

EMC[®] ViPR[®] Controller

Version 2.3

Virtual Data Center Requirements and Information Guide

302-002-068

REV 01

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

Published July, 2015

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

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

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

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

CONTENTS

Chapter 1	ViPR Controller VDC Requirements and Information Overview	5
Chapter 2	Role Requirements for the ViPR Controller Virtual Data Center	7
	Required VMware VirtualCenter role privileges.....	8
	Role requirements for VDC configuration.....	8
Chapter 3	Physical Asset Requirements and Information	9
	Storage systems.....	10
	Hitachi Data Systems	10
	Enable performance data logging of HDS storage arrays in Storage Navigator.....	11
	Create a new user using the Storage navigator on HDS	12
	Overview.....	12
	EMC VMAX.....	14
	SMI-S provider configuration requirements for VMAX.....	14
	SMI-S provider version requirements for SRDF.....	15
	Upgrade requirements.....	16
	VMAX storage system.....	16
	EMC VNX for Block	17
	SMI-S provider configuration requirements for VNX for Block.....	17
	VNX for Block storage system	18
	VPLEX	19
	Third-party block storage (OpenStack)	19
	Third-party block storage provider installation requirements.....	19
	Third-party block storage system support	20
	Supported ViPR Controller operations.....	20
	OpenStack configuration.....	21
	Cinder service configuration requirements.....	21
	Storage system (backend) configuration settings.....	21
	Volume setup requirements for OpenStack.....	22
	ViPR Controller configuration for third-party block storage ports.....	23
	Stand-alone EMC ScaleIO.....	24
	EMC XtremIO	24
	EMC VNXe.....	24
	General file storage system requirements.....	25
	EMC® Data Domain®	25
	EMC Isilon	25
	NetApp 7-mode	26
	NetApp Cluster-mode	26
	EMC VNX for File	27
	RecoverPoint systems.....	27
	Fabric Managers (switches).....	27
	ViPR Controller switch and network terminology.....	27
	Brocade.....	28
	ViPR Controller support for fibre channel routing with Brocade switches.....	30
	Brocade Fibre Channel routing configuration requirements.....	31
	Cisco.....	31

ViPR Controller support for Inter-VSAN Routing with Cisco switches	32
Cisco Inter-VSAN Routing configuration requirements	33
Vblock compute systems	34
Compute images	34
Hosts	34
AIX hosts and AIX VIO Servers	35
HP-UX	35
Linux	35
Windows	35
VMware® vCenter	37

Chapter 4	Virtual Asset Requirements and Information	39
Overview	40	
Plan to build your virtual array	40	
SAN zoning requirements	40	
Plan how to add the physical assets to the virtual array	40	
Virtual Array requirements for Vblock system services	42	
Block storage configuration considerations	43	
Hitachi Data Systems	43	
VMAX	44	
VMAX3	45	
EMC VNX for Block	45	
EMC VNXe for Block	46	
EMC VPLEX	46	
Third-party block (OpenStack) storage systems	46	
Block storage systems under ViPR Controller management	46	
File storage configuration considerations	47	
EMC® Data Domain®	47	
EMC VNX for File	47	
ViPR requirements for service profile templates	48	

CHAPTER 1

ViPR Controller VDC Requirements and Information Overview

This guide is for ViPR Controller System Administrators and Tenant Administrators to understand the information needed to configure the physical assets that will be added to the ViPR Controller Virtual Data Center (VDC), as well as the requirements and information to convert the ViPR Controller physical assets, discovered by the ViPR Controller, into the ViPR Controller networks, virtual arrays, and virtual pools of the ViPR Controller VDC.

Related documents

These guides explain how to configure VDC:

- *ViPR Controller User Interface Virtual Data Center Configuration Guide.*
- *ViPR Controller REST API Virtual Data Center Configuration Guide.*

Access both documents from the [ViPR Controller Product Documentation Index](#) .

The *ViPR Controller Support Matrix* provides the version requirements for the physical assets. You can find this document on the EMC Community Network at community.emc.com.

CHAPTER 2

Role Requirements for the ViPR Controller Virtual Data Center

This chapter contains the following topic:

- [Required VMware VirtualCenter role privileges](#) 8
- [Role requirements for VDC configuration](#) 8

Required VMware VirtualCenter role privileges

ViPR Controller requires that you have certain role privileges in VMware VirtualCenter to perform vCenter discovery and to create datastores.

Required datastore privileges

- Datastore.Allocate space
- Datastore.Browse datastore
- Datastore.Configure datastore
- Datastore.Remove datastore

Required host privileges

- Host.CIM.CIM Interaction
- Host.Configuration.Storage partition configuration

Role requirements for VDC configuration

The ViPR Controller System Administrators and Tenant Administrators are responsible for the configuration of the ViPR Controller VDC.

For complete details about ViPR Controller roles and access permissions, see the *ViPR Controller User Interface Tenants, Projects, Security, Users and Multisite Configuration Guide*.

ViPR Controller System Administrator operations

The ViPR Controller System Administrators are solely responsible for the following VDC set-up operations:

Adding and configuring these physical assets:

- Storage systems
- Data protection systems
- Fabric managers (Brocade, and Cisco switches)
- Networks
- Vblock compute systems
- Compute images

Adding and configuring these virtual assets:

- Networks
- Virtual arrays
- Block virtual pools
- File virtual pools
- Compute virtual pools

ViPR Controller Tenant Administrator operations

The Tenant Administrator role adds and configures hosts, clusters, and vCenters to the ViPR Controller physical assets.

CHAPTER 3

Physical Asset Requirements and Information

This chapter contains the following topics:

- [Storage systems](#)..... 10
- [Hitachi Data Systems](#) 10
- [EMC VMAX](#).....14
- [EMC VNX for Block](#) 17
- [VPLEX](#) 19
- [Third-party block storage \(OpenStack\)](#) 19
- [ViPR Controller configuration for third-party block storage ports](#)..... 23
- [Stand-alone EMC ScaleIO](#) 24
- [EMC XtremIO](#) 24
- [EMC VNXe](#)..... 24
- [General file storage system requirements](#).....25
- [EMC® Data Domain®](#)..... 25
- [EMC Isilon](#) 25
- [NetApp 7-mode](#) 26
- [NetApp Cluster-mode](#) 26
- [EMC VNX for File](#) 27
- [RecoverPoint systems](#)..... 27
- [Fabric Managers \(switches\)](#)..... 27
- [Vblock compute systems](#).....34
- [Compute images](#)..... 34
- [Hosts](#)..... 34
- [VMware® vCenter](#).....37

Storage systems

You can review the ViPR Controller requirements and information necessary for entering user credentials and setting up port-based metrics for storage systems, and review the information needed for adding and configuring a storage system in ViPR Controller.

Storage system user credentials

When adding storage systems to ViPR Controller, you specify the credentials of the user who will access the storage system from ViPR Controller. These credentials are independent of the user who is currently logged into ViPR Controller. All ViPR Controller operations performed on a storage system are executed as the user who was given access to the storage system.

ViPR Controller operations require that the ViPR Controller user has administrative privileges. If there are additional credential requirements for a specific type of storage system, it is explained in more detail in the following storage-specific sections of this guide.

Hitachi Data Systems

Before you add Hitachi Data Systems (HDS) storage to ViPR Controller, configure the storage as follows.

Gather the required information

Hitachi HiCommand Device Manager (HiCommand) is required to use HDS storage with ViPR Controller. You need to obtain the following information to configure and add the HiCommand manager to ViPR Controller.

Table 1 Hitachi required information

Setting	Value
A host or virtual machine for HiCommand Device Manager setup	
HiCommand Device Manager license	
HiCommand Device Manager host address	
HiCommand Device Manager user credentials	
HiCommand Device Manager host port (default is 2001)	

General configuration requirements

- HiCommand Device Manager software must be installed and licensed.
- Create a HiCommand Device Manager user for ViPR Controller to access the HDS storage. This user must have administrator privileges to the storage system to perform all ViPR Controller operations.
- HiCommand Device Manager must discover all Hitachi storage systems (that ViPR Controller will discover) before you can add them to ViPR Controller.
- When you add the HiCommand Device Manager as a ViPR Controller storage provider, all the storage systems that the storage provider manages will be added to ViPR Controller. If you do not want ViPR Controller to manage all the storage systems,

before you add the HiCommand Device Manager, configure the HiCommand Device Manager to manage only the storage systems that will be added to ViPR Controller.

Note

After you add the storage provider to ViPR Controller, you can deregister or delete storage systems that you will not use in ViPR Controller.

Configuration requirements for auto-tiering

ViPR Controller provides auto-tiering for the six standard HDS auto-tiering policies for Hitachi Dynamic Tiering (HDT) storage pools.

HDS auto-tiering requires the Hitachi Tiered Storage Manager license on the HDS storage system.

Configuration requirements for data protection features

HDS protection requires:

- Hitachi Thin Image Snapshot software for snapshot protection
- Hitachi Shadow Image Replication software for clone and mirror protection.

ViPR Controller requires the following is configured to use the Thin Image, and ShadowImage features:

- Requires Hitachi Replication Manager is installed on a separate server.
- The HiCommand Device Manager agent must be installed and running on a pair management server.
- To enable ViPR Controller to use the Shadow Image pair operations, create a ReplicationGroup on the HDS, named `ViPR-Replication-Group` using either the HiCommand Device Manager or Hitachi Storage Navigator .
- To enable ViPR Controller to use the Thin Image pair operations, create a SnapshotGroup on the HDS, named `ViPR-Snapshot-Group` using either the HiCommand Device Manager or Hitachi Storage Navigator .

Configuration requirements for ViPR Controller to collect HDS port metrics

Before ViPR Controller can collect statistical information for HDS storage ports, you must perform the following operations in the HDS Storage Navigator to enable performance logging on the storage system:

- [Enable performance data logging of HDS storage on page 11](#)
- [Create a new user using Storage Navigator on HDS on page 12](#)
- Enable the certificate on the HDS to access SMI-S in secure mode. For details, refer to the *Using the SMI-S function* section of the *Hitachi Command Suite Administrator guide*.

Enable performance data logging of HDS storage arrays in Storage Navigator

Before you can set up the metrics-based port selection for HDS, you must enable performance data logging for the HDS storage arrays in Storage navigator so that ViPR Controller can collect statistical information about the storage ports.

Procedure

1. Log into the Storage navigator.
2. Select **Performance Monitor**.
3. Click **Edit Monitor Switch**.
4. Enable **Monitor Switch**.

5. Set the **Sample Interval**.

Create a new user using the Storage navigator on HDS

To collect statistics using the embedded SMI-S provider, you must create a new user using the Storage Navigator on the Hitachi storage array.

Procedure

1. Log into the Storage Navigator.
2. Select **Administrator**.
3. Click **User Groups**.
4. Select **Storage Administrator (View Only) User Group**.
5. Click **Create User**.

Note

The username and password must be the same as the HiCommand Device Manager credentials that ViPR Controller uses to discover the storage provider.

6. Type a **User Name**.
7. Type a **Password**.
8. Type the **Account status**.
9. Type the **Authentication** details.

Overview

ViPR Controller System Administrators can learn about the Hitachi Data Systems Host Mode and Host Mode Option, how ViPR Controller sets the Host Mode and Host Mode Option, the ViPR Controller configuration requirements for ViPR Controller to apply the settings, and the steps to customize the Host Mode Option using the ViPR Controller UI.

Host Mode and Host Mode Options

Host Modes are HDS flags which are set on HDS host groups when an HDS storage volume is exported to the host group. The Host Mode is used to optimize the connection and communication between HDS storage and the host to which the HDS volume has been exported.

The Host Mode Options are a set of flags which can be enabled or disabled to further optimize the Host Mode set on the HDS host groups.

Refer to the Hitachi Data Systems documentation for details about the HDS Host Mode, and Host Mode Options.

How ViPR Controller sets the Host Mode

By default, when ViPR Controller is used to export an HDS volume to an AIX, ESX, HP-UX, Linux, or Windows host, ViPR Controller sets the following Host Mode on the host group by determining host operating system details.

Table 2 ViPR Controller default settings for Host Modes

Operating System	Default Host Mode
AIX	OF AIX

Table 2 ViPR Controller default settings for Host Modes (continued)

Operating System	Default Host Mode
ESX	21 VMware
HP-UX	03 HP
Linux	00 Standard
Windows	2C Windows Extension

Host Mode settings for "Other" hosts

When ViPR Controller is used to export an HDS volume to a host that was added to ViPR Controller as "Other," or to a host of which ViPR Controller cannot determine the type, the 00 Standard Host Mode is set on the host by the Hitachi HiCommand DM. 00 Standard is the HiCommand DM default Host Mode.

Changing the Host Mode

The ViPR Controller Host Mode settings are automatically set by ViPR Controller during an export operation, and you cannot change the ViPR Controller Host Mode settings. Once ViPR Controller has set the Host Mode on a host, you can use Hitachi Storage Navigator to change the Host Mode on the host.

Prior to ViPR Controller 2.2

ViPR Controller did not set the Host Mode for hosts prior to ViPR Controller 2.2. However, HiCommand Device Manager sets the 00 Standard Host Mode as the default on all HDS host groups.

If a host to which an HDS volume was exported by ViPR Controller, prior to ViPR Controller 2.2, is re-used by ViPR 2.2 or higher to export another HDS volume, the Host Mode will not be set by ViPR Controller on the host.

How ViPR Controller sets the Host Mode Option

By default, when ViPR Controller is used to export an HDS volume to an AIX, ESX, HP-UX, Linux, or Windows host, ViPR Controller sets the following Host Mode Options on the host group.

Table 3 ViPR Controller default settings for Host Mode Options

Operating System	Default Host Mode
AIX, Linux, Windows	2 => VERITAS Database Edition/Advanced Cluster 22 => Veritas Cluster Server
ESX	54 => Support option for the EXTENDED COPY command 63 => Support option for the vStorage APIs based on T10 standards
HP-UX	12 => No display for ghost LUN

Note

Refer to Hitachi storage system provisioning documentation for details about the Host Mode Options.

Host Mode Option settings for "Other," hosts

ViPR Controller does not set the Host Mode Option when ViPR Controller is used to export an HDS volume to a host that was added to ViPR Controller as "Other," or to a host of which ViPR Controller cannot determine the type.

Changing the Host Mode Option

The Host Mode Option is set on the host group the first time you export an HDS volume to the host. Once ViPR Controller has set the Host Mode Option on the host group, the Host Mode Option cannot be changed by ViPR Controller for example:

- If an HDS volume was exported to a host that was added to ViPR Controller as "Other," the Host Mode Option configured in ViPR Controller is not set on the host. If the type of host is changed in ViPR Controller from Other, to AIX, ESX, HP-UX, Linux, or Windows, and ViPR Controller is used to export another HDS volume to the same host, the Host Mode Options which are configured in ViPR Controller will not be set on the host by ViPR Controller, since it was not set on the host group the first time the volume was exported from ViPR Controller to that same host.
- If an HDS volume was exported to an AIX, ESX, HP-UX, Linux, or host, the Host Mode Options currently configured in ViPR Controller are set on the host. If the Host Mode Options are changed in ViPR Controller after the export, and ViPR Controller is used to export another HDS volume to the same host, the original Host Mode Options remain on the host, and are not configured with the new Host Mode Options.

Once the HDS volume has been exported to a host from ViPR Controller, you can change the Host Mode Setting from Hitachi Storage Navigator. If ViPR Controller is used to export another HDS volume to the same host after the Host Mode Option has been changed from Storage Navigator, ViPR Controller will reuse the Storage Navigator settings in the export.

Prior to ViPR 2.2

- The Host Mode options are only set to new host groups created using ViPR Controller 2.2 and higher.
- Prior to ViPR Controller 2.2, ViPR Controller did not set any Host Mode Options on HDS host groups.
- If a host group, that was created prior to ViPR Controller 2.2, is re-used by ViPR Controller, the ViPR Controller Host Mode Options are not set by ViPR Controller on the host group created prior to ViPR Controller 2.2.

EMC VMAX

ViPR Controller management of VMAX systems is performed through the EMC SMI-S provider. Your SMI-S provider and the VMAX storage system must be configured as follows before the storage system is added to ViPR Controller.

SMI-S provider configuration requirements for VMAX

You need specific information to validate that the SMI-S provider is correctly configured for ViPR Controller and to add storage systems to ViPR Controller.

Gather the required information

- SMI-S provider host address
- SMI-S provider credentials (default is admin/#1Password)
- SMI-S provider port (default is 5989)

Configuration requirements

SMI-S provider cannot be shared between ViPR Controller and any other application requiring an SMI-S provider to VMAX, such as EMC ViPR SRM.

Before adding VMAX storage to ViPR Controller, login to your SMI-S provider to ensure the following configuration requirements are met:

Note

The [ViPR Controller Support Matrix](#) has the most recent version requirements for all systems supported, or required by ViPR Controller. For specific version requirements of the SMI-S provider review the [ViPR Controller Support Matrix](#) before taking any action to upgrade or install the SMI-S provider for use with ViPR Controller.

-
- For VMAX storage systems, use either SMI-S provider 4.6.2 or SMI-S provider 8.0.3 but not both versions, however you must use 8.0.3 to use the new features provided with ViPR Controller 2.3. For new feature details, refer to the EMC ViPR Controller 2.3.0.0 Release Notes, which is available from EMC Support Zone (support.emc.com).
 - For VMAX3 storage systems, always use SMI-S provider 8.0.3.
 - The host server running Solutions Enabler (SYMAPI Server) and SMI-S provider (ECOM) differs from the server where the VMAX service processors are running.
 - The storage system is discovered in the SMI-S provider.
 - When the storage provider is added to ViPR Controller, all the storage systems managed by the storage provider will be added to ViPR Controller. If you do not want all the storage systems on an SMI-S provider to be managed by ViPR Controller, configure the SMI-S provider to only manage the storage systems that will be added to ViPR Controller, before adding the SMI-S provider to ViPR Controller.

Note

Storage systems that will not be used in ViPR Controller, can also be deregistered, or deleted after the storage provider is added to ViPR Controller. For steps to deregister or delete storage from ViPR Controller see either the *ViPR Controller User Interface Virtual Data Center Configuration Guide* or the *ViPR Controller REST API Virtual Data Center Configuration Guide*.

-
- The remote host, SMI-S provider (Solutions Enabler (SYMAPI Server) and EMC CIM Server (ECOM)) are configured to accept SSL connections.
 - The EMC storsrvd daemon is installed and running.
 - The SYMAPI Server and the ViPR Controller server hosts are configured in the local DNS server and that their names are resolvable by each other, for proper communication between the two. If DNS is not used in the environment, be sure to use the hosts files for name resolution (`/etc/hosts` or `c:/Windows/System32/drivers/etc/hosts`).
 - The EMC CIM Server (ECOM) default user login, password expiration option is set to "Password never expires."
 - The SMI-S provider host is able to see the gatekeepers (six minimum).

SMI-S provider version requirements for SRDF

When using SRDF in ViPR Controller, you must use specific SMI-S provider versions for data replication between the source and target sites.

When using SMI-S provider 8.0.3, you must enable the following properties regardless if you are running SRDF operations:

- SYMAPI_USE_GNS, SYMAPI_USE_RDFD under /var/symapi/config/options
- GNS_REMOTE_MIRROR under /var/symapi/config/daemon_options

VMAX to VMAX replication

SMI-S provider 4.6.2 or SMI-S provider 8.0.3 is required on both the source and the target sites when replicating data from a VMAX array to another VMAX array. You cannot have SMI-S provider 4.6.2 on one site and SMI-S provider 8.0.3 on the other site. They must be the same version on both sites.

VMAX3 to VMAX3 replication

SMI-S provider 8.0.3 is required on both the source and the target sites when replicating data from a VMAX3 array to another VMAX3 array.

VMAX to VMAX3 and VMAX3 to VMAX replication

SMI-S provider 8.0.3 is required on both the source and the target sites when replicating data from a VMAX3 array to a VMAX array or when replicating data from a VMAX3 array to a VMAX array.

Upgrade requirements

Review the following table to understand the upgrade requirements for VMAX storage systems.

Note

The [ViPR Controller Support Matrix](#) has the most recent version requirements for all systems supported, or required by ViPR Controller. For specific version requirements of the SMI-S provider review the [ViPR Controller Support Matrix](#) before taking any action to upgrade or install the SMI-S provider for use with ViPR Controller.

ViPR Controller 2.3 requires SMI-S Provider 4.6.2 for all VNX storage systems. Plan accordingly if you are using both VMAX and VNX storage systems in your environment.

When upgrading, you must upgrade the ViPR Controller to version 2.3 before you upgrade the SMI-S provider to 8.0.3.

To upgrade SMI-S Provider 8.0.3, you must contact EMC Customer Support.

Table 4 Upgrade requirements for VMAX storage systems

Upgrade from:		To:	
ViPR Controller	SMI-S provider	ViPR Controller	SMI-S provider
2.x	4.6.2	2.3	8.0.3
2.2	8.0.1	2.3	8.0.3

VMAX storage system

You prepare the VMAX storage system before adding it to ViPR Controller as follows.

- Create a sufficient amount of storage pools for storage provisioning with ViPR Controller (for example, SSD, SAS, NL-SAS).
- Define FAST policies.
Storage Tier and FAST Policy names must be consistent across all VMAX storage systems.

- It is not required to create any LUNs, storage groups, port groups, initiator groups, or masking views.
- After a VMAX storage system has been added and discovered in ViPR Controller, the storage system must be rediscovered if administrative changes are made on the storage system using the storage system element manager.
- For configuration requirements when working with meta volumes see *ViPR Controller Integration with VMAX and VNX Storage Systems User and Administration Guide*.

EMC VNX for Block

ViPR Controller management of VNX for Block systems is performed through the EMC SMI-S provider. Your SMI-S provider, and VNX for Block storage system must be configured as follows before the storage system is added to ViPR Controller.

SMI-S provider configuration requirements for VNX for Block

You need specific information to validate that the SMI-S provider is correctly configured for ViPR Controller and to add storage systems to ViPR Controller.

Gather the required information

- SMI-S provider host address
- SMI-S provider credentials (default is admin/#1Password)
- SMI-S provider port (default is 5989)

Configuration requirements

SMI-S provider cannot be shared between ViPR Controller and any other application requiring an SMI-S provider to VNX for Block, such as EMC ViPR SRM.

Before adding VNX for Block storage to ViPR Controller, login to your SMI-S provider to ensure the following configuration requirements are met:

- The host server running Solutions Enabler (SYMAPI Server) and SMI-S provider (ECOM) differs from the server where the VNX for Block storage processors are running.
- The storage system is discovered in the SMI-S provider.
- When the storage provider is added to ViPR Controller, all the storage systems managed by the storage provider will be added to ViPR Controller. If you do not want all the storage systems on an SMI-S provider to be managed by ViPR Controller, configure the SMI-S provider to only manage the storage systems that will be added to ViPR Controller, before adding the SMI-S provider to ViPR Controller.

Note

Storage systems that will not be used in ViPR Controller, can also be deregistered, or deleted after the storage provider is added to ViPR Controller. For steps to deregister or delete storage from ViPR Controller see either the *ViPR Controller User Interface Virtual Data Center Configuration Guide* or the *ViPR Controller RESTAPI Virtual Data Center Configuration Guide*.

- The remote host, SMI-S provider (Solutions Enabler (SYMAPI Server) and EMC CIM Server (ECOM)) are configured to accept SSL connections.
- The EMC storsrvd daemon is installed and running.
- The SYMAPI Server and the ViPR Controller server hosts are configured in the local DNS server and that their names are resolvable by each other, for proper

communication between the two. If DNS is not used in the environment, be sure to use the hosts files for name resolution (`/etc/hosts` or `c:\Windows\System32/drivers/etc/hosts`).

- The EMC CIM Server (ECOM) default user login, password expiration option is set to "Password never expires."
- The SMI-S provider host needs IP connectivity over the IP network with connections to both VNX for Block storage processors.

VNX for Block storage system

You prepare the VNX for Block storage system before adding it to ViPR Controller as follows.

Configuration requirements

- Create a sufficient amount of storage pools or RAID groups for storage provisioning with ViPR Controller.
- If volume full copies are required, install SAN Copy enabler software on the storage system.
- If volume continuous-native copies are required, create clone private LUNs on the array.
- Fibre Channel networks for VNX for Block storage systems require an SP-A and SP-B port pair in each network, otherwise virtual pools cannot be created for the VNX for Block storage system.
- For configuration requirements when working with meta volumes see *ViPR Controller Integration with VMAX and VNX Storage Systems User and Administration Guide*.

Configuration requirements for ViPR Controller to collect HDS port metrics

You must enable performance data logging in EMC Unisphere so that VNX for Block sends the required metrics to ViPR Controller before you can set up metrics-based port selection for VNX for Block . For steps to enable performance data logging in EMC Unisphere see: [Prerequisite configuration settings for VNX for Block on page 18](#).

Enable performance data logging for VNX for Block

You must enable performance data logging in EMC Unisphere so that VNX for Block sends the required metrics to ViPR Controller before you can set up metrics-based port selection for VNX for Block .

Procedure

1. Log into the EMC Unisphere.
2. Select **System** > **Statistics for Block**. Statistics for Block can be found in the **Monitoring and Alerts** section.
3. Click **Performance Data Logging**.
The **Data Logging** dialog is displayed.
4. If data logging is not already started:

Note

Do not change the default times for **Real Time Interval** and **Archive Interval**.

- a. Click **Start** to start data logging.
- b. Click **Apply**.

5. Click **OK**.

VPLEX

Before adding VPLEX storage systems to ViPR Controller, validate that the VPLEX environment is configured as follows:

- ViPR supports VPLEX in a Local or Metro configuration. VPLEX Geo configurations are not supported.
- Configure VPLEX metadata back-end storage.
- Create VPLEX logging back-end storage.
- Verify that the:
 - Storage systems to be used are connected to the networks containing the VPLEX back-end ports.
 - Hosts to be used have initiators in the networks containing the VPLEX front-end ports.
- Verify that logging volumes are configured to support distributed volumes in a VPLEX Metro configuration.
- It is not necessary to preconfigure zones between the VPLEX and storage systems, or between hosts and the VPLEX , except for those necessary to make the metadata backing storage and logging backing storage available.

Third-party block storage (OpenStack)

ViPR Controller uses the OpenStack Block Storage (Cinder) service to manage OpenStack supported block storage systems. Your OpenStack block storage systems must meet the following installation and configuration requirements before the storage systems in OpenStack, and the storage system's resources can be managed by ViPR Controller.

Third-party block storage provider installation requirements

Use the OpenStack Block Storage (Cinder) Service to add third-party block storage systems into ViPR Controller.

Supported Openstack installation platforms

Openstack installation is supported on the following platforms:

- Red Hat Enterprise Linux
- SUSE Enterprise Linux
- Ubuntu Linux

For a list of the supported platform versions, see Openstack documentation at: <http://docs.openstack.org>.

OpenStack installation requirements

You must install the following two components on either the same server or separate servers:

- OpenStack Identity Service (Keystone)
Required for authentication
- OpenStack Block Storage (Cinder)
The core service that provides all storage information.

For complete installation and configuration details, refer to the OpenStack documentation at: <http://docs.openstack.org>.

Third-party block storage system support

You configure third-party storage systems on the OpenStack Block Storage Controller node (Cinder service).

Supported third-party block storage systems

ViPR Controller operations are supported on any third-party block storage systems tested by OpenStack that use Fibre Channel or iSCSI protocols.

Non-supported storage systems

ViPR Controller does not support third-party storage systems using:

- Proprietary protocols, such as Ceph.
- Drivers for block over NFS.
- Local drivers such as LVM.

ViPR Controller does not support these OpenStack supported storage systems and drivers:

- LVM
- NetAppNFS
- NexentaNFS
- RBD (Ceph)
- RemoteFS
- Scality
- Sheepdog
- XenAPINFS

Refer to www.openstack.org for information about OpenStack third-party storage systems.

ViPR Controller native driver support and recommendations

ViPR Controller provides limited support for third-party block storage. To have full support of all ViPR Controller operations, it is recommended that you add and manage the following storage systems with ViPR Controller native drivers, and not use the OpenStack third-party block storage provider to add these storage systems to ViPR Controller:

- EMC VMAX
- EMC VNX for Block
- EMC VPLEX
- Hitachi Data Systems (with Fibre Channel only)

Add these storage systems directly in ViPR Controller using the storage system host address or the host address of the proprietary storage provider.

Supported ViPR Controller operations

You can perform ViPR Controller discovery and various service operations on third-party block storage systems.

You can perform these service operations on third-party block storage systems:

- Create Block Volume

- Export Volume to Host
- Create a Block Volume for a Host
- Expand block volume
- Remove Volume by Host
- Remove Block Volume
- Create Full Copy
- Create Block Snapshot
- Create volume from snapshot
- Remove Block Snapshot

Note

You cannot use the ViPR Controller **Create VMware Datastore** service to create a datastore from a block volume created by a third-party storage system. However, you can manually create datastores from third-party block volumes through the VMware vCenter.

OpenStack configuration

After installing the Keystone and Cinder services, you must modify the Cinder configuration file to include the storage systems to be managed by ViPR Controller. After modifying the configuration file, you must create the volume types to map to the backend drivers. These volume types are discovered as storage pools of a certain storage system in ViPR Controller.

OpenStack third-party storage configuration recommendations

Before adding the storage provider to ViPR Controller, configure the storage provider with only the storage systems to be managed by ViPR Controller through a third-party block storage provider. When the third-party block storage provider is added to the ViPR Controller Physical Assets, all storage systems managed by the OpenStack block storage service and supported by ViPR Controller are added to ViPR Controller.

If you included storage systems in the storage provider that will not be managed by ViPR Controller, you can deregister or delete those storage systems from ViPR Controller. However, it is a better practice to configure the storage provider for ViPR Controller integration before adding the storage provider to ViPR Controller.

Cinder service configuration requirements

You must add an entry in the Cinder configuration file for each storage system to be managed by ViPR Controller. This file is located in `/etc/cinder/cinder.conf`.

Cinder does not have any specific standards on backend driver attribute definitions. See the vendor-specific recommendations on how to configure the cinder driver, which may involve installing a vendor specific plugin or command line interface.

Storage system (backend) configuration settings

Cinder defines backend configurations in individual sections of the `cinder.conf` file to manage the storage systems. Each section is specific to a storage system type.

Edit the `cinder.conf` file:

Procedure

1. Uncomment `enabled_backends`, which is commented by default, and add the multiple backend names. In the following example, NetApp and IBM SVC are added as backend configurations.

```
enabled_backends=netapp-iscsi,ibm-svc-fc
```

2. Near the end of the file, add the storage system specific entries as follows:

```
[netapp-iscsi]
#NetApp array configuration goes here

[ibm-svc-fc]
#IBM SVC array configuration goes here
```

3. Restart the Cinder service.

```
#service openstack-cinder-volume restart
```

Volume setup requirements for OpenStack

ViPR Controller has specific setup requirements for volumes created in OpenStack.

ViPR Controller requires that you do the following for each volume in OpenStack:

- Map the volume to the backend driver.
- Indicate if the volume is set for thin or thick provisioning.

ViPR Controller-specific properties

You can create volume types through Cinder CLI commands or the Dashboard (OpenStack UI). The properties required by ViPR Controller in the Cinder CLI for the volume type are:

- `volume_backend_name`
- `vipr:is_thick_pool=true`

volume_backend_name

The following example demonstrates the Cinder CLI commands to create volume types (NetApp, IBM SVC) and map them to the backend driver.

```
cinder --os-username admin --os-password <password> --os-tenant-name
admin type-create "NetAPP-iSCSI"
cinder --os-username admin --os-password <password> --os-tenant-name
admin type-key "NetAPP-iSCSI" set volume_backend_name=NetAppISCSI

cinder --os-username admin --os-password <password> --os-tenant-name
admin type-create "IBM-SVC-FC"
cinder --os-username admin --os-password <password> --os-tenant-name
admin type-key "IBM-SVC-FC" set volume_backend_name=IBMSVC-FC

cinder --os-username admin --os-password <password> --os-tenant-name
admin extra-specs-list
```

vipr:is_thick_pool=true

By default, during discovery, ViPR Controller sets the provisioning type of OpenStack volumes to thin. If the provisioning type is thick, you must set the ViPR Controller-specific property for thick provisioning to `true` for the volume type. If the provisioning type of the volume is thin, you do not need to set the provisioning type for the volume in OpenStack.

The following example demonstrates the Cinder CLI commands to create a volume type (NetApp), and define the provisioning type of the volume as thick.

```
cinder --os-username admin --os-password <password> --os-tenant-name
admin --os-auth-url=http://<hostname>:35357/v2.0 type-create "NetAPP-
iSCSI"
cinder --os-username admin --os-password <password> --os-tenant-name
admin --os-auth-url=http://<hostname>:35357/v2.0 type-key "NetAPP-
iSCSI" set volume_backend_name=NetAppISCSI
cinder --os-username admin --os-password <password> --os-tenant-name
admin --os-auth-url=http://<hostname>:35357/v2.0 type-key "NetAPP-
iSCSI" set vipr:is_thick_pool=true

cinder --os-username admin --os-password <password> --os-tenant-name
admin --os-auth-url=http://<hostname>:35357/v2.0 extra-specs-list
```

Validate setup

Validate that OpenStack has been configured correctly to create volumes for each of the added storage systems.

1. Create volumes for each type of volume created in OpenStack.
Volumes can be created in the OpenStack UI or the Cinder CLI. The Cinder CLI command to create a volume is:

```
cinder --os-username admin --os-tenant-name admin --display-name
<volume-name> --volume-type <volume-type-id> <size>
```

2. Check that the volumes are getting created on the associated storage system. For example, NetApp-iSCSI type volumes should be created only on the NetApp storage system.

ViPR Controller configuration for third-party block storage ports

The OpenStack API does not provide the storage port World Wide Port Name (WWPN) for Fibre Channel connected storage systems and the IQN for iSCSI connected storage systems. As a result, ViPR Controller cannot retrieve the storage port WWPNs or IQNs during discovery.

After ViPR Controller discovers a third-party block storage array, a default storage port is created for the storage system, and appears in the **Storage Port** page with the name **Default** and the storage port identifier of **Openstack+<storagesystemserialnumber>+Port+Default**.

Fibre Channel configured storage ports

ViPR Controller export operations cannot be performed on an FC connected storage system, which was added to ViPR Controller without any WWPNs assigned to the storage port. Therefore, ViPR Controller system administrators must manually add at least one WWPN to the default storage port before performing any export operations on the storage system. WWPNs can be added to ViPR Controller through the ViPR Controller CLI and UI.

After the WWPN is added to the storage port, you can perform export operations on the storage system from ViPR Controller. At the time of the export, ViPR Controller reads the export response from the Cinder service. The export response will include the WWPN, which was manually added by the system administrator from the ViPR Controller CLI, and any additional WWPNs listed in the export response. ViPR Controller then creates a storage port for each of the WWPNs listed in the export response during the export operation.

After a successful export operation is performed, the **Storage Port** page displays any newly created ports and the Default storage port.

Each time another export operation is performed on the same storage system, ViPR Controller reads the Cinder export response. If the export response presents WWPNs, which are not present in ViPR Controller, then ViPR Controller creates new storage ports for every new WWPN.

For steps to add WWPNs to FC connected third-party block storage ports, see the *ViPR Controller CLI Reference Guide*, which is available from the [ViPR Controller Product Documentation Index](#).

iSCSI configured storage ports

The default storage port is used to support the storage system configuration until an export is performed on the storage system. At the time of the export, ViPR Controller reads the export response from the Cinder service, which includes the iSCSI IQN. ViPR Controller then modifies the default storage port's identifier with the IQN received from the Cinder export response.

Each time another export operation is performed on the same storage system, ViPR Controller reads the Cinder export response. If the export response presents an IQN, which is not present in ViPR Controller, then ViPR Controller creates a new storage port.

Stand-alone EMC ScaleIO

Your stand-alone ScaleIO should meet the following system requirements and be configured as follows before adding the storage to ViPR Controller.

- Protection domains are defined.
- All storage pools are defined.

EMC XtremIO

Before adding XtremIO storage to ViPR, ensure that there is physical connectivity between hosts, fabrics and the array.

EMC VNXe

Before you add VNXe storage systems to ViPR Controller review the following information.

Table 5 Gather the EMC Unisphere required information

Setting	Value
Unisphere host IP address	
Unisphere user credentials	
Unisphere host port (default is 443)	

Create a sufficient amount of storage pools for storage provisioning with ViPR Controller.

General file storage system requirements

When configuring your file storage system for ViPR Controller all file storage should be configured as follows.

- The NFS server must be configured in the storage system for ViPR Controller to create the NFS exports.
- The CIFS share must be configured in the storage system for ViPR Controller to create the CIFS shares.
- NETBIOS configuration needs to be done to access the share with NETBIOS name, if NETBIOS is present the mount point will be constructed with NETBIOS name instead of IP.

For EMC Isilon, NetApp 7-Mode, and NetApp Cluster-Mode

To add domain users to a share, the AD server should be configured on the storage system (CIFS server).

EMC® Data Domain®

Before adding Data Domain storage to ViPR Controller, configure the storage as follows.

- The Data Domain file system (DDFS) is configured on the Data Domain system.
- Data Domain Management Center (DDMC) is installed and configured.
- Network connectivity is configured between the Data Domain system and DDMC.

While adding Data Domain storage to ViPR Controller, it is helpful to know that:

- A Data Domain Mtree is represented as a file system in ViPR Controller.
- Storage pools are not a feature of Data Domain. However, ViPR Controller uses storage pools to model storage system capacity. Therefore, ViPR Controller creates one storage pool for each Data Domain storage system registered to ViPR Controller, for example, if three Data Domain storage systems were registered to ViPR Controller, there would be three separate Data Domain storage pools. One storage pool for each registered Data Domain storage system.

EMC Isilon

Before adding EMC Isilon storage to ViPR Controller, configure the storage as follows.

Gather the required information

The following information is needed to configure the storage and add it to ViPR Controller.

Table 6 Isilon required information

Setting	Value
IPv4 or IPv6 address	
Port (default is 8080)	
Credentials for the user to add Isilon storage systems to ViPR Controller, and to own all of the file systems created on the Isilon through ViPR Controller.	

Table 6 Isilon required information (continued)

Setting	Value
This user must have root or administrator privileges to the Isilon.	

Configuration requirements

- SmartConnect is licensed and configured as described in Isilon documentation. Be sure to verify that:
 - The names for SmartConnect zones are set to the appropriate delegated domain.
 - DNS is in use for ViPR Controller and provisioned hosts are delegating requests for SmartConnect zones to SmartConnect IP.
- SmartQuota must be licensed and enabled.
- There is a minimum of 3 nodes in the Isilon cluster configured.
- Isilon clusters and zones will be reachable from ViPR Controller Controller VMs.
- When adding an Isilon storage system to ViPR Controller, you will need to use either the root user credentials, or create an account for ViPR Controller users, that has administrative privileges to the Isilon storage system.
- The Isilon user is independent of the currently logged in ViPR Controller user. All ViPR Controller operations performed on the Isilon storage system, are executed as the Isilon user that is entered when the Isilon storage system is added to ViPR Controller.

NetApp 7-mode

Before adding NetApp 7-mode storage to ViPR Controller, configure the storage as follows.

- ONTAP is in 7-mode configuration.
- Multistore is enabled (7 Mode with Multi store is only supported by ViPR Controller).
- Aggregates are created.
- NetApp licenses for NFS, CIFS, and snapshots are installed and configured.
- vFilers are created and necessary interfaces/ports associated with them
- Setup NFS/CIFS on the vFilers

NetApp Cluster-mode

Before adding NetApp Cluster-mode storage to ViPR Controller, configure the storage as follows.

- ONTAP is in Cluster-mode configuration
- Aggregates are created.
- NetApp licenses for NFS, CIFS and snapshots are installed and configured.
- Storage Virtual Machines(SVMs) are created and the necessary interfaces and ports are associated with them.
- Setup NFS, and CIFS on SVMs.

- When discovering NetApp Cluster-mode storage systems with ViPR Controller, you must use the management IP . You cannot discover NetApp Cluster-mode storage systems using LIF IP.

EMC VNX for File

Before adding VNX for File storage to ViPR Controller, you must verify that VNX for File meets specific requirements.

These are the requirements for adding VNX for File storage to ViPR Controller:

- Storage pools for VNX for File are created.
- Control Stations are operational and reachable from ViPR Controller Controller VMs.
- VNX SnapSure is installed, configured, and licensed.

RecoverPoint systems

ViPR Controller System Administrators can review the ViPR Controller requirements and the information necessary to add and configure an EMC RecoverPoint system to ViPR Controller.

RecoverPoint site information

Review the information that is required to add a RecoverPoint system to ViPR Controller:

- RecoverPoint site management IPv4 or IPv6 address or hostname
- Port
- Credentials for an account that has the RecoverPoint administrator role to access the RecoverPoint site

Configuration requirements

To add a RecoverPoint system to ViPR Controller, do the following:

- Install and license the RecoverPoint systems.
- If ViPR Controller is not managing the SAN network, zone the RecoverPoint systems to the storage arrays and attach the RecoverPoint splitters.
- Establish IP connectivity between RecoverPoint and the ViPR Controller virtual appliance.

Fabric Managers (switches)

ViPR Controller supports Brocade and Cisco switches. System Administrators can review the ViPR Controller requirements and information necessary to add and configure Fabric Managers (switches) to ViPR Controller.

ViPR Controller switch and network terminology

ViPR Controller has three external interfaces - a REST API, a Command Line Interface (CLI), and a User Interface. The scripting and programming interfaces (API and CLI) use slightly different terminology to refer to the network elements of your Storage Area Network (SAN).

The terminology used in each ViPR Controller Interface is as follows:

- Cisco[®] switches and Brocade CMCNE management stations are called *fabric managers* in the user interface.

- In the ViPR REST API and the ViPR CLI, Brocade CMCNE management stations and Cisco switches are called *network-systems*. Cisco VSANs and Brocade fabrics are called *networks*.

Brocade

Your Brocade switch should meet the following system requirements and be configured as follows before adding the switch to ViPR Controller.

Software requirements

ViPR Controller requires that EMC Connectrix Manager Converged Network Edition (CMCNE) is installed. For supported versions, see the *EMC ViPR Controller Support Matrix* on the EMC Community Network (community.emc.com).

CMCNE can be downloaded from EMC Support Zone (<https://support.emc.com>).

CMCNE installation and configuration requirements

- CMCNE must be installed on a different host than the storage system SMI-S provider. CMCNE uses the same port as the SMI-S provider, and therefore causes a port conflict if CMCNE and the SMI-S provider are installed on the same machine.
- The SMI Agent (SMIA) must be installed and started as part of the CMCNE installation. [Validating CMCNE SMI Agent Installation on page 28](#) provides instructions to validate that the SMI Agent was installed correctly with CMCNE.
- CMCNE must have access to the switches with administrator privileges and the account must be configured with privileges to discover SAN topology, and to activate, create, and delete zones and zonesets..
- Use the CMCNE UI to prediscover the fabrics.
- You must have created the necessary fabrics that will be used by ViPR, assigned the ports to the fabrics, and configured any ISL links needed to connect a fabric between multiple switches.
- It is highly recommended that you implement redundant connections among ISL connections, so that the failure of a single ISL link does not partition fabrics.
- ViPR Controller has no restrictions on the number of fabrics.
- You do not need to create zones.

Validating CMCNE SMI Agent installation

The SMI Agent (SMIA) must have been installed and started as part of the CMCNE installation.

To validate that the SMI Agent was installed with the CMCNE installation, enter the CMCNE IP address in a browser to go to the EMC Connectrix Manager Converged Network Edition start page.

- If the SMIA provider was installed with CMCNE, a link to start the CMCNE SMIA configuration tool appears in the start page as demonstrated in the following image.



CMCNE - [Web Start the CMCNE Client](#)

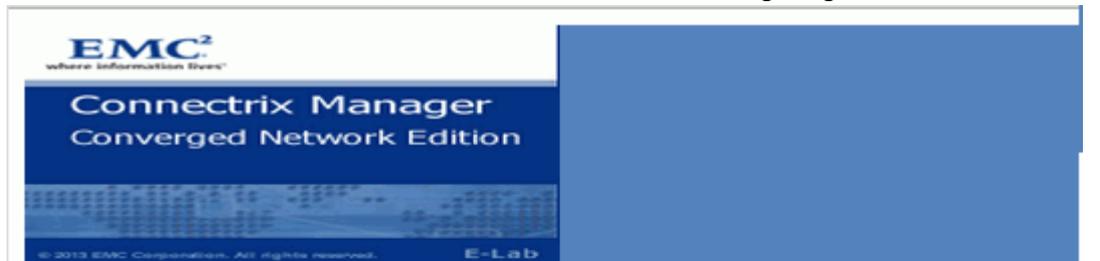
CMCNE SMIA - [Web Start the CMCNE SMIA Configuration Tool](#)

Download Java Runtime Environments (version 1.6.0_26)

- [Windows](#)

Download SNMP MIB Files

- If a link to start the CMCNE SMIA configuration tool does not appear in the start page if it was not installed with CMCNE, as demonstrated in the following image.



CMCNE - [Web Start the CMCNE Client](#)

Download Java Runtime Environments (version 1.7)

- [Windows 32-bit](#)

Download SNMP MIB Files

If the SMIA is installed, and not functioning as expected, check the C:\<installation path>\cimom\server\logs\smia-provider.log file for any errors.

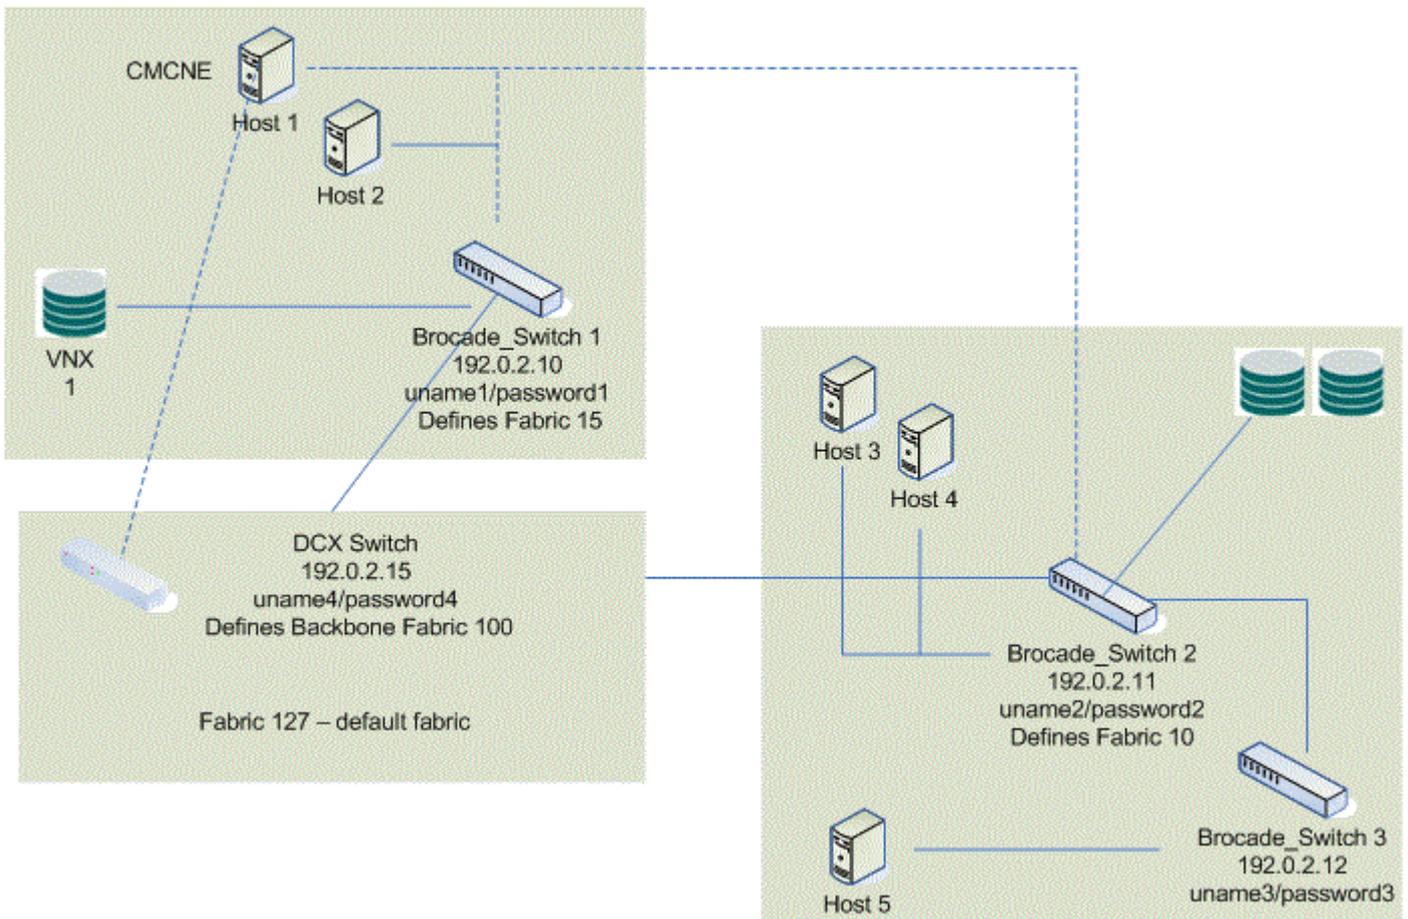
ViPR Controller support for fibre channel routing with Brocade switches

ViPR Controller includes support for Fibre Channel routing configurations using Brocade switches. Fibre Channel Routing (FCR) is designed to allow devices from separate fabrics to communicate without merging the fabrics.

Fibre Channel Routing allows separate autonomous fabrics to maintain their own fabric-wide services and provide the following benefits:

- Faults in fabric services, fabric reconfigurations, zoning errors, and misbehaving devices do not propagate between fabrics.
- Devices that cause performance issues in a fabric do not propagate those issues to other fabrics.

This support expands ViPR Controller network and switch support to include configurations that use a Brocade router to move traffic from one physical data center to another. The following figure shows an example of a ViPR Controller-supported datacenter configuration that uses Fibre Channel routing.



In this example, the DCX switch acts as the router between the devices in Fabric 15 and the devices in Fabric 10. In this simple example, Host 5 in Fabric 10 can consume storage on VNX 1 in Fabric 15.

Brocade Fibre Channel routing configuration requirements

The following guidelines apply to Brocade fibre channel routing configurations.

- As a part of the routing setup, at least one LSAN zone must be created between each pair of fabrics you expect ViPR Controller to consider. When ViPR Controller discovers the topology, it assumes that if an LSAN zone exists between two fabrics, then routing between the fabrics is allowed. If there is no LSAN zone between the fabrics, ViPR Controller assumes that routing is not allowed between the fabrics.
- CMCNE must discover the router and a seed switch in each participating fabric. In this example, CMCNE running on Host 1 would have to discover Brocade_Switch1, the DCX Switch, and Brocade Switch 2.
- ViPR must successfully discover CMCNE. Choose **Physical Assets > Fabric Managers** to add and discover CMCNE.
- Refer to the *EMC® Connectrix Manager Converged Network Edition User Guide* for more information on installing and configuring CMCNE
- Refer to the *CMCNE Release Notes* for a list of supported switch platforms and minimum firmware revisions.

Cisco

You will need the following information and your Cisco switches must be configured as follows before the switch can be added and used in ViPR Controller.

Cisco switch login credentials

Obtain the login credentials to use when adding the Cisco switch to ViPR Controller. The account must have privileges to provision (configure the mode of zone, zonesets, and VSAN) and to show the FCNS database.

Configuration requirements

Configure the switch as follows before adding it to ViPR Controller:

- Enable SSH.
- Create the VSANs that will be used by ViPR Controller on the Cisco switch.
- Assign the appropriate interfaces to the VSAN database.
- Configure any ISL links, and port channels, for multiple switch networks, and validate that the ISL links are operational. Doing so ensures that the FCNS database is correctly distributed and current on any switches that you add to ViPR Controller.

Configuration recommendations

It is highly recommended, but not required that you:

- Implement redundant connections between all ISL connections, so that the failure of a single ISL link does not partition VSANs.
- Enable Enhanced Zoning to avoid simultaneous modifications of the zoning database by ViPR and other applications from overwriting each other's changes. If Enhanced Zoning cannot be enabled (not a recommended configuration), it is advisable to limit the number of fabric managers in ViPR Controller to one fabric manager per VSAN.

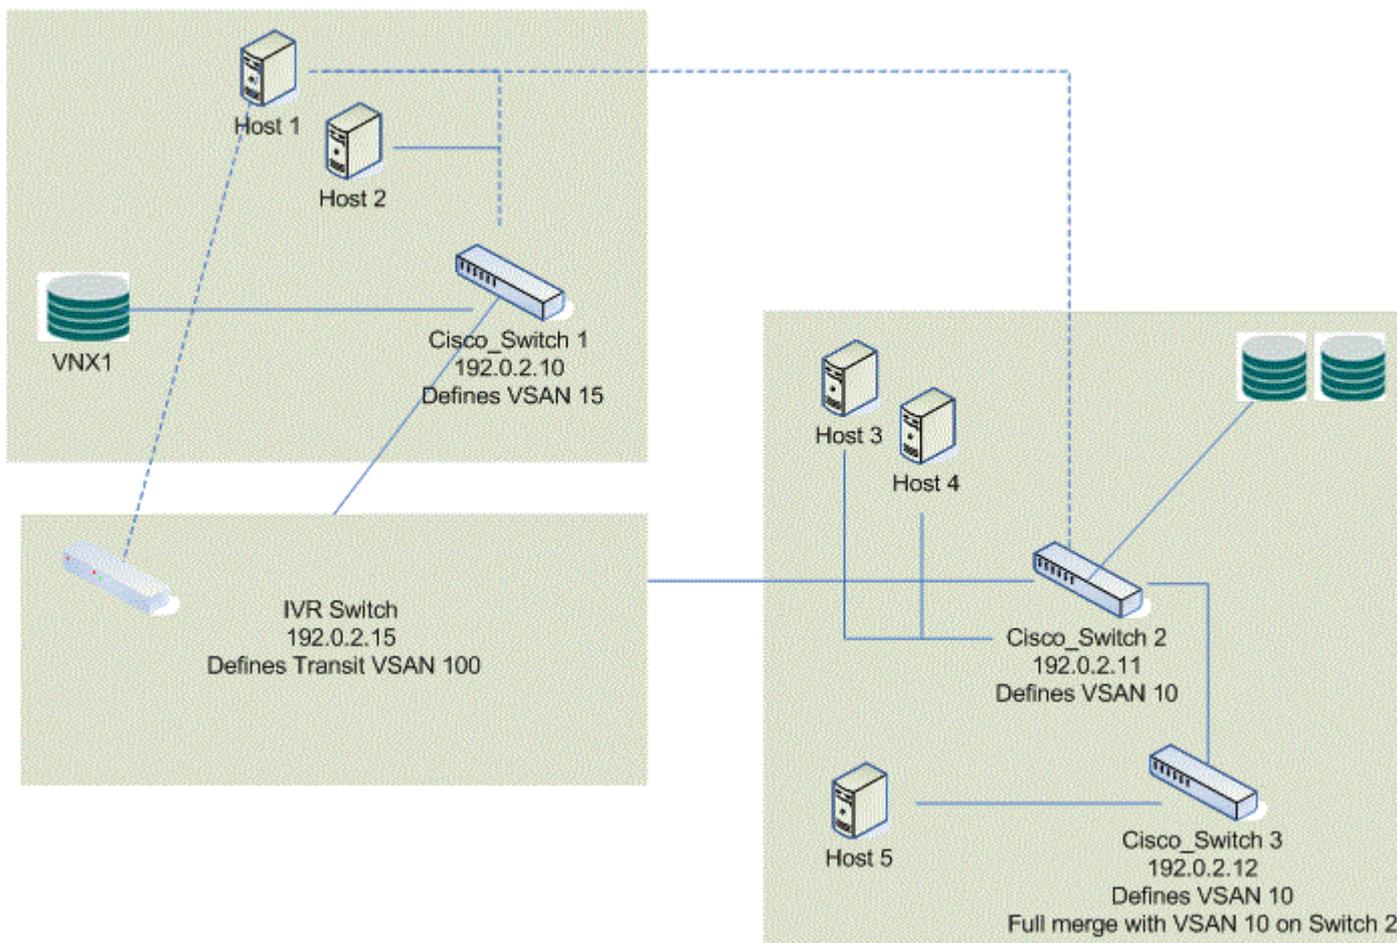
ViPR Controller support for Inter-VSAN Routing with Cisco switches

ViPR Controller includes support for Inter-VSAN Routing (IVR) configurations using Cisco switches. IVR is designed to allow hosts, storage arrays and other devices residing in separate VSANs to communicate without merging the VSANs.

IVR allows separate autonomous VSANs to maintain their own VSAN-wide services and provide the following benefits:

- Accesses resources across VSANs without compromising other VSAN benefits.
- Transports data traffic between specific initiators and targets on different VSANs without merging VSANs into a single logical fabric.
- Establishes proper interconnected routes that traverse one or more VSANs across multiple switches. IVR is not limited to VSANs present on a common switch.
- Shares valuable resources across VSANs without compromise. Fibre Channel traffic does not flow between VSANs, nor can initiators access resources across VSANs other than the designated VSAN.
- Provides efficient business continuity or disaster recovery solutions when used in conjunction with FCIP.
- IVR Is in compliance with Fibre Channel standards.

The following figure shows an example of a ViPR Controller-supported datacenter configuration that uses Inter-VSAN routing. In this example, the IVR switch acts as the router between the devices in VSAN 15 and the devices in VSAN 10. Host 5 in VSAN 10 can consume storage on VNX 1 in VSAN 15.



Isolated VSANS

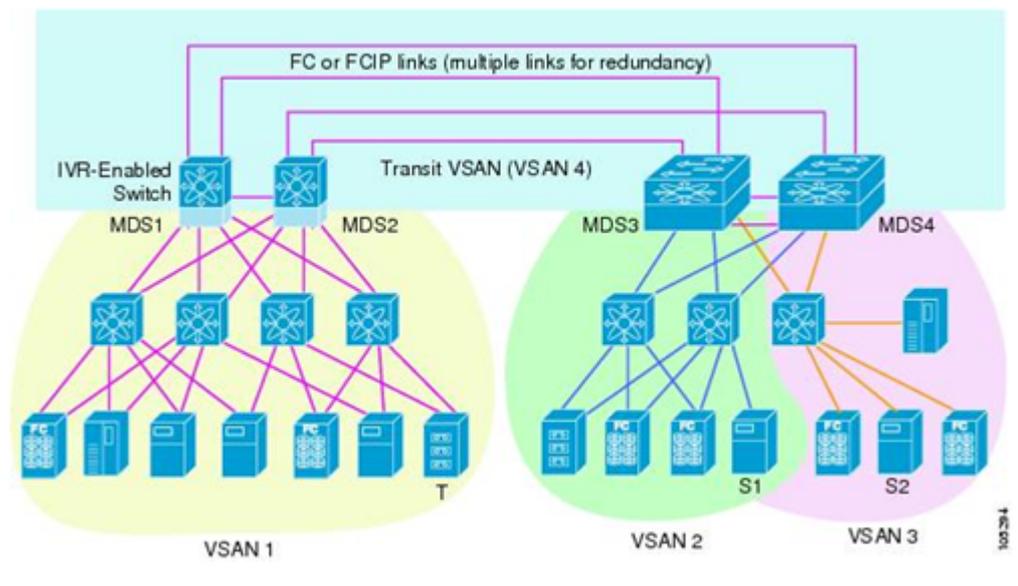
ViPR Controller can recognize and create IVR zones between isolated VSANs, whether they exist on a single switch or span multiple physical switches.

In order to accomplish and complete IVR zoning, ViPR Controller requires that the proper transit VSANs exist between the IVR routed VSANs. Additionally, ViPR Controller requires at least one IVR zone be created between each of the IVR routed VSANs. This allows ViPR Controller to associate and consider the IVR path as available and as an active provisioning path across VSANs.

When ViPR Controller discovers the topology, it assumes that when an IVR zone exists between the two VSANs, then routing between the VSANs is allowed and configured properly to allow host to storage access.

If there is not a pre-instantiated IVR zone between the IVR routed VSANs, ViPR Controller assumes routing is not allowed between the VSANs and does not create IVR zones.

For example, ViPR Controller can support this configuration.



(Graphic from *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide*)

Cisco Inter-VSAN Routing configuration requirements

The following guidelines apply to all fibre channel routing configurations.

- For each VSAN, ViPR Controller must successfully discover the Fabric Manager (Cisco switch).
- As a part of the routing setup, at least one IVR zone must be created between each pair of VSANs you expect ViPR Controller to consider. When ViPR Controller discovers the topology, it assumes that if an IVR zone exists between two VSANs, then routing between the VSANs is allowed. If there is no IVR zone between the VSANs, ViPR Controller assumes that routing is not allowed between the VSANs.
- ViPR Controller must also discover the IVR-enabled switches that enable routing between ViPR Controller managed VSANs in your data center.

Refer to the *Cisco MDS 9000 Family NX-OS Inter-VSAN Routing Configuration Guide* (http://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/5_0/configuration/guides/ivr/nxos/nxos_ivr.html) for more information on installing and configuring Cisco switches with IVR enabled, and creating IVR zones.

Vblock compute systems

ViPR Controller System Administrators can review the ViPR Controller requirements, and information necessary to add, and configure Vblock compute systems to ViPR Controller.

Cisco Unified Computing System Manager(UCSM)

- You will need the IP address, and user credentials with administrator privileges to the UCSM being used to manage the UCS.
- Service Profile Templates must be configured in UCSM for ViPR Controller to apply to the compute virtual pools when a cluster is created by ViPR Controller. Discuss which service profile templates should be used to provision with your UCS administrator. For the ViPR Controller configuration requirements for UCS service profile templates, see [ViPR Controller requirements for Service Profile Templates on page 48](#).

Private OS Install Network VLAN ID

You will need the VLAN ID for the OS Install Network. The OS Install Network is the second network provided when the compute image server was deployed.

The OS Install Network is a private VLAN for operating system (OS) installation. The OS Install Network is used by ViPR Controller during provisioning for communication between the hosts and the ViPR Controller compute image server. Since ViPR Controller utilizes a PXE boot process, a DHCP server is used and must be isolated from the customer network. During provisioning, the compute blades communicate with the image server and the operating system installation is performed over the OS Install Network. Once the OS installation is complete for a given host, the OS Install Network is no longer used to communicate to that host.

Compute images

ViPR Controller System Administrators can review the ViPR Controller requirements, and information necessary to add, and configure Compute Images to ViPR Controller.

Compute images are the operating installation images that can be installed by the ViPR Controller onto Vblock compute systems.

The Compute image server must have been deployed and configured in ViPR Controller. The compute image server must have been installed and configured in ViPR Controller before the compute images can be added. For instructions to deploy the compute image servers see: *ViPR Controller Installation, Upgrade, and Maintenance Guide*.

Hosts

ViPR Controller Tenant Administrators can review the ViPR Controller requirements, and information necessary to add, and configure hosts in ViPR Controller.

There are two ways to add hosts to ViPR Controller:

- Discoverable - once you have added the host, the ViPR Controller automatically discovers the host, and host initiators, and Windows clusters, and registers them to ViPR Controller. Only AIX[®], AIX VIO, Linux[®], and Windows hosts can be added as discoverable. Only Windows clusters can be automatically discovered in ViPR Controller.
- Undiscoverable - ViPR Controller does not discover, or register the host or host initiators. Any host that is not an AIX, AIX VIO, Linux, and Windows is added to ViPR Controller as undiscoverable. Optionally, AIX, AIX VIO, Linux, and Windows can be

added as undiscoverable as well. When an undiscoverable host has been added to ViPR Controller, you must manually add, and register the host initiators before using the host in a service operation.

AIX hosts and AIX VIO Servers

AIX hosts must be configured as follows to be discovered by ViPR Controller, and for ViPR Controller provisioning.

- Either EMC PowerPath , or AIX default MPIO (not both) must be enabled.
- EMC Inquiry (INQ) utility is required to be installed to match the volume World Wide Names (WWNs) to the host disks (hdisks).
- SSH must be installed and configured on the AIX hosts.

HP-UX

HP-UX hosts are added to ViPR Controller as undiscoverable. The HP-UX® operating system type option.

- Sets the Volume Set Addressing (VSA) flag, which is required for exporting EMC VMAX, and VPLEX volumes to HP-UX hosts.
- Is required to use the **Host Mode Option** when provisioning with HDS storage systems.
- When you add an HP-UX host to ViPR Controller, you will still need to manually add and register the host initiators in ViPR Controller.

Linux

Linux hosts must be configured as follows to be discovered by ViPR Controller, and used for ViPR Controller provisioning.

- SSH and LVM are enabled. ViPR Controller uses SSH for Linux hosts.
- EMC PowerPath or native Linux multipathing software is installed. Refer to the *EMC PowerPath for Linux Installation and Configuration Guide* and the *SUSE Linux Enterprise Server (SLES): Storage Administration Guide*, under "Configuring the System for Multipathing."
- Time synchronization is configured.
- In some cases, it may be necessary to install `lsb_release`. If host discovery fails due to compatibility, and logs indicate that the `lsb_release` command is not found, the package that includes that command must be installed.

Linux user requirements

When ViPR Controller storage is attached to a Linux host it needs to run commands on the host. To access the host, ViPR Controller uses the credentials entered at the time the host is added to ViPR Controller. These are usually the credentials for the root account. If you do not wish to give ViPR Controller root access to a Linux host, it is recommended to give the sudo user `ALL` privileges to run the commands ViPR Controller requires.

Windows

Windows hosts must be configured as follows to be discovered by ViPR Controller, and used for ViPR Controller provisioning.

Windows Remote Management (WinRM) must be enabled.

Refer to [Configuring WinRM on a Windows Host for ViPR on page 36](#).

Either EMC PowerPath, or Microsoft MPIO (not both) must be enabled.

For details see either the *EMC PowerPath and PowerPath/VE for Microsoft Windows Installation and Administration Guide* or the *Windows: Microsoft Multipath I/O Step-by-Step Guide*.

Time synchronization is configured.

If using LDAP or Active Directory domain account credentials, the domain user credentials must be in the same domain where the Windows host is located; otherwise the Windows host discovery will fail.

Configuring WinRM on a Windows host for ViPR Controller

Configures a Windows host to allow ViPR Controller to run commands on it.

Before you begin

- You must be logged in to the Windows host as administrator.
- For the ViPR Controller server to connect to Windows remote hosts, the host must accept remote Windows PowerShell commands. You can do this by enabling Windows remote access over HTTP.

Procedure

1. At an administrator command prompt on the Windows host, issue the following command:

```
winrm quickconfig
```

This starts up a listener on port 5985. The port on which you start the listener must be consistent with the port that you configure for the host in the host asset page.

2. You may need to make some configuration changes depending on how you want to connect to the host.
 - If you want ViPR Controller to connect to the host as a local user, you need to:
 - a. Check the winrm settings by running:

```
winrm get winrm/config/service
```

Basic Authentication and AllowUnencrypted must be set to true.
 - b. If basic authentication is not set to true, run:

```
winrm set winrm/config/service/auth @{Basic="true"}
```
 - c. If AllowUnencrypted is not set to true, run:

```
winrm set winrm/config/service @{AllowUnencrypted="true"}
```
 - d. Add the host to the ViPR Controller, **Physical Assets > Hosts** page.
 - If you want ViPR Controller to connect to the host as a domain user, you need to:
 - a. Ensure Kerberos is enabled. You can check using:

```
winrm get winrm/config/service
```
 - b. If you need to enable Kerberos, run:

```
winrm set winrm/config/service/auth @{Kerberos="true"}
```
 - c. A System Administrator must ensure that the domain has been configured as an authentication provider in ViPR Controller (**Security > Authentication Providers**).
 - d. A Tenant Administrator adds the host to ViPR Controller (**Physical Assets > Hosts**) page.

The credentials you supply for the host are of the form:

```
domain\username
```

3. Check that the host is displayed as valid in the table.

After you finish

After ViPR Controller is deployed, you can check that the host is displayed as valid in the **Physical Assets > Hosts** page. If you receive the following message WinRM may not be enabled or configured properly, or there may be a network problem.

```
Failed to connect to host. Please ensure the connection details are correct. [Error connecting: Connection refused]
```

VMware® vCenter

ViPR Controller System Administrators can review the ViPR Controller requirements, and information necessary to add, and configure vCenter to ViPR Controller.

vCenter role requirements

The role, which will be used by ViPR Controller to access vCenter, must have at least the following privileges enabled:

Datastore privileges:

- Allocate space
- Browse datastore
- Configure datastore
- Remove datastore

Host privileges:

- CIM
- CIM interaction
- Configuration

CHAPTER 4

Virtual Asset Requirements and Information

This chapter contains the following topics:

- [Overview](#)..... 40
- [Plan to build your virtual array](#)..... 40
- [Block storage configuration considerations](#)..... 43
- [File storage configuration considerations](#)..... 47
- [ViPR requirements for service profile templates](#)..... 48

Overview

ViPR Controller System Administrators can review the information to consider before configuring specific types of storage systems in ViPR Controller virtual arrays and virtual pools, and understand how ViPR Controller works with the storage system element managers once the volumes or file systems are under ViPR Controller management.

Plan to build your virtual array

Before you configure a virtual array:

- To learn about the components that make up a virtual array, and to see examples of different ways you can abstract the storage systems, and storage components in the virtual array, refer to the *ViPR Controller Concepts* article.
- To decide the type of SAN zoning to set on your virtual array refer to: [SAN zoning requirements on page 40](#).
- To decide the best method to use to set up the virtual array refer to: [Plan how to add the physical assets to the virtual array on page 40](#).
- If adding a Vblock system to ViPR Controller refer to: [Virtual array requirements for Vblock system services on page 42](#).

SAN zoning requirements

In the virtual array, you define whether ViPR Controller will automate SAN zoning at the time the storage is exported to the host from ViPR Controller, or if the SAN zoning will be handled manually outside of ViPR Controller operations. This discussion explains what to do when you chose manual SAN zoning.

If you chose manual SAN zoning:

- If there is an existing zone for the Host and Array:
After the ViPR provisioning operation completes, check the Port Group within the Masking View to identify the FA ports that ViPR selected for the provisioning request. Compare the FA ports in the zone to the FA ports in the Port Group. If they match, no further action is required. If they do not match, reconfigure the zone to use the same FA ports. Alternatively, a new zone can be created.
- If there is no existing zoning for the Host and Array:
After the ViPR provisioning operation completes, check the Port Group within the Masking View to identify the FA ports that ViPR selected for the provisioning request. Create a zone with the appropriate initiator and target ports.

Plan how to add the physical assets to the virtual array

At a minimum, a virtual array must include one network, and one storage system connected to the network.

When configuring the virtual array, you have the option to create the virtual array either by adding:

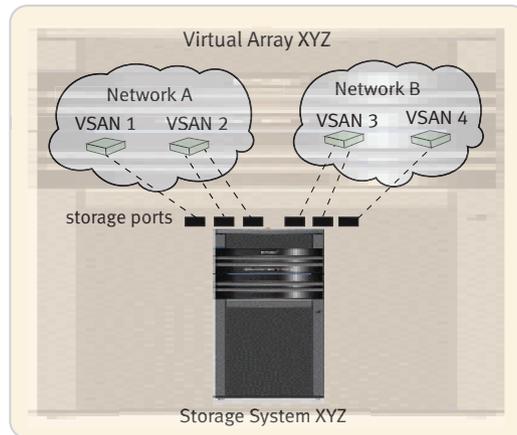
- Storage systems to the virtual array
- Storage ports to the virtual array

Add storage systems to the virtual array

You may want to add an entire storage system to the virtual array if you are planning to manage an entire storage system, or multiple storage systems in a single virtual array.

When you add an entire storage system to the virtual array, ViPR Controller automatically adds all of the registered networks, and storage ports associated with the storage system, to the virtual array. In the following example, when Storage System XYZ was added to the virtual array, all the storage ports, Networks A, Network B, VSAN 1, VSAN 2, VSAN 3, and VSAN 4 are all added to the virtual array.

Figure 1 Virtual array created by adding storage systems

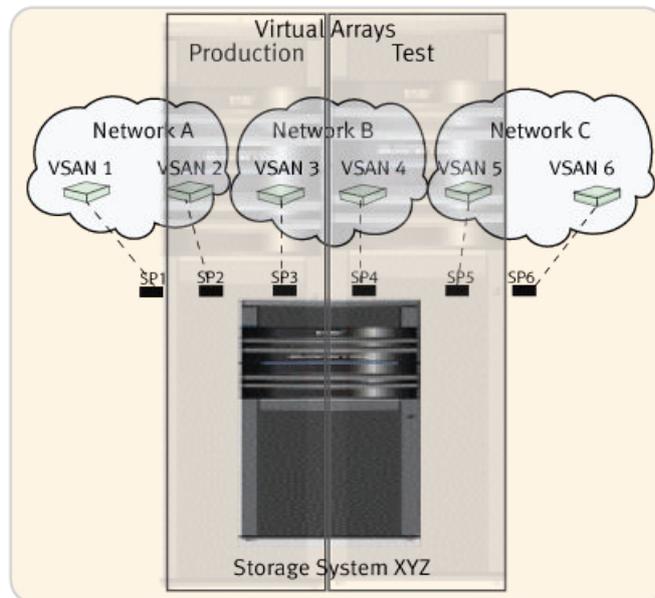


When you add an entire storage system to a virtual array, you will need to go back into the virtual array and remove any resources you don't want ViPR Controller to use.

For steps to create a virtual array by adding storage systems see, [Create a virtual array using storage system](#).

Add storage ports to the virtual array

If you want to partition a single storage system into multiple virtual arrays for example, allocate some of the storage system resources for testing and some for production, it maybe more useful to add the storage ports first to create the virtual array. If you choose to create the virtual array by first adding the storage ports, ViPR Controller will add only the networks, and storage systems associated to the storage ports, and you will start out with a more defined inventory in your virtual array. The following figure demonstrates how two virtual arrays were created from a single storage system by adding the ports first when creating the Production, and Test virtual arrays.

Figure 2 Virtual array created by adding ports

The Production virtual array was created by adding SP2, and SP3, which automatically adds Storage System XYZ, VSAN 2, VSAN 3, Network A, and Network B to the virtual array. While VSAN 1 is part of Network A, it is not added to the virtual array, because no storage ports, associated with VSAN 1 were selected to add to the virtual array.

The Test virtual array was created by adding SP4, and SP5, which automatically adds Storage System XYZ, VSAN 4, VSAN 5, Network B, and Network C to the virtual array. While VSAN 6 is part of Network C, it is not added to the virtual array, because no storage ports, associated with VSAN 6, were selected to add to the virtual array.

Furthermore, this image demonstrates how a network can be shared across two different virtual arrays. Since a storage port associated with Network B was added to each of the virtual arrays, Network B was added to both virtual arrays.

For steps to create a virtual array by adding storage ports see, [Create a virtual array using storage ports](#).

Virtual Array requirements for Vblock system services

For Vblock systems, storage must be accessible to compute systems through the virtual array. Vblock systems configured using the VCE logical build guide will have networks configured that connect the Cisco Unified Computing System™ (UCS) compute system to the storage via the SAN switches.

In ViPR Controller, virtual arrays should be created just as you would for any non-Vblock system. The networks that are defined in the virtual arrays will then determine whether the UCS systems have visibility to ViPR Controller storage.

The most effective thing is to do is discover all the Vblock system physical assets before defining virtual arrays. After discovering all components, consult with the UCS administrator to determine which networks (VSANs) will be used on a given Vblock system. Use those networks to define the ViPR Controller virtual arrays. On less complicated Vblock system configurations, for example, a single Vblock system, simply

adding the storage system to the virtual array may be enough. Once the virtual arrays are defined, they will be used by ViPR Controller for the following:

- ViPR Controller will automatically determine which UCS compute systems are available to compute virtual pools based on the selection of virtual arrays.
- ViPR Controller will automatically determine which blades to use to provision hosts based on the virtual arrays and compute virtual pools.

ViPR Controller makes these determinations by calculating which UCS compute systems have visibility to storage through the networks in the virtual arrays.

If working with updating service profile templates

When using updating service profile templates, it is recommended to create a dedicated virtual array that:

- Includes only the specific storage arrays that are intended to be used with the updating service profile template.
- Includes only the specific storage ports that are intended to be used with the updating service profile template.

Block storage configuration considerations

Before you create virtual arrays and virtual pools for block storage in ViPR Controller, review the following sections for storage system specific configuration requirements and recommendations:

- [Hitachi Data System \(HDS\) on page 43](#)
- [EMC VMAX on page 44](#)
- [EMC VMAX3 on page 45](#)
- [EMC VNX for Block on page 45](#)
- [EMC VNXe for Block on page 46](#)
- [EMC VPLEX on page 46](#)
- [Third-party block \(OpenStack\) storage systems on page 46](#)

Hitachi Data Systems

Review the following configuration requirements and recommendations before virtualizing your Hitachi Data Systems (HDS) in ViPR Controller.

Virtual pool considerations

ViPR Controller provides auto-tiering for the six standard HDS auto-tiering policies for Hitachi Dynamic Tiering (HDT) storage pools.

The HDS auto-tiering policy options in ViPR Controller are:

Policy name in ViPR Controller	Policy number	HDS level	Description
All	0	All	Places the data on all tiers.
T1	1	Level 1	Places the data preferentially in Tier1.
T1/T2	2	Level 2	Places the data in both tiers when there are two tiers, and preferentially in Tiers 1 and 2 when there are three tiers

Policy name in ViPR Controller	Policy number	HDS level	Description
T2	3	Level 3	Places the data in both tiers when there are two tiers, and preferentially in Tier 2 when there are three tiers.
T2/T3	4	Level 4	Places the data in both tiers when there are two tiers, and preferentially in Tiers 2 and 3 when there are three tiers.
T3	5	Level 5	Places the data preferentially in Tier 2 when there are two tiers, and preferentially in Tier 3 when there are three tiers.

VMAX

Review the following configuration requirements and recommendations before virtualizing your VMAX system in ViPR Controller.

Virtual Pool configuration requirements and recommendations

When VMAX is configured with Storage Tiers and FAST Policies:

- Storage Tier and FAST Policy names must be consistent across all VMAX storage systems.
- For more details about ViPR Controller management with FAST policies see: *ViPR Controller Integration with VMAX and VNX Storage Systems User and Administration Guide*
- Set these options when you build your virtual pool:

Option	Description
RAID Level	Select which RAID levels the volumes in the virtual pool will consist of.
Unique Auto-tiering Policy Names	VMAX only. When you build auto-tiering policies on a VMAX through Unisphere, you can assign names to the policies you build. These names are visible when you enable <code>Unique Auto-tiering Policy Names</code> . If you do not enable this option, the auto-tiering policy names displayed in the Auto-tiering Policy field are those built by ViPR.
Auto-tiering Policy	The Fully Automated Storage Tiering (FAST) policy for this virtual pool. FAST policies are supported on VMAX, VNX for Block, and VNXe. ViPR chooses physical storage pools to which the selected auto-tiering policy has been applied. If you create a volume in this virtual pool, the auto-tiering policy specified in this field is applied to that volume.
Fast Expansion	VMAX or VNX Block only. If you enable Fast Expansion, ViPR creates concatenated meta volumes in this virtual pool. If Fast Expansion is disabled, ViPR creates striped meta volumes.
Host Front End Bandwidth Limit	0 - set this value to 0 (unlimited). This field limits the amount of data that can be consumed by applications on the VMAX volume. Host front end bandwidth limits are measured in MB/S.

Option	Description
Host Front End I/O Limit	0 - set this value to 0 (unlimited). This field limits the amount of data that can be consumed by applications on the VMAX volume. Host front end I/O limits are measured in IOPS.

VMAX3

Review the following configuration requirements and recommendations before virtualizing your VMAX3 system in ViPR Controller.

Set these options when you build your virtual pool:

Table 7 VMAX3 Virtual Pool Settings

Field	Description
Provisioning Type	Thin. VMAX3 does not support thick volumes.
Protocols	FC.
System Type	EMC VMAX
Thin Volume Preallocation	0 or 100 . Other values would filter out the VMAX3 SRP pools. 0 - Volumes allocated using this pool are fully-thin. 100 - Volumes allocated using this pool are full-allocated.
Unique Auto-tiering Policy Names	Enabled.
Auto-tiering Policy	VMAX3 is delivered with pre-defined Storage Level Objectives, and workflows. You can specify the workflow and SLO you want applied to your volume during provisioning.
Expandable	Expanding VMAX3 volumes through ViPR Controller is not supported.
Host Front End Bandwidth Limit	0 - set this value to 0 (unlimited). This field limits the amount of data that can be consumed by applications on the VMAX3 volume. Host front end bandwidth limits are measured in MB/S.
Host Front End I/O Limit	0 - set this value to 0 (unlimited). This field limits the amount of data that can be consumed by applications on the VMAX3 volume. Host front end I/O limits are measured in IOPS.

EMC VNX for Block

Review the following configuration consideration before adding VNX for Block storage to the ViPR Controller virtual pools.

Virtual pool configuration considerations

- Fibre Channel networks for VNX for Block storage systems require an SP-A and SP-B port pair in each network, otherwise virtual pools cannot be created for the VNX for Block storage system.
- Prior to ViPR Controller version 2.2, if no auto-tiering policy was set on the virtual pool created from VNX for Block storage, ViPR Controller creates volumes from the virtual pools with auto-tiering enabled. Starting with ViPR Controller version 2.2, if no policy

is set on the virtual pool created for VNX for Block storage, ViPR Controller will create volumes from the virtual pool with the "start high then auto-tier" enabled on the new volumes created in the same virtual pool.

EMC VNXe for Block

It is recommended when exporting a VNXe for Block volume to a host using ViPR Controller that the host is configured with Fibre Channel only or iSCSI connectivity to the storage.

EMC VPLEX

Review the following configuration requirements, and recommendations before virtualizing your third-party block storage in VPLEX .

Virtual array configuration requirements and recommendations

While creating virtual arrays, manually assign the VPLEX front-end and back-end ports of the cluster (1 or 2) to a virtual array, so that each VPLEX cluster is in its own ViPR Controller virtual array.

Virtual pool configuration requirements and recommendations

When running VPLEX with VMAX, the Storage Tier and FAST Policy names must be consistent across all VMAX storage systems.

Third-party block (OpenStack) storage systems

Review the following configuration requirements, and recommendations before virtualizing your third-party block storage in ViPR Controller.

Virtual pool recommendations and requirements

If the discovered storage system is configured for multipathing, the values set in the virtual pool can be increased once the target ports are detected by ViPR Controller.

Block storage systems under ViPR Controller management

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

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

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

File storage configuration considerations

Review the following information before you add file storage systems to ViPR Controller virtual arrays and virtual pools, and before you use the file systems in a ViPR Controller service.

Virtual pool for configuration settings for all file storage systems

File systems are only thinly provisioned. You must set the virtual pool to Thin, when adding file storage to the virtual pool.

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.

Specific storage system configuration requirements

Before you create virtual arrays and virtual pools for File storage in ViPR Controller, review the following sections for storage system specific configuration requirements and recommendations:

- [EMC Data Domain on page 47](#)
- [EMC VNX for File on page 47](#)

EMC® Data Domain®

Review the following information before virtualizing the Data Domain storage in the ViPR Controller virtual arrays and virtual pools.

Virtual pool configuration requirement and considerations

When creating the file virtual pool for Data Domain storage, the **Long Term Retention** attribute must be enabled.

While configuring the file virtual pools for Data Domain storage systems it is helpful to know that:

- A Data Domain Mtree is represented as a file system in ViPR Controller.
- Storage pools are not a feature of Data Domain. However, ViPR Controller uses storage pools to model storage system capacity. Therefore, ViPR Controller creates one storage pool for each Data Domain storage system registered to ViPR Controller, for example, if three Data Domain storage systems were registered to ViPR Controller, there would be three separate Data Domain storage pools. One storage pool for each registered Data Domain storage system.

EMC VNX for File

When configuring a VNX file virtual pool that uses CIFS protocol, there must be at least one CIFS server on any one of the physical data movers.

ViPR requirements for service profile templates

The following sections explain the requirements to configure a service profile template for ViPR Controller provisioning operations.

Note

If existing service profile templates do not match the following requirements, clone one of the service profile template to create a new service profile template and alter the settings as required by ViPR Controller.

General properties

- The service profile template must not be associated to a server pool. Blade selection is performed by the ViPR Controller Compute Virtual Pools.
- UUID assignment must be from a valid UUID Suffix Pool set up in the UCS with available addresses.

Storage

ViPR Controller currently supports Fibre Channel boot for UCS servers. The following lists the Fibre Channel requirements:

- World Wide Node Name (WWNN) assignment must be from a valid UUID Suffix Pool set up in the UCS with available addresses.
- The Local Disk Configuration Policy must be set to a local disk configuration policy where the **Mode** is set to **No Local Storage**.
- There must be at least one vHBA interface.
- For each vHBA, the World Wide Port Name (WWPN) assignment must be from a valid WWPN pool set up in the UCS with available addresses.
- The VSAN set on each vHBA must be a valid network discovered by ViPR Controller. The VSAN must match one of the networks in a ViPR Controller virtual array.
- Policy settings on the vHBAs are not set by ViPR Controller provisioning and are at the administrator's discretion.

Network

- Policy settings on the vNICs are not set by ViPR Controller provisioning and are at the administrator's discretion.
- There must be at least one vNIC interface.
- For each vNIC, the MAC Address Assignment must be from a valid MAC pool that was set up in the UCS with available addresses.
- Each vNIC must have at least one VLAN.

Boot Policy and Boot Order

There are no Boot Policy requirements. ViPR Controller ignores all Boot Policy settings in the service profile template and overwrites any existing parameters when it creates service profiles.

Policies

ViPR Controller does not set any policies. The UCS administrator is responsible for setting the policies.

Updating service profile templates

If provisioning with updating service profile templates,

- The boot policy of the updating service profile template must specify SAN as the first boot device.
- If the boot policy of the updating service profile template enforces vNIC and vHBA names, the names of the vNICs and vHBAs in the service profile template must match those in its boot policy.
- The compute virtual pool with which the updating service profile template is being associated, must be associated to a virtual array that has storage ports on the VSANs that the vHBAs of the template use.
- If the boot policy of the updating service profile template specifies SAN boot target WWPNs, then compute virtual pool that the template is associated with must be associated with a virtual array that includes those storage ports on the appropriate VSANs.

