

Connecting EMC DiskXtender for Windows to EMC Centera

Best Practices Planning

Abstract

This white paper provides details on building the connection string that EMC[®] DiskXtender[®] for Windows uses to connect to EMC Centera[®]. Guidelines include recommendations on virtual pools, profiles, and replication considerations.

June 2008

Copyright © 2008 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 H5501

Table of Contents

| | |
|--|-----------|
| Executive summary | 4 |
| Introduction | 4 |
| Audience | 4 |
| Overview | 4 |
| Understanding the connection process | 5 |
| Access nodes | 5 |
| Pools and profiles | 6 |
| Understanding the impact of replication | 6 |
| Replication and DiskXtender files | 7 |
| Replication failover | 7 |
| Read-only failover | 7 |
| Read/write failover | 8 |
| No failover | 8 |
| Building the connection string | 8 |
| Connection configuration quick reference | 10 |
| Conclusion | 10 |
| References | 10 |

Executive summary

This white paper provides details on building the connection string that EMC® DiskXtender® for Microsoft Windows uses to connect to EMC Centera®. The contents of the connection string include the address information for the access nodes on EMC Centera and the profile that DiskXtender uses to connect to EMC Centera. In addition, a DiskXtender option controls the connection behavior when an EMC Centera cluster fails in a replication environment.

It is important to understand how DiskXtender connects to EMC Centera in order to determine the appropriate elements for the connection string and the correct selection for the replication option. Building the correct connection string ensures that files on a DiskXtender extended drive are migrated successfully to EMC Centera, and helps to prevent data loss.

Introduction

This white paper provides an explanation of how DiskXtender connects to EMC Centera, as well as how EMC Centera replication impacts the DiskXtender connection. In addition, it provides a step-by-step process for building the connection string and specifying the string in the DiskXtender File System Manager Administrator interface.

Audience

This white paper is intended for EMC customers, partners, and employees responsible for installing and configuring DiskXtender for Windows with EMC Centera. Readers should be well-versed in the following topics:

- DiskXtender and EMC Centera concepts and terminology
- Microsoft Windows network administration

Overview

DiskXtender for Windows is a storage management system that provides flexible data organization and rules-based file migration to multiple hardware and media types. DiskXtender extends the storage capabilities of Windows NT File System (NTFS) volumes by using the DiskXtender File System Manager component to move files from the NTFS volume to other storage media. The storage media is made available through communication with DiskXtender media services. These media services provide access to storage devices, such as an EMC Centera cluster or a tape library.

For example, users on the network may typically save data to a drive on a Microsoft Windows file server. If the drive is an NTFS volume, you can use DiskXtender File System Manager to significantly expand its file storage capabilities. The capacity expansion occurs without changing how a user views the data. File data saved to a drive extended by File System Manager is moved to media (for example, to EMC Centera) without affecting the file listing as seen by the end user.

EMC Centera is a content addressed storage (CAS) system, and is the first magnetic disk-based WORM device. DiskXtender and EMC Centera are tightly integrated to optimize file storage and retrieval performance. Several DiskXtender file migration and retention features are designed specifically to take advantage of the features available with an EMC Centera device.

Understanding the connection process

When you install or upgrade to DiskXtender 6.2 SP1, release 3.2 of the EMC Centera SDK is installed automatically on the DiskXtender server to enable communication between DiskXtender and an EMC Centera cluster.

When you create an EMC Centera media service in DiskXtender, you specify the connection string that the EMC Centera SDK uses to enable DiskXtender to connect to an EMC Centera cluster. The connection string includes information about the EMC Centera access nodes, as well as the Pool Entry Authorization (.pea) file containing the profile information that should be used for the DiskXtender connection. DiskXtender connects with the access nodes by using TCP/IP.

Access nodes

Each EMC Centera cluster contains two or more nodes with the access role. You should specify as many access nodes on the cluster in the connection string as possible.

When the EMC Centera SDK attempts to open a connection to the cluster for DiskXtender, it attempts to connect to the first access node listed in the connection string. If the connection to the first access node is successful, the SDK queries for and internally stores information about all available access nodes in the cluster. Communication between EMC Centera and DiskXtender is then automatically load balanced across this set of nodes.

If the connection to the first access node fails, however, then the SDK attempts to connect to the second access node listed in the connection string. If the connection to the second access node fails, then the SDK attempts to connect to the third access node listed in the connection string, and so on. This process continues until the SDK is able to successfully connect to an access node. Once the connection is made, information about all available access nodes is stored by the SDK, regardless of the number of access nodes listed in the connection string.

This connection process occurs each time the DiskXtender service restarts and each time the media service is set online. As a result, it is important to include as many access nodes on the connection string as possible, to ensure that DiskXtender can connect to EMC Centera even when one or more access nodes are offline.

In addition, you cannot edit access node information in DiskXtender after you create the media service. As a result, you cannot add or remove node information for an existing media service.

Because you cannot edit access node information in DiskXtender, equate each access node IP address with a node name alias before you create the media service. To configure node name aliases, either add the appropriate entries to the DNS server or edit the local HOSTS file on the DiskXtender server.

The connection string for a media service is limited to 128 characters. Keep this limit in mind when defining node name aliases. The aliases should be kept to a reasonable length. Otherwise, you may not be able to specify a sufficient number of access nodes on the connection string to enable reliable connections in the event of a node failure.

You can then specify the aliases instead of the IP addresses when you create the media service. If you later need to alter the access node IP addresses, you only need to change the HOSTS file or DNS entries. You might need to alter the access node IP addresses if you need to replace an access node.

If you add an access node to the EMC Centera cluster after you create the media service, you cannot add the address information to the media service connection string. However, DiskXtender can connect to the access node. This is because the EMC Centera SDK internally stores information about all available nodes after a successful connection is made to one of the current access nodes in the connection string.

If you are using EMC Centera replication, do not specify connection information for the access nodes of the target clusters. “Understanding the impact of replication” on page 6 provides additional details.

Pools and profiles

EMC Centera security is based on pools and application profiles. Each EMC Centera cluster can have multiple virtual pools and multiple access profiles.

- An application pool, or "virtual" pool, is a logical area on an EMC Centera device where applications can store their data. Virtual pools enable you to logically separate data on an EMC Centera device. This is particularly useful if there are multiple applications with different security needs writing data to an EMC Centera device.
- Access profiles provide access to one or more EMC Centera pools. Pools grant capabilities to applications that are accessing EMC Centera by using the profile. For DiskXtender, the profile should have the majority of available capabilities, including Write, Read, Delete, and Query. If you want to allow privileged deletes of retained files, then the profile should also have the Privileged Delete capability.
- A .pea file, generated while creating or updating an access profile, is a clear-text, XML-formatted, non-encrypted file that can be used by system administrators to communicate and distribute authentication credentials to application administrators.

The *EMC Centera Online Help* provides information on pools, access profiles, and .pea files, which should be configured by an EMC Centera technical representative.

When you create an EMC Centera media service in DiskXtender, specify a .pea file. You must specify the .pea file when you create the media service. You cannot edit a media service later to add a .pea file.

There should be a single .pea file for each media service. You cannot specify multiple .pea files for a single media service. If the virtual pool is being replicated to another EMC Centera cluster, the .pea files for the virtual pools on the two clusters are merged. This merged .pea file, which enables access to both pools with a single profile, should be configured by an EMC Centera technical representative.

If you have multiple, separate pools and access profiles—each pair with its own corresponding .pea file—then you can create multiple EMC Centera media services. Multiple media services enable you to further separate data you are writing through DiskXtender.

If you do not specify a .pea file when you create an EMC Centera media service, then DiskXtender uses the Anonymous profile to connect to EMC Centera.

The Anonymous profile is disabled by default in CentraStar® 3.1 and later. If CentraStar 3.1 or later is installed on the EMC Centera cluster, you *must* use an access profile when you create a media service in DiskXtender. If you are using an earlier release of CentraStar with DiskXtender (in other words, the media service already exists and connects through the Anonymous profile) and you upgrade to 3.1 or later, you can continue using the Anonymous profile.

Understanding the impact of replication

The EMC Centera replication feature protects against data corruption and loss by automatically copying data from one EMC Centera cluster to another. As an EMC Centera cluster acquires new content from an application, the replication mechanism ensures that this new content is automatically and transparently transferred across a WAN or LAN to a designated EMC Centera in another location.

Replication is used on an ongoing basis to keep two or more EMC Centera clusters synchronized with new content. In a typical replication setup, the EMC Centera clusters are geographically separate to ensure disaster recovery or to distribute the content for access from another location. For example, a company may replicate to a second EMC Centera cluster to enable recovery from the loss of the primary EMC Centera or to avoid multiple requests for the same content across a WAN connection.

The majority of EMC Centera environments with replication are configured for unidirectional replication. With unidirectional replication, one EMC Centera cluster updates another cluster with its content. For example, if content is written to cluster A, then unidirectional replication transfers the content to cluster B

so that it is located on both clusters. However, if content is written directly to cluster B, the content is not transferred to cluster A. As a result, there may be additional content on cluster B that does not exist on cluster A.

The *EMC Centera Online Help* provides additional information on replication. Replication should be configured by an EMC Centera technical representative.

Replication and DiskXtender files

The replication process itself is transparent to DiskXtender. In other words, after DiskXtender migrates files to EMC Centera, the files are replicated from the source cluster to the target cluster without any DiskXtender involvement.

Replication failover

In the past, DiskXtender supported three types of DiskXtender file behavior when the source EMC Centera cluster failed:

- Read-only
- Read/write
- No failover

Due to concerns about possible data loss, however, only read-only failover and no failover are now supported for DiskXtender 6.2 and later.

Read-only failover

When the source EMC Centera cluster fails, DiskXtender automatically attempts to set the failed media service online again. If DiskXtender is unable to set the media service online, the EMC Centera SDK provides a read-only connection to the target (replica) cluster after a brief pause. With read-only failover, files can be read from the target cluster, but additional files cannot be written to the target cluster.

Once the source cluster is set online again, DiskXtender automatically resumes normal read/write activity with the source cluster.

This scenario enables users to fetch files that have already been migrated and purged, even when the source cluster fails. However, new file migration activity (from new files that qualify for migration, as well as edits to and deletes of files that have already been migrated) must wait until the source cluster comes back online.

To enable this behavior, include the connection information for *only* the source cluster in the media service connection string, and select the Enable Read-Only Replica Failover checkbox when configuring the media service.

If the DiskXtender service restarts or if the media service is set offline while DiskXtender is connected to the target, then the connection to the target fails, and the media service is set offline. Purged files cannot be fetched until the source cluster is set online again. This is because DiskXtender passes the connection string to the EMC Centera SDK, and the connection string includes only addresses from the source cluster. If the source cluster is offline, then no connection can be made—not even to the target cluster. The address information for the target cluster is passed to DiskXtender only after a successful connection to the source cluster.

If the source cluster is offline for an extended period of time, contact EMC Customer Service. A Customer Support Representative can configure the environment to establish a read/write connection to the target cluster, and to ensure that files are replicated back to the source cluster once it is set online.

Read/write failover

Read/write failover, where DiskXtender both reads files from and writes files to the target cluster if the source cluster fails, is no longer supported due to concerns about possible data loss. If a failover occurs and files are written to the target cluster, then the files are not replicated back to the source cluster in a unidirectional replication environment, which includes the majority of replication environments. As a result, if the target cluster is removed from the environment, then the files that were written only to the target cluster are lost.

The *EMC DiskXtender Release 6.2 Microsoft Windows Version File System Manager Administrator's Guide* and *EMC DiskXtender Release 6.2 Microsoft Windows Version Best Practices Guide* incorrectly state that read/write failover is supported.

No failover

You can configure DiskXtender so that no failover to the target cluster occurs. In other words, files can neither be read from nor written to the target cluster. Both file fetches and file migration are disabled. Users can only access migrated files that were fetched prior to the failover.

To enable this behavior, include the connection information for *only* the source cluster in the media service connection string, and clear the Enable Read-Only Replica Failover checkbox when configuring the media service.

DiskXtender continues to attempt to set the media service back online until it is successful. Until then, the media service and its media remain offline.

Building the connection string

When you create an EMC Centera media service in DiskXtender, you specify the connection string that the EMC Centera SDK uses to enable DiskXtender to connect to an EMC Centera cluster. The connection string includes information about the EMC Centera access nodes, as well as the access profile that should be used for the DiskXtender connection.

To build the connection string for a single EMC Centera media service:

1. Collect the IP addresses for all access nodes on the primary EMC Centera cluster to which DiskXtender should write files.

If EMC Centera replication is configured in the environment, do *not* include the IP addresses for the access nodes on the target cluster.

2. Configure node name aliases for the access nodes. To configure node name aliases, use one of the following methods.
 - (Recommended) Add the appropriate entries to a common DNS server.
 - Edit the local HOSTS file on the DiskXtender server to include the IP addresses for the EMC Centera access nodes and the alias you want to use for each IP address.

The connection string for a media service is limited to 128 characters. Keep this limit in mind when defining node name aliases. The aliases should be kept to a reasonable length. Otherwise, you may not be able to specify a sufficient number of access nodes on the connection string to enable reliable connections in the event of a node failure.

3. On a single line in a text or word processor file, list the nodename aliases for the access nodes on the primary EMC Centera cluster, separating the aliases with a comma. For example, if there are four access nodes on the primary cluster and the nodename aliases for the nodes are defined as Node1, Node2, Node3, and Node4, then type:

Node1,Node2,Node3,Node4

4. Work with an EMC Centera technical representative to create one or more access profiles to provide access to the virtual pools.

In most cases, the profile should have the majority of available capabilities, including Write, Read, Delete, and Query. If you want to allow privileged deletes of retained files, then the profile should also have the Privileged Delete right.

This step should result in a .pea file. The EMC Centera technical representative should provide you the .pea file. Copy the file to a location on the DiskXtender server.

5. In the text or word processor file with the list of nodename aliases, add a question mark (?) after the nodename aliases, and then include the path on the DiskXtender server to the .pea file for the access profile that DiskXtender should use to connect to EMC Centera.

For example, if the .pea file is located on the system drive at C:\Centera\DXProfile.pea, then the connection string would appear as:

```
Node1,Node2,Node3,Node4?C:\Centera\DXProfile.pea
```

6. Ensure that the string with the nodename aliases and the path to the .pea file is fewer than 128 characters.
7. Type the connection string in the **Pool Address** text box on the **Centera Information** page when you create the media service, as illustrated in the following figure:

Centera Information

Enter the pool addresses for the Centera media service. Use commas to separate multiple addresses.

The addresses should be either a node name alias (equated to the front-end node TCP/IP addresses in either the local HOSTS file or on the DNS server) or a TCP/IP address.

Pool Address:

Node1,Node2?C:\Centera\DXProfile.pea

File Delete Audit String:

Embedded Blob Threshold in k (0-100): 0

Collision Avoidance Client Side ID (hash) Calculation

Enable Read-Only Replica Failover

< Back Next > Cancel Help

8. If you are using EMC Centera replication, specify how DiskXtender handles the failover when the primary cluster is unavailable.
 - (Recommended) To provide read-only access to files on the replica cluster, select the **Enable Read-Only Replica Failover** checkbox.
 - To prevent read access to files on the target cluster, clear the **Enable Read-Only Replica Failover** checkbox.

For details on the remaining options on the **Centera Information** page, refer to the File System Manager Administrator online help.

Connection configuration quick reference

The following table provides a quick reference to the required configuration of a connection between DiskXtender and EMC Centera, depending on the intended replication failover behavior:

| Failover behavior | Primary/source cluster access nodes on the connection string | Target/replica cluster access nodes on the connection string | Enable Read-Only Replica Failover checkbox selected |
|-------------------------|--|--|---|
| No replication | ✓ | | |
| No failover | ✓ | | |
| Read-only (recommended) | ✓ | | ✓ |

Conclusion

When building the connection string for DiskXtender and EMC Centera, include only the access nodes for the primary EMC Centera cluster, regardless of whether you are using EMC Centera replication. In addition, to ensure a read-only connection to an EMC Centera replica after a failover, select the Enable Read-Only Replica Failover checkbox for the media service.

References

- *EMC DiskXtender Release 6.2 and 6.2 SP1 Microsoft Windows Version Release Notes*
- *EMC DiskXtender File System Manager Release 6.2 Microsoft Windows Version Administrator's Guide*
- *EMC DiskXtender Release 6.2 Microsoft Windows Version Best Practices Guide*
- *EMC Centera Online Help*