

Migrating data to your PowerVault ME5

A technologies overview of different methods to migrate data to a PowerVault ME5.

Abstract

This whitepaper discusses different methods of migrating your data from existing storage to a Dell PowerVault ME5 storage array.

February 2022

The information in this publication is provided “as is.” Dell Inc. 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 software described in this publication requires an applicable software license.

Copyright © February 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [2/10/2022] [Whitepaper] [Migrating data to your PowerVault ME5]

Contents

Introduction	3
Migrating to your PowerVault ME5 on the filesystem level.....	4
VMware vSphere Storage vMotion.....	4
Microsoft Server storage live migration	4
Citrix XenServer Storage live migration	5
File copy between volumes	6
Migrating to your PowerVault ME5 on the block level	6
Linear and Virtual Storage.....	6
Migrating between Virtual Pools.....	7
Prerequisites.....	7
Configuring the Peer Connection	7
To create a peer connection from PowerVault ME4	8
Creating a replication set.....	8
Breaking a replication set	10
Presenting ME5 volumes	11
Conclusion	13

Introduction

This document explores various technologies used to migrate data to a PowerVault ME5, both on a file level utilizing live migration features of commonly used operating systems, as well as on a block level, when migrating from a PowerVault ME4 series to a ME5 series array.

It does not include guidance on sizing, performance, or design, instead it provides information about the feature set of the PowerVault ME5 series.

This document is intended for PowerVault ME5 Series administrators, system administrators, and anyone responsible for configuring ME5 Series systems. It is assumed the readers have prior experience with or training in SAN storage systems and limitations of their environment.

Migrating to your PowerVault ME5 on the filesystem level

Following different technologies used to perform live migrations between 2 storage locations, are being explored.

These require a certain level of licenses, network, and host configurations.

For specifics on these prerequisites we encourage the administrator to visit the support pages of their respective operating system provider.

VMware vSphere Storage vMotion

VMware vSphere offers 2 types of vMotions: Storage vMotion and Compute vMotion.

Compute vMotion allows the administrator to live migrate the computational ownership of virtual machines between VMware vSphere Hypervisors.

Whereas Storage vMotion, is a technology that allows you to move virtual machine files from one storage location to another, without having to perform a “cold migration”. This is done by copying the virtual machines home directory to the destination datastore, which results in a “shadow” virtual machine. Using changed block tracking, VMware vSphere continuously is copying data to the destination datastore, until only a small amount of data is left on the source datastore. At which point, VMware vSphere performs a Fast Suspend and Resume (FSR) on the virtual machine and transfer the remaining data over, thus completing the transfer to the new datastore. The users and virtual machine are oblivious to the FSR. Afterwards the virtual machine files in the source datastore are deleted.

vMotion, requires a vCenter Server instance and the ESXi host having access to both the source and destination datastore. Furthermore, vMotion is also a licensed OS feature, please consult with your OS provider for the appropriate license level.

Performing a Storage vMotion:

1. In the vSphere vCenter web client, navigate to the “Storage” view.
2. From the list of datastores select the one with VMs to migrate and switch to the “VM” list on the right.
3. “CTRL” + “left-click” to select all VMs in the same power state.
4. Right-click one of the VMs and select “**Migrate**”.
5. Click “**Yes**”
6. Select “**Change storage only**” and click “**Next**”.
7. Select the ME5 Datastore.
8. Click “**Finish**”.

For more details on VMware vSphere vMotion, please visit [VMware Knowledge base](#).

Microsoft Server storage live migration

Similar to VMware vSphere, Microsoft Server Hyper-V nodes offer 2 migration options: live migration and storage live migration.

Live Migration allows the administrator to live migrate the computational ownership of virtual machines between Microsoft Hyper-v servers.

Storage Live Migration, on the other hand migrates the virtual machines files from one storage location to another storage location, without the need of a “cold migration”. The process is similar to VMware vSphere Storage vMotion, wherein the Live Migration copies the initial state of the virtual machine to the destination storage. Thereafter the Storage Live Migration continuously copies any changes of the virtual machines data to the destination storage. Once only a small amount of data is left on the source storage, the virtual machine will be briefly stunned, and the remaining files moved.

Live Migration is a licensed OS feature, please consult with your OS provider for details on the required licenses.

Performing a Live Storage Migration:

1. In the Hyper-V Manager, Right-click the VM you wish to migrate.
2. Select **Move**.
3. Click **Next**.
4. Select “**Move the virtual machine’s storage**” and click **Next**.
5. Select “**Move all of the virtual machine’s data to a single location**” and click **Next**.
6. Click “**Browse...**” and navigate to your Cluster shared Volume and click “**Select Folder**”.
7. Click **Next**.
8. Confirm your changes and click **Finish**.

For more details on Microsoft Server storage live migration, please visit [Microsoft’s support page](#).

Citrix XenServer Storage live migration

Just as Microsoft Server and VMware vSphere, Citrix XenServer offers 2 live migration technologies: Live migration and storage live migration.

Live Migration allows the administrator to live migrate the computational ownership of virtual machines between XenServer nodes.

Whereas Storage Live Migration, migrates the virtual machines files from one storage location to another storage location, without the need of a “cold migration”. The Storage Live Migration copies the initial state of the virtual machine to the destination storage. Thereafter the Storage Live Migration continuously copies any changes of the virtual machines data to the destination storage. Once only a small amount of data is left on the source storage, the virtual machine will be briefly stunned, and the remaining files moved.

For Live Migration a XenCenter instance is required, furthermore it is a licensed feature, and an appropriate license is need. Please consult with your OS provider, for more details on the required license.

Steps to perform a Storage Live Migration:

1. In XenCenter, right-click the VM you wish to migrate and select **Migrate to Server**.
2. From the **Destination** menu, select your server.
3. From the **Home Server** menu, select a server to assign the VM to.
4. In the **Storage** tab, select the new storage repository to place the VM on and click **Next**.
5. Keep “**Place all migrated virtual disks on the same SR**” selected and click **Next**.
6. In the **Storage network** menu, select a network to use for the live migration and click **Next**.
7. Confirm your settings and click **Finish**.

For more details on XenServer Live Migration, please visit the [XenServer support page](#).

File copy between volumes

In case your application or operating system doesn't offer live migration technologies, you will need to "cold migrate" your data, by performing a file level copy from the source data location to your PowerVault ME5. This can be done, by either copy/pasting your data, or by using more intricate operations such as robocopy or xcopy.

Should you be migrating from a PowerVault ME4 series to a ME5 series array, you can make use of array-based replication, provided you are meeting the prerequisites mentioned in the section [Migrating between Virtual Pools](#). If your configuration does not meet the outlined prerequisites, you will need to perform a file level copy of your data, as well.

Please note: In order for to switch production from existing solutions to the PowerVault ME5, you will need to schedule a production halt. This is to ensure data integrity

Migrating to your PowerVault ME5 on the block level Linear and Virtual Storage

Both the PowerVault ME4 and PowerVault ME5 offer two different storage technologies, linear and virtual.

Virtual storage is a method of mapping logical storage requests to physical storage (disks). It inserts a layer of virtualization such that logical host I/O requests are mapped onto pages of storage. Each page is then mapped onto physical storage. Within each page the mapping is linear, but there is no direct relationship between adjacent logical pages and their physical storage.

A page is a range of contiguous Logical Block Addresses (LBAs) in a disk group, which is one of up to 16 RAID sets that are grouped into a pool. Thus, a virtual volume as seen by a host represents a portion of storage in a pool. Multiple virtual volumes can be created in a pool, sharing its resources. This allows for a high level of flexibility, and the most efficient use of available physical resources.

Some advantages of using virtual storage are:

- It allows performance to scale as the number of disks in the pool increases.
- It virtualizes physical storage, allowing volumes to share available resources in a highly efficient way.
- It allows a volume to be comprised of more than 16 disks.

Virtual storage provides the foundation for data-management features such as thin provisioning, automated tiered storage, SSD read cache, and the quick rebuild feature.

The linear method maps logical host requests directly to physical storage. In some cases, the mapping is one-to-one, while in most cases the mapping is across groups of physical storage devices, or slices of them. This linear method of mapping is highly efficient. The negative side of linear mapping is lack of flexibility. This makes it difficult to alter the physical layout after it is established.

Migrating between Virtual Pools

Prerequisites

PowerVault ME5 and PowerVault ME4 have the ability to replicate between each other, allowing for an easy method of migrating your data to your new storage system. Furthermore, the replication peer-relationship is bi-directional, which means that you can not only migrate your data of your PowerVault ME4, but you can also use your PowerVault ME4 as a disaster recovery target.

To facilitate the replication, following requirements must be met:

- Both systems must have iSCSI or FC host ports.
- Both systems must be using the same protocol.
- Both systems must be connected to the same Network or Fabric.
- Must be switch attached topology.
- Virtual Storage Pools configured on both sides.
- Sufficient free space for internal snapshots used in the replication.
- If CHAP authentication for iSCSI is used, the authentication must be valid.
- You will need to provide a user account with the “manage” role
- Maximum of 4 Peer connections

Please Note: Down time will be required to switch your production from your PowerVault ME4 series to your ME5 series array.

Configuring the Peer Connection

The replication peer connection can be configured from either the PowerVault ME5 or the PowerVault ME4 web interface. The peer connection only needs to be established from one of the systems. Once a peer connection is established, it enables bi-directional communication between a local system and a remote system to transfer data between the two systems. Creating a peer connection requires a name for the peer connection and either an IP address of a single available iSCSI host port on the remote system, or a WWN of a single available FC host port on the remote system. Only iSCSI and FC host ports are used for the peer connection. The peer connection is defined by the ports that connect the two peer systems, as well as the name of the peer connection. The local system uses the remote address to internally run the `query peer-connection` CLI command. The results of the query are used to configure the peer connection.

The peer connection is defined by the ports that connect the two peer systems, as well as the name of the peer connection. The local system uses the remote address to internally run the `query peer-connection` CLI command. The results of the query are used to configure the peer connection.

You can create a maximum of four peer connections per storage system. However, only one peer connection is allowed to a particular remote system. Attempting to create a second peer connection to the same system will fail.


While creating the peer connection, the local system receives information about all host ports and IPs on the remote system as well as the remote system's licensing and host port health. It also links host ports of the select host port type on the local system to those on the remote system, so all ports of that type are available as part of the peer connection. Once created, the peer connection exists on both the local and remote systems.

Replications use the bi-directional communication path between the systems when exchanging information and transferring replicated data. Once you create a peer connection, you can use it when creating any replication set. Because the peer connection is bi-directional, replication sets can be created from both systems with replication occurring from either direction.

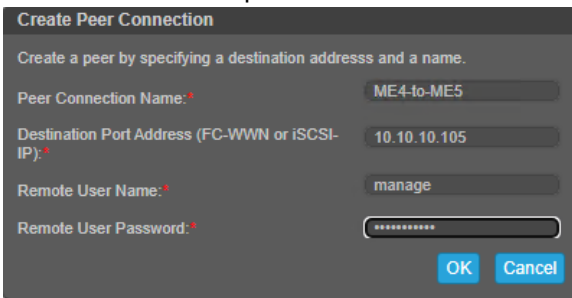
To create a peer connection from PowerVault ME4

The peer connection only needs to be created from one side, following is described how to create a peer connection using the PowerVault ME4 web interface:

1. In the Replications topic, select **Action > Create Peer Connection**. The Create Peer Connection panel opens.



2. Enter a name for the peer connection. The name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system or include the following: " , < \
3. Enter the destination port address for the remote system.
4. Enter the name and password of a user with the manage role on the remote system.

A screenshot of a 'Create Peer Connection' dialog box. The title bar says 'Create Peer Connection'. Below the title, there is a subtitle: 'Create a peer by specifying a destination address and a name.' There are four input fields: 'Peer Connection Name:' with the value 'ME4-to-ME5', 'Destination Port Address (FC-WWN or iSCSI-IP):' with the value '10.10.10.105', 'Remote User Name:' with the value 'manage', and 'Remote User Password:' with a masked password '*****'. At the bottom right, there are two buttons: 'OK' and 'Cancel'.

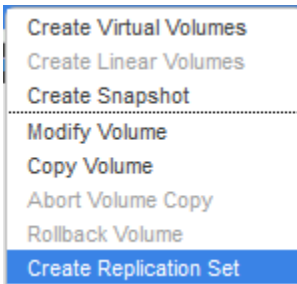
5. Click **OK**.
6. If the task succeeds, click **OK** in the confirmation dialog. The peer connection is created, and the Peer Connections table is updated.
If the task does not succeed, the Create Peer Connection panel appears with errors in red text. Correct the errors, then click **OK**.

Creating a replication set

A replication set, defines which volume or volume group you wish to replicate to the remote system. To create a replication set, open the PowerVault ME4 web interface and navigate to the Volume topic.

1. In the volumes table, select a volume or snapshot to use as the primary volume.

2. Click **Action > Create Replication Set**



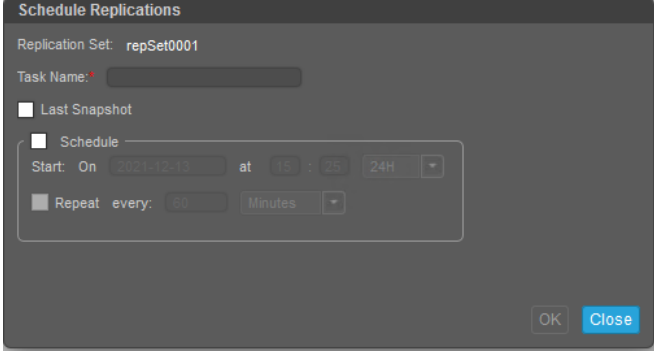
3. If the selected volume is in a volume group, source options appear.
 - To replicate the selected volume only, select **Single Volume**. This option is the default.
 - To replicate all volumes in the volume group, select **Volume Group**.
4. Enter a name for the replication set. The name is case sensitive and can have a maximum of 32 bytes. It cannot already exist in the system, include leading or trailing spaces, or include the following characters: " , < \
5. Optional: Select a peer system to use as the secondary system for the replication set.
6. Optional: Select a pool on the secondary system. By default, the pool that corresponds with the pool in which the primary volume resides is selected. The selected pool must exist on the remote system.
7. Optional: If **Single Volume** is selected, enter a name for the secondary volume. The default name is the name of the primary volume. The name is case sensitive and can have a maximum of 32 bytes. It cannot already exist on the secondary system or include the following: " , < \
8. Optional: Specify the **Queue Policy** action to take when a replication is running, and a new replication is requested.
9. Optional: Select the **Secondary Volume Snapshot History** check box to keep a snapshot history on the secondary system for the secondary volume.
 - Set the **Retention Count** to specify the number of snapshots to retain.
 - Modify the **Snapshot Basename** to change the snapshot name. The name is case sensitive and can have a maximum of 26 bytes. It cannot already exist in the system or include the following characters: " , < \
 - Set the **Retention Priority** to specify the snapshot retention priority.
 - Optional: Select the **Primary Volume Snapshot History** check box to keep a snapshot history for the primary volume on the primary system
10. Optional: Select the **Scheduled** check box to schedule recurring replications.

A screenshot of the 'Create Replication Set' dialog box. The 'Volume' field is 'ESXi_Datastore_Prod'. The 'Replication Set Name' is 'repSet0001'. The 'Select Peer Connection' dropdown is 'Me4-to-me5'. The 'Secondary Pool' has radio buttons for 'A' (selected) and 'B'. The 'Secondary Volume Name' is 'ESXi_Datastore_Prod'. There are radio buttons for 'Discard' and 'Queue Latest' (selected). A section for 'Secondary Volume Snapshot History' is expanded, showing 'Retention Count' as '1', 'Snapshot Basename' as 'ESXi_Datastore_Prod', and 'Retention Priority' as 'never-delete'. The 'Primary Volume Snapshot History' section is collapsed. The 'Scheduled (configure recurring replications)' checkbox is checked. At the bottom are 'OK' and 'Cancel' buttons.

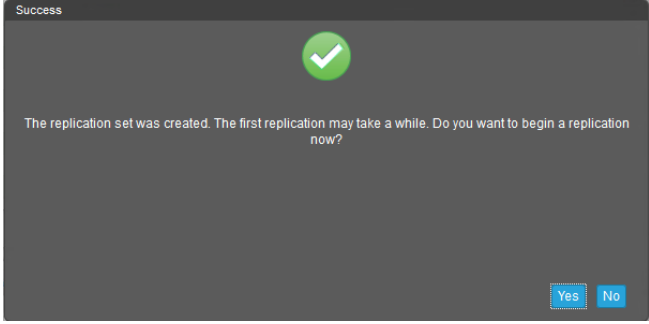
11. Click **OK**.

12. In the success dialog box:

- If you selected the Scheduled check box, click **OK**. The Schedule Replications panel opens, and you can set the options to create a schedule for replications. For more information on scheduling replications, see the PowerVault ME4 Administrator Guide - Initiating or scheduling a replication from the Volumes topic.



- Otherwise, you have the option to perform the first replication. Click **Yes** to begin the first replication, or click **No** to initiate the first replication later




Breaking a replication set

Once you have replicated all your data to the ME5 series array, schedule a time, where production can be temporarily taken offline. As soon as all services hosted by your PowerVault ME4 series array are shutdown, perform a final manual replica, to sync the most recent changes made.

Afterwards the replication set will need to be deleted. When you delete a replication set, all infrastructure created by the system (internal snapshots required to support replications) is also deleted. The primary and secondary volumes and volume groups no longer have restrictions and function like all other base volumes, volume groups, and snapshots.

1. In the Replications topic, select the replication set to be deleted in the Replication Sets table.
2. Select **Action > Delete Replication Set**.





3. Click **OK**. The replication set is deleted and the Replication Sets table is updated.



Presenting ME5 volumes


Now that you have deleted the replication set, you can present the ME5 volumes to your hosts and restore your services.

1. In the Provisioning Tab, select Volumes
2. Locate the Volume you wish to map and click the arrow pointing to the left: 
3. In the pop-up selected “**Attached Hosts**” and then “**Attach Host**”



No hosts attached to this volume.

4. Click **Continue**.
In case you already have a host initiator or host-group created, select the checkbox for **Attach volumes to existing hosts or host groups** and click **Continue**.



5. Enter a **Host Name** and select the corresponding **Initiator Id**, then click **Add Initiators to host**. If you have already created the host or host groups, select your host or host groups.

The screenshot shows a web interface titled "Attach to Hosts" with a breadcrumb trail "Hosts > Summary". Below the title is a "Create Hosts From Initiators" section. It contains a "Host Name" field with the value "Hyper-V01". Below this is a table with two columns: "Initiator Id" and "Nickname". The first row has an unchecked checkbox, the initiator ID "iqn.1988-11.com.dell:01.array.bc305bf0ce26", and the nickname "initiator0001". The second row has a checked checkbox, the initiator ID "iqn.1991-05.com.microsoft.win-8mqv26uk96a", and the nickname "initiator0002". Below the table is a blue button labeled "Add Initiators To Host". Underneath is a "New Hosts" section with a text box containing "No Hosts Created Yet". At the bottom are three buttons: "Continue", "Back", and "Cancel".

6. Click **Continue**.
7. Click **Continue**.

Conclusion

In conclusion, if the administrator intends to migrate their data live, without disruption to services, it is recommended to utilize file level Live Migration technologies, such as VMware vSpheres Storage vMotion or Microsoft Servers Live Migration.

Should the administrator not have these technologies available, they will need to perform a cold migration, either by copying, xcopy or robocopying the data from source to target system. If the administrator migrates from a PowerVault ME4 to a PowerVault ME5, they can make use of the virtual storage asynchronous bi-directional replication feature that only resides on a PowerVault ME5.