

EMC CLARiiON Storage Solutions: Microsoft Data Protection Manager 2007

Best Practices Planning

Abstract

This white paper discusses the initial replication and recovery practices for Microsoft Data Protection Manager 2007 on EMC[®] CLARiiON[®] storage systems.

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 H4260

EMC CLARiiON Storage Solutions:
Microsoft Data Protection Manager 2007
Best Practices Planning

Table of Contents

Executive summary	4
Introduction	4
Audience	4
Overview of Data Protection Manager 2007	4
Overview of CLARiiON storage systems	5
DPM initial replication	6
Using CLARiiON SnapView clones	6
Introduction to clones	6
CLARiiON storage system considerations	8
SnapView clone configuration guidelines.....	8
Creating a DPM initial replica with clones	9
SAN-based recovery using CLARiiON snapshots	16
Using CLARiiON SnapView snapshots	17
Introduction to snapshots	17
Reserved LUN pool configuration guidelines	19
SAN-based recovery using snapshots	19
Requirements/prerequisites	26
DPM server	26
Protected server.....	27
SAN (CLARiiON storage array requirements)	28
Conclusion	28
Appendix: Scripts	29

Executive summary

This white paper provides a set of recommended configuration and design details, as well as initial best practices, for using EMC® CLARiiON® storage systems with Microsoft Data Protection Manager 2007. DPM is a disk-based data protection software application that provides alternatives to tape backup solutions.

By following the best practices outlined in this paper customers can reduce time and eliminate the extra load on the local network during initial replication and recovery.

Introduction

This white paper gives an overview of Microsoft Data Protection Manager 2007 and EMC CLARiiON storage systems. It is intended to provide guidance on implementing DPM-based backup and recovery for applications utilizing storage from EMC CLARiiON storage systems. Readers are advised to consult the respective product documentation for more complete coverage of the features mentioned in this white paper.

Audience

This white paper is intended for IT administrators and system engineers who are looking for a solution to effectively manage backup and recovery of applications like Microsoft Exchange 2007, SQL Server 2005, SharePoint servers, and Windows volumes in their organization. The procedure depicted in this paper explains the performance and time-saving benefits gained by using Microsoft DPM 2007 with EMC CLARiiON arrays.

It is assumed that the reader has a general knowledge of SAN and CLARiiON features and terminologies.

Overview of Data Protection Manager 2007

Data Protection Manager (DPM) is a key member of the Microsoft System Center family of management products, designed to help IT professionals manage their Windows environment. DPM is the new standard for Windows backup and recovery—delivering continuous data protection for Microsoft application and file servers using seamlessly integrated disk and tape media. DPM enables rapid and reliable recovery through advanced technology for enterprises of all sizes.

Today's business climate is more challenging than ever and businesses are under constant pressure to lower costs while improving overall operational efficiency. In short, businesses are being asked to “do more for less.” One way that enterprises of all sizes can reduce costs and improve business agility is by changing the way data protection is managed. DPM provides the following benefits:

Continuous data protection – DPM captures data changes as they occur in real time and synchronizes every 15 minutes to ensure data and user productivity are protected—unlike other legacy methods that backup data only once a day—so that business-critical data can be protected at all times.

Lossless restores for applications – By seamlessly integrating a point-in-time database restore with the existing application logs, DPM delivers lossless recovery of Microsoft Exchange, SQL, and SharePoint servers without the need for constant replication or synchronization.

Superior application integration for Exchange Server, SQL Server, and SharePoint – DPM offers integrated support for Exchange Servers 2003 and 2007 with advanced cluster configurations, shorter SQL Server backup windows without the need for compression, and integrated restores for SharePoint. DPM also provides one-touch application restores with zero data loss.

Rapid recovery – Backing up data to disk provides the fastest way to recover data that's been lost due to user error or software and hardware corruption. With DPM, recovering information is as simple as browsing a share and copying directly from the DPM platform to the production server within seconds. By

restoring data from disk, DPM enables customers to recover data in minutes, versus the hours it takes to recover from tape.

Reliable recovery – DPM leverages disk-based backup to provide the highest level of reliability versus traditional—and often unpredictable—tape backup methods. All of the various failure points for tape backup, including corrupt indexes, broken media, misplaced cartridges, and human error, are all avoided by relying on disk storage as the primary restoration medium, while still leveraging tape for long-term archival storage.

Seamless disk and tape integration – DPM transparently leverages both disk and tape mediums to enable fast, multiple points-in-time-per-day restores from disk, while ensuring long-term retention and offsite portability with tape.

Unified protection policies across data types – DPM allows protection to be configured across heterogeneous applications and file-sharing platforms with a single policy. This allows you to manage logical groupings of data from a single UI—delivering Exchange, SQL, SharePoint and file data mixed, within a single policy, to any combination of disk and tape protection.

SLA-driven backup process – Protection policies are based on intent and SLAs, creating a layer of abstraction that insulates the user from the often confusing process of scheduling individual backup jobs in order to meet an overall SLA.

Block filter – With an efficient disk infrastructure and reduced network traffic, DPM’s volume filter changes how backups are achieved and maintained:

- The volume of full backups is reduced by as much as 90 percent, saving disk space and reducing backup time from hours to minutes.
- Express full backups and an enhanced network throttling mechanism allow for more granular management of bandwidth.

For additional information on DPM, go to the Microsoft website:

<http://www.microsoft.com/systemcenter/dpm/evaluation/default.aspx>

Overview of CLARiiON storage systems



EMC CLARiiON is a modular storage system that provides a high degree of performance and flexibility. The current generation of CLARiiON storage systems includes the CX300, CX3-10, CX3-20, CX3-40, and CX3-80 models,

providing up to 60, 60, 120, 240, and 480 disk drives, respectively. All models are available with Fibre Channel front-end connections. The CX3-10, CX3-20, and CX3-40 are also available as combo units with FC and iSCSI connections. All storage systems run on FLARE®, an operating environment designed and implemented exclusively for efficient, highly available, and highly reliable block storage access. They are managed via the Navisphere® user interface, which is Web-based and runs directly on the storage system. A single management interface can manage any number of CLARiiON storage systems. All these models contain dual storage processors and either storage processor can be accessed via integrated Ethernet connections to manage the storage system. Given the dual-redundant nature of CLARiiON storage processors, this provides an elegant, highly available approach to storage management.

For more in-depth information visit the [EMC CLARiiON family page](#) on EMC.com or [Powerlink®](#), EMC’s password-protected extranet for customers and partners.

DPM initial replication

Data Protection Manager (DPM) helps you manage the process of protecting and recovering data on the file and application servers in your network. This topic describes the high-level steps you need to perform to successfully protect and recover data in the DPM environment.

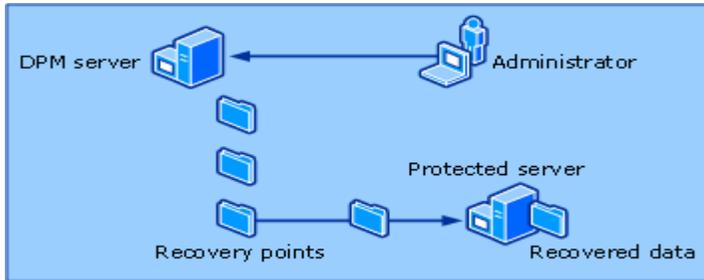


Figure 1. How to protect data

Protecting data, as illustrated in Figure 1, is a high-level process that involves the following steps:

1. Select data sources on a computer that you want to protect whether it is an application server or a file server or a workstation.
2. To start protecting data, DPM creates a full copy (referred to as a replica) of the selected data sources on the DPM server.
3. To continue protecting data, DPM synchronizes each replica with the data sources on a recurring schedule. When a replica is updated, it replaces the previous replica.
4. To support data recovery, DPM creates point-in-time copies (referred to as recovery points) of the replica on a recurring schedule. DPM maintains up to 64 recovery points for each volume (files) replica and for application data sources, DPM can store up to 448 express full backups and up to 96 incremental backups for each express full backup.

Using CLARiiON SnapView clones

Using CLARiiON SnapView™ clones is the recommended method to perform initial replication of protected data sources to the DPM server.

SnapView clones provide users the ability to create fully populated binary copies of LUNs within a single storage system. If more than one storage system is involved, then other replication technologies like SAN Copy™ and MirrorView™ could be used. Such a discussion is outside the scope of this white paper. Once populated, clones can be fractured from the source and presented to a secondary server to provide exact replicas of the source data.

Introduction to clones

This section discusses the features of clones, such as fracture, and the process for configuring clone groups. As of release 24, the Configure SnapView Clones Wizard walks the user through the initial configuration process. The underlying principles of setting up a clone group and allocating clone private LUNs are maintained, but many of the steps involved in these tasks are now automated for the user during initial setup. After initial configuration, clones are manipulated through Navisphere Manager or NaviCLI. For more information log on to Powerlink.

Basic clone features

The process of copying data from the source to the clone is called “synchronization,” and the other direction, from the clone to the source is called “reverse synchronization.” After the clone has been initially synchronized, it is generally fractured and presented to a secondary server. Upon fracturing the clone, the

fracture log keeps track of the regions (extents) that have been modified on either the clone or its source after the clone is fractured. The fracture log is a bitmap contained on disk – a region referred to as the clone private LUN – so this record is persistent even in the event of a power outage, ensuring protection and high availability. Subsequent synchronizations or reverse synchronizations are incremental in nature in that they copy only the extents that have changed on the source and/or the clone.

During both synchronization and reverse synchronization, server I/Os (read and write) can continue to the source. The clone, however, is not accessible for secondary server I/Os during either synchronizations or reverse synchronizations; the user must ensure that all server access to the clone is stopped (this includes ensuring that all cached data on the server is flushed to the clone) prior to initiating a synchronization or a reverse synchronization. This can be done by making sure no application is performing any reads or writes on the clone LUN. Failure to remove access will likely cause the server to report write errors or, in some cases, panic (especially if the clone was used as a server boot device).

Users can select the rate at which the data is synchronized to (or reverse-synchronized from) the clone. This way, users can prioritize between storage system CPU cycles directed to synchronization operations versus other I/O processing tasks.

Clone synchronization

Synchronization is the process of copying data from the source LUN to the clone. Upon creating the association of a clone with a particular source this translates to a full synchronization - all extents on the source LUN are copied to the clone to provide a completely redundant replica. Subsequent synchronizations involve only a copy of any data that has changed on the source since the previous synchronization - overwriting any writes that have occurred directly to the clone from any secondary server that had been accessing it while the clone was fractured. It is essentially an update for the clone. Once fully synchronized, the clone is ready to be fractured again to maintain the relevant point-in-time reference.

Clone fracture

A clone must be fractured from the source LUN to achieve the point-in-time representation of the data that the user is trying to obtain, and to make the clone available for secondary server access. Figure 2 shows the menu options to fracture, synchronize, and reverse-synchronize the clone, as well as the options to remove the clone from the clone group and view clone properties. All functions are accessed within the Navisphere tree by right-clicking the icon of the clone on which the user would like to perform the action.

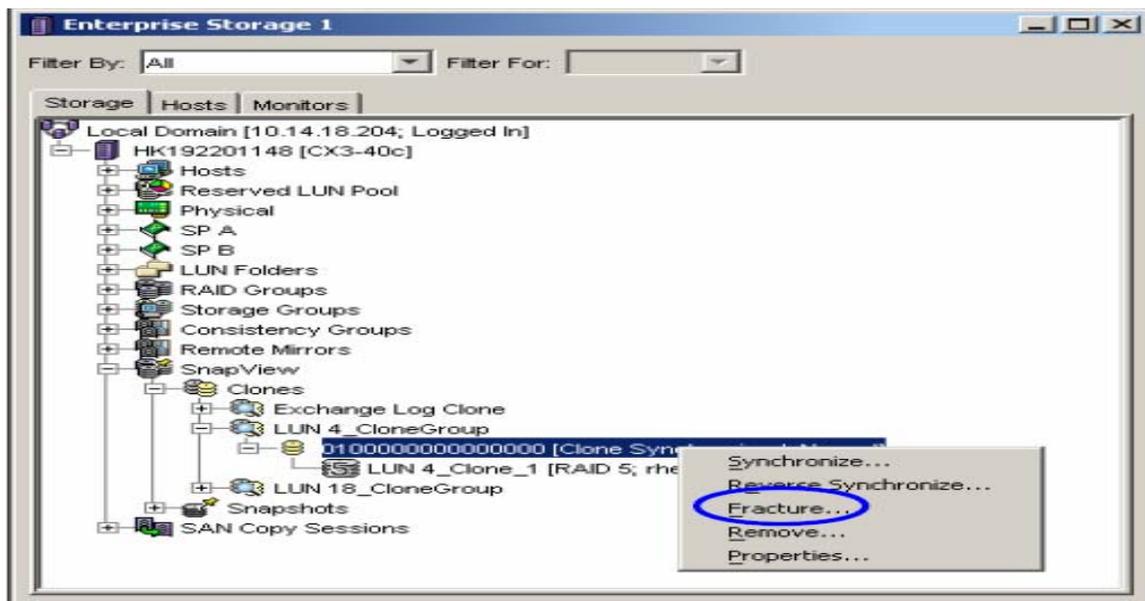


Figure 2. Menu options for the clone

Once fractured, clones can be presented to a secondary server and are available for read and write I/O. Usually, these writes are intended to be discarded, such as in a testing scenario. The next synchronization discards the writes, and also copies any data written to the source LUN while the clone is fractured. If the

user wants to copy the writes made to the clone back to the source LUN, a reverse synchronization would be used.

CLARiiON storage system considerations

Best performance and high availability of storage can be achieved by choosing the best layout for RAID groups and LUNs, fine-tuning the storage system features like write and read cache, file system alignment, choosing the right disk type based on workload type, providing redundant connections to the host, and so on.

In this section we will be discussing only a subset of those recommendations that will have direct impact on creating initial replica and performing SAN-based recovery in DPM 2007. For more detailed information on CLARiiON configuration best practices please refer to Powerlink.

Vault drives: Drives 0, 1, 2, 3 and 4 in the first DAE in a CLARiiON contain the vault. These are 10k or 15k rpm FC drives and can be of any capacity or back-end speed. The vault drives are special because they contain configuration information, the write cache dump area, and other information that is used by the CLARiiON. The vault drives are also the storage area that protects the write cache in case of a power failure. It is preferable to store less performance critical data in these drives.

Hot spares: Use at least one hot spare for every 30 drives.

High availability: Configure multiple paths between host servers and storage arrays using multiple HBAs and switches. Use at least two NICs or HBAs in each host. Install EMC PowerPath[®] software on all hosts, which provides failover and load balancing across available active paths.

Read-write cache: Make sure read and write cache are ON for all LUNs

SnapView clone configuration guidelines

The performance of the initial replication process relies on how effectively the cloning is done. The time taken for cloning the production LUN and the performance impact on the production LUN during cloning have a direct impact on the application response times and DPM environment.

This section discusses the best practice recommendation for configuring SnapView clone target LUNs (DPM replica volume). The same configuration guidelines apply for DPM recovery point volume as well.

Read-write cache: As a general rule, assign 20 percent of available cache to read cache and the remaining to write cache. Also be aware of the write cache limitations of a particular CLARiiON array model. This high write cache will be significantly useful while using less-expensive SATA drives as the clone target (Initial replica in our case).

LUN ownership: Clone LUNs must be assigned to the same SP as their source. (It is legal to configure clones on the peer SP, but the clone would be automatically trespassed during synchronizations.) Ensure that both the current SP owner and preferred owner are the same for both the source and target LUNs.

SATA drives: SATA drives are not the drives of choice for clone targets. However in a system that is not being stressed, the use of SATA drives as targets compared to FC drives will have a modest effect on system performance. But using SATA for a one-time cloning purpose like using it as a DPM replica has its own advantages. In such cases make sure to clone the data at a time when there is the least load on the production LUN. After the initial replication, DPM will be moving data from the production LUN to the clone replica (SATA LUN) in regular sequential intervals. Since SATA is good at handling this type of sequential loads, we won't see much of performance degradation. Also, since SATA drives are inexpensive compared to FC drives, there is the price advantage, too.

RAID types: For clone targets, use RAID 5 on a five- to nine-disk configuration for SATA disks. Increasing the number of spindles has a direct effect on cloning performance since RAID 5 of this size is relatively efficient in parity usage and also rebuilds in a reasonable time. However, blindly increasing the disk count to huge numbers will result in longer rebuild times.

MetaLUNS: EMC CLARiiON metaLUNs are a collection of individual LUNs that act in concert with and are presented to a host or application as a single storage entity. MetaLUNs allow users to expand existing volumes on the fly without disrupting host I/O to the volume. They are created by taking existing LUNs and combining them together.

The clone target (replica) can be a simple LUN or a metaLUN while cloning, fracturing, and assigning it to DPM as the initial replica. When there is a need to expand the replica volume on a DPM server, expand it by adding more LUNs to it. After a metaLUN is expanded in the storage system, the additional capacity shows up as unpartitioned space in Windows disk management. Using the Diskpart.exe utility can extend the file system to use the extra space available in the expanded metaLUN. This nondestructive expansion can be performed on the fly even when DPM is servicing the clients.

For more information about metaLUNs in CLARiiON storage arrays, refer to the white paper *EMC CLARiiON MetaLUNs: Concepts, Operations, and Management*.

For more information about clones in CLARiiON storage arrays, refer to the white paper *EMC CLARiiON SnapView Clones - A Detailed Review*.

Creating a DPM initial replica with clones

In Data Protection Manager (DPM), a replica is a complete copy of the protected data on a single volume, database, or storage group. Each member of a protection group is associated with a replica on the DPM server. A replica contains all the properties of the volume, including security settings, and sharing. The replica and the recovery point data are stored in separate volumes.

Before DPM can start protecting the data sources in a protection group, a replica of the data must be created. After a replica is created for each data source, changes to the protected data are transferred to DPM incrementally through synchronization, according to a set schedule.

To create a replica on the DPM server, you can either have DPM copy the data from the file or application server over the network or you can manually create a replica from a tape backup or other removable storage medium. Replicating the data over the network requires no intervention, but it can take several hours, depending on network bandwidth and the data size.

We recommend using a SnapView clone target to create the initial replica for DPM, reduce the time required for replica creation, and also eliminate the load on the local network. In the following section we'll demonstrate how the SnapView clones feature of the CLARiiON storage system is used to perform replica creation.

In the following sections, this paper explains the setup, requirements, and steps for creating the replica manually using CLARiiON clones. These instructions assume that the "custom volume" feature of DPM storage is used.

Replica volume (using CLARiiON clones)

A clone of the production data can be assigned as the replica on the DPM server, eliminating network flooding and the time required for building the replica over the network. Another important advantage of using clones is that the entire process of cloning and manual replica creation can be done on a live production volume without having to stop any I/O to the volume. This applies to DPM replica creation context because any potential inconsistencies between application state and physical storage will be nullified through the mandatory DPM consistency check. It is recommended to perform the cloning operation at a time when the production volume has the least load (usage), so as to reduce the impact on the application performance.

Recovery point volume

DPM creates point-in-time copies (referred to as recovery points) of the replica on a recurring schedule and maintains up to 64 recovery points for each volume (or file) replica; for application data sources, DPM can store up to 448 express full backups and up to 96 incremental backups for each express full backup. The point-in-time copies are stored in a volume called a recovery point volume. This data will be used during the recovery process.

While creating CLARiiON LUNs for the recovery point volume, in addition to the guidelines laid out in the section “[CLARiiON storage system considerations](#),” make sure the partition alignment and allocation unit size are set as per the following recommendations. For online volume expansion support refer to metaLUNs in the section “[SnapView clone configuration guidelines](#).”

Alignment: Windows Server 2003/Windows XP and earlier versions of the Windows operating systems will always create the first partition on that disk starting at the 64th sector, thus misaligning it with the underlying RAID stripe. This will cause disk crossings for a percentage of small I/O, resulting in slightly lower performance. The best performance can be achieved by aligning the partitions. For more information please refer to the white paper *Using diskpar and diskpart to Align Partitions on Windows Basic and Dynamic Disks*.

NTFS Allocation Unit Size: Choosing the right disk allocation unit is crucial because smaller read/writes on a large allocation unit can waste storage space while larger disk read/writes on a smaller allocation unit causes disk fragmentation, resulting in slower file access on subsequent read/writes. The default allocation unit size for Windows disks is 4 KB. Choose a 64 KB allocation unit size for DPM volumes. A 64 KB allocation size matches the CLARiiON internal storage structure size and alignment for optimal storage performance.

Replica creation steps

The example in this section demonstrates protecting an Exchange 2007 server storage group using DPM 2007. The storage group that we’re going to protect is “ESG”, which has its database file stored in volume H: and its log files stored in volume I:. These two volumes reside in separate CLARiiON LUNs. We’ve used SATA 7.2k rpm 500 GB disks for the clone target (replica) and recovery point volumes. All the LUNs are created with a RAID 5 configuration. Prior to following the procedure below, it is assumed that the production LUN (where the Exchange storage group resides) is cloned, fractured, and assigned to the DPM server as mentioned in the section “[Introduction to clones](#).” Upon fracturing the clone, remove the clone target LUN from the clone group.

Note: Remember that once a LUN is removed from a clone group, it loses its association with the clone group and cannot be reverse synchronized. However since DPM won’t be using this feature, it is safe to do it in this scenario.

Another assumption is that the appropriate DPM protection agent is installed on the production server.

From the Disk Management snap-in on the DPM server, mount the clone LUN (ESG_DB) to the folder “C:\Program Files\Microsoft DPM\DPM\Volumes\Replica\protected server’s fully qualified name\Microsoft Exchange Writer\Protected Exchange SG name”.

Create a LUN of size 100 GB (the actual size differs according to the size of protected data and the recovery point setting) and assign it to the DPM server. Create a partition and format with NTFS. Name it ESG_RecoveryPt and mount to the folder, “C:\Program Files\Microsoft DPM\DPM\Volumes\DiffArea\protected server’s fully qualified name\Microsoft Exchange Writer\ Protected Exchange SG name”. This will be used as the recovery point volume. Refer to the section “[Recovery point volume](#)” for the best practices while creating the recovery point LUN. The replica volume will be the clone LUN itself (Figure 3).

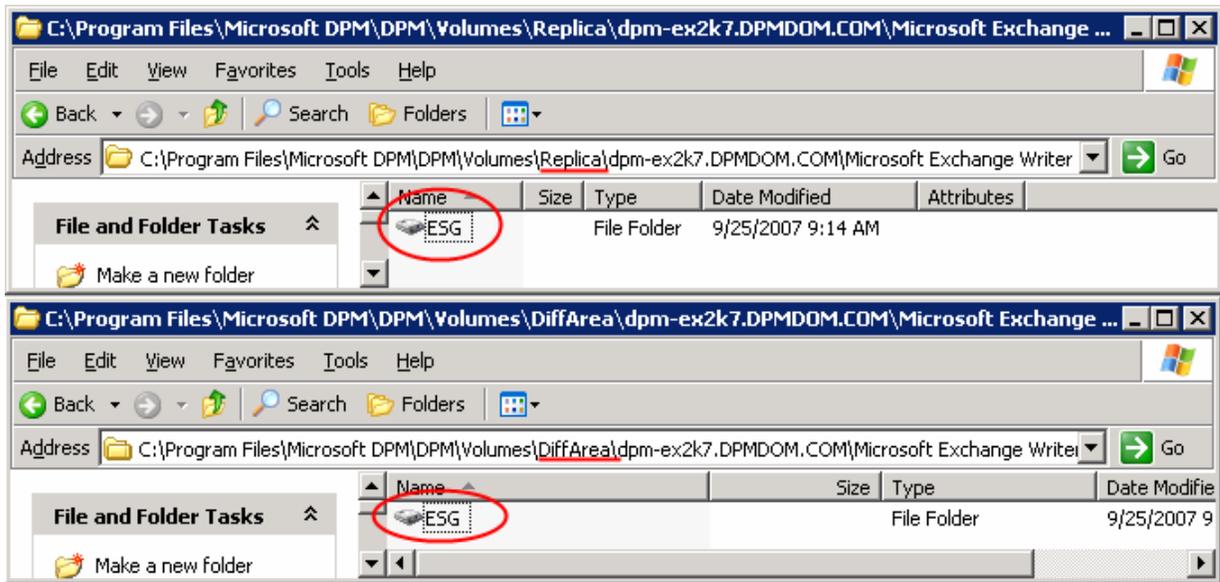


Figure 3. Creating the recovery point volume

The structure of the clone data (ESG_DB) is shown in Figure 4. This data will become the replica for the DPM protection group we're about to create.

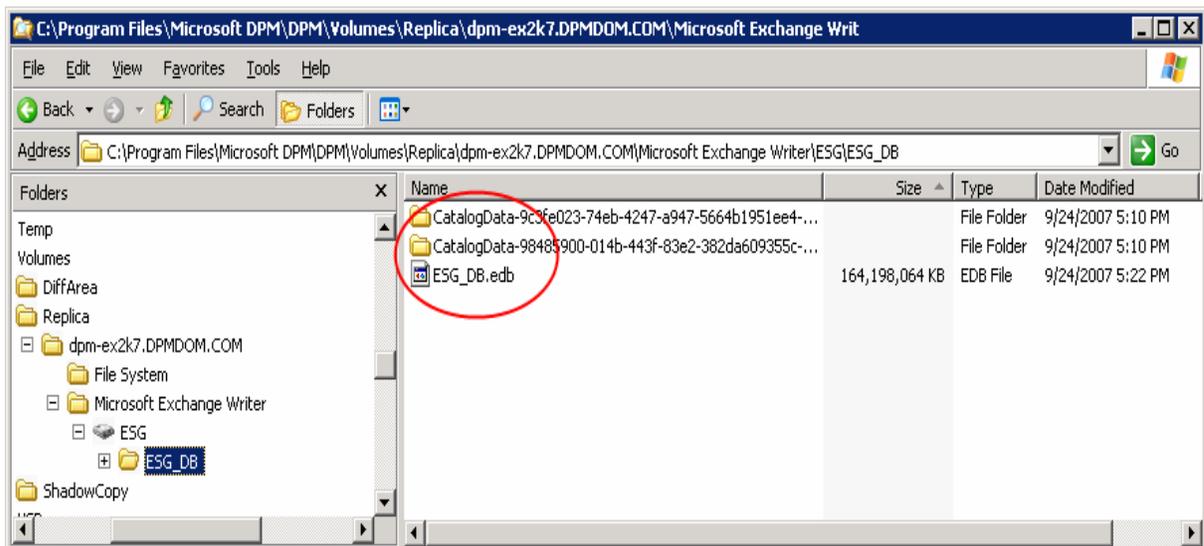


Figure 4. Clone data structure

Create the protection group (ESG_PG) by opening the Create New Protection Group Wizard in the DPM Administrator Console and selecting the data source to be protected (Figure 5).

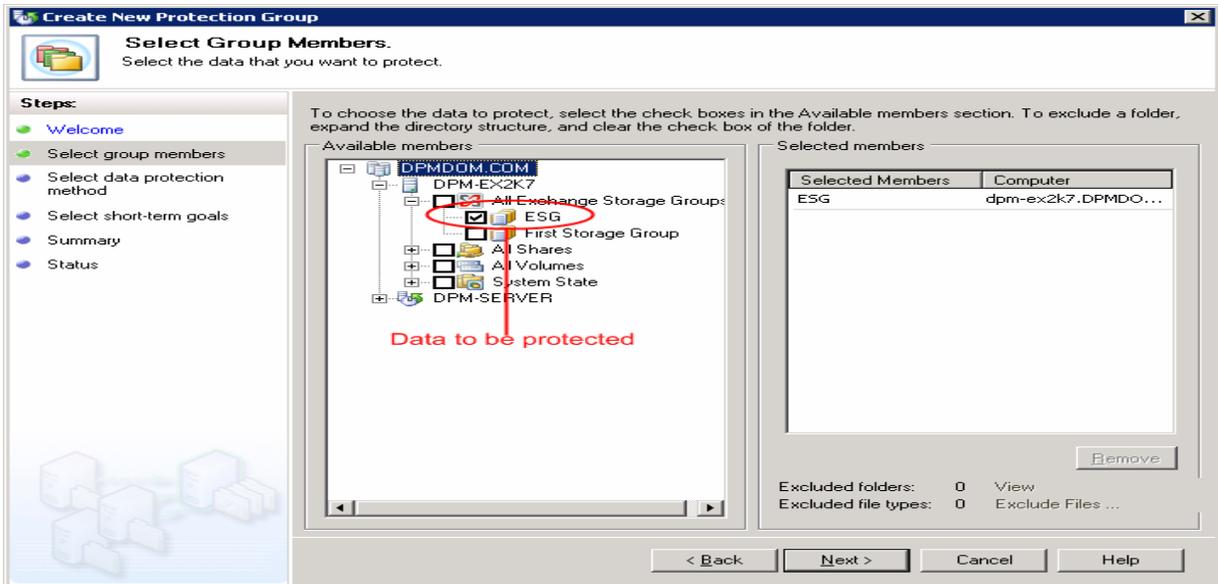


Figure 5. Create the Protection Group

The Review Disk Allocation page shows the disk space allocation for the protected members. In order to use the CLARiiON clone LUN as the replica volume, click **Modify**, select **Custom Volume** as the storage type, and choose ESG_DB as the replica volume and ESG_RecoveryPt as the recovery point volume. Also make sure to select the **Do not Format** option from the pull-down menu (Figure 6, 7, and 8).

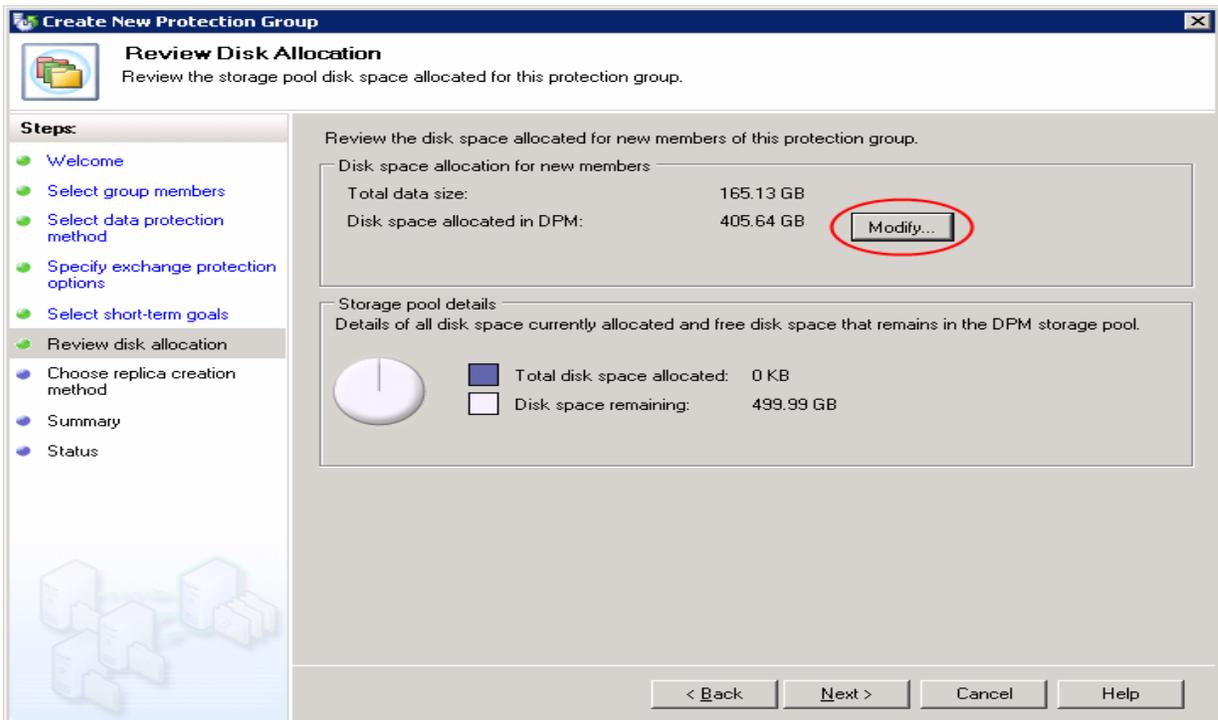


Figure 6. Modify the disk space allocation

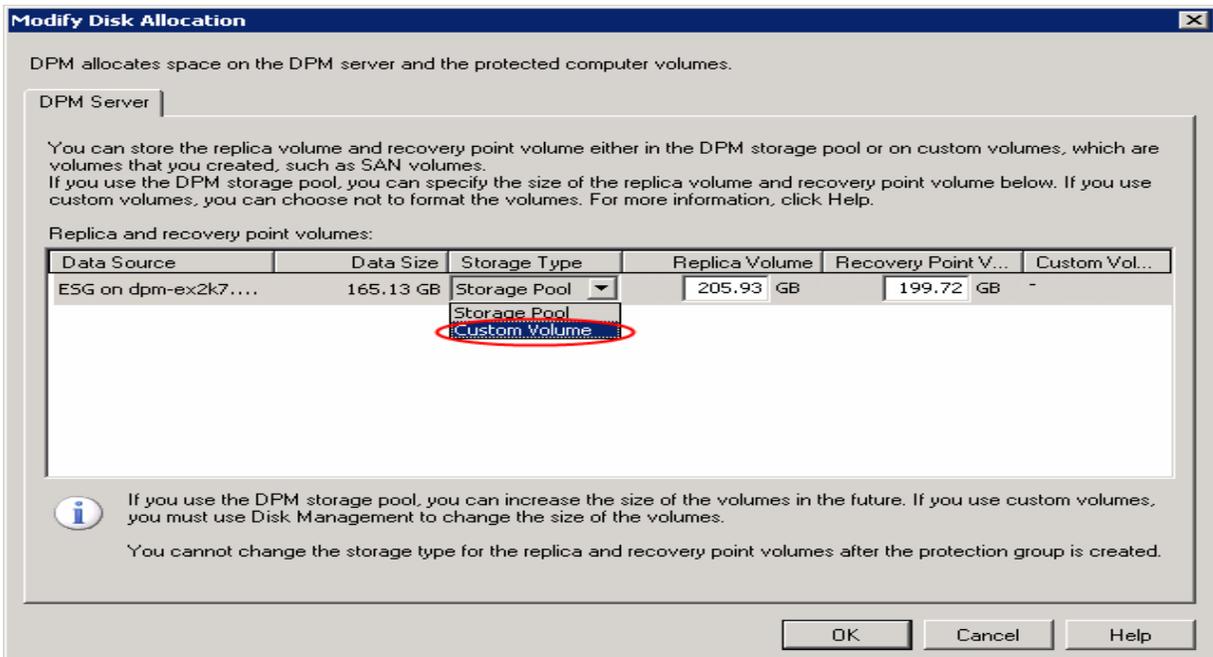


Figure 7. Choose the storage type

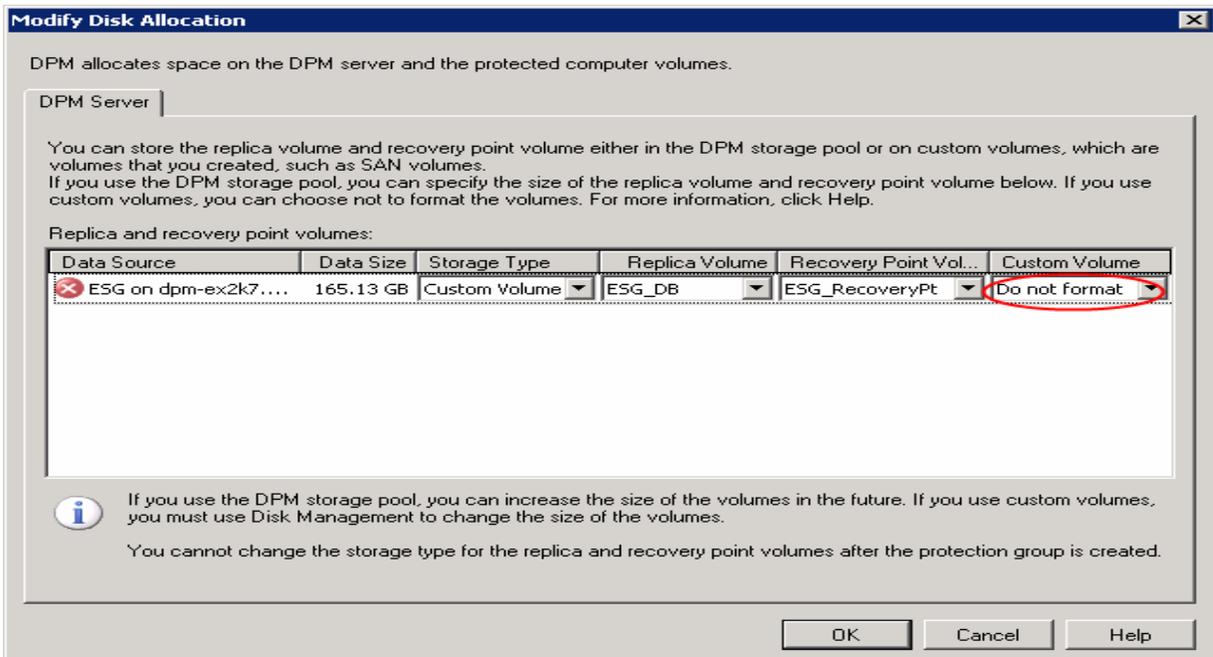


Figure 8. Choose Do not format

Because we assigned the CLARiiON clone LUN as the replica volume, in the Choose Replica Creation Method screen, only the **Manually** option will be selected. Click **Next** (Figure 9).

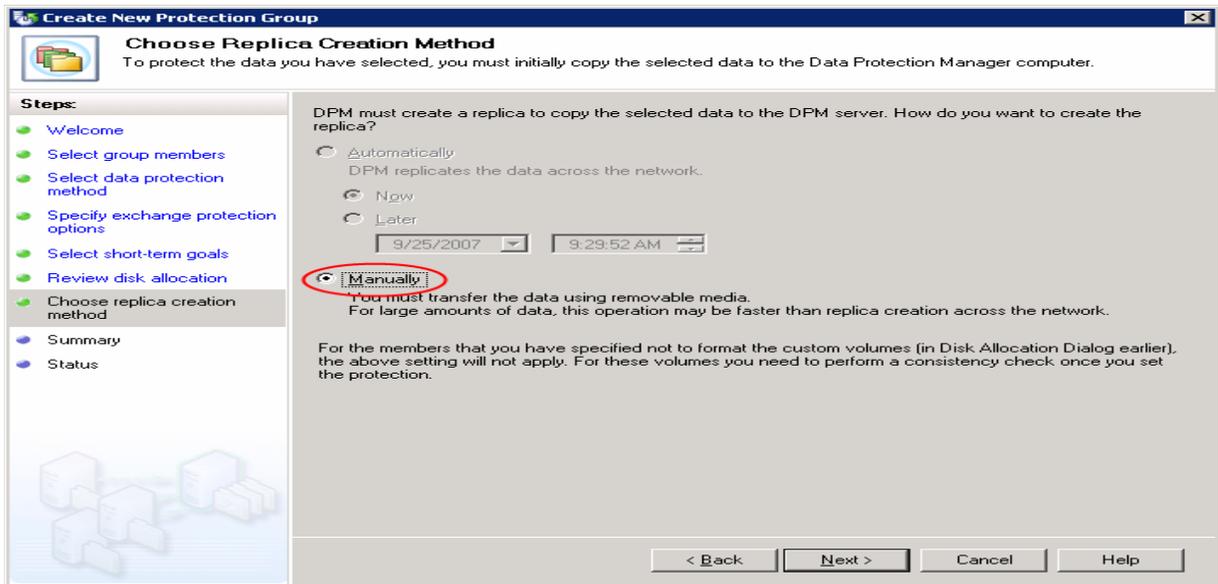


Figure 9. Create the replica manually

The Summary page opens. Review the settings and then click **Create Group** to create the protection group. Now the Status page opens and shows the status of the protection group creation task. When the task is complete click **Close** to close the wizard and notice how DPM created the directory structure (Figure 10).

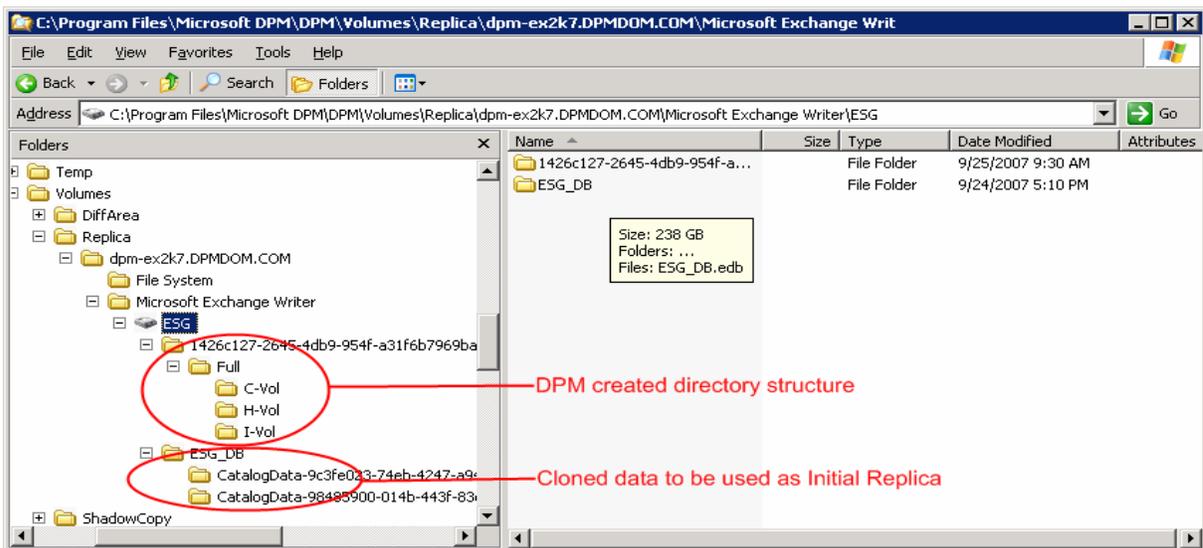


Figure 10. A directory structure is created

Move the data on the replica volume into the directory named Full\H-Vol (Figure 11 and Figure 12). This step is necessary because DPM stores the data to be protected into this directory. Also H-Vol corresponds to the actual drive letter H:, where the Exchange storage group data is located on the production server. So if the Exchange storage group data is located on the D: drive on the production server, then DPM would have created FULL\D-Vol and we would move the clone data into this folder.

Note: If we're protecting a volume rather than application data, then DPM will create just the "Full" folder and in such cases we must move the clone data directly into the Full folder and then perform the consistency check.

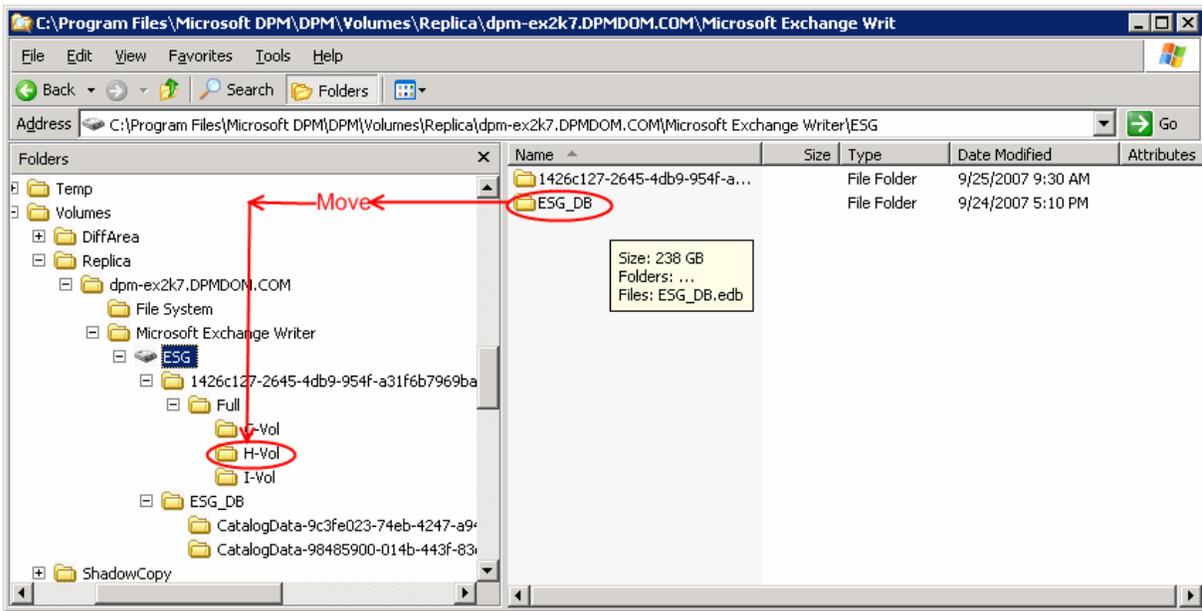


Figure 11. Moving data

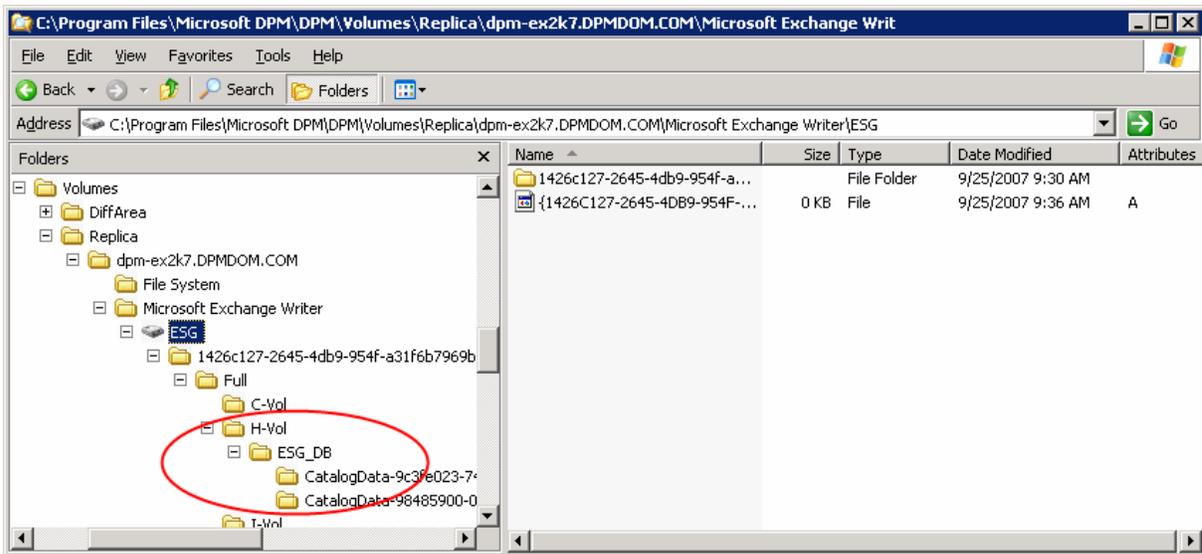


Figure 12. Moved data

In the Protection task area, run a consistency check on ESG_PG (Figure 13), which will take a lot less time and network traffic to synchronize since the replica created for this protection group member is a recent clone of the production data (ESG_DB) with very small differences, if any at all. However, the transaction logs and system logs corresponding to the storage group (ESG) that we're protecting will be copied to the folders FULL\I-Vol and FULL\C-Vol (Figure 10 on page 14) over the public LAN by DPM during the initial consistency check.

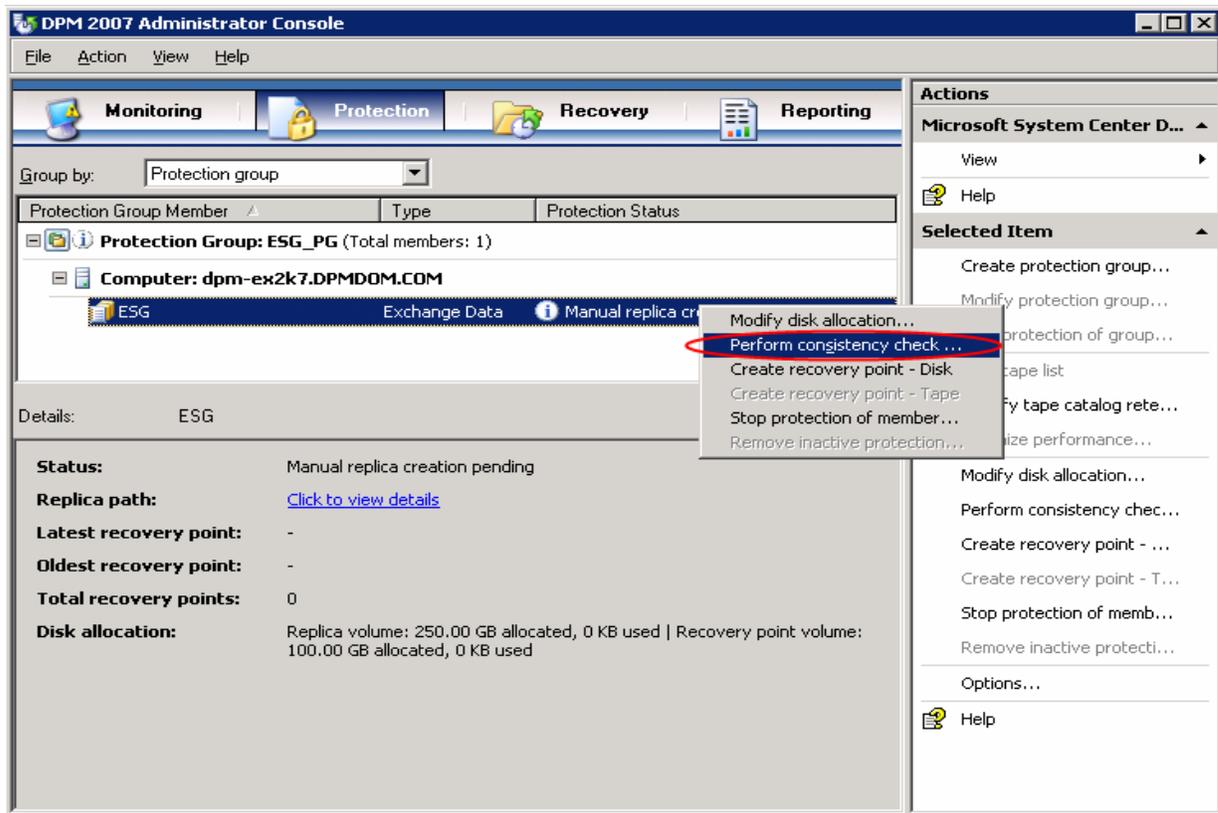


Figure 13. Perform a consistency check

After the consistency check is complete, the protection status becomes OK.

To summarize we just witnessed how easy it is to use CLARiiON clones as the replica for DPM 2007 and avoided flooding the network with a replica creation load. Compared to traditional backup techniques that take up almost all the network bandwidth, this is a very comprehensive solution to protect huge volumes of data without having to worry about the public network congestion. Also, manual replica creation is many times faster than the replica creation over LAN.

Note: When the size of data sources to be protected is considerably small and if network bandwidth is not an issue, **Automatic replica creation** can be used instead of **Manual replica creation**. In all other cases that involve huge data transfers, **Manual replica creation** mentioned previously is recommended.

SAN-based recovery using CLARiiON snapshots

Recovering data, as illustrated in Figure 14, is a high-level process that involves the following steps:

- Select the version of data that you want to recover from the recovery points on the DPM server.
- DPM restores a copy of the selected data to its point of origin on the server or to an alternate destination that you specify.

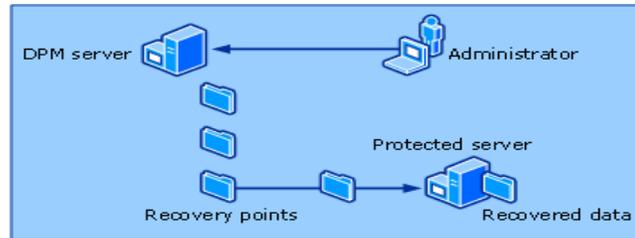


Figure 14. How to recover data

Using CLARiiON SnapView snapshots

To avoid network congestion and the longer time required for recovering data, the SAN solution provided by CLARiiON storage arrays can be used.

Imagine a scenario where the data required for recovery on a local volume on the protected server can be copied to the actual location locally instead of across the LAN. This immediately solves our problem by not just reducing but eliminating the network usage and reducing the time required for recovering the lost data. In the following sections, this white paper explains the setup, requirements, and steps for performing a SAN-based recovery using CLARiiON snapshot technology.

Introduction to snapshots

A snapshot is a virtual LUN that allows a secondary server to view a point-in-time copy of a source LUN. You determine the point in time when you start a SnapView session. The session keeps track of the source LUN's data at a particular point in time. Though a snapshot appears as a conventional LUN to other servers, its data does not reside on a disk like a conventional LUN. A snapshot is a composite of the unchanged data chunks on the source LUN and data chunks on the reserved LUN. The data chunks on the source LUN are those that have not been modified since you started the session. The data chunks in the reserved LUN pool are copies of the original source LUN data chunks that have been modified since you started the session.

During a session, the production server can still write to the source LUN and modify data. When this happens, the CLARiiON FLARE software stores a copy of the original point-in-time data on a reserved LUN in the reserved LUN pool. This operation is referred to as copy-on-first-write (COFW) because it occurs only when a data chunk is first modified on the source LUN.

As the session continues and additional I/O modifies other data chunks on the source LUN, the amount of data stored in the reserved LUN pool grows. If needed, you can increase the size of the reserved LUN pool by adding more LUNs to the LUN pool.

Important: An adequate number of reserved LUNs are essential since SnapView terminates sessions if the reserved LUN runs out of space and no additional LUNs are in the reserved LUN pool.

From a secondary server, you can view a session's point-in-time data by activating (mapping) a snapshot to the session. You can activate only one snapshot at a time to a session. If another point-in-time view is desired, you can deactivate (unmap) a snapshot from a session and activate it to another session of the same source LUN or you can create another snapshot and activate it to the second session.

Figure 15 shows an example of how snapshots work. The snapshot is a composite of the source LUN and the reserved LUN in the global reserved LUN pool.

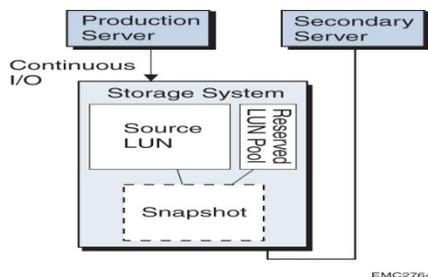


Figure 15. How snapshots work

To create a consistent replica across a set of LUNs, you can start a session in consistent mode. To read more about consistent mode and SnapView in general, refer to the *EMC SnapView for Navisphere Manager Administrator's Guide*.

Creating a snapshot using the SnapView Snapshot Configuration Wizard

Beginning in release 24, configuration of reserved LUNs and snapshots for one or more source LUNs can be automated through the SnapView Snapshot Configuration Wizard.

The Navisphere SnapView Snapshot Configuration Wizard consolidates all the steps required for configuring reserved LUNs and snapshots for one or more source LUNs. The Snapshot Configuration Wizard will:

- Bind and allocate new reserved LUNs for the selected source LUN(s).
- Create the snapshot(s) for the identified source LUN(s), if desired.
- Assign the snapshot(s) to a secondary server storage group, if desired.

Two reserved LUNs, which are usually allocated from different RAID groups from the source LUN, are created for each source LUN. Note that the wizard does not use LUNs that are already allocated from the reserved LUN pool. If you run the wizard a second time with the same selected source LUN, it attempts to remove the old reserved LUNs after creating the new ones, provided that the old LUNs are unallocated. The wizard will choose a RAID group for reserved LUNs if it meets the following criteria:

- Does not contain the source LUN
- Contains the fewest server-visible LUNs.
- Contains the fewest clones and mirror images.
- Contains the most free space.

The wizard also defaults the reserved LUN size to 30 percent of the size of the source LUN. This size can be changed by clearing the Accept Snapshot overhead values checkbox and selecting the appropriate size (with a minimum of 20 percent).

Starting a SnapView session

As noted earlier, the user can either create the snapshot first or start the snap session first. Since in this discussion we began by creating a snapshot, we will now discuss the process of starting a snap session.

When a session is started, a reserved LUN is assigned to each source LUN if no other session is already running on that source LUN. The storage system will start tracking the source LUN, invoking copy-on-first-write operations for data areas that are changed on the source LUN. When starting a session, you have two options:

- **Single-LUN session:** You may wish to have a session running on a single source LUN. If using the Navisphere UI, right-click the desired LUN and select **Start Session**. Enter a session name, then click **OK**, and the session will start.
- **Multi-LUN session:** For consistency, you may want to start a session containing multiple source LUNs when you have an interdependent data set (for example, a database consisting of several LUNs). With the addition of these consistent operations in FLARE release 19, users can now have consistent local replicas using SnapView sessions owned by either SP.

SAN-based DPM recoveries nearly always need this mode because replica and recovery point volumes nearly always reside on a different LUN.

As with other SnapView operations, start can be initiated in the Navisphere Manager GUI or Navisphere CLI. Admsnap can also be used to issue the start command. You also have the option of managing consistent operations manually (through Navisphere Manager or CLI), or to incorporate the consistent operations into a script (using Navisphere CLI or admsnap, as applicable). If using the Navisphere UI, right-click the storage system, and select **Start Session**. This opens a dialog box in which you select the source LUNs for the session. Give the session a name and click **OK** to start.

Activating the snapshot to a session

Once you have at least one SnapView session running and a snapshot created for a given source LUN, you can activate a snapshot to a session. This action essentially associates a snapshot to the point-in-time view provided by a particular session. Using the Navisphere UI, you can activate a snapshot by right-clicking the snapshot and selecting **Activate**.

Reserved LUN pool configuration guidelines

This section discusses the best practice recommendation for configuring LUNs for the SnapView snapshot reserved LUN pool.

Best practice recommendations for the reserved LUN pool configuration:

- Avoid placing reserved LUN pool LUNs on the same drives as the source LUNs you will snap. Write operations will result in very high seek times and disappointing performance. The same holds true for clone LUNs, too. Put them on disk groups separate from the LUNs they are cloning.
- It is recommended to use faster FC drives in the reserved LUN pool if available.
- Distributing LUNs to the reserved LUN pool across multiple RAID groups can minimize spindle contention and queuing due to copy-on-first-write or concurrent write activity for all SnapView sessions and snapshots.
- It is recommended to use RAID 5 as the RAID type choice for the reserved LUN since it is well equipped to handle the copy-on-first-write task performed by the reserved LUNs.
- A general guideline is that the reserved LUN typically needs to be 20 percent of the size of the source LUN. There are many other factors to consider when determining the size of the reserved LUNs. Key among these factors, however, is the duration of a given session and the amount of source data that changes during a session.

For more information about the reserved LUN pool's allocation considerations for CLARiiON storage arrays, refer to the white paper *EMC CLARiiON Reserved LUN Pool Configuration Considerations*.

SAN-based recovery using snapshots

This section describes the procedure and benefits of SAN-based recovery using the CLARiiON snapshot feature. In short, we're going to take the hardware snapshot of the replica volume and recovery point volume of the DPM server using CLARiiON SnapView. This snapshot will be mounted on to the production server to perform DPM's SAN-based recovery.

As an example we'll demonstrate the SAN-based recovery process for recovering a complete Exchange storage group (the same storage group that we protected using DPM in "[Using CLARiiON SnapView clones](#)") using CLARiiON snapshots.

Before creating the snapshot

Before creating the snapshot of the replica volume and recovery point volume, cancel any jobs that are executing by clicking the **Jobs** tab in the Monitoring task area in the DPM Administration Console, right-clicking on the jobs that are **In progress** and clicking **Cancel** (Figure 16). Also make sure that the Exchange database (ESG_DB) to recover has the overwrite flag set.

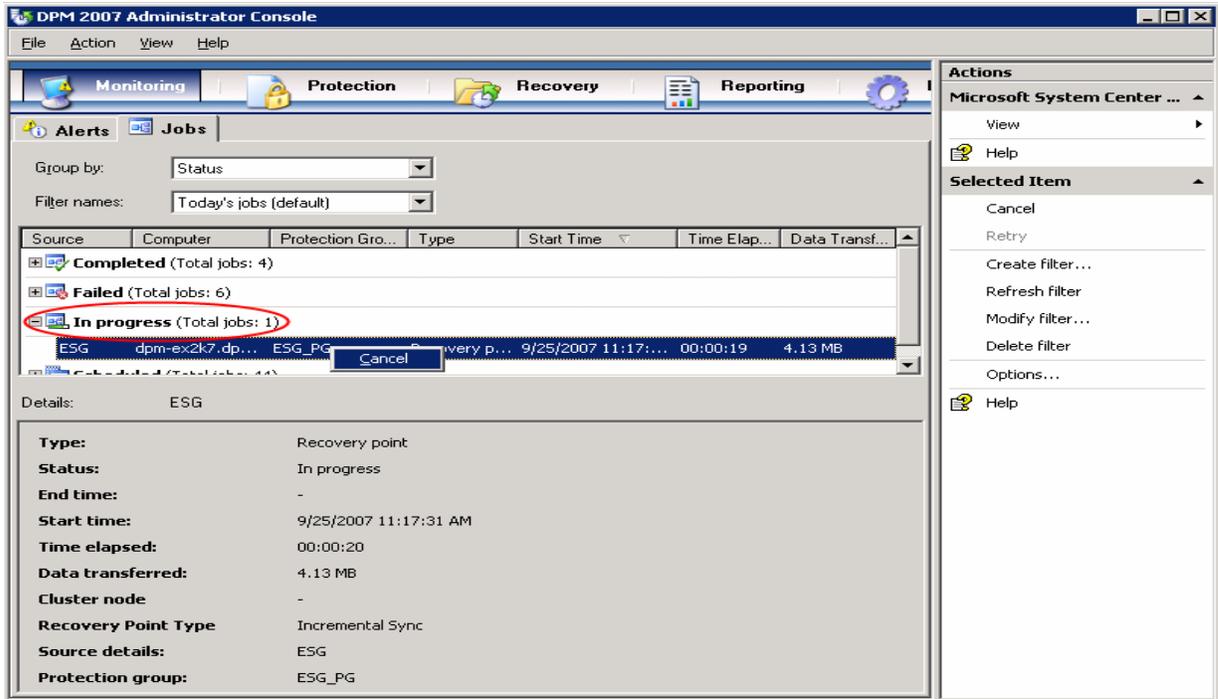


Figure 16. Click on the jobs In progress

Create a PowerShell script file named `CreateShadowCopy.ps1` (script contents are in the “Appendix”) in “C:\Program Files\Microsoft DPM\DPM\bin” on the DPM server. Execute the `CreateShadowCopy.ps1` PowerShell script on the DPM server (Figure 17). Provide the data source name and protection group name as the input parameters. In our example, the input parameters are `ESG` and `ESG_PG`.

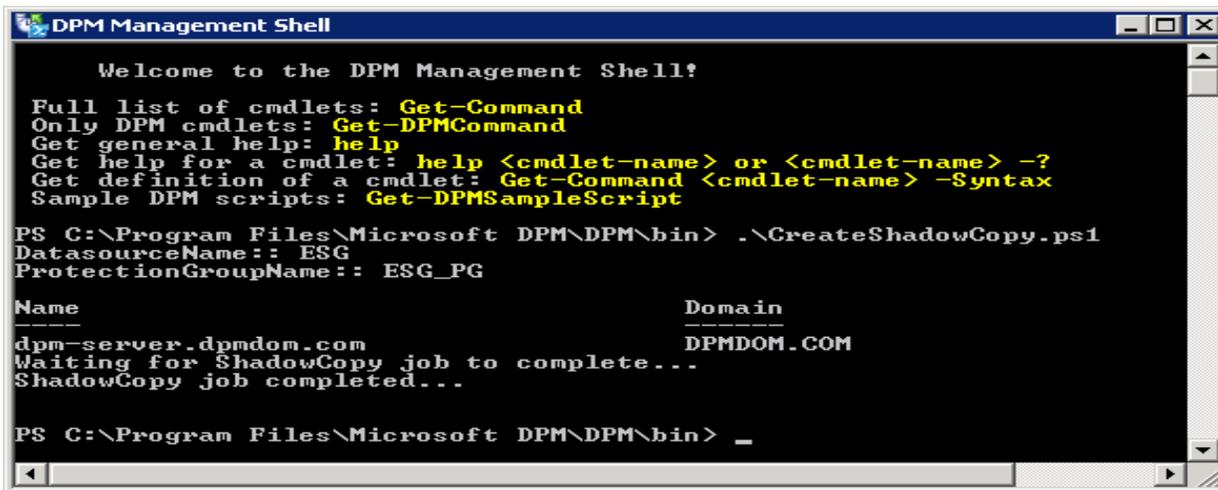


Figure 17. Execute the PowerShell script

Creating the SnapView snapshot of replica and recovery point volumes

Create the CLARiiON SnapView snapshot of the DPM replica volume and recovery point volume and add them to the production server’s storage group (Figure 18) using the Navisphere management UI as described in the section “[Creating a snapshot using the SnapView Snapshot Configuration Wizard.](#)”

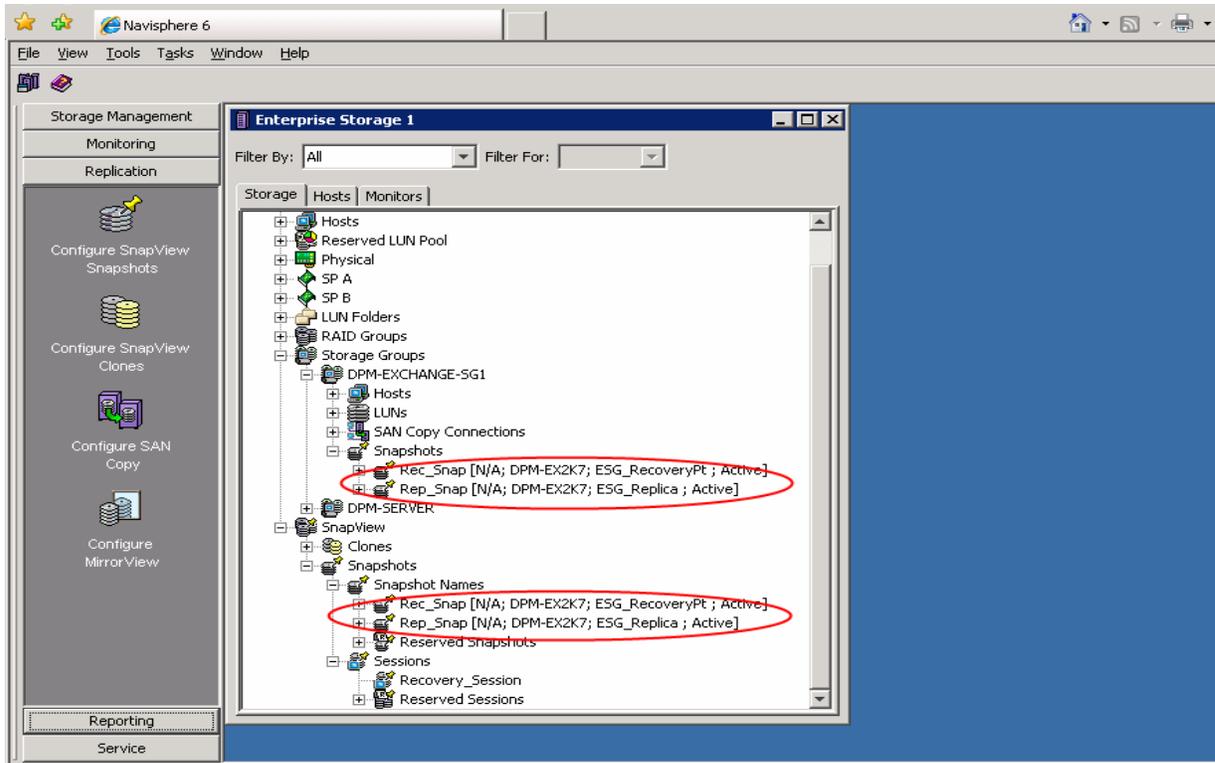


Figure 18. Add the snapshots to the storage group

Start a SnapView session and activate the snapshots to the session using the Navisphere management UI as described in the sections “[Starting a SnapView session](#)” and “[Activating the snapshot to a session](#).”

Navisphere also provides a secure CLI, NavisecCli, to perform all the previously mentioned operations from a command window.

For your convenience, if you do not want to go through all the manual steps involved in creating the snapshots, activating, and assigning them to the production server, we have developed a script, CreateDPMSnapshot.bat, to automate the process. The contents of this script can be found in the “Appendix” section.

Before executing this script or any Naviseccli commands on a host, add yourself to the security file on the current host (if you haven’t done this already on this host) by executing the following command:

```
Naviseccli -address StorageProcessorIP -AddUserSecurity -password NavisphereLoginPassword -scope 0 -user NavisphereLoginUsername
```

Now on this host you can enter CLI commands to any storage system on which you have an account that matches the username *NavisphereLoginUsername*, with the password *NavisphereLoginPassword* and global scope (scope 0).

CreateDPMSnapshot.bat

Syntax: CreateDPMSnapshot.bat *StorageProcessor_IPaddress Replica_LUN_Number Replica_snapshotname RecoveryPt_LUN_Number RecoveryPt_snapshotname Snapsession_name Production_Server_StorageGroup_Name*

Example: CreateDPMSnapshot *1.1.1.1 14 Replica_Snap 15 RecoveryPt_Snap Recovery_Session EXCH_SG*

Note: This gives you an idea about combining various Navisphere commands into a script and executing them sequentially. This has been tailored to the recovery scenario we’re dealing with in this paper. You can also create such scripts using Navisphere commands customized to suit your environment and application setup.

To read more about the CLI, its syntax, and usage refer to the *EMC SnapView Command Line Interfaces (CLI) Reference* on Powerlink.

At this point, the two snapshots will appear as two new disks with partitions (reflecting the DPM replica volume and recovery point volume) on the Microsoft disk management snap-in of the production server. Mount these partitions (can be mounted to a NTFS folder as shown in Figure 19 and Figure 20) using the Microsoft disk management snap-in on the production server.

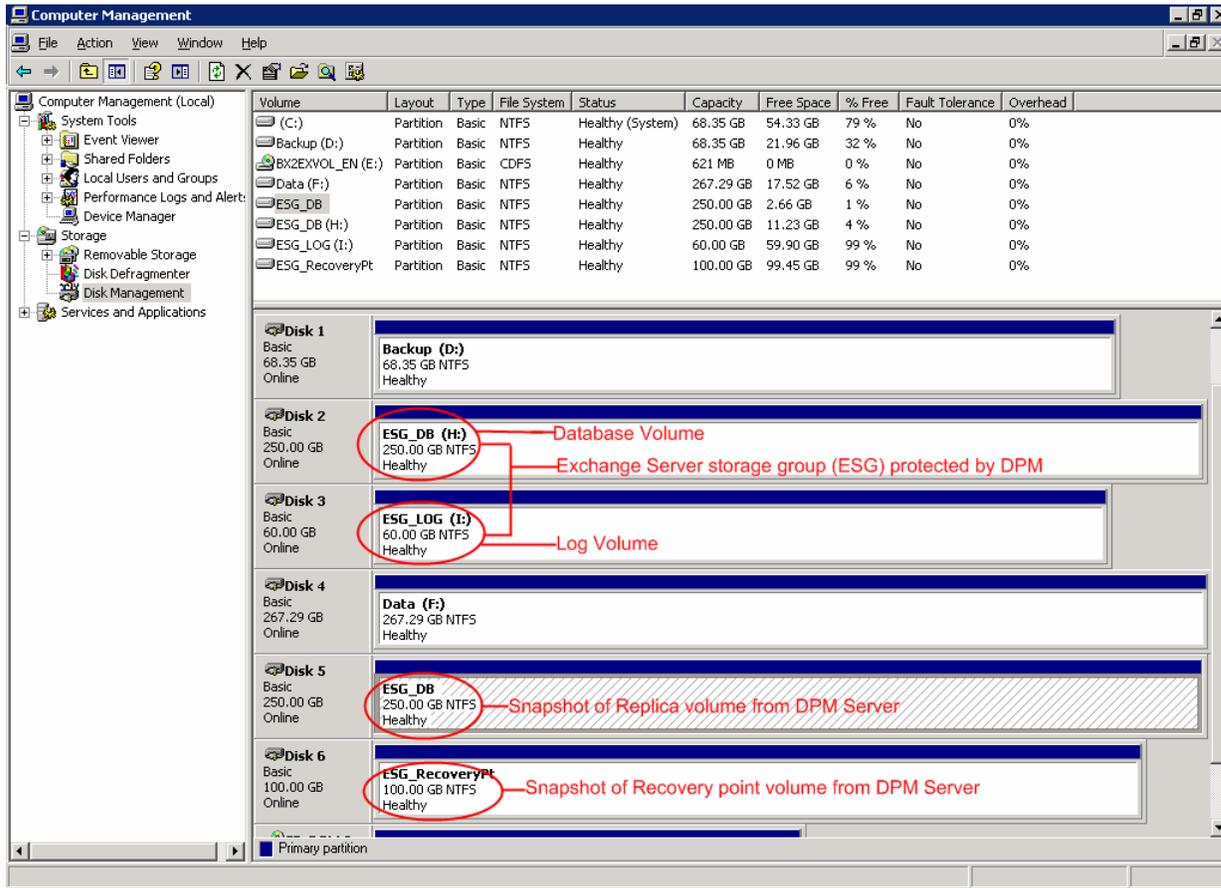


Figure 19. Snapshots of replica and recovery point volumes

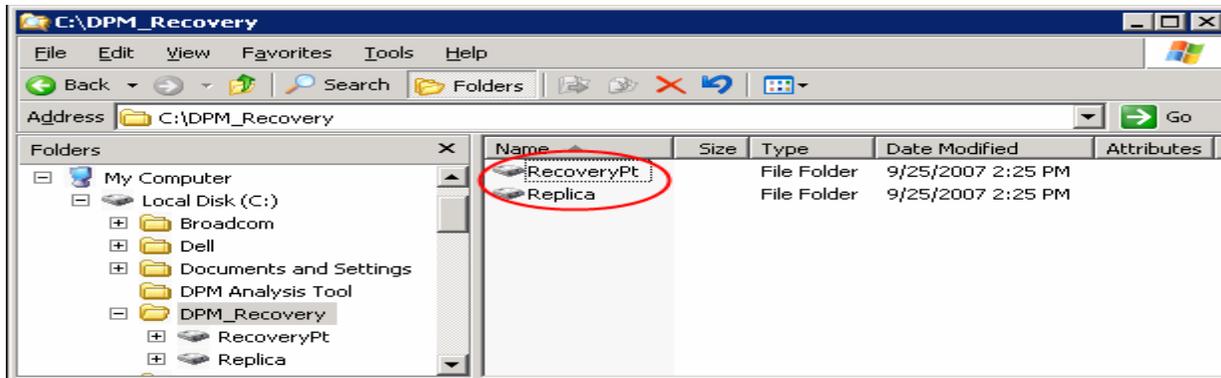


Figure 20. Mounted snapshots of replica and recovery point volumes

To read more about taking a snapshot of a LUN, activating it and assigning it to a host, refer to the *EMC SnapView for Navisphere Manager Administrator's Guide* on Powerlink.

The actual SAN-based recovery

In the DPM 2007 Administrator Console, in the Recovery task area, select the data to be recovered in the left pane. Now all available recovery points are listed in the **Recovery time** combo box in the center pane. This is why a snapshot of both replica and recovery point volumes is exposed to the protected server, to provide access to all recovery points as in any DPM recovery operation. Select a time from the list and click **Recover** in the actions pane (in the case of Exchange, you can restore an entire storage group as shown in Figure 21).

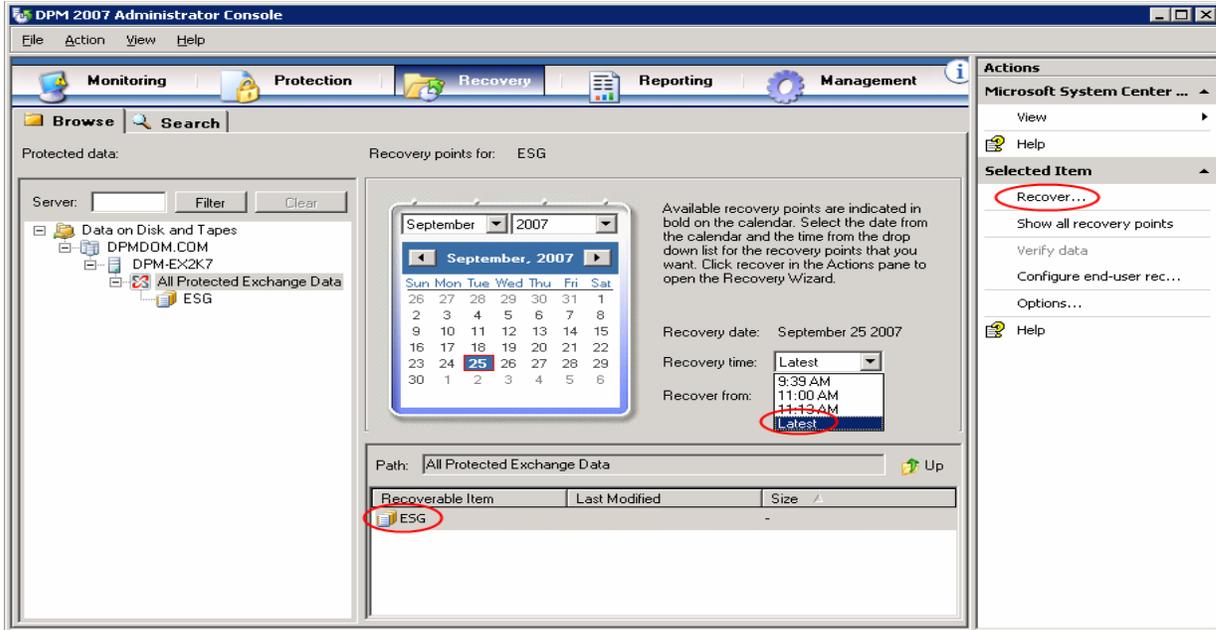


Figure 21. Select a recovery time in the DPM 2007 Administrator Console

The review recovery selection pane comes up. Verify the data to be recovered and click **Next**.

In the Select Recovery Type screen select the recovery type. As shown in Figure 22, the selection is **Recover to original Exchange Server location**. Click **Next**.

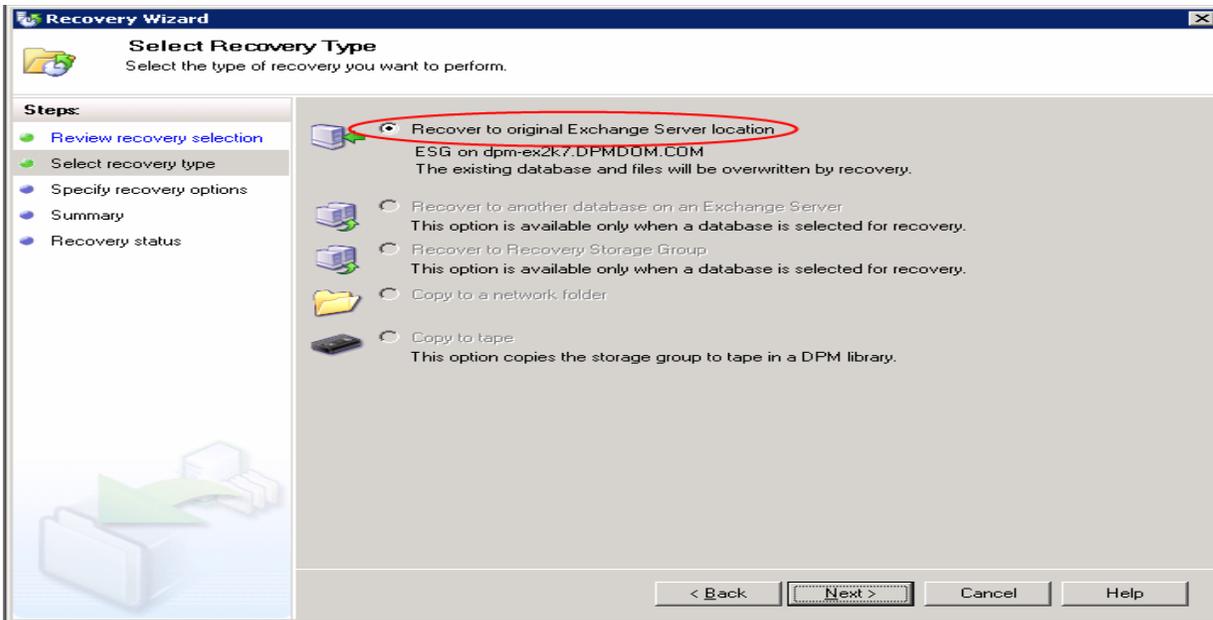


Figure 22. Selecting the recovery type

On the Specify Recovery Options page, select the checkbox **Enable SAN based recovery using hardware snapshots**. This is the important step that helps us to use the CLARiiON snapshots feature in performing SAN recovery. If recovering Exchange server data then also select the **Mount the database after they are recovered** option. Click **Next**.

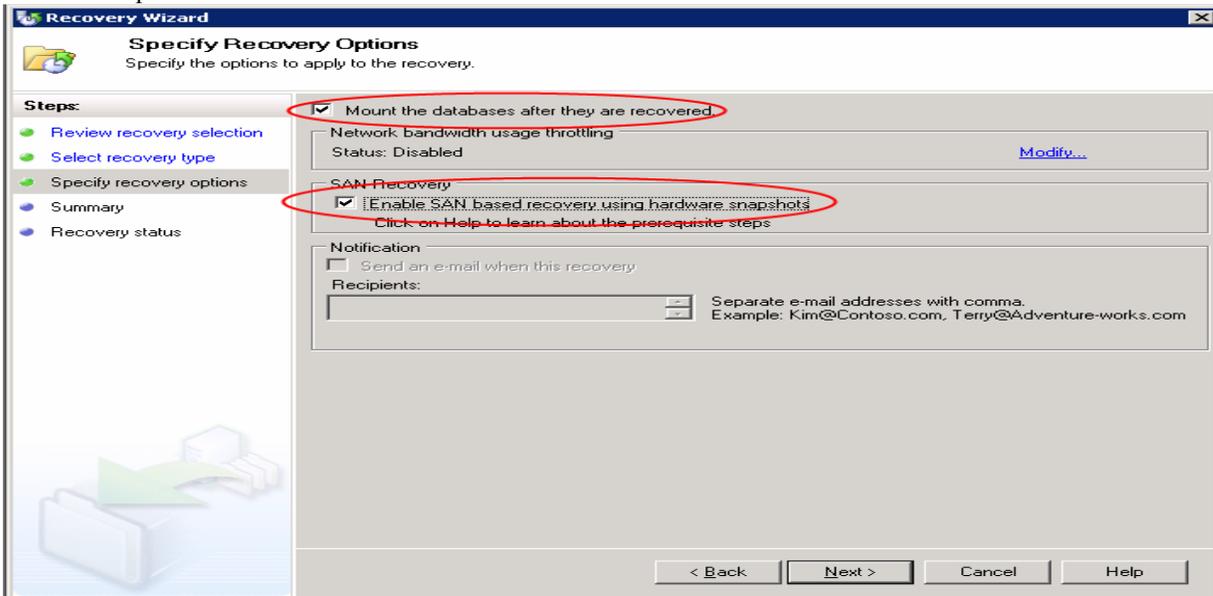


Figure 23. Selecting the recovery options

On the Summary page, click **Recover** to start the recovery process (Figure 24).

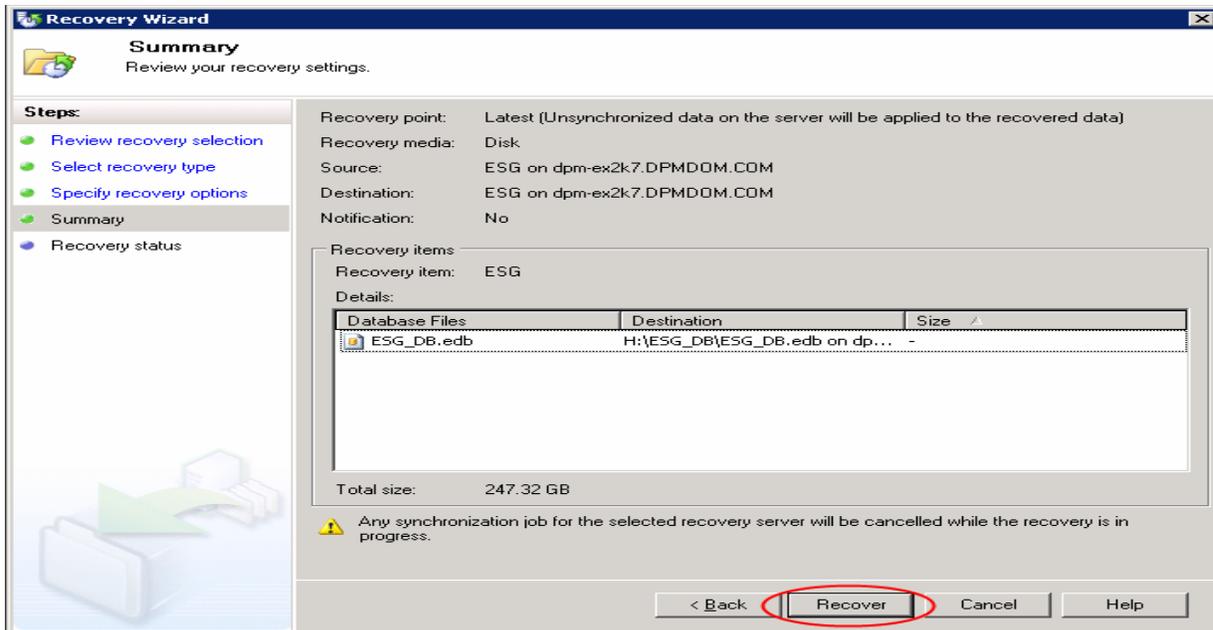


Figure 24. Start the recovery process

After the recovery is completed dismount the replica and recovery point snapshots from the production server. Use the DISKPART utility REMOVE command to accomplish this: DISKPART> REMOVE MOUNT=<path> DISMOUNT. Remove the replica and recovery point snapshots from the Exchange server storage group, delete them, and also stop the SnapView session that was created earlier. This must be done to free the LUNs in the reserved LUN pool so that they'll be available for future recoveries.

For your convenience, if you do not want to go through all the manual steps involved in removing the snapshots from the production server storage group, deleting the snapshots, and stopping the session, we have developed a script RemoveDPMSnapshot.bat to automate the process. The contents of this script can be found in the "Appendix" section. Execute this script after dismounting the snapshots from the production server using the Microsoft disk management snap-in.

Before executing this script or any Navisecli commands on a host, add yourself to the security file on the current host (if you haven't done this already on this host) by executing the following command:

```
Navisecli -address StorageProcessorIP -AddUserSecurity -password NavisphereLoginPassword - scope 0 -user NavisphereLoginUsername
```

Now on this host you can enter CLI commands to any storage system on which you have an account that matches the username NavisphereLoginUsername, with the password NavisphereLoginPassword and global scope (scope 0).

RemoveDPMSnapshot.bat

Syntax: RemoveDPMSnapshot.bat *StorageProcessor_IP Production_Server_StorageGroup_Name Replica_snapshotname RecoveryPt_snapshotname Snapsession-name*

Example: RemoveDPMSnapshot 1.1.1.1 EXCH_SG Replica_Snap RecoveryPt_Snap Recover_Session

Note: This gives you an idea about combining various Navisphere commands into a script and executing them sequentially. This has been tailored to the recovery scenario in this paper. You can also create such scripts using Navisphere commands customized to suit your environment and application setup.

When recovery is complete, the DPM replica will be inconsistent and DPM cannot resume protection of the recovered volume(s) because there is no way that DPM can guarantee post-recovery consistency in all cases without actually verifying that. In the Protection task area, run a consistency check on the

corresponding protection group (ESG_PG) as shown in Figure 13 on page 16. This will typically transfer very small amounts of data if any at all.

Since the SAN-based recovery option was selected, we were able to present the LUNs (snapshot of replica and recovery points) containing the recovery data as a local disk to the production host. This saves the network bandwidth to be used for some other purposes since DPM is no longer going to recover the data from a network resource. Even now DPM takes care of issuing appropriate commands and managing the process of recovering the data from the snapshot but with the added advantage that the data no longer has to travel through the LAN and suffocate it; rather the data gets copied from one LUN to another as a local disk-to-disk copy within the protected server. With the much improved throughput compared to LAN, CLARiiON aids DPM in performing the recovery in a considerably quicker rate compared to recovery over LAN.

Again, remember the time and network bandwidth that has been saved by selecting this SAN-based recovery with CLARiiON snapshot.

Requirements/prerequisites

DPM server

The System Center Data Protection Manager 2007 (DPM) server must be a dedicated, single-purpose server, and cannot be either a domain controller or an application server.

To install DPM successfully, you must install the software listed in Table 1 before installing the DPM application. DPM Setup will then install the additional required prerequisite software.

If you want to install the required prerequisite software manually, you must follow the steps detailed in "Manually Installing Prerequisite Software" in *Deploying Data Protection Manager 2007*.

Table 1 lists the DPM server prerequisite software that DPM Setup does not install. You must install the following software before you can install DPM.

Table 1. Prerequisite software (before installing DPM)

Software	Install from
Microsoft Management Console 3.0	http://go.microsoft.com/fwlink/?LinkId=55423
Hotfix 891957. This hotfix resolves an issue with Windows-based systems that may cause them to deplete their paged pool if the Volume Shadow Copy Service is placed under heavy load.	32-bit x86 operating systems: http://go.microsoft.com/fwlink/?LinkId=48584 x64-bit operating systems: http://go.microsoft.com/fwlink/?LinkId=75131
Windows PowerShell 1.0	32-bit x86 operating systems: http://go.microsoft.com/fwlink/?LinkId=658 x64-bit operating systems: http://go.microsoft.com/fwlink/?LinkId=65814

The following is the DPM server prerequisite software that DPM Setup installs before installing the DPM application.

- Microsoft .NET Framework 2.0
- Microsoft Software Quality Metrics (SQM)
- Single Instance Storage (SIS)

- Internet Information Services (IIS) 6.0

Protected server

Each server Microsoft System Center Data Protection Manager 2007 Beta2 (DPM) protects must meet the requirements in Table 2.

Table 2. Protected server requirements

Protected servers	Server requirements
File servers	<p>Windows Server 2003 with Service Pack 1 (SP1) Windows Server 2003 x64 Windows Server 2003 R2 Windows Server 2003 R2 x64 Windows Storage Server 2003 with Service Pack 1 (SP1) (To obtain SP1 for Windows Storage Server 2003, contact your original equipment manufacturer.) Windows Storage Server 2003 R2 Windows Storage Server 2003 R2 x64</p> <p>Note: DPM supports Standard and Enterprise Editions of all the required operating systems.</p>
Computers running SQL Server	<p>Microsoft SQL Server 2000 with Service Pack 4 (SP4) - OR - Microsoft SQL Server 2005 with Service Pack 1 (SP1) or Service Pack 2 (SP2)</p> <p>Note: DPM supports Standard, Enterprise, Workgroup, and Express Editions of SQL Server.</p> <p>Important: You must start the SQL Server VSS Writer Service on SQL Server before you can start protecting SQL Server data. By default, the SQL Server VSS Writer Service is turned off when you install SQL Server 2005.</p> <p>To start the SQL Server VSS Writer Service: Click Start, point to Administrative Tools, and then click Services. On the Services screen, scroll down and right-click SQL Server VSS writer, and then click Start.</p>
Computers running Exchange Server	<p>Exchange Server 2003 with Service Pack 2 (SP2) - OR - Exchange Server 2007</p> <p>Note: DPM supports Standard and Enterprise Editions of Exchange Server.</p>
Computers running Virtual Server	<p>Microsoft Virtual Server 2005 R2 Service Pack 1 (SP1)</p>
Windows SharePoint Services	<p>Windows SharePoint Services (WSS) 3.0 Microsoft Office SharePoint Server (MOSS) 2007</p> <p>Important: Before you can protect SharePoint data, you must start the SharePoint Writer service on the SharePoint server, and then provide the protection agent with credentials for the SharePoint Services farm. For more information, see <i>Starting and Configuring the SharePoint Writer Service</i>.</p>

Shared disk clusters	File servers SQL Server 2000 with Service Pack 4 (SP4) SQL Server 2005 with Service Pack 1 (SP1) Exchange Server 2003 with Service Pack 2 (SP2) Exchange Server 2007 Beta 2 Note: Only one Network Name resource can exist for the resource group that you are protecting. If there is more than one Network Name resource for a single resource group, DPM can support this configuration only if all dependent resources are associated with the same Network Name resource. For example, if you attempt to protect a SQL shared disk cluster, the physical disk resource that SQL uses must be associated with the same Network Name resource as the computer running SQL Server and the SQL Server Agent.
Non-shared disk clusters	Exchange Server 2007

SAN (CLARiiON storage array requirements)

- SnapView must be enabled on the CLARiiON array.
- The production server must use a CLARiiON array to store the production data.
- The DPM server must use a CLARiiON array to store the recovery point data.
- The DPM and production server must be registered with the CLARiiON array that holds the production data.
- The DPM and production server must have connectivity to the production array.
- To use SnapView clones and snapshots, it is mandatory that the production LUN and DPM LUNs must be provided by the same CLARiiON array.
 - The replica volume and recovery point volume must be two separate LUNs on the CLARiiON array.
 - The CLARiiON array used must have sufficient space to create clones (this is theoretically the same as the Production LUN size).
 - The CLARiiON array must have sufficient reserved LUNs to be used by SnapView snapshots.
 - Navisphere CLI must be installed on the host from which SAN-based recovery is done. This is necessary to execute the scripts (CreateDPMSnapshot.bat and RemoveDPMSnapshot.bat) for automating snapshot creation and removal.

Conclusion

Microsoft Data Protection Manager 2007 with EMC CLARiiON arrays gives you a comprehensive solution to protect your application data as well as fast recovery of data. DPM manages the application related functions like monitoring changed bits and performing synchronization, and CLARiiON provides faster replica creation and SAN-based recovery. Since you're using a CLARiiON storage array, you're automatically receiving not only the performance benefits but also high availability and on-the-fly storage expansion features that are very much required for business continuity. With the introduction of the CX3 UltraScale™ models, EMC CLARiiON arrays have achieved “five 9s”(qualifying for five 9s availability requires that there be no more than 10 hours of downtime for every 1 million runtime hours) availability. This has been achieved by the robust design of the CLARiiON CX3 UltraScale architecture and unique fault detection and isolation capabilities of the new FLARE operating environment. The white paper *EMC CLARiiON CX3 Best Practices for Achieving “Five 9s” Availability* can provide more information.

Appendix: Scripts

CreateShadowCopy.ps1	<pre>if(!\$args[0]) { if(!\$DSName) { \$DSName = read-host "DatasourceName:" } } else { if(("-"?"-help") -contains \$args[0]) { write-host Usage:: write-host CreateShadowCopy.ps1 DatasourceName ProtectionGroupName write-host Help:: write-host Creates a shadow copy for the given Datasource write-host exit 0 } else { write-host "Usage -? for Help" exit 1 } } if(!\$PGName) { \$PGName = read-host "ProtectionGroupName:" } \$dpmname = &"hostname" connect-dpmserver \$dpmname \$pg = get-protectiongroup -dpmservername \$dpmname if (!\$pg) { write-error "Cannot get the protectionGroup" disconnect-dpmserver \$dpmname exit 1 } \$mypg = \$pg where {\$_.FriendlyName -eq \$PGName} if (!\$mypg) { write-error "Cannot get the requested protectionGroup" disconnect-dpmserver \$dpmname exit 1 } \$ds = get-datasource -protectiongroup \$mypg if (!\$ds) { write-error "Cannot get the datasources for the PG" disconnect-dpmserver \$dpmname exit 1 } \$myds = \$ds where {\$_.Name -eq \$DSName} if (!\$myds) {</pre>
----------------------	--

	<pre>write-error "Cannot get the required Datasource" disconnect-dpmserver \$dpmname exit 1 } \$j = new-recoverypoint -datasource \$mysds -DiskRecoveryPointOption WithoutSynchronize -Disk if (!\$j) { write-error "Cannot get the required Datasource" disconnect-dpmserver \$dpmname exit 1 } \$jobtype = \$j.jobtype while (! \$j.hascompleted) { write-host "Waiting for \$jobtype job to complete..."; start-sleep 5 } if(\$j.Status -ne "Succeeded") { write-error "Job \$jobtype failed..." } Write-host "\$jobtype job completed..." disconnect-dpmserver \$dpmname exit</pre>
--	--

<p>CreateDPMSnapshot.bat</p> <div data-bbox="228 380 461 527" style="border: 1px solid green; padding: 5px; margin-bottom: 10px;"> <i>Use this for a 64-bit OS environment</i> </div> <div data-bbox="228 548 461 674" style="border: 1px solid blue; padding: 5px;"> <i>Use this for a 32-bit OS environment</i> </div>	<pre> :: This script will crate a snapshot of the DPM Replica and Recovery :: point LUNs, :: Creates a SnapView session, activate the snapshots on this session ::and assign :: these snapshots to the Production Server to be used by DPM for ::SAN- based recovery. cd "c:\Program Files (x86)\EMC\Navisphere CLI" (OR) cd "c:\Program Files\EMC\Navisphere CLI" NaviSECCLI.exe -h %1 snapview -createsnapshot %2 -snapshotname %3 NaviSECCLI.exe -h %1 snapview -createsnapshot %4 -snapshotname %5 NaviSECCLI.exe -h %1 snapview -startsession %6 -snapshotname %3,%5 - consistent NaviSECCLI.exe -h %1 snapview -activatesnapshot %6 -snapshotname %3 NaviSECCLI.exe -h %1 snapview -activatesnapshot %6 -snapshotname %5 NaviSECCLI.exe -h %1 storagegroup -addsnapshot -gname %7 -snapshotname %3 NaviSECCLI.exe -h %1 storagegroup -addsnapshot -gname %7 -snapshotname %5 </pre>
<p>RemoveDPMSnapshot.bat</p> <div data-bbox="228 978 461 1104" style="border: 1px solid green; padding: 5px; margin-bottom: 10px;"> <i>Use this for a 64-bit OS environment</i> </div> <div data-bbox="228 1125 461 1262" style="border: 1px solid blue; padding: 5px;"> <i>Use this for a 32-bit OS environment</i> </div>	<pre> :: Script to remove the snapshot and stop the snapshot session of the ::DPM Replica and Recovery point volumes cd "c:\Program Files (x86)\EMC\Navisphere CLI" (OR) cd "c:\Program Files\EMC\Navisphere CLI" NaviSECCLI.exe -h %1 storagegroup -removesnapshot -gname %2 - snapshotname %3 -o NaviSECCLI.exe -h %1 storagegroup -removesnapshot -gname %2 - snapshotname %4 -o NaviSECCLI.exe -h %1 snapview -rmsnapshot -snapshotname %3 -o NaviSECCLI.exe -h %1 snapview -rmsnapshot -snapshotname %4 -o NaviSECCLI.exe -h %1 snapview -stopsession %5 -o </pre>