

EMC ViPR Controller

Version 3.6

Service Catalog Reference Guide

302-003-713

REV 02

Copyright © 2013-2019 Dell Inc. or its subsidiaries. All rights reserved.

Published February 2019

Dell 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." DELL 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 DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners.
Published in the USA.

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

CONTENTS

Figures		7
Tables		9
Chapter 1	Service Catalog Overview	11
Chapter 2	Application Services	13
	Application services.....	14
	Application services configuration guidelines.....	14
	Protecting volumes used in applications.....	16
	Managing full copies of applications.....	16
	Managing snapshots of application volumes.....	16
	Creating an application.....	16
	Deleting applications.....	17
	Adding block volumes to an application.....	17
	Removing Block Volumes from an Application.....	19
	Creating a full copy from an application's block volumes.....	19
	Removing a full copy of an application's block volumes.....	20
	Detaching a full copy of an application's block volumes.....	20
	Restoring from a full copy of an application's block volumes.....	21
	Resynchronizing a full copy of an application's block volumes.....	21
	Creating an application snapshot.....	22
	Restoring an application's snapshot.....	22
	Resynchronizing an application snapshot.....	23
	Deleting an application snapshot.....	23
	Linking an application snapshot.....	24
	Unlinking an application snapshot.....	25
Chapter 3	ViPR Controller Block Storage Services	27
	Block storage services.....	28
	Block storage systems under ViPR Controller management.....	28
	Block storage provisioning support.....	29
	Create, mount, and format a block volume on a Windows cluster ...	33
	Expand Block Volume additional details.....	34
	Block storage services for unmanaged volumes.....	37
	Block storage services to change a volume, virtual pool and virtual array... 38	
	ViPR Controller support for removing RecoverPoint protection from	
	volumes.....	39
	Change virtual pool.....	40
	Change Volume Virtual Pool.....	44
	Change virtual array.....	48
	Export path adjustment.....	49
	Block storage protection services.....	51
	Add journal capacity	51
	Failover Block Volume.....	52
	Swap continuous copies.....	54

	Create block snapshot.....	56
	Link Block Snapshot.....	58
	Restore block snapshot.....	59
	Remove block snapshot.....	60
	Create full copy.....	61
	Remove full copies.....	62
	Restore from full copies.....	62
	Resynchronize full copies.....	62
	Detach full copies.....	63
	Create snapshot full copy.....	63
	Create continuous copy.....	64
	Export Continuous Copy to a Host.....	64
	Unexport Continuous Copy.....	65
	Remove continuous copies.....	65
	Export snapshot to a host.....	65
	Unexport snapshot.....	66
	Create VPLEX Volume from Block Snapshot.....	66
	Resynchronize Block Snapshot (VMAX2 and XtremIO only).....	67
Chapter 4	ViPR Controller File Storage Services	69
	File storage services.....	70
	File storage systems under ViPR Controller management.....	70
	File storage provisioning support.....	70
	Create and manage file systems.....	71
	Create and manage CIFS shares.....	72
	Create and manage NFS exports.....	74
	Create and remove file system quotas.....	77
	File provisioning services for VMware vCenter.....	78
	File storage services for unmanaged file systems.....	78
	File system protection support.....	80
	File system snapshots.....	81
	Export NFS snapshots.....	82
	Export CIFS snapshots.....	83
	Associate and dissociate a file policy with a file system.....	83
	Failover to and Failback from a file system.....	83
Chapter 5	ViPR Controller Object Storage Services	85
	Object storage services.....	86
	Creating a bucket.....	86
	Editing a bucket.....	87
	Deleting a bucket.....	87
Chapter 6	ViPR Controller Vblock System Services	89
	Vblock systems services.....	90
	Use the VCE Vblock System Services to create clusters on the Vblock compute system.....	91
Chapter 7	Migration Services	93
	Migration services.....	94
	VPLEX data migration.....	94
	Setting the type of transfer speed for VPLEX data migration.....	95
	Mobility group migration.....	96
	Adding resources to a mobility group.....	97

	Removing resources from a mobility group.....	97
Chapter 8	ViPR Controller Service Orders and Resources	99
	Working with service orders.....	100
	My Orders	100
	Order details.....	101
	Approvals.....	102
	All Orders.....	103
	Download orders.....	104
	Delete orders.....	105
	Scheduled orders	106
	Monitoring and troubleshooting a service order.....	107
	Working with service resources.....	108
	Applications.....	108
	Volume resources.....	109
	Snap Sessions.....	110
	Block snapshot resources.....	110
	Export group resources.....	111
	File system resources.....	111
	File snapshot resources	116
	vNAS server resources.....	117
	Bucket resources.....	118
	Resource tasks.....	120
	Actionable events.....	121
Chapter 9	ViPR Controller User Roles	123
	ViPR Controller user role requirements.....	124
Chapter 10	Troubleshooting Error Messages	129
	Troubleshooting ViPR Controller error messages.....	130

CONTENTS

FIGURES

1	RecoverPoint configuration before running the Failover Block Volume service.....	53
2	RecoverPoint configuration after running the Failover Block Volume service.....	54
3	RecoverPoint configuration before running the Swap Continuous Copies service.....	55
4	RecoverPoint configuration after running the Swap Continuous Copies service.....	56

FIGURES

TABLES

1	General block storage provisioning services.....	30
2	Block services for AIX.....	31
3	Block services for HP-UX	31
4	Block services for Linux	32
5	Block services for Windows	32
6	Block provisioning services for VMware vCenter.....	32
7	Block storage services for unmanaged volumes.....	37
8	File system provisioning services.....	70
9	Services to create and manage file systems.....	72
10	Services for creating and managing CIFS shares.....	73
11	Permissions for setting CIFS shares.....	73
12	Services for creating and managing NFS exports for file systems.....	74
13	Services for managing NFS exports.....	75
14	Permissions to set on an export.....	77
15	Services for creating and deleting file system quotas.....	77
16	Services for managing the quota directory.....	78
17	File system provisioning services for VMware.....	78
18	File storage services for unmanaged file systems.....	78
19	Resources ingested for a file system.....	80
20	File system protection services.....	81
21	Services for managing file system snapshots.....	81
22	Services for managing NFS mounted snapshots.....	82
23	Services for managing export rules.....	82
24	ViPR Controller Service Catalog, VCE Vblock System Services.....	90
25	Order status icons.....	100
26	Order actions.....	101
27	Order details areas.....	101
28	Affected resources.....	102
29	Order details.....	102
30	Order status icons.....	103
31	Order actions.....	103
32	Scheduled orders options and attributes.....	106
33	Areas of the Applications page.....	108
34	Application details.....	108
35	Areas of the Volumes page.....	109
36	Areas of the Snap Sessions page.....	110
37	Snapshot Session detail page information and options	110
38	Areas of the File System page for the selected file system.....	111
39	Areas of the File Snapshots page.....	116
40	Areas of the vNAS Servers page.....	117
41	Buckets attributes.....	118
42	Actionable Events page.....	121
43	VDC roles.....	124
44	Tenant roles.....	127
45	Troubleshooting tips for common error messages.....	130
46	Troubleshooting tips for Active Directory and LDAP.....	139
47	Troubleshooting tips for administrator tasks.....	140

TABLES

CHAPTER 1

Service Catalog Overview

The ViPR Controller Service Catalog provides a selection of services to automate provisioning, and protection of block, and file storage, as well as host provisioning services for Vblock compute systems.

ViPR Controller Storage system support

The storage systems, and hosts that are supported are listed in the *ViPR Controller Support Matrix* which is available on the EMC Community Network (community.emc.com).

ViPR Controller service interfaces

The services in the Service Catalog can be managed from anyone of the following ViPR Controller interfaces:

- Using the services from the Service Catalog at the ViPR Controller UI or from the Catalog API
- Using the ViPR Controller API
- Using the ViPR Controller CLI

This guide focuses on the capabilities provided by the Service Catalog to demonstrate the support provided by ViPR Controller. However, you can perform the same operations using the API or CLI.

User requirements

Service operations can be run by ViPR Controller administrators, and users, however when logged into ViPR Controller with a user role you can only create resources and perform operations on resources belonging to project that you are assigned to (or are the owner of). If you are a Tenant Administrator you can run all services and choose any project to be the owner of the resource.

ViPR Controller and Virtual Data Center configuration requirements

Services can only be run after the ViPR Controller System Administrator has:

- Installed the ViPR Controller.
- Configured users, and projects.
- Configured the ViPR Controller virtual data center.
For more details see the *ViPR Controller Installation, Upgrade, and Maintenance Guide*, which is available from the [ViPR Controller Product Documentation Index](#).

Virtual pool requirements

Service provisioning, and protection capabilities depend on the use of one or more virtual pool from which block volumes, file systems, and compute systems can be created.

Information on adding storage systems and on creating virtual pools is provided in the following guides which are available from the [ViPR Controller Product Documentation Index](#):

- *ViPR Controller User Interface Virtual Data Center Configuration Guide*

- *ViPR Controller REST API Reference*
- *ViPR Controller CLI Reference Guide*

CHAPTER 2

Application Services

This chapter includes the following topics:

- [Application services](#)..... 14
- [Protecting volumes used in applications](#)..... 16
- [Managing full copies of applications](#)..... 16
- [Managing snapshots of application volumes](#)..... 16
- [Creating an application](#)..... 16
- [Deleting applications](#)..... 17
- [Adding block volumes to an application](#)..... 17
- [Removing Block Volumes from an Application](#)..... 19
- [Creating a full copy from an application's block volumes](#)..... 19
- [Removing a full copy of an application's block volumes](#)..... 20
- [Detaching a full copy of an application's block volumes](#)..... 20
- [Restoring from a full copy of an application's block volumes](#)..... 21
- [Resynchronizing a full copy of an application's block volumes](#)..... 21
- [Creating an application snapshot](#)..... 22
- [Restoring an application's snapshot](#)..... 22
- [Resynchronizing an application snapshot](#)..... 23
- [Deleting an application snapshot](#)..... 23
- [Linking an application snapshot](#)..... 24
- [Unlinking an application snapshot](#)..... 25

Application services

An application is a logical grouping of volumes determined by the customer. With application services, you can create, restore, resynchronize, detach, or delete full copies or snapshots of the volumes that are grouped by application.

A single ViPR block consistency group represents consistency groups on all related storage and protection systems including RecoverPoint, VPLEX, and block storage arrays (such as VMAX and VNX). In previous releases, a single consistency group was limited, at most, to one consistency group on any one storage system. This prevented the creation of full copies or snapshots of subsets of RecoverPoint or VPLEX consistency groups. Now you can use Application services to create and manage sub groups of volumes in order to overcome this limitation.

Workflow

- Review the guidelines for configuring application services. [Application services configuration guidelines](#) on page 14 For the ViPR Controller user roles required to perform these operations see [ViPR Controller user role requirements](#).
- Create block volumes and ensure they are added to ViPR Controller block consistency groups.

Note

The multi-volume consistency flag must be checked when you create the virtual pool that holds the block volumes. From **Virtual > Block Virtual Pools > Edit Block Virtual Pool**, choose the **Hardware** tab and select **Multi-Volume Consistency**.

-
- Create applications. [Creating an application](#) on page 16
 - Add block volumes to applications. [Adding block volumes to an application](#) on page 17
 - Choose which protection services (full copy or snapshot) to perform. [Protecting volumes used in applications](#) on page 16

Application services configuration guidelines

Use the application services feature to make full copies or snapshots of volumes associated with an application.

Guidelines for configuring application services

- Applications are supported only on these array types: VNX, VMAX, VMAX3, and XtremIO.
- Do not mix or match volumes in consistency groups by adding them to both Application Group and Mobility Group operations. Treat these operations as exclusive.
- When creating new consistency groups that will be used in applications, the recommended practice is to uncheck the flag, **Enable Array Consistency**, for VPLEX or VPLEX /RecoverPoint consistency groups. Keep the flag checked for block storage array consistency groups that are used in applications.
- When relinking a target to a different SnapVX session for VMAX3 sub groups, choose only one sub group each time.

Guidelines when using VNX volumes in applications

- An application snapshot is not supported for VNX. When a VNX volume is added into an application replication group, it does not support the snapshot operation.
- When VNX volumes are added to an application, the VNX array side consistency group (ECOM replication group) is removed. The VNX volumes are mapped with the virtual group. ViPR Controller supports only full copies for VNX volumes in an application. You can add volumes with full copies to a VNX consistency group.

Guidelines for adding volumes with replicas

- Volumes being added to an application must not have any replicas (full copies, snapshots or snapshot sessions).
- When you remove a volume with replicas from an application, you must delete the snapshots or detach the clones before adding that volume back into the application.
- Automatic creation of replicas varies:
 - ViPR Controller will not create replicas automatically if you add volumes to an application that has either a new application sub group or a sub group without replicas.
 - ViPR Controller will create replicas automatically if you add volumes to an application where the sub group has existing replicas.
- When a VPLEX or RecoverPoint volume is removed from an application, it is removed from the backend array replication group. If it has a replica, the replica will be removed from the replication group too, even though it is still attached to the volume.
- After a volume or all volumes in the consistency group are added to an application, consistency group-level replica operations are not allowed.
 - If **Enable Array Consistency** is unchecked, any volume provisioned in the VPLEX or VPLEX /RecoverPoint consistency group will not be added into the backend replication group. This means you cannot create a snapshot for this consistency group.
 - If **Enable Array Consistency** is checked, the VPLEX or VPLEX /RecoverPoint consistency group will be added into the backend replication group. In this use case, you can create a snapshot even though the volumes are not added into application.
 - After a volume from a consistency group is added to an application, the **Enable Array Consistency** flag is automatically unchecked. This means from that point on, all VPLEX volumes provisioned to that consistency group will not have their backend volumes added to an array consistency group.

Guidelines when adding XtremIO volumes to applications

- When adding XtremIO volumes to an application with snapshots, new snapshots for the newly added volume will not be created and added to the existing snapshot groups. This is a limitation of XtremIO.

Note

When you add XtremIO volumes to an application with snapshots, ViPR Controller will not create snapshots of those new volumes. If you want to remove XtremIO volumes with snapshots from an application, you must delete the snapshots first. This is because XtremIO does not allow adding or removing snapshots from consistency groups.

Protecting volumes used in applications

Use the **Catalog > View Catalog > Application Services > Application Block Protection Services** pages to access the full copy and snapshot services.

Procedure

1. Select **Catalog > View Catalog > Application Services > Application Block Protection Services**.
2. Choose either **Full Copy** or **Snapshot** to access the full copy and snapshot replication services.

Managing full copies of applications

You may create, remove, detach, restore, or resynchronize full copies of the volumes associated with applications.

Procedure

1. Select **Catalog > View Catalog > Application Services > Application Protection Services**.
2. Choose **Full Copy** to access the full copy services.

Managing snapshots of application volumes

You may create, restore, resynchronize, or delete snapshots of the volumes associated with applications. For VMAX3, you may link or unlink snapshots.

Procedure

1. Select **Catalog > View Catalog > Application Services > Application Protection Services**.
2. Choose **Snapshot** to access the snapshot replication services.

Creating an application

Create the application name and description.

Before you begin

For the ViPR Controller user roles required to perform this operation see [ViPR Controller user role requirements](#).

Procedure

1. Select **Catalog > View Catalog > Application Services > Create an Application**.
2. Enter the application name.

3. Enter the application description.
4. Click **Order**.

The Orders page appears showing the progress of the order.

After you finish

[Adding block volumes to an application](#) on page 17

Deleting applications

Use the **Catalog > View Catalog > Application Services > Delete an Application** page to delete applications.

Procedure

1. Select **Catalog > View Catalog > Application Services > Delete an Application**.
2. Enter the name of the application to be deleted.
3. Click **Order**.

The Orders page appears showing the progress of the order.

Adding block volumes to an application

Associate volumes in one or more ViPR Controller block consistency groups to an application. All volumes must be assigned to sub groups in the application. You will be able to create full copies or snapshots of the volumes associated with each sub group.

Before you begin

- Create applications. [Creating an application](#) on page 16
- Create block volumes and ensure they are added to ViPR Controller block consistency groups.

Note

The multi-volume consistency flag must be checked when you create the virtual pool that holds the block volumes. From **Virtual > Block Virtual Pools > Edit Block Virtual Pool**, choose the **Hardware** tab and select **Multi-Volume Consistency**.

- For RP or VPLEX protected volumes, it is recommended to uncheck the **Enable Array Consistency** flag when adding those volumes to a new consistency group. Do not uncheck this flag when creating consistency groups for non-RecoverPoint or non-VPLEX protected block volumes.

Note

You can't change this flag for consistency groups that have been created already. This flag can only be set when the consistency group is created.

- Determine which volumes support which applications in your deployment.
- Restrictions:
 1. A mix of meta and non-meta volumes cannot be added to the same application sub group.

2. You cannot add or remove volumes from an application while a full copy or clone is being created.

A ViPR block consistency group represents any of the following types of consistency groups. Configuration of applications varies depending upon the type of consistency group in use:

- Block volumes for non-RecoverPoint or non-VPLEX protected volumes are grouped into array consistency groups. All volumes in the array consistency group have to be part of the same application. There is a one-to-one relationship between the ViPR consistency group and the array consistency group, so no further sub grouping is allowed.
- The volumes in a RecoverPoint consistency group can be divided into multiple array consistency groups by grouping the volumes by application sub group.
- VPLEX virtual volumes must be in VPLEX consistency groups in order for them to be added to an application.

Before configuring applications, determine which type of consistency group is in use for the volumes associated with the application.

Procedure

1. Select **Catalog > View Catalog > Application Services > Add Block Volumes to an Application**.

2. Select the application you want to configure.

3. Select the consistency group associated with the application data.

You must know whether the consistency group holds RecoverPoint or VPLEX-protected block volumes.

4. Select one or more volumes that you want to add to the application.

For Block Volumes (non-RecoverPoint , non-VPLEX), you have to add all the volumes in the ViPR consistency group, because the block volume (non-RecoverPoint , non-VPLEX) ViPR consistency group has a 1:1 mapping with the array consistency group.

For RecoverPoint or VPLEX-protected volumes, you can choose to add one or more volumes to the application sub group. The sub group groups the RecoverPoint or VPLEX backend array volumes into a new array consistency group. The original RecoverPoint or VPLEX consistency group is not changed.

5. Enter a sub group name.

This is a required field for RecoverPoint and VPLEX volumes. A sub group may not be used in more than one application.

- For block volumes that are not RecoverPoint or VPLEX-protected, you don't need to type a sub group name. Instead, choose the ViPR consistency group from the list and select all volumes. All volumes in the consistency group will be added.
- For RecoverPoint or VPLEX protected volumes, the sub group maps directly to the group name on the array. The name needs to be unique across the array.

6. Click **Order**.

The Orders page appears showing the progress of the order.

Removing Block Volumes from an Application

This service removes the volumes from the application and disassociates the volumes and full copies or snapshots (replicas) of the volume from the respective array consistency group. It does not delete the volumes or replicas.

- You must use this feature to remove volumes and replicas if you want to remove a sub group from an application or rename the sub group.
- If you remove all volumes that are part of an application sub group, ViPR Controller deletes the sub group name from the selection list in application services. After removal, the list of full copies or snapshots in **Resources > Applications** no longer displays these volumes or replicas.
- When you remove a volume with replicas from an application, you must delete the replicas before adding that volume back into the application. To delete the replicas after you remove the volumes from the application, [Remove full copies](#) on page 62 or [Remove block snapshot](#) on page 60.

Note

When you remove volumes and replicas, the volumes, full copies or snapshots still exist, but they are no longer part of an application or array consistency group. You cannot reuse the volume in another application unless you first remove the replicas.

- For VPLEX or RecoverPoint /VPLEX protected volumes that have XtremIO backend volumes, you will get an "Unsupported" exception if you try to remove an XtremIO volume from an application with snapshots. XtremIO does not allow removal of snapshots from groups. You must delete the snapshot first, then remove the volume from the application.
- For non-VPLEX or non-RecoverPoint volumes, all volumes from a consistency group must be removed in one order.

Procedure

1. Select **Catalog > View Catalog > Application Services > Remove Block Volumes from an Application**.
2. Select the application you want to remove.
3. Select the volume(s) you want to remove.
4. Click **Order**.

The Orders page appears showing the progress of the order.

Creating a full copy from an application's block volumes

Creates a full copy of all volumes associated with an application. Array consistency is at the application sub group level for all volumes in the same array.

- Use SnapVX snapshots for VMAX3 instead of application services full copies. This recommendation supports recent VMAX3 replica trends.

Procedure

1. Go to **Catalog > View Catalog > Application Services > Application Block Protection Services > Full Copy > Create Application Full Copy**.

2. Enter the application that has volumes to be copied.
3. Select the site.

This field is applicable only to RecoverPoint . Choose either Source or one of the targets.

Multiple target volumes are allowed, but you can copy only one site at a time. If non-RecoverPoint volumes belong to the application, leave this field blank.

4. Enter one or more application sub groups.
5. Enter a name for the full copy.

This name displays in **Resources > Applications > Full copies**. It is also used as a prefix for the full copy volumes. The name that displays consists of the copy name, the subgroup name, and a number.

6. Click **Order** to run the service.
7. Go to **Resources > Applications** and select the application to verify the full copy. Multiple copies display with a number suffix.

Removing a full copy of an application's block volumes

Removes full copies of an application's block volumes.

Before you begin

Procedure

1. Go to **Catalog > View Catalog > Application Services > Application Block Protection Services > Full Copy > Remove an Application Full Copy**.
2. Select the application.
3. Select the full copies to remove.
4. Select the application subgroup, if used.
5. Click **Order**.

The progress of the order displays.

Detaching a full copy of an application's block volumes

Removes the source and target relationship of a copy session for an application's volumes. This results in two independent volumes.

Concurrent restore / detach operations for full copies or link / unlink operations for snapshots acting on the same application or application sub group are not supported. The following message displays: *The following resources have pending task against them. Please retry when there are no pending tasks for the resources.*

Procedure

1. Go to **Catalog > View Catalog > Application Services > Application Block Protection Services > Full Copy > Detach an Application Full Copy**.
2. Select an application.
3. Select the full copy to detach.

After you order this service, the full copy becomes detached from its source volume and no longer appears in this list.

4. Select the application subgroup, if used.
5. Select **Order**.

The progress of the order displays. Once complete, you can expand, export, unexport, and delete a detached full copy volume.

Restoring from a full copy of an application's block volumes

Restores an application's source volume with the latest data from a full copy.

- If you added volumes to a subgroup after the last full copy was taken, the new volumes will return to the state they were in at the time of the addition to the subgroup.

Note

If you don't want these new volumes to be overwritten during the restore, you must remove them from the application before restoring. If you still want the new volumes, add them back in after the restore completes.

- Concurrent restore / detach operations for full copies or link / unlink operations for snapshots acting on the same application or application sub group are not supported. The following message displays: `The following resources have pending task against them. Please retry when there are no pending tasks for the resources.`
- You cannot add or remove volumes from an application while a restore is in progress.

Procedure

1. Go to **Catalog > View Catalog > Application Services > Application Block Protection Services > Full Copy > Restore from an Application Full Copy**.
2. Select the application with full copy source volume to restore.
3. Select the copy to restore.
All volumes in the copy will be restored.
4. Select the application sub group to restore, if used.
5. Select **Order**.
The progress of the order displays.

Resynchronizing a full copy of an application's block volumes

Copies the latest data from an application's source volume to its full copy.

- If volumes were added since the last full copy or resynchronization task completed, the volume will be included in the new copy.

Procedure

1. Go to **Catalog > View Catalog > Application Services > Application Block Protection Services > Full Copy > Resynchronize an Application Full Copy**.
2. Select the application with the full copy source volume containing the data to copy.

3. Select the full copy name to synchronize.
4. Select the application subgroup that will hold the full copy volume.
5. Select **Order**.

The progress of the order displays.

Creating an application snapshot

Create a snapshot from an application's block volumes.

Before you begin

Configure the associated virtual pool for snapshots.

Procedure

1. Go to **Catalog > View Catalog > Application Services > Application Block Protection Services > Snapshot > Create Application Snapshot**.
2. Select the application.
3. Select the site.

This field is applicable only to RecoverPoint . Choose either Source or one of the targets.

Multiple target volumes are allowed, but you can copy only one site at a time. If non-RecoverPoint volumes belong to the application, leave this field blank.

4. Enter one or more application sub groups.
5. Enter a name for the snapshot.

This name displays in **Resources > Applications > Snapshots**. It is also used as a prefix for the snapshot volumes. The name that displays consists of the copy name, the subgroup name, and a number.

6. If you want the XtremIO snapshot to be read only, check the box.

This option is for XtremIO snapshots only.

7. Click **Order** to run the service.
8. Go to **Resources > Applications > Snapshots** to view the application snapshot.

Restoring an application's snapshot

Restores an application's snapshot point-in-time data back to a source volume or a consistency group.

- If you added volumes to a subgroup after the last snapshot was taken, the new volumes will return to the state they were in at the time of the addition to the subgroup.

Note

If you don't want these new volumes to be overwritten during the restore, you must remove them from the application before restoring. If you still want the new volumes, add them back in after the restore completes.

- Concurrent restore / detach operations for full copies or link / unlink operations for snapshots acting on the same application or application sub group are not supported. The following message displays: *The following resources have*

pending task against them. Please retry when there are no pending tasks for the resources.

- You cannot add or remove volumes from an application while a restore is in progress.

Procedure

1. Go to **Catalog > View Catalog > Application Services > Application Block Protection Services > Snapshot > Restore Application Snapshot.**
2. Select the application.
3. Select the type of snapshot to restore.
For example, Snapshot/Linked Target or Snapshot Session.
4. Select the application snapshot copy to restore.
5. Select the application sub group, if used.
6. Click **Order**.

The progress of the order displays.

Resynchronizing an application snapshot

Copies the latest data from an application's source volume to its snapshot. Applicable for VMAX2, VMAX2 plus VPLEX, EMC Timefinder VPSnap, XtremIO, or XtremIO plus VPLEX backend arrays and for RecoverPoint protection of these arrays. This is NOT applicable for VNX or VMAX3 arrays.

- If volumes were added since the last snapshot or resynchronization task completed, the volume will be included in the new snapshot.

Note

XtremIO is an exception. A new snapshot is not created when adding XtremIO volumes to an application.

Procedure

1. Go to **Catalog > View Catalog > Application Services > Application Block Protection Services > Snapshot > Resynchronize Application Snapshot.**
2. Enter the application name.
3. Enter the name of the copy to be resynchronized.
4. Enter the application sub group if applicable.
5. Select **Order**.

The progress of the order displays.

Deleting an application snapshot

Deletes a snapshot of an application's block volumes.

Concurrent Restore, Detach, Link, or Unlink operations for full copies or snapshots acting on the same application or application sub group are not supported. The following message is shown: The following resources have pending task against them. Please retry when there are no pending tasks for the resources.

Procedure

1. Go to **Catalog > View Catalog > Application Services > Application Block Protection Services > Snapshot > Delete Application Snapshot**.
2. Select the application from which the snapshot will be deleted.
3. Select the type of snapshot.
For example, Snapshot/Linked Target or Snapshot Session.
4. Select the name of the copy to be deleted.
5. Select the application sub group, if applicable.
6. Click **Order**.
The Orders page is displayed with the progress of the order.
7. At the **User > Resources** page, observe that the host has been removed from the Hosts field for the resource.

Linking an application snapshot

Link an application snapshot to an application session. This is applicable for VMAX3 only.

Concurrent restore / detach operations for full copies or link / unlink operations for snapshots acting on the same application or application sub group are not supported. The following message displays: The following resources have pending task against them. Please retry when there are no pending tasks for the resources.

Procedure

1. Go to **Catalog > View Catalog > Application Services > Application Block Protection Services > Snapshot > Link Application Snapshot**
2. Select the application.
3. Select the copy to be linked.
4. Select the application subgroup associated with the copy.
5. Select the snapshot target to be linked.

This relinks existing Linked Snapshot Targets to the same or a different Snap Session for the selected application sub group.

Note

When you try to relink an application VMAX3 subgroup's linked target to a different SnapVX snapshot session, you must relink only one sub group at a time.

6. Select **Advanced** to link a new snapshot to a snapshot session.
Snapshot sessions are applicable to VMAX3 only.
7. Select **Order**.
The progress of the order displays.

Unlinking an application snapshot

Unlink an application snapshot from a snapshot session. This feature is applicable for VMAX3 only.

Concurrent restore / detach operations for full copies or link / unlink operations for snapshots acting on the same application or application sub group are not supported. The following message displays: The following resources have pending task against them. Please retry when there are no pending tasks for the resources.

Procedure

1. Go to **Catalog > View Catalog > Application Services > Application Block Protection Services > Snapshot > Unlink Application Snapshot**
2. Select the application.
3. Select the copy to be unlinked.
4. Select the application subgroup associated with the copy.
5. Select the snapshot target to be unlinked.

This will unlink an existing Linked Snapshot Target from the Snap Session for the selected Application Sub Group.

6. Check the **Delete Targets** box if you want to delete the target volume snapshots.
7. Select **Order**.

The progress of the order displays.

CHAPTER 3

ViPR Controller Block Storage Services

This chapter includes the following topics:

- [Block storage services](#)..... 28
- [Block storage provisioning support](#)..... 29
- [Block storage services for unmanaged volumes](#)..... 37
- [Block storage services to change a volume, virtual pool and virtual array](#).....38
- [Export path adjustment](#).....49
- [Block storage protection services](#)..... 51

Block storage services

ViPR Controller supports the provisioning of volumes and the protection of file storage using snapshots, full copies, and continuous copies.

- The block storage services include:
 - [Block storage provisioning services](#)
 - [Block storage services for unmanaged volumes](#)
 - [Block storage services to change a volume, virtual pool, or virtual array](#)
 - [Block protection services](#)

Block storage systems under ViPR Controller management

Using only ViPR Controller to manage the volume prevents conflicts between the storage system database and the ViPR Controller database, and avoids concurrent lock operations being sent to the storage system.

After a volume is placed under ViPR Controller management and is provisioned or exported to a host through a ViPR Controller service, do not use the storage system element manager to provision or export the volume to hosts. Here are some examples of failures that could occur when the element manager and the ViPR Controller database are not synchronized:

- 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 Controller on the same storage system, the storage system may be locked by the operation run from the element manager, causing the ViPR Controller "Create a Volume" operation to fail.
- After a volume is exported to a host through ViPR Controller, the same masking view, which was used by ViPR Controller during the export, was changed on the storage system through the element manager. When ViPR Controller attempts to use the masking view again, the operation fails because what ViPR Controller has in the database for the masking view is not the same as the actual masking view reconfigured on the storage system.

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

How ViPR Controller deletes masking views for VMAX and VNX for block storage

ViPR Controller sets the Solutions Enabler `DeleteWhenBecomcesUnassociated` flag to true when it is used to provision VMAX or VNX for block storage.

Note

Non-ViPR Controller created groups have this flag set to FALSE because this an SMI-S Provider API exposed parameter.

When ViPR Controller is used to delete or decommission the storage that was previously created and provisioned by ViPR Controller, the masking views are deleted or decommissioned as follows:

Example 1

Deleting a masking view with the `DeleteWhenBecomcesUnassociated` group flag set to `TRUE` at the parent group level only deletes the parent because the children do not have the flag set to `TRUE`.

Deleting a Masking View with the `DeleteWhenBecomcesUnassociated` group flags set at the parent group levels, deletes the groups that are not part of any other masking view.

Example 2

Deleting a masking view with the `DeleteWhenBecomcesUnassociated` group flag set to `TRUE` at both parent and child group levels delete all groups at all levels.

Deleting a masking view with the `DeleteWhenBecomcesUnassociated` group flags set to `TRUE` at both parent and child group levels **ONLY** delete these groups if they are **NOT** part of any other masking view.

Example 3

Deleting a masking view with the `DeleteWhenBecomcesUnassociated` group flags set to `TRUE` at the child group level, only deletes the child group and leaves the parent intact.

Deleting a masking view with the `DeleteWhenBecomcesUnassociated` group flags set to `TRUE` at the child group level, only deletes the child groups if they are not part of any other masking view.

Note

If the storage group is FAST managed, it does not get deleted.

Block storage provisioning support

The ViPR Controller Service Catalog provides access to a set of predefined services, which includes high-level operations that carry out common provisioning activities, such as creating a block storage volume and exporting it to a host or cluster, and "building block" services that perform more granular operations, such as creating a ViPR Controller volume or exporting storage to a host.

General block storage provisioning services

ViPR Controller services are organized in categories. There is a general Block Storage Services category used to create block storage volumes and export them to any type of host.

The **Service Catalog > Block Storage Services** pages display only unmounted volumes. You cannot select a mounted block volume from these pages and delete it. Instead, use the Block Services for Linux, Block Services for HP-UX, Block Services for VMware vCenter, and so forth to first unmount the volume. Then you can safely unexport or delete the volume.

Note

The API and CLI commands will not allow you to unexport a mounted volume.

To access these services, go to **Service Catalog > View Catalog > Block Storage Services**.

Table 1 General block storage provisioning services

Service name	Description
Create Block Volume for a Host	Creates one or more volumes of a specified size on a selected virtual array and virtual pool, and then exports these volume(s) to a host or cluster.
Expand Block Volume	Increases the amount of provisioned storage to the host or cluster. For details see: Expand Block Volume additional details .
Create Block Volume	Creates one or more volumes of a specified size on a selected virtual array and virtual pool. The Create block volume service can also be used to add a volume to a consistency group. When snapshots are enabled on a consistency group, you can only add a volume to the consistency group with the following storage system types: <ul style="list-style-type: none"> • VMAX • VNX for block (when array consistency is disabled) • XtremIO • VPLEX with VMAX, VNX, or XtremIO backing volumes
Export Volume to Host	Creates the exports from the volume to the host or cluster. Note On an XtremIO array, an Initiator can be part of only one Initiator Group (IG). In that case, when ViPR Controller tries to export a Volume to a Host (Host-A), if the Host's initiators (Host-A) are already registered on the array and that initiator group contains other host's initiators, it has to reuse the existing initiator group. As a result, the Volume is exposed to other host's initiators as well.
Unexport Volume	Removes a volume from an export. This volume is still available using another export.
Unexport Multiple Volumes	Removes multiple volumes from all of their exports. The volumes will no longer be accessible from any exports.
Remove Block Volumes	Removes unexported block volumes or consistency groups. Deletes only volumes that have nothing attached.

Table 1 General block storage provisioning services (continued)

Service name	Description
Unexport Remove Block Volumes	Removes block volumes or consistency groups and associated exports.
Remove Volume by Host	Removes an unmounted block volume assigned to a host from all of its exports and deletes the volume. The deleted volume is not available from any host.

Block provisioning services for hosts and VMware vCenter

To access these services, go to **Service Catalog > View Catalog > Block Storage Services for AIX**.

Table 2 Block services for AIX

Service name	Description
Create and Mount Block Volume	Creates, mounts, and formats a block volume on an AIX host.
Mount Existing Volume on AIX	Mounts and formats an existing volume that was already exported to an AIX host or a cluster.
Unmount Volume on AIX	Unmounts a block volume from an AIX host, leaving the storage intact.
Unmount and Delete Volume	Unmounts and deletes a block volume from an AIX host.

To access these services, go to **Service Catalog > View Catalog > Block Services for HP-UX**.

Table 3 Block services for HP-UX

Service name	Description
Create and Mount Block Volume	Creates, mounts, and formats a block volume on an HP-UX host.
Unmount Volume on HP-UX	Unmounts a block volume from an HP-UX host, leaving the storage intact.
Unmount and Delete Volume	Unmounts and deletes a block volume from an HP-UX host.
Mount Existing Volume on HP-UX	Mounts and formats a previously created block volume on an HP-UX host.
Expand Volume on HP-UX	Expands an HP-UX host by expanding a mounted volume.

To access these services, go to **Service Catalog > View Catalog > Block Services for Linux**.

Table 4 Block services for Linux

Service name	Description
Create and Mount Block Volume	Creates, mounts, and formats a block volume on a Linux host.
Unmount and Delete Volume	Unmounts and deletes block volume from a Linux host.
Mount Existing Volume on Linux	Mounts and formats a previously created block volume on a Linux host.
Unmount Volume on Linux	Unmounts a block volume from a Linux host, leaving the storage intact.
Expand Linux Mount	Expands a Linux host by expanding a mounted volume.

To access these services, go to **Service Catalog > View Catalog > Block Services for Windows**.

Table 5 Block services for Windows

Service name	Description
Create and Mount Volume	Creates, mounts, and formats a block volume on a Windows host.
Unmount and Delete Volume	Unmounts and deletes block volume from a Windows host.
Mount Volume on Windows	Mounts and formats a previously created and exported block volume on a Windows host.
Unmount Volume on Windows	Unmounts a block volume from a Windows host, leaving the storage intact.
Expand Volume on Windows	Expands a block volume mounted on a Windows host.

To access these services, go to **Catalog > View Catalog > Block Services for VMware vCenter**.

Table 6 Block provisioning services for VMware vCenter

Service name	Description
Create Volume for VMware	Creates one or more volumes of a specified size for a selected virtual array and virtual pool and exports these volumes to an ESX/ESXi host or cluster.
Create Volume and Datastore	Creates one or more volumes of a specified size for a selected virtual array and virtual pool, creates a VMware datastore on each new volume, and then assigns them to an ESX/ESXi host or cluster.
Remove Volume for VMware	Removes a volume from VMware.
Create VMware Datastore	Creates a VMware datastore from an existing volume.
Delete VMware Datastore	Removes a VMware datastore leaving the storage intact.
Export Volume for VMware	Exports a volume to a vCenter host or cluster, and then rescans the HBAs on the vCenter host or cluster.

Table 6 Block provisioning services for VMware vCenter (continued)

Service name	Description
Unexport Volume for VMware	Unexports a volume from a vCenter host or cluster, and then rescans the HBAs on the vCenter host or cluster.
Extend Datastore with New Volume	Creates a new volume to increase the amount of storage allocated to the datastore.
Extend Datastore with Existing Volume	Uses an existing volume to increase the amount of storage allocated to the datastore.
Expand Volume and Datastore	Expands the size of a VMware datastore and its underlying volume.
Create Block Snapshot	Creates a snapshot of a volume on an ESX/ESXi host or cluster.

Block storage services for hosts and vCenters unsupported for Openstack Third-party storage systems

The OpenStack Cinder API does not provide the storage volume world wide name (WWN) that is required to perform some of the operations for Block Storage services for hosts and vCenters. Due to this issue, you cannot use the following services for third-party block storage systems:

- **Block Services for AIX > Create and Mount Volume**
- **Block Services for AIX > Mount Volume on AIX**
- **Block Services for HP-UX > Create and Mount Block Volume**
- **Block Services for HP-UX > Mount Existing Volume on HP-UX**
- **Block Services for Linux > Create and Mount Volume**
- **Block Services for Linux > Mount Volume on Linux**
- **Block Services for Windows > Create and Mount Volume**
- **Block Services for Windows > Mount Volume on Windows**
- **Block Services for VMware vCenter > Create Volume and VMware Datastore**
- **Block Services for VMware vCenter > Create VMware Datastore**

Create, mount, and format a block volume on a Windows cluster

The following procedure shows how to create, mount, and format a volume on a Windows cluster.

Before you begin

All hosts must be powered on to provision volumes.

Procedure

1. Go to the **Service Catalog > Block Services for Windows > Create and Mount Volume** service.
2. Select **Shared** for the **Storage Type**.

Shared is to provision the storage across the cluster.

Exclusive is to provision the storage to a single host.

3. Select the cluster for which the storage will be provisioned.
4. Select the virtual array from which the storage will be allocated.

Note

Available virtual arrays are based on ViPR Controller's visibility to the host initiator. If the virtual array you are wanting to select is not available in the selection list, it may be that the host is powered down.

5. Select the virtual pool from which the storage will be allocated.
6. Select the project to which the resources will be assigned once the storage is provisioned.
7. Provide a name for the volume.
8. If required, select the volume consistency group.
9. Enter the size the volume will be after provisioned.
10. In the **File System Type** field, select the volume format.
11. Enter the drive letter for the mount point.
If left blank, the next available letter will be used.
12. Enter a drive label.
If left blank, ViPR Controller assigns the volume name as the drive label..
13. Click **Order**.
The Orders page appears with the progress of the order.

Expand Block Volume additional details

The ViPR Controller **Expand Block Volume** services can be used to increase the storage capacity of a volume under ViPR Controller management.

To increase the storage capacity of any volume in ViPR Controller, run one of the following ViPR Controller services from the **Catalog**:

- To expand a volume mounted on a Linux host, run **Block Services for Linux > Expand Linux Mount**.
- To expand a volume mounted on a Windows host, run **Block Services for Windows > Expand Volume on Windows**.
- To expand an unmounted volume, run **Block Storage Services > Expand Block Volume**.

Expand block volume support and limitations

You can use ViPR Controller to expand any volume you have provisioned or ingested in ViPR Controller with the following exceptions:

You cannot expand the following volumes with ViPR Controller:

- RecoverPoint target volumes. When you expand RecoverPoint source volumes through the services, automatically ViPR Controller expands target volumes.
- RecoverPoint journal volumes.
- Volumes with replicas (snap, clone, mirror)
- VMAX3 volumes that are in a replication state which includes mirrors, clones, and snapshots. Although ViPR Controller does not support the native expansion of

volumes in an SRDF relationship, it does support this workflow to expand VMAX3 volumes:
 SRDF links broken > volume natively expanded > SRDF links re-established

Note

You can expand the SRDF source and target volumes when the source and targets are both on VMAX3 storage systems, however, the expand operation is disruptive.

- SRDF target volumes. (When the source volume is expanded, the target volume is automatically expanded as well.)
- VMAX3 volumes that are less than 25 MB in size.

Expand the block volume

You can expand the storage capacity of a block volume using the **Expand Block Volume** service.

Before you begin

You must be logged in as a user with admin rights to the project that contains the volume to expand.

The volume must meet all the prerequisites in [Expand block volume support and limitations](#).

Procedure

1. From the ViPR Controller Service Catalog, run **Block Storage Services > Expand Block Volume**.
2. In the **Project** field, select the project that contains the volume.
3. In the **Volume** field, select the volume to expand.
4. In the **Size** field, enter a size number larger than the current size of the volume.
5. Click **Order**.

After you finish

According to RecoverPoint best practices, your RecoverPoint journal volumes must be 2.5 times as large as the volumes they protect. To comply with best practices, when you expand a VPLEX virtual volume with RecoverPoint protection, expand the RecoverPoint journal volumes.

To adjust the journal volume size, use Unisphere for RecoverPoint. For more information, see the *EMC RecoverPoint Administrator's Guide*.

Expand a block volume mounted on a Linux host

The **Expand Linux Mount** service expands a volume that is mounted on a Linux host.

Before you begin

Run a ViPR service such as **Create Block Volume** to create the volume. Alternatively, call **Ingest Unmanaged Volumes** to bring an external volume under ViPR management.

Discover the Linux host. To add a host, select **Physical Assets > Hosts**.

You must be logged in as a tenant administrator with access to the Linux host where the volume is mounted.

Do not access storage during an expand operation. Users that access storage while the service is running can experience an access disruption.

Procedure

1. Log in as a tenant administrator with access to the Linux host where the target volume is mounted.
2. Choose **Block services for Linux > Expand Linux Mount** from the ViPR Controller service catalog.
3. In the **Linux Host** field, select the Linux host where the volume is mounted.
4. In the **Volume** field, select the volume that you want to expand.
5. Choose the new size of the volume.
6. Click **Order**.

Expand a block volume mounted on a Windows host

The **Expand Volume on Windows** service non-disruptively expands a volume that was mounted on a Windows host.

Before you begin

Run a ViPR Controller service such as **Create Block Volume** to create the volume. Alternatively, call **Ingest Unmanaged Volumes** to bring an external volume under ViPR Controller management.

Discover the Windows host. To add a host, select **Physical Assets > Hosts**.

You must be logged in as a tenant administrator with access to the Windows host where the volume is mounted.

Procedure

1. Log in as a Tenant Administrator with access to the Windows host where the target volume is mounted.
2. From the ViPR Controller service catalog, select **Block services for Windows > Expand Volume on Windows**.
3. Select the storage type: An **Exclusive** volume means that this volume is mounted on a specific host. A **Shared** volume indicates that the volume is mounted on a Windows cluster.
4. Choose the Windows host or cluster where the volume is mounted.
5. In the **Volume** field, select the volume to expand.
6. In the **New Size** field, enter the new size of the volume.
7. Click **Order**.

Orchestration: Expand a VPLEX virtual volume with RecoverPoint protection

ViPR Controller supports the expansion of VPLEX virtual volumes that have RecoverPoint protection. An *orchestration* is a series of functions performed by ViPR Controller in a specific order that accomplishes a requested task.

Before you begin

To build the volume, follow the instructions in *EMC ViPR Controller Integration with RecoverPoint and VPLEX User and Administration Guide*.

To expand the volume, run one of the services described earlier from the ViPR Controller service catalog. The sequence that follows describes the orchestration ViPR Controller performs after you submit the expansion request.

Procedure

1. Remove the RecoverPoint replication set for the volume.

2. Use native array tools to expand the physical volume.

ViPR expands the RecoverPoint source and target volumes. It cannot expand the RecoverPoint journal volumes.

3. Re-create the RecoverPoint replication set.

Results

ViPR Controller expands the volume to the size requested by the user. The expand service does not change the worldwide name (WWN) of the volume.

Block storage services for unmanaged volumes

Unmanaged volumes are volumes that exist on a storage system discovered by ViPR Controller but were not created or discovered by ViPR Controller. The discover and ingest services are used by ViPR Controller to place the unmanaged volumes under ViPR Controller management.

When working with services for unmanaged volumes:

- For the ViPR Controller user roles required to perform this operation see [ViPR Controller user role requirements](#).
- You can perform the ingest operations from the ViPR Controller UI, API and CLI.
- You must run the Discover Unmanaged Volumes service on the storage system before running an ingest service.
- Ingested volumes are assigned to a project. You must belong to a selected project and have write permissions on that project.
- If the virtual array or virtual pool was modified since the last time the unmanaged volumes were discovered, rerun Discover Unmanaged Volumes prior to running the ingest operation to ensure volumes are assigned to the correct virtual array and virtual pool.

To access these services, go to **Service Catalog > View Catalog > Block Storage Services**.

Table 7 Block storage services for unmanaged volumes

Service category	Service name	Description
Block storage services	Discover Unmanaged Volumes	Finds block volumes that are not under ViPR Controller management and matches them to a ViPR Controller virtual pool. When performing discovery it is important to note: <ul style="list-style-type: none"> • 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.
	Ingest Unexported Unmanaged Volumes	Ingests volumes that were created on a storage system but were not exported to hosts. When ingesting unexported unmanaged volumes from a VPLEX storage system, you have two options:

Table 7 Block storage services for unmanaged volumes (continued)

Service category	Service name	Description
		<ul style="list-style-type: none"> • Full Ingestion including Backend Volumes which ingests both the VPLEX virtual volumes and the backend volumes on VMAX, VNX, XtremIO, and VMAX3 for both VPLEX local and metro. • Ingest only Virtual Volume to ingest just the VPLEX virtual volume. <p>After the unmanaged volumes are ingested into ViPR Controller, you can export them to a host and mount them, or use them for other functions such as SRDF mirror volumes.</p> <p>For detailed information on ingestion, see <i>ViPR Controller Ingest Services for Existing Environments Guide</i>.</p>
	Ingest Exported Volumes	<p>Brings previously discovered unmanaged block volumes that were already exported to hosts under ViPR Controller management.</p> <p>When ingesting exported unmanaged volumes from a VPLEX storage system, you have two options:</p> <ul style="list-style-type: none"> • Full Ingestion including Backend Volumes which ingests both the VPLEX virtual volumes and the backend volumes on VMAX, VNX, XtremIO, and VMAX3 for both VPLEX local and metro. • Ingest only Virtual Volume to ingest just the VPLEX virtual volume. <p>After the unmanaged volumes are ingested into ViPR Controller, you can export them to a host and mount them, or use them for other functions such as SRDF mirror volumes.</p> <p>For detailed information on ingestion, see <i>ViPR Controller Ingest Services for Existing Environments Guide</i>.</p>

Block storage services to change a volume, virtual pool and virtual array

You can change the virtual array or virtual pool of volumes.

When working with services for changing the virtual array or virtual pool of volumes:

- For information on which ViPR Controller roles can order the services for changing the virtual array or virtual pool, see [ViPR Controller user role requirements](#).
- The operations can be performed from the ViPR Controller UI, API and CLI.

- When performing the operation, **VPLEX Data Migration**, you can set the data migration transfer speed. See [Setting the type of transfer speed for VPLEX data migration](#).

To access these services using the ViPR Controller UI, go to **Service Catalog > View Catalog > Block Storage Services**

ViPR Controller support for removing RecoverPoint protection from volumes

The **Block Storage Services > Change Virtual Pool** and **Block Storage Services > Change Volume Virtual Pool** services can also be used to remove RecoverPoint protection from volumes.

When RecoverPoint protection is removed from volumes, the source volumes remain intact, but the target volumes are deleted. In addition, if these volumes are the last volumes in the consistency group, then journal volumes are also deleted.

The target virtual pool must be identical to the original virtual pool, with the exception that the target virtual pool does not include RecoverPoint protection.

You can remove the RecoverPoint protection from VPLEX source volumes, only if those volumes do not have any snapshots.

Note

If you have run the **Block Protection Services > Swap Continuous Copies** service to make the RecoverPoint target become the source, you cannot remove the RecoverPoint protection, until you run the **Swap Continuous Copies** service again to reverse the personalities of the source and target virtual pools.

Configuration requirements to use ViPR Controller to remove RecoverPoint protection from volumes

Volume configuration	Requirements to remove protection from volumes
RecoverPoint protection — CDP, CRR, CLR	<ul style="list-style-type: none"> • Cannot be in a swapped state. • Target virtual volume must be identical to original virtual pool.
RecoverPoint protection with VPLEX — CDP, CRR, CLR	<ul style="list-style-type: none"> • Cannot be in a swapped state. • All snapshots need to be deleted. • Target virtual volume must be identical to original virtual pool.
MetroPoint protection — CDP, CRR	<ul style="list-style-type: none"> • Cannot be in a swapped state. • All snapshots need to be deleted. • Target virtual volume must be identical to original virtual pool except for RecoverPoint protection.

Change virtual pool

The Change Virtual Pools service moves a set of volumes from one virtual pool to another.

Before you begin

- This operation requires the Tenant Administrator role in ViPR Controller.
- You can move a maximum of 100 volumes at a time into a different virtual pool for volumes that belong to the same array group. However, if the volumes are not in a consistency group, there is no maximum volume limitation when moving volumes into a different virtual pool.

In previous versions, all volumes in the same VPLEX consistency group resided in the same backend array group. But, with application services, you can assign volumes in the same VPLEX consistency group to a different array group.

- The *ViPR Controller Support for VPLEX and VPLEX with EMC Data Protection User and Administration Guide* and the *ViPR Controller Service Catalog Reference Guide* found on the [ViPR Controller Product Documentation Index](#) provide further details.

Procedure

1. Go to **Service Catalog > Block Storage Services > Change Virtual Pool**.
2. Select the project to which the virtual pool belongs.
3. Select the current virtual pool.
4. Select the operation to perform to move the virtual pool volumes to another virtual pool.

Option	Description
Change from VPLEX Local to VPLEX Distributed	<p>The volumes in the virtual pool are moved from a local to distributed VPLEX .</p> <hr/> <p>Note</p> <p>Changing the virtual pool from VPLEX local to VPLEX distributed is not supported when the VPLEX volume is in a consistency group.</p> <hr/>
VPLEX Data Migration	<p>By changing the volumes of the virtual pool on the backend storage system, you can perform functions such as; change the backend volumes from thin to thick, change the storage system from VMAX to VNX or vice versa.</p> <p>You can configure the speed of the data migration using Assets > Controller Config > VPLEX and then adding a new configuration for <code>Data Migration Speed</code>.</p> <p>You can also use the REST API and the CLI to retrieve a list of data migrations, show the details of a data migration, pause, resume, and cancel a migration. See the <i>ViPR Controller REST API Reference</i>, available as a zip file from the ViPR Controller Product Documentation Index and the <i>ViPR Controller CLI Reference Guide</i> which can be found on the ViPR Controller Product Documentation Index.</p>

Option	Description
	<p>RecoverPoint protected VPLEX volumes or MetroPoint (VPLEX Metro only) volumes are eligible for VPLEX Data Migration too. For these volumes, the original virtual pool is compared to the target virtual pool and migrations are based on changes in</p> <ul style="list-style-type: none"> • Source virtual pool • Source journal virtual pool • Target virtual pools • Target journal virtual pools <p>Targets and Journals can be implicitly migrated if there are changes in the new virtual pool when compared to other virtual pools. (The other virtual pools must be eligible for migration.)</p> <p>RecoverPoint protected VPLEX volumes or MetroPoint (VPLEX Metro only) volumes that are in consistency groups with array consistency enabled OR are in Applications will be grouped together for migration.</p> <p>The same rules apply to all virtual pools when determining whether or not a migration will be triggered.</p> <p>RecoverPoint or MetroPoint (VPLEX Metro only) Target volumes that are in Applications will be grouped together for migration.</p>
Move Into VPLEX	<p>Moves non-virtualized volumes in a virtual pool from a VNX for Block or VMAX storage system to a VPLEX . Once moved to a VPLEX , the volumes are virtualized. The storage system from which the volume will be moved, must have connectivity to a VPLEX .</p>
Add RecoverPoint Protection	<p>Adds RecoverPoint or MetroPoint (VPLEX Metro only) protection to the virtual pool volumes. Volumes cannot be part of an existing consistency group.</p>
Remove RecoverPoint Protection	<p>Removes RecoverPoint protection from volumes. The source volumes remains intact, but the target volumes are deleted. In addition, if these volumes are the last volumes in the consistency group, then journal volumes are also deleted.</p> <p>The target virtual pool must be identical to the original virtual pool, with the exception that the target virtual pool does not include RecoverPoint protection.</p> <p>You can remove the RecoverPoint protection from VPLEX source volumes, only if those volumes do not have any snapshots.</p>

Option	Description
	<p>Note</p> <p>If you have run the Swap Continuous Copies service to make the RecoverPoint target become the source, you cannot remove the RecoverPoint protection, until you run the Swap Continuous Copies service again to reverse the personalities of the source and target virtual pools.</p>
<p>Change RecoverPoint Protection to MetroPoint</p>	<p>Non-disruptive upgrade of an existing RecoverPoint +VPLEX CRR configuration to MetroPoint CRR. The virtual pool change affects the entire consistency group and all volumes in the consistency group will be moved to the target virtual pool.</p> <p>The target virtual pool must be identical to the original virtual pool, with the exception that the target virtual pool has:</p> <ul style="list-style-type: none"> • A Data Protection setting of VPLEX Distributed. • The target virtual array and virtual pool for the RecoverPoint copy defined in Data Protection > RecoverPoint Copies > Add Copy. Optionally, you can also set a separate virtual array and virtual pool for the journal volume of the RecoverPoint copy. • Protect Source Site selected in Data Protection > RecoverPoint Advanced settings • Protect HA Site selected in Data Protection > RecoverPoint Advanced settings
<p>Add SRDF Protection</p>	<p>Adds SRDF protection to the volumes in the virtual pool from one VMAX to another VMAX.</p>
<p>Add Continuous Copy Protection</p>	<p>Adds Continuous Copy protection to VNX block, VMAX virtual pool volumes, or VPLEX local volumes. When added to the VPLEX local volume a local mirror of the volumes are also created.</p>
<p>Change Export Path Parameters</p>	<p>Moves the virtual pool volumes to a virtual pool with a different number of path attributes.</p>
<p>Change Auto-tiering Policy or Host IO Limits</p>	<p>Changes the auto-tiering policy on the volumes.</p> <ul style="list-style-type: none"> • If you are changing SLOs for all VMAX3 backend volumes for VPLEX volumes within the masking view, then they need not be in parent/child relationship. It can be a flat storage group. • For VMAX2 it does not matter if it is a cascaded storage group (SG) or child, the policy change is requested for all volumes in a SG. If there are phantom SGs (SGs that are non-FAST and non-cascaded), then this restriction is not applicable. <p>Restrictions:</p> <ul style="list-style-type: none"> • If you are changing the SLOs of only a subset of the VMAX3 backend volumes for VPLEX volumes within the

Option	Description
	<p>masking view , then the VMAX3 backend volumes for VPLEX must all be contained in a child storage group, under a parent SG(cascaded). The parent SG should be associated to a Masking View (MV).</p> <hr/> <p>Note</p> <p>If the above restrictions are not followed, an error similar to the following is encountered:</p> <pre data-bbox="772 569 1465 932">Error 12000: An error occurred while executing the job, Op: updateStorageGroupPolicyAndLimits with message None of the Storage Groups on ExportMask BE_Vplex242vmax3_1035_MV1 is updated with new FAST policy or Host IO Limits. Because the given Volume list is not same as the one in Storage Group (or) any of the criteria for 'moveMembers' didn't meet in case of VMAX3 volumes. Please check log for more details.</pre> <hr/> <p>In this example, the VPLEX Backend (BE) masking view is a shared masking view:</p> <ol style="list-style-type: none"> User1 creates 2 Bronze SLO VPLEX virtual volumes which were added to the existing BE Bronze Cascade MV1/Cascade SG1 User2 creates 3 Bronze SLO VPLEX virtual volumes where were added to existing BE Bronze Cascade MV1/Cascade SG1 User 1 tries to change the SLO from Bronze to Gold. The order will fail, and generate the Error 12000 since the Bronze BE SG1 is shared among all users. User 1 is trying to change the SLO of only a subset of the VPLEX virtual volumes in the Bronze SG, instead of all of the volumes in the SG.
<p>Update the replication mode (link policies) on an entire consistency group</p>	<p>When you create a RecoverPoint protected volume, all consistency group link policies are initialized according to the setting in the replication mode field in the virtual pool. Use the change virtual pool operation to update the replication mode (link policies) on the entire consistency group. Before you begin:</p> <ol style="list-style-type: none"> Create a virtual pool with RecoverPoint data protection set to asynchronous replication mode. Create a volume using this virtual pool. Notice that the consistency group link policy in RecoverPoint is set to asynchronous. Create a duplicate virtual pool and change the replication mode to synchronous.

Option	Description
	<p>d. Perform a change virtual pool operation specifying the asynchronous virtual pool as the source and the synchronous virtual pool as the target. After this step, the link policy for the consistency group changes to synchronous.</p> <hr/> <p>Note</p> <p>Changing the replication mode to synchronous can only be performed on a consistency group with no more than one remote copy.</p>

5. Select the target virtual pool.

In addition to the properties to support the selected operation, the following VNX and VMAX block volume attributes can also be changed by moving the volume to a new virtual pool:

- Type of provisioning: Thin or Thick
- FAST policy
- Raid Types

6. If adding RecoverPoint protection, select the consistency group.
7. Select how you want to filter the volumes. The filter retrieves the volumes, sorts them alphabetically, and lists them.
8. Select the volumes.
9. Select **Order**.

Change Volume Virtual Pool

Moves the volume into a different virtual pool.

Before you begin

This operation requires a Tenant Administrator role in ViPR.

Procedure

1. Go to **Catalog > Block Storage Services > Change Volume Virtual Pool**.
2. Select the project in which the volume is located.
3. Select the volume to move.
4. Select the operation to perform by moving the volume to another virtual pool.

Option	Description
<p>Change from VPLEX Local to VPLEX Distributed</p>	<p>The volume is moved from a local to distributed VPLEX .</p> <hr/> <p>Note</p> <p>Changing the virtual pool from VPLEX local to VPLEX distributed is not supported when the VPLEX volume is in a consistency group.</p>

Option	Description
VPLEX Data Migration	<p>By changing the volume on the backend storage system, you can perform functions such as: change the backend volume from thin to thick, change the storage system from VMAX to VNX or vice versa.</p> <p>You can configure the speed of the data migration using Assets > Controller Config > VPLEX and then adding a new configuration for <code>Data Migration Speed</code>.</p> <p>You can also use the REST API and the CLI to retrieve a list of data migrations, show the details of a data migration, pause, resume, and cancel a migration. See the <i>ViPR Controller REST API Reference</i>, available as a zip file from the ViPR Controller Product Documentation Index and the <i>ViPR Controller CLI Reference Guide</i> which can be found on the ViPR Controller Product Documentation Index.</p> <p>RecoverPoint protected VPLEX volumes or MetroPoint (VPLEX Metro only) volumes are eligible for VPLEX Data Migration too. For these volumes, the original virtual pool is compared to the target virtual pool and migrations are based on changes in</p> <ul style="list-style-type: none"> • Source virtual pool • Source journal virtual pool • Target virtual pools • Target journal virtual pools <p>Targets and Journals can be implicitly migrated if there are changes in the new virtual pool when compared to other virtual pools. (The other virtual pools must be eligible for migration.)</p> <p>The same rules apply to all virtual pools when determining whether or not a migration will be triggered.</p> <p>RecoverPoint protected VPLEX volumes or MetroPoint (VPLEX Metro only) volumes that are in consistency groups with array consistency enabled OR are in Applications will be grouped together for migration.</p> <p>RecoverPoint or MetroPoint (VPLEX Metro only) Target volumes that are in Applications will be grouped together for migration.</p>
Move Into VPLEX	<p>Moves a non-virtualized volume from a VNX for Block or VMAX storage system to a VPLEX . Once moved to a VPLEX , the volume is virtualized. The storage system from which the volume will be moved, must have connectivity to a VPLEX .</p>
Add RecoverPoint Protection	<p>Adds RecoverPoint or MetroPoint (VPLEX Metro only) protection to the virtual pool volumes. Volumes cannot be part of an existing consistency group.</p>
Remove RecoverPoint Protection	<p>Removes RecoverPoint protection from volumes. The source volumes remains intact, but the target volumes are deleted. In addition, if these volumes are the last volumes in</p>

Option	Description
	<p>the consistency group, then journal volumes are also deleted.</p> <p>The target virtual pool must be identical to the original virtual pool, with the exception that the target virtual pool does not include RecoverPoint protection.</p> <p>You can remove the RecoverPoint protection from VPLEX source volumes, only if those volumes do not have any snapshots.</p> <hr/> <p>Note</p> <p>If you have run the Swap Continuous Copies service to make the RecoverPoint target become the source, you cannot remove the RecoverPoint protection, until you run the Swap Continuous Copies service again to reverse the personalities of the source and target virtual pools.</p>
Change RecoverPoint Protection to MetroPoint	<p>Non-disruptive upgrade of an existing RecoverPoint +VPLEX CRR configuration to MetroPoint CRR. The virtual pool change affects the entire consistency group and all volumes in the consistency group will be moved to the target virtual pool.</p> <p>The target virtual pool must be identical to the original virtual pool, with the exception that the target virtual pool has:</p> <ul style="list-style-type: none"> • A Data Protection setting of VPLEX Distributed. • The target virtual array and virtual pool for the RecoverPoint copy defined in Data Protection > RecoverPoint Copies > Add Copy. Optionally, you can also set a separate virtual array and virtual pool for the journal volume of the RecoverPoint copy. • Protect Source Site selected in Data Protection > RecoverPoint Advanced settings • Protect HA Site selected in Data Protection > RecoverPoint Advanced settings
Add SRDF Protection	<p>Adds SRDF protection to the volume from one VMAX to another VMAX.</p>
Add Continuous Copy Protection	<p>Adds Continuous Copy protection to VNX block, VMAX volumes, or VPLEX local volumes. When added to the VPLEX local volume a local mirror of the volume is also created.</p>
Change Export Path Parameters	<p>Moves the volume to a virtual pool with a different number of path attributes.</p>
Change auto-tiering Policy, Host IO Limits, or Compression	<p>Can be used to change the auto-tiering policy on the volume.</p> <ul style="list-style-type: none"> • If you are changing SLOs for all VMAX3 backend volumes for VPLEX volumes within the masking view,

Option	Description
	<p>then they need not be in parent/child relationship. It can be a flat storage group.</p> <ul style="list-style-type: none"> For VMAX2 it does not matter if it is a cascaded storage group (SG) or child, the policy change is requested for all volumes in a SG. If there are phantom SGs (SGs that are non-FAST and non-cascaded), then this restriction is not applicable. <p>Restrictions:</p> <ul style="list-style-type: none"> If you are changing the SLOs of only a subset of the VMAX3 backend volumes for VPLEX volumes within the masking view, then the VMAX3 backend volumes for VPLEX must all be contained in a child storage group, under a parent SG(cascaded). The parent SG should be associated to a Masking View (MV). <hr/> <p>Note</p> <p>If the above restrictions are not followed, an error similar to the following is encountered:</p> <pre data-bbox="772 926 1465 1293"> Error 12000: An error occurred while executing the job, Op: updateStorageGroupPolicyAndLimits with message None of the Storage Groups on ExportMask BE_Vplex242vmax3_1035_MV1 is updated with new FAST policy or Host IO Limits. Because the given Volume list is not same as the one in Storage Group (or) any of the criteria for 'moveMembers' didn't meet in case of VMAX3 volumes. Please check log for more details. </pre> <hr/> <p>In this example, the VPLEX Backend (BE) masking view is a shared masking view:</p> <ol style="list-style-type: none"> User1 creates 2 Bronze SLO VPLEX virtual volumes which were added to the existing BE Bronze Cascade MV1/Cascade SG1 User2 creates 3 Bronze SLO VPLEX virtual volumes where were added to existing BE Bronze Cascade MV1/Cascade SG1 User 1 tries to change the SLO from Bronze to Gold. The order will fail, and generate the Error 12000 since the Bronze BE SG1 is shared among all users. User 1 is trying to change the SLO of only a subset of the VPLEX virtual volumes in the Bronze SG, instead of all of the volumes in the SG. <p>Can be used to move VMAX3 volumes to a virtual pool where compression is enabled, or where compression is enabled, and the compression ratio set on the storage pools</p>

Option	Description
	in the virtual pool matches the ratio set on the volume being moved.
Update the replication mode (link policies) on an entire consistency group	<p>When you create a RecoverPoint protected volume, all consistency group link policies are initialized according to the setting in the replication mode field in the virtual pool. Use the change virtual pool operation to update the replication mode (link policies) on the entire consistency group. Before you begin:</p> <ol style="list-style-type: none"> a. Create a virtual pool with RecoverPoint data protection set to asynchronous replication mode. b. Create a volume using this virtual pool. Notice that the consistency group link policy in RecoverPoint is set to asynchronous replication mode. c. Create a duplicate virtual pool and change the replication mode to synchronous. d. Perform a change virtual pool operation specifying the asynchronous virtual pool as the source and the synchronous virtual pool as the target. After this step, the link policy for the consistency group changes to synchronous. <hr/> <p>Note</p> <p>Changing the replication mode to synchronous can only be performed on a consistency group with no more than one remote copy.</p>

5. Select the target virtual pool.

In addition to the properties to support the selected operation, the following VNX and VMAX block volume attributes can also be changed by moving the volume to a new virtual pool:

- Type of provisioning: Thin or Thick
- FAST policy
- Raid Types

6. If adding RecoverPoint protection, select the consistency group.

7. Click **Order** to run the service.

Change virtual array

The Change Virtual Array service is designed to use in a VPLEX environment.

Before you begin

- This operation requires the Tenant Administrator role in ViPR Controller.
- You can move a maximum of 100 volumes at a time into a different virtual array for volumes that belong to the same array group. However, if the volumes are not in a consistency group, there is no maximum volume limitation when moving volumes into a different virtual array.

In previous versions, all volumes in the same VPLEX consistency group resided in the same backend array group. But, with application services, you can assign volumes in the same VPLEX consistency group to a different array group.

- This service is only supported in these VPLEX Metro configurations:
 - on local virtual volumes that have not been exported to a host
 - on local virtual volumes that have been exported to a host when the volumes' exported paths (VPLEX Storage Ports and Host Initiators) are also in the target virtual array.
- The new block storage volume is created from the same virtual pool as the original block storage volume. Therefore the new virtual array being chosen must be configured with the same virtual pool as the original virtual array that was used for the VPLEX virtual volume.

Procedure

1. Select **User > Service Catalog > Block Storage Services > Change Virtual Array**.
2. Select the project to which the volume belongs.
3. Select the volume to move.
4. Select the target virtual array; the virtual array to which the volume will be moved.

Export path adjustment

Adjust the fiber channel paths on an existing VPLEX or VMAX export. You can increase or decrease the number of paths, the number of paths per initiator, or the minimum number of paths. Use the Preview option before finalizing changes to paths. This feature is supported for VPLEX and VMAX exports.

Before you begin

This operation requires the Tenant Administrator role in ViPR Controller.

If you want to adjust the path for an export where there are pre-existing masking views for a host or cluster in the export group, first ingest any volumes exported to the host. Then adjust paths for the exports. If path adjustment is done without ingestion, there is no way to decrease the number of paths in the export group. You can only increase the number of paths (ports). After deleting an export group, the ports or zones added by ViPR Controller are not touched. The ports and zones continue to exist in the pre-existing masking view. This indirectly increases the number of paths for pre-existing volumes.

Procedure

1. Go to **Service Catalog > Block Storage Services > Export Path Adjustment**.
2. Select the project to which the export group belongs.
3. Select the storage type.

Choose **Exclusive** for exports to a single host. Choose **Shared** for exports to a cluster.
4. Choose the export group where you want to adjust paths.
5. Enter the minimum and maximum number of paths allowed from the initiator to the Storage Port on the storage array.

ViPR Controller tries to provision the maximum number of paths possible without exceeding the setting in the Maximum Paths field. (Factors influencing

the number of paths include limitations of the number of initiators, the Paths Per Initiator setting, and the number of usable Storage Ports.)

The maximum number of paths provides an upper boundary on the number of Storage Ports that will be used. For Storage Ports to be usable, they must:

- be in the selected virtual array
- be accessible to the initiator across the network
- be registered on the switch
- be online showing a good status
- not exceed a user-provisioned usage ceiling. (The settings for initiator or volume ceilings set in the **Physical > Controller configurations** pages will not be overridden by changes in path adjustment settings.)

6. Enter the number of storage port paths that may be provisioned for each initiator.

This controls how many ports will be assigned to each initiator. If too few ports are available, the initiator will not be provisioned.

7. Enter **Yes** to use the existing export paths and/or add new paths.

When set to **Yes**, you can add paths. You cannot decrease "Maximum Paths" or "Paths Per Initiator." When adding new Storage Ports due to overutilization of existing ports, set this to **Yes** so that more paths will be added to balance utilization. Paths will not be removed with this option.

Enter **No** if you want new export paths to be calculated using current port usage data. This may cause existing paths to be removed. New paths will be added first, followed by the removal of paths that are no longer needed. Click **Preview** to view the impact on existing paths. Ensure hosts are able to use the newly provisioned paths before completing the order.

8. Enter the storage system which needs a path adjustment.

When exports from multiple storage systems are in the export group, only the paths of this storage system will be adjusted.

9. Select the virtual array containing the export group to be adjusted.

For VPLEX cross-connected exports, you may select the virtual array of either VPLEX cluster. The adjustment occurs only for paths from the corresponding VPLEX cluster.

10. Do not select any ports unless specific ones are to be used in the export.

If ports are selected, they will be added to the usable storage ports to determine the final list of ports available for allocation. If an insufficient number of ports are selected, the order may fail.

11. Check **Suspend before removing paths** if you want to verify that the new paths are usable on the hosts before completing the order.

If selected, masking and zoning operations are performed to add the new paths to be provisioned. A host rescan will be performed if possible. After verifying the hosts have connectivity by way of the new paths, click **Resume** to remove any paths which are no longer needed. Or click **Cancel** to stop the order. This action cancels the steps in the order to remove unneeded paths, leaving the paths as they are after addition of any new paths.

Note

ViPR Controller can't initiate a host re-scan on hosts that were manually added to ViPR Controller. To avoid loss of connectivity, turn on **Suspend before removing paths** when the service is used for manually discovered hosts.

Results

View the changes made to the paths by looking at **Resources > Volumes**.

Block storage protection services

The ViPR Controller block protection services use various technologies, such as RecoverPoint, SRDF and TimeFinder, to replicate and secure data on storage systems.

These services provide protection for volumes. Before you run any of these protection services, export the volumes to a host.

To access these services, go to **Service Catalog > View Catalog > Block Protection Services**. After entering the required parameters, select **Order** to order the service.

Note

When remounting a source or target volume using one of the mounting existing volume services, make sure you deselect the **Format Volume** checkbox to ensure that ViPR Controller does not remove data from these volumes.

Add journal capacity

Adds RecoverPoint journal capacity to a RecoverPoint consistency group. You can use a different virtual array and virtual pool from the ones used for the original copy creation.

Before you begin

When logged into ViPR Controller with a user role you can only perform operations on resources belonging to projects that you are assigned to (or are the owner of). If you are a Tenant Administrator you can run all user services and choose resources from any project.

Procedure

1. Go to **Service Catalog > Block Protection Services > Add Journal Capacity**.
 2. Select the project that owns the consistency group to which you want to add journal capacity.
 3. Select the consistency group.
 4. Select the copy name.
-

Note

Only the copies associated with the selected consistency group are available for selection.

5. Select the virtual array.
6. Select the virtual pool.

7. Select the number of volumes being added.
8. Select the size of each volume.
9. Select **Order**.

Failover Block Volume

After building a block volume, you can choose to provide access to an image at the remote site using the Failover Block Volume service.

Procedure

1. Navigate to **Catalog > View Catalog > Block Protection Services > Failover Block Volume**.
2. Select the project that owns the volume or consistency group.
3. Select either `Volume` or `Consistency Group` for the **Storage Type** field.
4. Select the volume or consistency group, depending on your selection in step 3.
5. In the **Failover Target** field, select the protection array.
6. In the **Image to Access** field, select the image to failover.
 - If **A Specific Point In Time** is selected in the **Image to Access** field, the **Point in Time** field will be used. The date/time value specified in the **Point in Time** field is ALWAYS specified in local time (browser's timezone). If no value is provided in the **Point in Time** field, the current date/time will be used.

Note

The UI transforms the date/time value to GMT/UTC in the format required by the API. For API or CLI calls, the `Point in Time` parameter takes any point in time used for failover, specified in GMT/UTC. Allowed values: `"yyyy-MM-dd_HH:mm:ss"` with the formatted date or datetime in milliseconds.

- The **Image to Access** and **Point in Time** fields are ignored for non-RecoverPoint protected volumes.
- Select **Enable Direct Access** for RecoverPoint volume or consistency group failover if you are worried about the journal volume running out of room during the failover. By default, ViPR Controller failover (using the RecoverPoint test copy) enables image access for a target copy in logged access mode. In this mode, all new writes are written to the replica volume and undo information is stored in the image access log, which is located within the journal. With direct access, the journal is not kept and a full sweep is done after direct access mode is complete. This is an option for long term tests where the journal may not have enough space for long term image access mode.

Note

Ensure you deselect this choice when you fail back to the original volume or consistency group.


7. Select **Order**.

Image access changes with the Failover Block Volume service

The **Failover Block Volume** catalog service enables image access on an image at the remote site for a specific point in time.

The following figure illustrates how the system appears to RecoverPoint before running the **Failover Block Volume** catalog service.


Figure 1 RecoverPoint configuration before running the Failover Block Volume service



The following figure illustrates how the system appears to RecoverPoint after running the **Failover Block Volume** catalog service.

Note

Exporting the target volumes to the host at the remote site is a separate operation, and it can be done before the failover.

Figure 2 RecoverPoint configuration after running the Failover Block Volume service**Note**

For ViPR Controller managed RecoverPoint protected volumes, if you mark a RecoverPoint bookmark on a target volume outside of ViPR Controller, return the volume to its original state before continuing to manage this resource using ViPR Controller.

Swap continuous copies

After building a block volume, you can choose to reverse the personalities of the RecoverPoint or SRDF source and target, making the source become the target, and target become the source.

Procedure

1. Navigate to **Catalog > Block Protection Services > Swap Continuous Copies**.
2. Select the **Project**.
3. Select either **Volume** or **Consistency Group** for the **Storage Type**.
4. Select the source volume or consistency group from the **Volume/Consistency Group** field, depending on your selection in step 3.


5. Select the protection array in the **Failover Target** field.
6. Click **Order**.

Source and target changes with the Swap Continuous Copies service

The Swap Continuous Copies service reverses the personalities of the source and target.

The following figure illustrates the RecoverPoint configuration before running the **Swap Continuous Copies** service.


Figure 3 RecoverPoint configuration before running the Swap Continuous Copies service



The following figure illustrates the RecoverPoint configuration after running the **Swap Continuous Copies** service. After running the service, New York is now the source volume and Boston is the target volume.

Note

Exporting the target volumes to the host at the remote site is a separate operation (and it can be done before the swap).

Figure 4 RecoverPoint configuration after running the Swap Continuous Copies service**Note**

For ViPR Controller managed RecoverPoint protected volumes in case of a datacenter disaster: If for any reason you perform a RecoverPoint failover of volumes outside of ViPR Controller, return volumes to original state before continuing to manage these resources using ViPR Controller.

Create block snapshot

Use this server to create a point-in-time copy of a volume or a consistency group, or to create a snapshot session for VMAX3 storage systems.

Procedure

1. Go to **Catalog > Block Protection Services > Create Block Snapshot**.
2. Select the project.
3. Select the type of block storage.
4. Select the volume or consistency group on which to create the snapshot.
5. Select the snapshot type.

Option	Description
None	Use if relinking the snapshot to the same target, or to a different target.
Local Array Snapshot	To create a snapshot on the same storage system on which the source volume or consistency group resides.
Snapshot Session	To create a snapshot session for VMAX3 storage systems.

6. Check **Read only** if you want to create the snapshot with read only permissions. The option is only applicable in XtremIO, or VPLEX with XtremIO environments.
7. Enter a name to identify the snapshot, or snapshot session.
8. Use the **Advanced** options, for VMAX3 only, to create, a new target device for the snapshot session and to link the source volume or consistency group to the snapshot session target.

You may also link a snapshot session at a later time using the Link Block Snapshot service.

Note

Do not use the **Advanced** options when creating a VMAX3 block snapshot for a volume or consistency group. The creation and linking to a target in no-copy mode happens automatically. A separate step to link the snapshot session is not required. Using the **Advanced** options in this situation will have no affect.

Option	Description
Linked Snapshot Target Name	The name of the target snapshot that will be created.
Linked Snapshot Target Count	The number of snapshots to create, and link to the snapshot session.
Copy Mode	Can be Copy or No Copy .

9. Optionally, expand the Scheduler area to schedule the order, and set the following options:

Option	Description
Enable Scheduler	Select to schedule the order.
Start Date/Time	Enter the date and time for the first time the order will be run.
Frequency	Select how often to run the service order.
Recur every	If you have scheduled the service order run indefinitely, or a recurrence, select how often during the frequency to run the order, for example 1 Day would be daily, 2 Days, would schedule it to run every other day.

Option	Description
Number of recurrences	<p>If you selected End after recurrences for the Frequency, you will enter the number of recurrences to schedule here. Once you have reached the given number of recurrences, no more snapshots will be taken for this order.</p> <hr/> <p>Note</p> <p>You cannot set a recurring schedule if you are using the VMAX3, Snapshot Target feature outside of ViPR Controller. The VMAX3 Snapshot Target feature is not supported by the ViPR Controller, Scheduler.</p>
Automatic Expiration	<p>Maximum number of snapshots to keep. Once the retention is met, the oldest snapshot is removed before the new one is created.</p> <p>It is recommended that you do not manually remove any of the snapshots or copies being managed by the ViPR Controller scheduler.</p>

Once a scheduled order is set, it can only be edited by a Tenant Administrator.

Scheduled snapshots are labeled according the name defined in the service order by the user, which can include the '{datetime}' pattern for example:

mysnapshotname-2016-09-21_09:30:06

The date and time will reflect ViPR Controller server side date and time when the scheduled snapshots were created.

10. If you have added a schedule to the order, click **Schedule Order**, otherwise click **Order** to run the order immediately.

This progress of the order displays.

Link Block Snapshot

Use this service to link a target volume to a volume or consistency group snapshot session. You can link the snapshot session to a new target volume you are creating, relink the snapshot session to the same target that it was previously linked to, or link to a different target that is already created in ViPR Controller.

Before you begin

- This service is only available to VMAX3 storage systems, or storage systems on which VMAX3 is used as the backend device such as with VPLEX and RecoverPoint storage systems.
- You must have created the snapshot session in **ViPR Controller** using the Create Block Snapshot service.

A source volume or consistency group can also be relinked to a snapshot session target from the **Snapshot Sessions > Resources** pages.

Procedure

1. Go to **Catalog > Block Protection Services > Link Block Snapshot**.
2. Select the project.

3. Select the type of block storage.
4. Select the volume or consistency group.
5. Select the snapshot session.
6. In the **Linked Snapshot** area, select a snapshot target that was already linked to the same snapshot session, or to a different snapshot session.
7. Use the **Advanced** options to create, a new target device for the snapshot session, and to link the new target device to the volume or consistency group.

Option	Description
Linked Snapshot Target Name	The name of the target snapshot that will be created.
Linked Snapshot Target Count	The number of snapshots to create, and link to the source.
Copy Mode	Can be Copy or No Copy mode.

8. Click **Order**.

This progress of the order displays.

Restore block snapshot

Restores a snapshot's point-in-time data back to a source volume or a consistency group.

Before you begin

- When VNXe Block restores a snapshot, it creates a new snapshot for the volume. As this behavior is specific to VNXe arrays, ViPR Controller also displays these new snapshots.
- When RecoverPoint volumes were ingested:
The ingestion process will not automatically group replicas that were created in a RecoverPoint device/replication group before ingestion. These replicas are consistent with each other and are most conveniently managed as a group, but ViPR Controller cannot identify replicas that were created in a group due to provider limitations.

After ingesting a RecoverPoint consistency group with replicas, all management operations should be performed on all of the snapshots in ViPR Controller to get a consistent result and to ensure that data integrity is upheld. In order to restore the entire snapshot set, restore each snapshot to its respective volume at the same time.

You can still access, mount, delete, restore, and manage the individual snapshots. This is possible because:

- Replicas are ingested with the consistency group.
- Relationships are established in ViPR Controller between each individual replica and its originating volume.

The *ViPR Controller Ingest Services for Existing Environments* provides more information.

- For XtremIO storage systems, when a volume that is not in a consistency group is restored, a snapshot set instance is created on the XtremIO array and it indicates that the snapshot has been restored. However, when all volumes in a CG / CG

snap are restored, there is no snapshot set instance created on the XtremIO array. This is the expected behavior.

Procedure

1. Go to **Catalog > Block Protection Services > Restore Block Snapshot**.
2. Select the project.
3. Select the volume or consistency group to restore.
4. Select the snapshot type.

Option	Description
Local	To restore a local snapshot.
Remote	To restore a remote snapshot.
Snapshot Session	To restore a snapshot session that was previously deleted.

5. Select the snapshot.
6. Click **Order**.

The progress of the order appears.

Remove block snapshot

Use this server to unlink, and delete a point-in-time copy of a volume or a consistency group, or to unlink and delete a snapshot session for VMAX3 storage systems.

Procedure

1. Go to **Catalog > Block Protection Services > Remove Block Snapshot**.
2. Select the project.
3. Select the type of block storage.
4. Select the snapshot type.

Option	Description
None	Use if removing the snapshot from the same target, or from a different target.
Local Array Snapshot	To remove a snapshot from the same storage system on which the source volume or consistency group resides.
Recoverpoint Bookmark	To remove the Recoverpoint Bookmark.
Snapshot Session	To: <ul style="list-style-type: none"> • Unlink the target volume from the snapshot session and also delete the target volume. • Delete the snapshot session.

5. Select the volume or consistency group from which to remove the snapshot.
6. Select the snapshot, or snapshot session to delete.

Create full copy

Creates a full copy of a volume or consistency group.

Before you begin

Configure the virtual pool to allow native snapshots.

Procedure

1. Go to **Service Catalog > Block Protection Services > Create Full Copy**.
2. Select the project.
3. Indicate whether this is a volume or consistency group full copy.
4. Select the volume or consistency group to copy.
5. Enter the name of the full copy.
6. Specify the number of copies to create.
7. Optionally, expand the Scheduler area to schedule the order, and set the following options:

Option	Description
Enable Scheduler	Select to schedule the order.
Start Date/Time	Enter the date and time for the first time the order will be run.
Frequency	Select how often to run the service order.
Recur every	If you have scheduled the service order run indefinitely, or a recurrence, select how often during the frequency to run the order, for example 1 Day would be daily, 2 Days, would schedule it to run every other day.
Number of recurrences	If you selected End after recurrences for the Frequency , you will enter the number of recurrences to schedule here. Once you have reached the given number of recurrences, no more copies will be taken for this order.
Automatic Expiration	<p>Maximum number of copies to keep. Once the retention is met, the oldest copy, or copies are removed before the new one is created.</p> <hr/> <p>Note</p> <p>If you have set the order to create multiple copies, Automatic Expiration applies to all the copies created for the order for example: If you set the order to create 3 full copies, and you set automatic expiration to keep 5 full copies, ViPR Controller, will maintain 15 (3x5) full copies, and delete the three oldest copies after the retention is met.</p>

Once a scheduled order is set, it can only be edited by a Tenant Administrator.

Scheduled full copies are labeled according the name defined in the service order by the user, which can include the '{datetime}' pattern for example:

mycopyname-2016-09-21_09:30:06

The date and time will reflect ViPR Controller server side date and time when the scheduled full copies were created.

8. If you have added a schedule to the order, click **Schedule Order**, otherwise click **Order** to run the order immediately.

This progress of the order displays.

9. Go to **User > Resources > Block Volumes** to verify the full copy. Multiple copies appear with a number suffix.

Remove full copies

Removes full copies from a volume or consistency group.

Procedure

1. Go to **Service Catalog > Block Protection Services > Remove Full Copies**.
2. Select the project.
3. Select the volume or consistency group from which to remove the full copies.
4. Type the name of the full copies to remove.
5. Click **Order**.

Restore from full copies

Restores a source volume or consistency group with the latest data from a full copy.

When a volume is linked to a full copy, you cannot expand the source volume or the full copy.

Procedure

1. Go to **Service Catalog > Block Protection Services > Restore From Full Copies**.
2. Select the project.
3. Select the full copy source volume or consistency group to restore.
4. Select the full copy containing the data to restore.
5. Select **Order**.

The progress of the order displays.

Resynchronize full copies

Copies the latest data from a source volume or a consistency group to a full copy.

When a volume is linked to a full copy, you cannot expand the source volume or the full copy.

Procedure

1. Go to **Service Catalog > Block Protection Services > Resynchronize Full Copies**.
2. Select the project.
3. Select the full copy source volume or consistency group containing the data to copy.
4. Select the full copy to synchronize.
5. Select **Order**.

The progress of the order displays.

Detach full copies

Removes the source and target relationship of a copy session for volumes or consistency groups. After you order this service, you cannot run the restore full copy service or the synchronize full copy service against the source volume, consistency group, or the full copy.

Before you begin

After ViPR Controller detaches and removes clones from a VMAX3 volume, it cannot create continuous copies on this volume due to an SMI-S Provider limitation.

Procedure

1. Go to **Catalog > Block Protection Services > Detach Full Copies**.
2. Select the project.
3. Select a source volume or consistency group.
4. Select the full copy to detach.

After you order this service, the full copy becomes detached from its source volume or consistency group and no longer appears in this list.

5. Select **Order**.

The progress of the order displays. Once complete, you can expand, export, unexport, and delete a detached full copy volume.

Create snapshot full copy

Creates one or more full copies of a snapshot.

Before you begin

Configure the associated virtual pool to allow for native snapshots.

Note

You cannot create a full copy of a volume or a snapshot if it is in a consistency group.

Procedure

1. Go to **Service Catalog > Block Protection Services > Create Snapshot Full Copy..**
2. Select the project.
3. Select the volume on which the snapshot was created.
4. Type the name of the snapshot from which the full copy will be created.
5. Type a name of the full copy.
6. Type the number of full copies to create.
7. Click **Order**.

This progress of the order displays.

Create continuous copy

Creates ongoing mirroring of data from a source volume to a target volume.

Before you begin

Configure the associated virtual pool for continuous copies.

Procedure

1. Go to **Service Catalog > Block Protection Services > Create Continuous Copy**.
2. Select the project.
3. Select the volume.
4. Specify the number of copies to create.
5. Select **Order**.
The progress of the order displays.
6. Go to **User > Resources > Block Volumes** to view the continuous copies.

Export Continuous Copy to a Host

Export a continuous copy (target) volume, of a block volume (source) to a host.

Procedure

1. Select **Catalog > Block Protection Services > Export Continuous Copy to a Host**.
2. Select the storage type:

Option	Description
Exclusive	Exports the continuous copy to a single host.
Shared	Exports the continuous copy to an entire cluster.

3. Select the host or cluster to which the continuous copy will be exported.
4. Select the project from which the continuous copy will be selected.
5. Select the source volumes for which the continuous copies were created.
6. Enter one or more continuous copies to export to the host to make the copy available to the host.
7. Use the default host logical unit number or type another number.
For IBM XIV storage systems, do not specify the value of zero.
8. Optionally, expand **Advanced** to override the values set for Minimum Paths, Maximum Paths, and Paths Per Initiator on the volume's virtual pool. These values are used when exporting the volume to the host or cluster.
9. If you would like to schedule the service order to run at a later time, enable the Scheduler, define the time, and click **Schedule Order**, otherwise, click **Order** to run the service order now.

Unexport Continuous Copy

Remove a block volume continuous copy from an export. The volume will still be accessible using another export.

Procedure

1. Select **Catalog > Block Protection Services > Unexport Continuous Copy**.
2. Select the **Project** to which the volumes belong.
3. Select the source volumes for which the continuous copies were created.
4. Enter the continuous copy to unexport.
5. Select the **Export** from which to remove the continuous copy.
6. If you would like to schedule the service order to run at a later time, enable the **Scheduler**, define the time, and click **Schedule Order**, otherwise, click **Order** to run the service order now.
7. When the order is complete, view Exports from the **Resources** pages to observe that the continuous copies have been removed.

Remove continuous copies

Stops the ongoing mirroring of data from a source volume to a target volume.

Before you begin

Unexport the volume storing the continuous copies and then pause continuous copies.

Procedure

1. Go to **Service Catalog > Block Protection Services > Remove Continuous Copy**.
2. Select the project.
3. Select the volume for which the continuous copy was created.
4. Enter the continuous copy to remove.
5. Select **Order**

Export snapshot to a host

Exports a snapshot to a host or cluster.

Procedure

1. Go to **Catalog > Block Protection Services > Export Snapshot to a Host**.
2. Select the storage type:

Option	Description
Exclusive	Exports the snapshot to a single host.
Shared	Exports the snapshot to an entire cluster.

3. Select the host or cluster on which to export the snapshot.
4. Select the project in which the snapshot belongs.
5. Select the snapshot to export.
6. Use the default host logical unit number or type another number.

For IBM XIV storage systems, do not specify the value of zero.

7. Click **Advanced**. You have the option to override the values set for Minimum Paths, Maximum Paths, and Paths Per Initiator on the volume's virtual pool. These values are used when exporting the snapshot to the host or cluster.
8. If you would like to schedule the service order to run at a later time, enable the **Scheduler**, define the time, and click **Schedule Order**, otherwise, click **Order** to run the service order now.

Unexport snapshot

Remove a volume snapshot from an export. The volumes will still be accessible using another export.

Before you begin

When logged into ViPR Controller with a user role you can only perform operations on resources belonging to projects that you are assigned to (or are the owner of). If you are a Tenant Administrator you can run all user services and choose resources from any project.

Procedure

1. Select **User > Service Catalog > Block Protection Services > Unexport Snapshot**.
2. Select the **Project** to which the volumes belong.
3. Select the **Snapshot** you want to remove from an export.
4. Select the **Export** from which to remove the snapshot.
5. Click **Order**.

The Orders page is displayed with the progress of the order.

6. When the order is complete, view Exports from the **Resources** pages to observe that the snapshots have been removed.

Create VPLEX Volume from Block Snapshot

The **Create VPLEX Volume from Block Snapshot** service allows you to create a VPLEX virtual volume from a block snapshot. You may then export this VPLEX virtual volume using the **Export Block Volume** service. In this way, you can export the snapshot to the host/cluster through the VPLEX rather than through the backend storage array.

Before you begin

- You can only create a VPLEX volume from a block snapshot when the source volume for the block snapshot is a VPLEX backend volume. In other words, the block snapshot must be a snapshot of a VPLEX volume.
- The VPLEX virtual volume so created can only be exported, unexported, and deleted. No other ViPR Controller service are supported on the VPLEX virtual volume.
- If a VPLEX volume is created from a block snapshot, the VPLEX volume must be deleted prior to deleting the block snapshot.
- If the block snapshot is of a VPLEX local virtual volume, then using the new service to create a VPLEX virtual volume from this snapshot will result in a local virtual volume. If the block snapshot is of a VPLEX distributed virtual volume, then using this new feature to create a VPLEX volume from the snapshot will result is a VPLEX distributed virtual volume.

- Note**

When a snapshot is created of a VPLEX distributed volume, ViPR Controller will only create a native block snapshot of the source-side backend volume. The source-side backend volume is the backend volume in the same virtual array as the VPLEX virtual volume and is the virtual array specified when the VPLEX volume was created.

Procedure

1. Select **Catalog > View Catalog > Block Protection Services > Create VPLEX Volume from Block Snapshot**.
2. Select the project containing the snapshot.
3. Select the **Storage Type: Volume or Consistency Group**.
4. If you selected **Volume** in step three, choose a volume from the list of VPLEX virtual volumes to choose from the list of the backend snapshots associated with the VPLEX virtual volume in step 5.

If you selected **Consistency Group** in step 3, choose a consistency group from the list to select from the snapshots associated with the consistency group in step 5.
5. Select the **Snapshot** from which to create the VPLEX volume.

Resynchronize Block Snapshot (VMAX2 and XtremIO only)

Use this service to resynchronize a snapshot of a block volume or consistency group for VMAX2 or XtremIO storage systems.

Procedure

1. Go to **Catalog > Block Protection Services > Resynchronize Block Snapshot**.
2. Select the project.
3. Select the storage type.
4. Select the volume or consistency group.
5. Select the snapshots.
6. Select **Order**.

The progress of the order displays.

CHAPTER 4

ViPR Controller File Storage Services

This chapter includes the following topics:

- [File storage services](#)70
- [File storage provisioning support](#)70
- [File provisioning services for VMware vCenter](#)78
- [File storage services for unmanaged file systems](#)78
- [File system protection support](#)80

File storage services

ViPR Controller supports the provisioning of file systems and the protection of file storage using snapshots.

The file storage services include:

- [File storage provisioning support](#) on page 70
- [File provisioning services for VMware vCenter](#) on page 78
- [File storage services for unmanaged file systems](#)
- [File system protection support](#) on page 80

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.

File storage provisioning support

ViPR Controller enables file systems to be created and made available as CIFS shares or NFS exports.

The services provided in the Service Catalog enable a file system to be created first and subsequently shared using CIFS or NFS protocols, alternatively, file systems can be created and shared in a single operation.

To access these services, go to **Catalog > View Catalog > File System Services**.

Table 8 File system provisioning services

Category	Service function	Services
File Storage Services	Create and manage file systems on page 71	Create File System
		Expand File System
		Remove File System
	Create and manage CIFS shares on page 72	Create File System and CIFS Share
		Create CIFS Share for File System
		Remove CIFS Share for File System
	Create and manage NFS exports on page 74	Create File System and NFS Export

Table 8 File system provisioning services (continued)

Category	Service function	Services
		Create NFS Export for File System
		Remove NFS Export for File System
	Create and remove file system quotas on page 77	Create File System Quota Directory
		Remove File System Quota Directory

Set the smart quota on an Isilon file system

You can use ViPR Controller to set the following smart quota limits at the file system, and quota directory level of Isilon storage systems managed by ViPR Controller.

- **Advisory Limit (%)** — Once this limit is reached, ViPR Controller displays an alert on the file system, Resources page, and a notification is sent to the configured users from the storage system.
- **Soft Limit (%)** — ViPR Controller issues a, **Capacity Exceeded**, warning when this limit is reached. Writes to the file system are allowed to continue until the grace period reached.
- **Grace Period (Days)** — The number of days after the soft limit is reached before writes to the file system will be stopped. Once the grace period is exceeded writes to the file system are stopped.

Smart quotas can be set from the **Advanced** area of the following file system provisioning services.

- Create File System
- Create File System and CIFS Share
- Create File System and NFS Export

All alerts are displayed in the **Details** area of the **Resources** page for the file system.

Create and manage file systems

ViPR Controller provides the ability to create file systems and to make them available as CIFS shares and NFS exports, or both.

ViPR Controller supports concurrent file system provisioning operations in VNX for one or multiple arrays for the following services:

- Create a File System and CIFS Share
- Create a File System and NFS Export
- Remove a File System (with CIFS/NFS)

The following services are provided to enable the creation and management of file systems.

Table 9 Services to create and manage file systems

Service	Description
Create a File System	Enables you to create a new file system from a specified file virtual pool. The file system can be made available as a CIFS share or NFS export, or both.
Expand File System	Enables you to expand an existing file system.
Remove File System	Enables you to remove a file system. You can delete the file system from the ViPR Controller database (Inventory Only) or from both the ViPR Controller database and its backend storage system (Full). A Full delete removes the file system and all objects referencing the file system, such as CIFS shares, snapshots, and quota directory, from the ViPR Controller database and its backend storage system. An Inventory Only delete removes the file system and all objects referencing the file system from the ViPR Controller database.

If you run the Create a File System service, the file system will create either CIFS shares, or NFS exports, or both depending on the settings in your virtual pool. If the virtual pool is set to both CIFS, and NFS, but the storage system is not enabled for NFS, then at the time the service is run the CIFS share is created, no NFS export is created, and the following error is returned:

```
com.emc.vipr.client.exceptions.ServiceErrorException: Error 1034
(http: 400): An error occurred while finding a suitable placement to
handle the request. No Storage Port was assigned to virtual array...
```

The **Resources > File Systems** area enables you to view information about a file share and how it has been made available as a CIFS share or NFS export.

Expanding File Systems on EMC Isilon

When expanding a file system that resides on an EMC Isilon array, you can specify a file system size that is larger than the available capacity.

All Isilon file systems are thinly provisioned and do not consume any capacity when created. Hence, when extending a file system, a client may violate the limit, but the array will always alert when the array is near full utilization.

Provisioning error for vNAS servers

If a vNAS server is in an invalid state, such as the unloaded state, or was deleted from its storage system, ViPR Controller is unable to detect this until the next array discovery. ViPR Controller still selects these vNAS servers for provisioning, resulting in an error. You can run a provisioning operation again after the storage system has been rediscovered.

Create and manage CIFS shares

You can expose a previously created file system as a CIFS share, or you can create a file system and a CIFS share in a single operation. The service catalog also allows you to delete a CIFS share.

The following services support the creation and management of CIFS shares.

Table 10 Services for creating and managing CIFS shares

Service	Description
Create File System and CIFS Share	Enables you to create a new file system from a file virtual pool and to share the file system as a CIFS share.
Create CIFS Share for File System	Enables you to share an existing file system as a CIFS share.
Remove CIFS Share for File System	Enables the sharing of a file system using CIFS to be removed. The Resources > File Systems page also enables the CIFS shares associated with a file system to be listed and individual shares to be deleted.

Notes on creating CIFS shares are provided in [CIFS share notes](#) on page 73.

CIFS share notes

The following notes apply to the creation of CIFS shares from file systems and file system snapshots.

Share names

Valid characters that can be used for CIFS share names are alphanumeric characters, and "_" (underscore).

CIFS shares and permissions

You can set these permissions for CIFS shares:

Table 11 Permissions for setting CIFS shares

Array	File System Export	Snapshot Export
Isilon	Read,Change,FullControl	Not supported
VNX	Read,Change,FullControl	Read
NetApp 7-mode	Read,Change,FullControl	Read
NetApp Cluster-mode	Read,Change,FullControl	Read
VNXe	Read,Change,FullControl	Read
Data Domain	Read,Change,FullControl	Not supported

Access Permissions

For EMC Isilon, NetApp 7-mode, and NetApp Cluster-mode file storage systems the following Access Control List (ACL) functionality is supported:

- Use ViPR Controller to add, modify, and delete permissions for a user or group on CIFS share.
- ACLs are discovered, and ingested with discovery, and ingestion of unmanaged filesystems.

For Data Domain, VNX for File, and VNXe:

- Default access permissions are enforced when creating CIFS shares from the ViPR Controller.
- Access permissions for CIFS shares must be configured using Access Control Lists on the storage system that provides the file system.

- ACLs are not discovered and ingested with discovery of unmanaged filesystems.

Mapping a CIFS Share

To map the Windows share as a network drive, you can look at the order or at the **Resources > File Systems** page to find its location.

The share location is shown in the format:

```
\\ComputerNameFQDNorIP\SharedFolder
```

or

```
\\netbios name\SharedFolder
```

For example:

```
\\vipr-isi6132.lss.emc.com\tcshare01
```

is the FQDN location, where

```
\\LGLW6204\etapp7thickjune18shr1
```

is the netbios name configured on the array. The netbios name is only used if it is configured on the storage system. If it is not configured on the storage system, then the FQDN or CS IP will be shown as mount path.

Create and manage NFS exports

You can expose a previously created file system as an NFS export or you can create a file system and NFS export in a single operation.

The following services are provided to enable you to create and manage NFS exports.

Table 12 Services for creating and managing NFS exports for file systems

Service	Description
Create NFS Export for a File System	Enables you to create an NFS export for an existing file system.
Create File System and NFS Export	Enables you to create a new file system from a file virtual pool and to make the file system available as an NFS export.
Remove NFS Export for a File System	Enables you to remove an NFS export for a file system.
Mount an NFS export	Mount a previously created NFS Export of a file system to a Linux host.
Unmount an NFS export	Unmount a previously mounted NFS Export on a Linux host.
Create a file system, NFS export, and mount it	Create a new file system, create an NFS export, and mount it to a Linux host.

In addition the following operations can be performed from the **Resources > File Systems** page.

Table 13 Services for managing NFS exports

Service	Description
Add an Export Rule	Enables you to add additional export rules to those that already exist.
Modify (an Export Rule)	Enables you to change an existing export rule. You can add additional hosts (endpoints), change the security type (including selecting multiple security types), and change the permissions associated with the host.
Delete (an Export Rule)	Enables you to delete an export rule.

General notes on creating NFS exports are provided in [NFS export notes](#) on page 75 and any constraints on creating exports for each file storage system supported by ViPR are described in [NFS export rules and permissions](#) on page 76.

NFS export notes

The following notes apply to the creation of NFS exports from file systems (and file system sub-directories) and file system snapshots.

Sub-directory Exports

You can export sub-directories of a file system. The sub-directory must already exist and cannot be created from within ViPR Controller.

Snapshots of file system sub-directories are not supported.

ViPR Controller requirement for SUSE Linux Enterprise Server (SLES)12

If you plan to use ViPR Controller UI to export or mount a volume to a Linux host running SLES 12, you must first manually create and mount one `ext3` or `ext4` filesystem on the host, prior to exporting or mounting a volume using the ViPR Controller UI .

Note

Manually creating an `ext3` or `ext4` filesystem is only required if you are using the ViPR Controller UI to export or mount a volume. You do not need to create the filesystem prior to exporting or mounting a volume, if you are using the ViPR Controller REST API.

Adding and Modifying an Export Rule

In general, you can add an export rule for each file system or file system snapshot for each security type. The export rules that can be created for an export depend on the storage system. See [NFS export rules and permissions](#) on page 76 for details.

To add an export rule you need to:

1. Specify the hosts that can access the export. To specify more than one Export Host for a rule, the IP addresses or FQDNs of the hosts can be entered as a comma separated list.
2. Specify the Security Type for the export rule.
3. Set permissions to specify access that clients will have to the export: Read/Write (`rw`), Read Only (`ro`), or Root (`root`).

Modifying an export rule allows you to:

- Add additional endpoints that can access the share

- Specify the effective user id of anonymous users
- Specify the permissions for each export. This allows more granular setting of permissions than those configured when the NFS export is created.

Mounting an NFS Export

To mount the NFS export, you can look at the order or at the **Resources > File Systems** page or, for a snapshot, at the **Resources > File Snapshots** page, to find its mount point.

For example:

```
vi-pr-isi6132.lss.emc.com:/ifs/vi-pr/Isilon_Pool/Provider_Tenant/ProjectA/myNFSExport
```

If you have created a mount-point directory on your system (for example, `/mnt/my nfs`) the export can then be mounted using:

```
mount -t nfs vi-pr-isi6132.lss.emc.com:/ifs/vi-pr/Isilon_Pool/Provider_Tenant/ProjectA/myNFSExport /mnt/my nfs
```

Unmounting an NFS Export outside of ViPR Controller

It is recommended that you do not use applications outside of ViPR Controller to manage file systems outside of ViPR Controller, once a file system, and its exports are under ViPR Controller management. However, in the event that an NFS Export, which is under ViPR Controller management has been unmounted outside of ViPR Controller, you will need to also unmount the NFS export in ViPR Controller to ensure that the storage system database, and ViPR Controller database is in a consistent state.

NFS export rules and permissions

File systems, file system sub-directories, and file system snapshots can be exported as NFS exports and access to an exported file system depends on the security type and the permissions assigned.

The security types supported and the rules that can be created on each supported array are detailed below.

Isilon

Supports `sys`, `krb5`, `krb5p`, `krb5i` security types, but allows only one rule to be set. If you have set a rule for `sys`, for example, you cannot set a further rule for another security type.

VNX

Supports `sys`, `krb5`, `krb5p`, `krb5i` security types and allows one rule per security type.

NetApp

Supports `sys`, `krb5`, `krb5p`, `krb5i` security types and allows one rule per security type.

VNXe

Supports `sys` security type, with one rule.

Data Domain

Supports `sys` and `krb5` security type, with one rule per security type.

You can set these permissions on an export:

Table 14 Permissions to set on an export

Array	File System Export	Snapshot Export
Isilon	rw, ro, root	Not supported
VNX	rw, ro, root	ro
NetApp 7-mode	rw, ro, root	ro
NetApp Cluster-mode	rw, ro, root	Not supported
VNXe	rw, root (See Note)	ro
Data Domain	rw, ro, root	Not supported

Note

For VNXe, a file system or sub-directory export must have root permission in order to mount it on a host and write data to it. If an export has read-write permissions, you can mount it but cannot write data to it.

Create and remove file system quotas

You can create a quota directory at the root of a file system.

In NetApp this is a QTree, for VNX File this is a Quota Tree, and for Isilon this is a Sub-directory with Quota. The size of the directory is a Group Quota Hard limit.

In NetApp Cluster-mode Data ONTAP, for an NFS client to mount a qtree, the NFS client must have read-only permissions at all the parent junction paths up to the Storage Virtual Machines (SVMs) root file system junction path (that is, /). For NFS clients to mount qtrees, the qtrees must belong to a file system that has read-only permissions. Without the read-only permissions at the file system, the NFS clients cannot mount the qtree.

The Service Catalog provides the following services to enable the creation and deletion of file system quotas:

Table 15 Services for creating and deleting file system quotas

Service	Description
Create File System Quota Directory	Enables you to create a quota directory at the root of the file system.
Remove File System Quota Directory	Enables you to remove a quota directory.

In addition, the **Resources > File Systems** page enables the quota directories associated with a file system to be displayed and for the following operations to be performed.

Table 16 Services for managing the quota directory

Operation	Description
Modify (Quota Directory)	Enables you to modify the setting for a quota directory: set the security style and change the size.
Delete (Quota Directory)	Enables you to remove a quota directory.

Modifying a quota directory enables the size to be changed and, for NetApp, enables the Security Style and Read/Write Cache (oplock) to be enabled.

File provisioning services for VMware vCenter

You can create a file systems and attach it as an VMware datastore on an ESX host as well as create a datastore from an existing NFS export.

To access these services, go to **Service Catalog > View Catalog > File Services for VMware vCenter**.

Table 17 File system provisioning services for VMware

Service	Description
Create File System and NFS Datastore	Creates an NFS export and mounts it to an ESX host as a datastore.
Create VMware NFS Datastore	Creates a datastore from an existing NFS export.
Delete NFS Datastore and File System	Removes a datastore, its underlying file system, and the associated NFS export.
Delete VMware NFS Datastore	Deletes a VMware datastore, leaving the NFS export intact.

File storage services for unmanaged file systems

Unmanaged file systems, are file systems that exist on a storage system discovered by ViPR Controller, but were not created or discovered by ViPR Controller. The discover, and ingest services are used by ViPR Controller to get the unmanaged file systems under ViPR Controller management.

Table 18 File storage services for unmanaged file systems

Service category	Service name	Description
File storage services	Discover unmanaged file systems	<p>Finds file systems which are not under ViPR Controller management and matches them to a ViPR Controller virtual pool. When performing discovery it is important to note:</p> <ul style="list-style-type: none"> The virtual array and virtual pool into which you want to ingest the storage pools must exist when the discovery is performed.

Table 18 File storage services for unmanaged file systems (continued)

Service category	Service name	Description
		<ul style="list-style-type: none"> There must be at least one virtual pool in ViPR Controller that matches the physical storage pool that contains the volume.
	Ingest file systems	Used to ingest file systems. Refer to the Ingest File System Support table below.
	Change file virtual pool	Used to move file systems enabled for replication (continuous copy) into a replication enabled file virtual pool. Support for replication and continuous copies for Isilon file systems with SyncIQ was introduced in ViPR Controller 3.0. If you had created or ingested file systems in previous version of ViPR Controller that were enabled for SyncIQ, you can use this service to move the replication enabled file systems to a replication enabled virtual pool.

When discovering or ingesting unmanaged file systems:

- The services for unmanaged file systems can only be performed by ViPR Controller system administrators.
- The operations can be performed from the ViPR Controller UI, API and CLI.
- You must run the Discover Unmanaged File Systems service on the storage system before running an ingest service.
- 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 each discovered file system matches.
- To be ingested, the unmanaged file systems must be in physical pools which are already associated with a ViPR Controller virtual storage pool.
- 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.
- ACLs are discovered, and ingested with discovery, and ingestion of unmanaged filesystems for EMC Isilon, NetApp 7-mode, and NetApp Cluster-mode file storage systems.
- ViPR Controller will also discover and ingest access controls set on the sub-directories of Isilon file systems enabled with NFSv4. Once ingested ViPR

Controller allows you edit the permissions of ingested access controls, add more access controls to the Access Control List (ACL), and delete access controls from the list.

Ingest file system support

The following table lists the resources that are ingested for a file system for each type of file storage system.

Table 19 Resources ingested for a file system

Resource	Isilon	VNX File	NetApp 7-Mode	NetApp Cluster-Mode	Data Domain	VNXe	Unity
FileSystem without Exports	Yes	Yes	Yes	Yes	Yes	Yes	Yes
FileSystem with Exports	Yes	Yes	Yes	Yes	Yes	Yes	Yes
NFS Exports and rules	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CIFS Shares	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CIFS Share ACL	Yes	No	Yes	Yes	No	No	No
Sub Directory/QD Exports	Yes	Yes	Yes	No	Yes	Yes	Yes
Sub Directory/QD Shares	Yes	Yes	Yes	No	Yes	Yes	Yes
Quota Directory	No	Yes	No	No	No	No	Yes
Snapshots	No	No	No	No	No	No	No
Snapshot NFS Exports	No	No	No	No	No	No	No
Snapshot CIFS Shares	No	No	No	No	No	No	No

Note

File systems with more than one export rule of same security type for an export, and exports other than supported security types (sys, krb5, krb5p, krb5i) will not be ingested.

File system protection support

ViPR Controller supports the creation of file system snapshots and enables previously created file system snapshots to be made available as CIFS shares or NFS exports. you can also associate/dissociate file policies to file systems, create and remove replication copies, and failover/failback file systems.

To access these services, go to **Catalog > View Catalog > File Protection Services**.

Table 20 File system protection services

Category	Service Function	Service
File Protection Services	File system snapshots on page 81	Create File System Snapshot
		Restore File System Snapshot
		Remove File System Snapshot
	Export CIFS snapshots on page 83	Create CIFS Share for Snapshot
		Remove CIFS Share for Snapshot
	Export NFS snapshots on page 82	Create NFS Export for Snapshot
		Remove NFS Export for Snapshot
	File Policy	Associate File Policy to File System
		Dissociate File Policy from File System
	Failover/Failback	Failover File System
Failback File System		

File system snapshots

You can create snapshots of file systems and restore the snapshot.

The following services are provided:

Table 21 Services for managing file system snapshots

Service	Description
Create File System Snapshot	<p>Enables you to create a snapshot of an existing file system.</p> <p>The following prerequisites apply:</p> <ul style="list-style-type: none"> You must have access to the project to which the file system you want to snapshot belongs. The snapshot that you create will also belong to this project. A file system must exist and the file virtual pool from which it was created must be enabled for snapshots. A System Administrator can access the file virtual pool at Virtual Assets > File Virtual Pools, and the number of snapshots allowed for a selected file virtual pool can be set at the Data Protection > Maximum Snapshots field.
Restore File System Snapshot	<p>Enables you to restore a previously created file system snapshot. You must have access to the project to which the file system and its snapshots belong.</p>
Remove File System Snapshot	<p>Enables you to remove a previously created snapshot. You must have access to the project to which the file system and its snapshots belong.</p>

Virtual Pool Maximum Snapshots

To be allowed to create a snapshot of a file system, the file virtual pool that the file system is assigned to must be enabled for snapshot. At the UI, you can access the file virtual pool at **Virtual Assets > File Virtual Pools**, and the number of snapshots allowed for a selected file virtual pool can be set at the **Data Protection > Maximum Snapshots** field.

Snapshot names

Valid characters for the snapshot name are numbers, the English alphabet, and the underscore (_).

Export NFS snapshots

You can make snapshots available as NFS exports.

The following services support the management of NFS mounted snapshots:

Note

Create, modify, and remove NFS exports is not supported for NetApp Cluster-Mode.

Table 22 Services for managing NFS mounted snapshots

Service	Description
Create NFS Export for Snapshot	Enables you to select a previously created snapshot and make it available as an NFS export.
Remove NFS Export for Snapshot	Enables you to remove an NFS export for a snapshot.

In addition the following operations can be performed from the **Resources > File Snapshots** page.

Table 23 Services for managing export rules

Service	Description
Add an Export Rule	Enables you to add additional export rules to those that already exist.
Modify (an Export Rule)	Enables you to change an existing export rule.
Delete (an Export Rule)	Enables you to delete an export rule.

Note

For VNXe, it is not possible to delete an export for a snapshot; the delete operation causes the export to be deleted from ViPR, however, the VNXe has not deleted it. Subsequent attempts to add an export will fail.

General notes on creating NFS exports are provided in [NFS export notes](#) on page 75 and the support for creating exports for each file storage system supported by ViPR Controller is described in [NFS export rules and permissions](#) on page 76.

Export CIFS snapshots

You can make a file system snapshot available as a CIFS share.

The following services support the creation and management of CIFS shares.

Service	Description
Create CIFS Share for File System	Enables you to select a previously created snapshot and make it available as a CIFS share. ViPR Controller does not support creating CIFS share for File System for Isilon storage systems.
Remove CIFS Share for Snapshot	Enables you to remove an CIFS share for a snapshot.

Notes on creating CIFS shares are provided in [CIFS share notes](#) on page 73.

Associate and dissociate a file policy with a file system

You can associate/dissociate a file snapshot/replication schedule policy with a file system.

The following services support the file policy operations.

Service	Description
Associate File Policy to File System	Associate a file snapshot/replication schedule policy to a file system.
Dissociate File Policy from File System	Remove a file policy from a file system.

Failover to and Failback from a file system

You can perform Disaster Recovery Failover and Failback operations using a file system.

The following services support the Failover and Failback operations.

Service	Description
Failover File System	Allows you to failover to the continuous copy (target) of the file system. You can also replicate the system configuration (the default), which includes CIFS shares, ACLs, NFS exports, and export rules. Once the issue found on the original source file system is resolved, you can use this service to failback to the original source file system.
Failback File System	After a failover file system operation, allows you to failback to the original source file system, once the issue found on the original source file system is resolved. You can also replicate the system configuration (the default), which includes CIFS shares, ACLs, NFS exports, and export rules.

CHAPTER 5

ViPR Controller Object Storage Services

This chapter includes the following topics:

- [Object storage services](#)..... 86
- [Creating a bucket](#).....86
- [Editing a bucket](#)..... 87
- [Deleting a bucket](#)..... 87

Object storage services

The Object Storage services enable you to create and manage EMC Elastic Cloud Storage buckets in ViPR Controller. Buckets are containers for object data. A bucket belongs to an ECS namespace and object users are also assigned to an ECS namespace. Each object user can create buckets only in the namespace to which they belong.

After ECS is under ViPR Controller management and its storage is provisioned using ViPR Controller, do not use the storage system element manager to manage the buckets created in ViPR Controller.

ViPR Controller does not support ingestion of buckets from ECS or object protocols, such as S3. You must create these buckets in ViPR Controller.

You can do the following using the object storage services:

- Create buckets and assign a valid ECS namespace object owner to each bucket.
- Modify the quota and retention period values of a bucket.
- Remove buckets.

For the ViPR Controller user roles required to perform these operations see [ViPR Controller user role requirements](#).

Creating a bucket

To create a bucket on a ViPR Controller virtual pool, set its quota and retention period attributes, and assign it an ECS object owner, use the **Object Storage Services > Create Bucket** service.

Before you begin

The ECS namespace that you specify on this page is mapped to a ViPR Controller tenant. When a bucket is created in ViPR Controller, it is added to the ViPR Controller project within a given tenant, and then added to the ECS namespace associated with that tenant.

Procedure

1. Select the ViPR Controller virtual array.
2. Select the ViPR Controller object virtual pool associated with the virtual array.
3. Select the project.
4. Type the name of the bucket.
5. Set the soft quota for the bucket.

Soft quotas log events to inform you when this quota limit is reached. This represents the Notifications Quota on the ECS Bucket Management page.

6. Set the hard quota for the bucket.

Hard quotas provide a hard limit on the amount of object storage to use for the bucket. This represents the Max Quota on the ECS Bucket Management page.

7. Set the retention period for the bucket.

This setting prevents data from being modified within the specified retention period. This value cannot exceed the retention value set on the bucket's object virtual pool.

8. Optionally, in **Owner**, type a valid ECS object owner.
9. Optionally, expand the **Access Control List** to add access control to the new bucket, and enter the following information

Option	Description
Type	The type of access control list: User, Group, or Custom Group.
Name	The domain name, which is configured on the ECS for the user name, group name, or custom group name. When adding access control to buckets, you cannot use spaces in the user, group, or custom group names.
Domain	Used when the user, group, or custom group has been configured as a domain on the ECS.
Permissions	The access permission that will be allowed for the user, group, or custom group assigned to this ACL: <ul style="list-style-type: none"> • Full Control — grants the users for this access control entry all permissions. • None — will give the users no access to the bucket. • Type the permission in the entry box to locate it in the list of possible permissions, or scroll the list of permissions, and check the ones to apply to the entry. Refer to ECS documentation for permission details.

10. Click **Order**.

A page appears showing the status of this order.

Editing a bucket

To edit the quota and retention period attributes of a bucket, use the **Object Storage Services > Edit Bucket** service.

Procedure

1. Select the project.
2. Select the bucket.
3. Change the quota and retention values as desired.
4. Click **Order**.

A page appears showing the status of this order.

Deleting a bucket

To remove a bucket from ViPR Controller and ECS, use the **Object Storage Services > Delete Bucket** service.

Procedure

1. Select the project.
2. Select the bucket.
3. Click **Order**.

A page appears showing the status of this order.

CHAPTER 6

ViPR Controller Vblock System Services

This chapter includes the following topics:

- [Vblock systems services](#)..... 90
- [Use the VCE Vblock System Services to create clusters on the Vblock compute system](#)..... 91

Vblock systems services

ViPR Controller services automate the following operations on Vblock systems that have been virtualized in ViPR Controller:

Table 24 ViPR Controller Service Catalog, VCE Vblock System Services

Service	Description
Provision Clusters	Provisions clusters of hosts on the compute system (Cisco Unified Computing System™ (UCS)), installs the operating system (ESX) on the hosts, and adds the cluster to a vCenter data center.
Add Host(s) to Cluster	Creates new hosts on the compute systems, installs the operating system on the hosts, adds the host to a cluster previously created in ViPR Controller, and adds the host to the vCenter datacenter if the vCenter was previously added to ViPR Controller.
Provision Bare Metal Cluster	Provisions clusters on the compute system, without installing an operating system.
Add Bare Metal Hosts to Cluster	Adds a new host to a cluster without installing an operating system on the host.
Update vCenter Cluster	Updates the vCenter with new cluster or host. If creating a new cluster, the cluster is added to the vCenter. If a new host was added to the cluster, and the cluster is already in the vCenter, only the new host is added to the cluster in vCenter.
Decommission Host(s) from Cluster	Decommissions one or more hosts from a cluster that was provisioned by ViPR Controller.
Decommission Cluster	Decommissions an entire cluster that was provisioned by ViPR Controller.

ViPR block and file storage services

Additionally, ViPR Controller Block and File Storage services can be used to manage Vblock storage systems.

ViPR Controller operations not supported for Vblock systems

ViPR Controller does not support ingestion of Vblock compute system blades that are being used outside of ViPR Controller management. The blades are discovered as unavailable to ViPR Controller, and will not be used by ViPR Controller for Vblock system provisioning or decommissioning services. However, you can add those hosts to the ViPR Controller physical assets, and UCS will discover the hosts from those compute systems through the operating system layer, and then ViPR Controller can export storage to those hosts.

ViPR Controller does not automate layer 2 network configuration. Layer 2 network configuration is managed by the UCS service profile templates assigned to the compute virtual pool.

Use the VCE Vblock System Services to create clusters on the Vblock compute system

Before you begin

Before you can successfully provision a cluster, ViPR Controller System Administrators, and Tenant Administrators must have configured the necessary projects, physical assets, and virtual assets.

The **Provision Cluster** service performs the following tasks in a single ViPR Controller operation:

- Creates a cluster in ViPR Controller, provisions hosts on the compute system (UCS).
- Creates boot volumes and zones them to their respective hosts.
- Installs the operating system (ESX) on the hosts.

Note

ESXi 5.1 and later requires a persistent `/scratch` partition. The ESXi installer creates a datastore named `datastore1` for this purpose using the remaining volume storage on the boot LUN allocation.

- Adds the cluster to a vCenter data center.

Procedure

1. Go to the **Service Catalog > VCE Vblock System Services > Provision Cluster** service.
2. Complete the order form as follows:

Option	Description
Project	The project in which the cluster will be added after it is created.
Name	The name of the cluster that will be created on the Vblock system and vCenter.
Virtual Array	The virtual array in which the Vblock networks, and boot LUN storage are included.
Compute Virtual Pool	The compute pool containing available compute elements to use in provisioning.
VCE Vblock Compute Image (OS)	The operating system that will be installed on the cluster hosts. A ViPR Controller System Administrator must add the compute image to the ViPR Controller Physical Assets, before the service can be ordered.
Hosts	The host information for each host created in the cluster. Click Add for each additional host to add to the cluster.

Option	Description
Host Name (FQDN)	The fully qualified domain name that will be assigned to each host created in the cluster.
IP Address	The IP Address that will be assigned to each host created in the cluster.
Netmask	Netmask that will be used by all of the hosts in the cluster.
Gateway	Gateway that will be used by the hosts.
Management Network	The Management VLAN over which ViPR Controller will use to communicate with the hosts for management purposes. The management network is required to add the ESX hosts to a vCenter and for any communication between ViPR Controller and the ESX hosts.
NTP Server	NTP Server that will be used by all of the hosts in the cluster.
DNS Servers	Comma separated list of DNS Servers that will be used by the hosts.
Host Root Password	The password that will be assigned to the root user for all the hosts in the cluster.
Block Virtual Pool	The block virtual pool from which the boot LUN will be created.
Size of Boot Volume	The size of the volumes that will be used to boot the hosts.
vCenter	The vCenter in which the datacenter, to which the cluster will be added, resides. A ViPR Controller Tenant Administrator must add the vCenter to ViPR Controller before the service can be ordered.
Datacenter	The vCenter Datacenter to which the cluster will be added. A ViPR Controller Tenant Administrator must add the vCenter to ViPR Controller before the service can be ordered.

3. Click **Order**.

The Orders page is displayed with the progress of the order.

CHAPTER 7

Migration Services

The Service Catalog provides the following areas for migration services.

- [Migration services](#).....94

Migration services

ViPR Controller provides services to migrate data.

The migration services include:

- [VPLEX Data Migration](#)
- [Mobility group migration](#) on page 96
- [Adding resources to a mobility group](#) on page 97
- [Removing resources from a mobility group](#) on page 97

VPLEX data migration

Move a volume from one virtual pool to another to perform a VPLEX data migration. Use this service to change volumes from VPLEX local to VPLEX distributed. You may also use this service to change the virtual pool for VPLEX source, target, or journal volumes that are protected by RecoverPoint.

Before you begin

If a VPLEX Migration order has both RecoverPoint Source and Targets and there is an error which rolls back the RecoverPoint Source migrations but the RecoverPoint Target migrations succeed, then order can be placed again without any consequences, as the second order will only migrate the previously failed RecoverPoint Source volumes.

Procedure

1. Navigate to **Service Catalog > Migration Services > VPLEX Data Migration**
2. Select the **Project** in which the volume is located.
3. Select the **Virtual Pool** containing the volume being moved.
4. Select the **Operation** to perform by moving the volume to another virtual pool. If you wish to move multiple volumes, use [Mobility group migration](#) on page 96.

Option	Description
VPLEX Data Migration	<p>By changing the volume on the backend storage system, you can perform functions such as: change the backend volume from thin to thick, change the storage system. You can configure the speed of the data migration using Physical Assets > Controller Config > VPLEX and then adding a new configuration for <code>Data Migration Speed</code>. See Setting the type of transfer speed for VPLEX data migration.</p> <p>You can also use the REST API and the CLI to retrieve a list of data migrations, show the details of a data migration, pause, resume, and cancel a migration. See the <i>ViPR Controller REST API Reference</i>, available as a zip file from the ViPR Controller Product Documentation Index and the <i>ViPR Controller CLI Reference Guide</i> which can be found on the ViPR Controller Product Documentation Index.</p>

Option	Description
	<p>RecoverPoint protected VPLEX volumes or MetroPoint (VPLEX Metro only) volumes are eligible for VPLEX Data Migration too. For these volumes, the original virtual pool is compared to the target virtual pool and migrations are based on changes in</p> <ul style="list-style-type: none"> • Source virtual pool • Source journal virtual pool • Target virtual pools • Target journal virtual pools <p>Targets and Journals can be implicitly migrated if there are changes in the new virtual pool when compared to other virtual pools. (The other virtual pools must be eligible for migration.)</p> <p>The same rules apply to all virtual pools when determining whether or not a migration will be triggered.</p> <p>RecoverPoint protected VPLEX volumes or MetroPoint (VPLEX Metro only) volumes that are in consistency groups with array consistency enabled OR are in Applications will be grouped together for migration.</p> <p>RecoverPoint or MetroPoint (VPLEX Metro only) Target volumes that are in Applications will be grouped together for migration.</p>
Change from VPLEX Local to VPLEX Distributed	<p>The volume is moved from a local to distributed VPLEX .</p>

5. Select the **Target Virtual Pool**.
6. Select the **Volume**.
7. Check the **Suspend** box if you want to suspend migration before committing and deleting the original source, target, or journal volumes. This allows you to manually verify data integrity.
8. Set the **Display Journal** field to **yes** if you want to migrate only RecoverPoint - protected journal volumes.
9. You can schedule the time for performing the migration by selecting **Enable Schedule** and entering the start date and time.
10. Select **Order** to run the service.

Setting the type of transfer speed for VPLEX data migration

You can set the type of transfer speed to be used when executing the VPLEX Data Migration operation in the VPLEX Data Migration, Change Volume Virtual Pool, and Change Virtual Pool services. This transfer speed also applies to the Change Virtual Array service.

Procedure

1. Navigate to **Physical > Controller Config**.

2. Select the **VPLEX** tab.
3. Select **Data Migration Speed** in the drop-down list.
4. Click **Add**
5. Select the **Scope Type**.
6. Select the **Scope Value**.
7. Type the **Value** of your required data migration speed.

Option	Description
Lowest	Mapping transfer size is 128 KB. Note Less impact on the host I/O but data migration takes longer to complete.
Low	Mapping transfer size is 2 MB.
Medium	Mapping transfer size is 8 MB.
High	Mapping transfer size is 16 MB.
Highest	Mapping transfer size is 32 MB. Note Data migration completes more quickly, but there is a greater impact on host I/O.

Mobility group migration

Enables migration of multiple VPLEX volumes with one order. You can group volumes by host, cluster, or by an explicit list of volumes.

Procedure

1. Select **Catalog > View Catalog > Migration Services > Mobility Group Migration**.
2. Select a mobility group name.
This is the name created with the **Virtual > Mobility Groups** page.
3. Select the migration method.
 - Select **Migrate Only** when you want to migrate the mobility group volumes into the target virtual pool.
 - Select **Ingest and Migrate** when you want to ingest first and then migrate the volumes. Select **Ingest and Migrate** only when the mobility group is grouped by host or cluster.
4. Select the project that will hold the volumes when using the **Ingest and Migrate** option.
5. Select the virtual array that will hold the volumes when using the **Ingest and Migrate** option.
6. Select the virtual pool that will hold the volumes when using the **Ingest and Migrate** option.

7. Select the target virtual pool that will hold the volumes that are being migrated.
8. Click **Order**.

The Orders page appears showing the progress of the order.

Results

You can monitor the status of the migration from the **Resources > Volumes** page. Depending on the state of the migration, there will be buttons to pause, resume, or cancel the migration.

Adding resources to a mobility group

Add volumes, hosts, or clusters to a mobility group.

Procedure

1. Enter the mobility group name.
2. Select the volumes, hosts, or clusters that will be migrated.
Adjust the type of group to migrate (volumes, hosts, or clusters) by editing the mobility group settings in **Virtual > Mobility Groups** .
3. Click **Order**.

The Orders page appears showing the progress of the order.

Removing resources from a mobility group

Remove volumes, hosts, or clusters from a mobility group.

Procedure

1. Enter the mobility group name.
2. Select the volumes, hosts, or clusters to be removed.
3. Click **Order**.

The Orders page appears showing the progress of the order.

Results

View the updated mobility group resources from the **Virtual > Edit Mobility Groups** page.

CHAPTER 8

ViPR Controller Service Orders and Resources

This chapter includes the following topics:

- [Working with service orders](#)..... 100
- [Monitoring and troubleshooting a service order](#)..... 107
- [Working with service resources](#)..... 108

Working with service orders

An order is placed anytime a service is run from ViPR Controller. When placing a service order it is helpful to be aware of the following.

Checking the status of an order

Once you place a service order the progress of the order is automatically presented in the UI. You can also view the status of an order by clicking on a placed order from anyone of the following UI pages:

- **Catalog > My Orders**
- **Catalog > All Orders**
- **Catalog > Scheduled Orders**

Wait for orders to complete

After using a service to create a volume, file system, or replication of the volume or file system, wait for the entire order to complete, before using the newly created device. In some situations, not waiting for a ViPR Controller order to complete, before performing operations on the devices, could create a data unavailability situation.

For example, if a create order was run to add a volume to an existing SRDF consistency group, in which SRDF volumes, and R2 side replicas such as full copies, snaps, snap sessions, or continuous copies already existed, a new R1 device, and its R2 device is created, and its corresponding replica is also created. If a failure occurs while the replica is being created, ViPR Controller automatically executes a rollback on the entire order. If you do not wait for the order to complete, and you are unaware of the unexpected failure while creating the replica, you could begin to export the R2 device to a host. As part of the ViPR Controller rollback, the host having the exported R2 device will encounter a data unavailability situation as it has started using the R2 device, which was removed as part of the rollback from the failure.

My Orders

The **Catalog > My Orders** page lists orders submitted by your user account.

The page shows the order number, which is a unique identifier assigned to the order to reference it as required, the order status, a summary of the order request, the date and time the order was submitted, and the date and time the order is scheduled to execute, if applicable.

Table 25 Order status icons

Icon	Meaning
✓	Order was processed successfully.
🕒	Order is scheduled. Open the order to see the execution windows in which it will be executed.
✗	Order failed.
👍	Order is waiting for approval.

You can use the following features to research orders.

Table 26 Order actions

Action	Description
To search for a specific order.	Type an order number in the Search text box.
To limit the list of orders to specific timeframes.	<p>Choose a value from the Showing orders from the last drop-down list:</p> <ul style="list-style-type: none"> • Day • Week • 2 Weeks • 30 Days • 90 Days • Date Range—Specify the range by clicking the calendar icons that appear and selecting both the start and end dates. <hr/> <p>Note</p> <p>You must select both dates. Even though the end date calendar is pre-populated with the current date, you must click that date to confirm the selection.</p> <hr/>
To show more details about an order.	Click the order row. The Order Details page for the order appears.

Order details

The order details shows the order request information, details of the selected order, and a summary of the execution steps.

Table 27 Order details areas

Order area	Contains
Summary	<p>The top part of the order displays the order identity and status on the left-hand side, and displays the parameters passed to the operation on the right-hand side. You will see all of the parameters passed to the service, even if they were hidden on the service form.</p> <p>If the order is scheduled to run in an execution window, the status displays <code>Order Scheduled</code>. If the order requires approval, the status displays <code>Pending Approval</code>.</p>
Affected Resources	Shows the details of the resource that was created as a result of the operation or in which the operation was run.
Order Details	Displays detailed information about the steps taken to process the order.

The following table lists the affected resource types and the information provided with each type.

Table 28 Affected resources

Affected resource	Description
Block Storage	
Block Export	If the operation exports a block volume to a host, this entry displays the details of the host.
Volume	Shows details of the block volume that has been created. If the operation fails and rollback requires the volume to be deleted, displays: "volume has been deleted."
File Storage	
File System	Shows details of the file system created. If the operation fails and rollback requires the file system to be deleted, displays: "file system has been deleted."
CIFS Share	Shows details of the CIFS share. Indicates the mount point that can be used when mapping the drive on a Windows host.
NFS Export	Shows details of the NFS export. Indicates the mount point that can be used when mounting the file system on a Linux host.

The following table describes each of the order detail types.

Table 29 Order details

Category	Content
Logs	Shows the messages written to the log file during execution of the order.
Precheck	Shows the pre-checks carried out before executing the order. The pre-checks ensure that the host to which storage will be attached is available and that there are the required paths between the host HBA and the selected array.
Execution Steps	Lists the execution steps that were carried out in executing the order.
Rollback	If an order fails, the steps carried out to rollback any command executed before failure are listed.

Approvals

Tenant approvers can access the **Approvals** page from the **Catalog > Approvals** menu. Only Tenant Approvers can see the Approvals menu.

Approval table

The **Approvals** page shows all approvals that have been submitted, the user who submitted the order, the status of the order's approval request, the time it was submitted and, if it has been approved, who approved it.

Approving or rejecting orders

Note

Tenant Administrators can configure the approvals feature to send a notification email when an order is waiting for approval.

1. From the **Catalog > Approvals** page, click the service name in the table to open the approval panel.
2. Optionally, enter a reason for approving or rejecting the order.
3. Click **Approve** or **Reject**.

Note

If you created the order, the **Approve** button is not available. It is a security issue to permit the same user to create an order and then approve it. You can only reject your own orders.

All Orders

An order is a record of a request to run a service. Administrators use the **Catalog > All Orders** page to view orders created by users within the tenant.

The **All Orders** page differs from the **My Orders** page in the User view in the following ways:

- It shows the user who submitted each order
- It includes the Download and Delete buttons which help administrators manage orders in the database

The order records the details of the request: which service was requested, what parameters were specified in the service request, who requested it, whether the order is scheduled, and the outcome of the order submission.

The summary tells you the service that was executed and the date and time at which it was executed. The time is displayed in local time using the time zone set in your browser.

In addition, the page shows you the status of the order and displays an order ID that provides a unique identity for the order that provisioning users can use when requesting information from you (Tenant Administrator) when there is a problem with the execution of an order.

The following table describes the order status icons.

Table 30 Order status icons

Icon	Meaning
✓	Order was processed successfully.
🕒	Order is scheduled. Open the order to see the execution windows in which it will be executed.
✘	Order failed.
👍	Order is waiting for approval.

You can use the following features to research orders.

Table 31 Order actions

Action	Description
To search for a specific order.	Type an order number in the Search text box.

Table 31 Order actions (continued)

Action	Description
To limit the list of orders to specific timeframes.	<p>Choose a value from the Showing orders from the last drop-down list:</p> <ul style="list-style-type: none"> • Day • Week • 2 Weeks • 30 Days • 90 Days • Date Range—Specify the range by clicking the calendar icons that appear and selecting both the start and end dates. <hr/> <p>Note</p> <p>You must select both dates. Even though the end date calendar is pre-populated with the current date, you must click that date to confirm the selection.</p>
To show more details about an order.	Click the order row. The Order Details page for the order appears.
To create an offline copy of orders	Select orders to copy using the checkboxes and then click Download .
To delete orders from the database	Select orders to delete using the checkboxes and then click Delete .

Download orders

On the **Catalog > All Orders** page, the **Download** button lets you save an offline copy of the information in selected orders. This action requires the Tenant Administrator role.

The **Download** button triggers a job that creates a zip file named `orders_n-DDMMYY-HHMM`. The zip file contains one text file named `orders`. The `orders` file contains all available information about all of the orders that you selected for the current job.

The downloaded information is for reference only. It cannot be uploaded back into the ViPR Controller database.

Procedure

1. Select **Catalog > All Orders**.
2. Define a subset of orders using the options in the **Show all orders from last** field.

For example, select **30 Days**, or select **Date Range** and specify both the start and end dates.

Note

The end date calendar pre-populates the current date; however, you must click the pre-populated value to confirm the selection.

3. Select the orders that you want to download using the following actions:

Option	Description
Select all orders from the most recent query	<p>Click the check box in the header.</p> <hr/> <p>Note</p> <p>This action selects all orders that satisfy the most recent query. For example, if a query returns 9,000 orders, even though only 6,000 orders can display on the UI, the select-all checkbox applies to all 9,000 results for the purposes of using the Download and Delete buttons.</p> <hr/>
Select individual orders	Click the check box in each row.
Deselect all checked items	Click the check box in the header.

4. Click **Download**.
5. Follow your browser prompts to save the file.

The default location of the zip file is your local `Downloads` folder.

Results

An informational message about the job appears at the top of the page. The informational message includes the job id and describes how many orders are included in the job, how long the job has been running, and the expected remaining time until completion. Only one download or delete job can run at a time. Both the **Download** and **Delete** buttons are disabled for the duration of the job run.

Delete orders

On the **Catalog > All Orders**, the **Delete** button lets you remove selected orders from the database. This action requires the Tenant Administrator role. The expected workflow is to archive orders using the **Download** button before deleting them.

The **Delete** button triggers a job that processes each selected order, determining if the order is a candidate for deletion. To be deleted, an order must meet the following criteria:

- The order status must be one of the following: SUCCESS, PARTIAL_SUCCESS, ERROR, CANCELLED, REJECTED. Orders associated with scheduled events have a status of SCHEDULED and therefore can not be deleted.
- The order must be older than 30 days.

When a selected order meets those criteria, the job deletes the order and all objects and logs associated with it from the database.

The **Delete** button is an effective tool for managing database space.

Procedure

1. Select **Catalog > All Orders**.
2. Define a subset of orders using the **Date Range** option in the **Show all orders from last** drop-down list.

Note

Remember that orders must be at least 30 days old to be deleted.

3. Select the orders that you want to delete using the following actions:

Option	Description
Select all orders from the most recent query	<p>Click the check box in the header.</p> <hr/> <p>Note</p> <p>This action selects all orders that satisfy the most recent query. For example, if a query returns 9,000 orders, even though only 6,000 orders can display on the UI, the select-all checkbox applies to all 9,000 results for the purposes of using the Download and Delete buttons.</p>
Select individual orders	Click the check box in each row.
Deselect all checked items	Click the check box in the header.

4. Click **Delete**.

Results

An informational message about the job appears at the top of the page. The informational message includes the job id and describes how many orders are included in the job, how long the job has been running, and the expected remaining time until completion. Only one download or delete job can run at a time. Both the **Download** and **Delete** buttons are disabled for the duration of the job run.

No more than 20,000 order deletions can be processed at a time. If a delete job contains more than that, the first 20,000 orders are deleted initially, and the job continues 5 days later to delete the next 20,000 orders. This processing continues until all of the requested orders are deleted. If a job contains 50,000 orders, the job takes more than 10 days to complete. The job continuation occurs automatically without any user interaction. The 5-day delay is an internally set database configuration parameter.

Scheduled orders

Use the **Catalog > Scheduled Orders > Scheduled Orders** page to review orders that are scheduled to run in an Execution Window, or to view orders scheduled by users through the protection services.

Additionally, Tenant Administrators can edit or cancel a scheduled order, edit a recurring schedule, or edit the schedule expiration policy from this page.

Table 32 Scheduled orders options and attributes

Column name	Description
Selection column	Tenant Administrators can select one or more rows and click Cancel to cancel the order, or click Edit Schedule to edit

Table 32 Scheduled orders options and attributes (continued)

Column name	Description
	a scheduled order, edit a recurring schedule, or edit the schedule expiration policy.
Number	The unique number to identify the order.
Status	The status of the operation.
Service	The name of the ordered service.
Summary	Service description.
Submitted by	User that placed the order.
Date Submitted	The date the order was placed.
Scheduled Time	Displays the time the order is scheduled to run.
Execution Window	The date, time, and how often the order is scheduled to run. The execution window is configured in the service by the Tenant Administrator.

Monitoring and troubleshooting a service order

Once a service is ordered, you can watch the progress of the service, and troubleshoot issues with the service from the ViPR Controller user interface, **Orders** page, or the **My Orders** page.

After ordering a service, the order page opens displaying the progress of the order.

Procedure

- Keep the order page open to continue to watch the progress of the service operation.
If you have left the orders page, go to the **Service Catalog > My Orders** page, and you will see the order in the list.
- Click the order to view the details of the order.
- Expand the **Logs** section to view the logs entered for the service.
Red text indicates that an error occurred while the service was running, and provides details of the error.
- Expand **Precheck Steps** to view the steps ViPR performs before executing the service.
The order will not proceed if any of the precheck steps fail.
- Expand **Execution Steps** to view the steps ViPR performs to complete the service order.
Red text indicates an error occurred during the service execution and the point in the execution steps where an error occurred.
- Expand **Tasks** to view the details of each specific task run to complete the order. If an order failed, the Tasks will show the specific task where the order to fail.

Working with service resources

Once a service operation has completed successfully, the resources are put into a project that was assigned in the service order. You can view and delete these resources using the **Resources** pages.

Before you begin

Provisioning users can only select the projects to which they are assigned and can only view the resources in those projects.

Tenant administrators can see all projects and project resources.

Procedure

1. Open **Resources** and select the type of resource to view.
2. Select the project in which the resource belongs.

A list of the resources of the selected type appear in the table. For example, if you selected **File System**, a list of file systems that were provisioned in the selected project appears in the **File System** table.

If you selected **Volumes**, a list of block volumes that were provisioned in the selected project appears in the **Volumes** table.

3. To delete a resource, check the box in the row of the resource and then click **Delete**.
4. Click anywhere in a row to see more details about a resource.

Applications

Use the **Resources > Applications** page to view and manage volumes in a selected application.

View the volumes, full copies, snap sessions, and snapshots associated with an application.

Note

To manage the volumes, full copies, or snapshots associated with application sub groups, use the **Catalog > View catalog > Application Services** pages.

Table 33 Areas of the Applications page

Area	Description
Application Name	Lists the name of the application.
Description	Displays the description added when you first created the application.

Table 34 Application details

Area	Description
Volumes	Provides summary information about the volume, including its name, size, virtual array, virtual pool, and application sub group.

Table 34 Application details (continued)

Area	Description
	When you select a volume, you can view details about the Volume resources on page 109.
Full Copies	Provides the creation time of the application subgroup full copy. When you select the full copy instance, you can view details for volumes in the application sub group such as size, virtual array, and virtual pool.
Snap Sessions	When you select the snap session, you get a summary of that snap session where you can drill down to get more details.
Snapshots	Lists the snapshot sets created for the application sub groups.

Volume resources

Use the **Resources > Volumes** page to view and manage volumes in a selected tenant and project. You can view the access state and exports of a volume as well as delete a volume from the ViPR Controller inventory and its backend storage system using this page.

The **Volumes** page contains these areas.

Table 35 Areas of the Volumes page

Area	Description
Summary	You can filter volumes by project or by application. This page provides summary information about the volume, including its WWN, size, associated ViPR Controller virtual array and virtual pool, consistency group, storage system name, device label, compression ratio for VMAX3 arrays, native ID, access state, volume ID, creation date, and any tags. In Actions , you can delete the volume from the ViPR Controller database (Inventory Only) or from both the ViPR Controller database and its backend storage system (Full). A Full delete removes the volume and all objects referencing the volume, such as exports and snapshots, from the ViPR Controller database and its backend storage system. An Inventory Only delete removes the volume and all objects referencing the volume from the ViPR Controller database.
Exports	Lists the exports created for the volume.
Snapshots	Lists the snapshots created on the volume.
Full Copies	Lists the full copies created on the volume.
Continuous Copies	Lists the continuous copies created on the volume.
Tasks	Lists the tasks performed on the volume.

Snap Sessions

Use the **Resources > Snap Sessions** page to view the details of snapshot sessions in a selected tenant and project, or click the snapshot session name to view the details of a snapshot session.

Snap Sessions page

The **Snap Session** page lists the snapshot sessions, and the following snapshot session attributes.

Table 36 Areas of the Snap Sessions page

Column	Description
Selection column	Use the selection column to select the snap sessions to delete, and then click Delete to delete the snapshot session for the source volume. After which, all the targets linked to the snapshot session must be unlinked before the snapshot session can be deleted.
Name	The name of the snapshot session created in ViPR Controller. Click the name to see the details of a specific snap session.
Volume	The volume or consistency group for which the snap was created.
Date created	The date the snap session was created in ViPR Controller.

Snapshot Session details page

The **Snapshot Session** details page provides the following information and options.

Table 37 Snapshot Session detail page information and options

Area	Description
Actions	Delete Snap Session — delete the snapshot session for the source volume. After which, all the targets linked to the snapshot session must be unlinked before the snapshot session can be deleted.
Snapshot Targets	List the snapshot targets, and allows you to perform the following actions: <ul style="list-style-type: none"> Click the name of the snapshot target to open the details page for that target device. Click Relink to relink the snapshot session target with the source devices. Click Unlink and with Delete to unlink the snapshot session, and then delete the target. Click Unlink and without Delete to unlink the target from the snapshot session while continuing to use ViPR Controller to manage the target as a volume without a snapshot session.

Block snapshot resources

Use the **Resources > Block Snapshots** page to view and delete snapshots created on volumes of a selected tenant and project.

Export group resources

Use the **Resources > Export Groups** page to view and delete export groups of a selected cluster, tenant, or project.

Export groups are ViPR Controller constructs that are created and updated to reflect export mask-like constructs on the storage systems. ViPR Controller creates or updates export groups as a result of other operations to reflect volume access to hosts and clusters.

Also see [Actionable events](#) on page 121 for information about actions related to export groups that affect clusters, hosts, initiators, and volumes. These actions must be approved or disapproved by an administrator with correct permissions.

Note

While deleting an export group (if the last volume is requested to be removed), exporting the volume to the same resource is not supported.

File system resources

Use the **Resources > File Systems** page to view the list of file systems and file system attributes, to delete a file system from the ViPR Controller inventory and from its backend storage system, and to access the resource page for a specific file system.

File System > [file system name] page

When you click on a file system from the **File System** page, you gain access to the file system specific page, which contains these areas.

Table 38 Areas of the File System page for the selected file system

Area	Description
Mount Path	The file system mount path.
Size	The size of the file system, whether or not it was provisioned, and if provisioned, how it was provisioned (Thin, Thick).
Protocols	The protocols associated with this file system.
Virtual Array	The virtual array from which the file system was created.
Virtual Pool	The virtual pool from which the file system was created.
Actions	<p>Provides the action to delete a file system. Use:</p> <ul style="list-style-type: none"> • Full to delete the file system from both the ViPR Controller database and its backend storage system. A Full delete removes the file system from ViPR Controller database and its backend storage system. • Inventory Only to remove the file system and all objects referencing the file system from the ViPR Controller database. <p>If you use the delete action from this page, you must remove all the references to that file system such as, file export rules, or ACEs, before deleting the file system. However, you do not have to remove the references if you</p>

Table 38 Areas of the File System page for the selected file system (continued)

Area	Description
	remove the file system using the Remove File System service, which is available from the Service Catalog. The Remove File System will provide the same functionality of a full delete, but will also automatically delete the file system references.
Click More Details to see:	
Storage System	The physical storage system on which the file system was created.
Storage Pool	The physical storage pool from which the file system was created.
Storage Port	The storage port used by the file system.
Advisory Limit (%)	Once this limit is reached, ViPR Controller displays an alert on the file system, Resources page, and a notification is sent to the configured users from the storage system.
Soft Limit (%)	ViPR Controller issues a, Capacity Exceeded , warning when this limit is reached. Writes to the file system are allowed to continue until the grace period reached.
Grace period (Days)	The number of days after the soft limit is reached before writes to the file system will be stopped. Once the grace period is exceeded writes to the file system are stopped.
Native ID	The mount path id used by the native storage system to identify the mount path.
ID	The file system id used by ViPR Controller to identify the file system in its database.
Created	Date the file system was created.
Tags	List of tags, if applicable, assigned to the file system.
Export Rules	Lists the NFS exports created for the file system and the security associated with each rule. You can use the displayed mount point to mount the NFS export on the host on which it was exported. Provides operations that enable you to: <ul style="list-style-type: none"> • Modify an export rule by adding additional hosts (endpoints), changing the security type (including selecting multiple security types for Isilon), and changing the permissions associated with the host. • Add an export rule. • Delete an export rule.
Shares	Lists the CIFS shares associated with a file system. A file system can be shared with a number of names. The Shares area provides options to perform the following tasks:

Table 38 Areas of the File System page for the selected file system (continued)

Area	Description
	<ul style="list-style-type: none"> • Delete — to delete the CIFS share • Share — to add a CIFS share at the subdirectory level. The subdirectory must already exist on the file system. • Access Control — access control can be added, modified, or deleted from the share for Isilon, and NetApp 7-mode, and NetApp Cluster-mode storage systems. Two types of users, or groups can be used for ACLs <ul style="list-style-type: none"> ▪ Domain users or groups — must be registered on the domain controller and the storage system's data mover has to be added to the domain. ▪ Local users or groups — can be a local user or group, or part of the Authentication Provider (For example, AD, LDAP), which is configured on the storage system.
Snapshots	<p>Lists the snapshots created on the file system. Snapshots area only lists the file snapshots that were created on demand. Expand the Snapshot Policies area, and click Snapshot List, to see the list of scheduled snapshots created for the file system.</p>
Quota Directory	<p>You can also delete and modify this directory. The modify action allows you to change the size of the quota directory and set its security style.</p> <ul style="list-style-type: none"> • Lists any quota directories that were created on the file system. • Lists the Smart Quota's set for the quota directory if applicable. • Allows you to perform the following operations on the quota directories for the file system: <ul style="list-style-type: none"> ▪ Modify the size of the quota directory, and set its security style. ▪ Set the smart quota limits on the quota directory. Smart quota limits include: <ul style="list-style-type: none"> – Advisory Limit (%) — ViPR Controller provides an informational alert once this limit is reached on the quota directory. – Soft Limit (%) — ViPR Controller issues a, Capacity Exceeded, warning when this limit is reached. Writes to the quota directory will be allowed to continue until the grace period reached. – Grace Period (Days) — The number of days after the soft limit is reached before writes to

Table 38 Areas of the File System page for the selected file system (continued)

Area	Description
	<p>the quota directory will be stopped. Once the grace period is exceeded writes to the quota directory are stopped. ViPR Controller issues a critical alert when the Grace Period is exceeded.</p>
NFS Access Control	<p>For Isilon storage systems configured with NFSv4 protocol you can expand to:</p> <ul style="list-style-type: none"> • View the list of Mount Paths on which NFS Access Controls can be configured. • Click Add Access Controls to open the Add Access Control Entry page. • Click Access Control to navigate to the Access Control List page. While in the ACL page: <ul style="list-style-type: none"> ▪ Select the check box next to an ACE, and click Delete to delete the ACE. ▪ Click on an ACE to edit the type of permission set (Allow, Deny), and which permissions were set on the file system or sub directory. <p>For steps to add a new ACE see: Add access control to NFSv4 file systems.</p>
File Mirrors	<p>Lists the continuous copy targets created for the file system and allows you to perform the following actions on the mirrors:</p> <ul style="list-style-type: none"> • Refresh — to manually refresh the mirror status. • Stop — to manually stop synchronization between the file system and the continuous copy of the file system. • Pause — to temporarily pause synchronization between the file system and the continuous copy of the file system. • Resume — to resume synchronization of the file system and the continuous copy of the file system after synchronization has been paused.
File System Protection Policies	<p>Lists the protection policies that have been added to the file system, and, for protection policies applied at the file system level, allows you to assign or unassign them to/from a target virtual array.</p> <p>If the protection policy is applied at the file system level, click Assign Policy to assign it to a target virtual array, or Unassign Policy to unassign it from a target virtual array.</p>
Tasks	<p>Lists the tasks performed on the file system.</p>

Add access controls to NFSv4 file systems

Access controls can be added to file systems, or a file system's subdirectory on Isilon storage systems configured with NFSv4 protocol.

Before you begin

- The ViPR Controller user must be able to access the mount path on which the access control will be added.
- The user or group can be a domain or local user that has been configured on the Isilon storage system.
- If multiple Access Control Entries (ACEs) are entered on a file system for the same user or groups, the lowest assigned permission will be the one used.
- If ACEs with the same user or group are assigned to a file system, and the file system subdirectory, the lowest assigned permission will be the one used.
- Sub directories are not created by ViPR Controller, you must know the name of a sub directory that already exists within the file system on the Isilon storage system and enter it exactly as it appears on the Isilon storage system.

Procedure

1. Go to **Resources > File Systems** page and select the file system on which you will be adding the ACE.
2. Expand **NFS Access Controls**, and click **Add Access Controls**.
3. Enter the following information

Option	Description
Sub Directory	Optionally, enter the sub directory on which the permissions will be set.
Type	Select User , if it will be an individual user, or Group if it will be a group of users.
Name	Enter the user or group name exactly as it has been configured on the Isilon storage system.
Domain	If the user or group has been configured as a domain user or group on the Isilon storage system, enter the domain here.
Permission Type	Select Allow , if you are granting the permissions to the user, or Deny if you are not giving the user the permissions.
Permissions	Select the permissions to allow or deny: Read, Write, Execute, or Full Control

4. Optionally, click **Add** to add another ACE to the access control list, and enter the ACE information.

If you entered a sub directory in step 3, the ACE will be applied to the same sub directory entered at the top of the page.

5. Click **Save**.

You are returned to the details page of the file system

6. Click **Access Control** to view the list of ACEs you just added to the file system.

File snapshot resources

Use the **Resources > File System Snapshots** page to view and manage the snapshots created on file systems for a selected tenant and project.

Viewing snapshots created on demand vs. scheduled

The File Snapshots page only lists the file snapshots which were created on demand. The File Snapshot page does not present snapshots which were scheduled to be created. To see the list of scheduled snapshots created for a specific file system, you must:

1. Go to the **Resources > File systems** page.
2. Select the file system.
3. Expand the Snapshot Policies area.
4. Click Snapshot List.

File Snapshots page

The **File Snapshots** page contains these areas.

Table 39 Areas of the File Snapshots page

Area	Description
Summary	Provides summary information for the file system: Its size, supported protocols (NFS and/or CIFS), mount path and the ViPR Controller virtual array and virtual pool that it belongs to.
Details	Provides additional details about the snapshot.
Export Rules	<p>Lists the NFS exports that were created for the file snapshot and the security associated with each rule. The mount point displayed can be used to mount the NFS export on a host to which it has been exported.</p> <p>Provides operations that enable:</p> <ul style="list-style-type: none"> • Modify the rule by adding or removing allowed hosts and changing the permissions associated with a host. • Delete a rule. • Add an Export Rule. <hr/> <p>Note</p> <p>Adding an export rule is not supported for Isilon storage systems</p> <hr/>
Shares	<p>Lists the CIFS shares associated with a snapshot. The Shares area provides options to perform the following tasks:</p> <ul style="list-style-type: none"> • Delete — to delete the CIFS share snapshot • Add Share — to add a CIFS share to the snapshot at the subdirectory level. The subdirectory must already exist on the file system.

Table 39 Areas of the File Snapshots page (continued)

Area	Description
	<ul style="list-style-type: none"> • Access Control — access control can be added, modified, or deleted from the snapshot for Isilon, NetApp 7-mode, and NetApp Cluster-mode storage systems. Two types of users, or groups can be used for ACLs <ul style="list-style-type: none"> ▪ Domain users or groups — must be registered on the domain controller and the storage system's data mover has to be added to the domain. ▪ Local users or groups — can be a local user or group, or part of the Authentication Provider (For example, AD, LDAP), which is configured on the storage system.
Tasks	Lists the tasks associated with the snapshot that were performed. The task details can be displayed, which in turn, can be used to link back to the original order.

Deleting a file snapshot

You can delete a file snapshot from the **Resources > File Snapshots** page if you have removed all of its associated exports and shares. However, even if you have removed all of the Export Rules, you must also ensure that any NFS exports have been removed using the **Service Catalog > View Catalog > File Protection Services > Remove NFS Export for Snapshot** service.

vNAS server resources

Use the **Resources > vNAS Servers** page to view the details of vNAS servers in a selected tenant and project, or click the **vNAS** button next to any VNX for File or Isilon storage system on the **Storage Systems** page to access this information. You can also view the performance metrics, such as used storage capacity and average percentage busy, for workloads on the **vNAS Servers** page.

The **vNAS Servers** page contains these areas.

Table 40 Areas of the vNAS Servers page

Area	Description
Name	The name of the vNAS server that was discovered.
Registered	Indicates whether the vNAS server and its attributes were successfully discovered and registered in ViPR Controller.
Protocols	The protocol being used by the vNAS server during the provisioning operation.
Parent NAS Server	The physical NAS server on which the vNAS server was created.
Domain	The name of the domain to which the vNAS server belongs.
State	The state of the vNAS server, which can be mounted, loaded, and unknown.

Bucket resources

Use the **Resources > Buckets** page perform the following operations on Elastic Cloud Storage (ECS) buckets in a selected tenant and project:

- View list of buckets, and bucket attributes.
- View the details of an individual bucket.
- Delete the bucket from both ViPR Controller, and the ECS system, or just from the ViPR Controller inventory.
- Add, or edit access control to a bucket.

View the list of buckets and bucket attributes

The **Buckets** presents the list of buckets created in ViPR Controller, and the following attributes for each bucket.

Table 41 Buckets attributes

Area	Description
Bucket Name	The name of the bucket, which consists of the namespace, project, and bucket, respectively. For example: provider_education_bucket6
Hard Quota	Hard quota assigned to the bucket. Hard quotas provide a hard limit on the amount of object storage to use for the bucket. This represents the Max Quota in the ECS Bucket Management page.
Soft Quota	Soft quota assigned to the bucket. Soft quotas log events to inform you when this quota limit is reached. This represents the Notifications Quota in the ECS Bucket Management page.
Virtual Array	The virtual array in which the bucket belongs.
Object Virtual Pool	The virtual pool associated with the virtual array in which the bucket belongs.
Protocols	The protocol assigned to the bucket's object virtual pool, which can be S3, Atmos, and Openstack Swift.

View the details of an individual bucket

To view the details of an individual bucket, and to see the tasks performed from ViPR Controller on the bucket:

1. Go to the **Resources > Buckets** page.
2. Click the bucket name.

Delete the bucket

To delete a bucket:

1. Go to the **Resources > Buckets** page.
2. Check the selection box next to the bucket, or click the bucket name to go to the bucket details page.
3. Open the **Delete Buckets** menu.
4. Select:
 - **Full** — to delete the bucket from the storage system, and ViPR Controller.

- **Inventory only** — to delete the bucket only from ViPR Controller. You will need to remove the bucket from the storage system outside of ViPR Controller.

Note

If a bucket needs to be deleted from both the ViPR Controller database, and the ECS a full delete should be used. If the full delete fails, because ViPR Controller did not detect the bucket on the ECS, then you can use the **Inventory Only** delete to remove the bucket from the ViPR Controller database.

Add (or edit) access controls for a bucket

The user, group, or custom group must have been configured for the ECS prior to assigning the user access control to the bucket from ViPR Controller.

To add (or edit) access control to a bucket:

1. Go to the **Resources > Buckets** page.
2. Click the bucket name.
3. Click **Manage ACLs**.
4. Click **Add** to add new access controls to the bucket, or click the access control name in the list to edit the access control permissions.
5. Enter the following information:

Option	Description
Type	The type of access control list: User, Group, or Custom Group.
Name	The domain name, which is configured on the ECS for the user name, group name, or custom group name. When adding access control to buckets, you cannot use spaces in the user, group, or custom group names.
Domain	Used when the user, group, or custom group has been configured as a domain on the ECS.
Permissions	The access permission that will be allowed for the user, group, or custom group assigned to this ACL: <ul style="list-style-type: none"> • Full Control — grants the users for this access control entry all permissions. • None — will give the users no access to the bucket. • Type the permission in the entry box to locate it in the list of possible permissions, or scroll the list of permissions, and check the ones to apply to the entry. Refer to ECS documentation for permission details.

6. If you want to add multiple access control entries, click **Add**, and repeat step 5.
7. Click **Save**.

Resource tasks

Use the **Resources > Tasks** page to view the list of system-level or tenant-level tasks and their progress. Any user with access to a tenant can use this page to view the tasks associated with that tenant.

The **Tasks** page contains the following information about tenant-level tasks or system-level tasks.

Column or tab name	Description
Tenant	<p>Lists the tenant-level tasks, such as create a host or create a volume. Includes a count of:</p> <ul style="list-style-type: none"> • pending tasks • tasks that stopped due to an error • tasks that completed successfully <p>Any user who has access to a tenant can view the tasks associated with that tenant.</p>
System	<p>Lists the system-level tasks that are not associated with any tenant, such as discovering an array. Includes a count of:</p> <ul style="list-style-type: none"> • pending tasks • tasks that stopped due to an error • tasks that completed successfully <p>Only System Administrators can view system tasks.</p>
Name	The name of the operation generating the task.
Resource	The name of the resource for which the task was created. Selecting the resource provides access to the page to edit the resource.
Progress	<p>The percentage complete for the task. The color of the bar indicates the success or failure of the task:</p> <ul style="list-style-type: none"> • Green indicates that the task completed successfully. • Red indicates that there was an error.
State	<p>The state of the task:</p> <ul style="list-style-type: none"> • Complete • Pending • Error
Start	Indicates when the task was started.
Elapsed	The time it took for the task to finish.

Actionable events

Use the **Resources > Actionable Events** page to review the list of Pending, Failed, Approved, or Declined events. In addition to viewing events, you can click the event and delete associated tasks.

If changes occur during vCenter or Host discovery, ViPR Controller may need to trigger export group updates in order to maintain the correct state among hosts, clusters, and their export groups. Instead of performing these updates automatically, a list of actionable events is generated. A Tenant Administrator can approve or decline each actionable event.

Multiple actionable events for the same export group can be approved and run concurrently.

Default behavior for export group updates varies:

- For vCenter, Windows, Linux, AIX, or HP-UX-discovered hosts, you must use the **Resources > Actionable Events** page to manage export group updates. Automatic export is turned off.
- Automatic export is on by default for manual or user-created clusters when moving hosts between clusters in the UI.
- Automatic export is off by default when using the CLI commands.
- There is no automatic export for NFS exports. Actionable events are created only if the host is in a shared block export group and is being removed/added to a cluster. An actionable event is created if the host moved to a different datacenter or if it was added/removed from vCenter.
- If a host is removed from vCenter (not discoverable at all through vCenter and not the last one in a cluster), then an actionable event is created even if it doesn't have any block exports. Approving the event will unassign the host from vCenter but not perform any block export updates.

Click > to access more detailed information for each actionable event.

Table 42 Actionable Events page

Column or tab name	Description
Warning	Lists the actionable event that occurred, such as addition or removal of a host or initiator to a cluster, a cluster down, and so forth. Any user who has access to a tenant can view the actionable events associated with that tenant.
Resource	The name of the resource for which the actionable event was created. Selecting the resource provides access to the page to edit the resource.
Event	The name of the operation generating the actionable event.
Status	The state of the actionable event: <ul style="list-style-type: none"> • Pending • Failed • Approved • Declined

Table 42 Actionable Events page (continued)

Column or tab name	Description
Created At	Indicates when the actionable event started.

Service catalog behavior associated with pending events

The "horn" icon in the UI banner displays the number of pending and failed events. This number is refreshed every 10 seconds. You must Approve or Decline these events before you can successfully process orders for the affected hosts or clusters. If the event fails, you must correct the underlying problem and Approve or Decline the event before submitting new orders.

Click the > button to open the details panel for the actionable event. The "If Approved" and "If Declined" text explains what will happen if you choose "Approve" or "Decline." If a link is available in the Event Tasks row, click it to get more information or to confirm your action. When you Approve or Decline an event, you must type `confirm` in the dialog in order for the action to take place. If the action fails, there will not be a link associated with Event Task. Use the Event ID URL to obtain more information or to report the problem for troubleshooting.

CHAPTER 9

ViPR Controller User Roles

The following topics are contained in this chapter:

- [ViPR Controller user role requirements](#)..... 124

ViPR Controller user role requirements

ViPR Controller roles fall into two groups: roles that exist at the ViPR Controller virtual data center level, and roles that exist at the tenant level.

Note

Access to different areas of the ViPR Controller UI is governed by the actions permitted to the role assigned to the user. The actions authorized when you access ViPR Controller from the UI can differ (be more constrained) from those available when you use the REST API or CLI.

Virtual data center-level roles

VDC roles are used to set up the ViPR Controller environment which is shared by all tenants. The following table lists the authorized actions for each user role at the virtual data center level.

Table 43 VDC roles

VDC Role	Authorized Actions
Security Administrator	<ul style="list-style-type: none"> • Manages the authentication provider configuration for the ViPR Controller virtual data center to identify and authenticate users. Authentication providers are configured to: <ul style="list-style-type: none"> ▪ Use Active Directory/Lightweight Directory Access Protocol (AD/LDAP) user accounts/domains to add specified users into ViPR Controller. ▪ Register ViPR Controller as block storage service in Openstack (Keystone). <hr/> <p>Note</p> <p>Security Administrator role is required to add Keystone, but Keystone users cannot be added into ViPR Controller.</p> <hr/> <ul style="list-style-type: none"> • Creates ViPR Controller User Groups. • Assigns VDC and Tenant roles. • Sets ACL assignments for Projects, and Service Catalog. • Sets ACL assignments for virtual arrays, and virtual pools, from the ViPR Controller API and CLI. • Update vCenter Tenants (ACLs) and Datacenter Tenant from ViPR Controller REST API and CLI (Only System Administrators can perform any of these functions from the ViPR Controller UI). • Creates, modifies, and deletes sub-tenants. • Assigns the tenant quotas, and user mappings. • Manages ViPR Controller virtual data center software and license updates. • Configures the repository from which ViPR Controller upgrade files will be downloaded and installed. • Manages SSL, and trusted certificates.

Table 43 VDC roles (continued)

VDC Role	Authorized Actions
	<ul style="list-style-type: none"> • Can change IPs for ViPR Controller nodes deployed on VMware without a vApp, and Hyper-V. • Schedule backups of ViPR Controller instances. • Reset local user passwords. • Configures ACLs. • Restores access to tenants and projects, if needed. (For example, if the Tenant Administrator locks himself/herself out, the Security Administrator can reset user roles to restore access.) • Can add or change ViPR Controller node names. • Initiate a minority node recovery from the ViPR Controller REST API, and CLI. • View the minority node recovery status from the ViPR Controller CLI. • Make changes to the ViPR Controller, General Configuration, Security settings. • Shuts down, reboots, and restarts ViPR Controller services from the ViPR Controller REST API/CLI. • Manages IPsec actions, such as rotate IPsec key, check IPsec status. <p>The Security Administrator must also be assigned a System Administrator role to perform the following operations from the ViPR Controller UI:</p> <ul style="list-style-type: none"> • Shut down, reboot, and restart ViPR Controller nodes or services. • Set ACL assignments for virtual arrays, and virtual pools. • Initiate a minority node recovery. <p>In Geo-federated Environment:</p> <ul style="list-style-type: none"> • Has Security Administrator privileges on authentication providers, which are global resources.
System Administrator	<ul style="list-style-type: none"> • Performs system upgrades. • Creates system backups • Add ViPR Controller licenses. • Send support requests. • Add, edit, delete, disconnect, and reconnect virtual data centers (VDCs). • Sets up the physical storage infrastructure of the ViPR Controller virtual data center and configures the physical storage into two types of virtual resources: virtual arrays and virtual pools. Authorized actions include: <ul style="list-style-type: none"> ▪ Adding, modifying, and deleting the following physical storage resources into ViPR Controller such as storage systems, storage ports, and storage pools, data protection systems, fabric managers, networks, compute images, Vblock compute systems, and vCenters.

Table 43 VDC roles (continued)

VDC Role	Authorized Actions
	<p data-bbox="730 352 783 378">Note</p> <p data-bbox="730 396 1385 453">System Administrators cannot add, delete, or modify hosts or clusters.</p> <ul style="list-style-type: none"> <li data-bbox="694 480 1458 537">▪ Updating vCenter cascade tenancy and vCenter tenants (ACLs) and Datacenter Tenant from the ViPR Controller REST API, UI and CLI. <li data-bbox="694 558 1458 615">▪ Associate a vNAS server to one or more projects (Requires both the System and Tenant Administrator roles). <li data-bbox="694 636 967 661">▪ Creating virtual pools. <li data-bbox="694 682 975 707">▪ Creating virtual arrays. <li data-bbox="694 728 999 753">▪ Creating mobility groups. <li data-bbox="694 774 1238 800">▪ Create/Modify File Protection Policy templates <li data-bbox="694 821 1235 846">▪ File Protection Policy assignment at vPool level <li data-bbox="694 867 1430 924">▪ File Protection Policy assignment at project and file system levels (also need System Monitor privileges) <li data-bbox="651 945 1453 1001">• Manages the ViPR Controller virtual data center resources that tenants do not manage. <li data-bbox="651 1022 1358 1079">• Retrieves ViPR Controller virtual data center status and health information. <li data-bbox="651 1100 1402 1157">• Retrieves bulk event and statistical records for the ViPR Controller virtual data center. <li data-bbox="651 1178 1129 1203">• View the Database Housekeeping Status. <li data-bbox="651 1224 1430 1249">• View the minority node recovery status from the ViPR Controller CLI. <p data-bbox="651 1270 983 1295">In Geo-federated Environment:</p> <ul style="list-style-type: none"> <li data-bbox="651 1316 1227 1341">• Adds a VDC to create Geo-federated environment <li data-bbox="651 1362 1166 1388">• Add, disconnect, reconnect, or delete a VDC <li data-bbox="651 1409 1430 1465">• Has System Administrator privileges on global virtual pools, which are global resources. <li data-bbox="651 1486 1466 1543">• Sets ACL assignments for virtual arrays, and virtual pools, from the ViPR Controller API
System Monitor	<ul style="list-style-type: none"> <li data-bbox="651 1577 1453 1667">• Has read-only access to all resources in the ViPR Controller virtual data center. Has no visibility into security-related resources, such as authentication providers, ACLs, and role assignments. <li data-bbox="651 1688 1402 1745">• Retrieves bulk event and statistical records for the ViPR Controller virtual data center. <li data-bbox="651 1766 1358 1822">• Retrieves ViPR Controller virtual data center status and health information. <li data-bbox="651 1843 1453 1900">• (API only) Can create an alert event, with error logs attached, as an aid to troubleshooting. The alert event is sent to ConnectEMC. <li data-bbox="651 1921 1129 1946">• View the Database Housekeeping Status.

Table 43 VDC roles (continued)

VDC Role	Authorized Actions
	<ul style="list-style-type: none"> View the minority node recovery status from the ViPR Controller UI, and CLI. List backups from external server. Check upload status of a backup. Check restore status. Create/Modify File Protection Policy templates Create vPools File Protection Policy assignment at project and file system levels (also need System Administrator privileges)
System Auditor	Has read-only access to the ViPR Controller virtual data center audit logs.

Tenant-level roles

Tenant roles are used to administrate the tenant-specific settings, such as the service catalog and projects, and to assign additional users to tenant roles. The following table lists the authorized actions for each user role at the tenant level.

Table 44 Tenant roles

Tenant-Level Role	Authorized Actions
Tenant Administrator	<ul style="list-style-type: none"> Becomes Tenant Administrator of created tenant. A single-tenant enterprise private cloud environment has only one tenant, the Provider Tenant, and Tenant Administrators have access to all projects. Modifies the name and description of the tenants. Add vCenters to ViPR Controller physical assets in their own tenant. Manages tenant resources, such as Hosts, Clusters vCenters, and Projects. Configures ACLs for projects and the Service Catalog in their tenant. Assigns roles to tenant users. (Can assign Tenant Administrator or Project Administrator roles to other users.) Create Schedule Policies. Associate a vNAS server to one or more projects (Requires both the System and Tenant Administrator roles). Manage application services. Accept or decline actionable events. Edit service order schedules. Download orders (for archiving). Delete orders.

Table 44 Tenant roles (continued)

Tenant-Level Role	Authorized Actions
	<ul style="list-style-type: none"> • Adjust the fiber channel paths on an existing VPLEX or VMAX export. <hr/> <p>Note</p> <p>A user or group of users can be configured to have a Tenant Administrator role for Multiple Tenants. This user/group of users must belong to the Provider Tenant. However, they do not have to have the Tenant Administrator role in the provider tenant. This functionality can be used in multi-tenant environments in cases where a group of users needs to perform provisioning operations for multiple tenants and they do not want to use root user for these operations.</p> <hr/> <p>In Geo-federated Environment:</p> <ul style="list-style-type: none"> • Has Tenant Administrator privileges on tenants, which are global resources.
Tenant Approver	<ul style="list-style-type: none"> • Approves or rejects Service Catalog orders in their tenant. • Views all approval requests in their tenant.
Project Administrator	<ul style="list-style-type: none"> • Creates projects in their tenant and obtains an <code>OWN</code> ACL on the created project.

CHAPTER 10

Troubleshooting Error Messages

This chapter contains the following topics:

- [Troubleshooting ViPR Controller error messages](#).....130

Troubleshooting ViPR Controller error messages

Review this information for common ViPR Controller error messages and their resolutions.

Troubleshooting common error messages

Table 45 Troubleshooting tips for common error messages

Error message	Description	Resolution/Workaround
UI: Failed command to provision storage resource	The provisioning operation failed because: <ul style="list-style-type: none"> the network connection between ViPR and the storage array was lost Solutions Enabler is offline 	Do the following: <ul style="list-style-type: none"> Check your network connections Restart Solutions Enabler if it is offline.
API: Error Message/Code = ...	The provisioning operation failed because: <ul style="list-style-type: none"> the network connection between ViPR and the storage array was lost Solutions Enabler is offline 	Do the following: <ul style="list-style-type: none"> Check your network connections Restart Solutions Enabler if it is offline.
Logs: ConnectException: Connection refused: ... while sending command to the storage system	The provisioning operation failed because: <ul style="list-style-type: none"> the network connection between ViPR and the storage array was lost Solutions Enabler is offline 	Do the following: <ul style="list-style-type: none"> Check your network connections Restart Solutions Enabler if it is offline.
The target namespace does not exist. (Invalid namespace root/brocadel)	The SMI-S discovery for an array or switch failed because an array provider was added instead of a switch provider.	Delete the array provider and enter the IP address and port information for the correct switch provider.
Config change failed could not find disks that satisfy our mirror/raid policy	Creating a volume failed because the VMAX storage pool does not have a disk with a matching SymWin policy.	Add more disks to the storage pool.

Table 45 Troubleshooting tips for common error messages (continued)

Error message	Description	Resolution/Workaround
Failed to get array system info (Authorization failed)	The NetApp discovery failed because the user account does not have administrative privileges.	Add administrative privileges to the users account using the NetApp CLI.
Storage Array: 'FOO' is not registered. It can not be edited	Discovery failed because the storage array is not registered and can not be edited.	Register the storage array.
Dashboard (if accessible) may show network or VIP ERROR (System Health tab -> Diagnostics)	The system network virtual IP address, or a Controller VM IP address, is incorrect or invalid, resulting in the user being unable to login after deployment and all management and provisioning actions fail.	Redeploy the ViPR virtual appliance, or change the system IP addresses of the virtual appliance using Edit Settings in vCenter.
Invalid Username or Password	The username or password is incorrect. A username must have a domain suffix and passwords are case sensitive.	Retry your username and password.
Manager authentication with LDAP server failed. Please contact your administrator if the problem persists	The authentication provider is registered incorrectly, or the password of the user registering the authentication provider has expired or was changed.	Contact the system administrator to update the authentication provider with the correct manage domain name and valid password.
[MiscStage:1] ERROR CassandraDaemon.java (line 164) Exception in thread Thread[MiscStage:1,5,main] java.lang.NullPointerException [GossipTasks:1] INFO Gossiper.java (line 768) InetAddress /	A known issue for ViPR installations utilizing three nodes.	Ignore the error.

Table 45 Troubleshooting tips for common error messages (continued)

Error message	Description	Resolution/Workaround
xx.xxx.xx.xxx is now dead		
<pre>svcuser@node1:/etc> ./diagtool sed: can't read /etc/ovf-env.properties: Permission denied</pre>	<p>A permissions error when the svc user executes the diagnostic tool (diagtool).</p>	<p>When executing the diagtool, the svc user should use the sudo command. For example:</p> <pre>sudo /etc/diagtool</pre>
Certificate error	<p>Unable to log in using a browser after an upgrade or property reconfiguration because of SSL certificate changes.</p>	<p>Do the following:</p> <ul style="list-style-type: none"> • Clear your certificates, cookies, cache, and history, and then restart your browser. • If the error is received after restarting your browser, restart the system running the browser.
N/A	<p>An SMI-S Provider can be registered twice.</p>	<p>Do not register SMI-S Providers more than once.</p>
No Storage Found	<p>The Storage Pools list is empty in a virtual storage pool, or provisioning failed when no storage was found. These errors are caused because the available networks are not assigned to the associated virtual storage array.</p>	<p>Do the following:</p> <ol style="list-style-type: none"> 1. Ensure all required switches are discovered. 2. Ensure the necessary IP network is created, and the storage ports are assigned to it. 3. Ensure the network is assigned to the corresponding virtual storage array.
N/A	<p>After deleting an SMI-S Provider managed storage array, the storage array is not rediscovered and is marked for permanent exclusion from ViPR.</p>	<p>To use a storage system not managed by ViPR:</p> <ol style="list-style-type: none"> 1. De-register the storage array. 2. Register the storage array with ViPR.
<pre>2013-08-29 12:32:18,242 [GossipStage:1] INFO Gossiper.java (line 754) InetAddress / a.b.c.d is now UP 2013-08-29 12:32:55,971 [GossipTasks:1]</pre>	<ul style="list-style-type: none"> • Multiple ViPR nodes have the same IP address • There is a high load on ViPR and the CPU or memory is almost exhausted • The network is unstable, the connection between nodes is turning off and on 	<p>Determine which of the problems is occurring. Depending on the problem, you may need to redeploy ViPR.</p>

Table 45 Troubleshooting tips for common error messages (continued)

Error message	Description	Resolution/Workaround
<pre>INFO Gossiper.java (line 768) InetAddress / a.b.c.d is now dead.</pre>	<ul style="list-style-type: none"> There are too many concurrent create and delete operations on the database The disk space is exhausted or almost exhausted 	
<pre>Connection refused or authentication failed</pre>	<p>The Windows host was not added to ViPR after configuring WinRM.</p>	<p>Set the following properties in the WinRM configuration file for local user authentication:</p> <ul style="list-style-type: none"> winrm get winrm/config/service winrm set winrm/config/service/auth @{Basic="true"} <p>For domain authentication with Kerberos,</p> <ul style="list-style-type: none"> winrm get winrm/config/service winrm set winrm/config/service/auth @{Kerberos="true"}
<pre>Run date on each nodes, the time is not the same among nodes</pre>	<p>The ViPR node times are not synchronized. This can be caused by:</p> <ul style="list-style-type: none"> The NTPD service is down. The /etc/ntp.conf file contains an invalid NTP server. 	<p>Do the following:</p> <ol style="list-style-type: none"> Run an NTP diagnostic test. Resolve the problem based on the test results: <ul style="list-style-type: none"> UNCONFIGURED — Configure the NTP setting in System > Configuration > Network. CONFIGURED UNREACHABLE — Check the NTP settings and the status of the NTP server. CONFIGURED DEGRADED — Check the NTP settings and the status of the NTP server.
<pre>An error occurred while finding a suitable placement to handle the request (code: 1034). no IP networks found</pre>	<p>The host IP address is not set in the virtual storage array network settings.</p>	<p>Do the following:</p> <ol style="list-style-type: none"> In the virtual storage array settings, click Edit Network. Type the file host IP address. Click OK.

Table 45 Troubleshooting tips for common error messages (continued)

Error message	Description	Resolution/Workaround
<p>The vSphere HA agent on host 'hostname' failed to quiesce file activity on datastore '/vmfs/volumes/[id]'. To proceed with the operation to unmount or remove a datastore, ensure that the datastore is accessible, the host is reachable and its vSphere HA agent is running.</p>	<p>The vSphere HA agent failed to unmount or remove a datastore. The datastore is not accessible or the vSphere HA agent is not running.</p>	<p>Download vCenter Server 5.1 Update 1a. You can download the latest version from the <i>VMware vCloud Suite Download Center</i>.</p>
<p>ViPR virtual appliance is not accessible or status remains at Degraded.</p>	<p>Invalid IPv4 network netmask or network gateway.</p>	<p>Shutdown the ViPR virtual appliance, and update the system IP address and netmask of the virtual appliance using Edit Settings in vCenter.</p>
	<p>Invalid IPv6 prefix length or network gateway.</p>	<p>Shutdown the ViPR virtual appliance, and update the system IP address and netmask of the virtual appliance using Edit Settings in vCenter.</p>
<p>Service Unavailable (6503) The service is currently unavailable because a connection failed to a core component. Please contact an administrator or try again later.</p>	<p>The ViPR UI was opened before all ViPR services were started.</p>	<p>Wait 5 minutes after ViPR controller deployment before running the UI.</p>
<p>ViPR virtual appliance remains in Syncing state</p>	<p>Credentials for an account with insufficient privileges were used to download the img file during upgrade.</p>	<p>1. Use the ViPR CLI to check the virtual appliance state. Make sure current version is still 1.0.0.7.1065 (V1.0) or whatever the pre-upgrade version should be, and the CLUSTER_STATE is</p>

Table 45 Troubleshooting tips for common error messages (continued)

Error message	Description	Resolution/Workaround
		<p>SYNCING. # ./viprcli system get-cluster-state</p> <p>2. Using remove-image command with force flag (-f), remove the image that failed to download: #./viprcli system remove-image -f vipr-1.0.0.8.103</p> <p>3. At this point the ViPR virtual appliance should return to Stable, and you should be able to upgrade after supplying credentials with correct permissions.</p>
<p>Error 999 (http: 500): An unexpected error occurred, please check the ViPR logs for more information.</p>	<p>A user attempts to create a bucket in the ViPR user interface although no datastores are in the services virtual pool, resulting in a failed operation.</p>	<p>Before creating a bucket, ensure the services virtual pool providing the storage for the bucket contains at least one datastore.</p>
<p>Error 16000: Error occurred running an SMIS command. The job has failed: string ErrorDescription = "Volume Delete failed: C:ERROR_CLASS_SOFTWARE F:ERROR_FAMILY_FAILED R:1000086 L: 2 C:ERROR_CLASS_SOFTWARE F:ERROR_FAMILY_FAILED R:1000086 Failed to acquire the requested lock : \"Unable to write-protect selected device \" : 2 : 2550 : \"Unable to</p>	<p>Unable to delete a volume on a VMAX storage array.</p>	<p>The error message indicates there is a lock on the volume because another user is accessing it. Wait and perform the delete operation again once no other users are accessing the volume.</p>

Table 45 Troubleshooting tips for common error messages (continued)

Error message	Description	Resolution/Workaround
<pre> acquire the Symmetrix device lock\" @ [1] com.emc.cmp.osls. se.osl.Device.Sto rDeviceDelete(): 150 [0] com.emc.cmp.osls. se.array.job.JOB_ VolDelete.run(): 136 "; Rollback error: The job has failed: string ErrorDescription = "Volume Delete failed: C:ERROR_CLASS_SOF TWARE F:ERROR_FAMILY_FA ILED R:1000086 L: 2 C:ERROR_CLASS_SOF TWARE F:ERROR_FAMILY_FA ILED R:1000086 Failed to acquire the requested lock : \"Unable to write-protect selected device \" : 2 : 2550 : \"Unable to acquire the Symmetrix device lock\" @ [1] com.emc.cmp.osls. se.osl.Device.Sto rDeviceDelete(): 150 [0] com.emc.cmp.osls. se.array.job.JOB_ VolDelete.run(): 136 "; </pre>		
<pre> ERROR Error 40009 (http: 400): "Invalid bucket </pre>	<p>The bucket name contains invalid characters.</p>	<p>Rename the bucket using valid characters.</p>

Table 45 Troubleshooting tips for common error messages (continued)

Error message	Description	Resolution/Workaround
<pre>name". Invalid bucket Name test this com.emc.vipr.clie nt.exceptions.Ser viceErrorException: Error 40009 (http: 400): "Invalid bucket name".</pre>		
<pre>ERROR HDFS service failed java.io.IOExcepti on: ClientApi failed to initialize, status=ERROR_INTE RNAL HDFS service failed java.io.IOExcepti on: ClientApi failed to initialize, status=ERROR_INTE RNAL</pre>	<p>After initial deployment of ViPR, errors appear when switching to LOG view.</p>	<p>This error occurs when the HDFS service starts up faster than the services. Ignore the error.</p>
<pre>createExportMask failed - maskName: urn:storageos:Exp ortMask:d101e3a5- 146b-4a26-916e- f3bc5112a62c:vdcl WBEMException: CIM_ERR_FAILED (A general error occurred that is not covered by a more specific error code. (com.emc.cmp.osls .se.osl.Masking.S torEndptGroupCrea te():1872 C:ERROR_CLASS_SOF TWARE F:ERROR_FAMILY_FA</pre>	<p>A duplicate network was discovered by ViPR and caused ViPR to reuse the same ports to recreate the initiator groups.</p>	<p>Remove the physical assets from the masking view, and then add the physical assets back to the masking view.</p>

Table 45 Troubleshooting tips for common error messages (continued)

Error message	Description	Resolution/Workaround
<pre>ILED R:1000124 L: 2 C:ERROR_CLASS_SOFTWARE F:ERROR_FAMILY_FAILED R:1000124 The specified WWN is already in use : "StorEndptGroupCreate failed" : 2 : 3568 : "The specified WWN is already in use"))</pre>		
<pre>Host operation failed: Host <ESX/ESXi host> not reachable in state UNREACHABLE - Ensure host is powered on and responsive. Can be caused by intermittent or temporary connectivity issue thus retry</pre>	<p>During the VCE Vblock System Service, Provision Cluster operation, ViPR:</p> <ol style="list-style-type: none"> 1. Creates the ESX hosts. 2. Creates the cluster in vCenter. <p>During the create the cluster in vCenter operation, ViPR adds the newly created ESX hosts to the vCenter cluster. When ViPR attempts to add the ESX hosts to the vCenter cluster before one or more of the ESX hosts have been started, the <code>Host not reachable</code> error occurs because the hosts have not completely rebooted and are not ready to be added to the cluster until they have been started.</p>	<p>To resolve the issue, use the Update vCenter Cluster service from the ViPR Service Catalog to update the vCenter cluster with the newly created hosts.</p> <p>Optionally, to avoid the error during future operations, increase the ViPR default vCenter host operation timeout value.</p> <p>To increase the timeout value:</p> <ol style="list-style-type: none"> 1. Get a list of all configuration properties from the ViPR REST API. GET on <code>https://<ViPR Host>:4443/config/properties</code> 2. Change the property for vCenter host operation timeout. PUT to <code>https://<ViPR Host>:4443/config/properties</code> <p>Allowed values, specified in seconds, are: 60, 150, 300, 450, 600, 750, 900, 1800</p> <hr/> <p>Note</p> <p>Default value is 450 seconds (7.5 minutes). See example following this table showing how to increase the timeout value.</p> <hr/>

Table 45 Troubleshooting tips for common error messages (continued)

Error message	Description	Resolution/Workaround
Error 12025: Export operation failed due to existence of non FAST volumes in storage group.. While attempting to export a FAST volume, an existing Storage Group PRGDC_2 was found on the array with non-FAST volumes in it. Adding FAST volumes to this Storage Group is not permissible.	Creating a block volume on a virtual pool with a FAST VP policy, failed.	Create two cascaded storage groups: <ul style="list-style-type: none"> FAST VP volumes non-FAST VP volumes <hr/> <p>Note</p> <p>This solution is an offline operation for VMAX w5876 code, if the storage group to be reconstructed is part of a masking view.</p> <hr/>
Error 1013 (http: 400): Bad request body. Cannot change the virtual pool pathsPerInitiator parameter for ExportGroup somexyz.exp.grp.com ExportMask somexyzexpgrpcom.	Moving volumes from one virtual pool to another fails if there is more than one target per initiator.	This operation is not supported.

Example for increasing the timeout value:

```
<property_update>
<properties>
  <entry>
    <key>vcenter_host_operation_timeout</key>
    <value>900</value>
  </entry>
</properties>
</property_update>
```

Troubleshooting Active Directory and LDAP

Table 46 Troubleshooting tips for Active Directory and LDAP

Symptom	Cause	Resolution/Workaround
Access forbidden: Authentication required, and log contains ERROR	Invalid format of username	Match the username with the searchfilter used. For example: userName=%u

Table 46 Troubleshooting tips for Active Directory and LDAP (continued)

Symptom	Cause	Resolution/Workaround
CustomAuthenticationManager.java (line 99) Unsupported credentials admin \adc34103		means a username of the format foo@bar.com.
Search failed while trying to find user in ldap tree	User not found because user name does not exist within the searchbase.	Be sure you have specified the searchbase at the correct location in the tree.
	User not found because user name types do not match the filter.	Be sure you are using %u versus %U properly to match complete versus local part of name.
	There is more than one match, based on the filter.	Check the value of the search filter.
Bind problems when adding a new authentication provider	Special characters exist in the managerDN name.	To specify the managerDN value, copy the contents of the user's distinguishedName value from Active Directory Users and Computers, Properties, Attribute Editor. That value will have the proper escape characters.
Authentication issue and log contains: LDAP: error code 49 - 80090308: LdapErr: DSID-0Cxxxxxx, comment: AcceptSecurityContext error, data xxx, vece	xxx is an Active Directory error code.	Refer to Active Directory documentation for the error code.

Troubleshooting administrator tasks

Table 47 Troubleshooting tips for administrator tasks

Symptom	Resolution/Workaround
No matching storage pools displayed when creating a virtual pool for IP connected file storage.	Ensure a file array has been added to a network in the virtual array.
No IP network found to satisfy user request.	If a user is attaching provisioned storage to an IP-connected host, the host IP address or hostname must be added to the IP network.
MultiVolumeConsistency is set to true but no consistency group is provided.	If consistency groups are enabled on a virtual host, a resource is not created unless a user selects a consistency group to add it to.

Table 47 Troubleshooting tips for administrator tasks (continued)

Symptom	Resolution/Workaround
No volumes are displayed when a user attempts to create a snapshot.	The virtual storage pool must have the maximum number of snapshots set to at least 1.
RAID groups created with unbound RAID levels cannot be used in ViPR because the capacity provider is reporting 0 free capacity.	Do the following: <ol style="list-style-type: none"> 1. Create a RAID group with unbound RAID levels. 2. Create a small volume on the RAID group.
Unable to login when IPv6 prefix is set to the wrong value.	Update the system settings of the ViPR virtual appliance using Edit Settings in vCenter.

