

EMC DOCUMENTUM xPLORE 1.1 DISASTER RECOVERY USING EMC NETWORKER

Abstract

The objective of this white paper is to describe the architecture of and procedure for configuring EMC® Documentum® xPlore 1.1 disaster recovery using EMC NetWorker®.

May 2011

Copyright © 2011 EMC Corporation. All Rights Reserved.

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

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

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

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

All other trademarks used herein are the property of their respective owners.

Part Number h8249

Table of Contents

Executive summary	4
Introduction	4
Audience.....	4
Back up and restore Documentum xPlore with NetWorker	4
Architecture description	4
Overview of the xPlore backup process	6
Supported backup technologies, levels, and techniques.....	7
File backup	8
File, full, cold backup for xPlore federations	8
File, full, warm backup for xPlore federations	10
Native-xDB backup	11
Native-xDB, full, cold backup for xPlore federations	12
Native-xDB, full, hot backup for xPlore federations	12
Native-xDB, full, incremental backup for xPlore federations.....	12
Native-xDB, full, warm backup for xPlore domains.....	13
Native-xDB, full, warm backup for xPlore collections	13
Volume backups	14
Volume, full, cold backup for xPlore federations.....	14
Volume, full, warm backup for xPlore federations.....	14
Conclusion	15

Executive summary

Backup and recovery is a key component of business continuity. “Business continuity” is the umbrella term that covers all efforts to keep critical data and applications running despite any type of interruption (both planned and unplanned). Planned interruptions include regular maintenance or upgrades. Unplanned interruptions could include hardware or software failures, data corruption, natural or man-made disasters, viruses, and human errors. Backup and recovery is essential for operational recovery; that is, recovery from errors that can occur on a regular basis but are not catastrophic—for example, data corruption or accidentally deleted files. Disaster recovery is concerned with catastrophic failures.

Introduction

Enterprise content management has become a key component for global companies to conduct business efficiently. Powerful search is one of the critical requirements, including metadata searches as well as full-text searches of the content itself. In this white paper, the EMC® Documentum® xPlore high-availability deployment is described along with a detailed setup example using Microsoft Clustering Services (MSCS).

Audience

This paper is written for IT staff who are considering how to back up and restore Documentum xPlore 1.1 using EMC NetWorker®.

Back up and restore Documentum xPlore with NetWorker

NetWorker provides the ability to protect enterprises against data loss. NetWorker software is a cross-platform, client-server application that remotely manages all NetWorker clients and servers from a web-enabled, graphical interface.

xPlore is a multi-instance, scalable, high-performance, full-text index server that can be configured for high availability and disaster recovery. xPlore provides command-line utilities for backup and restore.

The focus of this white paper is on integrating NetWorker with xPlore to achieve disaster recovery.

Architecture description

You can negatively impact performance if you concentrate the content processing, indexing, and query activities on a single instance in a large xPlore deployment with a high volume of documents to index. To improve performance, you can set up multiple secondary xPlore instances on the same or different machines. This distributed

configuration improves performance by enabling parallel index processing. Multi-instance deployments are best used where large volumes of data are indexed and searched and high performance is required. In this specific deployment, two instances (the primary and one secondary) are deployed.

The following diagram shows the topology for the components and connections.

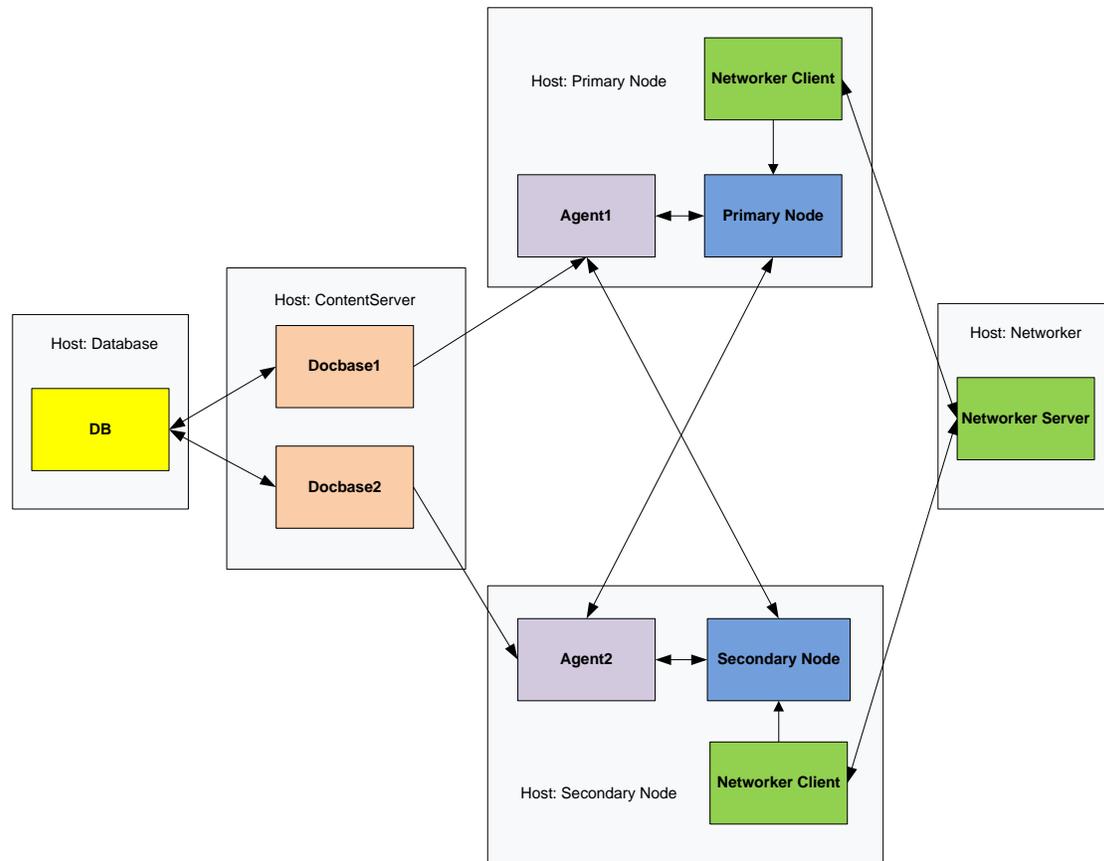


Figure 1. xPlore disaster recovery deployment topology

In this deployment:

- There is one connection broker and two repositories (Docbase1 and Docbase2).
- For the xPlore deployment, two hosts are deployed: One hosts the xPlore primary instance, and the other one hosts the secondary instance. On the primary host, an index agent is deployed and used by Docbase1, while the index agent on the secondary host is used by Docbase2.

For details on deployment of multi-instance xPlore, please refer to the *EMC Documentum xPlore Deployment Guide*.

- NetWorker Server is deployed on a separate host where the EMC NetWorker server, NetWorker console, and NetWorker storage reside.
- A NetWorker client is installed on each xPlore host.

For details on installing EMC NetWorker, please refer to the *EMC NetWorker 7.6 Installation Guide*.

Overview of the xPlore backup process

The xPlore backup process is usually composed of the following steps:

1. Pre-processing (optional)
2. An xPlore backup with the xPlore command line (required only when using native-xDB backup)
3. A NetWorker backup
4. Post-processing (optional)

Pre-processing is usually required for cold and warm backups. This pre-processing consists of suspending or stopping the xPlore application or service to prepare it for the actual backup process.

An xPlore backup is required only when performing a native-xDB backup. xPlore backup commands are used to generate a dump file based on the current xPlore data files, log files, and configuration files.

A NetWorker backup is a backup of the xPlore files or the dump file generated by xPlore native backup with NetWorker command line or console.

Post-processing is required for cold and warm backups to resume the xPlore application or service after the backup has completed. In many cases, it also cleans up the files that have been backed up to NetWorker storage.

The following table lists the typical backup scenarios. The [Supported backup technologies, levels, and techniques](#) section provides explanations of the individual backup technologies, levels, and techniques.

Table 1. Backup scenarios

	Scenarios by backup technology, level, and technique
Scenario 1	File, full, cold backup for xPlore federation
Scenario 2	File, full, warm backup for xPlore federation
Scenario 3	Native-xDB, full, cold backup for xPlore federation
Scenario 4	Native-xDB, full, hot backup for xPlore federation
Scenario 5	Native-xDB, full, incremental backup for xPlore federation
Scenario 6	Native-xDB, full, warm backup for xPlore domain
Scenario 7	Native-xDB, full, warm backup for xPlore collection
Scenario 8	Volume, full, cold backup for xPlore federation
Scenario 9	Volume, full, warm backup for xPlore federation

Note: Restore is not included with the listed backup scenarios since it does not require automatic execution with NetWorker as is required for backup scenarios.

Supported backup technologies, levels, and techniques

The following tables describe the backup technologies, levels, and techniques supported by xPlore.

Table 2. Backup technologies

Backup technology	Description
File	A backup of individually listed files or entire directories (and their contents). Involves copying the files from their current location to a backup staging area. The most basic form of backup, but has limited support on Windows.
Volume	A backup that is based on the underlying blocks in the file system. Typically does not have knowledge of the file system. Typically requires some extra technology to assist (for example, the Linux <code>dd</code> command).
Native-xDB	Backup files created by the xPlore command-line interface.

Table 3. Backup level

Backup level	Description
Full	A complete copy of all of the required files.
Incremental	A backup in which only the files or blocks that have been <i>changed</i> since the last backup are saved (incremental or full). You must first apply the last full backup and all incremental backups in sequence up to the most recent incremental backup. Therefore, the recovery will be successful only up to the last good incremental backup (because each incremental backup depends on the previous ones).
Differential	A rollup of all incremental backups since the last full backup. The restore time for a differential backup is shorter than an incremental, but the amount of data backed up (per differential) is larger.

Table 4. Backup technique

Backup technique	Description
Cold	xPlore is down for the duration of the backup. The simplest and safest form of backup.
Warm	xPlore is up and servicing queries, but index agents do not send data during the duration of backup.
Hot	xPlore services queries and ingests data.

File backup

NetWorker directly backs up specified xPlore files. Hot backup is not supported in this scenario because indexing during the backup process could change these files and make the backup inconsistent.

At least the following directories must be backed up:

- /config
- /data
- /dblog

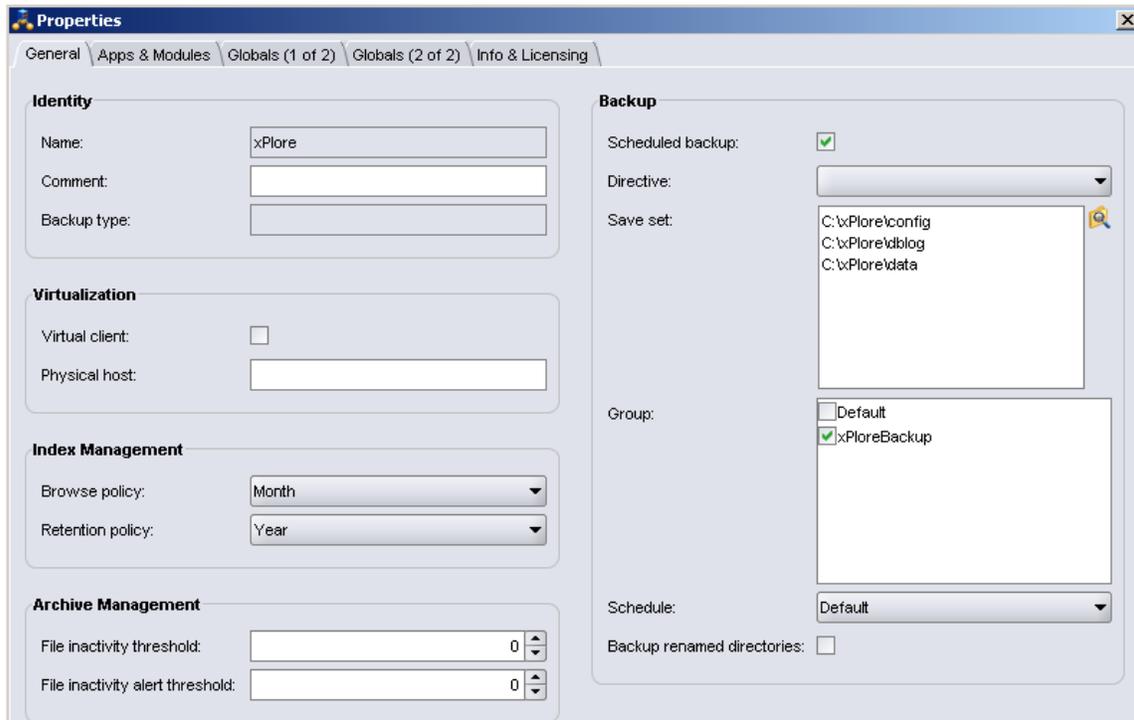
File, full, cold backup for xPlore federations

Use NetWorker to configure a full, cold, file backup of an xPlore federation as follows:

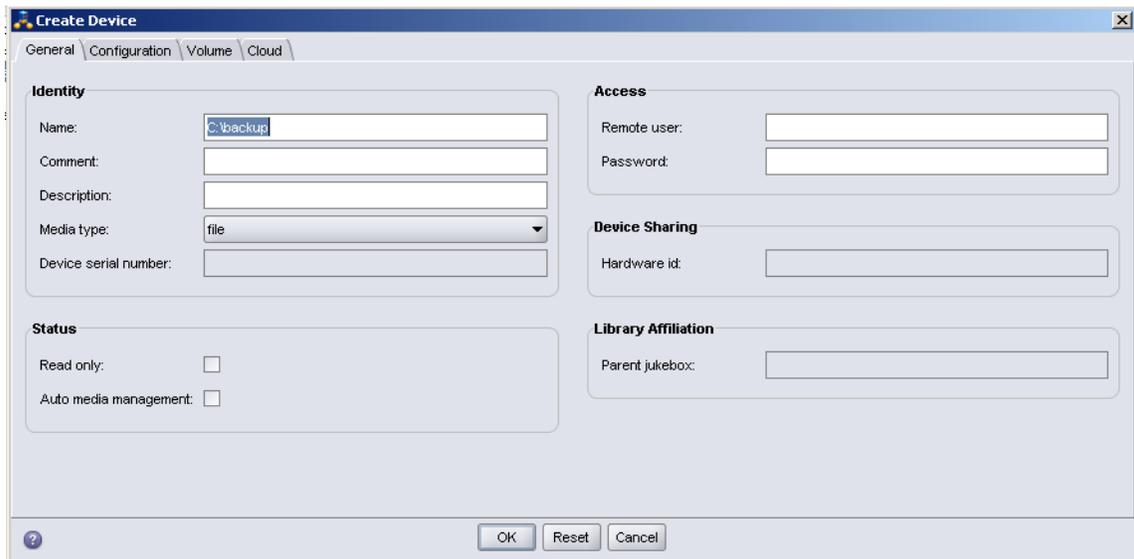
1. In NetWorker, create a group, **xPloreBackup**, and in the properties of the group, set a start time for running the backup.



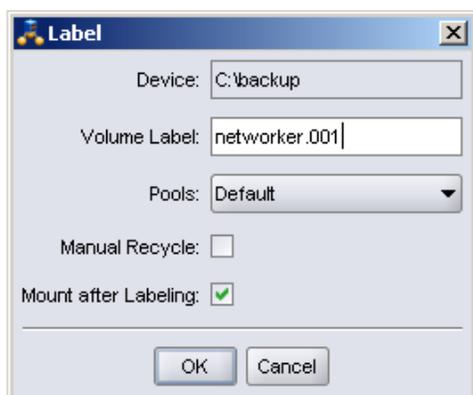
2. Create a new client in the **xPloreBackup** group, specify a **Save set**, and add at least the /config, /data, /dblog directories to the save set.



3. Create a **device** based on a Windows folder, C : \backup, and set **Media type** to **file**.



4. Label this device as `networker.001`.



Note: The previous steps are required for every backup scenario in which NetWorker is used as the backup application.

Because you are configuring a cold backup, you must shut down xPlore and the index agent manually before the backup. Otherwise, at the scheduled start time, NetWorker automatically starts to back up the three specified xPlore directories.

Follow these steps to restore xPlore from the NetWorker backup.

Note: If you are performing a federation restore, clean up all xPlore data, configuration, and log files before performing these steps.

1. On the NetWorker client machine that hosts the xPlore primary instance, execute this command syntax on the NetWorker command line to connect to the NetWorker server:

```
recover -s <NetWorker_server> -c <NetWorker_client> <backup_folder>
```

For example:

```
recover -s networker -c xPlore C:\xPlore
```

2. After connecting to the NetWorker server, execute the following commands to select the files you want to recover:

```
add C:\xPlore\config
```

```
add C:\xPlore\data
```

```
add C:\xPlore\dblog
```

3. Start the restore by executing the following command:

```
recover
```

4. Start xPlore.

File, full, warm backup for xPlore federations

Because indexing is stopped but the query service is allowed to run during the warm backup process, you must run a script that suspends the disk(s) used by xPlore. For instructions, follow the procedure in the [Native-xDB backup](#) section.

Native-xDB backup

In a native-xDB backup, NetWorker is configured to call a customized script. The following steps describe the overall process:

1. On the NetWorker client machine (for example, the xPlore primary node), create a NetWorker command-line script file (*.bat/*.sh) to suspend the disk and back up the files.

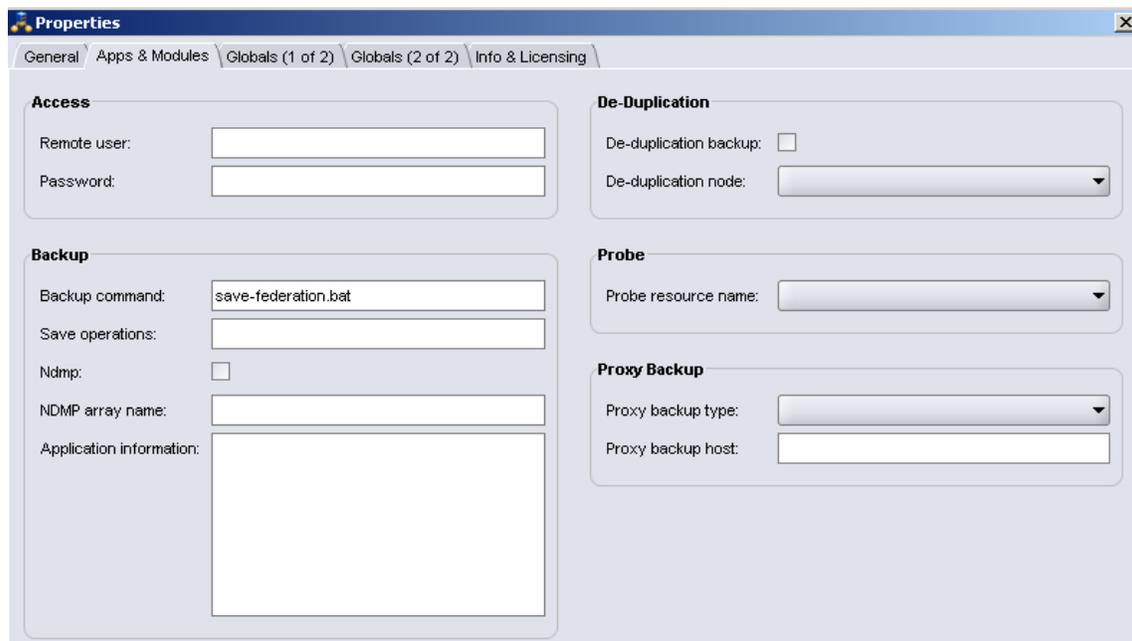
The file must meet the following requirements:

- The name of the script must begin with save.
- The file must be located in the /NetWorker_install_dir/nsr/bin directory. It is recommended that you copy xplora.sh/bat to that directory for the script to easily call the CL command line. Since the xPlore command must read an xplora.properties file, you must add the following to JAVA_OPTS in xplora.sh:

```
-Dxplora.cli.config=<Path to xplora.properties>
```

- The NetWorker save command must be called in the script. On Linux, you must append \$@ to save.
2. On the NetWorker console, on the client's **Properties** page, select the **Apps & Modules** tab, and input the name of the script file in the **Backup command** field.

With this configuration, when the group starts a backup, the NetWorker server calls the script file instead of the default backup method.



The screenshot shows the 'Properties' dialog box with the 'Apps & Modules' tab selected. The 'Backup' section is highlighted, showing the 'Backup command' field set to 'save-federation.bat'. Other sections include 'Access', 'De-Duplication', 'Probe', and 'Proxy Backup'.

Section	Field Name	Value / State
Access	Remote user:	[Empty]
	Password:	[Empty]
De-Duplication	De-duplication backup:	<input type="checkbox"/>
	De-duplication node:	[Dropdown]
Backup	Backup command:	save-federation.bat
	Save operations:	[Empty]
	Ndmp:	<input type="checkbox"/>
	NDMP array name:	[Empty]
Probe	Probe resource name:	[Dropdown]
	Application information:	[Text Area]
Proxy Backup	Proxy backup type:	[Dropdown]
	Proxy backup host:	[Empty]

The following sample script executes a warm, file backup:

Note: This script must include NetWorker backup commands.

```
#!/bin/bash

#start keep disk in suspend state
/usr/sbin/xplore.sh "suspendDiskWrites()"

#NetWorker internal command to backup directories to server
save "$@" /root/dss/config /root/dss/data /root/dss/dblog

#stop the disk on suspend state
/usr/sbin/xplore.sh "resumeDiskWrites()"
```

Native-xDB, full, cold backup for xPlore federations

This sample script executes a native-xDB, full backup. In this script, an xPlore backup command is executed before the NetWorker backup and post-processing that delete the xPlore backup dump files are performed.

Note: This script runs on Linux. The \$@ characters are appended to save.

```
#!/bin/bash

# backup federation with xPlore CL
/usr/sbin/xplore.sh "backupFederation '$BACKUP',false,null"

# backup the dump file with NetWorker's CL
save "$@" /$BACKUP/federation

# Clean the dump file on the NetWorker client machine.
rm -r /$BACKUP/federation/*
```

To restore from a native-xDB backup:

1. Restore from the NetWorker backup.
2. Restore xPlore using the xPlore command line.

For details, refer to the *EMC Documentum xPlore Administration Guide*.

Native-xDB, full, hot backup for xPlore federations

A hot backup allows both querying and indexing to run during the backup. You can use the same cold federation backup script for this backup process.

Native-xDB, full, incremental backup for xPlore federations

Because an incremental backup must be based on a full backup, the restore must first be performed from a full backup.

This sample script performs a native-xDB, incremental backup.

```
#!/bin/bash

# backup federation with xPlore CL
/usr/sbin/xplore.sh "backupFederation '$BACKUP',true,null"

# backup the dump file with NetWorker's CL
save "$@" /$BACKUP/federation

# Clean the dump file on the NetWorker client machine.
rm -r /$BACKUP/federation/*
```

Native-xDB, full, warm backup for xPlore domains

A domain is a subset of a federation in xPlore. Backing up and restoring a domain does not impact any other domains. Domain backups only allow hot or warm backups. Cold backups are not supported.

This sample script executes a domain backup. The domain must be specified.

```
#!/bin/bash

# backup domain with xPlore CL
/usr/sbin/xplore.sh "backupDomain '$DOMAIN','$BACKUP'"

# backup the dump file with NetWorker's CL
save "$@" /$BACKUP/federation

# Clean the dump file on the NetWorker client machine.
rm -r /$BACKUP/federation/*
```

Native-xDB, full, warm backup for xPlore collections

A collection is a subset of a domain in xPlore. Backing up and restoring a collection does not impact any other collections. Collection backups only allow hot or warm backups. Cold backups are not supported.

This sample script executes a collection backup. The collection and its associated domain must be specified.

```
#!/bin/bash

# backup collection with xPlore CL
/usr/sbin/xplore.sh "backupCollection collection('$DOMAIN','$COLLECTION'), '$BACKUP'"

# backup the dump file with NetWorker's CL
save "$@" /$BACKUP/federation

# Clean the dump file on the NetWorker client machine.
rm -r /$BACKUP/federation/*
```

Volume backups

Volume backup is only supported in Linux. The `dd` command is used to back up a physical or logical partition.

Note: Because the `dd` command generates a considerable amount of data and the backup can be quite slow, volume-based backups should be considered only after you have determined that file and native-xDB backups do not meet your requirements.

To use the `dd` command to back up a partition, EMC recommends that you:

- Store all xPlore data and configuration files on a stand-alone disk.
- Back up the data to another stand-alone disk; that is, not the partition on which the system files reside.

To begin the backup using the `dd` command, stop the `nfs` service and unmount the device to keep the disk in a no-access state in case of a cold backup.

NetWorker backs up the file exported by the `dd` command to a NetWorker storage location. For a volume-based backup, NetWorker must use the customized backup script located on the xPlore primary host.

The `dd` command is also used to restore the backup. Before restoring from the volume-based backup, you must first stop the `nfs` service and unmount the device.

Volume, full, cold backup for xPlore federations

This sample script executes a volume, full, cold backup.

```
#!/bin/bash

# stop nfs service
stopnfs.sh

# dd backup the device which store the xPlore data.
dd if=/dev/sdc1 of=/mnt/sdb/coldnw.img

# NetWorker backup CL
save "$@" /mnt/sdb/coldnw.image

# Clean the dump file on the NetWorker client machine.
rm -f /mnt/sdb/coldnw.img
```

Volume, full, warm backup for xPlore federations

For warm backups, you are not required to stop the `nfs` service and unmount the device. The `suspend-disk-write` script should be included to ensure that the backup is consistent.

This sample script executes a volume, full, warm backup:

```
#!/bin/bash

# start keep disk in suspend state
/usr/sbin/xplore.sh "suspendDiskWrites()"

# dd backup the device which store the xPlore data.
dd if=/dev/sdcl of=/mnt/sdb/warmnw.image

save "$@" /mnt/sdb/warmnw.image

# stop disk on suspend state
/usr/sbin/xplore.sh "resumeDiskWrites()"

# Clean the dump file on the NetWorker client machine.
rm -f /mnt/sdb/warmnw.img
```

Conclusion

This white paper has described the different backup levels, techniques, and technologies when using EMC NetWorker to back up xPlore 1.1.