



Technical report:
**IBM FileNet Image Services and IBM
System Storage N series
with SnapLock**

Data Archival

• • • • • • • • •

Document NS3335-0

October 12, 2007



Table of contents

Abstract	3
Overview	3
Purpose and Scope	4
Introduction	4
FileNet IS 4.0	4
Assumptions	4
Infrastructure	5
FileNet IS 4.0 on Windows 2000 Server.....	5
FileNet IS 4.0 and the Oracle Database Environment.....	6
Installation Tasks	6
Installing FILENET IS Software	6
Configuring the MSAR Storage Library	6
Upgrading to IS 3.6 ESE from IS 3.6	7
Using the MSAR Storage Library.....	7
Configuring MSAR surfaces	8
Integration of IS Connector for SnapLock	9
Preinstallation Requirements	9
FileNet IS Server Configuration.....	10
Install IS Toolkit (WAL)	10
IBM N series Configuration Review.....	10
Windows Server Platform.....	10
Verify the Installation Directory.....	11
Configuring SSAR Storage Connector Software.....	11
Configuration of the hosts File.....	11
Environment Variable PATH Setup	11
Ownership and Permissions on UNIX Platforms.....	11
SnapLock Configuration Information	12
Setting the IS User and Password in a Secure Method	13
Destination Configuration	13
Attribute Settings	14
Index Database Settings	14
Archiving Documents to a SnapLock Volume	14
Creating/Reporting Family, Class, and Index.....	15
FileNet IDM Desktop Client	17
Adding Documents to FileNet Library and Archiving to a SnapLock Volume.....	18
Starting the Database Server	18
Starting FileNet IS	18
Conclusions and Caveats	24
Trademarks and Special Notices	25



Abstract

IBM System Storage N series with SnapLock is the IBM implementation of high-performance disk-based, nonerasable, nonrewritable, magnetic WORM storage. The primary objective of this combination is to provide secure, disk-based storage-enforced data retention functionality, via open file protocols, for protecting data in regulatory environments. Other configurations of SnapLock can be deployed for nonregulated or more flexibly regulated environments. IBM FileNet Image Services software supports SnapLock for archiving data to a SnapLock volume. This paper discusses the procedure for integrating FileNet Image Services and IBM N series with SnapLock unified storage.

Overview

To manage content successfully, an enterprise must be able to deliver fast access to fixed content objects that may number from the hundreds of thousands to the hundreds of millions. It must also be able to support rapid content growth while still providing adequate performance and high data availability. Fixed content typically includes electronic documents, faxes, images, and rich media files for large numbers of users. In addition to these requirements, strong security and data protection solutions must be part of the permanent storage solution for such critical business information. These stringent requirements demand an efficient content management solution.

FileNet® has a range of products that provide business solutions for document management and content management needs. The most common and best known of these is FileNet Image Services (IS). Optical storage and retrieval (OSAR) was once the preferred (and only) storage media choice for IS, but improvements in technology led FileNet to consider the advantages of magnetic disk media over optical storage, and FileNet IS now supports magnetic disk media using its Magnetic Storage and Retrieval (MSAR) software.

IBM® System Storage™ N series devices are integrated hardware- and software-based network storage systems that serve data using any of the multiple storage protocols supported for both storage area network (SAN) and network-attached storage (NAS) environments. These storage devices are known as IBM N series filers or IBM N series fabric-attached storage (FAS) servers and act on application or server requests for data and process those requests by writing data to or retrieving data from the storage system. Fast and reliable operations, as well as advanced data protection options, are made possible by the IBM N series microkernel operating system called Data ONTAP®.

Write-once-read-many (WORM) media allow data to be written only once, and written data can never be overwritten or erased. In some cases this is a property of the physical media (such as with WORM optical platters), while in other cases the physical media is rewritable but integrated hardware and software codes controlling access to the media prevent such overwrites (such as with WORM magnetic tape and disk-based SnapLock® software).

In regulated environments, the regulating body mandates that business-generated data be archived on WORM media in order to maintain a nonerasable and nonrewritable electronic paper trail that can be used for the purposes of discovery or investigation. Even in nonregulated environments, customers are finding advantages in protecting their data using secure storage while exploiting the benefits offered by WORM capability on magnetic disks. IBM N series offers two types of data protection: SnapLock Compliance and SnapLock Enterprise. SnapLock Enterprise allows System Administrators to manage



WORM volume storage—for example, by deleting all the data on a volume in order to recover storage space. SnapLock Compliance does not allow modification or deletion of any unexpired WORM volumes.

Purpose and Scope

This paper describes the steps necessary to integrate FileNet IS software with IBM N series with SnapLock to archive documents as nonerasable, nonrewritable WORM-volume storage. Configurations and procedures covered in this report will also apply to using IS products on UNIX® platforms with IBM N series storage using NAS architecture. The information in this document should be taken only as a starting point, and customers should consult FileNet and IBM to determine the actual configuration needs of their environments. FileNet IS archives to SnapLock volumes using Common Internet File System (CIFS) and Network File System (NFS) protocols. SnapLock is not supported on block devices that use Fibre Channel (FC). FileNet IS works at the document level, hence the requirement for CIFS or NFS.

Introduction

Content management is a complex challenge. Integration of FileNet IS 4.0 with IBM N series storage offers an effective solution. This section describes the configurations used to integrate IS 4.0 and IBM N series unified storage.

FileNet IS 4.0

FileNet Enterprise Content Management (ECM) solutions allow customers to build and sustain competitive advantage by managing content throughout their organizations, automating and streamlining their business processes and providing the full spectrum of connectivity needed to simplify both critical and everyday decision making. FileNet ECM solutions deliver a comprehensive set of capabilities that integrate with existing information systems to provide cost-effective solutions that solve real-world business problems. FileNet IS 4.0 delivers faster access to a large number of fixed objects (billions of them) such as documents, reports, print streams, faxes, e-mail, and rich media content. FileNet IS provides organizations with the ability to:

- Improve the operational effectiveness of content information.
- Allow high availability and yet ensure the security of information assets.
- Increase the content access experience of users.
- Continue access while preventing data corruption and providing data security.

Currently, FileNet IS configuration is supported on Microsoft® Windows 2000. FileNet supports both UNIX and Windows® platforms; supported UNIX platforms are Sun™ Solaris™ 8, HP/UX, and IBM AIX®. The unified IBM N series storage solution enables industry-leading enterprise database and business applications that provide key benefits to customers:

- Operational efficiency
- Resource utilization
- Support for technology and partnership experience.

Assumptions

The reader should be familiar with the operation of IBM N series storage devices and SAN and NAS concepts. The reader needs a system administrator's knowledge of FileNet IS server software. All needed



FileNet IS components should be installed and configured. The term *filer* refers to the IBM N series storage equipment. SnapLock works either on a fabric-attached server or on IBM System Storage N series with NearStore® feature storage devices using either CIFS or NFS protocols.

Infrastructure

A sample system configuration was selected to install and test the IS Connector for SnapLock (ISCS) product to validate the information in this document. The purpose was to show the procedure for integrating ISCS and FileNet IS with an IBM N series storage device using the SnapLock protocol. ISCS is now known as SnapLock Storage and Retrieval (SSAR).

FileNet also supports Windows 2003 and later, our test setup, however, used Oracle 9.2.0.5 database server and FileNet IS 4.0 on a Windows 2000 host with a filer as listed below:

- FileNet IS 4.0 Server, Windows 2000 Server
- IBM N series filer or IBM N series with NearStore with CIFS, iSCSI and FCP/NFS capability
- Network for management and dataflow
- FileNet IS for environment settings such as user accounts, software, storage appliance mount points and/or SAN storage.

This configuration used the following infrastructure to install the IS software:

- The name of the filer `boy`
- The filer command prompt is shown as `boy>`
- The Solaris host name `beavis`
- Operating system users referenced in this document `'fnsw' 'orcl' 'root'`
- Oracle database username `oracle`
- Command output is displayed in courier font.

FileNet IS 4.0 on Windows 2000 Server

FileNet supports both UNIX and Windows 2000 platforms (again, FileNet also supports later Windows versions). Among other versions, supported UNIX platforms are Solaris 2.8, HP/UX, and AIX. Our test configuration used a Windows 2000 system. Similar IBM N series storage configurations will also work for IS on the HP/UX, AIX or Solaris platforms, though specific settings will be different for other versions of UNIX. Patch requirements for UNIX platforms should be checked in FileNet IS documentation. Any discussion of UNIX platforms is beyond the scope of this document.

FileNet IS requires a specific version of the operating system, and an additional patch installation is necessary. FileNet 4.0 supports Oracle 9.2.0.2 or later. Again, this test environment used Windows 2000 Server, Oracle 9.2.0.5, and Data ONTAP with FileNet IS 4.0.

Archiving documents to WORM media using SnapLock requires the CIFS protocol on Windows and NFS protocol on UNIX systems. Even a regular network infrastructure works fine with the SSAR configuration. To take full advantage of the performance available with this system, it's best to use fast network connectivity such as Gigabit Ethernet.



FileNet IS 4.0 and the Oracle Database Environment

Currently FileNet supports database servers such as Oracle 9.2.0.2, IBM DB2[®], and SQL Server 2000. Our test setup used Oracle 9.2.0.5 on Windows 2000 with iSCSI and IBM System Storage N series with SnapDrive[®] 3.0 or higher.

Installation Tasks

This section briefly explains the steps required to install and configure FileNet IS. A user account and appropriate groups must be created as described in the FileNet IS installation documents.

Once these tasks have been completed, local disks must be set up and configured. This may involve a SAN configuration using either FCP or iSCSI protocols:

- Install the Oracle database
- Perform post installation tasks
- Install FileNet IS
- Perform post installation tasks
- Configure MSAR storage libraries
- Install IS Connector for SnapLock.

Installing FILENET IS Software

To continue with the installation and configuration of ISCS software, the supported database and FileNet IS must be installed and configured on the server. FileNet IS must be upgraded using a new Hot Fix Pack (HFP).

During HFP setup, FileNet configuration files will be overwritten on the target system and any changes to the original configuration must be manually incorporated from the backup directory into the updated file. This will complete upgrading of FileNet IS to 4.0 SP01 release.

If FileNet IS and database are already installed and configured on the server, continue with setup of the ISCS software.

Configuring the MSAR Storage Library

FileNet IS 4.0 supports MSAR directly without having to upgrade the FileNet IS server. In July 2002, FileNet added MSAR to its IS Release 3.6 ESE (Extended Storage Edition) to support the storage of images and documents on magnetic storage. MSAR can be used along with OSAR or can replace optical storage entirely. Because of the benefits available with the MSAR storage library, many IS customers are adopting MSAR instead of using optical media. Some of the benefits of using MSAR compared to optical media are listed below:

- Increased storage scalability and density
- Faster performance due to:
 - Faster drive performance on magnetic drives than on optical
 - Elimination of delays due to media swapping and spinup
- Higher availability due to elimination of library robotics
- Instant data backup and quick recovery with IBM System Storage N series with Snapshot[™] and IBM System Storage N series with SnapRestore[®] features

- File backup and recovery with IBM N series Snapshot and SnapRestore features
- Fewer end-user delays and complaints
- Improved system manageability.

The minimum MSAR surface size is 1GB and the maximum MSAR surface size is 32GB. The maximum number of slots is 1024 per MSAR storage library. Types of MSAR storage library are Model 16, Model 128, Model 256, Model 512, and Model 1024 (the model number corresponds to the number of licenses purchased). The maximum capacity of MSAR storage is 128TB. (These MSAR licensing options are subject to change by FileNet in future versions of IS.)

Our test configuration used the Model 512 MSAR storage library to configure MSAR surfaces.

Upgrading to IS 3.6 ESE from IS 3.6

Both IS 3.6 SP2 and IS 3.6 ESE were independent releases. IS 3.6.xx was the non-MSAR IS field release. It's important not to confuse these releases.

The following diagram, duplicated from FileNet documents, shows the progression of the service pack release for IS 3.6 in relation to IS 3.6 ESE.

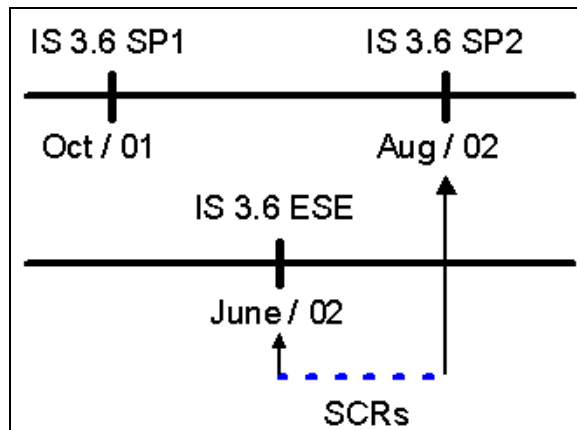


Figure 1. FILENET IS releases in relation to IS 3.6 ESE.

1. If IS 3.6 ESE is installed on a system that has already been upgraded to IS 3.6 SP2, the system will essentially revert back to an older release. If that happens, all the installed patches will be removed and will have to be reinstalled. The required SCRs are listed in Section 15.
2. To bring an MSAR system with IS 3.6 ESE up to the latest SCR level similar to IS 3.6 SP2, all the required SCRs must be manually installed.
3. The IS 4.0 release has a combined MSAR and GA release for IS.

The delta SCRs has new functionality for the IS 3.6 ESE release and the patches released during the release date of IS 3.6 ESE and IS 3.6 SP2. After installing all the required SCRs to the IS 3.6 ESE, install the additional patches that have appeared since the IS 3.6 ESE release date.

Using the MSAR Storage Library

MSAR surface creation using IBM N series storage is simple. It is created in the MSAR creation directory. Its path information is stored in the database, which has one creation directory dedicated for storing the MSAR surface. New MSAR surfaces can be created when the existing MSAR surface is full or

approaches the configured size. All existing MSAR surfaces will remain intact, and the old MSAR files do not automatically move over to the newly created MSAR surfaces. Migrating them to the newly created surface requires intentional surface data movement. Data written to MSAR on NFS volumes is written synchronously to guarantee data integrity.

Configuring MSAR surfaces

Use the FileNet configuration editor (`fn_edit.exe`) to configure a new MSAR surface or delete an existing MSAR surface library (Figure 2). To add an MSAR surface, choose Configuring an MSAR Storage Library, click Run, and follow the instructions. Note that the number of disks for favoring writes can be ignored and set to zero, as shown in Figure 2.




Figure 2. Configuring the MSAR surface library—number of drives favoring writes.

Once configuration is completed, save the changes before exiting the configuration editor and rebuild the datasets by running the `fn_build -a` command.

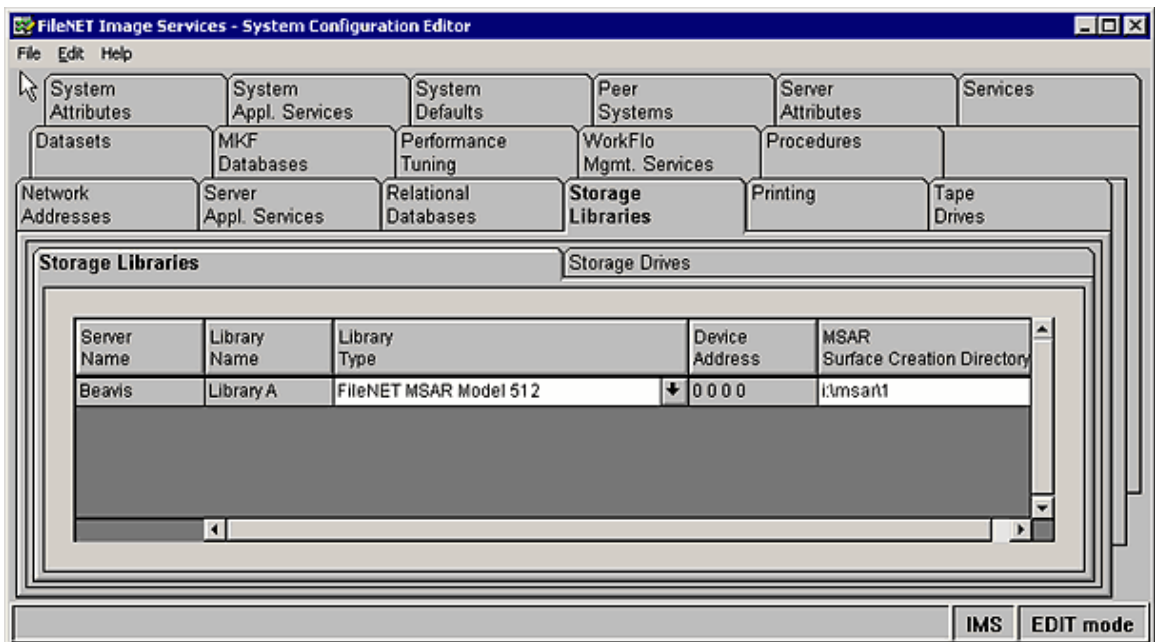


Figure 3. FILENET IS Configuration Editor.

Integration of IS Connector for SnapLock

To take advantage of WORM capability on magnetic media, FileNet has announced support for IBM N series with SnapLock via ISCS. ISCS works with both IS 3.6 ESE and IS 4.0 servers. To take advantage of native retrieval capabilities, this paper recommends upgrading to IS 4.0 SP01 software.

Support of ISCS on IS 3.6 ESE is available only as a FileNet Professional Services consulting project, and requires FileNet clients to make changes to retrieve the files stored on the SnapLock volume. This paper strongly recommends upgrading to IS 4.0 SP2 or later releases. IS 4.0 and later releases support native retrieval of files and hence no changes to FileNet clients are needed. Starting with IS 4.0, the native retrieval is supported with SSAR configuration.

SnapLock configuration allows customers to archive data onto permanent, nonerasable, nonrewritable magnetic media while taking advantages of IBM N series technology for archival, backup, and disaster recovery.

Nearline storage (NLS) is a part of the Single Document Storage (SDS) configuration. It provides a way to archive documents from FileNet IS to systems using Hierarchical Storage Management (HSM) architecture. The NLS version supports IBM N series with SnapLock as an additional storage configuration.

This paper discusses the procedure used to install and configure IS in Single Document Storage (SDS) and NLS modules. The paper is intended for people are familiar with the relevant operating system such as UNIX and Windows. Note that NLS modules work similarly in both the UNIX and Windows environments, though the NLS installation procedure differs.

Preinstallation Requirements

The NLS module currently uses IS Toolkit (WAL) version 3.6 with post qualification for WAL 4.0. All hardware and software support requirements of FileNet IS must be in place, including the storage space and database server configuration.

The FileNet IS server must be upgraded with the necessary patches. Special Contingency Requirements (SCRs) are listed in the file `readmeNLSxxx.txt`. These files must be installed before running NLS modules. A sample from this file is shown below.



FileNet IS Server Configuration

It is important to have correct entries in the host file. NLS software will not work properly if an entry is missing or is incorrect.

```
tcPIP      systemname domainname domainname-filenet-nch-server

A sample file from a test setup is shown below:
# Additionally, comments (such as these) may be inserted on individual
# lines or following the machine name, denoted by a '#' symbol.
#
# For example:
#
#      102.54.94.97      rhino.acme.com      # source server
#      38.25.63.10      x.acme.com      # x client host
#      127.0.0.1        localhost
#      0.10.90.18              boy
#      10.10.10.34          gangoor-l2l  ntap-filenet-nch-server
#      10.10.90.6          venus          venus-filenet-nch-server
#      192.168.252.129      65R2sim
```

Note that here the domain name for FileNet IS is case-sensitive in relationship to runtime files. Verify that the domain-domainname entry in the `NLS.cfg` file matches the entry in the `hosts` file. IS requires a unique user or SysAdmin user to process the work. It is suggested to configure at least 10 concurrent logins.

Install IS Toolkit (WAL)

Install and configure the IS Toolkit software to complete the NLS configuration. After installing the WAL toolkit, apply the HFP to fix some known issues. Installation and upgrading of the WAL Toolkit requires the Windows server to be restarted.

IBM N series Configuration Review

Before proceeding further, verify that the IBM N series storage device has necessary services such as CIFS or/and NFS licensed and started.

Windows Server Platform

- Log in as a user with Administrator privileges
- Insert software (CD-ROM) and double-click `PS_install.exe` if the autorun program is disabled.
- Click Continue.
- At the software installation prompt select the WAL version installed on the system from the Available Releases frame and click Install:

```
C:\>stamp c:\fnsw\client\shobj\*SysV*
C:\fnsw\client\shobj\wal_sysv.dll (NT bin):
  system 4.0.10.39(0) (lib, Thu Aug 21 12:19:33 2003)
  developer 4.0.0.0.10 (lib, Thu Aug 21 12:19:30 2003)
  SubSys: mv, Rel_type: wal_nt, SCR#: 180566, mode: 100666, size: 435764
C:\>
```



Verify the Installation Directory

Select the installation directory for installing the NLS module. The default installation for the NLS module is the IS and WAL installation directory. A standard installation path on Windows would be `\fnsw_loc\bin`. It's therefore convenient to set the user environment to include this path but doing so is optional. The installation of NLS software will complete quickly after the required files are copied. After installing the NLS module, you must configure it to work properly with the SnapLock volumes.

Configuring SSAR Storage Connector Software

Prior to the first use of ISCS software, the following steps must be completed. The first step is to edit the entries in the hosts file.

Configuration of the hosts File

The hosts file is located on Windows 2000 systems as `\WINNT\system32\drivers\etc\hosts`. On UNIX platforms, `/etc/hosts` is the file that must have correct entries to inform the WAL applications to execute the WAL logon call. An example of a hosts file is shown below.

```
127.0.0.1      localhost
# 10.10.11.18  boy
10.10.10.14    gangoor-121  ntap-filenet-nch-server
# 10.10.90.6   venus venus-filenet-nch-server
192.168.252.129 65R2sim
192.168.23.129  achor
```

Environment Variable PATH Setup

The variable PATH must be set up in the user environment to enable WAL calls to access WAL shared libraries rather than the IS shared libraries. This paper suggests distinguishing in setting up the environment variable PATH. On UNIX, the PATH variable can be set using the `.profile` file or using by shell commands such as `setenv` in `csh` mode. On Windows the PATH variable must be appended.

Ownership and Permissions on UNIX Platforms

It is important to set the correct UNIX permission and ownership settings on the WAL modules. The IS server must have the user and groups used to run the NLS modules. In addition to this, the IS user must have read, write, and execute permissions. It is necessary to provide the IS username and password in the NLS configuration file and these parameters are specified in the NLS configuration file called `NLS.cfg`.

Configuration settings may be made in a specific file called `NLS.cfg`. It can be updated or modified using a text editor. A sample file from our test environment is shown below, in the next section.



SnapLock Configuration Information

To successfully set the retention date and archive to a SnapLock volume, the `NLS.cfg` file must be configured properly. The NLS archival program will look for this file and set the configuration using the values in it. Note that a commented line starts with a ";" symbol and is ignored by the SnapLock connector software, which instead uses the default values for that attribute. A sample `NLS.cfg` file is listed below.

```
*****
LogonAttribute {
    UserName="SysAdmin"
    Password="SysAdmin"
    Domain="ntap"
    Organization="FileNet"
}
; General information
PgmAttribute {
; WorkingDirectory=/NLS
    WorkingDirectory= L:\
; 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)
;
; The following time-based keywords all have a minimum of 1 second
; and a maximum of 3600 seconds. If set to zero or a negative number
; it will default to 5 seconds.
; FetchSleep=(in secs, default=600)
; ArchSleep=(in secs, default=600)
; FetchTimeout=(in secs, default=3600)
;
; Display format for the year %Y (default) is 'xxxx', %y is 'xx'.
; This affects display of a date only and extends the length of
; filenames.
; 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
;
; IBM N series SnapLock Keyword
; To engage the retention rules the SnapLock keyword must be
; uncommented below and set to one of the following values.
; 0 = No retention (initially set to 5 mins)
; -1 = Infinite Retention Mode
; 1 or greater than 1 = set retention using input value in hours
; SnapLock={Default=0}
; SnapLock = 1
```

The above example shows the SnapLock archive directory on for Windows.



Setting the IS User and Password in a Secure Method

It is possible to specify the FileNet IS logon information using simple ASCII characters, but in certain environments this approach may not be acceptable for security reasons. To provide an encrypted password capability, the NLS module has a utility called `PS_Password`. This program will create a `.ps_passwd` file that contains the encrypted records. The system administrator must provide a password to allow creation of the `.ps_passwd` file. However, the system administrator must remember this password in order to retain the ability to modify the password at a later time.

Once the encrypted password configuration is completed, edit the `NLS.cfg` file and comment out the line that says Password in the LogonAttribute section in the configuration file.

Destination Configuration

The `NLS.cfg` file lets you specify the working directory as well as the log directory. `WorkingDirectory` specifies the path where the documents will be archived. The necessary directory structure will be created under this name using the system serial number. All files will be created under the directory following this unique serial number.

The `LogDirectory` directory contains the log files created each time the NLS module is run. Journal files include `NLS Archive`, `DumpQ`, `Add2Q` and `NLS_fetch` log information.

A sample entry from a journal file is shown below:

```
PlatformWindows w/HSM Support
2004/05/02 22:49:48 <gangoor> (003108) System Shell Win32 Cmd
2004/05/02 22:49:48 <gangoor> (003108) PS Lib Vers 2.3.2
2004/05/02 22:49:48 <gangoor> (003108) NLS Pgm Vers 1.5.7
2004/05/02 22:49:48 <gangoor> (003108) WAL Release 4.0.0
2004/05/02 22:49:48 <gangoor> (003108)
IDMIS Release 4.0.0 (ntap:FileNet)
2004/05/02 22:49:48 <gangoor> (003108) Starting NLS_Archive
2004/05/02 22:50:17 <gangoor> (003108)
Terminate received. Waiting for active children to finish
2004/05/02 22:50:17 <gangoor> (003108) Ending NLS_Archive
Total Time = 29 Secs
Sleep Time = 26 Secs
Query Time = 3 Secs
Completed 3 of 3 items
2004/05/02 22:50:17 <gangoor> (003108)
NLS_Archive has terminated as requested
```

The destination directory can be set using the `NLS.cfg` configuration file. This working directory will allow the WAL applications and the NLS module to archive the documents to a SnapLock volume. This volume will have WORM capability.

Verify that the IBM N series storage device CIFS and/or NFS services are running. This can be verified by IBM System Storage N series with FilerView[®] graphical user interface (GUI) utility. If CIFS Services are set up and running, configure a network share disk. This task requires Administrator or root user privileges.

Attribute Settings

The last part of NLS.cfg configuration concerns attribute parameters. These parameters are explained in FileNet product documentation and the following section describes setting those parameters relevant to the SnapLock environment.

To set the retention on a document and archive it to a SnapLock volume, edit the line that says: SnapLock = 1 with the entry as shown in the following figure.

SnapLock = 0	Will set no retention	This option will not configure the NLS module not to set the document for retention.
SnapLock = -1	Will set the retention forever	If the document is intended to be archived, set the negative value
SnapLock = 1 (or greater than 1)	Will set the retention for 1 hour or the number of hours specified	After the expiry of 1 hour, document is eligible to be updated or deleted
		F_DELETEDATE F_ARCHIVEDATE parameters override this value

Figure 4. SnapLock archival enable options.

Once the WAL Toolkit application archives to a SnapLock volume, it sets the retention date and commits as a WORM document.

Index Database Settings

An index must be created called `FNP_ARCHIVE`. It is a user-defined index of 'date' type with default attributes.

Functionally, the NLS module can retrieve the documents using a native client-based method and has integrated document management (IDM) desktop and web retrieval approaches.

Archiving Documents to a SnapLock Volume

- Create/verify the new MSAR device configuration
- Use the Configuration Editor.
- Verify the configuration parameters such network address.
- Run `fn_build -a`

Creating/Reporting Family, Class, and Index

Start IS. Verify that the database server is up and running and start IS. Verify that the locally configured disks have been created. On Windows, verify that the IBM N series storage device has been configured and the necessary CIFS Shares are mounted as network share disks. On UNIX, NFS service must be enabled to mount the SnapLock volume. To create a new family, start the application executive and select database maintenance. On our test setup, new families called 'compliance,' 'emails' and 'nonregulated' are shown below.

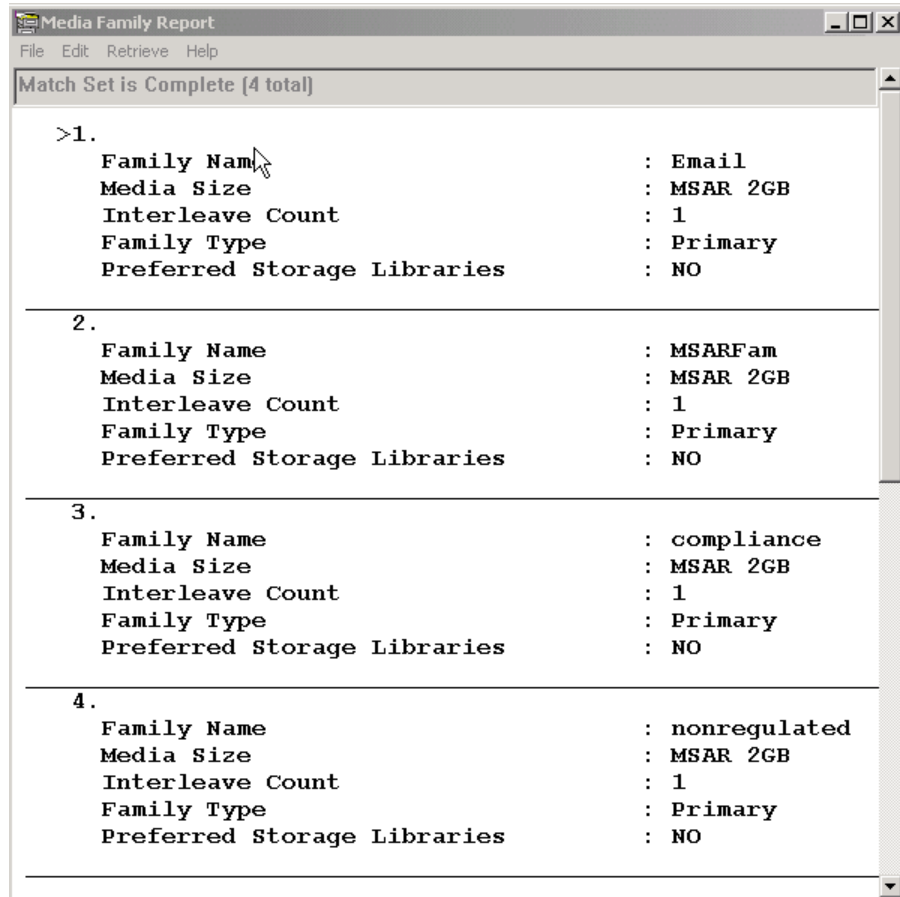


Figure 5. Media family report.

Create a class for the document family configured. On our test setup, we created a class with its own retention parameters for each family. The summary of available classes is shown below. Note that we have to create an index called `FNP_ARCHIVE` that allows the documents to be archived onto a SnapLock volume. Configuring the class requires addition of an index with the name `FNP_ARCHIVE`.

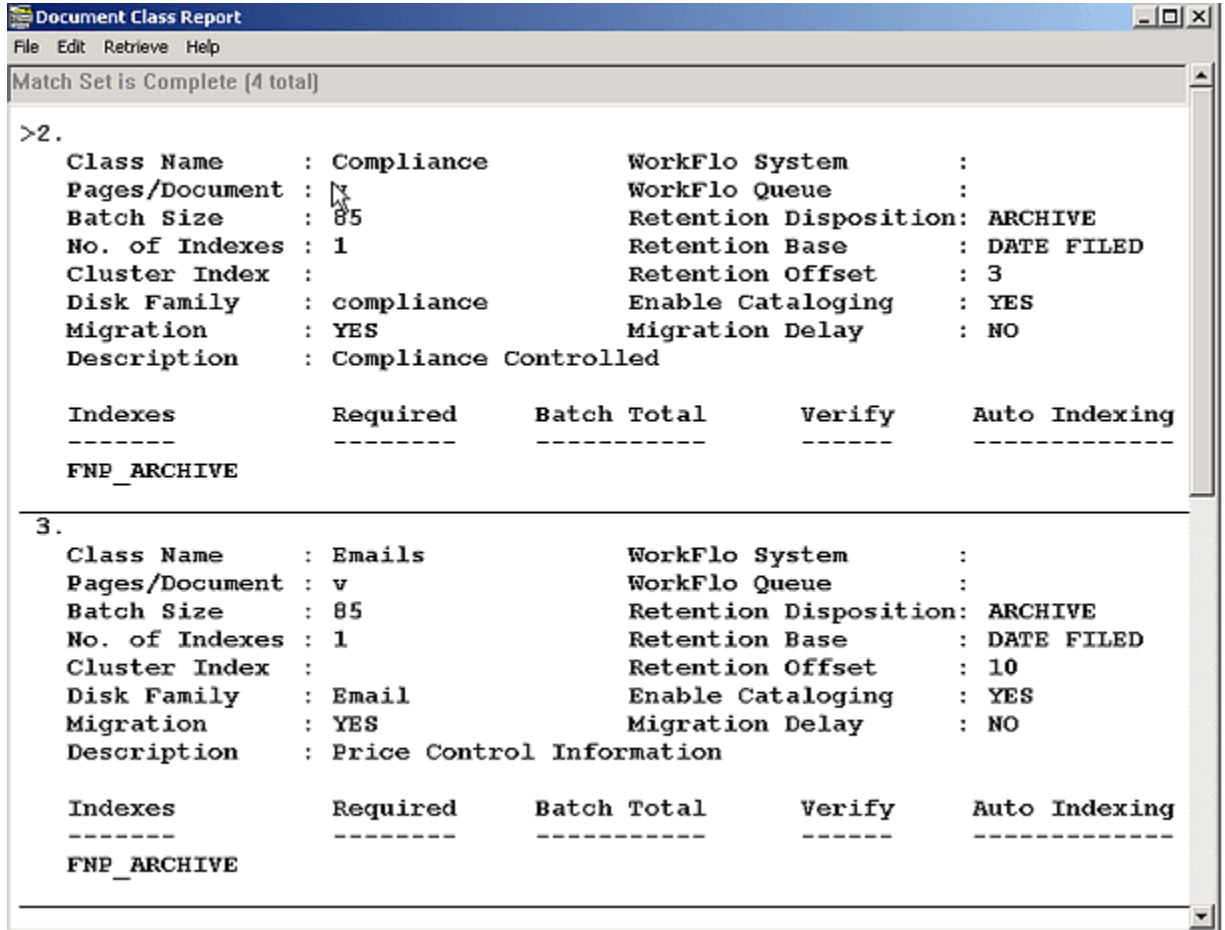


Figure 6. Document class report.

FileNet IDM Desktop Client

Client software must be installed and configured. On our test setup, we installed FileNet IDM desktop client software on a client machine. It is important to install the Hot Fix Pack (HFP) available. .

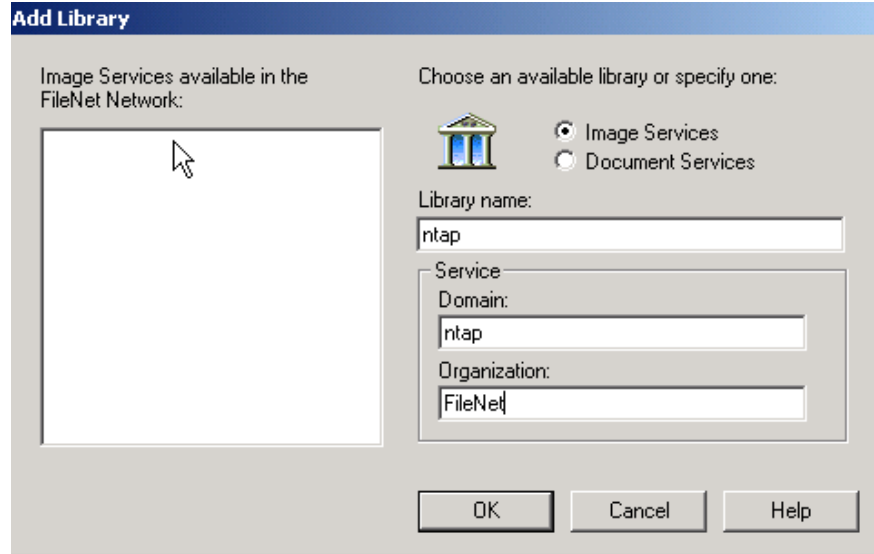


Figure 7. Adding a FileNet Library.

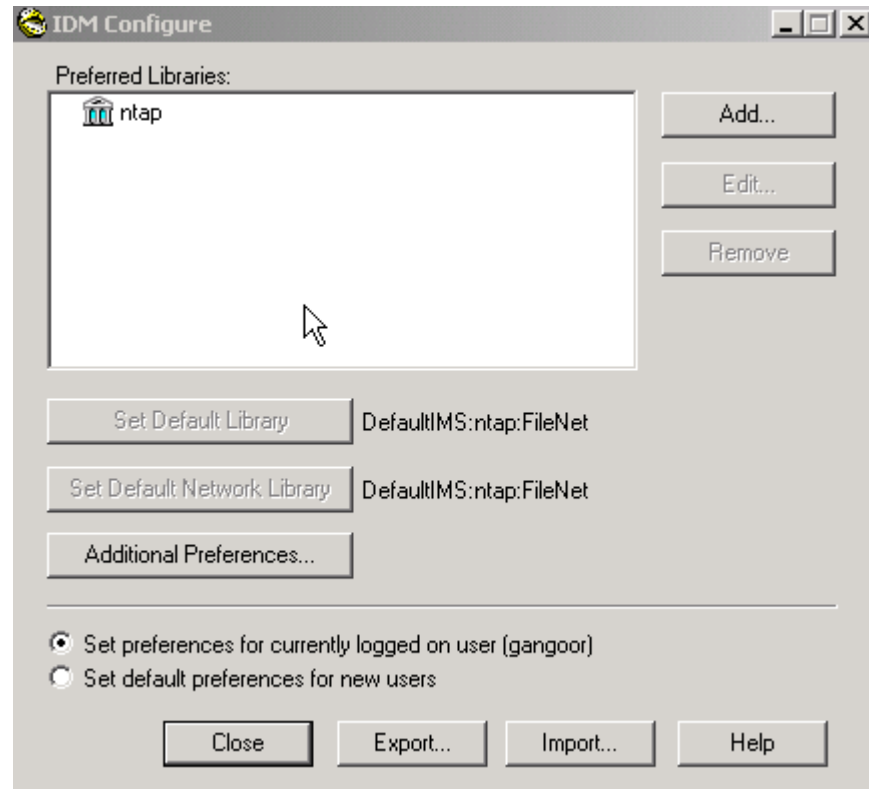


Figure 8. IDM desktop configuration.



Adding Documents to FileNet Library and Archiving to a SnapLock Volume

We have mapped a network share of the SnapLock volume as the path configured in the `NLS.cfg` file. In our setup, the archival destination was specified as `L:\`.

Now start the database server if it is not already started. On our test setup, the database server startup is shown below.

Starting the Database Server

If the database server is not yet started, start it before attempting to start IS. On our test setup, we started the Oracle server by using the `sqlplus` utility as shown below.

```
C:\FNSW>sqlplus "/as sysdba"
SQL*Plus: Release 9.2.0.4.0 - Production on Thu Jul 22 13:49:51 2004
Copyright (c) 1982, 2002, Oracle Corporation. All rights reserved.
Connected to an idle instance.
SQL> startup
ORACLE instance started.
Total System Global Area      44588428 bytes
Fixed Size                    454028 bytes
Variable Size                  41943040 bytes
Database Buffers               2048000 bytes
Redo Buffers                   143360 bytes
Database mounted.
Database opened.
SQL>
```

Starting FileNet IS

If it's not already running, start FileNet IS. Note that the site-controlled database server must be running before attempting to start IS. Before starting the FileNet IS service by issue the following commands.

```
C:\fnsww>initfnsww stop
C:\fnsww> killfnsww -SAD
```

Now start IS by using the `initfnsww start` command and check the event log if make sure that the services are up and running properly.

```
C:\FNSW>initfnsww start
Terminating processes...
Initializing FileNET software...
Starting index database...
Starting permanent database...
Starting transient database...
Starting security database...
Starting Courier...
Starting NCH_daemon...
Starting the Security Daemon...
Starting INXbg...
Starting INXu...
Starting document services...
Starting batch entry services...
Starting print services...
Startup of FileNET software initiated. See event log for detailed status.
C:\FNSW>
```

If IS starts with no error message, continue with the document services. Open the FileNet IDM client and log on to the system library. On our system it appears as below.




Figure 9. Open FileNet Neighborhood.

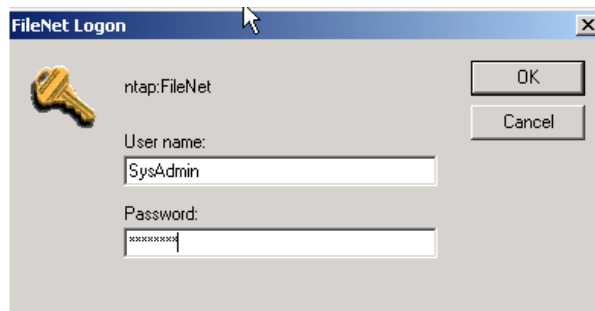


Figure 10. FileNet logon screen.

Double-click the library name and go the FileNet library folder to which you want to add documents. If this is the first time you have added documents, create the appropriate folders after logging in to the FileNet library. In the following example, three folders will be created. All documents related to compliance matters will be stored in the `Compliance` folder, e-mails will go into the `Emails` folder, and other nonregulated information that needs to be archived into a SnapLock volume will be added to a folder called `NonRegulated`.




Figure 11. Folders within the FileNet Library.

Verify that the NLS service is installed. It can be checked in the Services section in the Control Panel. If the service is present, start the service using `NLS_start`. (The `NLS_start` program can be found in the NLS installation directory.) Below is the output of `NLS_start` program.

```
C:\FNSW_LOC\bin>nls_start
NLS Start Requested - Removing Flags File
Starting C:\FNSW_LOC\bin\NLS_Archive program succeeded
Starting C:\FNSW_LOC\bin\NLS_Fetch program succeeded
Starting C:\FNSW_LOC\bin\NLS_Dispatch program succeeded
Starting C:\FNSW_LOC\bin\NLS_DumpQ program succeeded
Startup of NLS in C:\FNSW_LOC\bin has succeeded

C:\FNSW_LOC\bin>
```

Once you commit a document into FileNet library, SnapLock Storage and Retrieval uses the configuration parameters to set the retention date and then archives the document to a SnapLock destination specified in the `NLS.cfg` configuration file. Archival to a SnapLock volume is done on a separate background, and frequency of checking for the document archival can be configured. If there are no documents to commit to SnapLock volumes, this process goes into sleep mode. It wakes up according to a specified schedule and commits any queued documents to the destination path. If the queued documents must be committed before the end of the NLS sleep time, run the `NLS_stop` and `NLS_start` programs. After restarting, all queued documents will be archived to SnapLock volumes.

Once the documents are archived to a SnapLock volume, even the System Administrator cannot modify the attributes or delete the document. This configuration is shown in our setup below.

On our test setup, we committed different documents to corresponding folders depending on the type of document.

Before committing a document classified as compliance, the FileNet neighborhood under the `Compliance` folder displayed the following document numbers already added to the FileNet Library.

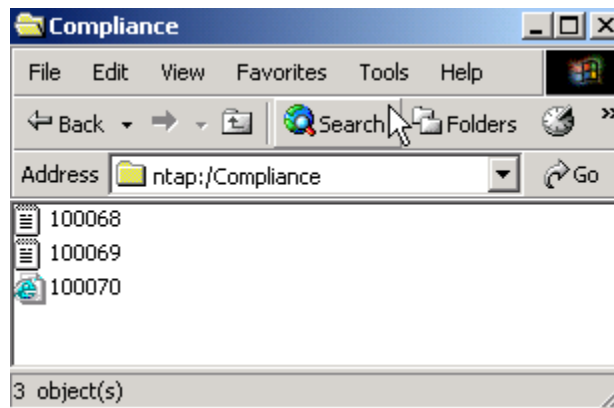


Figure 12. Archived documents for a specific class.

We will add a document to the `Compliance` folder in the FileNet neighborhood and IS services will give that document ID number 100073, as shown below.

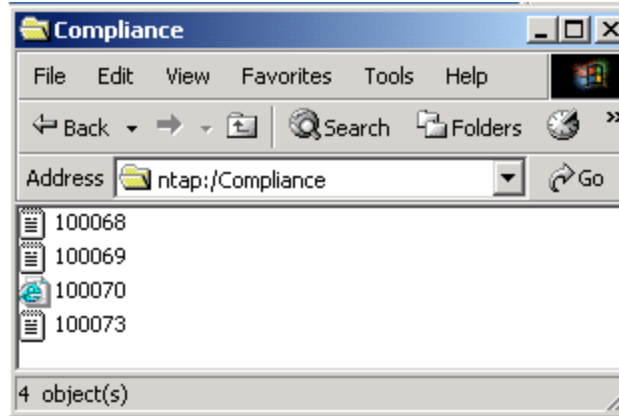


Figure 13. After archiving a new document to the same class.

Start the SnapLock archiving program using `NLS_start` command. This will allow SSAR to set the retention date and other parameters and archive to a SnapLock destination directory. The destination directory for archival is specified in the `NLS.cfg` file. Once a document is added to the library and archived to a SnapLock directory, the necessary directory structure will be created for each defined document class.

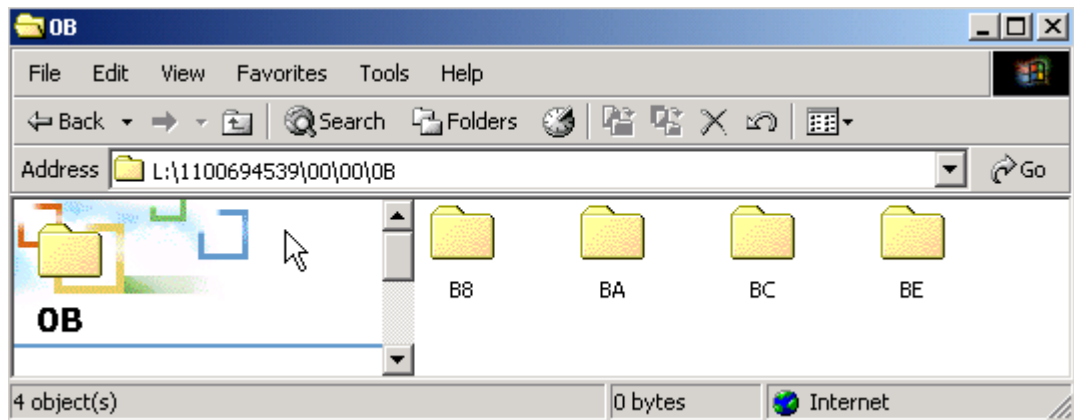


Figure 14. Folders created for each document class.

In the above example, each folder corresponds to the class of documents where it is archived in the SnapLock volume.

Since we added a document called `Confidential` into compliance folder, we will be able to verify that the document with the number 100073 is archived. Recently committed document(s) will be archived to this SnapLock volume. In our example, a document with ID 100073 should be available on the destination directory.

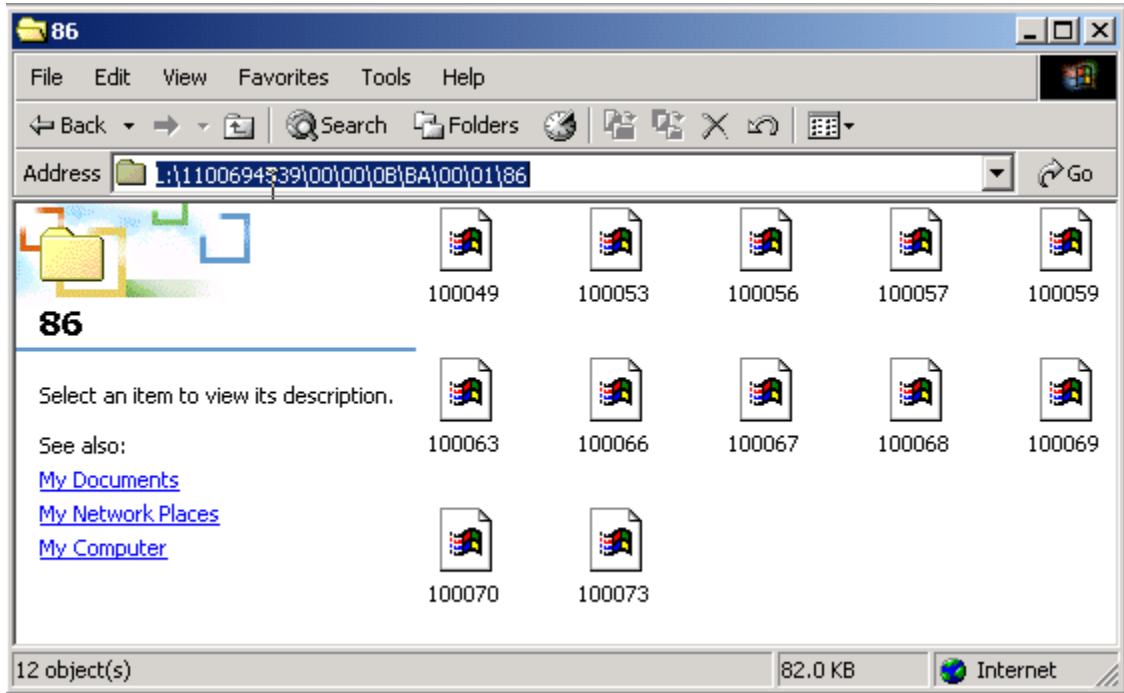


Figure 15. SnapLock archived documents of a specific document class.

A document called `Confidential` added to a FileNet library in the `Compliance` folder is also archived to nonrewritable storage. The archived file has a read-only attribute, meaning that the archived file cannot be modified or deleted by any users, including privileged users. Once the retention date has expired, the document may be destroyed, or the retention period can be extended for compliance or other purposes.




Figure 16. Archived document properties.



Conclusions and Caveats

FileNet supports magnetic media in its IS environment through its MSAR feature. By using IBM N series unified storage customers are able to exploit the benefits of magnetic media storage in the FileNet environment. Recent regulations requiring data protection have forced customers to search for a robust solution for document content protection. By using a connector software module, documents can be archived onto a nonrewritable magnetic media using SnapLock software on IBM N series storage devices. These documents will have retention parameters set according to the organization's retention policy. In addition, IBM N series storage system integration provides FileNet customers with the quick backup and recovery capabilities, simplifying both data replication and disaster recovery planning. This paper demonstrates a simple procedure for integrating FileNet IS, SSAR connector software with the IBM N series storage solution using SnapLock.

All possible combinations of hardware platforms and storage architecture and software options have not been tested. If you use a different server OS, a different version of IS, or a different database, significant differences in your configurations may alter the procedures necessary to achieve the objectives outlined in this document.



Trademarks and Special Notices

© International Business Machines 1994-2007. IBM, the IBM logo, System Storage, and other referenced IBM products and services are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. All rights reserved.

References in this document to IBM products or services do not imply that IBM intends to make them available in every country.

Network Appliance, the Network Appliance logo, Data ONTAP, FilerView, NearStore, SnapDrive, SnapLock, SnapRestore and Snapshot are trademarks or registered trademarks of Network Appliance, Inc., in the U.S. and other countries.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

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.