DART Application

DART – Document Archive/Retrieval Transport

Installation and Configuration Guide

June 7, 2007
Version 1.8.1

Notices

Table of Contents

1 INTRODUCTION & OVERVIEW

Installation by a FileNet Certified Professional (FCP) Recommended. For more information on the FCP program, refer to the FileNet Web site (http://filenet.com), Customer Center > Global Learning Services > Certification Programs.

This document describes the installation, configuration, and use of the Document Archive and Retrieval Transport (DART).

DART provides a method of backing up documents located on optical disks, thereby facilitating a complete disaster recovery opportunity. To achieve this, DART stores documents from IS on a magnetic disk directory. Once written to magnetic disk, documents can be backed up using a variety of options including tape or network backup facilities. These documents are then available to recover the system in the event of the physical loss of an optical disk and its documents.

While DART captures the document index values at the time of backup, it does not backup annotations or highlights, nor does it backup post-DART index changes. These are backed up during normal database backup procedures, independently of DART.

This document is intended to be read by people who are familiar with the UNIX or the Windows operating systems and IS software.

NOTE **THE DESCRIPTIONS GIVEN IN THIS DOCUMENT FOR THE OPERATION OF DART ARE FOR FUNCTIONALITY AS IT IS CURRENTLY IMPLEMENTED. FILENET RESERVES THE RIGHT TO MODIFY THE INTERNAL OPERATION OF DART AT ANY TIME. THEREFORE USERS SHOULD NOT RELY ON THIS DESCRIPTION OR USE DART IN WAYS OTHER THAN DESCRIBED.**

There is no difference in DART functionality when running on a UNIX platform or a Windows platform. However, the method for installing DART does differ on the two platforms. See the section on installation specific to the target Operating System for details.

2 DEFINITIONS

DART – Document Archive Retrieval Transport

Domain – The name of the Image Services

IS – Image Services

ISTK – Image Services Toolkit (WorkFlo Application Library)

SCR – Software Change Request

OS – Operating System

3 HARDWARE & SOFTWARE REQUIREMENTS

DART provides releases compiled for currently supported ISTK versions. When installing on UNIX, the DART installation program will determine the ISTK version release and install the DART version, which coincides with the ISTK version found on the server. (Example: DART 3.6 will be installed for ISTK 3.6 and DART 4.0 will be installed for ISTK 4.0.) When installing on Windows the user must know the version of ISTK they have installed. The section on installation will explain how to determine your ISTK version. Any other requirements depend on the ISTK and IS supported hardware and software requirements. IS and ISTK must already be installed, configured, and running before installing DART. See IS and ISTK documentation for those requirements.

The DART software installation requires a minimum of 10 MB of hard disk space. Runtime directories' space usage will be dependent upon size of user files, volume of batches, and other factors such as file retention schedules.

> *You must provide enough hard disk space to equal the storage capacity of the optical surface being backed up.*

4 SOFTWARE INSTALLATION

4.1 UNIX Platform

This section describes the installation procedure for UNIX systems **only**. Refer to the next section, "Windows Platform" for installation instructions when installing DART on Windows.

NOTE Installing from the web provides a local installer and specific instructions for that procedure. Follow the readmeUNIXinstall.txt document that can be found on the web in the Utilities Documents directory.

The following procedure installs DART into the default installation location or the "home" directory **/fnsw/local/bin**.

1. Log in as root, on the server where you will be installing DART.
2. Insert the CD into the CD-ROM drive, execute the appropriate mount command for your UNIX Operating System, and then change the directory to the CD-ROM mount point, where the "ps_install" executable is located.

NOTE The utilities are released on a CD that can be read by all supported operating systems. HP operating systems require a special mount command to read this CD.

Example:

```
mount –o cdcase /dev/cd0 /cdrom
```

Run the installation program for DART using the following command.

```
>csh  ./ps_install.
```

NOTE: The periods are needed before and after the command.

NOTE: You must specify the "csh" shell in the command is needed if you are not using the csh shell. If you are already using the csh shell, including "csh" in the command will not cause any problems for the install.

NOTE There can be more than one csh (c shell) available on a UNIX system. Using csh (c shell) when executing the ps_install program, can fail depending on the default csh. If a failure occurs during installation, issue the ./ps_install. command without specifying a shell command or by specifying the path to an alternate csh on the system.

3. Verify that the following files were written to the install directory (/fnsw/local/bin):

```
CenteraVerify.exe
CheckSSN
DART (AIX, HP/UX, or Solaris)
DART.EXE (Win32)
DART_cfg.sample
FPLibrary.dll
FPToolbox.dll
PS_Password.exe
PSs (AIX)
LibPSs.sl (HP/UX)
```

```
LibPSs.so (Solaris)
PSs.dll (Win32)
timer.awk
```

Note: **DART_cfg.sample** is a sample configuration file. Modify this file for your environment, then copy or rename it to **DART.cfg,** to be used as the configuration file.

4.2 Windows Platform

NOTE Installation from the web requires specific instructions for that procedure. Follow the readmeWindowsinstall.txt document that can be found on the web in the Utilities Documents directory

1. Log in as a user with Administrator privileges on the server where you will be installing DART.
2. Insert the CD into the CD-ROM drive and the installation will being automatically using Autorun. If Autorun is disabled, double-click on PS_Install.exe (on the CD), or select Run from the Start menu, and type the following command:

   `<CD ROM drive letter>: PS_Install.exe`

3. Click **Continue** on the Application Package Installer dialog box.
4. Select the ISTK version installed on your system from the **Available Releases** frame on the product installation screen. Click **Install**.
   - To check your ISTK version level on a Windows Server, type the following command at the Windows Server prompt to stamp the following module.

     `stamp <drive>:\fnsw\client\shobj\*SysV*`

Example of an ISTK 4.0 module stamp:

```
D:\FNSW\CLIENT\shobj>stamp wal_sysv.dll
  system 4.0.5.7(0) (lib, Wed Oct 15 08:38:16 2003)
  developer 4.0.0.0.11 (lib, Wed Oct 15 08:38:14 2003)
  SubSys: mv, Rel_type: wal_nt, SCR#: 184822, mode: 100666, size: 447160
```

5. Click ***Yes*** to verify the installation location.

   The initial installation location for DART is: `C:\fnsw_loc\bin`

   This installation location drive letter is usually changed to a location where the IS software and ISTK software is normally installed. To specify this preferred area for the install location, click ***No***.

6. If you clicked ***No*** above, specify/select a new installation location in the ***Destination Path Selection*** dialog box. Select/create the new installation location and then click ***Accept***.

   A standard installation location for DART is:

   `<drive letter>:\fnsw_loc\bin`

NOTE Do not install DART in a directory that contains spaces in its name.

7. Click ***OK*** and ***Quit*** when the installation has completed.

4.3 Completing Installation on UNIX and Windows Operating Systems

Complete the following steps before configuring and executing DART for the first time. This section details how to update the system "hosts" file, verify that DART can communicate with FileNet ISTK Runtime, and verify that you can logon to the FileNet system.

4.3.1 "hosts" file entry

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

UNIX: /etc/hosts

Windows: <drive>:\WINDOWS\system32\drivers\etc\hosts

ISTK applications require specific entries in the "hosts" file to correctly execute the ISTK logon call. This entry is specific to ISTK applications and must be added to the "hosts" file before running a FILENET utility. The FILENET domain name entry must appear exactly as it appears in the FILENET Application Executive. If the name contains capital letters and underscores, use the same case and character composition.

Type the following:

8. The TCPIP address first, followed by a tab.
9. The server name, followed by a space.
10. The FILENET domain name, followed by a space.
11. The 3-part NCH server name.

Examples:

192.48.11.12 servername domainname domainname-filenet-nch-server

183.52.10.11 SERVERNAME DOMAIN_NAME DOMAIN_NAME-filenet-nch-server

NOTE If IS requires a domain name entry of a different case, add two 3-part NCH server names to the line entry for your FILENET server.

4.3.2 Environment PATH

You must set up the PATH in the user environment (usually the fnsw user) so that ISTK calls the ISTK shared libraries, not the IS shared libraries. The UNIX configuration of the PATH differs from the Windows Server configuration of the PATH. Do not confuse the Windows configuration with the UNIX environment configuration. Do not limit the PATH to the examples below as there are also directories needed by other software and OS requirements found in each environment PATH that must be retained. The examples below are to only show the ISTK vs. IS PATH hierarchy.

To configure the PATH in UNIX:

```
/fnsw/client/bin:/fnsw/local/bin:/fnsw/bin
```

To configure the PATH in Windows:

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

NOTE IS systems running eProcess where there is a requirement to expose the IS shared library path could have conflicts with the ISTK shared libraries. Should this occur, contact CSS for details on alternate ways to run DART that will avoid this conflict.

4.3.3 ISTK file ownership and permissions on UNIX

ISTK provides an installation readme that outlines the UNIX environment permissions and ownership settings on the ISTK modules. This MUST be followed to ensure correct access of ISTK modules during application runtime. The ISTK modules can NOT be owned by root:system. There are two modules that are owned by root:fnusr. Refer to the ISTK readme for complete instructions.

4.3.4 CheckSSN (optional)

A simple executable is delivered with the utilities to verify that you can successfully logon to the FILENET server and access ISTK Runtime.This program is only a checkpoint and will not determine the ultimate success or failure of the utility to run.

If **CheckSSN** fails with undetermined problems, configure and run the DART utility with debug options to detect if a problem actually exists or if the failure is limited to **CheckSSN**.

NOTE If **CheckSSN** fails with the error: “No IP reference” you must complete the install of the utility and check the login through the utility itself. There is an anomaly that occurs in unspecified environments that will cause **CheckSSN** to fail.

Run **CheckSSN** as described below. You will need the correct Domain, Organization, User name, and password for your FILENET system as input for the **CheckSSN** program.

NOTE The input for **CheckSSN** is CASE SENSITIVE. Ensure that the NCH Domain name entries used for Domain and Organization are the correct case as reported back from IS through the FILENET Application Executive.

Example:

On **UNIX** default location installs:

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

On **WINDOWS** default location installs:

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

> **CheckSSN**

Sample input values:

```
Enter IMS Domain:         profserv
Enter IMS Organization: FileNet
Enter User name:      SysAdmin
Enter Password:       SysAdmin
Enter Application name: DART

Logging onto 'profserv:FileNet' as 'SysAdmin/SysAdmin' @'SSNCHECK'
Getting SSN of DocServices
SSN for IMS System 'profserv:FileNet' is 1100202180
  Primary Security File='./PSLSF
```

```
  SSN=1100202180,YY=100,DDD=5,NumSSN=126,FirstK=5, Pgm='DART'
License for 'DART' Registered Copy
>
```

If you receive an error running **CheckSSN**, display the error tuple message to determine why the failure occurred. Correct the problem by installing missing software or configuring access to IS, and then re-run **CheckSSN**.

NOTE When your FileNet Service organization name is not "FileNet", enter the entry for "domain" using the two-part naming convention. Use upper case letters and underscores as the naming convention is actually seen on the FileNet Service.

In the example below, the user is not prompted to enter the "Organization name." This is the normal prompting sequence when the domain name and organization name are entered on the same line in the two-part format.

Example:

```
Enter IMS Domain:       profserv:MSS_ISR
Enter User name:        SysAdmin
Enter Password:         SysAdmin
Enter Application name: DART
```

NOTE The DART user MUST have at least 1 concurrent login configured/available through the IS Security Administration.

CheckSSN must have the application name input in all upper case characters.

Incorrect value input or any configuration changes can cause **CheckSSN** to continue to fail when subsequent attempts to run it are completed with correct value input. Should this occur, run **wal_purge** on the UNIX environment or recycle the FILENET IS software on the Windows environment.

5 CONFIGURING DART

Once the software is installed, you must set search paths and user security, and modify the ***DART.cfg*** configuration file.

5.1 Working Directory

The ***WorkingDirectory*** is the top-level directory where DART will write its working files.

5.1.1 UNIX Platform

The default working directory for UNIX is /fnsw/local/bin. This directory must be owned by ***fnsw*** and the group set to ***fnusr***. Set the permissions on this directory as read/write/execute for all. ***(chmod a=rwx path)***

Files from DART are written as:

Owner - fnsw
Group - fnusr

NOTE If using a remote server, create an ***fnusr*** group with an ***fnsw*** user, with group and user IDs that match those of the corresponding IS library. If the directory is to be remotely mounted, it must be owned by ***fnsw:fnusr***, with ***777*** permissions ***(a=rwx)***.

To check if the remote file system is mounted correctly, log in as the fnsw user and do the following:

`df <mountpoint>`	returns the mountpoint and remote directory information
`ls –ld <mountpoint>`	returns the permissions and owner information
`touch <mountpoint>/<SSN>/mtcheck`	creates a file in the `<mountpoint>/<SSN>` directory

5.1.2 WINDOWS Platform

The default working directory for WINDOWS is <drive>:\fnsw_local\bin. The owner permission that exists in UNIX environments is not applicable on the WINDOWS platform.

5.2 User environments

UNIX Operating System User only

Although the ***fnsw*** user is often used to run DART, a separate UNIX user can be created to run DART. This user must have read, write, and execute privileges for all directories, files, and programs used by DART. The UNIX user must be a member of the groups: ***fnusr, fnop, fnadmin, dba,*** and ***root.*** Refer to the “Environment PATH” section to setup this *newuser’s* PATH correctly. The DART modules in the installation location must also have the “sticky bit” set to provide for proper module access.

Use the UNIX chmod command on all DART runtime modules as follows:

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

```
chmod g+s *
chmod u+s *
```

Use the UNIX chmod command on **wal_purge** and **wal_daemon** modules as follows:

```
cd /fnsw/client/bin
chmod g+s wal_purge wal_daemon
chmod u+s wal_purge wal_daemon
```

All ISTK and DART modules should still be owned by ***fnsw:fnusr***. When running DART you will login as the *newuser*. If another user was used to run DART, you must check permissions on all data input and output files to ensure proper file access.

WINDOWS and UNIX Image Services User

An IS user must be created that has read, write, and execute privileges for documents on the IS library. This user must also have been configured with at least 3 concurrent logins. Enter the name and password of this IS user in the DART.cfg file. The default user for this entry is SysAdmin. To configure a special user for DART, consult your implementation specialist.

5.3 DART.cfg

DART.cfg is a text file, which can be modified using any text editor. The following shows the format of the sample file. Note that the line beginning with a semi-colon (;) is a comment line and is ignored by the program. In addition, blank lines are not valid in the **DART.cfg** file.

DART_cfg.sample file:
profserv(root)/fnsw/local/bin> cat DART.cfg

```
;*************************************************************************
;  This module contains the configuration information for the XXX.        *
;*************************************************************************
; $Author: giffj $
; $Date: 10/.0/.0 .1:.3:.0 $
; $Header: ART_cfg.sample,v 1.2 10/.0/.0 .1:.3:.0 giffj Exp $
; $Locker:  $
; $Revision: 1.2 $
; $Source: /usr/ProfServ/lbush/ART/src/RCS/ART_cfg.sample,v $
; $State: Exp $
;*************************************************************************
; INSTALL HISTORY:
;       INSTALLED BY                DATE                DESCRIPTION
;------------------------------ -MM/DD/YYYY-    --------------------------
;
;*************************************************************************
;    Logon Information
```

```
;*************************************************************************
LogonAttribute {
     UserName="SysAdmin"
     PassWord="SysAdmin"
     Domain="profserv"
     Organization="FileNet"
}
; General information
PgmAttribute {
; HOME indicates the directory where this DART.cfg file is located
; Directory used for export of files _ MUST CHANGE for your environment
          WorkingDirectory=/Xtra1/DART
;         LogDirectory={Default=/{HOME}/journals}
;    DocReportFrequency CAN be changed to a maximum of 999.
          DocReportFrequency=999
;   The following time based keywords all have a minimum of 1 seconds
;   and a maximum of 21600 seconds. If set to zero or a negative number
;   it will default to 5 seconds.
;         FetchSleep=(in secs, default=600)
;         OSNice=(in secs, no default)
;
;         WalkBack={Default=100000}
;         CleanCache={Default=False}
;         Timing={Default=False}
;         Primary={Default=True}
;         DDExim={Default=True}
;   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"
;   Minumum for the following items are 0 seconds
;      OSNice={Default=0}
;
}
;END OF CONFIGURATION FILE
```

5.3.1 LogonAttribute

The logon attributes should match the IS user and standard logon information for the IS system. For information on user privileges, see the section on **Users**.

5.3.2 Configuring Encrypted Password (optional)

To provide values for an IS login through the IS Utility Application, set the values of the UserName and PassWord in the **LogonAttributes** of the **DART.cfg** file. Enter this information using one of the following methods:

- Enter the values for the UserName and PassWord in ascii text form. This means that the password for the user will be visible to anyone with read access to the **DART.cfg** file.
- Configure these login attributes to use the Encrypted Password functionality. Follow the instructions below to use this functionality.

The System Administrator will use the **PS_Password** program to create the ".ps_passwd" file which will contain the encrypted records. The first time the PS_Password program is run, provide a password for the creation of the ".ps_passwd" file as well as the password(s) for any IS Utility Application and User Name. Remember the .ps_passwd file password or you will not be able to run the **PS_Password** program a second time to change the .ps_passwd file.

Login as **fnsw** and run the **PS_Password** program:

UNIX:

```
cd /fnsw/local/bin
>PS_Password
```

Windows Server:

Open a DOS window

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

When the **PS_Password** program first runs, you will be notified that it is running in "*Initialization Mode*". You will then be asked to Enter and ReEnter (validate) the password for the *Password Application Key*. (The character limit for this password is 8 characters and it is not changeable once it has been set.) The *Password Application Key* is the password that will be used to run the **PS_Password** program a second time. Subsequent runs of the program will require you to enter this password, and the *Initialization Mode* message will not display.

The **PS_Password** program will then ask for the *PS Application Name*, a *Password* for the login being used and the *User Name* of that login. You MUST provide a valid user ID that has SysAdmin privileges within the IS domain.

Once the application name, password, and user name have been entered, the program will display the application and user name with validation that the password has been set. Another application password can be set at this time or you can click **Enter** at the next prompt**,** to stop the **PS_Password** program, and display a list of all applications and user names entered.

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       : DART
        Application (39)'DART'
        Enter New Password for Application DART        :
        ReEnter Password for Application DART        :
          Enter User Name       : SysAdmin
          PS Program 'DART': User=SysAdmin, Password Set
        Enter PS Application Name       :
**** Current Applications Set ****
**  Application (39)'DART' [User='SysAdmin']
*********************************
```

Now that the **PS_Password** program has been run and the *.ps_passwd* file has been created, configure the **UserName** and **PassWord** values in the **DART.cfg** file.

Set the **UserName** value in the configuration file to the application name entered into the .ps_passwd file. In this case that would be **UserName="DART".**

Remove or comment out the **PassWord** value in the **LoginAttribute** section of the **DART.cfg** file. If the DART application sees the **PassWord** parameter in the **DART.cfg** file during start-up it will NOT use the encrypted password that was entered when the **PS_Password** program was run and DART will not run.

Example of the **DART.cfg LogonAttribute** section when using the encrypted password functionality:

```
;*********************************************************
;       LOGON INFORMATION
;*********************************************************
;DART.cfg.sample
LogonAttribute {
        UserName="DART"
;       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 domain
        Domain="profserv" (MUST change to correct domain)
        Organization="FileNet"
}
```

NOTE If you forget or lose the password that allows access to the **.ps_passwd** file through the **PS_Password** program, you must delete and recreate the **.ps_passwd** file with the **PS_Password** program Initialization Mode.

5.3.3 PgmAttribute

Directories

All the directory attributes are configurable. The defaults shown above are for example only. The ***WorkingDirectory*** is mandatory. All others will be created by DART if they do not exist. All directories must have read, write, and execute privileges for the UNIX user (usually ***fnsw***) who is running DART.

- ***WorkingDirectory*** is the directory that specifies the top-level path where documents will be stored. The value for this parameter must be a fully qualified path. This Keyword name was formerly known as the "ParentDirectory". This naming convention can still be found in some systems' *.cfg file and is backwards compatible.
- ***LogDirectory*** is the directory where log files are created each time DART is run. The directory structure is comprised of the default <home> path followed by a "journals" directory or under the defined path found in the ***LogDirectory*** value followed by the "journals" directory. You can read the log files found in this structure, **drt*yymmdd***, while the DART program is operating. Again, if the ***LogDirectory*** is not specified, the journals directory and log files will be created under the DART home directory.

  This Keyword name was formerly known as the "JournalsDirectory." This naming convention can still be found in some systems' *.cfg file, and is backwards compatible.

  The installation location/home directory, including the journals directory defaults are:

  UNIX: ***/fnsw/local/bin/journals***

  WINDOWS: ***<drive letter>:\fnsw_loc\bin\journals***

NOTE The ***WorkingDirectory*** path along with internal commands and spaces has a character limit of 64. (A path exceeding 35 characters could exceed the path limit and cause the 202,0,9 error. The HP Operating System seems to be the most sensitive to this issue.)

Other Attributes

- **DocReportFrequency**
  DocReportFrequency defines the document *set size*. This *set size* affects the amount of memory used, the size of the sets created, and the frequency of status reporting. The minimum is 4; the maximum, which is the default, is 999. The value should be altered based on site conditions.

  DART collects documents in sets, which are then processed in a batch-like method. All of the arrays and memory buffers used are based on the *set size*. (The bigger the *set size*, the more memory is used and the faster DART will backup.) Another memory consideration is the number of pages per document. If the "average pages per document" is high, you should lower the ***DocReportFrequency*** setting, depending on your environment. Once the

documents have been archived, the *set size* has no meaning.
In addition, Archived set sizes of one value do not affect Recovery sizes of another. This allows small systems to use less memory by having small set sizes. These changes will affect DART's performance and can measurably affect the system performance by using less or more memory.

The FileNet Object (BLOb) created by DART consists of the following attributes that can affect the FileNet Object size:

- 1 kb header
- image
- index information

The header size can vary and will depend on how many docs there are on the surface. Many very large docs can cause the header size to be smaller than 1kb.

Index information can vary the size of an Object. Few indexes with very long values and large amounts of indexes can both cause the FileNet Object size to be larger.

- **CleanCache**
  During archiving, if ***CleanCache*** is set to true, and the document in cache is ageable, the document will be deleted after DART has backed it up. The ***CleanCache*** default is False.

- **Timing**
  If ***Timing*** is set to true, DART writes timing information to the log files. Normally ***Timing*** will be set to false (the default) 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 to the screen various timing figures.

  Usage for UNIX: awk –f timer.awk journals/drt20010328

  Usage for WINDOWS: awk –f timer.awk journals\drt20010328
  (you must have awk for WINDOWS)

- **Primary**
  The ***Primary*** keyword controls how the surface id portion of the object path is created.

  The default value of TRUE causes the surface id portion of the object path to be represented by an 8 character hex number of the Primary surface id (MKF's surface_id_1). This 8 character hex number is divided into 4 2-character directory levels.

  The value of FALSE causes the surface id portion of the object path to be represented by the actual numerical number of the Secondary surface id (MKF's surface_id_2).

Example:

```
Doc Id       = 1234567,  hex = x0012D687
Surface Id 1 = 3010,     hex = x00000BC2
Surface Id 2 = 5120,     hex = x00001400

Primary=TRUE   Path = <wd>/<ssn>/00/00/0B/C2/00/12/D6, File = 1234567
Primary=FALSE  Path = <wd>/<ssn>/5120/00/12/D6/, File = 1234567
```

- **WalkBack**

  DART uses 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 ID's that exceeds this default setting of 100,000. This can occur when there is more than one document entry device, each with its own unique document ID numbering scheme. It can also occur when throughput is very large and there is an OSAR migration delay that could present a situation where documents are not migrated in a timely manner.

  The **.dartrc** file contains information that assists DART in restarting from a reasonable point. When DART is started again it will check this information and automatically walk back 100,000 document ID's to begin its daily backup. If there are known environment issues which would cause this walk-back sequence to miss documents, set the DART ***WalkBack*** keyword value to any number larger than 100,000, to assist DART in identifying which document ID number to begin its backup.

NOTE The **.dartrc** file is created and/or updated at the completion of a successful set during Incremental (default) mode and Consolidated mode only.

Example:

Setting the ***WalkBack*** equal to 1 million will cause DART to begin at the document ID that is 1 million back from the last entry in the **.dartrc** file.
WalkBack=1,000,000
If the resulting Document ID returned is less than the valid minimum document ID on the system, DART will function as if this was running for the first time.

- **FetchSleep**

  The ***FetchSleep*** keyword allows the user 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 ***DocReportFrequency*** keyword value. The sleep will automatically adjust for sets of 10 items or less with the maximum sleep to be automatically set at 10 seconds. The standard set size for most production systems (assuming the system has resources to handle such a load) is 999. If the ***FetchSleep*** keyword is not set, then the default sleep between sets would be 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 will default to the minimum of 5 seconds) and the maximum is 21600 (6 hours). The default=600 <10 minutes> when the *set size* is not set to a figure between 1 and 10.

- **OSNice**

  OSNice is the minimum time period between sets. Enter the number of seconds you want DART to *sleep* so other processes get processing time. This parameter is useful when DART is run during normal production time.

- **YearFmt**

  YearFmt is the display format for the year portion of the date which is appended to the DART working files. This parameter can change the length of a file name. The default is '%Y' which displays in the 4-digit format "yyyy". To display a 2-digit "yy" format, change the default to '%y'

- **DDExim** (only available in a UNIX environment)
  When set to TRUE, ***DDExim*** creates a text description of the Data Dictionary for backup. The default=True<ON>.

  For more details on ***DDExim*** see the ***DDExim*** section of the IS documentation.

NOTE The ***DDExim*** feature has been known to cause DART to hang on some systems when DART is first started. Turning the value to FALSE will elevate this problem. Check the IS documentation for other ways to run ***DDExim***.

6 FUNCTIONAL DESCRIPTION

With DART, you can back up documents from IS to magnetic disk. These documents can then be transferred to other media, using the customer's preferred backup method. If the documents must be restored back to the IS library, DART will provide the ability to recover the images from the backup. DART is comprised of two main server-based components:

- The archiving function, which identifies documents to be archived, reads them from the IS library, converts each document into a storage object, and writes the object to magnetic disk.
- The retrieval function, which identifies documents to be restored, reads the storage object from magnetic disk, and re-creates the documents on the IS library.

6.1 Backing up Documents

DART identifies documents to be backed up by one of the following three specifications. See Section 6.3 for additional details.

- **Disk Family -** Specifying a Disk Family name causes DART to determine the currently active surface(s) and internally generate a list of documents to back up. If not specified, the family default is ***tranlog***.
- **Optical Disk Surface Id -** Specifying an optical disk surface Id causes DART to internally generate a list of documents to back up from the specified disk surface.
- **Document List -** Specifying a list of documents causes DART to back up the documents contained in the list. The process for backing up documents is depicted in the following diagram.

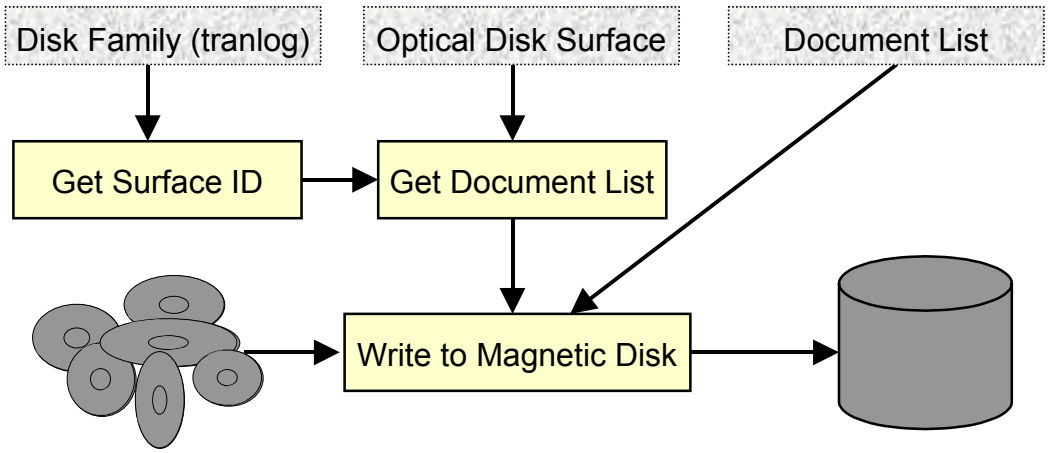

6.1.1 Backup Modes

DART has four backup modes:

Archive	The Archive mode provides for two tracking files to determine what will be backed up. It updates both the Mark Checkpoint File ('xxxxxxxx.mark') and the Surface Checkpoint File ('xxxxxxxxx.surf'). Archive is the mode that will occur when DART is run for the very first time on a system. The Archive mode starts a backup from the document identified in the Mark Checkpoint file. If the Mark Checkpoint file is empty, it creates a backup of everything specified by Disk Family or Surface ID. Typically, the Archive mode can be used in conjunction with the Consolidated mode described under the Consolidated description below. To use Archive with Incremental mode: Monday - Incremental (backs up Monday documents) Tuesday - Incremental (backs up Tuesday documents) Wednesday - Incremental (backs up Wednesday documents) Thursday - Incremental (backups up Thursday documents) Friday - Incremental (backs up Friday documents) Sunday – Archive (backs up Monday through Sunday documents) ♦ Only applies if a Disk Family or Surface ID is specified ♦ Starts at the Document identified in the corresponding Mark Checkpoint File ♦ Documents for the applicable optical disk surface are backed up in sets ♦ Updates the Surface and Mark Checkpoint files after each set ♦ .dartrc file is NOT updated in this mode

Consolidated	The Consolidated mode uses the Mark Checkpoint File to determine what has already been backed up and then creates a backup for the difference. It differs from the incremental method, in that it uses the Mark Checkpoint File to create its starting point but DOES NOT update the Mark Checkpoint File after it completes. Because of this, it can be used to provide a cumulated backup since the last Archive. Consolidate will back up the same documents as well as any new documents over and over again until an Archive is done which will update the Mark Checkpoint file. An example of using consolidated would be: Start with empty Mark Checkpoint File Monday - Consolidated (backs up Monday documents) Tuesday - Consolidated (backs up Monday and Tuesday) Wednesday - Consolidated (backs up Monday, Tuesday, Wednesday) Thursday - Consolidated (backups up Monday, Tuesday, Wednesday, Thursday) Friday - Consolidated (backs up Monday, Tuesday, Wednesday, Thursday, and Friday) Sunday - Archive (backs up everything since last Archive, which in this case would be the full week) ♦ Only applies if a Disk Family or Surface ID is specified ♦ Starts at the Document identified in the corresponding Mark Checkpoint File ♦ Documents for the applicable optical disk surface are backed up in sets ♦ Updates the Surface Checkpoint File after each set, but not the Mark Checkpoint File ♦ .dartrc file is updated in this mode at the completion of a successful set

Incremental	With Incremental mode, which becomes the default mode after the first DART run, all documents associated with the Disk Family or Surface ID will be selected for back up. The Surface Checkpoint File ('xxxxxxxx.surf') contains the list of document IDs that have been backed up from that surface. It is used to compare what was backed up previously against the selected list and the difference is the actual backup list processed. The Surface Checkpoint File is updated to contain the additional document IDs. An example of using incremental would be: Monday - Incremental (backs up Monday documents) Tuesday – Incremental (backs up Tuesday documents) Wednesday – Incremental (backs up Wednesday documents) ♦ Only applies if a Disk Family or Surface ID is specified ♦ Starts at the Document identified in the corresponding Surface Checkpoint File ♦ Documents for the applicable optical disk surface are backed up in sets ♦ Updates the Surface Checkpoint File after each set ♦ .dartrc file is updated in this mode at the completion of a successful set
FileList	The FileList mode is used to back up a specific set of documents. ♦ Only applies if a Document List is specified ♦ Starts at the first Document in the list and processes the list sequentially ♦ There is no checkpointing in FileList mode ♦ .dartrc file is NOT updated in this mode

6.2 The Retrieval (Restore) Function

DART identifies documents to be restored by one of the following three specifications. See Section 6.4 for additional details.

- **Disk Family -** Specifying a Disk Family name causes DART to determine the currently active surface(s) and internally generate a list of documents to restore.
- **Optical Disk Surface Id -** Specifying an optical disk surface Id causes DART to internally generate a list of documents to restore.
- **Document List -** Specifying a document list causes DART to restore the documents from that list. The process for retrieving documents is depicted in the following diagram.

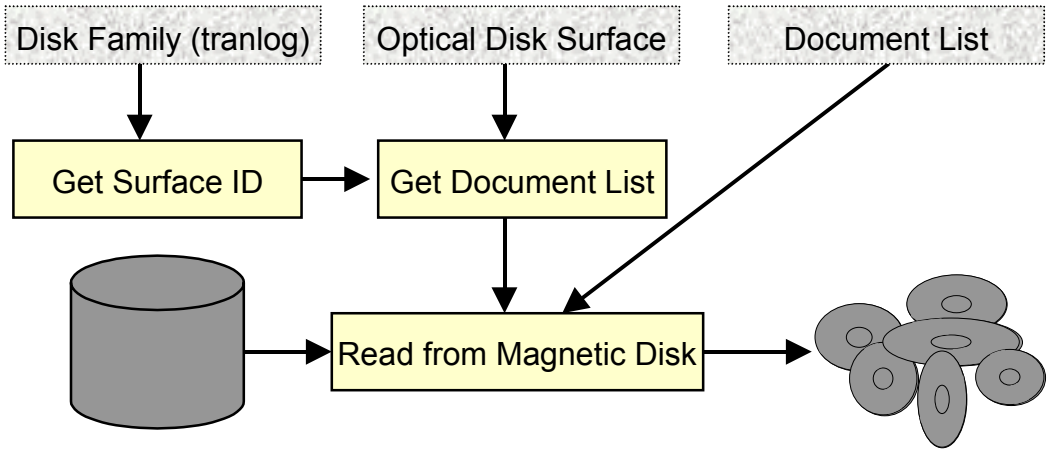

6.2.1 Retrieval (Restore) Modes

DART has three recovery modes. You must understand all characteristics of a mode before selecting a recovery mode. The disposition of the recovered documents will depend on the mode selected.

DART –r without the use of any other input command flag options, will cause DART to recover documents assigned to the current "tranlog" write surface as indicated by the IS database. The tranlog must be named the default "tranlog" for this type of DART recovery to succeed.

Standard	♦ Only applies if a Disk Family or Surface ID is specified ♦ Processes all documents for the applicable Surface ♦ Only recovers a Document if it exists in the Index Database ♦ There is no checkpointing
Forced	♦ Only applies if a Document List is specified ♦ Starts at the first Document in the list and processes the list sequentially ♦ If the Document exists in the Index Database, it is recovered ♦ If the Document does not exist in the Index Database, it is created ♦ Requires identical Document Class Definitions to the system that created the backup
Scan	♦ Applies if a Disk Family, Surface ID, or Document List is specified ♦ Processes all documents for the applicable Surface or in the Document List ♦ Assigns new document ids ♦ Creates new documents on the system being recovered ♦ Requires identical Document Class Definitions to the system that created the backup ♦ There is no checkpointing

6.3 Scenarios / Examples of possible site procedures

There are many ways to backup and restore files. The procedure you select will depend on your site requirements. Refer to the following examples for ideas, when forming a strategy for your site.

6.3.1 Create incremental backups and a full backup

Monday-Saturday - DART -t <family> -c (Consolidated Backup)

Sunday - DART -t <family> -a (Archive Backup)

Use this procedure to create a full backup once a week with incremental backups during the week. This is beneficial if it takes a month or more to fill a tranlog at your site. Should your site experience a surface failure at the end of the month when the tranlog is almost full, you will restore three Sunday Archive backups and the most current Consolidated backup. When using the Archive mode, a *mark file is created to tell the Consolidated mode to start the next backup from the last Consolidated backup.

NOTE Using Archive mode will backup documents that were not backed up the night before. This method could leave updated indexing information off your backups if your site re-indexes documents.

6.3.2 Create a full backup every night

Monday-Sunday - DART -t <family> -c (Consolidated Backup)

Use this procedure to create a full backup every night. These backups will always have current indexing information. However, the time to do the backup and the disk space required to house the backup will need to be large enough to contain the contents of the entire disk. Consolidated mode will always do a full backup if Archive mode is never done, since the *.mark file will always be empty. Thus DART will always go back to the beginning of the surface when performing a backup. Should your site experience a surface failure, you will only restore the last Consolidated backup.

7 THE MAGNETIC DISK DIRECTORY STRUCTURE

When DART backs up documents, by default, it uses a directory structure similar to the following.

/WorkingDirectory

/<ssn>

/<SurfaceID#>

/<DocumentID#>

/[DocumentID]

Within this directory structure, the following meanings apply:

WorkingDirectory	The DART mount point specified in ***DART.cfg***
<ssn>	The system serial number of the IS system where the documents reside
<SurfaceID#>	A four level directory structure based on the Surface ID
<DocumentID#>	A three level directory structure based on the Document ID
[DocumentID]	The files that store the IS Document - There is one file per document with the original IS Document ID as the filename.

If the **Primary** option variable is set to FALSE, this structure will change according to the Secondary Surface ID (tranlog) structure.

7.1 Managing the Magnetic Disk backups

How the customer chooses to manage the Magnetic Disk backups will vary depending on operational needs, available space, and the techniques used for securing the magnetic disk backups.

In general, once a transaction log (surface) is full and has been removed from the OSAR, the need for the backup copy will cease. If the backup is stored on tape or a remote magnetic RAID, for example, the media could be recycled. If the current in-OSAR tranlog is lost, DART will need access to a magnetic disk with the relevant documents and directory structure intact.

8 DART OPERATION

There is no difference in the DART functionality when DART is running on UNIX or WINDOWS. The starting procedure will differ slightly but will use the same startup flags to specify the functionality options.

UNIX: DART is started on a command line or can be set up as a *cron* job.

WINDOWS: DART is started in a DOS shell window with the DART command.

8.1 Starting DART

With the IS software running, perform the following steps:

- Log on as ***fnsw*** or another member of the ***fnusr*** group. This should correctly set the path variable and the shell on a UNIX system and allow for the correct user properties on a WINDOWS system. Do NOT run as the ***root*** user.
- Go to the default directory ***/fnsw/local/bin*** on a UNIX system and to the default directory **<drive letter>:\fnsw_loc\bin** on an WINDOWS system
- Run DART using the appropriate options described below

Type ***DART –h*** to see all allowable options, as follows:

```
hq-pearl(fnsw)/fnsw/local/bin> DART -h
    Setting Default 'drt' log file
    2003/05/16 14:12:42 <fnsw> (001933)
    Switching to Journal Log ./journals/drt20030516
    Starting 'drt' Journal Log ./journals/drt20030516
Usage:  -------------------- DART Options --------------------
        -f <doc list>   File containing Document Ids
        -surf <surfid>  Document list is from surface
        -both           Use both sides of platter when –surf is used
        -ssn <ssn>      Specifies an ssn to backup to/restore from
                        Default=<ssn of domain in 'cfg' file>
        -t <family>     Specify a family. Default='tranlog'
        ** These options are mutually exclusive
        ** Default is Incremental Backup
        -i              Incremental backup
        -a              Archive Backup
        -c              Consolidated Backup
        -r              Recover - Original Document Ids
        -s              Scan - Associate New Document Ids
        ----------- PS Application General Options -----------
        -h              This help message
```

```
-h <homedir>    Home Directory
                Location of PSLSF and DART.cfg
-q              Quiet mode
-v [domain]     Display Version - Option domain
                Provide domain to include IDMIS version
----------- Special - use only as directed -----------
-S              Spy mode
-D              Debug mode
-M              Debug Memory mode
-C <mins>       Clock Timeout
```

Usage: ***DART -h {homedir} -v -q -S –D –M –C <mins>***

NOTE When a file list is created for use with DART on Windows, the very last DocID listed in the file must be followed by a carriage return. If this carriage return is not present, the last DocID in the list will be omitted when the documents are backed up and restored.

-f	Document list is from an input filelist
-surf <surfid>	Document list will be created by DART using database information for specified surface
-both	Tells DART to use both sides of platter and is only used with -surf
-ssn <ssn>	Specifies the source **ssn** for backups or destination **ssn** for restores Default = <ssn of domain in 'cfg' file>
-t <family>	Specifies a family – Default='tranlog'

***The following options are exclusive. If none are specified, the default is for an incremental backup.

-i	Incremental backup
-a	Archive backup
-c	Consolidated backup
-r	Recover – Original Document Ids
-s	Scan – Associate New Document Ids

---- PS Application General Options ----

-h <homedir>	Location of ***Dart.cfg***
-h	This help message
-q	Quiet mode
-v	Display Version

The following options are special use options. Use only as directed by FILENET CSS.

-S	Spy mode
-D	Debug mode
-M	Debug Memory mode
-C <mins>	Clock Timeout

NOTE The ***-f, -surf***, and ***–t*** options are mutually exclusive. If more than one of these options are specified, the precedence will be ***-f <doc list>,*** followed by ***-surf <surf id>,*** and finally ***-t <family>.***

DART reads the configuration file ***DART.cfg*** for the assigned operating parameters. Even though these programs are run on the server, they function like FILENET client programs, and log on to IS using the logon ID and password found in the DART configuration file.

8.2 Stopping DART

DART will stop when it completes the work specified by the user. If you need to stop DART before the completion of the work specified, follow one of the procedures below.

1. On UNIX: ***Ctrl-C*** or ***kill –term*** will allow DART to complete the current process and set the check files correctly before exiting.

   If DART is stopped with UNIX ***kill -9*** command, it could corrupt check files and leave the ISTK environment in an indeterminate state. If ***kill –9*** is issued, run ***wal_purge*** to clean up the ISTK environment. Restart DART with discretion using the exact same parameters. DART will normally read the working files and restart from where it left off. However, the working files could have been left in an indeterminate state.

2. On Windows: If you do an ***End Process*** then restart DART, the program should start from where it left off. Restart DART with discretion using the exact same parameters.

NOTE If DART is stopped unexpectedly by the user, some IS processes might continue to run. The Prefetch and “write to cache” will continue for the set that was “in process”.

8.3 Backing up Documents

To back up documents, DART operates in one of four modes—Archive, Incremental, Consolidated, and File List. For more information on these backup modes, see Section 6.1.

8.3.1 Archive Mode

Syntax: ***DART -a*** {valid switches}

Valid switches in Archive mode are:

-h <homedir>	Points to the location of ***Dart.cfg***
-t <family>	Specifies a family name
-surf <surfid>	Tells DART which surface’s database information to use to create the document list
-both	Instructs DART to use both sides of platter and is only used with -surf
-ssn <ssn>	Indicates the source ***ssn*** for the backup

8.3.2 Incremental Mode

Syntax: ***DART -i*** {valid switches}

Valid switches in Incremental mode are:

-h <homedir>	Points to the location of ***Dart.cfg***
-t <family>	Specifies a family name

-surf <surfid>	Tells DART which surface's database information to use to create the document list
-both	Instructs DART to use both sides of platter and is only used with -surf
-ssn <ssn>	Indicates the source ***ssn*** for the backup

8.3.3 Consolidate Mode

Syntax***: DART –c*** {valid switches}

Valid switches in Consolidate mode are:

-h <homedir>	Points to the location of ***Dart.cfg***
-t <family>	Specifies a family name
-surf <surfid>	Tells DART which surface's database information to use to create the document list
-both	Instructs DART to use both sides of platter and is only used with -surf
-ssn <ssn>	Indicates the source ***ssn*** for the backup

8.3.4 File List Mode

Syntax: ***DART –f<Filename>*** {valid switches}

Valid switches in File List mode are:

-h <homedir>	Points to the location of ***Dart.cfg***
-ssn <ssn>	Indicates the source ***ssn*** for the backup

8.4 Recovering Documents

When recovering documents, DART is designed to operate in one of three modes—Standard Recovery, Forced Recovery, and Scan.

8.4.1 Standard Recovery Mode

Syntax: ***DART –r*** {valid switches}

Valid switches in Standard Recovery mode are:

-h <homedir>	Points to the location of ***Dart.cfg***
-t <family>	Specifies a family name
-surf <surfid>	Tells DART which surface's database information to use to create the document list
-both	Instructs DART to use both sides of platter and is only used with -surf
-ssn <ssn>	Indicates the destination ***ssn*** for the restore

8.4.2 Forced Recovery Mode

Syntax: ***DART –r –f<Filename>*** {valid switches}

Valid switches in Forced Recovery mode are:

-h <homedir>	Points to the location of ***Dart.cfg***
-ssn <ssn>	Indicates the destination ***ssn*** for the restore

8.4.3 Scan Mode

Syntax: ***DART –s*** {valid switches}

Valid switches in Scan mode are:

-h <homedir>	Points to the location of ***Dart.cfg***
-f <doc list>	Identifies the file containing the document Id list
-t <family>	Specifies a family name
-surf <surfid>	Tells DART which surface's database information to use to create the document list
-both	Instructs DART to use both sides of platter and is only used with -surf
-ssn <ssn>	Indicates the destination ***ssn*** for the restore

9 TROUBLESHOOTING

9.1 Utility / ISTK version relationship

The DART version selected will depend on the ISTK version in use on the IS Server. DART has been compiled to run with ISTK 3.5 as well as with ISTK 3.6. The installation CD will detect the version of ISTK and install the correct version of DART when installing on a UNIX system. When installing on a WINDOWS system you must know the ISTK version and select the correct DART version when asked during the setup procedure. To download a version of DART off of the CSS download site, verify the ISTK version by running a stamp command on an ISTK module in the ISTK install directory.

Example:

UNIX system:

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

WINDOWS system:

```
cd <drive letter>:\fnsw\client\shobj
stamp *SysV*
```

9.2 Error Logging Files

Errors that could occur during the operation of DART can be reported in several places, depending on the error type. When an IS related error is encountered, the FILENET error tuple and error text, if available, will be reported. If a DART or File System error is encountered, a DART error code will be reported. The DART error codes are generally related to an incorrect configuration, bad input file format, or a problem with the file system. If file system errors (Code 1024) occur, the System Administrator is normally responsible for fixing them.

The following is an outline of error types and locations related to DART:

1. IS error logs – all IS related messages and errors (OSAR, database, security, and so on)

   Default UNIX - `/fnsw/local/tmp/logs or /fnsw/local/logs`

   Default WINDOWS - `\fnsw_loc\logs\elogs`

2. ISTK error logs – any internal ISTK error (an IS error log specifically for the ISTK environment)

   Default UNIX - `/fnsw/client/tmp/wal`*`YYYYMMDD`*

   Default WINDOWS - `\fnsw\CLIENT\tmp\wal`*`YYYYMMDD`*

3. DART error logs – all batch progress and general error messages are reported in the journal files

   Journal files will be entitled as follows:

   DART = drt.*mmddyyyy*

Journal log files are located in the “journals” directory. The path to this “journals” directory is designated by the default path or configured path found in the ***LogDirectory*** value in the ***DART.cfg*** file

Default UNIX - `/fnsw/local/bin/journals`

Default WINDOWS - `<drive>:\fnsw_loc\bin\journals`

9.3 Procedure for Reporting Problems

If you encounter a problem with the functionality of DART, recreate the problem with a minimum amount of data. If necessary, add the debug option '-D', the spy option '-S' and/or the memory option '-M'. When in doubt of which flag to use, use all three. These debugging messages will be added to the journal log file. Open a call with CSS, and provide the following data:

1. Site name
2. System Configuration
   - Operating System platform and version
   - Database platform and version

     Run `fn_util whichfn` – supply output to CSS
   - IS Domain, Organization, and ssn

     Run `nch_check` – supply output to CSS
   - IS version and stamp for module:

     `stamp /fnsw/lib/shobj/*SysV*`
   - ISTK version and stamp for module:

     `stamp /fnsw/client/shobj/*SysV*`
   - Directory listing of ISTK to verify permissions

     `(ls –alR /fnsw/client > filename)`
   - Is ISTK/Utility installed directly on IS or on a remote system? If remote, supply the Operating System for remote environment.
   - Version of DART on Windows (use “DART –v”)
   - Version of DART on UNIX (use “stamp DART”)
   - Directory listing of DART with permissions (ls –al /fnsw/local/bin > *filename*)
3. CheckSSN output from display in DOS prompt or Command window
4. Full description of problem
5. Steps to reproduce problem
6. Exact command used to start/stop DART
7. Output from display in DOS prompt or Command window of DART runtime
8. Utility Journal logs (/fnsw/local/bin/journals/*)
9. DART Configuration file (***DART.cfg***)
10. Input files used (DocID list files, eob files, transact.dat files, images, and so on)

11. ISTK Error logs (/fnsw/client/tmp/*walyyyymmdd*)
12. IS Error Logs (elog#########)
13. Stack trace of core file (if applicable)

9.4 When to use wal_purge – UNIX ONLY

The ISTK program, **wal_purge,** can clean up numerous problems but if used incorrectly when ISTK programs are running, can cause application/program failures. Only run **wal_purge** when ALL ISTK applications are stopped. Run a script immediately after bringing up the FILENET software but before starting any ISTK applications/programs.

Run **wal_purge** to clean up an application problem in any of the following situations:

- If the IS software is shut down before DART is stopped
- If you make any changes to the ***DART.cfg*** file
- If DART fails to recognize a newly created doc class or index
- If you encounter security problems with DART

9.5 Using a File List on Windows

When you create a file list to use with DART, insert a carriage return after the last DocID in the list. If this carriage return is not present, that last DocID will be omitted when the documents are backed up and restored.

9.6 OS path issue

Correctly set all ISTK paths in /etc/profile to avoid problems with DART.

9.7 HP CDROM mounting for ISO9660 CD type

DART is released on a CD type that is created for all supported operating systems. Use the following mount command on HP operating systems:

```
mount –o cdcase /dev/cd0 /cdrom
```

9.8 DART – restore by surface or family

Remove the following DART backup files before you use DART restore with the surface "-surf" or family "-t" flags.

```
*.surf
*.chkpt
*.mark
*.temp
```

9.9 Error codes known to need specific actions

15,16,17 – This is an IS network error and must be resolved by the Network Administrator or FILENET CSS.

80,1,61 – "Duplicate write request" errors occur when DART is run twice using the same Doc IDs. To resolve, wait for the first write requests to be complete, or delete the pending write requests before re-running the second request.

0,2,#### – folder/file path does not exist

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

156,0,24 - <NCH,0,24> "A network related error was encountered."

The most common reasons for this error are:

- An incorrect "nch-server" entry in the /etc/hosts file.
- An incorrect Domain or Organization name in the ***DART.cfg*** file.
- The FILENET IS software is stopped.
- The network is down.

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

This error usually means that the ISTK libraries are not installed properly and DART is accessing the IS libraries instead of the ISTK libraries.

202,0,9 – SysV segmentation violation: Problem between IS and ISTK 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 – ISTK install incomplete (Modules are not copied down entirely or the configuration is incomplete.)

202,0,9 – ISTK and Utility compile mismatch (DART is compiled for each ISTK version. See section on Hardware and Software requirements.)

202,0,9 – Utility installation incomplete (Not all modules were copied down or permissions on DART modules are incorrect on UNIX.)

202,0,9 – DART WorkingDirectory path is too long (A path exceeding 35 characters could exceed the path limit and cause the 202,0,9 error. The HP OS seems to be the most sensitive to this issue.)

202,0,9 – SysV module from IS is being accessed instead of the SysV module in the ISTK installed directory. Check your PATH to see that the ISTK path is listed before the IS path. Tests have found problems where the IS shobj path has been input into another PATH value that causes the IS SysV to load when an ISTK call is made thus causing a segmentation violation.

Error output that can be found in the nohup log (/tmp/MRII_import*yyyymmdd*) file:

```
sys_log: fnc_get_process_name error: program is not registered.
pid = 15701

Process aborting due to segmentation violation...
```

9.10 Error Codes

This table is a general list of errors that could be encountered when using DART. Included is a description of their meaning where appropriate. These error code numbers can change in subsequent builds, and might not be reflected in this document.

Error Number	Error Text	Error Information
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 2nd 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	
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	
Error Number	**Error Type**	**Qualified Error**
1030	IO error	Could not create open or read directory. Usually a permissions problem or the file system ran out of space or an NFS mount was lost
1031	Critical IO error – Program Abort	
1032	Can't kill a child process	
1033	Line too long	
1034	malloc failed	Unable to get an allocation of memory from the operating system.
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	ISR error on attempt to complete Activity Logging Event on an annotation.
1040	Document does not contain any page	
1042	One or more critical resources are not available	
1043	Not a valid FileNET format. See SC_convert.out for details	
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	
Error Number	**Error Type**	**Qualified Error**
1053	Too many docs returned by query. Must query on unique index	
1054	Batch contains more than one family	Multiple Families are NOT supported.
1055	too many pages in doc	Document exceeds the max page limit of %d. Ln='%s'
1057	Transact file contains incorrect format	
1058	Required input file was not provided	
1059	PS Shared Library Mismatch-Contact FileNET Support	
1060	EMC-Centera Support not yet available-No ClipId Produced	
1061	EMC-Centera device IO error	
1062	ISR Validation Failure	
1064	Preliminaries Complete	

10 APPENDIX A–EPHEMERAL PORT SETTINGS

A TCP/IPv4 connection consists of two endpoints, and each endpoint consists of an IP address and a port number. Therefore, when a client user connects to a server computer, an established connection can be thought of as the 4-tuple of: (server IP, server port, client IP, client port). Usually three of the four are readily known – the client machine uses its own IP address and when connecting to a remote service, requires the server machine's IP address and service port number.

What is not immediately evident is that when a connection is established, t the client side of the connection uses a port number. If a client program does not explicitly request a specific port number, an ***ephemeral*** port number is used. Ephemeral ports are temporary ports assigned by a machine's IP stack, and are assigned from a designated range of ports for this purpose. When the connection terminates, the ephemeral port is available for reuse, although most IP stacks won't reuse that port number until the entire pool of ephemeral ports have been used. So, if the client program reconnects, it will be assigned a different ephemeral port number for its side of the new connection.

Similarly, for UDP/IP, when a datagram is sent by a client from an unbound port number, an ephemeral port number is assigned automatically so the receiving end can reply to the sender.

Symptom and Description:

The default settings that come pre-configured with the OS are often insufficient for high-volume FileNet activity. As a result, the OS can intermittently run out of free socket ports and not be able to open any new TCP connections, and activity will halt. The system might not be able to recover once free ports become available again. To avoid these operational interruptions, you must tune the ephemeral ports.

The system might be susceptible to this problem if your application mix involves a high number of short server connections.

- Symptoms of the problem could include:
    - Client intermittently receives "Method of object failed" during Doc.GetCachedFile
    - IS Toolkit logs get chronic SPP type errors (for example, 15,16,17):

      COR_Open: connect failed with errno 10048
    - Workflo Queue intensive app gets:

      "COR_Open: connect to 30.34.192.237 [32769] failed with errno 227"

- Use either of the two methods to tune the ephemeral ports:
    - Increase the range of ports available to the OS
    - Decrease the "TIME_WAIT" period for which the OS cannot reuse a closed port number. (Reducing TIME_WAIT is NOT CONFIGURABLE on AIX and HPUX 10.20)

  Both methods are platform-dependent. Refer to the sections below for the correct method for your OS.

  Both methods are independent of each other - you can do one, the other, or both.

NOTE We recommend increasing the ephemeral port range when you install the IS server software. We don't recommend modifying TIME_WAIT unless the need arises.

There are no hard and fast guidelines for decreasing TIME_WAIT. The default value is usually 4 minutes. It can often safely be reduced to as little as 2 minutes or even 30 seconds. TIME_WAIT is designed to prevent any "lost packets" from an old connection from having the same port number as a current connection. If you're in a high-latency environment (for example, WAN or satellite traffic), your TIME_WAIT should be longer (for example, the full 4 minutes). If you're certain that all packets will be received promptly (within milliseconds), then you can safely reduce TIME_WAIT.

10.1 Windows platform

While the ephemeral ports must be adjusted for all platforms, the Windows platform is the most affected if the ports are not modified.

- Make the modifications on all the servers that are part of the FileNet domain.
- Use the Registry Editor (regedt32.exe) to make the modifications.

10.2 MaxUserPort

The *MaxUserPort* determines the highest port number TCP can assign when an application requests an available user port from the system. Typically, ephemeral ports (those used briefly) are allocated to port numbers 1024 through 5000.

NOTE Windows does not add this entry to the registry. You can add it by editing the registry or by using a program that edits the registry.

Location: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

Data Type: REG_DWORD

Range: 5,000-65,534 (port number)

Default Value: 5000

Recommended value: 65534 (65534 DEC)

10.3 TcpMaxConnectTransmissions

Location: HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

Data Type: REG_DWORD

Recommended value: 5 (5 DEC)

10.4 TcpMaxConnectRetransmissions

The ***TcpMaxConnectRetransmissions*** determines the number of times TCP will retransmit an unanswered request for a new connection. TCP retransmits new connection requests until they are answered or until this value expires

Location:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

Data Type: REG_DWORD

Range: 0–255 (*retransmission attempts*)

Default Value: 2

Recommended value: 5 (5 DEC)

10.5 TcpTimedWaitDelay

The ***TcpTimedWaitDelay*** determines the time that must elapse before TCP can release a closed connection and reuse its resources. This interval between closure and release is known as the TIME_WAIT state or 2MSL state. During this time, the connection can be reopened at much less cost to the client and server, than establishing a new connection.

RFC 793 requires that TCP maintain a closed connection for an interval at least equal to twice the maximum segment lifetime (2MSL) of the network. When a connection is released, its socket pair and TCP control block (TCB) can be used to support another connection. By default, the MSL is defined to be 120 seconds, and the value of this entry is equal to two MSLs, or 4 minutes. However, you can use this entry to customize the interval.

Reducing the value of this entry allows TCP to release closed connections faster, providing more resources for new connections. However, if the value is too low, TCP might release connection resources before the connection is complete, requiring the server to use additional resources to reestablish the connection.

NOTE Normally, TCP does not release closed connections until the value of this entry expires. However, TCP can release connections before this value expires if it is running out of TCP control blocks (TCBs). The number of TCBs the system creates is specified by the value of the MaxFreeTcbs <58770.asp> entry.

Windows does not add this entry to the registry. You can add it by editing the registry or by using a program that edits the registry.

Location:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

Data Type: REG_DWORD

Range: 30-300 seconds

Default Value: 240 seconds = 4 minutes

Recommended value: 45 seconds (45 DEC)

10.6 FN_COR_QLEN

15,16,17 errors – WSAECONNREFUSED 10061 – indicates that the NLS_Archiver can not connect to the IS COR_Listen due to unavailability of COR queue space.

To resolve the 15,16,17 error listed here, on Windows, add the variable ***FN_COR_QLEN***, in **My Computer >Advance > Environment Variables** with a value of 20-25 (the default on UNIX).

10.7 UNIX Platform

Possible issues with ephemeral ports on UNIX platforms

On the IS server check to see if the UDP low ephemeral port is set incorrectly. We've only seen this on AIX so far, but it might happen on Solaris as well.

- On AIX enter 'no -a | grep udp_ephemeral'
- On Solaris enter ' ndd /dev/udp udp_smallest_anon_port' , then ' ndd /dev/udp udp_largest_anon_port'
- On HP 10.30/11.x enter ' ndd /dev/udp udp_smallest_anon_port' , then ' ndd /dev/udp udp_largest_anon_port'
- On HP 10.20, this shouldn't be a problem, but you can check to see if high_port_enable flag has been set by entering '/usr/contrib/bin/nettune -l | grep high' (the switch is a lower case L)

In all cases the number for the low end of the range should be higher than 32771. By default most UNIX implementations start this value at 32768, and this can allow other programs to come in and step on NCH_daemon.

10.8 IBM AIX

On AIX, check the /etc/rc.net file (or on AIX 5.2 in the /etc/tunables/nextboot file) for the following lines:

/usr/sbin/no -o udp_ephemeral_high=65535

/usr/sbin/no -o udp_ephemeral_low=42767

For a detailed list of steps to modify the /etc/rc.net file, see the IS 4.0 Installation and Configuration Procedures for AIX/6000.

- Perform the steps in this section on all servers. Backup the /etc/ rc.net file. As root user, make sure you have write permission on the above files by entering:
- chmod 754 /etc/ rc.net

Use your preferred text editor (such as vi) to modify the /etc/ rc.net file.

Locate the following 'if' statement, near the end of the file.
##

 # The socket default buffer size (initial advertised TCP window) is being

 # set to a default value of 16k (16384). This improves the performance

 # for ethernet and token ring networks. Networks with lower bandwidth

```
# such as SLIP (Serial Line Internet Protocol) and X. 25 or higher bandwidth
# such as Serial Optical Link and FDDI would have a different optimum
# buffer size.
# (OPTIMUM WINDOW = Bandwidth * Round Trip Time )
#####################################################
if [ -f /usr/ sbin/ no ] ; then
        /usr/sbin/ no -o tcp_sendspace= 16384
        /usr/sbin/ no -o tcp_recvspace= 16384
fi
```

Add the following lines:

```
/usr/sbin/ no -o tcp_keepidle= 80
/usr/sbin/ no -o tcp_keepintvl= 20
/usr/sbin/ no -o tcp_ephemeral_ high= 65535
/usr/sbin/ no -o tcp_ephemeral_ low= 42767
/usr/sbin/ no -o udp_ephemeral_ high= 65535
/usr/sbin/ no -o udp_ephemeral_ low= 42767
```

so the "if" statement looks like this:

```
if [ -f /usr/sbin/ no ] ; then
        /usr/sbin/ no -o tcp_sendspace= 16384
        /usr/sbin/ no -o tcp_recvspace= 16384
        /usr/sbin/ no -o tcp_keepidle= 80
        /usr/sbin/ no -o tcp_keepintvl= 20
        /usr/sbin/ no -o tcp_ephemeral_ high= 65535
        /usr/sbin/ no -o tcp_ephemeral_ low= 42767
        /usr/sbin/ no -o udp_ephemeral_ high= 65535
        /usr/sbin/ no -o udp_ephemeral_ low= 42767
fi
```

Save your changes and exit from the file.

Reboot the server (s)

10.9 HP-UX

HPUX 10.20 ONLY

Perform the steps in this section on all servers. You can make your FileNet system run more efficiently by making changes to the /etc/rc.initfnsw file. The modification expands the number of available ephemeral ports. These modifications are not required, but have been found to be optimal when running FileNet software. So unless you have set these options for other system reasons, we recommend that you make these changes.

Ephemeral ports are temporary ports assigned by a server's IP stack, and are assigned from a designated range of ports for this purpose. When network traffic is extremely heavy, it's possible to run out of ephemeral ports unless you specify the high_ port_ enable option in /etc/rc.initfnsw. Backup the /etc/rc.initfnsw file. As root user, make sure you have write permission on the this file by entering:

```
chmod 754 /etc/rc.initfnsw
```

Use your preferred text editor (such as vi) to modify the /etc/rc.initfnsw file.

Locate the following statement near the end of the file: **# Set up network options HPUX (parameters in half- seconds)**

Add the following line somewhere after that statement:

```
/usr/contrib/bin/nettune -s tcp high_ port_ enable 1
```

Save your change and exit from the file.

Reboot server (s)

HPUX 10.30/HPUX 11.x

As root user, vi /sbin/rc2.d/S340net =>

```
ndd -set /dev/udp udp_smallest_anon_port 42767
ndd -set /dev/udp udp_largest_anon_port  65535
ndd -set /dev/tcp tcp_smallest_anon_port 42767
ndd -set /dev/tcp tcp_largest_anon_port  65535
ndd -set /dev/tcp tcp_time_wait_interval 30000 (for HPUX
10.30/11.x)
```

Save your change and exit from the file.

Reboot server (s)

10.10 Sun Solaris

Perform the steps in this section on all servers. On Solaris, you set these parameters in /etc/rc2.d/S69inet - double check the install documentation for syntax, or contact CSS Network Support for assistance.

You can make your FileNet system run more efficiently by making changes to the /etc/rc2.d/S69inet file. The modification expands the number of available ephemeral ports and reduces the time-out delay. These modifications are not required, but have been found to be

optimal when running FileNet software. So unless you have set these options for other system reasons, we recommend that you make these changes.

Ephemeral ports are temporary ports assigned by a server's IP stack, and are assigned from a designated range of ports for this purpose. When network traffic is extremely heavy, it's possible to run out of ephemeral ports unless you specify a wider range of port numbers in /etc/rc2.d/S69inet.

The tcp_close_wait_interval parameter determines the length of time the server waits before reusing a closed ID socket. Although the default value is typically around 240000 milliseconds (four minutes), this parameter can safely be reduced to as little as 30000 milliseconds (30 seconds) on high-speed networks.

Make a backup copy of the /etc/rc2.d/S69inet file before you modify it.

As root user, make sure you have write permission on the file by entering:

```
chmod 754 /etc/rc2.d/S69
```

Use your preferred text editor (such as vi) to modify the /etc/rc2.d/S69inet file.

Add the following lines somewhere near the end of the file:

```
ndd -set /dev/udp udp_smallest_anon_port 42767
ndd -set /dev/udp udp_largest_anon_port 65535
ndd -set /dev/tcp tcp_smallest_anon_port 42767
ndd -set /dev/tcp tcp_largest_anon_port 65535
ndd -set /dev/tcp tcp_close_wait_interval 30000 (for Solaris 2.x only)
ndd -set /dev/tcp tcp_time_wait_interval 30000 (for Solaris 8 and above)
```

Save your change and exit from the file.

Reboot server (s)