy**". If a date that has passed is used, the retention defaults to the date of archiving.

Examples:

If **SdsRetention=48** and **F_DELETEDATE** and **F_ARCHIVEDATE** are **not** set. The document has retention set for 2 days.

If **SdsRetention=48** and **F_DELETEDATE** is set, the 48 hours is ignored and retention is set for the amount of hours based on **F_DELETEDATE**.

If **SdsRetention="12/31/2009"** (or **SdsRetention="2009/12/31"**) and **F_DELETEDATE** and **F_ARCHIVEDATE** are NOT set, the docu-

ment has retention set for the number of days/hours/minutes/seconds until 12/31/2009.

The **read only** attribute is set on the file when the SdsRetention keyword is uncommented in the NLS.cfg file and the value is set to 0, -1, >=1, or date.

To obtain a universal unique Identifier (UUID) file for the FileNet object, the SnapLock keyword must be uncommented in the NLS.cfg file. If the file has been written to a NetApps volume, and NLS_Archive can get the UUID from the NetApps server via the SNMP call, a file containing the UUID is created. It also contains the SSN, Surface Id, and Document Number of the FileNet object. The file has its Modified and Access times, as well as the **read-only** attribute, set the same as the FileNet object.

Your NLS.cfg file might look similar to the following:

```

;*****
; This module contains the configuration information for Snaplock.      *
;*****
; $Author: giffj $
; $Date: 10/.0/.2 .1:.2:.4 $
; $Header: NLS_cfg,v 1.4 10/.0/.2 .1:.2:.4 giffj Exp $
; $Locker: $
; $Revision: 1.4 $
; $Source: /usr/ProfServ/ISR_200/NLS/src/RCS/NLS_cfg.sample,v $
; $State: Exp $
;*****
; INSTALL HISTORY:
;          INSTALLED BY          DATE          DESCRIPTION
;-----
;
;
; (continued on next page)

```

```

(continued from previous page)
;*****
;   Logon Information
;*****
LogonAttribute {
    UserName="SysAdmin"
    PassWord="SysAdmin"
    Domain="serverA"
    Organization="YourCo"
}
;   General information
PgmAttribute {
    WorkingDirectory=/NLS/Snaplock
    LogDirectory=/fnsw/local/logs/Snaplock
;   Control for set size
;   MaxFetchDocs={Default=1000}
;   MaxQueryDocs={Default=1000}
;   MaxFetchChildren controls the maximum number of concurrent Add2Q processes
;   MaxFetchChildren={Default=3}
;
;   Control for set size
;   MaxFetchDocs={Default=1000}
;   MaxQueryDocs={Default=1000}
;   MaxFetchChildren controls the maximum number of concurrent Add2Q processes
;   MaxFetchChildren={Default=3}
;
;   MinDocId=(Default=100000)
;   MaxDocID=(Default=3999999999)
;   Walkback=(Default=100000)
;   AppMaxDocID=(Default=3999999999)
;   MinDate=(Default=980101)
;
;   CleanCache=(Default=False)
;   Timing=(Default=False)
;
(continued on next page)

```

```
(continued from previous page)
; The following time based keywords all have a minimum of 1 seconds
; and a maximum of 3600 seconds. If set to zero or a negative number
; it will default to 5 seconds.
;   FetchSleep=(in sec, default=600)
;   ArchSleep=(in sec, default=600)
;   FetchTimeout=(in sec, default=3600)
;
; Display format for the year %Y (default) is 'xxxx', %y is 'xx'.
; This effects displaying of a date only and extends the length of
; file names.
;   YearFmt="%Y"
;   Format is Year, month, day. Year can be 2(19xx) or 4 digits
;   NOTE: Values that are out of range take on max value for that item
;
; SDS-Keywords
;   Provides the type of interface for SDS
;   Centera = EMC/Centera Protected Storage
;   Tivoli = IBM/Tivoli Storage Manager
;   Snaplock = NetApps-style Protect Storage File System
;               (SnapLock,StorageTek-Trinity,PermaBit, etc...)
;   Filesystem= Non-Protected Storage File System
;               (StorageTek-ASM)
;   SdsDevice={Default="Filesystem"}
;   SdsDevice="Snaplock"
;   Provides the DNS and/or IP information for SDS device.
;   SdsDomain="212.3.248.36"
;   Sets the buffer size(k) for IO to SDS device.
;   SdsBufSize={Default=10240}
;   Sets the network timeout value to SDS device, in seconds.
;   SdsTimeout={Default=120}
;   Sets the number of retries to SDS device.
;   SdsRetry={Default=6}
;   Sets the length of sleep between retires, in seconds.
;   SdsSleep={Default=120}
(continued on next page)
```

```
(continued from previous page)
;   Sets the Retention policy for FileNet objects where 0 or No Retention
;   is the Default.
;   0 = No Retention mode (actually set to 5 mins)
;   -1 = Infinite Retention mode
;   1 or greater than 1 = set retention using input value in hours
;   NOTE: 875520 = 100 yrs (100 yrs * 12 mths * 30.4 dys * 24 hrs)
;   See Documentation for details
;   SdsRetention={Default=0}
;   OR a date is valid in the form of "yyyy/mm/dd" or "mm/dd/yyyy"
;   SdsRetention="2005/08/23"
;   Sets the disposition of the FileNet object (a FileNet
;   document with index values) in the Working directory
;   as well as the location of the clip id information
;   SdsDisposition={Default="FS"}
;   SdsDisposition="IDX,FS,DEL"
;END OF CONFIGURATION FILE
```

Verify the sds_conf file

In [Chapter 4, “Configuring SDS File System Retrieval,” on page 61](#) you created the sds_conf file for your SnapLock device. Make sure that it looks similar to the following example:

```
SDS 1 {
  sds_name      NLS SDS
  worker        SDS_FS_reader 20
  info          "WorkingDirectory=/NLS/Snaplock Domain=212.3.248.36
  sds_priority  medium
  debug        off
}
```

If your SnapLock device has more than one physical volume, the sds_conf file might look like this:

```
SDS 1 {  
  sds_name      NLS_SDS  
  worker        SDS_FS_reader 20  
  info          "WorkingDirectory=/NLS2/Snaplock;/NLS/Snaplock Domain=212.3.248.36"  
  sds_priority  medium  
  debug         off  
}
```

Whenever you add a new physical volume, add its file system to the beginning of the WorkingDirectory list. The SDS reader searches for content to retrieve from the first volume listed in the WorkingDirectory series. If it doesn't find it there, it searches the second volume listed, and so on. The new volume will contain the newest information, which will most likely be retrieved more frequently than from older volumes.

Note The domain in the sds_conf file must be the same as the SdsDomain in the NLS.cfg file.

Restart the Image Services Software

Restart the Image Services software on all servers to put the sds_conf file changes into effect.

Define a Document Class

NLS Archive modules will update an index field called FNP_ARCHIVE. Therefore, any Document Class that is to be archived must have this index field included in its definition.

- FNP_ARCHIVE is a user defined Index of 'date' type. It must have default attributes.

Note To define and create document classes that use the above index fields, use the Database Maintenance application in the FileNet Application Executive. Refer to the *System Administrator's Handbook* for more information.

Complete the following steps on the Image Services Root/Index server:

- 1 Through the Database Maintenance function of the Application Executive (Xapex), create a family that uses MSAR or optical disk. Create one index:
 - **FNP_ARCHIVE**, type date
- 2 Create a documentation class that uses the one index.
- 3 Commit documents to the document class.
- 4 Run **NLS_Archive -f <docidfile>** or **NLS_start** so the documents are archived to the SnapLock device.
- 5 Verify that the document is archived by checking that the FNP_ARCHIVE field is not NULL.
- 6 Return to **[Chapter 5, “Administering and Running NLS,” on page 70](#)** for information on running the NLS_Archive function and archiving documents using the NetApp SnapLock storage device.

Appendix E – 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.

Note HP Integrity servers do not support IBM DR550.

What is IBM Tivoli Storage Manager?

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).

Install the Tivoli API Client Package

The Tivoli API client package software must be installed on all Image Services servers that use SDS to access Tivoli 5.5 libraries.

- 1 To download the TSM software to your local server, go to:

<ftp://ftp.software.ibm.com/storage/tivoli-storage-management/maintenance/client/v5r5/>
- 2 Open the appropriate folder for your operating system, open the folder for your server type (if necessary), and locate the **v550** folder.

- 3 Download the README file to your local server.
 - On UNIX servers, download the tar.Z file (UNIX) and extract the files.
 - On Windows servers, download the .exe file. (This file is a self-expanding archive.)
- 4 The README file contains general installation instructions.

See the *IBM Tivoli Storage Manager Quick Start Guide* for installation instructions and important links. The Quick Start Guide is available on your Quick Start CD or download image.

For detailed installation instructions, see "Chapter 1. Installing Tivoli Storage Manager" of the [***IBM Tivoli Storage Manager Backup-Archive Clients Installation and User's Guide***](#).

Note Install the 32-bit version of the API client, for example:

```
tivoli.tsm.client.api.32bit
```

Install the gsk* modules (AIX servers only)

Depending on your current gsk version, you might also need to install the following two modules:

gskta.rte

SSL Toolkit 32-bit (required by the backup-archive client and its components).

gksa.rte

SSL Toolkit 64-bit (required by the 64-bit client API).

These modules are included in the .tar.Z file for AIX:

5.5.0.0-TIV-TSMBAC-AIX.tar.Z

See the subsection, “Installing the AIX client” in “Chapter 1. Installing Tivoli Storage Manager” of the [***IBM Tivoli Storage Manager Backup-Archive Clients Installation and User's Guide***](#).

Set Up TSM Library and Error Message Files

From the NLS software package, you need to copy the TSM library file and TSM error message files onto the Storage Library servers that want to retrieve documents from an IBM DR550. Then you need to copy the TSM library file and TSM error message file to a specified location. Follow these steps:

- 1 Load and mount the PS NLS CD into the CD-ROM drive on your IS server.
- 2 Depending upon your operating system type complete one of the following bullets:
 - On a UNIX system:
 - a Type the appropriate mount command for your UNIX operating system. For example, on an AIX server you might enter:
mount -rv cdrfs /dev/cd0 /cdrom
 - b Change to the p8im directory on the CD:
cd /cdrom/p8im
 - c Execute the im_install script:

. /im_install

This copies all of the TSM library and error message files to the appropriate locations as show in the table below.

- On a Windows system:
 - a Change directory to the CD-ROM drive (D drive, for example) by entering a command similar to the following in the command window:

```
D:\> \P8IM
```

- b Execute the im_install script and specify the drive (c, for example) where the fnsw directory is located so the TSM library is placed in the correct directory:

```
D:\P8IM> .im_install c
```

This copies all of the TSM library and error message files to the appropriate locations as show in the table below.

Platform	Destination Location
AIX	/fnsw/lib/shobj/libApiDS.a /fnsw/lib/IBM_DR550/en_US/dsmclientV3.cat
HP-UX	/fnsw/lib/shobj/libApiDS.sl /fnsw/lib/IBM_DR550/en_US/dsmclientV3.cat
Solaris	/fnsw/lib/shobj/libApiDS.so /fnsw/lib/IBM_DR550/en_US/dsmclientV3.cat
Windows	C:\fnsw\lib\shobj\tsmapi.dll C:\fnsw\lib\shobj\tsmutil1.dll C:\fnsw\lib\shobj\dsmntapi.dll C:\fnsw\lib\IBM_DR550\dscenu.txt

Define a Document Class

NLS Archive modules update an index field called FNP_ARCHIVE. Therefore, any Document Class that is to be archived must have this index field included in its definition.

- FNP_ARCHIVE is a user defined Index of 'date' type. It must have default attributes.

When the NLS Archive modules are run using an IBM DR550 storage unit, you can store the CLIPIDs in the Index database, or in the NLS File System.

If the Index database is used, you must configure the Document Class to contain the Index field FNP_CLIPID to access this attribute.

- FNP_CLIPID is a 65 character 'string' type with default values.

Note To define and create document classes that use the above index fields, use the Database Maintenance application in the FileNet Application Executive. Refer to the *System Administrator's Handbook* for more information.

Complete the following steps on the Image Services Root/Index server:

- 1 Start the Image Services software on your IS server.

Tip Make sure the sds_conf file that uses the SDS_ISAR_reader does not exist when you bring up the IS software prior to creating the registered node for your IS domain on the DR550. The steps for implementation are: 1) Create indexes and document classes. 2) Run ddexim and FTP

the files to DR550. 3) Create a registered node on DR550. 4) Create an sds_conf file and recycle IS.

- 2 Through the Database Maintenance function of the Application Executive (Xapex), create a family that uses MSAR or optical disk. Create two indexes:
 - **FNP_CLIPID**, type string, 65 character length
 - **FNP_ARCHIVE**, type date
- 3 Create a documentation class that uses the two indexes.

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.

Set up the NLS software for IBM DR550

In [Chapter 2, “Installing NLS Software,” on page 31](#), you installed the NLS software on the Image Services Root/Index server. The NLS_Archive package contains the following two scripts: TSM_SetupSvr and TSM_SetupSvr.awk. To set up an instance of IBM DR550, complete the following steps:

- 1 Run ddexim on the IS server that is going to communicate with the IBM DR550 device:

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

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

- 2 FTP the following files to the IBM DR550: inxdb.<ssn>, TSM_SetupSvr, TSM_SetupSvr.awk.
- 3 On the IBM DR550 console, execute the TSM_SetupSvr script with the <ssn> argument as shown below to generate three files: FN<ssn>, FN<ssn>.scr, and FN<ssn>.rpt - where <ssn> is the ssn of the Root/Index server.

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 domain, policysset and management class and pipe the output to a log file:

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

where <ssn> is the system serial number 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 other(s) 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

(continued on next page)
```

(continued from previous page)

```
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 the above 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 following example)

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 next example, 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 other(s) 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, the **TSM_SetupSvr** script should be re-run and the new document classes added to the IBM DR550 environment by repeating the steps in this section.

Verify the NLS.cfg File

In [Chapter 3, “Configuring NLS Software,” on page 38](#) you created the NLS.cfg file for your IBM DR550 device. Make sure that the appropriate keyword values have been set.

- Make sure the **WorkingDirectory** and **LogDirectory** are set appropriately.
- Make sure the **SdsDevice** is set to "Tivoli".
- Make sure the **SdsDisposition** is “**IDX,FS,DEL**” so that the document access is to the DR550 instead of NLS file system. If you want to retrieve from the file system, the **SdsDisposition** is **IDX, FS**. The **IDX** means CLIP index is saved in the index database. The **FS** means the CLIP index is stored in a CLIP file.
- Make sure the **SdsRetention** value is properly set for your environment and business requirements. By default, with **SdsRetention** not set at all (the line is commented with a ‘), no retention value is enforced, which means it’s not used. The Read-Only state of the document is not triggered by NLS. If the SDS file system has manufacturer defaults implemented, then those defaults are enforced by the file system itself.

Setting the **SdsRetention** value to 0 sets a default retention of 5 minutes (for Tivoli it is 1 day, the smallest allowed). Setting the value to some other value will set the retention as explained in the following paragraphs.

SdsRetention sets the Document retention to adhere to a “Compliance” or “Non Compliance” model. The FileNet IS software controls this document setting. The **SdsRetention** keyword value in the NLS.cfg file overrides all other default values for retention on the IBM DR550 repository unless the FileNet IS **F_DELETEDATE**

and/or **F_ARCHIVEDATE** values are set. (Specifics to this exception are defined below.)

Note If the FileNet IS software is not configured correctly for your site's environment model, the document objects could end up with non-compliance settings.

0 = set default retention

The minimum retention for Tivoli is one day.

-1 = set to infinite retention (WORM)

All documents entered with **SdsRetention=-1** are WORM compliant for an infinite length of time.

number greater than 0 (zero)

Sets retention using input value in hours. This minimum setting is only used when the IS **F_DELETEDATE** and/or **F_ARCHIVEDATE** are **not** present. If either **F_DELETEDATE** or **F_ARCHIVEDATE** are set, this date takes precedence over the **SdsRetention=(number greater than 0)** NLS.cfg keyword setting. If both dates are set, the **F_DELETEDATE** takes precedence over the **F_ARCHIVEDATE** and the **SdsRetention= (number greater than 0)** NLS.cfg keyword setting.

Valid date formats are either "**yyyy/mm/dd**" or "**mm/dd/yyyy**". If a date that has passed is used, the retention defaults to the date of archiving.

Examples:

If **SdsRetention=48** and **F_DELETEDATE** and **F_ARCHIVEDATE** are **not** set. The document has retention set for 2 days.

If SdsRetention=48 and F_DELETEDATE is set, the 48 hours is ignored and retention is set for the amount of hours based on F_DELETEDATE.

If SdsRetention="12/31/2009" (or SdsRetention="2009/12/31") and F_DELETEDATE and F_ARCHIVEDATE are NOT set, the document has retention set for the number of days/hours/minutes/seconds until 12/31/2009.

Your NLS.cfg file might look similar to the following:

```

;*****
; This module contains the configuration information for the IBM DR550.*
;*****
; $Author: giffj $
; $Date: 10/.0/.2 .1:.2:.4 $
; $Header: NLS_cfg,v 1.4 10/.0/.2 .1:.2:.4 giffj Exp $
; $Locker: $
; $Revision: 1.4 $
; $Source: /usr/ProfServ/ISR_200/NLS/src/RCS/NLS_cfg.sample,v $
; $State: Exp $
;*****
; INSTALL HISTORY:
;          INSTALLED BY          DATE          DESCRIPTION
;-----
;*****
; Logon Information
;*****
LogonAttribute {
    UserName="SysAdmin"
    PassWord="SysAdmin"
    Domain="serverA"
    Organization="YourCo"
}
(continued on next page)

```

```
(continued from previous page)
;
;   General information
PgmAttribute {
    WorkingDirectory=/NLS/IBM_DR550
    LogDirectory=/fnsw/local/logs/IBM_DR550
;   Control for set size
;   MaxFetchDocs={Default=1000}
;   MaxQueryDocs={Default=1000}
;   MaxFetchChildren controls the maximum number of concurrent Add2Q processes
;   MaxFetchChildren={Default=3}
;
;   Control for set size
;   MaxFetchDocs={Default=1000}
;   MaxQueryDocs={Default=1000}
;   MaxFetchChildren controls the maximum number of concurrent Add2Q processes
;   MaxFetchChildren={Default=3}
;
;   MinDocId=(Default=100000)
;   MaxDocID=(Default=3999999999)
;   Walkback=(Default=100000)
;   AppMaxDocID=(Default=3999999999)
;   MinDate=(Default=980101)
;
;   CleanCache=(Default=False)
;   Timing=(Default=False)
;
;   The following time based keywords all have a minimum of 1 seconds
;   and a maximum of 3600 seconds. If set to zero or a negative number
;   it will default to 5 seconds.
;   FetchSleep=(in sec, default=600)
;   ArchSleep=(in sec, default=600)
;   FetchTimeout=(in sec, default=3600)
;
(continued on next page)
```

(continued from previous page)

```
;
; Display format for the year %Y (default) is 'xxxx', %y is 'xx'.
; This effects displaying of a date only and extends the length of
; file names.
;   YearFmt="%Y"
;   Format is Year, month, day. Year can be 2(19xx) or 4 digits
;   NOTE: Values that are out of range take on max value for that item
;
; SDS-Keywords
;   Provides the type of interface for SDS
;   Centera = EMC/Centera Protected Storage
;   Tivoli = IBM/Tivoli Storage Manager
;   Snaplock = NetApps-style Protect Storage File System
;               (SnapLock,StorageTek-Trinity,PermaBit, etc...)
;   Filesystem= Non-Protected Storage File System
;               (StorageTek-ASM)
;   SdsDevice={Default="Filesystem"}
;   SdsDevice="Tivoli"
;   Provides the DNS and/or IP information for SDS device.
;   SdsDomain="212.3.248.36"
;   Sets the buffer size(k) for IO to SDS device.
;   SdsBufSize={Default=10240}
;   Sets the network timeout value to SDS device, in seconds.
;   SdsTimeout={Default=120}
;   Sets the number of retries to SDS device.
;   SdsRetry={Default=6}
;   Sets the length of sleep between retires, in seconds.
;   SdsSleep={Default=120}
;
(continued on next page)
```

```
(continued from previous page)
;   Sets the Retention policy for FileNet objects where 0 or No Retention
;   is the Default.
;   0 = No Retention mode (actually set to 1 day for Tivoli)
;   -1 = Infinite Retention mode
;   1 or greater than 1 = set retention using input value in hours
;   NOTE: 875520 = 100 yrs (100 yrs * 12 mths * 30.4 dys * 24 hrs)
;   See Documentation for details
SdsRetention=-1
;   OR a date is valid in the form of "yyyy/mm/dd" or "mm/dd/yyyy"
;   SdsRetention="2005/08/23"
;   Sets the disposition of the FileNet object (a FileNet
;   document with index values) in the Working directory
;   as well as the location of the clip id information
;   SdsDisposition={Default="FS"}
SdsDisposition="IDX,FS,DEL"
;END OF CONFIGURATION FILE
```

Return to [Chapter 4, “Configuring SDS File System Retrieval,” on page 61](#) to create and configure the sds_conf file.

Verify the sds_conf file

In [Chapter 4, “Configuring SDS File System Retrieval,” on page 61](#) you created the sds_conf file for your IBM DR550 device. Make sure that it looks similar to the following example:

```
SDS 1 {
  sds_name      NLS SDS
  worker       SDS_ISAR_reader 20
  info         "workingdirectory=/NLS/IBM_DR550 domain=212.3.248.36"
  sds_priority  medium
  debug        off
}
```

If your IBM DR550 device has more than one physical volume, the `sds_conf` file might look like this:

```
SDS 1 {  
  sds_name      NLS SDS  
  worker        SDS_ISAR_reader 20  
  info          "WorkingDirectory=/NLS2/IBM_DR550;/NLS/IBM_DR550 Domain=212.3.248.36"  
  sds_priority  medium  
  debug         off  
}
```

Whenever you add a new physical volume, add its file system to the beginning of the `WorkingDirectory` list. The SDS reader searches for content to retrieve from the first volume listed in the `WorkingDirectory` series. If it doesn't find it there, it searches the second volume listed, and so on. The new volume will contain the newest information, which will most likely be retrieved more frequently than from older volumes.

Note The domain's IP address in the `sds_conf` file must be the same as the `SdsDomain`'s IP address in the `NLS.cfg` file. The directory paths must also match.

Restart the Image Services Software

Restart the Image Services software on all servers to put the new `sds_conf` file into effect.

Verify the TSM API Environment Variables

The `dsm` files must be set up in order 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 Image Services, you do not need to do much. By default, Image Services creates the `dsm.sys` (UNIX only) and `dsm.opt` (UNIX and Windows) configuration files in the following locations when the Image Services software is brought up, depending upon your operating system:

- `/fnsw/lib/IBM_DR550/dsm.sys` (UNIX)
- `/fnsw/lib/IBM_DR550/dsm.opt` (UNIX - placeholder file)
- `<drive>:\fnsw\lib\DR550\dsm.opt` (Windows)

Note There are both a `dsm.sys` and a `dsm.opt` file 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 180](#).

These files are expected by the initial TSM API call (`dsmSetup`). By default, the APIs work fine. API library error messages are logged to one of the following files:

```
/fnsw/local/logs/IBM_DR550/rdr_TSM_err_yyyymmdd (UNIX)
<drive>:\fnsw_loc\logs\IBM_DR550\rdr_TSM_err_yyyymmdd (Windows)
```

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 Image Services, 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 Image Services 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 from the `sds_conf` file.

<code>Servername</code>	<code>FileNet_SDS</code>
<code>Commmethod</code>	<code>tcpid</code>
<code>Tcpserveraddress</code>	<code>212.3.248.36</code>
<code>EnableArchiveRetentionProtection</code>	<code>yes</code>

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 180](#).

Set the MaxSessions Parameter

It is important to consider how many systems are going to be in communication with the IBM DR550 and 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
                Server URL:
                Crossdefine: Off
                Server Password Set: No
Server Installation Date/Time: 02/26/04    17:40:30
Server Restart Date/Time: 11/09/05    16:01:25
                Authentication: On
                Password Expiration Period: 90 Day(s)
                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 Day(s)
Activity Log Number of Records: 5804
                Activity Log Size: <1 M
Activity Summary Retention Period: 30 Day(s)
                License Audit Period: 30 Day(s)
                Last License Audit: 12/04/05    16:01:46
                Server License Compliance: Valid
                Central Scheduler: Active
                Maximum Sessions: 256
                Maximum Scheduled Sessions: 128
                Event Record Retention Period: 10 Day(s)}
```

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.

Restart the Image Services Software

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

Setup NLS_Archive

- 1 Commit documents to the document class you created earlier.
- 2 Run **NLS_Archive -f <docidfile>** or **NLS_start** so the documents are archived to the DR550.
- 3 Verify that the documents are archived by checking that the FNP_ARCHIVE field is not NULL. Verify that the FNP_CLIPID is not null. It actually has the following format:

```
0:xxxxx:/aa/bb/cc/dd/ee/ff/gg/yyyyy
```

where:

0 is the high order object ID
 xxxxx is the low order object ID
 /aa/bb/cc/dd/ee/ff/gg/ is the file system path
 yyyyyy is the doc_id

For example:

```
cat ./00/00/52/0C/01/05/BE/17153743.clipid

hqpsmb20(fnsw)/usr/ProfServ/TestEnv/NLStest/11001120> cd ./00/00/52/0C/01/05/BE

hqpsmb20(fnsw)/usr/ProfServ/TestEnv/NLStest/11001120/00/00/52/0C/01/05/BE>ll
-rw-rw-r-- 1 fnsw  fnusr  10240 Nov  5 17:07 ./00/00/52/0C/01/05/BE/17153743
-r--r----- 1 fnsw  fnusr   65 Nov 10 10:44 ./00/00/52/0C/01/05/BE/17153743.clipid

hqpsmb20(fnsw)/usr/ProfServ/TestEnv/NLStest/11001120> cat 17153743.clipid
0:3434497:/00/00/52/0C/01/05/BE/17153743 11001120:21004:17153743
```

In the example above, the CLIP ID is in bold text. The 11001120:21004:17153743 after the CLIPID is additional information: ssn:surf_id:doc_id.

- 4 Return to [Chapter 5, “Administering and Running NLS,” on page 70](#) for information on running the NLS_Archive function and archiving documents using the IBM DR550.

Set API Environment Variables for Tracing (Optional)

Both IBM and FileNet software provide tracing methods that you can use for troubleshooting the IBM DR550 device.

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

- You can optionally set the ISAR_CMD_TRACE environment variable to take advantage of FileNet's ISAR API tracing.

Note Be aware that these tracing methods 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, you need to 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 on the My Computer icon > select Properties from the list > select the Advanced tab > click on 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 located 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 it has content 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 sds_conf 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 located in the **<api trace file name>** you specify in the dsm.opt file.

Restart the Image Services Software

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

Set FileNet ISAR Tracing Method

You can set the **ISAR_CMD_TRACE** environment variable if you wish to capture a list of the IBM DR550 API calls to a log file. Note that this method does not provide you with as much detail as the TSM method.

- 1 Make sure the Image Services software is stopped.

initfnsw -y stop

- 2 Set the environment variable by entering one of the following commands:

On UNIX platforms:

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

export ISAR_CMD_TRACE=true

- For C shell users, enter the following command:

setenv ISAR_CMD_TRACE true

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

- a Right-click on the My Computer icon > select Properties from the menu > select the Advanced tab > click on 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 **ISAR_CMD_TRACE** in the Variable Name field and **true** in the Variable Value field and click **OK**.
- 3** Stop the IS software with a **killfns** and restart the Image Services software to put the FileNet ISAR Tracing into effect.

The following is the location of the trace file:

```
/fnsw/local/logs/IBM_DR550/rdr_cmd_trace_yyyymmdd (UNIX)
<drive>:\fnsw_loc\logs\IBM_DR550\rdr_cmd_trace_yyyymmdd (Windows)
```

The following sample is the output of a trace file:

```
2005/10/07_16:03:51<<pid=10847>> dsmSetUp: (0='success')
2005/10/07_16:05:07<<pid=10847>> dsmInitEx: (0='success')
2005/10/07_16:05:07<<pid=10847>> dsmQuerySessOptions: (0='success')
2005/10/07_16:05:07<<pid=10847>> dsmQuerySessInfo: (0='success')
2005/10/07_16:05:08<<pid=10847>> dsmBeginGetData: (0='success')
2005/10/07_16:05:08<<pid=10847>> dsmGetObj: (121='ANS0272I (RC121) The operation
is finished.')
2005/10/07_16:05:08<<pid=10847>> dsmEndGetObj: (0='success')
2005/10/07_16:05:08<<pid=10847>> dsmEndGetData: (0='success')
2005/10/07_16:05:08<<pid=10847>> dsmTerminate: (0='success')
```

- 4** Return to [Chapter 5, “Administering and Running NLS,” on page 70](#) for information on running the NLS_Archive function and archiving documents using the IBM DR550.

Appendix F – Sun StorageTek 5320 Storage and Retrieval

This appendix contains information specific to the SUN 5320 storage and retrieval (SunSAR) option for FileNet Image Services.

What is Sun StorageTek™ 5320?

The Sun StorageTek 5320 NAS (network-attached storage) architecture provides full system redundancy with 4 GB FC (fibre channel) RAID arrays delivering high levels of availability and performance in most open environments.

The Sun StorageTek 5320 NAS Appliance supports:

- Compliance Volumes
- File System Journaling
- Checkpointing
- Remote Mirroring (Replication)
- Clustering

Designed to scale in capacity and performance, this powerful NAS appliance can easily be scaled to 134 TB of raw FC storage, and as much as 224 TB of SATA (serial advanced technology attachment) storage. A choice of FC or SATA drives inside the same system means that information that is frequently accessed could be placed on FC drives for exceptional performance, while archive data could be placed on higher-capacity, less expensive SATA disks. Tape backup is also supported using an NDMP (network data management protocol) client.

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 Sun StorageTek 5320 discussed in the following sections.

Set up the NLS software for Sun StorageTek 5320

In [Chapter 2, “Installing NLS Software,” on page 31](#) you installed the NLS software on the Image Services Root/Index server.

Verify the NLS.cfg File

In [Chapter 3, “Configuring NLS Software,” on page 38](#) you created the NLS.cfg file for your Sun StorageTek device.

PgmAttribute

In the PgmAttribute section, make sure that the appropriate keyword values have been set.

- Make sure the **WorkingDirectory** and **LogDirectory** are set appropriately.
- Make sure the **SdsDevice** is set to “**Sunsar**”

SDS-Keywords

In the SDS-Keywords section, make sure that appropriate keyword values have been set.

SdsDevice

Provides the type of interface for Single Document Storage.

Example:

```
SdsDevice="Sunsar"
```

SdsRetention

Sets the Document retention to adhere to a “Compliance” or “Non Compliance” model. The FileNet IS software controls this document setting. The **SdsRetention** keyword value in the NLS.cfg file overrides all other default values for retention on the file system repository unless the FileNet IS **F_DELETE** and/or **F_ARCHIVE** values are set. (Specifics to this exception are defined below.)

Note

If the FileNet IS software is not configured correctly for your site’s environment model, the document objects could end up with non-compliance settings.

0 = set no retention

When setting this value to zero, FileNet sets the document retention to 5 minutes before the “no retention” value setting.

-1 = set to infinite retention (WORM)

All documents entered with **SdsRetention=-1** are WORM compliant for an infinite length of time.

number greater than 0 (zero)

Sets retention using input value in hours. This minimum setting is only used when the IS **F_DELETE** and/or **F_ARCHIVE** are **not** present. If either **F_DELETE** or **F_ARCHIVE** are set, this date takes precedence over the **SdsRetention=(number greater than 0)** NLS.cfg keyword setting. If both dates are set, the **F_DELETE** takes precedence over the **F_ARCHIVE** and the **SdsRetention= (number greater than 0)** NLS.cfg keyword setting.

Valid date formats are either "**yyyy/mm/dd**" or "**mm/dd/yyyy**". If a date that has passed is used, the retention defaults to the date of archiving.

Examples:

If SdsRetention=48 and F_DELETEDATE and F_ARCHIVEDATE are **not** set. The document has retention set for 2 days.

If SdsRetention=48 and F_DELETEDATE is set, the 48 hours is ignored and retention is set for the amount of hours based on F_DELETEDATE.

If SdsRetention="12/31/2009" (or SdsRetention="2009/12/31") and F_DELETEDATE and F_ARCHIVEDATE are NOT set, the document has retention set for the number of days/hours/minutes/seconds until 12/31/2009.

The **read only** attribute is set on the file when the SdsRetention keyword is uncommented in the NLS.cfg file and the value is set to 0, -1, >=1, or date.

Your NLS.cfg file might look similar to the following:

```

;*****
; This module contains the configuration information for Sun StorageTek.*
;*****
; $Author: giffj $
; $Date: 10/.0/.2 .1:.2:.4 $
; $Header: NLS_cfg,v 1.4 10/.0/.2 .1:.2:.4 giffj Exp $
; $Locker: $
; $Revision: 1.4 $
; $Source: /usr/ProfServ/ISR_200/NLS/src/RCS/NLS_cfg.sample,v $
; $State: Exp $
;*****
(continued on next page)

```

```

(continued from previous page)
; INSTALL HISTORY:
;      INSTALLED BY          DATE          DESCRIPTION
;----- -MM/DD/YYYY- -----
;
;*****
;      Logon Information
;*****
LogonAttribute {
    UserName="SysAdmin"
    Password="SysAdmin"
    Domain="serverA"
    Organization="YourCo"
}
;      General information
PgmAttribute {
    WorkingDirectory=/NLS/Sunsar
    LogDirectory=/fnsw/local/logs/Sunsar
;      Control for set size
;      MaxFetchDocs={Default=1000}
;      MaxQueryDocs={Default=1000}
;      MaxFetchChildren controls the maximum number of concurrent Add2Q processes
;      MaxFetchChildren={Default=3}
;
;      Control for set size
;      MaxFetchDocs={Default=1000}
;      MaxQueryDocs={Default=1000}
;      MaxFetchChildren controls the maximum number of concurrent Add2Q processes
;      MaxFetchChildren={Default=3}
;
;      MinDocId=(Default=100000)
;      MaxDocID=(Default=3999999999)
;      Walkback=(Default=100000)
;      AppMaxDocID=(Default=3999999999)
;      MinDate=(Default=980101)
;
;      CleanCache=(Default=False)
;      Timing=(Default=False)
;
(continued on next page)

```

```

(continued from previous page)
; The following time based keywords all have a minimum of 1 seconds
; and a maximum of 3600 seconds. If set to zero or a negative number
; it will default to 5 seconds.
;   FetchSleep=(in sec, default=600)
;   ArchSleep=(in sec, default=600)
;   FetchTimeout=(in sec, default=3600)
;
; Display format for the year %Y (default) is 'xxxx', %y is 'xx'.
; This effects displaying of a date only and extends the length of
; file names.
;   YearFmt="%Y"
;   Format is Year, month, day. Year can be 2(19xx) or 4 digits
;   NOTE: Values that are out of range take on max value for that item
;
; SDS-Keywords
;   Provides the type of interface for SDS
;   Centera = EMC/Centera Protected Storage
;   Tivoli = IBM/Tivoli Storage Manager
;   Snaplock = NetApps-style Protect Storage File System
;               (SnapLock,StorageTek-Trinity,PermaBit, etc...)
;   Sunsar = Sun StorageTek 5320
;   Filesystem= Non-Protected Storage File System
;               (StorageTek-ASM)
;   SdsDevice={Default="Filesystem"}
;   SdsDevice="Sunsar"
;   Provides the DNS and/or IP information for SDS device.
;   SdsDomain="212.3.248.36"
;   Sets the buffer size(k) for IO to SDS device.
;   SdsBufSize={Default=10240}
;   Sets the network timeout value to SDS device, in seconds.
;   SdsTimeout={Default=120}
;   Sets the number of retries to SDS device.
;   SdsRetry={Default=6}
;   Sets the length of sleep between retries, in seconds.
;   SdsSleep={Default=120}
(continued on next page)

```

```
(continued from previous page)
;   Sets the Retention policy for FileNet objects where 0 or No Retention
;   is the Default.
;   0 = No Retention mode (actually set to 5 mins)
;   -1 = Infinite Retention mode
;   1 or greater than 1 = set retention using input value in hours
;   NOTE: 875520 = 100 yrs (100 yrs * 12 mths * 30.4 dys * 24 hrs)
;   See Documentation for details
;   SdsRetention={Default=0}
;   OR a date is valid in the form of "yyyy/mm/dd" or "mm/dd/yyyy"
;   SdsRetention="2005/08/23"
;   Sets the disposition of the FileNet object (a FileNet
;   document with index values) in the Working directory
;   as well as the location of the clip id information
;   SdsDisposition={Default="FS"}
;   SdsDisposition="IDX,FS,DEL"
;END OF CONFIGURATION FILE
```

Verify the sds_conf file

In [Chapter 4, “Configuring SDS File System Retrieval,” on page 61](#) you created the sds_conf file for your Sun StorageTek device. Make sure that it looks similar to the following example:

```
SDS 1 {
  sds_name      NLS SDS
  worker        SDS_FS_reader 20
  info          "WorkingDirectory=/NLS/Sunsar Domain=212.3.248.36
  sds_priority  medium
  debug         off
}
```

If your SunSAR device has more than one physical volume, the `sds_conf` file might look like this:

```
SDS 1 {  
  sds_name      NLS_SDS  
  worker        SDS_FS_reader 20  
  info          "WorkingDirectory=/NLS2/Sunsar;/NLS/Sunsar Domain=212.3.248.36"  
  sds_priority  medium  
  debug         off  
}
```

Whenever you add a new physical volume, add its file system to the beginning of the `WorkingDirectory` list. The SDS reader searches for content to retrieve from the first volume listed in the `WorkingDirectory` series. If it doesn't find it there, it searches the second volume listed, and so on. The new volume will contain the newest information, which will most likely be retrieved more frequently than from older volumes.

Note The domain in the `sds_conf` file must be the same as the `SdsDomain` in the `NLS.cfg` file.

Restart the Image Services Software

Restart the Image Services software on all servers to put the `sds_conf` file changes into effect.

Define a Document Class

NLS Archive modules will update an index field called `FPN_ARCHIVE`. Therefore, any Document Class that is to be archived must have this index field included in its definition.

- FNP_ARCHIVE is a user defined Index of 'date' type. It must have default attributes.

Note To define and create document classes that use the above index fields, use the Database Maintenance application in the FileNet Application Executive. Refer to the *System Administrator's Handbook* for more information.

Complete the following steps on the Image Services Root/Index server:

- 1 Through the Database Maintenance function of the Application Executive (Xapex), create a family that uses MSAR or optical disk. Create one index:
 - **FNP_ARCHIVE**, type date
- 2 Create a documentation class that uses the one index.
- 3 Commit documents to the document class.
- 4 Run **NLS_Archive -f <docidfile>** or **NLS_start** so the documents are archived to the Sun StorageTek 5320 device.
- 5 Verify that the document is archived by checking that the FNP_ARCHIVE field is not NULL.
- 6 Return to **[Chapter 5, “Administering and Running NLS,” on page 70](#)** for information on running the NLS_Archive function and archiving documents using the Sun StorageTek 5320 storage device.

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.

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

SC19-2687-00

