

EMC ViPR Controller

Version 2.3

ViPR Controller Ingest Services for Existing Environments Guide

302-002-073

REV 01

Copyright © 2013-2015 EMC Corporation. All rights reserved. Published in USA.

Published July, 2015

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.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

EMC Corporation
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.EMC.com

CONTENTS

| | | |
|------------------|--|-----------|
| Chapter 1 | Using ViPR Controller with Existing Environments | 5 |
| | Overview | 6 |
| | Block storage systems under ViPR Controller management..... | 7 |
| | File storage systems under ViPR Controller management..... | 8 |
| | Managing exported volumes created outside of ViPR Controller..... | 8 |
| | | |
| Chapter 2 | Ingesting Unmanaged Block Volumes into ViPR Controller | 11 |
| | Introduction..... | 12 |
| | Snapshots, full copies and continuous copies..... | 12 |
| | Ingestion of VMAX3 volumes with snapshots, full copies and continuous copies..... | 13 |
| | Ingesting block volumes into ViPR Controller: notes..... | 14 |
| | Ingestion constraints..... | 14 |
| | Before Discovery - VPLEX virtual volumes..... | 15 |
| | Ingest unexported unmanaged volumes..... | 16 |
| | Before discovery - unexported volumes and metavolumes..... | 16 |
| | Discover unmanaged volumes..... | 16 |
| | Ingest unexported unmanaged volumes..... | 17 |
| | Ingest exported unmanaged volumes..... | 18 |
| | Before discovery - exported volumes and metavolumes..... | 18 |
| | Ingest exported unmanaged volumes..... | 20 |
| | Block volume ingestion use cases..... | 21 |
| | Ingest SRDF protected volumes..... | 21 |
| | Ingest VPLEX volumes..... | 23 |
| | Move data from an unsupported array to supported array using VPLEX | 24 |
| | | |
| Chapter 3 | Ingesting Unmanaged File Systems into ViPR Controller | 27 |
| | Introduction..... | 28 |
| | Ingest file systems from the ViPR Controller UI..... | 28 |
| | Discover unmanaged file systems..... | 28 |
| | Ingest unmanaged file systems..... | 29 |

CONTENTS

CHAPTER 1

Using ViPR Controller with Existing Environments

This chapter contains the following topics:

- [Overview](#) 6
- [Block storage systems under ViPR Controller management](#)..... 7
- [File storage systems under ViPR Controller management](#)..... 8
- [Managing exported volumes created outside of ViPR Controller](#)..... 8

Overview

ViPR Controller enables storage arrays to be added to a virtual data center and for their storage pools to be used as the target for block and file storage provisioning operations. Where storage pools already host block volumes, or file systems, that were created outside ViPR Controller, and are not currently managed by ViPR Controller, ViPR Controller provides mechanisms for bringing them under ViPR Controller management.

The mechanism for bringing storage under ViPR Controller control is referred to as *ingestion*.

ViPR Controller also provides the ability to ingest volumes that are part of a VPLEX configuration, or that are protected using RecoverPoint (RP) or the Symmetrix Remote Protection Facility (SRDF).

The way in which existing storage has been provisioned on an array affects the ability of ViPR Controller to ingest block volumes or file systems from the array, and can affect future provisioning operations. The following table lists the import mechanisms that are provided and identifies any limitations.

Table 1 Existing environment support

| Operation | Systems | Notes and Limitations |
|---|-----------------------------------|--|
| Ingest Block Volumes Supports the ingestion of block volumes from block storage systems that are managed by ViPR Controller. See Ingesting Unmanaged Block Volumes into ViPR Controller on page 11 . | VMAX | VMAX volumes - both exported and unexported - can be ingested. However, where masking views exist, the provisioning of standard or FAST LUNs can be affected. The following white paper describes the effects of masking views and zones. <i>EMC ViPR in Existing VMAX Environments</i> |
| | VMAX and associated SRDF volumes. | Supports both SRDF/S and SRDF/A protected volumes. |
| | VPLEX | A VPLEX volume not in storage view can be ingested. However, back end operations are not supported. |
| | VNX | VNX volumes - both exported and unexported - can be ingested. |
| | VMAX, VNX, and VPLEX | ViPR Controller does not support the ingestion of volumes that are in a consistency group. |
| | XtremIO | All hosts and clusters must be discovered and configured in ViPR Controller. |
| Ingest File Systems Supports the ingestion of file systems from file storage systems that are managed by ViPR Controller. See Ingesting Unmanaged File Systems into ViPR Controller on page 27 . | Isilon/VNX File/NetApp | To be ingested by ViPR Controller, a file system must be exported. When ingesting unmanaged file systems into ViPR Controller, there are a number of limitations on the handling of file system exports that can cause the ingest operation to |

Table 1 Existing environment support (continued)

| Operation | Systems | Notes and Limitations |
|---|-------------|--|
| | | <p>fail or impact the way exports are accessed. These are:</p> <ul style="list-style-type: none"> • An export can only have a single export rule. • ViPR Controller only supports exports with single security type. For example, an NFS export on an array is allowed to have multiple security types (sys, krb, krbi, krpb, etc.) while ViPR Controller only allows for one security type. • All endpoints within an export have one set of host permissions (read-write, read-only and root) type. • ViPR Controller does not handle multiple use cases surrounding root user mappings. |
| Import and move LUNs using Host Migration Utility | | <p>Can be used to onboard a non-ViPR Controller LUN by migrating the data to a ViPR Controller-managed LUN.</p> <p>Can change the virtual storage pool of a LUN by migrating a source LUN to a new target.</p> <p>Documentation for this utility can be found here: ViPR Controller Migration Services - Host Migration Utility</p> |
| Use Existing Masking Views | VPLEX -VMAX | <p>When creating VPLEX volumes, existing masking views must meet the ViPR Controller criteria. The following White Paper discusses this scenario:</p> <p>EMC ViPR Controller Enhancement: Discovery and Use of VPLEX-VMAX Pre-Existing Multiple Masking Views</p> |

Block storage systems under ViPR Controller management

Once a volume is under ViPR Controller management, and has been provisioned or exported to a host through a ViPR Controller service, you should no longer use the storage system element manager to provision or export the volume to hosts. Using only ViPR Controller to manage the volume will prevent conflicts between the storage system database and the ViPR Controller database, as well as avoid concurrent lock operations being sent to the storage system. Some examples of failures that could occur when the element manager and ViPR database are not synchronized are:

- If you use the element manager to create a volume, and at the same time another user tries to run the "Create a Volume" service from ViPR on the same storage system, the storage system may be locked by the operation run from the element manager, causing the ViPR "Create a Volume" operation to fail.

- After a volume was exported to a host through ViPR, the same masking view, which was used by ViPR during the export, was changed on the storage system through the element manager. When ViPR attempts to use the masking view again, the operation will fail because what ViPR has in the database for the masking view is not the same as the actual masking view reconfigured on the storage system.

You can, however, continue to use the storage system element manager to manage storage pools, add capacity, and troubleshoot ViPR Controller issues.

File storage systems under ViPR Controller management

Once a filesystem is under ViPR Controller management, and has been provisioned or exported to a host through a ViPR Controller service, you should no longer use the storage system element manager to provision or export the filesystem to hosts. Using only ViPR Controller to manage the volume will prevent conflicts between the storage system database and the ViPR Controller database, as well as avoid concurrent lock operations being sent to the storage system. You can however continue to use the storage system element manager to manage storage pools, add capacity, and troubleshoot ViPR Controller issues.

Managing exported volumes created outside of ViPR Controller

ViPR Controller has the ability to manage exported volumes that it did not create or ingest. This discussion explains how ViPR Controller processes these volumes and the associated limitations.

For hosts with existing exported volumes that cannot be fully brought under ViPR Controller management, ViPR Controller makes a significant effort to co-exist by allowing new volumes to be created and exported to these hosts from within ViPR Controller. ViPR Controller attempts to reuse existing export entities (such as export masks and storage groups) on the storage system to do the following:

- Adhere to certain storage array restrictions that prevent new export entities from being created. For example, VMAX and VNX limit the number of initiator groups for a given initiator to one, resulting in ViPR Controller reusing the existing initiator groups.
- Avoid depleting the available resources on a storage system by reusing export entities instead of creating new ones. By reusing existing export entities on a storage system, ViPR Controller makes the following assumptions:
 - The existing export is correct and complete. ViPR Controller does not attempt to add additional storage ports or zones to these exports, and does not validate the exports against the virtual pool that is associated with the volumes being added.
 - ViPR Controller does not automatically try to ingest, manage or make visible pre-existing volumes in the export.

If an exported volume does not have any export entities, you can create and manage them in ViPR Controller.

Limitations

Review the following limitations to managing volumes that were created outside of ViPR Controller:

- If ViPR Controller uses an existing storage group, no new zones are created, even if there are existing zones.
- Adding or removing an initiator does not create a new zone if there are not existing zones.

- ViPR Controller does not check for other hosts in a storage group. If you export a new volume to a host, other hosts in the same storage group can see the new volume.
- ViPR Controller does not detect ports in a storage group, even without a connection, that are out of the virtual array.
- ViPR Controller does not verify initiator types. For example, it does not detect that the failover mode is not 4.
- ViPR Controller does not check the maximum paths and paths per initiator values set for the virtual pool before reusing masking views.

CHAPTER 2

Ingesting Unmanaged Block Volumes into ViPR Controller

This chapter contains the following topics:

- [Introduction](#)..... 12
- [Ingest unexported unmanaged volumes](#)..... 16
- [Ingest exported unmanaged volumes](#)..... 18
- [Block volume ingestion use cases](#)..... 21

Introduction

Existing block volumes can be brought under ViPR Controller management by ingesting them using services provided in the Service Catalog. The operation is a two stage process and requires the use of two services, a discovery service followed by an ingestion service. The services are for use by System Administrators only and are not visible to normal catalog users. This article describes how to use these services and provides additional information to support some ingestion use cases.

Once under ViPR Controller management, ingested storage resources can be managed by provisioning users in the same way as if they had been created using ViPR Controller, allowing them to be exported to hosts, expanded, protected using snapshot and copy techniques, etc. All functions, such as auto tiering policies, are supported on ingested volumes.

Any volume that is brought under ViPR Controller management must be assigned to a block virtual pool. So, before you can ingest volumes, you must ensure that a virtual pool exists whose criteria match those of the volumes to be ingested. For example, if you want to ingest volumes located on solid state devices (SSDs), you must ensure that virtual pools exist that allow SSDs. Similarly, if you want to ingest VMAX volumes that have associated SRDF volumes, you must ensure that there is a block virtual pool that can have VMAX volumes that are SRDF protected.

The discovery process finds volumes on a specified array and matches them with block virtual pools. When you use the ingest service, only volumes that were previously discovered are ingested.

Snapshots, full copies and continuous copies

When ViPR Controller ingests VMAX, VMAX3 and VNX volumes, it can also ingest their snapshots, full copies, and continuous copies. This ingestion can occur for exported and unexported volumes.

Always use the same project to ingest the source volume and its associated snapshots, full copies, and continuous copies volumes into ViPR Controller. Access is prohibited to these associated volumes when they are in a different project from their parent source volume.

ViPR Controller can ingest the following:

Snapshots

Snapshot targets are ingested as block snapshots. The ingested target has a corresponding session if one is available.

Note

For VNX volumes, ViPR Controller cannot ingest cascade snapshots.

Full copies

Differential full copies with active sessions to the source volume regardless of their current state are ingested. Detached full copies are ingested as independent volumes without maintaining association to the actual source.

Continuous copies

All continuous copies with active links to the source volume regardless of their current state are ingested. Any VNX continuous copies in a fractured or a split state are ingested as full copies. Detached continuous copies are ingested as independent volumes without maintaining association to the actual source volume.

Ingestion constraints

ViPR Controller does not ingest the following:

- Remote full and continuous copies
- Group snapshots
- Nested replicas, such as a snapshot of a snapshot
- VMAX3 SnapVX sessions
- VNX SnapView sessions

Post ingestion operations

After snapshots, full copies, and continuous copies are ingested, ViPR Controller can perform the following operations on them:

- Snapshots:
 - Restore
 - Delete
 - Export
 - Unexport
- Full copies:
 - Restore
 - Resynchronize
 - Detach
 - Delete
- Continuous copies:
 - Pause
 - Resume
 - Stop
 - Delete

Note

Before running the Remove Continuous Copy service, you must unexport the volume storing the continuous copies and then pause continuous copies. If not, this message appears: *Remove Continuous Copies for ingested mirror vol failed because mirror vol exported on array side before ingest.*

Ingestion of VMAX3 volumes with snapshots, full copies and continuous copies

When a VMAX3 array has volumes with snapshots, full copies or continuous copies, ViPR Controller ingests the source volume but not the associated replica volume and marks the source volume as internal. This also includes SRDF R1 and R2 volumes. Since ViPR Controller does not view these volumes as stand-alone volumes, you cannot perform any operations on these volumes until you ingest the associated replica volume.

Ingesting block volumes into ViPR Controller: notes

Before you ingest a volume, review the following notes and limitations.

- ViPR Controller can only ingest volumes on VMAX, VNX, and XtremIO arrays.
- ViPR Controller can ingest unexported VPLEX Local and VPLEX distributed virtual volumes. Ingest of exported VPLEX volumes is not supported in ViPR Controller.
- In a VPLEX environment, back-end storage is not ingested when you ingest the VPLEX virtual volume.
- Datastore information built from an ESX cluster is not ingested into ViPR Controller when the volume underlying the datastore is ingested. Only the volume presented to ESX is ingested.
- ViPR Controller cannot ingest volumes exposed to hosts with both iSCSI and Fiber Channel initiators present.
- After SRDF/A volumes set with the consistency exempt option are ingested, ViPR Controller can only perform the Create, Delete, Stop, and Suspend SRDF operations on them. All other SRDF operations act on the RDF group as a whole.
- ViPR Controller can ingest exported and unexported volumes and snapshots from XtremIO.
- For XtremIO, if a snapshot has multiple export masks (exported to multiple hosts) and ViPR Controller ingests only one of them, and if you delete this snapshot from the ingested volume, then the snapshots are also deleted from all export masks. In other words, all hosts containing the exported snapshots lose access to the ingested snapshot and are removed.
- For all block arrays, when the host's initiator group has mixed initiators (iSCSI and Fibre Channel), ViPR Controller is unable to ingest the masking views.
- For XtremIO iSCSI ingestion, all array storage ports (iSCSI and Fibre Channel) must exist in the virtual array for ViPR Controller to ingest the unmanaged export mask.

Ingestion constraints

ViPR Controller performs a set of tests before ingesting volumes.

- Both the virtual pool and virtual array specified in the service order must be active and accessible to the requesting tenant and project.
- You cannot ingest a volume or metavolume that is in a consistency group. If you request three volumes to be ingested, and two of them belong to consistency groups on the array, one volume will be successfully ingested and the two volumes in the consistency groups will be skipped.

Note

VPLEX virtual volumes that are part of a consistency group can be ingested.

- If the volume is exported to a host or cluster, it can be ingested if you run the service **Ingest Exported Unmanaged Volumes**.
- If the volume is RecoverPoint-protected, it cannot be ingested.
- If the volume's auto-tiering policy does not match the auto-tiering policy specified in its virtual pool, it cannot be ingested.
- If, by ingesting the volumes you requested, the capacity quota set on the virtual pool would be exceeded, the ingestion will fail.

- The SAN Multipath parameters in the virtual pool must satisfy the constraints established in the volume's masking view and zoning. See [SAN Multipath settings for VNX volume ingestion on page 19](#) for more information.
- If you are ingesting a VNX volume that was exported to a host, verify that the host is registered with ViPR Controller using its Fully Qualified Domain Name (FQDN). If the host is registered with ViPR Controller using its IP address, use the ViPR Controller User Interface to change the Host IP address to an FQDN.
- Ingestion of an exported volume will fail if the ports in its masking view/storage group are not part of the virtual array into which it will be ingested.
- If you ingest a volume, the Host LUN ID for the volume is not ingested from the export group on the array. To synchronize the Host LUN ID for a volume, you can unexport then export the volume. The Host LUN ID will then display properly in the ViPR Controller user interface.

Ingestion constraint: masking view with no zones

Follow the procedure below to ingest a volume that is masked, but has no zones.

If you are trying to ingest a volume that has been exported to a host, you could encounter this error:

```
Checking numpath for host rmhostc22.lss.emc.com
Initiator 20:00:00:25:B5:16:C1:20 of host rmhostc22.lss.emc.com is
not assigned to any ports.
Host rmhostc22.lss.emc.com
(urn:storageos:Host:a380810f-0962-4a69-952a-bbe2a498cf1a:vdcl)
has fewer ports assigned 0 than min_paths 4
```

This error indicates that there are no zones established for the masking view. To fix this problem, follow these steps:

1. Using your switch software, create zones for the masking view.
2. From the ViPR Controller user interface, choose **Physical Assets > Fabric Managers** and rediscover the switch.
3. Run **Block Storage Services > Discover Unmanaged Volumes** to discover the unmanaged volumes on the array.
4. Rerun **Block Storage Services > Ingest Exported Unmanaged Volumes**.

Before Discovery - VPLEX virtual volumes

ViPR enables you to ingest unmanaged VPLEX virtual volumes.

Before you ingest an unmanaged VPLEX virtual volume, complete the following procedures.

- Build a virtual array that includes network access to the VPLEX device and the VPLEX cluster on which the volume is configured.
- Build a virtual pool that has the same **High Availability** setting as the volume. (**VPLEX Local** or **VPLEX Distributed**)
- if the VPLEX virtual volume is in a consistency group, the virtual pool must have **Multi-Volume Consistency** enabled.

Backing storage for the VPLEX virtual volume is not ingested - only the VPLEX virtual volume.

RecoverPoint-protected volumes are not ingested.

Ingest unexported unmanaged volumes

To ingest a volume that is not exported to a host, run these services in this order:

1. Discover Unmanaged Volumes
2. Ingest Unexported Unmanaged Volumes

Before discovery - unexported volumes and metavolumes

Before you run the Discover Unmanaged Volumes service on an unmanaged volume or metavolume, follow this procedure.

Before you begin

From ViPR Controller, discover the storage array where the volume resides. This brings the physical storage pools on the array under ViPR management.

Procedure

1. Examine the volume in Unisphere to determine the storage pool in which the volume resides.
2. From ViPR Controller, build a virtual array that includes connectivity to the physical storage array on which your volume resides.
3. From ViPR Controller, build a virtual pool that matches the physical storage pool where the volume resides.
4. If the volume resides in a thin pool on the array, be sure that the virtual pool that matches the volume's physical storage pool has **Provisioning Type: Thin**.

Discover unmanaged volumes

The ViPR Controller service catalog provides a Discover Unmanaged Volumes service that finds block volumes which are not under ViPR Controller management and matches them to a ViPR virtual pool. The operation is also supported from the ViPR API and CLI.

Before you begin

The following prerequisites are applicable:

- This operation requires the System Administrator role in ViPR Controller.
- The virtual array and virtual pool into which you want to ingest the storage pools must exist when the discovery is performed. There must be at least one virtual pool in ViPR Controller that matches the physical storage pool that contains the volume.

Procedure

1. Select **Service Catalog > View Catalog > Block Storage Services > Discover Unmanaged Volumes**.
2. Select the physical block storage system on which you want to discover unmanaged volumes. You can select more than one storage system.
3. Select **Order**.

The orders page is displayed and shows the progress of the request.

Was the discovery successful?

This topic describes steps to take to see if your discovery service was successful, and how to respond to a discovery that does not succeed.

There are two ways to determine if the discovery is successful:

- Check the ViPR log ({vipr home}/logs/controllersvc.log) for the string SUPPORTED_VPOOL_LIST after unmanaged volume discovery. The matching virtual pools will be listed in the logs. If no pools are listed in the logs, the SUPPORTED_VPOOL_LIST will be missing in the logs for a given unmanaged volume.
- Call the following ViPR REST API:

```
GET
/vdc/unmanaged/volumes/{id}
```

The SUPPORTED_VPOOL_LIST section of the unmanaged volume feed should contain the names of the matching virtual pools. If this list is empty, then the volume can not be ingested as there are no supported virtual pools.

Ingest unexported unmanaged volumes

The ViPR Controller service catalog provides an Ingest Unexported Unmanaged Volumes service that brings previously discovered unmanaged block volumes under ViPR Controller management. The operation is also supported from the ViPR Controller API and CLI.

Before you begin

- This operation requires the System Administrator role in ViPR Controller.
- You must run [Discover Unmanaged Volumes on page 16](#) on the array from which the block volumes are ingested.
- Ingested volumes are assigned to a project. You must belong to the selected project and have write-permissions on that project.

Note

If the virtual array or virtual pool was modified since the last time the unexported unmanaged volumes were discovered, rerun Discover Unmanaged Volumes before running the ingest operation to ensure volumes are assigned to the correct virtual array and virtual pool.

Procedure

1. At the ViPR Controller UI, select **Service Catalog > View Catalog > Block Storage Services > Ingest Unmanaged Volumes**.
2. Select the storage system from which you want to ingest the block volumes.
3. Select a virtual array that contains physical array storage pools that you want to import. The storage system might contribute physical storage pools to a number of virtual pools. If you want to ingest volumes that match other virtual pools you will need to run the service again for the other virtual pools.

It is possible that not all of the array physical pools are included in the virtual array or arrays that form part of your virtual data center. For that reason, you don't want to ingest all unmanaged block volumes on the array, just those in physical array pools that form part of the virtual array.
4. From the array physical storage pools that form part of the virtual array, select a virtual pool that matches the physical pool where the unexported unmanaged volume reside.
5. Select a project. ViPR Controller assigns the unexported unmanaged volumes to the selected project.
6. Select **Order**.

The orders page shows the progress of the request. If the order is successfully fulfilled, you can look at the **Resources** page to see the imported volumes.

After you finish

After the unexported unmanaged volumes are ingested into ViPR Controller, you can export them to a host and then mount them, or use these volumes for other functions such as SRDF mirror volumes.

To export a volume to either a Windows or Linux host, use:

- **Service Catalog > View Catalog > Block Storage Services > Export Volume to a Host.**

To mount the volumes on a host:

- For Linux hosts use: **Service Catalog > View Catalog > Block Service for Linux > Mount Existing Volume on Linux.**
- For Windows hosts use: **Service Catalog > View Catalog > Block Service for Windows > Mount Volume on Windows.**

Ingest exported unmanaged volumes

To ingest a volume that is exported to a host, you must run these services in this order:

1. Discover Unmanaged Volumes
2. Ingest Exported Unmanaged Volumes
3. For a VMAX3 replica volume, repeat step 2 to ingest this volume

Before discovery - exported volumes and metavolumes

Before you run the Discover Unmanaged Volumes service on an unmanaged volume or metavolume that has been exported to a host or cluster, follow this procedure.

Before you begin

From ViPR Controller, discover the storage array where the volume resides.

From ViPR Controller, discover the host or cluster to which the volume has been exported. Use the FQDN of the host for discovery. If you discover the host using its IP address, you may have encounter problems ingesting exported VNX volumes.

Procedure

1. Examine the volume in Unisphere to determine the storage pool in which the volume resides.
2. Build a project or choose an existing project to which you have write access.
3. From ViPR Controller, build a virtual array that includes connectivity to the physical storage array on which your volume resides, and any hosts to which your target volume has been exported.
4. From ViPR Controller, build a virtual pool that matches the physical storage pool where the volume resides.
5. If the volume resides in a thin pool on the array, be sure that the virtual pool that matches the volume's physical storage pool has **Provisioning Type: Thin**.

Before Discovery - Checks for exported VMAX volumes

Before you ingest an unmanaged VMAX block volume that was exported to a host, you must collect some information from the masking views on the array and the fabric on the switch.

From the masking view, collect this information:

- Determine if the Masking View has a storage group (SG) associated with a specific FAST policy. If so, the FAST policy must be specified in the virtual pool that you use for the ingestion.
- For VMAX3, use Unisphere or the `symsg` command to determine the Service Level Option (SLO), the Workload Level and the Storage Resource Pool (SRP) settings for the storage group in the masking view. The following `symsg` commands provide this information:
 - `symsg -sid <SID> list`
 - `symsg -sid <SID> show <SG_Name>`
- If the host or cluster has multiple masking views, and those masking view have different FAST policies assigned, you will have to build multiple virtual pools - one for each FAST policy.
- If the storage group in the masking view specifies Host IO Limit settings, the virtual pool must specify Host IO Limit Settings that match the Storage Group.

Use your switch software to collect information about the fabric that enables communication between your host and your array.

- If your host initiator is zoned to 2 front end adapter ports, the virtual pool you use for ingest must have Minimum Paths set to 2.
- Set the Maximum Paths value in the virtual pool to the number of paths established for all of the host's initiators in all zones on the switch (or a greater value).
- If a host has multiple initiators, all of those initiators must be zoned to the same number of array ports. The number of paths per initiator is set in the virtual pool you use for ingest.

SAN Multipath settings for VNX volume ingestion

If you are ingesting an exported VNX volume, set SAN Multipath fields according to the guidelines below.

- Check the host in UniSphere (or ViPR or another tool) for the initiators associated with the host.
- Use the switch software (such as CMCNE or an equivalent CLI) to count the number of paths from each of the host's initiator ports to the array. This value is the number of array ports to which all the host's initiators have been zoned. For a VNX, the number of paths is determined based on the existing zones between the host initiators and the array ports that are in the VNX storage group. For example, if a host has two initiators, and each initiator is zoned to two array ports, the Maximum Paths field in the ingest virtual pool is set to 4.
- Check to be sure that each of the host's initiators is zoned to the same number of array ports. If the host's initiators are zones to different numbers of array ports, the ingest service will fail.

You will require this information when you build the virtual pool for the ingest.

Setting paths in zones and on the array before ViPR ingest

When computing the number of paths a host has to a VNX storage group or Vmax masking view, ViPR uses the zoning paths between the host initiators and the storage ports that are in that VNX storage group or Vmax masking view.

For some arrays such as a VPLEX or VNX, the same paths defined in the zones must also be defined on the storage array for the paths to become effective. It is possible for the two paths sets (the switch-defined zone paths, and the array-defined paths) to not match. This is considered to be a mis-configuration that should be corrected.

For example, assume the following:

- A storage group has 2 initiators: I1 and I2
- The storage group has two storage ports P1 and P2
- VNX paths are defined as I1 --> P1 and I2 --> P2
- All initiator ports are in a single zone: I1, I2, P1 and P2.

When ViPR generates a path count, it will count 4 paths:

- I1 --> P1
- I1 --> P2
- I2 --> P1
- I2 --> P2

If you assume that only 2 paths exist (and enter 2 in the Maximum Paths field of the virtual pool you use for the ingest operation) but ViPR counts 4, your ingest may fail. Remove un-used zoning paths before you try to ingest the volume.

FAST policy settings for ingested volumes

If your volume has a FAST policy assigned to it on the array, you must set the **Auto-tiering Policy** field to that policy name in the virtual pool you use for ingestion.

Ingest exported unmanaged volumes

The ViPR service catalog provides an Ingest Exported Unmanaged Volumes service that brings previously discovered unmanaged block volumes that have already been exported to hosts under ViPR management. The operation is also supported from the ViPR API and CLI.

Before you begin

This operation requires the System Administrator role in ViPR Controller.

Procedure

1. From the ViPR Controller user interface, select **Service Catalog > Block Storage Services > Ingest Exported Unmanaged Volumes**.

Enter the information in the following table into the service order form.

Table 2 Ingest Exported Unmanaged Volumes Service Settings

| Field | Description |
|---------------|---|
| Storage Type | Choose Exclusive or Shared . Exclusive means that this volume is exported to a single host. Shared means that this volume is exported to a cluster. |
| Host | Choose the Host or Cluster to which the volume has been exported. |
| Virtual Array | Choose the virtual array that you built, which contains the physical storage array and the host. |
| Virtual Pool | Choose the virtual pool that matches the physical storage pool that contains the volume. |
| Project | Choose the project you built. |

Table 2 Ingest Exported Unmanaged Volumes Service Settings (continued)

| Field | Description |
|---------|---|
| | Always use the same project to ingest the source volume and its associated snapshots, full copies, and continuous copies volumes into ViPR Controller. Access is prohibited to these associated volumes when they are in a different project from their parent source volume. |
| Volumes | Click the check box next to the volumes you want to ingest. |

2. Click **Order**.

The orders page shows the progress of the request. If the order is successfully fulfilled, the ingested volume appears on the Resources page. For VMAX3 volumes with an associated replica volume, an error displays in the Status field and a message similar to the following appears:

```
Error 64000: Unmanaged Volume Ingestion. Replicas
Symmetrix+000196700583+VOLUME+0060A not ingested for unmanaged
volume 008CF.
The unmanaged volume 008CF has been ingested but the ViPR object
has been
marked as internal.
```

The ingested source volume, which is marked as internal, does not appear on the Resources page until you ingest the VMAX3 replica volume.

3. Repeat steps 1 and 2 to ingest the VMAX3 replica volume.

If the order is successfully fulfilled, the ingested source and replica volumes appear on the Resources page.

Block volume ingestion use cases

For some ingestion use cases, the ingestion procedure requires additional steps or additional information about the configuration required in order to perform ingestion.

The following use cases are described:

- [Ingest SRDF protected volumes on page 21](#)
- [Ingest VPLEX volumes on page 23](#)
- [Move data from an unsupported array to supported array using VPLEX on page 24](#)

Ingest SRDF protected volumes

For volumes in an SRDF relationship, you ingest both the source volume (R1) and the target volume (R2) using a project whose name matches its RDF group. The source virtual pool must contain physical storage pools on the source array. The target virtual pool must contain physical storage pools on the target array. You perform two ingest operations because the source and target volumes are ingested using different virtual pools.

If you have not yet configured SRDF protection between two sites, see the *EMC ViPR Controller Integration with VMAX and VNX Storage Systems Guide*.

Note

Both the source and target virtual pools can belong to the same virtual array. However, it is recommended that you configure them into different arrays to ensure the disaster recovery relationship is clear.

To ingest an SRDF-protected volume, do the following:

1. Check the name of the RDF group for the SRDF pair you want to ingest. The ViPR Controller project you create for the ingest operation must have the same name as the RDF group. To check the RDF group name, you can use an element manager such as SMC, or use the following SYMCLI command:

```
# symcfg -sid <id> list -rdfg all
```

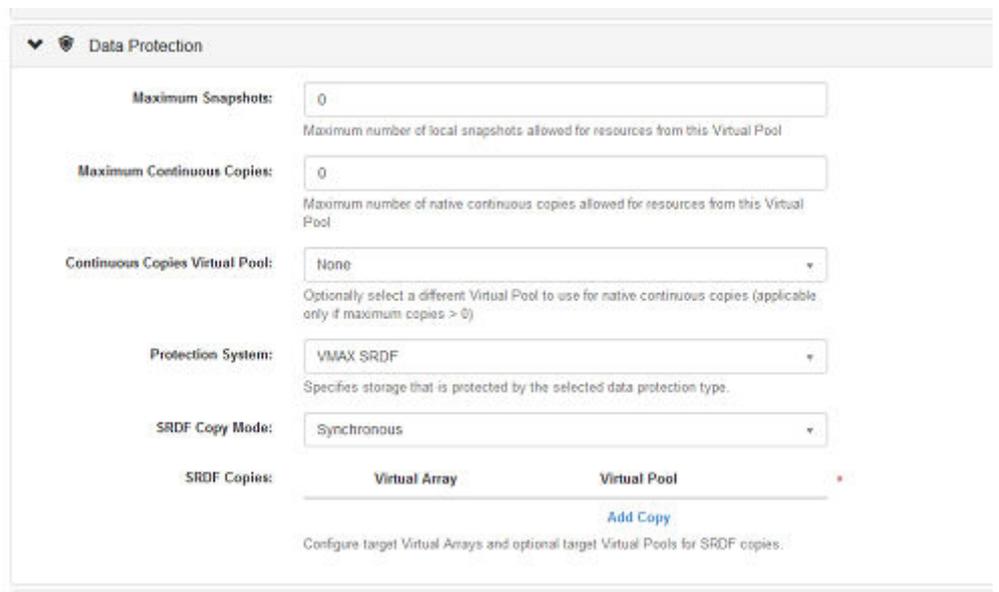
2. Ensure that you have set up a virtual array with a virtual pool that has SRDF protection configured. This virtual pool is used to ingest the source volume. You set the source virtual pool with the following SRDF settings, which you can find at the **Virtual Assets > Block Virtual Pools > Data Protection**.

Table 3 SRDF data protection settings

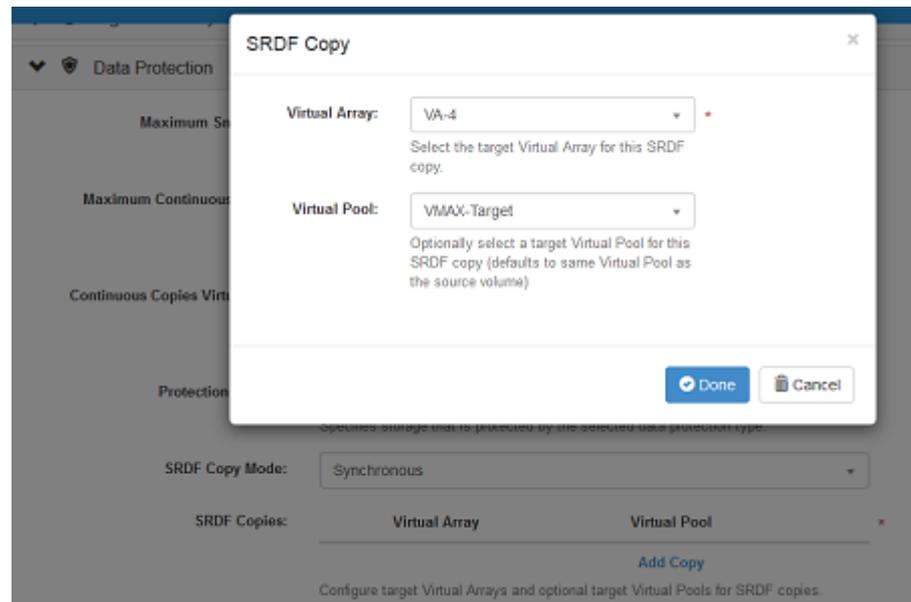
| Data Protection Field | Data Protection Setting |
|-----------------------|---|
| Protection System | VMAX SRDF . Only use this setting in the virtual pool during ingestion of the source volume. |
| SRDF Copy Mode | Synchronous or asynchronous mode. |
| SRDF Copies | Set the virtual array and virtual pool to host the target volume. |

The following shows the virtual pool Data Protection panel with the appropriate selections.

Figure 1 SRDF protection on source virtual pool



When you select **Add Copy**, specify the target virtual array and target virtual pool. Only virtual arrays and their associated pools that can act as the target for the source VMAX appear in this screen.

Figure 2 SRDF protection: specify target virtual pool

3. Ensure that you set up a virtual array with a virtual pool for ingesting the target volume. The virtual pool must *not* be SRDF protected.
4. Run the discovery process and specify both the source and target arrays, as explained in [Discover unmanaged volumes on page 16](#).
5. After successful discovery on the source and target arrays, run the ingestion process on the source array followed by the target array, as explained in [Ingest unexported unmanaged volumes on page 17](#).
The SRDF parent and SRDF targets field is populated which indicates the SRDF relationship is formed within ViPR.

Ingest VPLEX volumes

Ingestion of VPLEX volumes can be achieved using the discover and ingest services from the ViPR service catalog, or using the ViPR API or CLI. Because ViPR does not support all of the array back end types supported by VPLEX, it can ingest the VPLEX virtual volume and perform operations that do not depend having control of the back-end array.

The following notes are applicable:

- A volume can be ingested only if the VPLEX operations supported by ViPR can be performed. For example Export, Unexport, Mirror, Migrate etc.
- A virtual volume can be ingested only if it not already exported (and if it does not have an local mirror attached).
- Operations that require back end operations: Expand, Snap, Clone etc are not supported.

Ingested volumes can also be moved to a different virtual pool. This is useful where you want to migrate the back-end storage of a an ingested VPEX volume from an array that is not supported by ViPR to an array that is supported. When performing this virtual pool migration, the following apply:

- A virtual pool change to migrate the back-end storage can be performed on all ingested Local VPLEX volumes. For distributed VPLEX volumes, the ingested volume must consist of a single extent on a single storage volume.

- A virtual pool change of an ingested local VPLEX volume will result in a change to the name of the volume on the VPLEX. The VPLEX automatically changes the name of the local virtual volume to be the name of the local device on which it is built with the "_vol" suffix appended. Since it is migrated to a new local device created by ViPR, with the ViPR naming convention, the new name has the same convention as a ViPR created volume. This does not happen for a distributed volume, which uses VPLEX extent migration, because VPLEX device migration is not supported for distributed volumes.

A typical procedure for ingesting and moving an ingested volume is provided in: [Move data from an unsupported array to supported array using VPLEX on page 24](#).

Move data from an unsupported array to supported array using VPLEX

You can use a VPLEX storage system managed by ViPR to move data from an array that is not supported by ViPR to a supported array. The procedure requires you to connect the unsupported array to VPLEX and ingest the source

Before you begin

- Ensure that a VPLEX system has been added to ViPR and has been successfully discovered.
- Ensure that a VPLEX virtual pool exists. This will be used to create a target volume to which data will be moved. The virtual pool must specify High Availability as VPLEX Local.
- You must have the Tenant Administrator role in ViPR in order to access the Change Volume Virtual Pool Service.
- You must have the System Administrator role to perform any physical or virtual asset operations and to run the discovery and ingestion services.

Procedure

1. Connect VPLEX to the origin array.
2. Encapsulate origin array volumes into VPLEX using the VPLEX UI or Unisphere.
3. At the ViPR UI, create a virtual pool that includes the VPLEX physical storage pool that hosts the volume from which you want to ingest and migrate.
4. In the UI user mode, run the discovery service **Service Catalog > View Catalog > Block Storage Services > Discover Unmanaged Volumes** on the VPLEX array.

Details of running this service are provided in ([Discover unmanaged volumes on page 16](#)).

5. After successful discovery, run the ingestion process ([Ingest unexported unmanaged volumes on page 17](#)).

You will need to specify the virtual pool that you created and that has the origin volume.

6. If you have not already configured VPLEX to use the destination array, and configured a virtual pool that includes suitable storage pools from the destination array, you will need to:
 - a. Bring the destination array online and connect it to the VPLEX array.
 - b. Add the VPLEX array to ViPR and make sure that it has been discovered by ViPR.
 - c. Create a virtual pool (in an existing or new virtual array) that includes the VPLEX virtual pool to which you want to migrate the data. The VPLEX virtual pool be a storage pool located on the destination backing array.

7. Run the Change Volume Virtual Pool service to migrate the data. The service is located in the service catalog at **Service Catalog > View Catalog > Block Storage Services > Change Volume Virtual Pool** and specify the Operation as VPLEX Data Migration.

More information on using the Change Volume Virtual Pool service can be found in: [Change the ViPR Virtual Pool used in a VPLEX Environment](#).

More information on using the Change Volume Virtual Pool service can be found in: *Change the ViPR Virtual Pool used in a VPLEX Environment*.

This step will create volumes on both the VPLEX and the destination array and ViPR converts the underlying volume on the destination array to a ViPR-managed volume.

8. At the ViPR UI, delete the origin volume using the **Service Catalog > View Catalog > Block Storage Services > Remove Block Volumes** service or using the **Delete** operation from the **Resources > Volumes** page. Specify the Deletion Type as Inventory Only.

This will delete the volumes from ViPR only, the volumes remain on physical devices.

9. The VPLEX and origin array can now be removed, if desired, since data is on the new array.

Results

Data from the unsupported array is now successfully migrated to the supported storage array.

CHAPTER 3

Ingesting Unmanaged File Systems into ViPR Controller

This chapter contains the following topics:

- [Introduction](#)..... 28
- [Ingest file systems from the ViPR Controller UI](#)..... 28

Introduction

Existing file systems can be brought under ViPR Controller management by ingesting them using services provided in the service catalog. The operation requires two services: a discovery service followed by an ingestion services. This article describes how to use these services. These services are for use by System Administrators only and are not visible to normal catalog users.

Once under ViPR Controller management, the ingested storage resources can be managed by provisioning users in the same way as if they had been created using ViPR Controller, allowing them to be exported to hosts, expanded, protected using snapshot and copy techniques, etc.

Any file system that is brought under ViPR management must be assigned to a file virtual pool. So, before you can ingest volumes, you must ensure that a virtual pool exists whose criteria match those of the file systems to be ingested.

The discovery process finds file systems on a specified array and matches them with file virtual pools. When you use the ingest service, only file systems that were previously discovered to belong to a specified file virtual pool are ingested.

Ingest file systems from the ViPR Controller UI

To ingest file systems using the ViPR Controller Service Catalog, perform these tasks.

Procedure

1. Discover all unmanaged file systems on an array.
See [Discover unmanaged file systems on page 28](#).
2. Ingest discovered block volumes.
See [Ingest unmanaged file systems on page 29](#).

Discover unmanaged file systems

The ViPR service catalog provides a Discover Unmanaged File Systems service that finds file systems which are not under ViPR management. The operation is also supported from the ViPR API and CLI.

Before you begin

The following prerequisites are applicable:

- This operation requires the System Administrator role in ViPR Controller.
- The virtual array and virtual pool into which you want to ingest the storage pools must exist when the discovery is performed.

The discovery process finds storage pools on a selected storage system and identifies the virtual array and virtual pool that discovered each file system matches with.

- File systems will only be discovered on file storage system that have been added to ViPR as physical assets.

Procedure

1. Select **Service Catalog** > **View Catalog** > **File Storage Services** > **Discover Unmanaged File Systems**

2. Select the physical file storage systems from which you want to discover unmanaged file systems. You can select more than one file storage system.
3. Select **Order**.

The orders page is displayed and shows the progress of the request. If the order is successfully fulfilled, you can use the Ingest Unmanaged File Systems to bring them under management by ViPR.

Ingest unmanaged file systems

The ViPR Controller Service Catalog provides an Ingest Unmanaged File Systems service that brings previously discovered unmanaged file systems under ViPR Controller management. The operation is also supported from the ViPR Controller API and CLI.

Before you begin

The following prerequisites are applicable:

- Only System Administrators can ingest unmanaged file systems.
- To be ingested, the unmanaged file systems must be in physical pools which are already associated with a ViPR Controller virtual storage pool.
- The Discover Unmanaged File Systems service must have been run on the file storage system.
- Rerun the Discover Unmanaged File Systems service if the virtual array or virtual pools have been modified since the last time the Discover Unmanaged File Systems service was run.
- Ingested file systems will be assigned to a project. You must belong to the selected project and have write-permission on the project.
- To ingest an Isilon unmanaged file systems:
 - The virtual pool must have been configured with the `Provisioning Type` set to Thin. Isilon file systems are thinly provisioned. Thin resources can only be created in "thin" v pools.
 - The Isilon file system exports must be in either `/ifs/sos` or `/ifs/vipr`. File systems exported to other locations will not be ingested by ViPR Controller.

Procedure

1. Select **Service Catalog > View Catalog > File Storage Services > Ingest Unmanaged File Systems**
2. Select the storage system from which you want to ingest file systems.
3. Select the virtual array whose virtual pools contain the storage system physical pools that host the file systems you want to import. The storage system might contribute physical storage pools to a number of virtual pools. If you want to ingest to all virtual pools, you will need to run the service again for the other virtual pools.

It is possible that not all of the storage system physical pools are included in the virtual array that forms part of your virtual data center. For that reason, you don't want to ingest all unmanaged file systems on the storage system, just those in physical storage pools that meet the criteria of your virtual pools.

4. Select the virtual pool that the unmanaged volumes are in.
5. Select the project that you want the unmanaged file systems to be assigned to.
6. Select **Order**.

The orders page is displayed showing the progress of the request. If the order is successfully fulfilled, you can look at the **Resources > File Systems** page to see the imported file systems.