

EMC Solutions for Microsoft Exchange 2007 NS Series iSCSI

Applied Technology

Abstract

This white paper presents the latest storage configuration guidelines for Microsoft Exchange 2007 on the Celerra[®] NS Series using iSCSI. It provides step-by-step instructions for creating file systems and iSCSI LUNs on Celerra.

July 2007

Copyright © 2007 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 H2879

Table of Contents

Executive summary	4
Introduction	4
Audience	4
Terminology	4
Celerra file system creation	5
File system layout	5
iSCSI LUN creation	17
Optional configuration step	33
Conclusion	33
References	33

Executive summary

E-mail is a mission-critical application for most businesses today. It is important to plan a Microsoft Exchange solution that can grow while maintaining optimum performance, high availability, and disaster recovery.

Introduction

This white paper is meant to be a technical guide for configuration of Exchange 2007 storage on EMC® Celerra® via iSCSI. It also provides step-by-step guidelines for creating file systems and creating iSCSI LUNs.

Audience

This white paper is part of the NS Series documentation set, and is intended for use by Exchange 2007 administrators. Readers are expected to be familiar with the following topics:

- Microsoft Exchange 2007
- Celerra NS Series
- iSCSI networking

Terminology

Active Directory: An advanced directory service introduced with Windows 2000 Server. It stores information about objects on a network and makes this information available to users and network administrators through a protocol such as LDAP.

Automatic Volume Management (AVM): A feature of the Celerra Network Server that creates and manages volumes automatically, without manual volume management by an administrator. AVM organizes volumes into pools of storage that can be allocated to file systems.

Data Mover: A Celerra Network Server cabinet component running the data access in real time (DART) operating system that retrieves files from a storage device and makes the files available to a network client.

disk volume: On Celerra systems, a physical storage unit as exported from the storage array. All other volume types are created from disk volumes.

iSCSI (Internet SCSI): A protocol for sending SCSI packets over TCP/IP networks.

iSCSI initiator: An iSCSI endpoint, identified by a unique iSCSI-recognized name that begins an iSCSI session by issuing a command to the other endpoint (the iSCSI target).

iSCSI target: An iSCSI endpoint, identified by a unique iSCSI-recognized name that executes commands issued by the iSCSI initiator.

MC/S: Multiple connections per session.

RAID: Redundant array of independent disks. A method for storing information where the data is stored on multiple disk drives to increase performance and storage capacities and to provide redundancy and fault tolerance.

RAID 1: RAID method that provides data integrity by mirroring (copying) data onto another disk. This RAID type provides the greatest assurance of data integrity at the greatest cost in disk space.

RAID 5: Data is striped across disks in large stripes. Parity information is stored so data can be reconstructed if needed. One disk can fail without data loss. Performance is good for reads but slower for writes.

RAID group: The CLARiiON® storage-system term for a Celerra disk group.

SP: Storage processor on a CLARiiON storage system. On a CLARiiON storage system, a circuit board with memory modules and control logic that manages the storage-system I/O between the host's Fibre Channel adapter and the disk modules.

SP A: Storage processor A. Generic term for the first storage processor in a CLARiiON storage system.

SP B: Storage processor B. Generic term for the second storage processor in a CLARiiON storage system.

Volume Shadow Copy Service (VSS): A Windows service and architecture that coordinates various components to create consistent point-in-time copies of data called "shadow copies."

Celerra file system creation

File system layout

This section describes the steps performed to create file systems required for building block configuration on a Celerra system. In Figure 1, the two Celerra disk volumes in the RAID Group 8, RG8, are placed in a Celerra user-defined storage pool named LogSG1_SG2. The Celerra file system used for storing the Exchange log files for Storage Groups 1 and 2 is created from the LogSG1_SG2 storage pool. If the desired file system size is larger than the capacity of a single disk volume, Celerra automatically concatenates the two disk volumes.

For the Exchange database, the two Celerra disk volumes in RG9 are concatenated in to a Celerra metavolume. Similarly, the two Celerra disk volumes in RG10 are concatenated in to a Celerra metavolume. A Celerra stripe volume with a stripe size of 32KB is then created across these two new metavolumes. This stripe volume is placed in a Celerra user-defined storage pool named DbSG1_SG2. The Celerra file system used to store the Exchange database for Storage Groups 1 and 2 is created from the DbSG1_SG2 storage pool.




Figure 1. Physical spindle layout

The following is the step-by-step procedure for creating the Microsoft Exchange log file system and Exchange database file system.

Volumes nasadmin@10.6.24.38

Show Volumes of Type: Storage Systems:

Search

Name	Type	Used by	Storage Capacity	Storage Used(%)	Disk Type	Disk Group	Data Mover	Directors	Storage Systems	Storage Protection	Storage Pool
d3	disk	md3	2.0 GB		CLSTD	0000	server_2,s...	SP-A	CLARiiON CX3-40 APM00064501699	RAID5(4+1)	
d4	disk	md4	2.0 GB		CLSTD	0000	server_2,s...	SP-A	CLARiiON CX3-40 APM00064501699	RAID5(4+1)	
d5	disk	md5	2.0 GB		CLSTD	0000	server_2,s...	SP-A	CLARiiON CX3-40 APM00064501699	RAID5(4+1)	
d6	disk	md6	2.0 GB		CLSTD	0000	server_2,s...	SP-A	CLARiiON CX3-40 APM00064501699	RAID5(4+1)	
d159	disk		52.3 GB		CLSTD	0000	server_2,s...	SP-A	CLARiiON CX3-40 APM00064501699	RAID5(4+1)	
d174	disk		52.3 GB		CLSTD	0000	server_2,s...	SP-B	CLARiiON CX3-40 APM00064501699	RAID5(4+1)	
d160	disk		66.8 GB		CLSTD	0008	server_2,s...	SP-A	CLARiiON CX3-40 APM00064501699	RAID1	
d175	disk		66.8 GB		CLSTD	0008	server_2,s...	SP-B	CLARiiON CX3-40 APM00064501699	RAID1	
d161	disk		66.8 GB		CLSTD	0009	server_2,s...	SP-A	CLARiiON CX3-40 APM00064501699	RAID1	
d176	disk		66.8 GB		CLSTD	0009	server_2,s...	SP-B	CLARiiON CX3-40 APM00064501699	RAID1	
d162	disk		66.8 GB		CLSTD	0010	server_2,s...	SP-A	CLARiiON CX3-40 APM00064501699	RAID1	
d177	disk		66.8 GB		CLSTD	0010	server_2,s...	SP-B	CLARiiON CX3-40 APM00064501699	RAID1	
d163	disk		66.8 GB		CLSTD	0011	server_2,s...	SP-A	CLARiiON CX3-40 APM00064501699	RAID1	
d178	disk		66.8 GB		CLSTD	0011	server_2,s...	SP-B	CLARiiON CX3-40 APM00064501699	RAID1	
d164	disk		66.8 GB		CLSTD	0012	server_2,s...	SP-A	CLARiiON CX3-40 APM00064501699	RAID1	
d179	disk		66.8 GB		CLSTD	0012	server_2,s...	SP-B	CLARiiON CX3-40 APM00064501699	RAID1	
d165	disk		66.8 GB		CLSTD	0013	server_2,s...	SP-A	CLARiiON CX3-40 APM00064501699	RAID1	
d180	disk		66.8 GB		CLSTD	0013	server_2,s...	SP-B	CLARiiON CX3-40 APM00064501699	RAID1	
d166	disk		66.8 GB		CLSTD	0014	server_2,s...	SP-A	CLARiiON CX3-40 APM00064501699	RAID1	

46 items displayed.

New Delete

Figure 2. Disk volume list

To create the Exchange log file system, perform the following steps:

1. To identify disk volumes that belong to the same RAID group, do the following:
 - a. From the Celerra Manager > Storage > Volumes page, select **disk** from the **Show volume of Type** list box.
 - This lists all disk volumes on your Celerra system.
 - b. To identify the disk volumes that belong to the same RAID group on your system, sort the list by clicking on the **Disk Group** column heading.
 - Figure 2 and Figure 3 show examples of the disk volumes from the Celerra Manager.

Show Volumes of Type:

Storage Systems:

Name	Type	Used by	Storage Capacity	Storage Used(%)	Disk Type	Disk Group
d3	disk	md3	2.0 GB		CLSTD	0000
d4	disk	md4	2.0 GB		CLSTD	0000
d5	disk	md5	2.0 GB		CLSTD	0000
d6	disk	md6	2.0 GB		CLSTD	0000
d159	disk		52.3 GB		CLSTD	0000
d174	disk		52.3 GB		CLSTD	0000
d160	disk		66.8 GB		CLSTD	0008
d175	disk		66.8 GB		CLSTD	0008
d161	disk		66.8 GB		CLSTD	0009
d176	disk		66.8 GB		CLSTD	0009
d162	disk		66.8 GB		CLSTD	0010
d177	disk		66.8 GB		CLSTD	0010
d163	disk		66.8 GB		CLSTD	0011

Figure 3. Disk volume list items

The following are examples of disk volume information as shown in Figure 3.

- Disk volumes d160 and d175 are part of Disk Group 8. Disk Group is another term for RAID group.
- Disk Group 9 uses d161 and d176. Disk Group 10 uses d162 and d177.
- Exchange logs will be held on d160 and 175. Exchange databases will be held on d161, d162, d176, and d177.
- SPA owns d160, d161, and d162. SPB owns d175, d176, and d177.

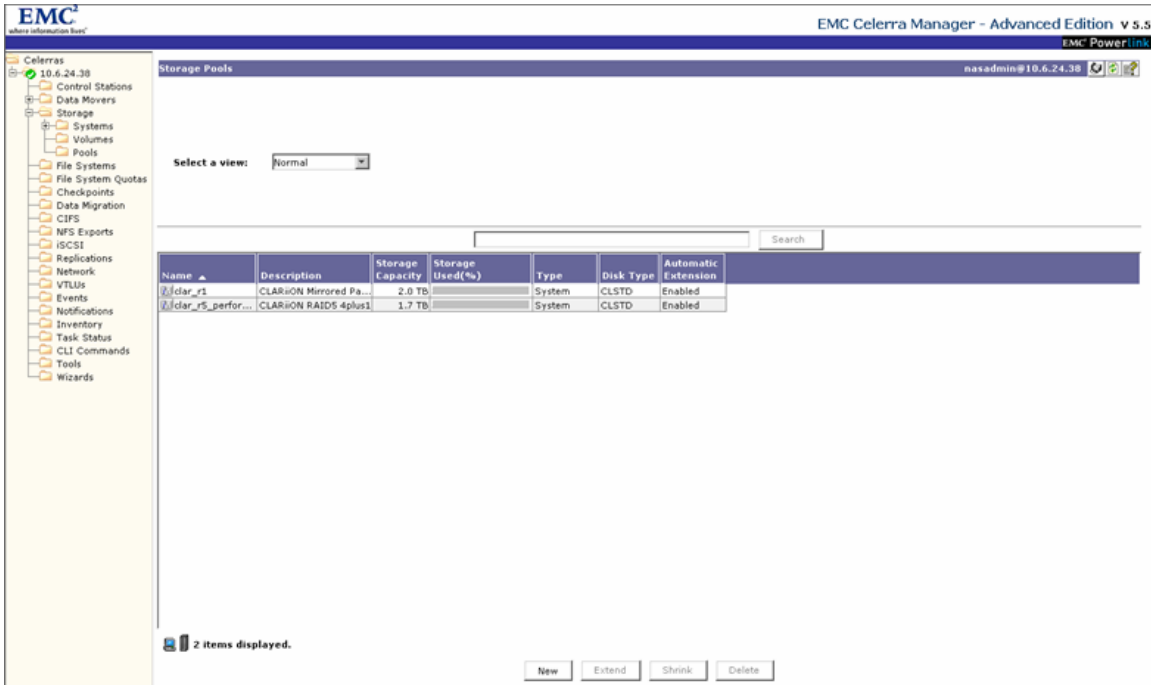


Figure 4. Storage Pools

2. To create a storage pool for the Exchange log file system, from the Celerra Manager > Storage > Pools page, click **New** in the Storage Pools page as shown in Figure 4.
 - In this example, disk volumes d160 and d175 will be added to a storage pool and the file system for the Exchange logs will be created from this storage pool.

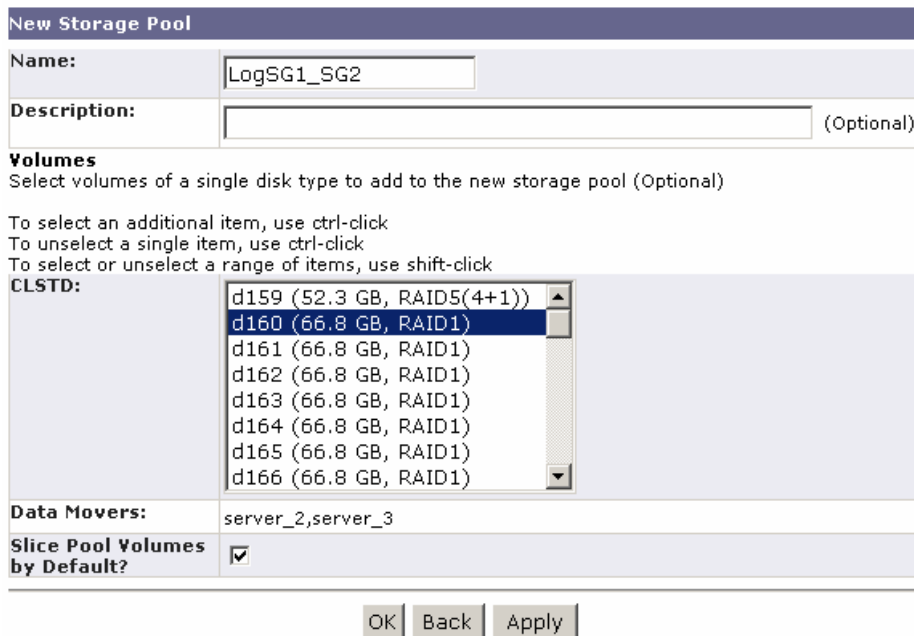


Figure 5. Storage pool creation

3. Enter the name of the storage pool (in this example it is **LogSG1_SG2**), and click **d160** in the **CLSTD** volumes list.

New Storage Pool

Name: LogSG1_SG2

Description: (Optional)

Volumes
 Select volumes of a single disk type to add to the new storage pool (Optional)

To select an additional item, use ctrl-click
 To unselect a single item, use ctrl-click
 To select or unselect a range of items, use shift-click

CLSTD:

- d171 (267.3 GB, RAID5(4+1))
- d172 (267.3 GB, RAID5(4+1))
- d173 (267.3 GB, RAID5(4+1))
- d174 (52.3 GB, RAID5(4+1))
- d175 (66.8 GB, RAID1)**
- d176 (66.8 GB, RAID1)
- d177 (66.8 GB, RAID1)
- d178 (66.8 GB, RAID1)

Data Movers: server_2,server_3

Slice Pool Volumes by Default?

OK Back Apply

Figure 6. Storage pool creation

4. Scroll down the **CLSTD** volumes list and hold the **Ctrl** key while selecting **d175** so that both **d160** and **d175** are highlighted, and then click **OK** as shown in Figure 6.

EMC Celerra Manager - Advanced Edition v 5.2

nasadmin@10.6.24.38

Storage Pools

✓ Create storage pool LogSG1_SG2. OK

Select a view: Normal

Name	Description	Storage Capacity	Storage Used(%)	Type	Disk Type	Automatic Extension
/clar_r1	CLARiiON Mirrored Pa...	1.8 TB		System	CLSTD	Enabled
/clar_r5_perfor...	CLARiiON RAID5 4plus1	1.7 TB		System	CLSTD	Enabled
/LogSG1_SG2		133.6 GB		User	CLSTD	N/A

Figure 7. Storage pool list

The new storage pool LogSG1_SG2 is listed as shown in Figure 7. LogSG1_SG2 will be used to create the file system for the Exchange logs.

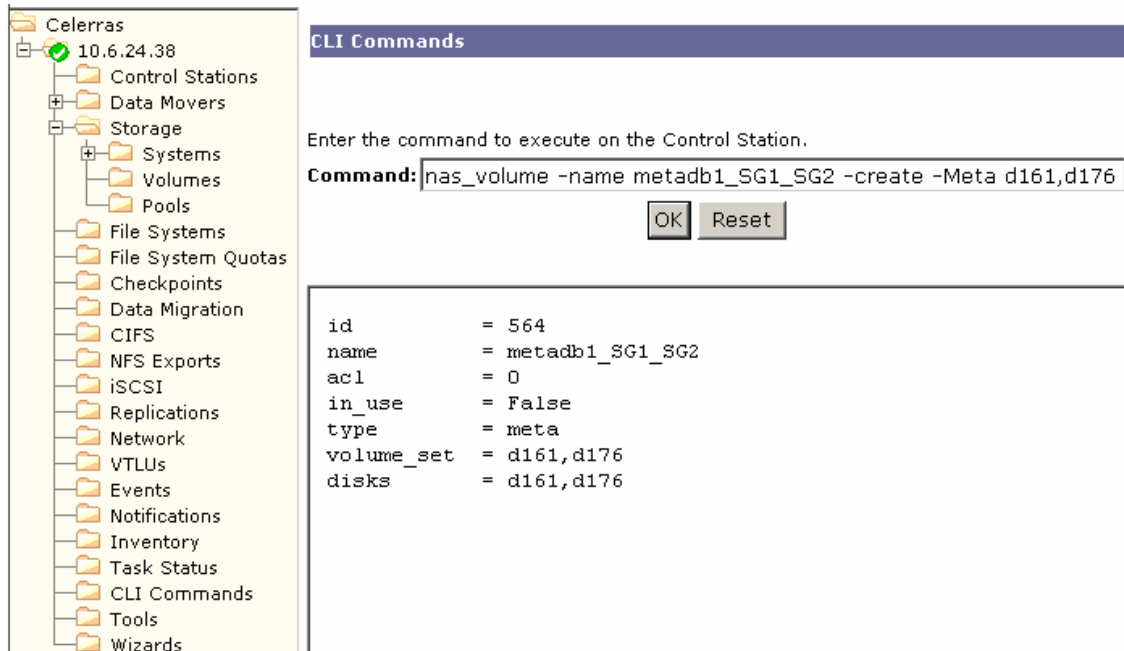


Figure 8. Metavolume creation

In creating the Exchange database file system, note the following two points:

- The Exchange database file system is created from a storage pool made up of a stripe volume created from two metavolumes.
- In creating two metavolumes, it is essential that SPA owns one and SPB owns the other metavolume.

To create the Exchange database file system, perform the following steps:

5. Create two metavolumes (Figure 8).
 - a. Navigate to Celerra Manager > CLI Commands in the tree structure.
 - b. In the **Command** text box, type in the command using your Celerra disk volumes:
nas_volume -name metadb1_SG1_SG2 -create -Meta d161,d176
 - c. Click **OK**.
 - A new metavolume is created.
 - In this example, metavolume metadb1_SG1_SG2 is owned by SPA. This is because of the order of the disk volumes, d161 and d176, which are listed in the `nas_volume` command. SPB owns the other metavolume. SPB owns d177, so d177 will be entered into the command before d162.

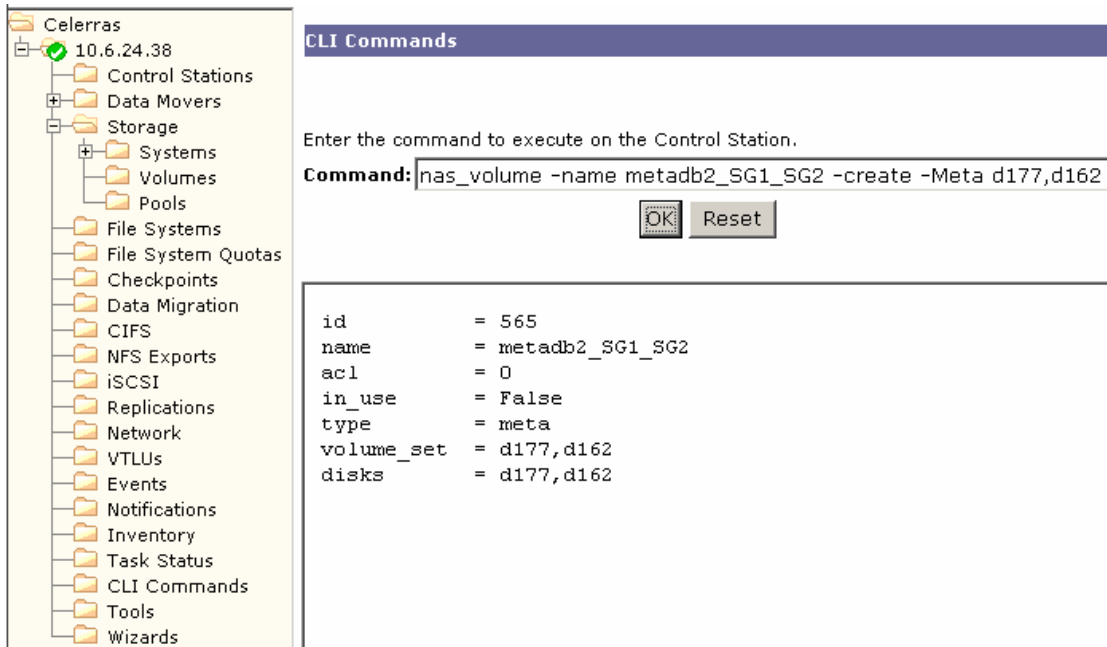


Figure 9. Metavolume creation

- d. In the **Command** text box, type in the command using your Celerra disk volumes (Figure 9):
`nas_volume -name metadb2_SG1_SG2 -create -Meta d177,d162`
- e. Click **OK**.
 - SPA and SPB each own one of the two new metavolumes.
 - Metavolume creation must be performed from the CLI as Celerra Manager does not support specifying the order of disk volumes.

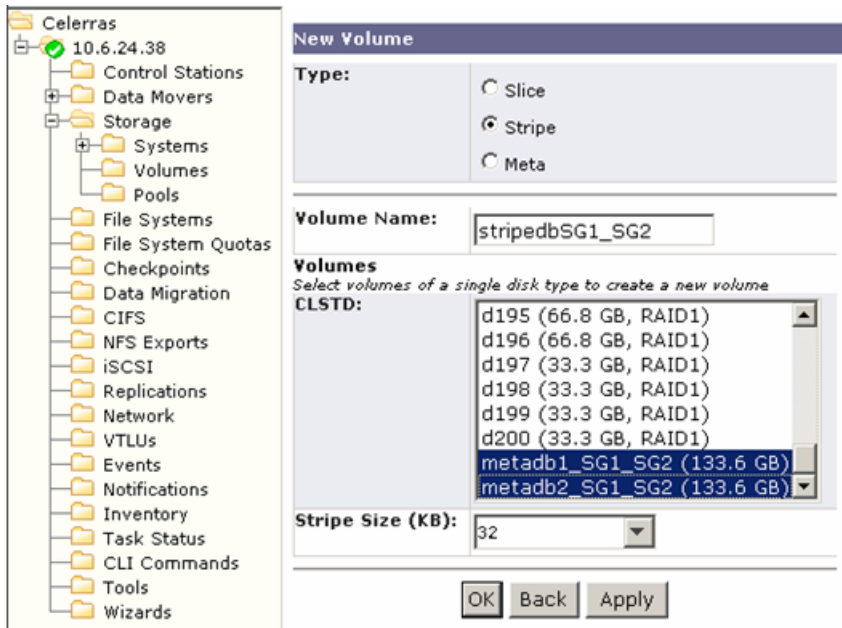


Figure 10. Stripe volume creation

6. Create the stripe volume from the two metavolumes as shown in Figure 10.
 - a. Navigate back to Storage > Volumes, and click **New** located at the bottom of the page.
 - b. In the volume **Type**, select the **Stripe** radio button.
 - c. In the **Volume Name** text box, type the new stripe volume: **stripedbSG1_SG2**
 - d. Scroll to the bottom of the **CLSTD** volumes list and **Ctrl-click** the two new metavolumes that you created: metadb1_SG1_SG2 and metadb1_SG1_SG2.
 - e. Click **OK**.

EMC
where information lives!

EMC Celerra Manager - Advanced Edition v 5.5
EMC Powerlink

nasadmin@10.6.24.30

✓ Create volume stripedSG1_SG2. OK

Show Volumes of Type: All Volumes

Storage Systems: All Systems

Name	Type	Uses Volumes	Used by	Storage Capacity	Storage Used(%)	Disk Type	Storage Systems
d190	disk			66.8 GB		CLSTD	CLARIION CX3-40 APM00064501r
d191	disk			33.3 GB		CLSTD	CLARIION CX3-40 APM00064501r
d192	disk			33.3 GB		CLSTD	CLARIION CX3-40 APM00064501r
d193	disk			33.3 GB		CLSTD	CLARIION CX3-40 APM00064501r
d194	disk			33.3 GB		CLSTD	CLARIION CX3-40 APM00064501r
d195	disk			66.8 GB		CLSTD	CLARIION CX3-40 APM00064501r
d196	disk			66.8 GB		CLSTD	CLARIION CX3-40 APM00064501r
d197	disk			33.3 GB		CLSTD	CLARIION CX3-40 APM00064501r
d198	disk			33.3 GB		CLSTD	CLARIION CX3-40 APM00064501r
d199	disk			33.3 GB		CLSTD	CLARIION CX3-40 APM00064501r
d200	disk			33.3 GB		CLSTD	CLARIION CX3-40 APM00064501r
md3	meta	d3	root_fs_d3	2.0 GB		CLSTD	CLARIION CX3-40 APM00064501r
md4	meta	d4	root_fs_d4	2.0 GB		CLSTD	CLARIION CX3-40 APM00064501r
md5	meta	d5	root_fs_d5	2.0 GB		CLSTD	CLARIION CX3-40 APM00064501r
md6	meta	d6	root_fs_d6	2.0 GB		CLSTD	CLARIION CX3-40 APM00064501r
metadb...	meta	d161 d176	stripedbSG1_SG2	133.6 GB		CLSTD	CLARIION CX3-40 APM00064501r
metadb...	meta	d177 d162	stripedbSG1_SG2	133.6 GB		CLSTD	CLARIION CX3-40 APM00064501r
stripedb...	stripe	metadb1_SG1_SG2 metadb2_SG1_SG2		267.3 GB		CLSTD	CLARIION CX3-40 APM00064501r

53 items displayed.

New Delete

Figure 11. Stripe volume list

The new stripe volume, stripedbSG1_SG2, is listed as shown in Figure 11.

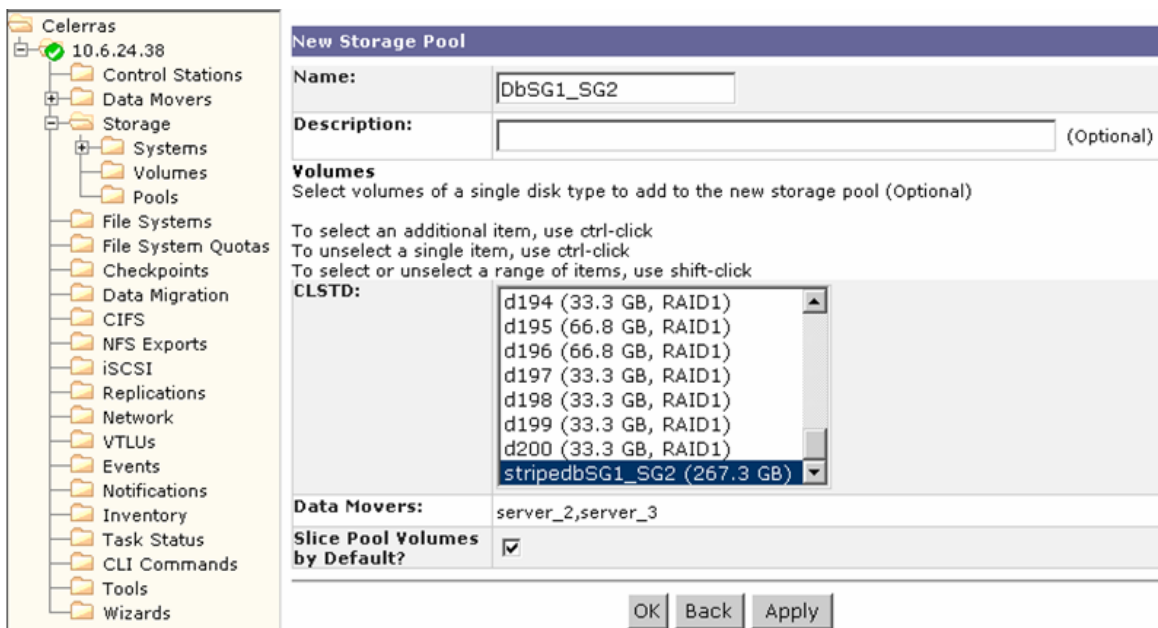


Figure 12. Storage pool creation

7. Create the storage pool from your new stripe volume (Figure 12).
 - a. From the Celerra Manager > Storage > Pools page, click **New**.
 - b. Type the new storage pool in the **Name** text box; in this example it is DbSG1_SG2.
 - c. Scroll to the bottom of the **CLSTD** volumes list and select stripedbSG1_SG2.
 - d. Click **OK**.

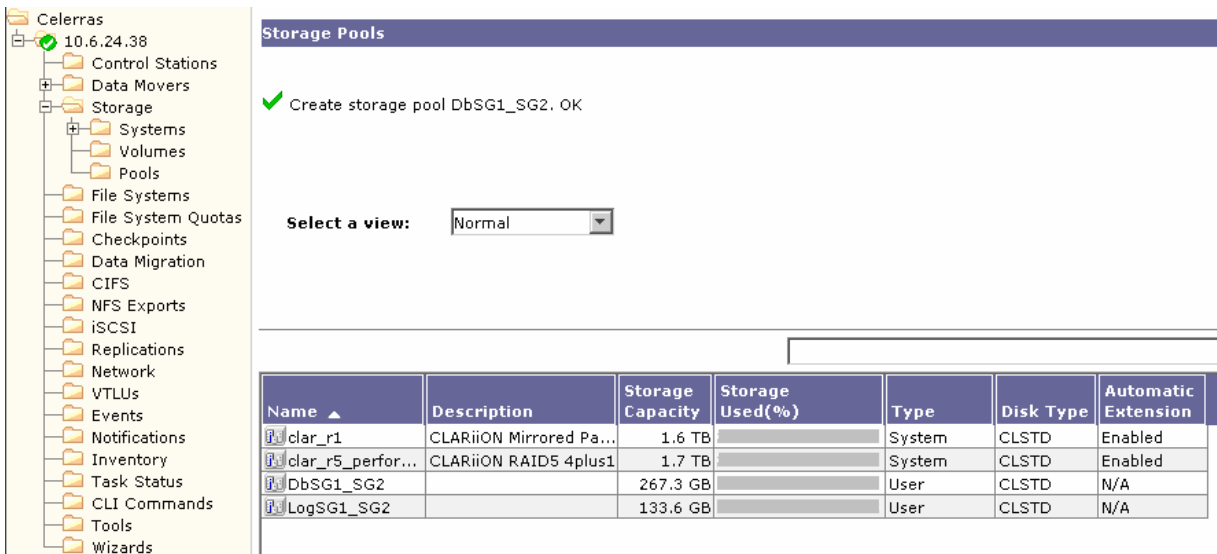


Figure 13. The two storage pools are listed

The new storage pool DbSG1_SG2 is created as shown in Figure 13.

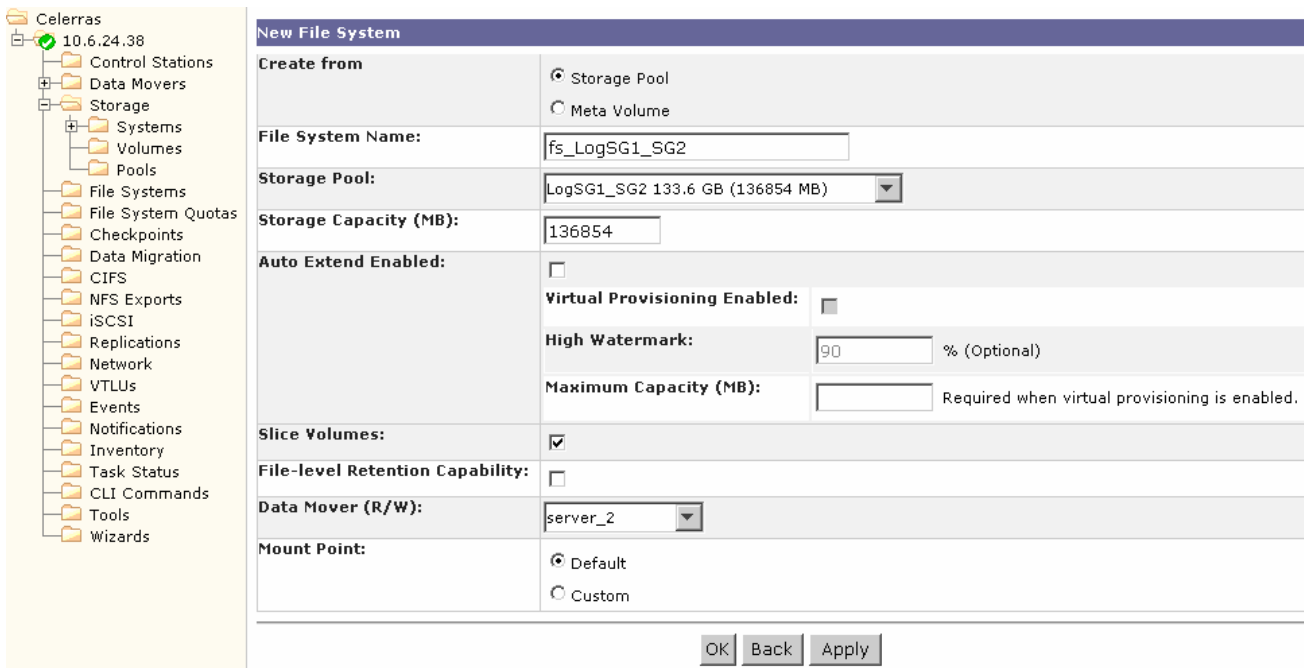


Figure 14. File system creation

8. Create the Exchange log file system using the log storage pool (Figure 14).
 - a. From the Celerra Manager > File Systems page, click **New**.
 - b. Enter the file system name in the **File System Name** text box.
 - c. Select a storage pool from the **Storage Pool** list box.
 - d. Enter the MB size for the file system in the **Storage Capacity (MB)** text box.
 - e. Click **OK**.

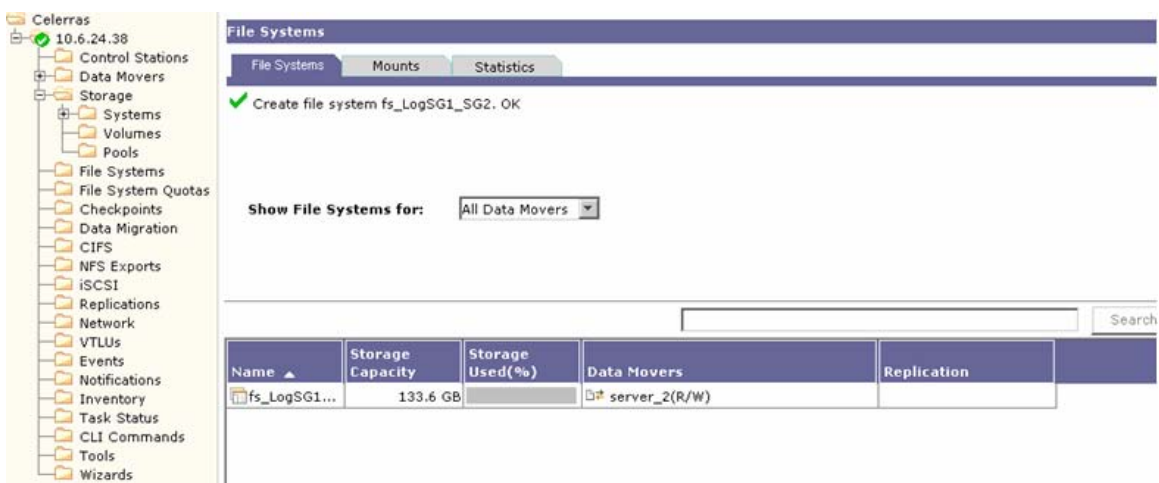


Figure 15. File system list

The new file system fs_LogSG1_SG2 is created as shown in Figure 15.

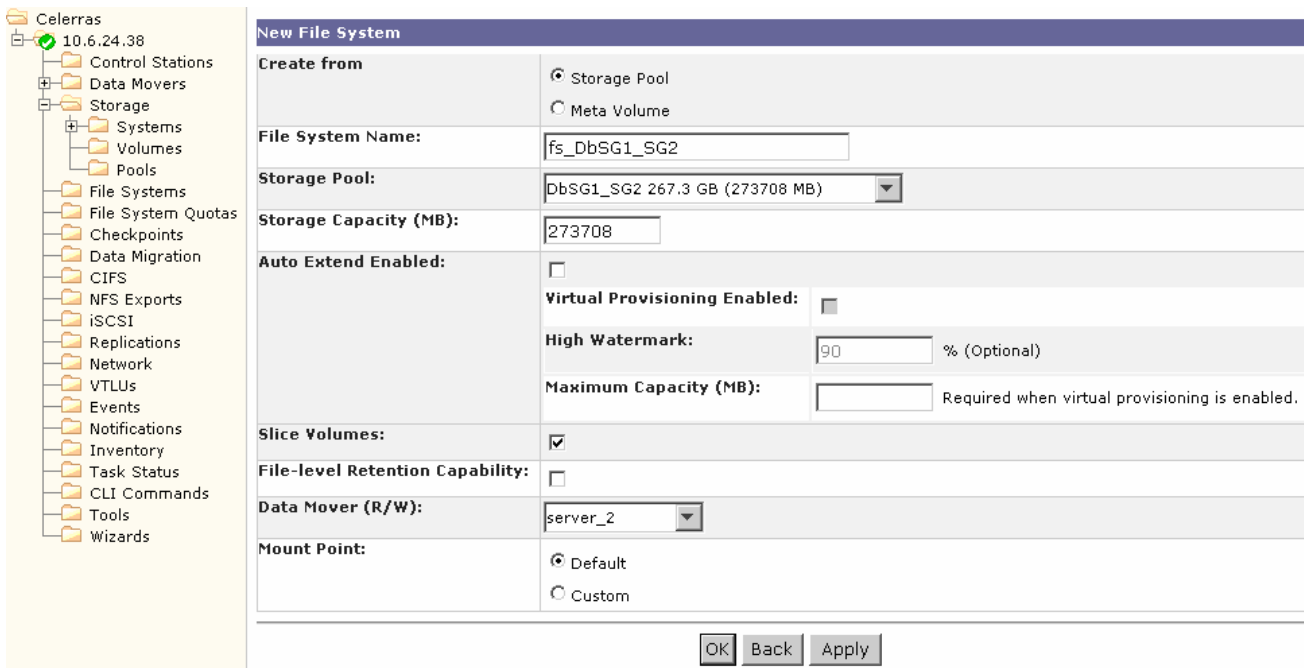


Figure 16. Create second file system

9. Create the Exchange database file system using the database storage pool (Figure 16).
 - a. Click **New** (on the File Systems page) to create the database file system.
 - b. Enter the file system name in the **File System Name** text box.
 - c. Select a storage pool from the **Storage Pool** list box.
 - d. Enter the MB size for the file system in the **Storage Capacity (MB)** text box.
 - e. Click **OK**.

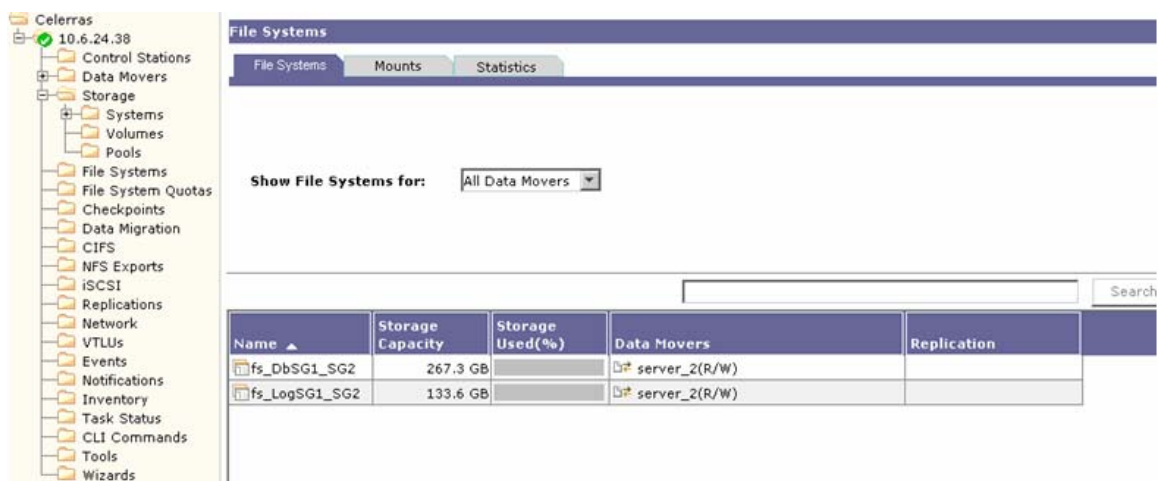


Figure 17. File system list

The two new file systems are listed. The iSCSI LUNs will be created from these file systems.

iSCSI LUN creation

This section describes the steps performed to create iSCSI LUNs from the file systems created on Celerra. Two Exchange storage groups are used for user mailboxes. Two iSCSI LUNs are needed for the database and logs of each storage group. A total of four iSCSI LUNs will be created for this configuration. It is important to first create the iSCSI target, and then create four iSCSI LUNs.

1. From Celerra Manager > Wizards, click **Create an iSCSI Target** (Figure 18).



Figure 18. Celerra Manager Wizards

2. From the Wizard pop-up window, click **Next**.

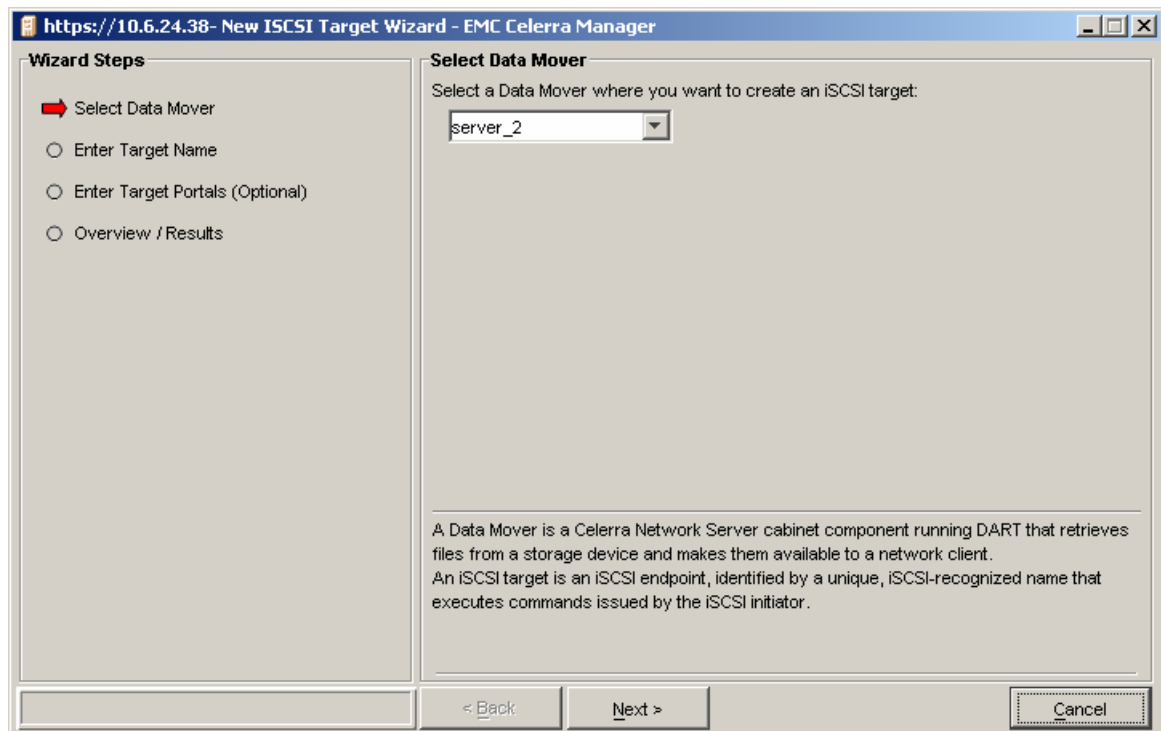


Figure 19. iSCSI Target Wizard, Select Data Mover

3. Type in a target alias name in the **Enter Target Alias Name** text box, and click **Next**.

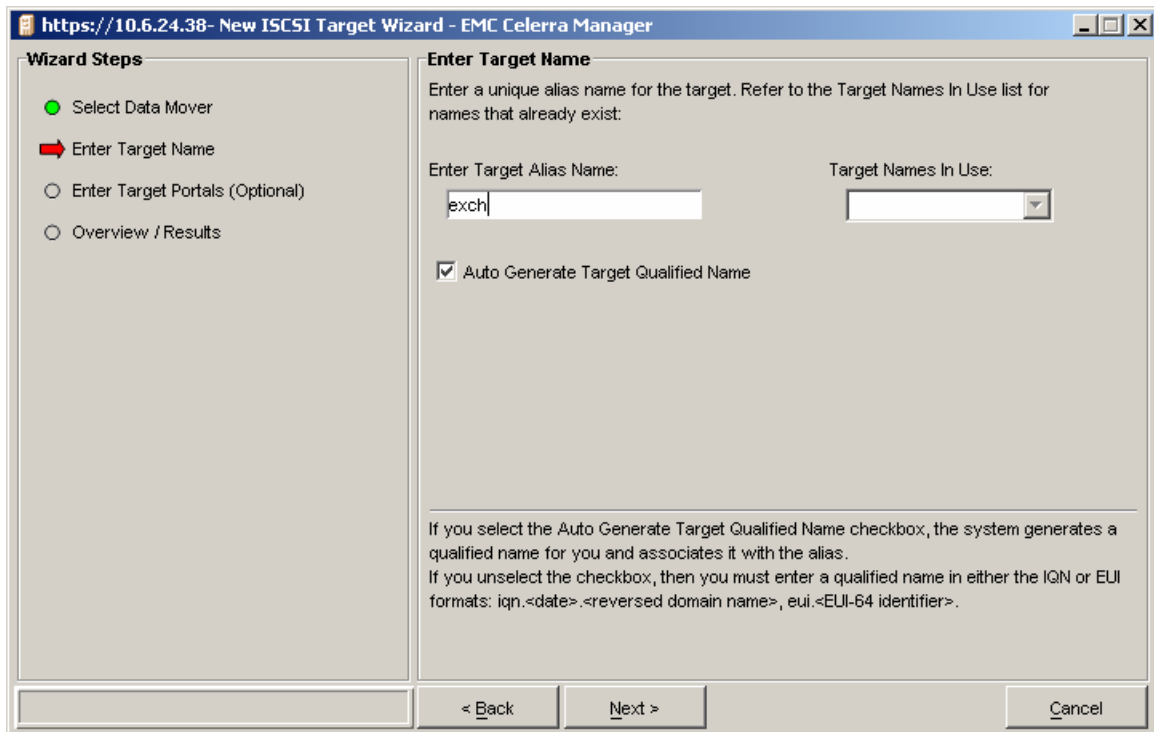


Figure 20. iSCSI Target Wizard, Enter Target Alias Name

4. Select the interfaces that you want to use for iSCSI target portals as shown in Figure 21.

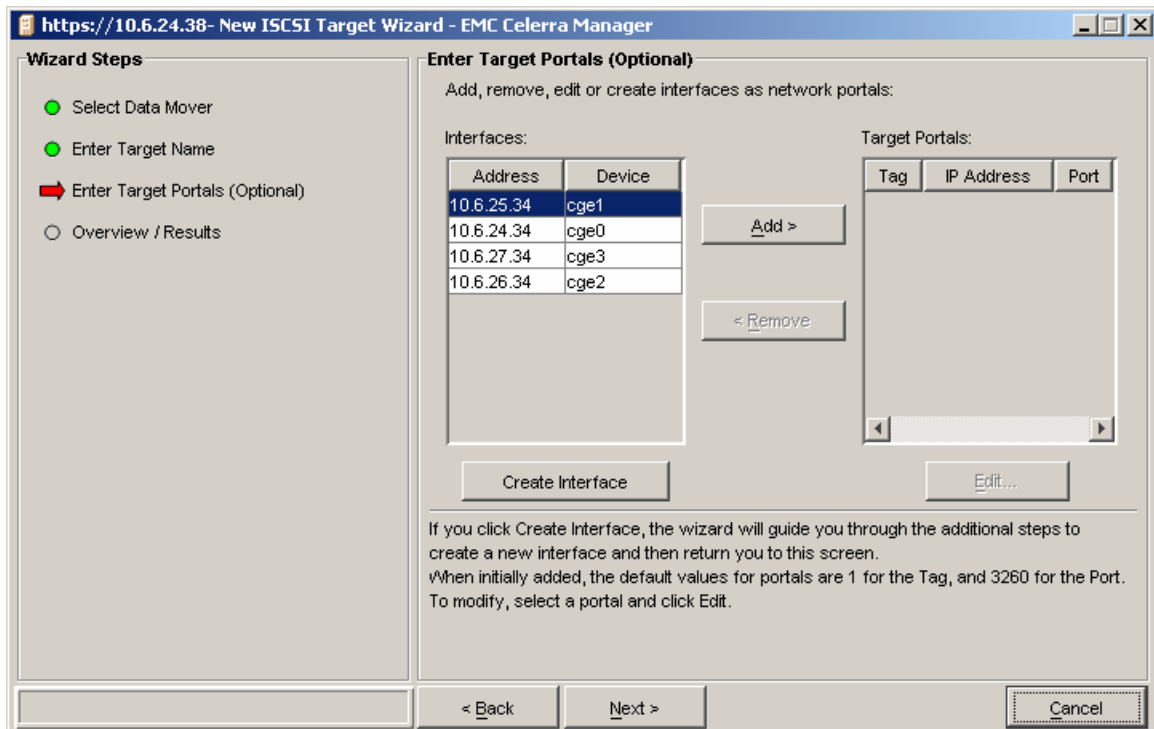


Figure 21. iSCSI Target Wizard, Enter Target Portals

5. Click **Add** to make your selections, and then click **Next** as shown in Figure 22.

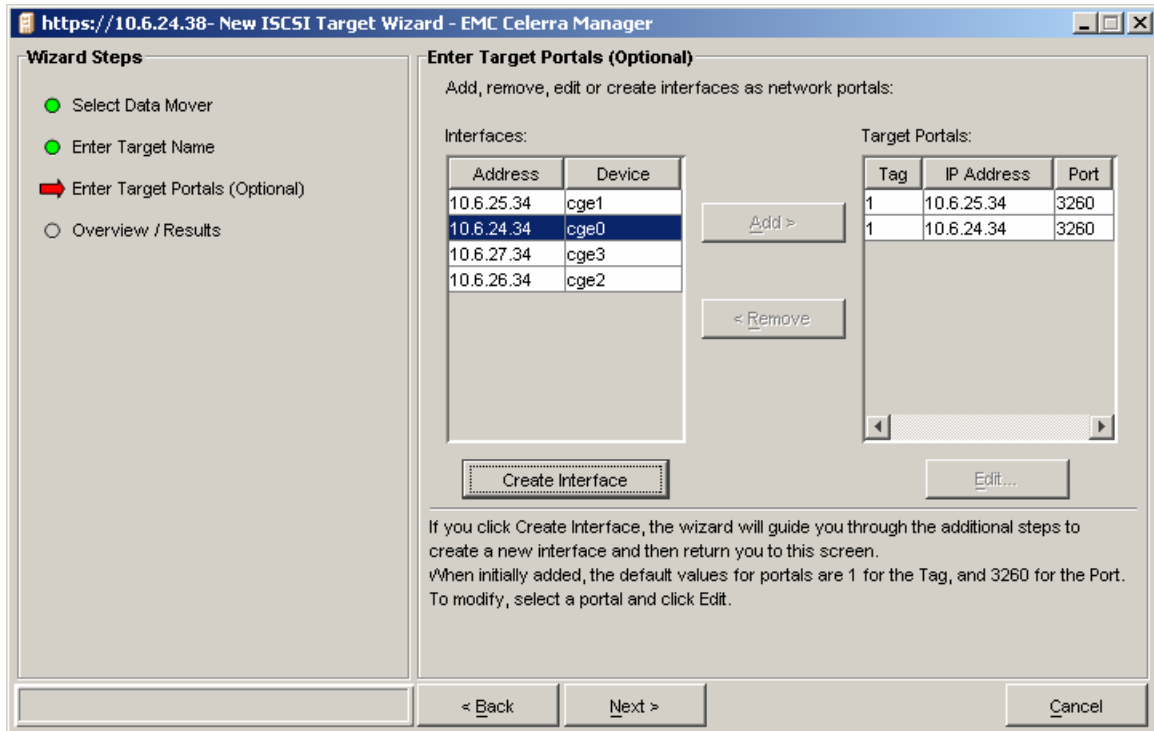


Figure 22. iSCSI Target Wizard, Enter Target Portals

6. Click **Finish**.

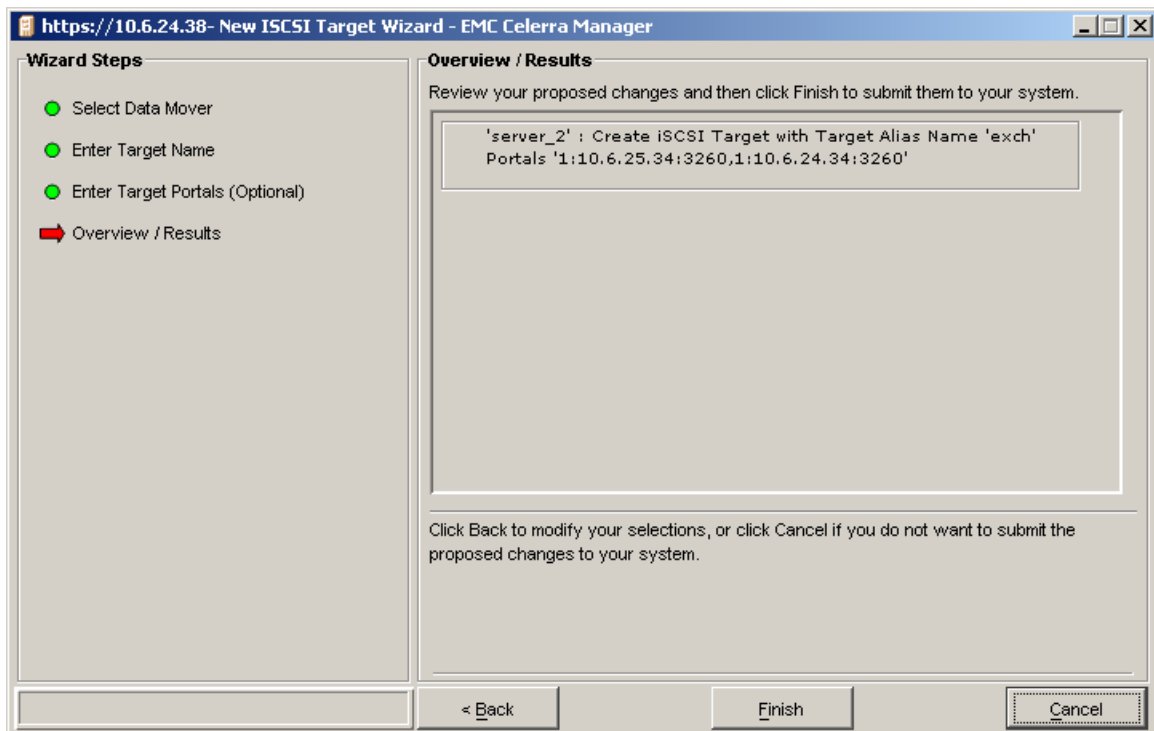


Figure 23. iSCSI Target Wizard, Results

7. Click **Close** (Figure 24).

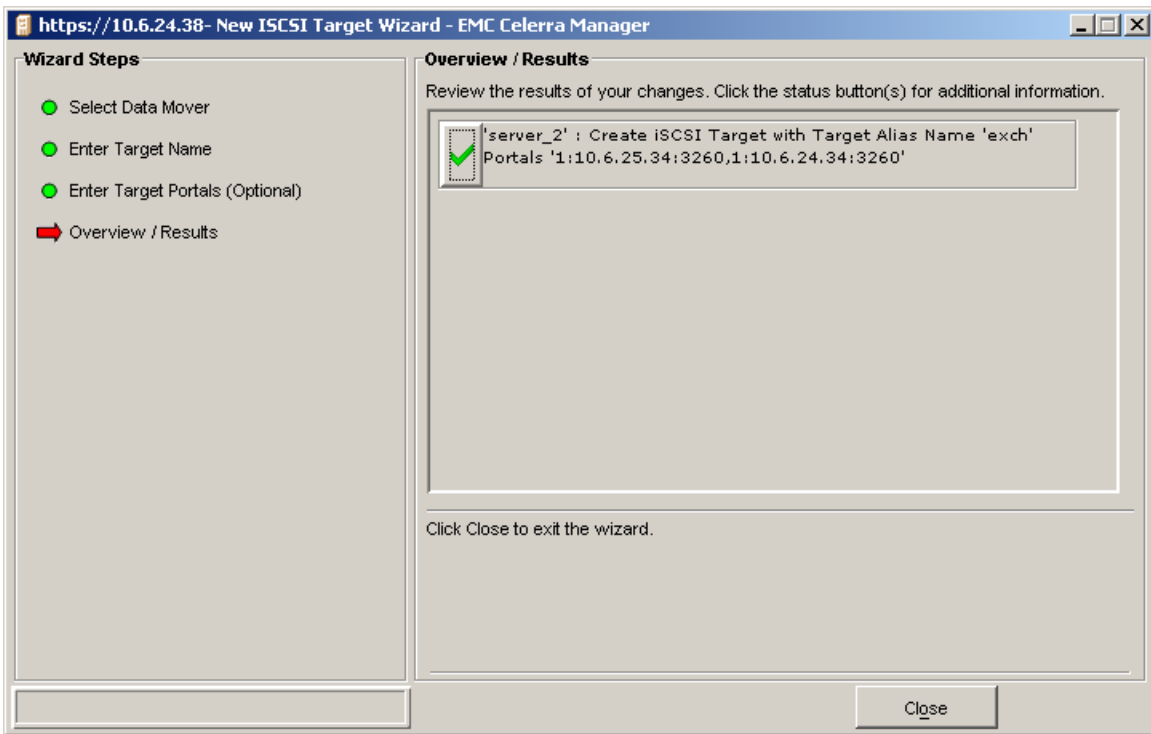


Figure 24. iSCSI Target Wizard, Results

8. From the Celerra Manager > Wizards, click **Create an iSCSI Lun**.



Figure 25. Celerra Manager Wizards

9. From the Wizard pop-up window, click **Next** as shown in Figure 26.

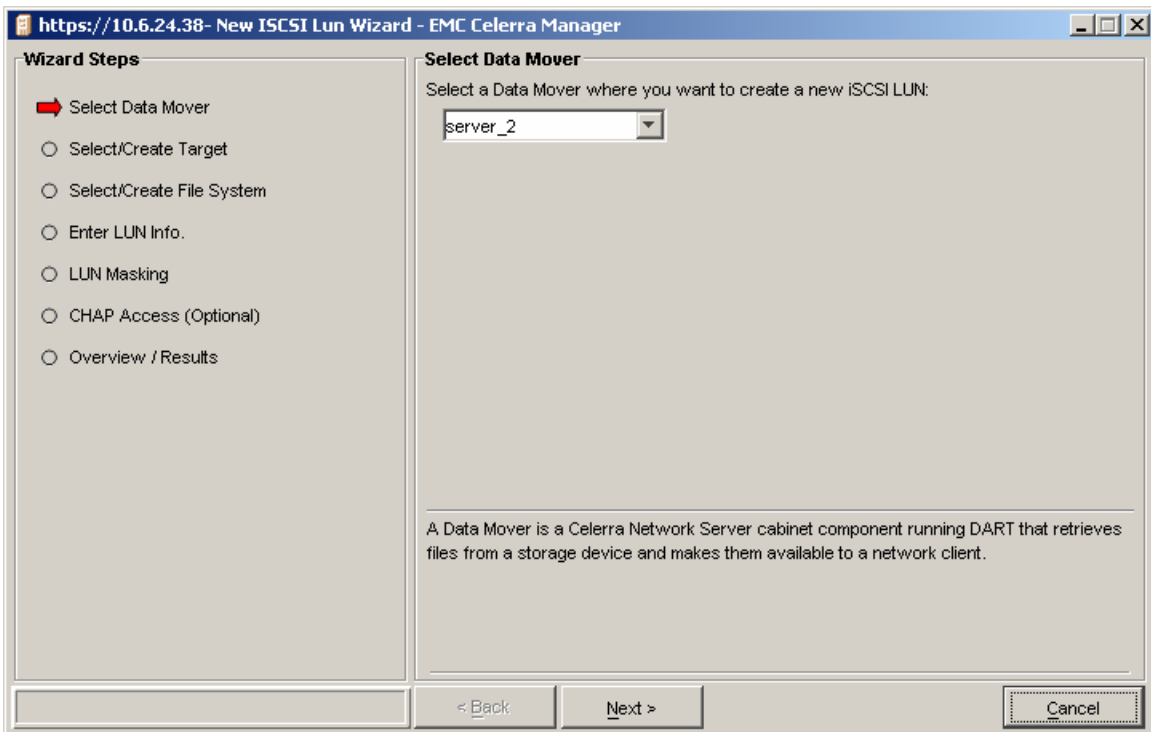


Figure 26. iSCSI Lun Wizard, Select Data Mover

10. Click **Next**.

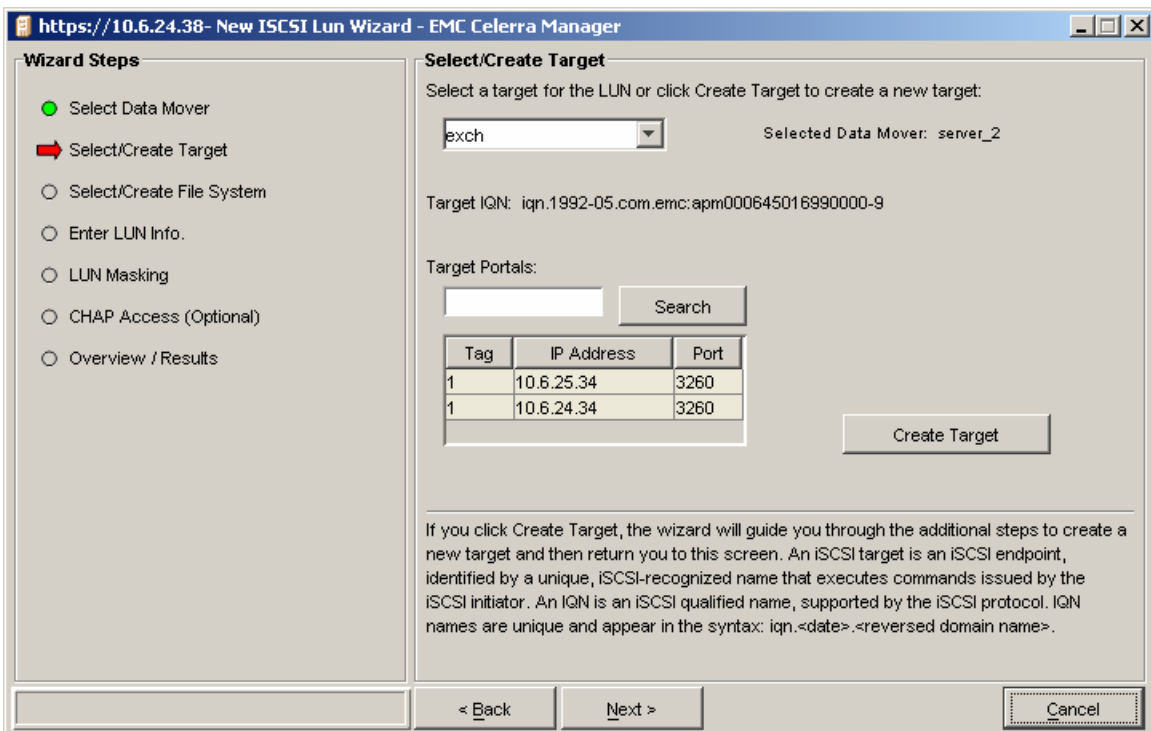


Figure 27. iSCSI Lun Wizard, Select Target

11. Select a file system (log file system), and click **Next** as shown in Figure 28.

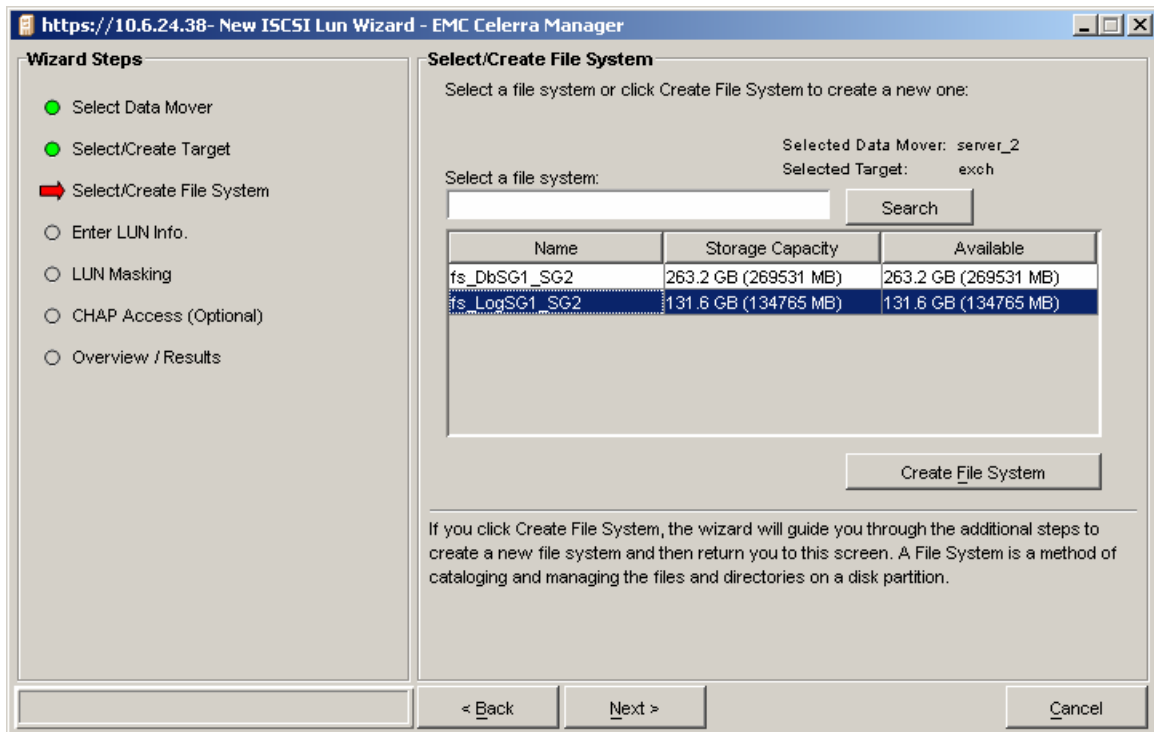


Figure 28. iSCSI Lun Wizard, Select File System

12. On the **Enter LUN Info.** section in Figure 29, do the following:
 - a. Select the checkbox **Create Multiple LUNs**.
 - b. Enter the number of LUNs to create.
 - c. Enter the size of each LUN in MB.
 - d. Click **Next**.

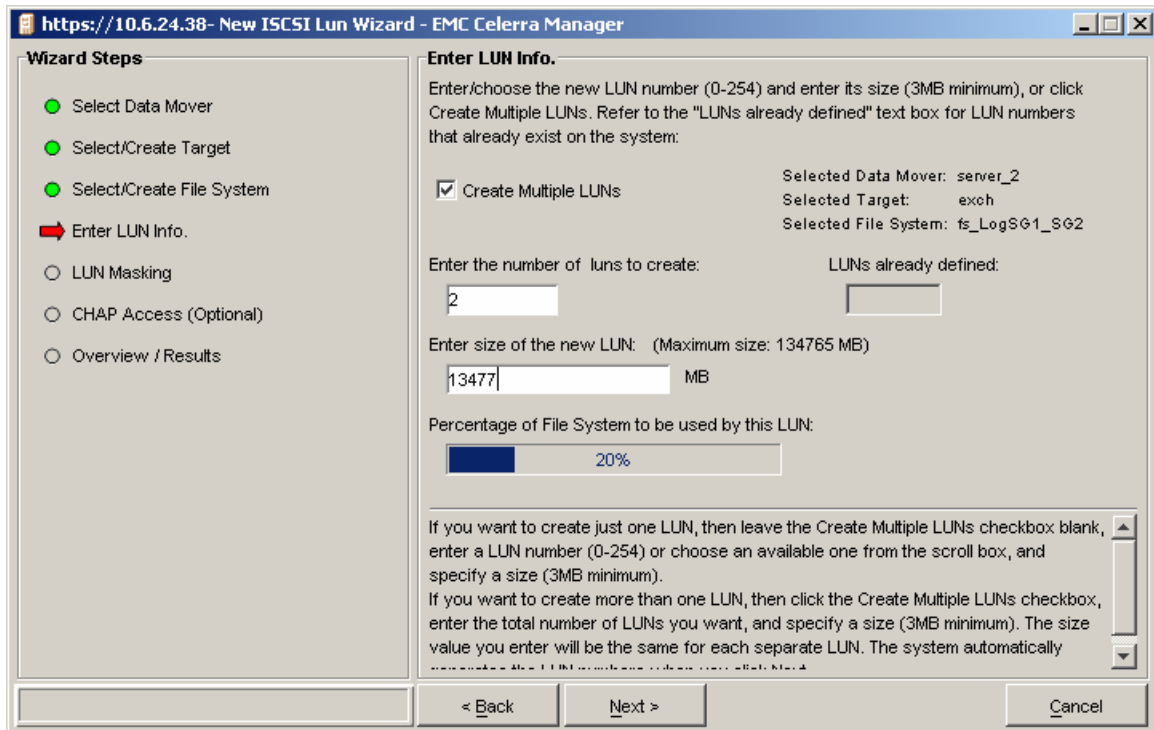


Figure 29. iSCSI Lun Wizard, Enter LUN Info.

13. Click **Add New**.

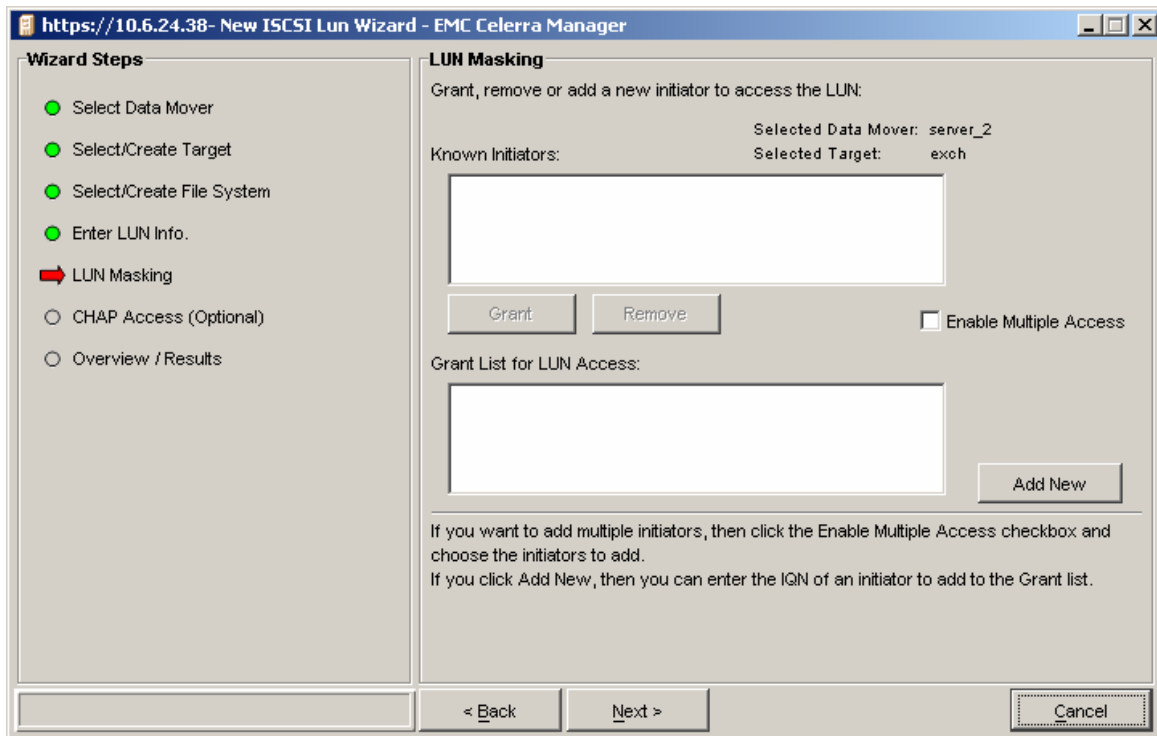


Figure 30. iSCSI Lun Wizard, LUN Masking

14. In the Add New Initiator window, type the new initiator qualified name in the **Enter New Initiator Qualified Name** textbox.

- You can find the initiator qualified name by navigating in Celerra Manager to iSCSI and then click the tab **Target**.

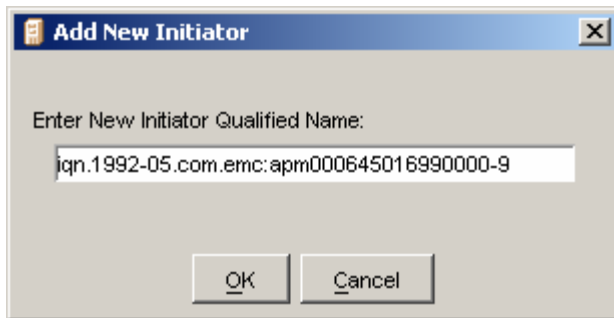


Figure 31. iSCSI Lun Wizard, Add New Initiator

15. Click **Next** as shown in Figure 32.

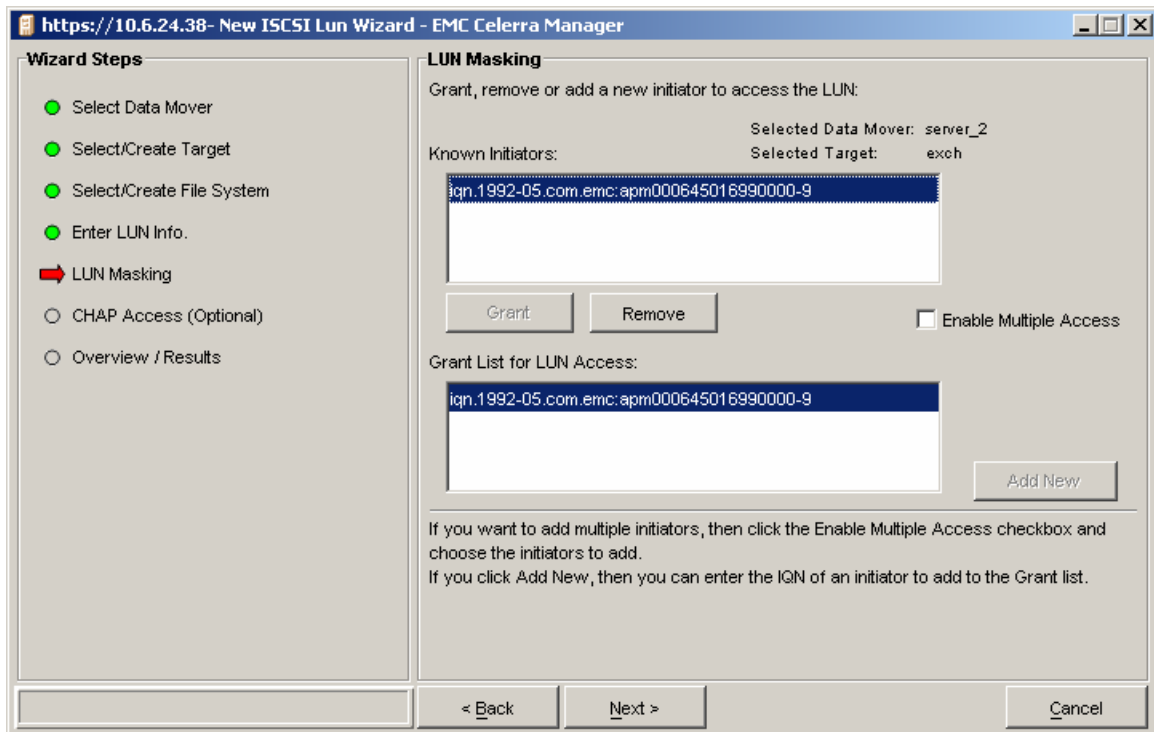


Figure 32. iSCSI Lun Wizard, LUN Masking

16. The **CHAP Access (Optional)** section appears, and click **Next**.

- **NOTE:** If you want to use CHAP access, you can configure it as this time before clicking **Next**.

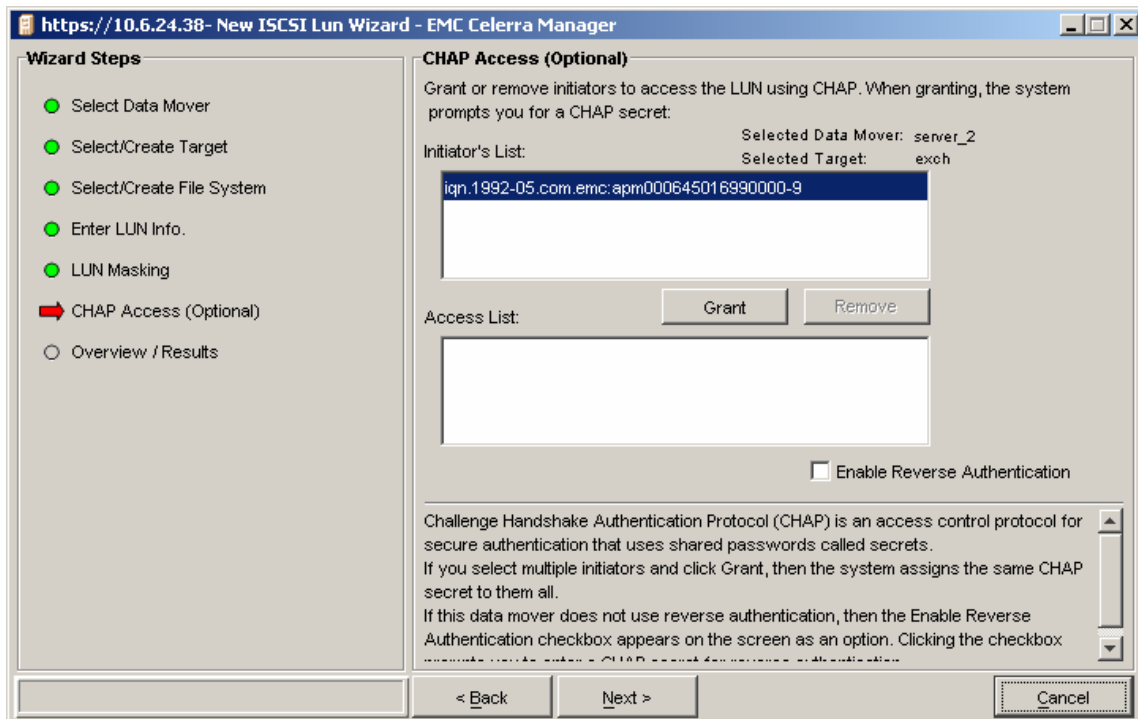


Figure 33. iSCSI Lun Wizard, CHAP Access

17. Click **Finish** (Figure 34).

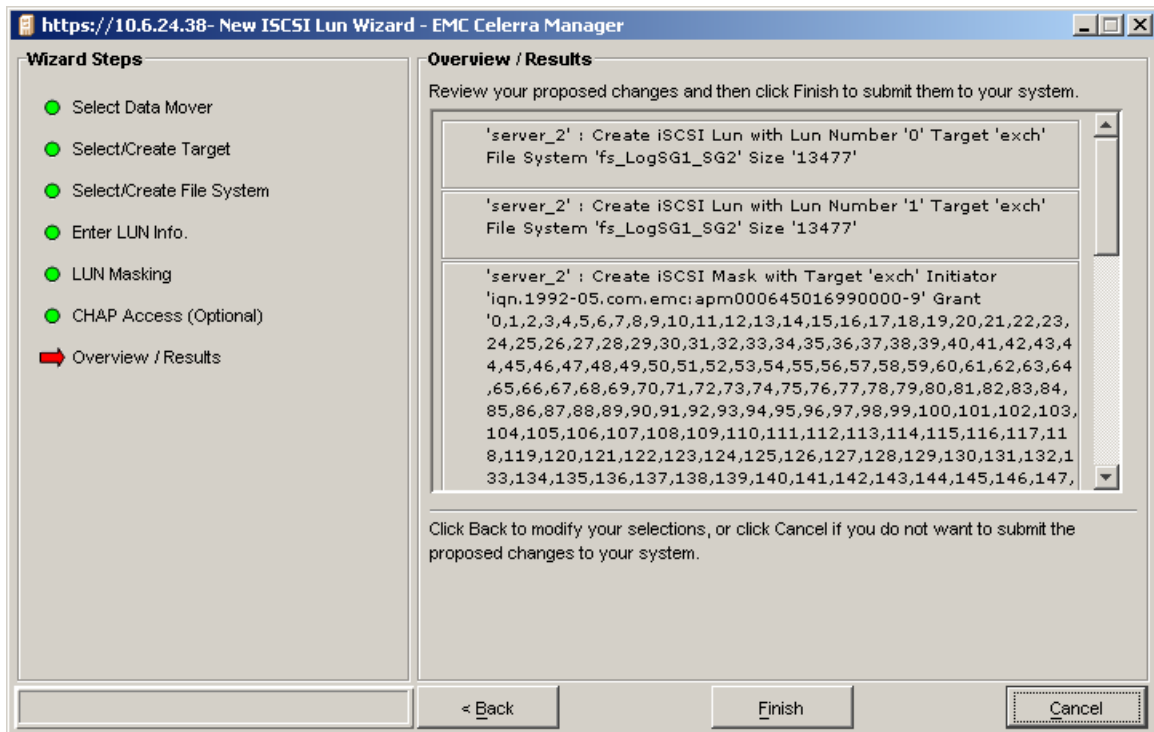


Figure 34. iSCSI Lun Wizard, Results

18. Click **Close**.

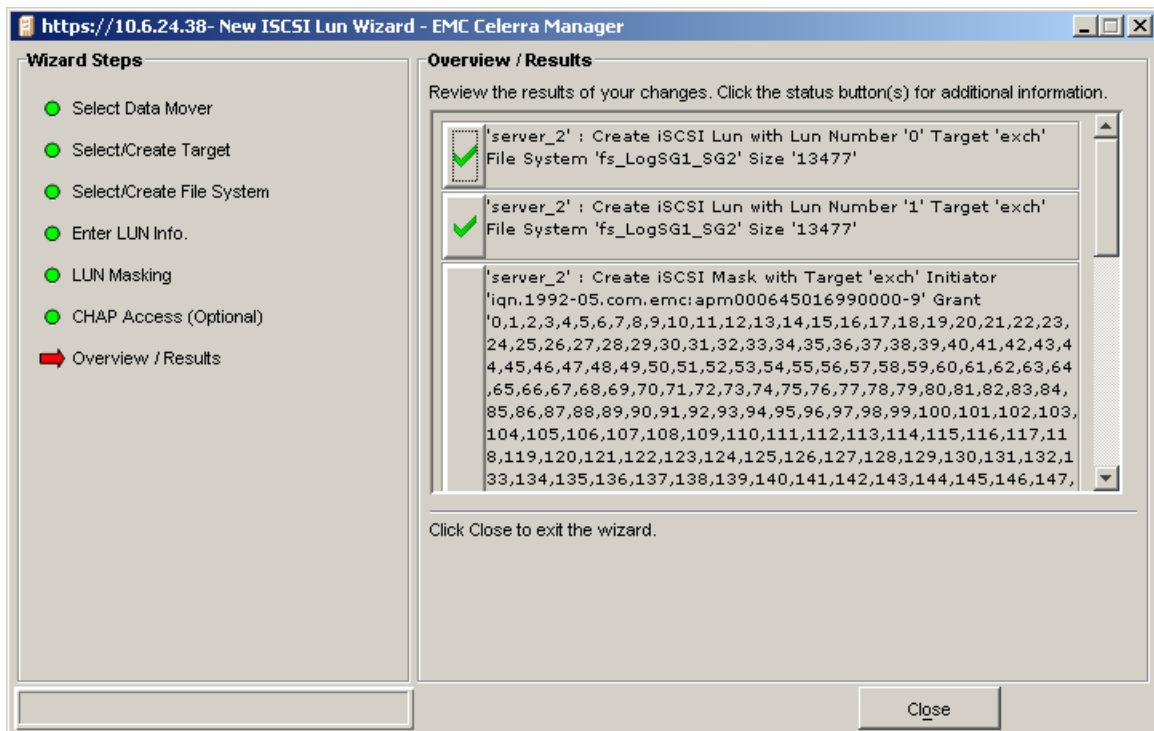


Figure 35. iSCSI Lun Wizard, Results

19. After you observe that the two iSCSI LUNs for the log file systems are listed in the **LUNs** tab, go to **Wizards** and start the **Create an iSCSI Lun** wizard to create the database iSCSI LUNs.

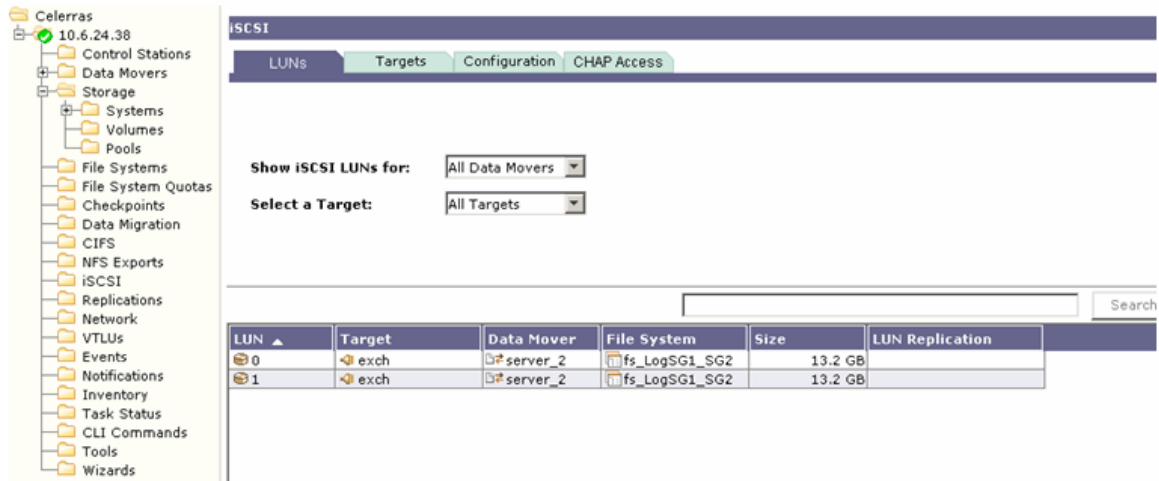


Figure 36. Celerra Manger Wizards

20. Click **Next**.

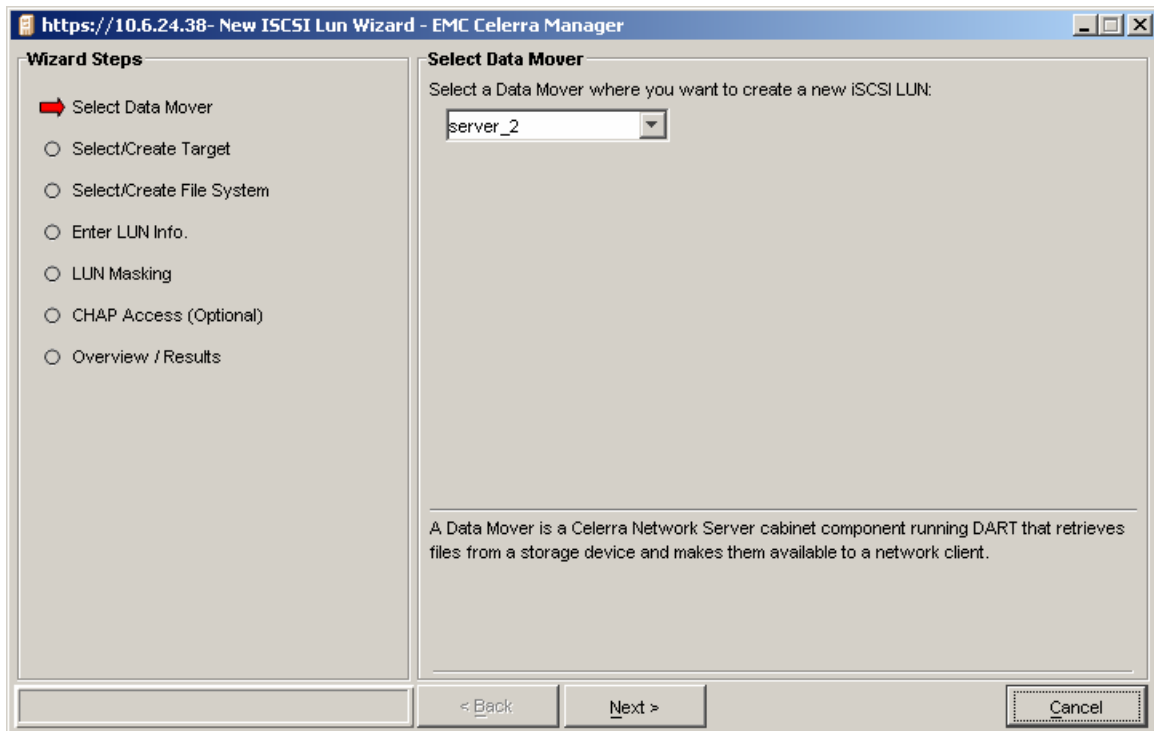


Figure 37. iSCSI Lun Wizard, Select Data Mover

21. Click Next.

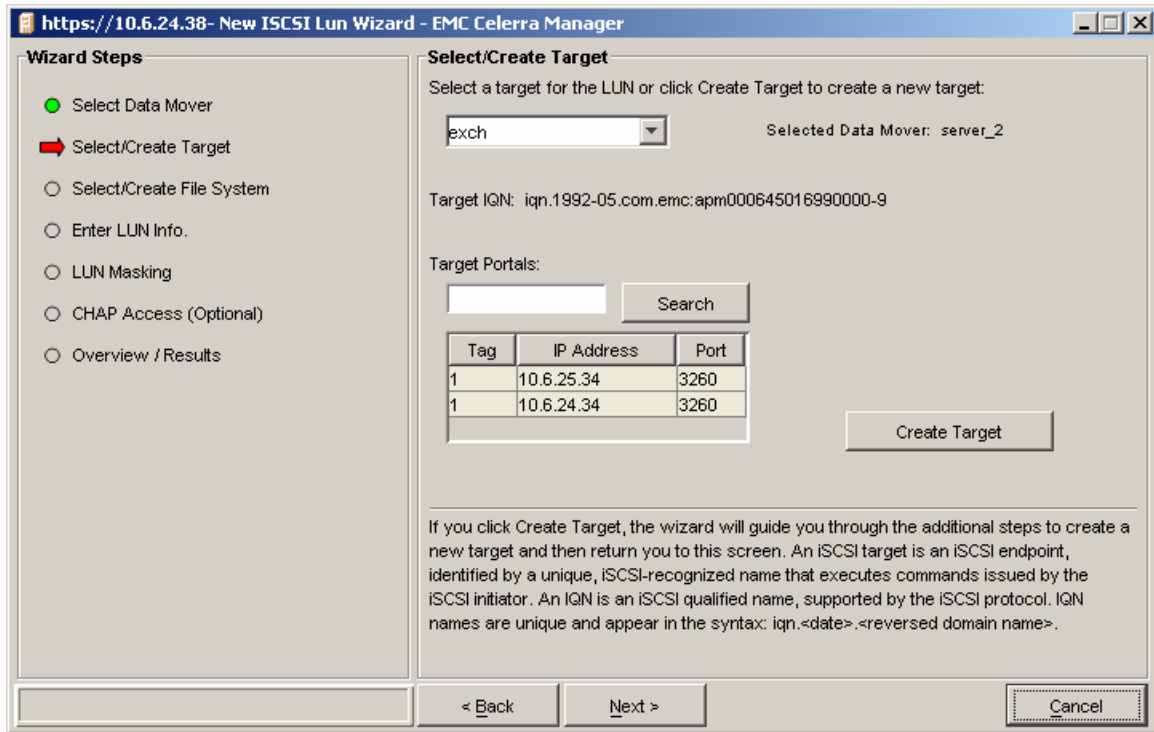


Figure 38. iSCSI Lun Wizard, Select Target

22. Select the file system for these iSCSI LUNs (the database file system as shown in Figure 39), and click Next.

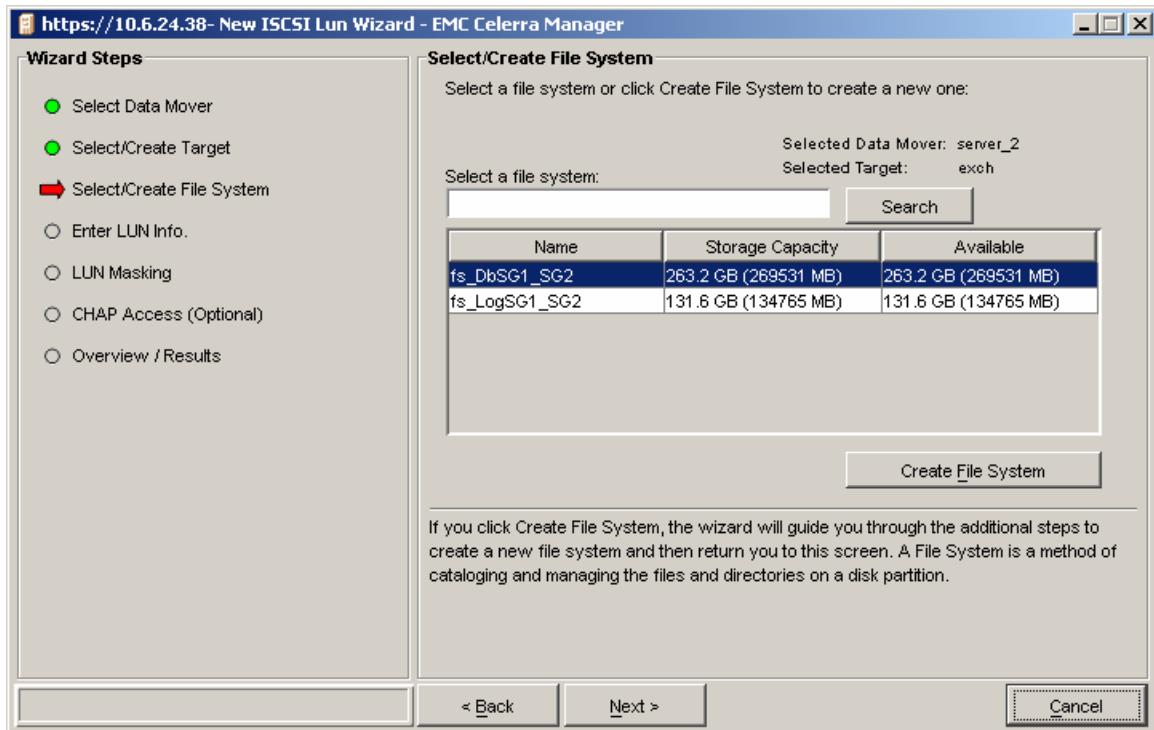


Figure 39. iSCSI Lun Wizard, Select File System

23. On the **Enter LUN Info.** section, do the following as shown in Figure 40:

- e. Select the checkbox **Create Multiple LUNs**.
- f. Enter the number of LUNs to create.
- g. Enter the size of each LUN in MB.
- h. Click **Next**.

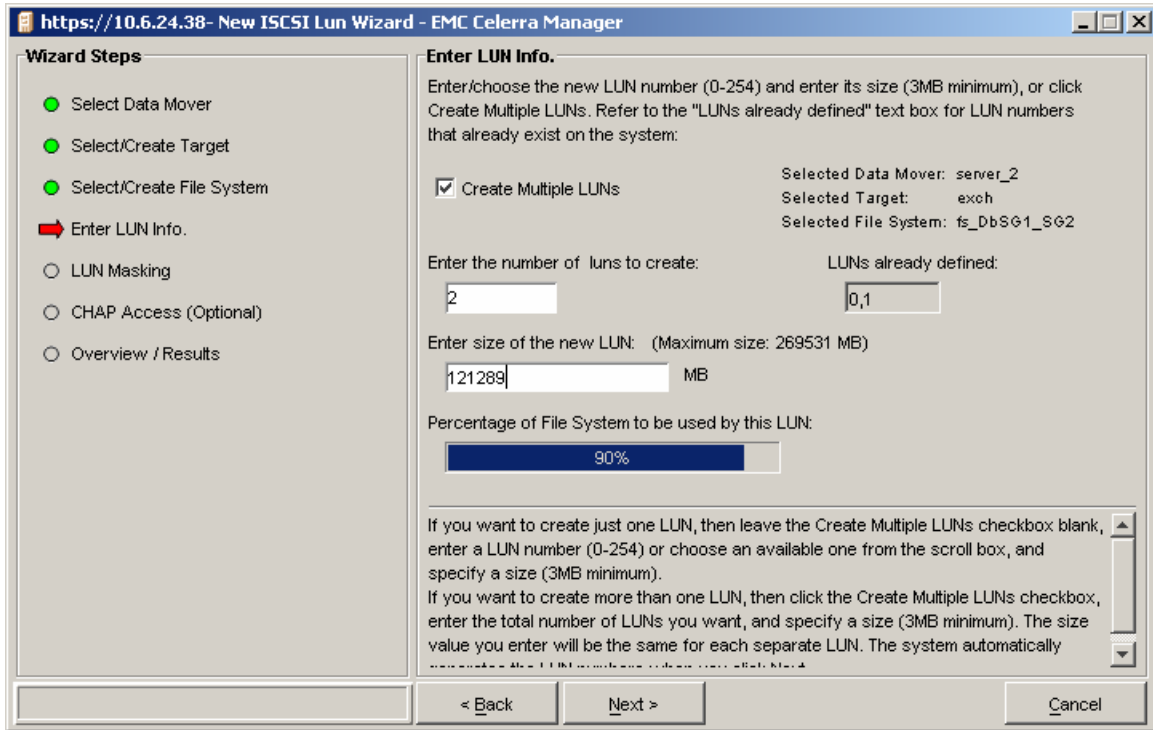


Figure 40. iSCSI Lun Wizard, Enter LUN Info.

24. Click Next.

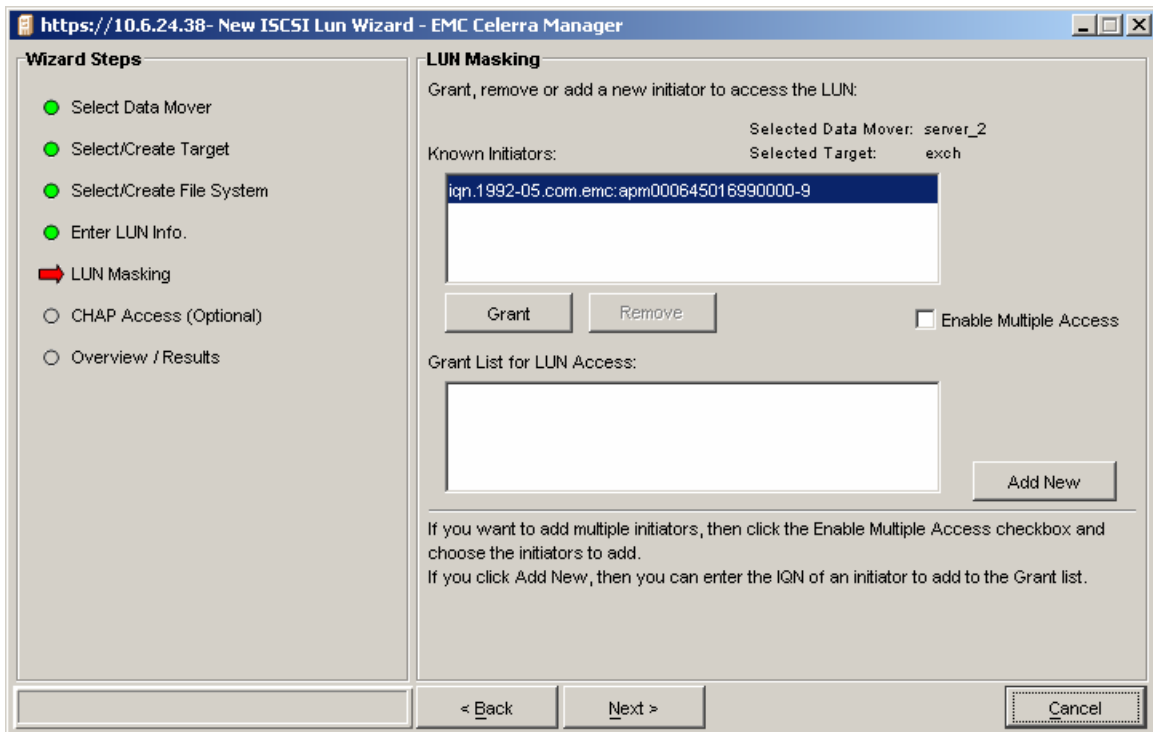


Figure 41. iSCSI Lun Wizard, LUN Masking

25. Click **Next**.

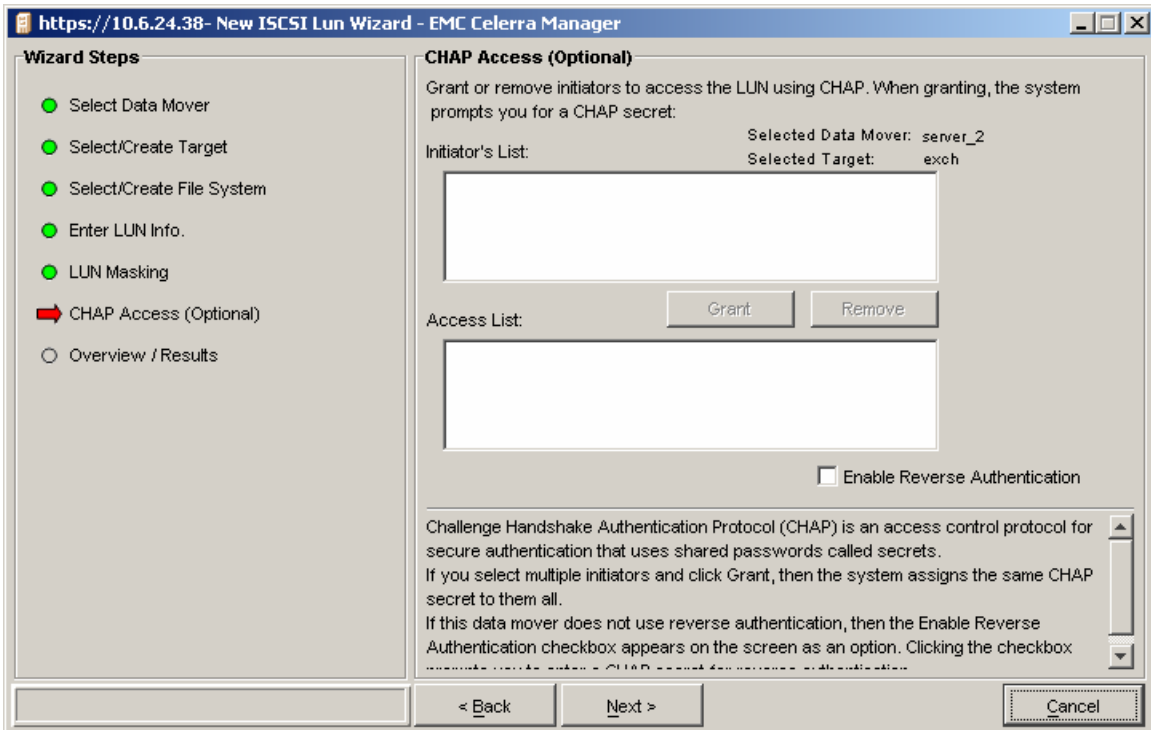


Figure 42. iSCSI Lun Wizard, CHAP Access

26. Click **Finish**.

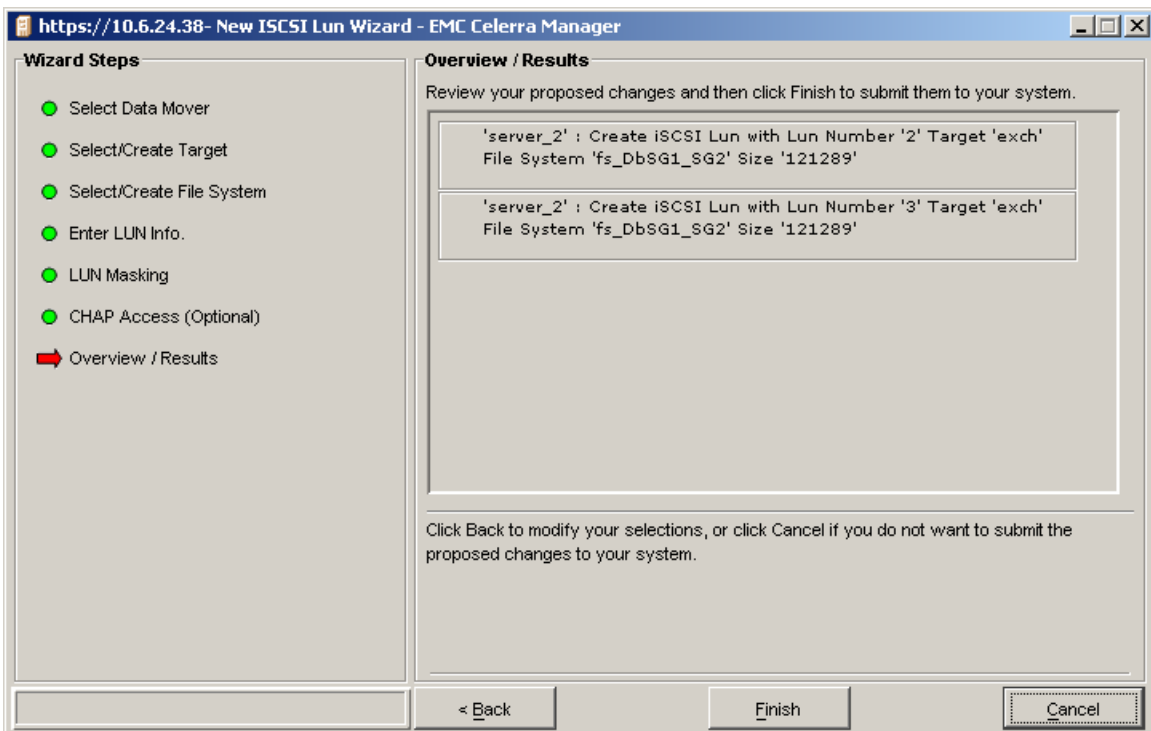


Figure 43. iSCSI Lun Wizard, Results

27. Click **Close**.

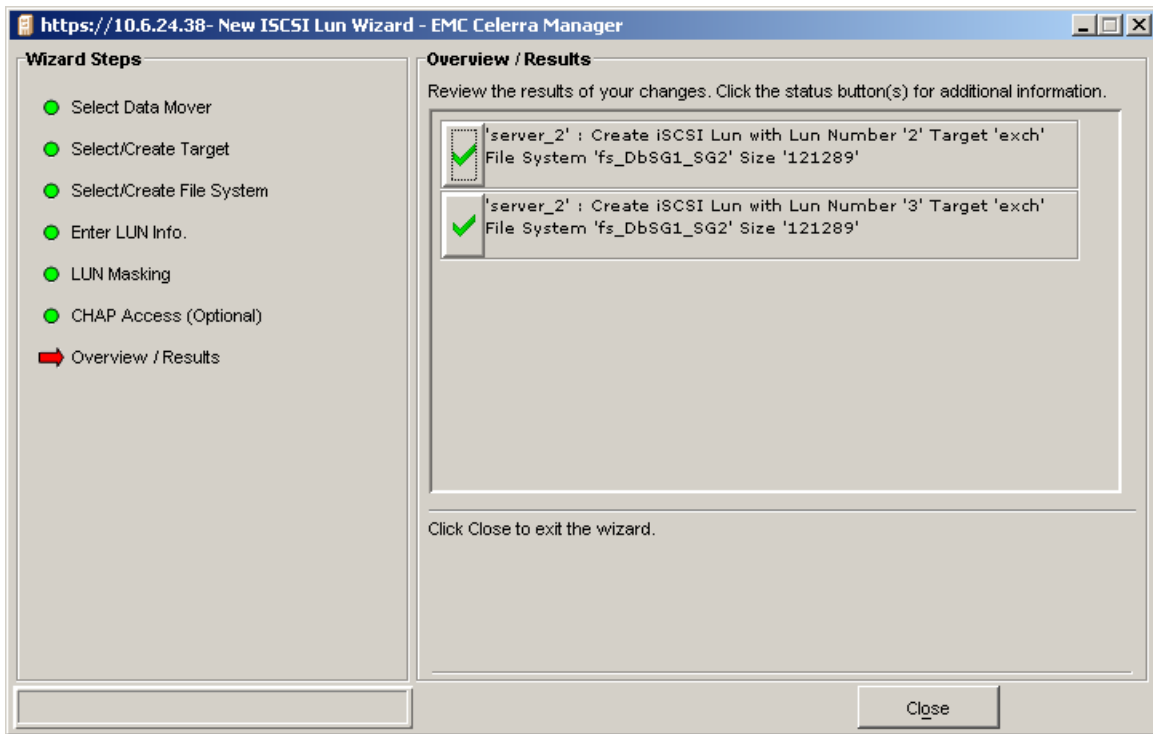


Figure 44. iSCSI Lun Wizard, Results

28. Go to Celerra Manager > iSCSI, and observe that the four new iSCSI LUNs are listed as shown in Figure 45.

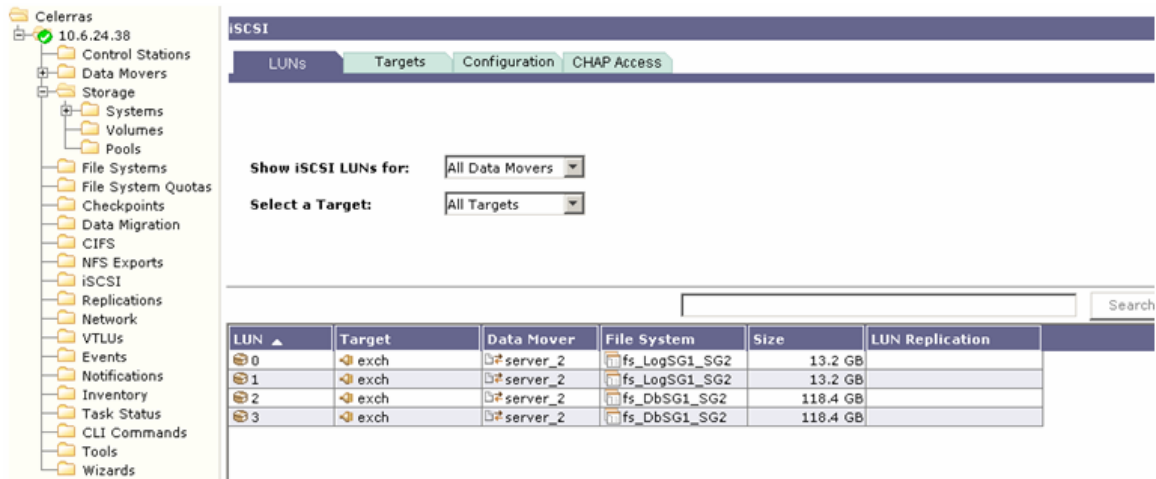


Figure 45. iSCSI LUN list

Optional configuration step

Multiple connections per session is an advanced iSCSI networking option to increase throughput of iSCSI network traffic and availability of the iSCSI LUNs. Detailed instructions can be found on Powerlink at emc.powerlink.com in the white paper *Best Practices for Celerra iSCSI: Considerations to Understand When Deploying Celerra iSCSI within Your Environment*.

Conclusion

This step-by-step guide covers creating Celerra file systems, Celerra iSCSI targets, and iSCSI LUNs. Celerra Manager wizards can be utilized during setup for ease of use.

References

The following documents provide additional, relevant information:

- *EMC Solutions for Microsoft Exchange 2007 NS Series iSCSI - Reference Architecture*
- *EMC Solutions for Microsoft Exchange 2007 NS Series iSCSI - Best Practices Planning* white paper
- *Best Practices for Celerra iSCSI: Considerations to Understand When Deploying Celerra iSCSI within Your Environment*