

EMC ViPR In Existing VMAX Environments

Abstract

This white paper explains how EMC ViPR operates in customer SANs where host provisioning occurred before the introduction of ViPR. These are known as brownfield environments. This paper describes what to expect with VMAX Masking Views and Fibre Channel switch zoning. It also provides solutions to smoother integration.

April 2014

Copyright © 2014 EMC Corporation. All Rights Reserved.

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

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com.

VMware is a registered trademark of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners.

Part Number: H13059

Table of Contents

Executive summary.....	4
Audience.....	4
Understanding Symmetrix Masking Views	5
ViPR Masking View in a Greenfield Environment	6
Physical Environment	6
ViPR Configuration.....	7
Remediation.....	10
ViPR Masking View in a Brownfield Environment.....	12
Physical Environment	12
Brownfield Use Case 1: Provisioning a New LUN from a FAST Pool in ViPR.....	13
Result.....	14
Removing ViPR Provisioned Volumes.....	15
Brownfield Use Case 2: Modifying an Existing Brownfield Environment to Accommodate ViPR	16
Result.....	16
Remediation.....	16
Brownfield Use Case 3: Provisioning a ViPR Non-FAST LUN to a Host with an Existing Masking View that does not have Cascaded Storage Group	18
Result.....	19
Brownfield Use Case 4: Provisioning a ViPR FAST LUN to a Host with an Existing Masking View that does not have Cascaded Storage Group	19
Result.....	20
Remediation.....	20
Brownfield Use Case 5: Provisioning a ViPR LUN to a Host with an Existing Masking View - Virtual Pool has More Paths than the Existing Port Group.....	23
Result.....	24
Brownfield Use Case 6: Provisioning a ViPR LUN to a Host with an Existing Masking View - Virtual Pool has Fewer Paths than the Existing Port Group.....	26
Result.....	27
Brownfield Use Case 7: Using Fibre Channel Switch Zoning to Limit Paths	29
Result.....	30
Conclusion.....	32

Executive summary

EMC® ViPR™ is a software-defined storage platform that abstracts, pools, and automates a datacenter's underlying physical storage infrastructure. It provides a single control plane for heterogeneous storage systems to data center administrators.

Once ViPR is deployed in a data center, ViPR discovers the physical infrastructure including storage systems, Fibre Channel storage area networks (SANs), and hosts so that ViPR can understand the full topology of the data center. ViPR discovers the physical storage pools and the storage ports for each storage system registered with ViPR. The physical pools and ports on the storage systems are used by ViPR to make the storage devices visible to the hosts.

Prior to EMC ViPR, storage provisioning was performed by SAN administrators using a variety of methods including product element managers, APIs, CLI scripting or third-party management applications. They would generally use vendor best practices for configuring Fibre Channel switches and storage arrays. In some cases, administrators develop unique standards that are customized for their organizations.

When EMC ViPR is introduced, it inevitably changes the way provisioning operations are done. In existing environments, EMC VMAX Masking Views and Fibre Channel zones are created by the element manager applications. In many cases, ViPR is compatible with these existing customer configurations, referred to as *brownfield* environments.

However, there are brownfield environments that may require modification for ViPR to provision. In comparison to some SAN configurations, there are also differences with the way that ViPR structures Masking Views and zoning.

This paper describes how ViPR 1.1 builds Masking Views and Fibre Channel switch zones in greenfield and brownfield environments.

Audience

This white paper is intended for storage administrators and architects, customers, and EMC field personnel who want to understand the implementation of storage provisioning with ViPR to a host or a cluster in a Symmetrix VMAX environment.

Understanding Symmetrix Masking Views

A Symmetrix Masking View (MV) is a logical grouping of related objects and the creation of a view that associates the related groups together. When a VMAX MV is created, the necessary mapping and masking operations are performed automatically to provision storage. Once a MV has been created, any changes to the grouping of initiators, ports, or storage devices are propagated throughout the view and the mapping and masking are automatically updated as required.

The full provisioning process involves multiple steps performed by the SAN administrator. The following are the general steps performed on the Symmetrix when storage is provisioned to a new server.

1. Identify and select appropriate devices based on configuration type and capacity.
2. Identify and verify paths between storage and HBAs.
3. Map devices to appropriate front-end director ports.
4. Mask devices to specific HBAs through specific ports.
5. Fabric switch zoning.

Careful analysis and planning is required before implementation.

The following are definitions for these objects.

- An *initiator group* is a logical grouping of up to 32 Fibre Channel initiators or eight iSCSI names or a combination of both. An initiator group may also contain the name of another initiator group to allow the groups to be cascaded to a depth of one.
- A *port group* is a logical grouping of Fibre Channel and/or iSCSI front-end director ports. The only limit on the number of ports in a port group is the number of ports in the Symmetrix VMAX; however, it is likely that a port group will contain a subset of the available ports in order to isolate workloads to specific ports.
- A *storage group* is a logical grouping of up to 4,096 Symmetrix devices. LUN addresses are assigned to the devices in the storage group when the MV is created using the dynamic LUN addressing feature.
- A *masking view* (MV) defines an association between one initiator group, one port group, and one storage groups. When a MV is created, the devices in the storage group are mapped to the ports in the port group and masked to the initiators in the initiator group. Depending on the server and application requirements, each server or group of servers may have one or more MVs that associate a set of Symmetrix devices to an application, server, or cluster of servers.

Note: One (and only one) MV is supported for a host.

ViPR Masking View in a Greenfield Environment

A greenfield environment refers to a newly built SAN infrastructure where there is no existing Fibre Channel switch zoning or array-based MVs.

The following section details an environment where there is no existing SAN configuration for the host. The host has been physically connected to the SAN, but there is no MV or Fibre Channel switch zones.

Physical Environment

Host

ESXi 5.0 (Managed by vCenter 5.1)
10.247.188.13

Emulex HBA:

vmhba2 10:00:00:00:c9:77:48:01
vmhba3 10:00:00:00:c9:77:46:d5

Storage Array

VMAX

000195700999
5876.251.161
7E0 50:00:09:73:00:0f:9d:18
7E1 50:00:09:73:00:0f:9d:19
9G0 50:00:09:73:00:0f:9d:a0
9G1 50:00:09:73:00:0f:9d:a1

Brocade

Connectrix Manager Converged Network Edition 12.0.1

Fabric A (Brocade.170)

ESX HBA ports

vmhba3 10:00:00:00:c9:77:46:d5

VMAX FA ports

7E0 50:00:09:73:00:0f:9d:18
9G0 50:00:09:73:00:0f:9d:a0

Fabric B (Brocade.171)

ESX HBA ports

vmhba2 10:00:00:00:c9:77:48:01

VMAX FA ports

7E1 50:00:09:73:00:0f:9d:19
9G1 50:00:09:73:00:0f:9d:a1

ViPR Configuration

Virtual Array

VA1 includes:


- Brocade switches - Brocade.170 (Fabric A) and Brocade.171 (Fabric B)
- 1 HBA port and 2 storage ports.
- Automatic selected for SAN zoning

Virtual Pool

VP1 includes:



- VA1
- 1 minimum path
- 4 maximum paths
- 1 path per initiator
- 1 Solid State Drive pool from VMAX999

ESX host, 10.247.188.13, currently has no storage provisioned to it on VMAX999. Using the ViPR Service Catalog, the **Create Volume and Datastore** is used to provision storage. Here is the image of the successful ViPR Order:



Create Volume and Datastore

Create block volume and VMware Datastore.

Order Number:	4	Datastore Name:	Test021114B
Date Submitted:	Feb 11th 2014, 11:39:36 am	vCenter:	vCenter.22
Submitted By:	root	Datacenter:	CSE 176 Lab
Status:	✓ Order Successfully Fulfilled	Storage Type:	Exclusive
Execution Time:	4 minutes	ESX Host/Cluster:	10.247.188.13
Precheck Steps:	7 	Virtual Array:	VSA1
Execution Steps:	8 	Virtual Pool:	VSP1
		Project:	Human Resources
		Name:	Test021114B
		Consistency Group:	
		Size (GB):	3
		HLU:	-1

VMAX Masking View

ViPR successfully provisions a 3GB LUN to the ESX host. MVs and FC switch zones are created. The following is created:

Masking View

Name: 10_247_188_13_1392136950099__10_247_188_13
 <FQDN/IP> <Time Stamp> <Export Label>

The format for the naming is the <FQDN/IP>_<Time Stamp>_<Export Label>. In the case of the ViPR UI, the <Export Label> is hostname/IP used. Note that the naming convention across all MV components are the same except for a few differences.

Storage Group

Parent Name: 10_247_188_13_1392136950099__10_247_188_13
Cascaded Name: 10_247_188_13_1392136950099__10_247_188_13_SG_NonFast

If the Virtual Pool contains a VMAX FAST pool, then it would be appended at the end of the Storage Group Cascaded name. If there is no FAST policy associated, then the Storage Group Cascaded includes NonFast.

Initiator Group

Parent Name: 10_247_188_13_1392136950099__10_247_188_13_CIG
Cascaded Name: 10_247_188_13_IG
Initiators:
 20:00:00:00:c9:77:48:01 10:00:00:00:c9:77:48:01
 20:00:00:00:c9:77:46:d5 10:00:00:00:c9:77:46:d5

The Parent name of the Initiator Group is appended with CIG. This is short for *Cascaded Initiator Group*. ViPR adds initiators to the Cascaded Initiator Group.

Port Group

Name: 10_247_188_13_1392136950099__10_247_188_13_PG

The Virtual Pool attributes included minimum paths of 1, maximum paths of 4, and paths per initiator of 1. Since this ESX host (10.247.188.13) has 2 HBA ports, this equates to a total of 2 paths; one path using vmhba2 and the other using vmhba3 in order to achieve redundancy. Because the Virtual Pool has 1 path per initiator as an attribute, 2 FA ports are needed.

VMAX FA Ports selected:

7E0	50:00:09:73:00:0f:9d:18 (Fabric A)
9G1	50:00:09:73:00:0f:9d:a1 (Fabric B)

Host I/O Limits

By default, ViPR does not set Host I/O limits or Child I/O limits.

Volume

Name: 2846:Test021114B

Fibre Channel Switch Zoning

In this greenfield environment with automatic selected for SAN zoning, ViPR performs all zoning operations. The initiators and targets within the zone are chosen based on ViPR's port selection algorithm. This is covered in more detail in another paper.

ViPR zones are created using WWPNs (World-Wide Port Names). ViPR does not create aliases for initiators or targets. The ViPR zone name consists of the following:

- Begins with "SDS"
- Hostname or IP
- HBA WWN
- Last four digits of the array serial number
- Storage port designation

Example:

FC Switch Zoning

Brocade.170

Name: SDS_10_247_188_13_0000C97746D5_0999_FA7E0
Initiator: 10:00:00:00:c9:77:46:d5
Target: 50:00:09:73:00:0f:9d:18

Brocade.171

Name: SDS_10_247_188_13_0000C9774801_0999_FA9G1
Initiator: 10:00:00:00:c9:77:48:01
Target: 50:00:09:73:00:0f:9d:a1

Paths Per Initiator Setting

The Virtual Pool settings directly affect zoning and selected storage ports. In the case of the `Paths Per Initiator` setting, redundant fabric paths may be impacted because the paths per initiator take precedence. Take the following example configuration:

- Host
 - **2 initiators** (1 in Fabric A and 1 Fabric B)
- VMAX FA ports
 - 4 available (2 in Fabric A and 2 Fabric B)
- VP1
 - **Max paths 2**
 - Min paths 1
 - **Paths per initiator 2**

The `Paths Per Initiator` setting takes precedence over redundant fabrics. As you can see, there are 2 paths per initiator yet only 2 max paths. ViPR will look at the 2 paths per initiator

setting first. Therefore, it will choose two FA ports and create two zones in one fabric. The HBA with the initiator on Fabric A will have 2 zones. This meets the 2 paths per initiator requirement. Since the max paths setting is set to 2, ViPR will create no more zones. This configuration is not highly available and is not recommended.

Remediation

There are two options to remedy this situation.

1. *Prior* to provisioning with that Virtual Pool, ensure that the Virtual Pool settings will yield redundant paths.
2. Use the ViPR API to change Virtual Pools.

Ensuring that redundant paths will be created by ViPR requires minor changes to the Virtual Pool. However, they should be made prior to the first provisioning operation against this pool. You must consider the following:

1. Number of initiators on the host
2. Max paths value
3. Paths per initiator value

In the example configuration below, Host1 could potentially be provisioning storage from VP1 or VP2. With the settings below, both Virtual Pools will successfully create a highly available path configuration for Host1.

- Host1
 - 2 initiators (1 in Fabric A and 1 Fabric B)
- VMAX FA ports
 - 4 available (2 in Fabric A and 2 Fabric B)
- VP1
 - Max paths 2
 - Min paths 1
 - Paths per initiator 1
- VP2
 - Max paths 4
 - Min paths 1
 - Paths per initiator 2

Changing Virtual Pools requires using the ViPR API. This function is not available in the GUI. Refer to the ViPR API documentation for more information.

Here is an example of changing Virtual Pools.

- Host1
 - 4 initiators (2 in Fabric A and 2 Fabric B)
- VMAX FA ports
 - 4 available (2 in Fabric A and 2 Fabric B)
- VP1
 - Max paths 2

- Min paths 1
 - Paths per initiator 1
- VP2
 - Max paths 4
 - Min paths 1
 - Paths per initiator 2

If a provisioning operation is performed using VP1, ViPR will create an Initiator Group within the MV that will include all 4 initiators. Since the max paths for VP1 is 2, the Port Group within the MV will have only 2 FA ports. This means that only 2 zones will be created. One initiator and one FA port in each zone (1 in Fabric A and 1 in Fabric B). This is an expected and highly available result since the paths per initiator is 1.

Now, the ViPR administrator wants to increase the number of paths. This can be accomplished through a Virtual Pool change.

Using the ViPR API, the administrator changes the Virtual Pool of a volume in the Storage Group of a MV from VP1 to VP2. (See API docs for exact syntax.)

Since the max paths value has increased from 2 to 4, ViPR can double the number of paths. It does this in several steps:

1. No changes to the IG since all 4 initiators were already added.
2. 2 more FA ports are added to the Port Group.
3. 2 more zones are added (1 to each fabric).

Note: ViPR will *never* zone two initiators to the same FA port.

ViPR Masking View in a Brownfield Environment

A brownfield environment refers to a SAN infrastructure with existing MVs and Zoning. ViPR makes use of existing MVs and Fibre Channel switch zoning when possible.

The following section details an existing provisioning environment that is leveraged by ViPR for new provisioning operations. Several use cases for existing environments are covered in the subsequent sections.

Physical Environment

Host

ESXi 5.0 (Managed by vCenter 5.1)
10.247.188.13

Emulex HBA: vmhba2 10:00:00:00:c9:77:48:01
 vmhba3 10:00:00:00:c9:77:46:d5

Storage Array

VMAX

000195700999
5876.251.161
7E0 50:00:09:73:00:0f:9d:18
7E1 50:00:09:73:00:0f:9d:19
9G0 50:00:09:73:00:0f:9d:a0
9G1 50:00:09:73:00:0f:9d:a1

Brocade

Connectrix Manager Converged Network Edition 12.0.1

Fabric A (Brocade.170)

ESX HBA ports
 vmhba3 10:00:00:00:c9:77:46:d5
VMAX FA ports
 7E0 50:00:09:73:00:0f:9d:18
 9G0 50:00:09:73:00:0f:9d:a0

Fabric B (Brocade.171)

ESX HBA ports
 vmhba2 10:00:00:00:c9:77:48:01
VMAX FA ports
 7E1 50:00:09:73:00:0f:9d:19
 9G1 50:00:09:73:00:0f:9d:a1

Brownfield Use Case 1: Provisioning a New LUN from a FAST Pool in ViPR

In this case, an administrator has used the Unisphere for VMAX storage provisioning wizard to present a LUN to an ESX server from a standard thin pool.

The ViPR administrator will provision a new LUN from a FAST pool.

Existing Masking View

Name: ESXServer1_MV

Existing Storage Group

Parent Name: ESXServer1_SG

Cascaded Name: BrownfieldSG

Existing Initiator Group

Name: ESXServer1

Cascaded Name: <none>

Initiators: 20:00:00:00:c9:77:48:01 10:00:00:00:c9:77:48:01
20:00:00:00:c9:77:46:d5 10:00:00:00:c9:77:46:d5

Existing Port Group

Name: ESXServer1_PG

VMAX FA Ports selected: 7E0 50:00:09:73:00:0f:9d:18 (Fabric A)
9G0 50:00:09:73:00:0f:9d:a0 (Fabric A)
7E1 50:00:09:73:00:0f:9d:19 (Fabric B)
9G1 50:00:09:73:00:0f:9d:a1 (Fabric B)

Existing Volume

Name: 2846:Test021114B

Existing FC Switch Zoning

Fabric A

zone: ESXServer_46d5_7E0_9G0
10:00:00:00:c9:77:46:d5
50:00:09:73:00:0f:9d:18
50:00:09:73:00:0f:9d:a0

Fabric B

zone: ESXServer_4801_7E1_9G1
10:00:00:00:c9:77:48:01
50:00:09:73:00:0f:9d:19
50:00:09:73:00:0f:9d:a1

ViPR Virtual Array

VA1 includes:

- Brocade switches - Brocade.170 (Fabric A) and Brocade.171 (Fabric B)
- 2 HBA ports and 4 storage ports.

- Automatic selected for SAN zoning.

ViPR Virtual Pool

VP3 includes:

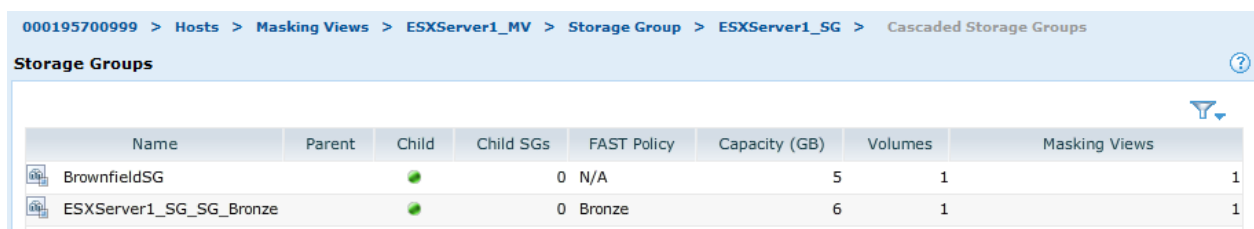
- VA1
- 1 minimum path
- 4 maximum paths
- 2 paths per initiator
- 1 FAST pool from VMAX999

Result

As a result of this particular ViPR provisioning operation, two things occur.

1. A new Cascaded Storage Group is created.
2. New zones on both fabrics are created.

ViPR leverages the existing MV. However, since ViPR created a new LUN from a FAST pool and the existing LUN provisioned was from a non-FAST pool, ViPR created a new Cascaded Storage Group on the VMAX.



Name	Parent	Child	Child SGs	FAST Policy	Capacity (GB)	Volumes	Masking Views
BrownfieldSG		●	0	N/A	5	1	1
ESXServer1_SG_SG_Bronze		●	0	Bronze	6	1	1

The newly created Cascaded Storage Group is named `ESXServer1_SG_SG_Bronze`. The Parent name of the existing Storage Group was `ESXServer1_SG`. ViPR appended `SG_Bronze` to reflect the new Storage Group from the Bronze FAST pool.

Even though there were existing zones on each fabric, ViPR doesn't rely on the zoning created by the SAN administrator when SAN zoning in the Virtual Array is configured for automatic. ViPR creates single initiator and single target zones. The attributes of VP3 state that storage provisioned from this Virtual Pool will have 2 paths per initiator. The host has two HBA ports, one connected to each fabric. For VMAX, that means each HBA port will be zoned to two VMAX FA ports for a total of four active I/O paths.

In the case of VNX, the number of storage ports double due to the nature of the optimized and non-optimized pathing on an ALUA array. In this scenario, when 2 paths per initiator is configured in the Virtual Pool, that means each HBA port will be zoned to four VNX Storage Processor ports for a total of four optimized I/O paths and four non-optimized I/O paths.

New zones on Fabric A

```
zone: SDS_10_247_188_13_0000C97746D5_0999_FA7E0
      10:00:00:00:c9:77:46:d5
      50:00:09:73:00:0f:9d:18
zone: SDS_10_247_188_13_0000C97746D5_0999_FA9G0
      10:00:00:00:c9:77:46:d5
      50:00:09:73:00:0f:9d:a0
```

New zones on Fabric B

```
zone: SDS_10_247_188_13_0000C9774801_0999_FA7E1
      10:00:00:00:c9:77:48:01
      50:00:09:73:00:0f:9d:19
zone: SDS_10_247_188_13_0000C9774801_0999_FA9G1
      10:00:00:00:c9:77:48:01
      50:00:09:73:00:0f:9d:a1
```

Removing ViPR Provisioned Volumes

When the ViPR administrators remove volumes, ViPR takes the following action:

1. Volumes are deleted from the Cascaded Storage Group created by ViPR. The Cascaded Storage Group is not removed though.
2. ViPR-created zones are removed. For each provisioned volume, ViPR adds a reference structure for the new volume(s) using the zone that it created. When new volumes are provisioned, the same zone is used. The zone will only be removed once the last volume using that zone is deleted.

Brownfield Use Case 2: Modifying an Existing Brownfield Environment to Accommodate ViPR

In brownfield environments, there are several configurations that require modifications to the existing MV.

1. Initiators in Multiple MVs.
2. MV without Cascaded Initiator Groups
3. Provision a ViPR non-FAST LUN to a host with an existing MV that does NOT have Cascaded Storage Group.
4. Provision a ViPR FAST LUN to a host with an existing MV that does NOT have Cascaded Storage Group.
5. Provision a ViPR LUN to a host with an existing MV. VP has more paths than the existing Port Group has.

Initiators in Multiple Masking Views

In this brownfield scenario, a SAN administrator has provisioned two volumes to a host. However, two MVs were used. Both MVs share a common set of initiators.

Existing Masking View 1

Name: ESXServer1_MV

Existing Initiator Group 1

Name: ESXServer1

Initiators: 10:00:00:00:c9:77:48:01
10:00:00:00:c9:77:46:d5

Existing Masking View 2

Name: ESXServer1_B_MV

Existing Initiator Group 2

Name: ESXServer1

Initiators: 10:00:00:00:c9:77:48:01
10:00:00:00:c9:77:46:d5

Result

As a result of this particular ViPR provisioning operation:


1. A ViPR error is generated.

Remediation

To complete this provisioning operation:




1. MVs must be consolidated.

In a brownfield configuration like this, ViPR will generate an error if the administrator attempts to provision to the same host on the same array with the same initiators. As mentioned earlier, ViPR will leverage existing MVs when possible. In the case where the initiators are attached to multiple MVs, ViPR does not which to select.



Create Volume and Datastore

Create block volume and VMware Datastore.

Order Number:	12	Datastore Name:	DS021414A
Date Submitted:	Feb 14th 2014, 11:51:06 am	vCenter:	vCenter.22
Submitted By:	root	Datacenter:	CSE 176 Lab
Status:	✖ Error Occurred Processing Order	Storage Type:	Exclusive
Execution Time:	3 minutes	ESX Host/Cluster:	10.247.188.13
Precheck Steps:	7 	Virtual Array:	VSA1
Execution Steps:	3 	Virtual Pool:	VSP1
Rollback Steps:	1 	Project:	Human Resources
		Name:	vol021414A
		Consistency Group:	
		Size (GB):	2
		HLU:	-1
		Re-Submit:	<button>Order</button>

Create Export

Name: 10.247.188.13-1392396849936, HostName: 10.247.188.13, Volume: urn:storageos:Volume:1ffcd10a-bced-4ba5-9497-a0cf08d5e7e6:, HLU: -1

11 seconds

com.emc.vipr.client.exceptions.ServiceErrorException: Error 12023: Export create operation failed. Encountered an error in export create: While attempting do an export operation, more than one existing mask was found for the same compute resource: 10000000C97746D5=[urn:storageos:ExportMask:2fe90c3c-3263-43ca-8c47-1f11f8789543:, urn:storageos:ExportMask:e67fe4ea-e9f8-4df8-b365-7ddb66efb696:], 10000000C9774801=[urn:storageos:ExportMask:2fe90c3c-3263-43ca-8c47-1f11f8789543:, urn:storageos:ExportMask:e67fe4ea-e9f8-4df8-b365-7ddb66efb696:]. ViPR will be unable to export volumes in this configuration.

In order to resolve the issue, the initiators must be part of *only one* MV. Combining MVs is an offline event. In order to move a Cascaded Storage Group from one Parent Storage Group to another, the Cascaded Storage Group must first be removed from one. The volume will not be available to the host until it is added to the target Parent Storage Group. Once the initiators are members of a single MV, ViPR will be able to provision.

Brownfield Use Case 3: Provisioning a ViPR Non-FAST LUN to a Host with an Existing Masking View that does not have Cascaded Storage Group

In this brownfield scenario, a SAN administrator has an existing MV with one or more LUNs provisioned to the host. However, this existing MV was built with a Storage Group *without* Cascaded Storage Groups. That is, the LUN(s) are directly under the Primary Storage Group.

Existing Masking View

Name: ESXServer1_MV

Existing Storage Group

Parent Name: ESXServer1_SG

Cascaded Name: <none>

Existing Initiator Group

Name: ESXServer1

Initiators: 10:00:00:00:c9:77:48:01

10:00:00:00:c9:77:46:d5

Existing Port Group

Name: ESXServer1_PG

VMAX FA Ports selected:	7E0	50:00:09:73:00:0f:9d:18 (Fabric A)
	9G0	50:00:09:73:00:0f:9d:a0 (Fabric A)
	7E1	50:00:09:73:00:0f:9d:19 (Fabric B)
	9G1	50:00:09:73:00:0f:9d:a1 (Fabric B)

Existing Volume

Name: 28D1:TestVolume

ViPR Virtual Array

VA1 includes:


- Brocade switches - Brocade.170 (Fabric A) and Brocade.171 (Fabric B)
- 2 HBA ports and 4 storage ports.
- Automatic selected for SAN zoning.



ViPR Virtual Pool

- VP2 includes:
- VA1
- 1 minimum path
- 4 maximum paths
- 2 paths per initiator
- 1 non-FAST pool from VMAX999

Result


As a result of this particular ViPR provisioning operation, in a brownfield configuration like this, ViPR supports this configuration. No errors are generated.

 **Create Volume and Datastore**
Create block volume and VMware Datastore.


Order Number: 35
Date Submitted: Feb 24th 2014, 11:59:45 am
Submitted By: root
Status: ✓ Order Successfully Fulfilled
Execution Time: 6 minutes
Precheck Steps: 7 
Execution Steps: 8 

Datastore Name: DS0224A
vCenter: vCenter.22
Datacenter: CSE 176 Lab
Storage Type: Exclusive
ESX Host/Cluster: 10.247.188.13
Virtual Array: VSA1
Virtual Pool: VSP2
Project: Human Resources
Name: volume0224A
Consistency Group:
Size (GB): 9
HLU: -1

Affected Resources


Block Export

Block Export: urn:storageos:ExportGroup:a6bb50e5-0825-447a-8956-0d9ebab5e605:
Name: 10.247.188.13-1393261469811
Virtual Array: VSA1


Volume

Volume: urn:storageos:Volume:0afc03a9-74a9-476e-8885-704b4cedf3cc:
Name: volume0224A
WWN: 60000970000195700999533032383634
LUN: 1
Size: 9.00 GB Provisioned / 9.00 GB Requested
Virtual Array: VSA1
Virtual Pool: VSP2

Brownfield Use Case 4: Provisioning a ViPR FAST LUN to a Host with an Existing Masking View that does not have Cascaded Storage Group

In this brownfield scenario, a SAN administrator has an existing MV with one or more LUNs provisioned to the host. However, this existing MV was built with a Storage Group without Cascaded Storage Groups. That is, the LUN(s) are directly under the Primary Storage Group.

Existing Masking View

Name: ESXServer1_MV

Existing Storage Group

Parent Name: ESXServer1_SG

Cascaded Name: <none>

Existing Initiator Group

Name: ESXServer1

Initiators: 10:00:00:00:c9:77:48:01

10:00:00:00:c9:77:46:d5

Existing Port Group

Name: ESXServer1_PG

VMAX FA Ports selected:	7E0	50:00:09:73:00:0f:9d:18 (Fabric A)
	9G0	50:00:09:73:00:0f:9d:a0 (Fabric A)
	7E1	50:00:09:73:00:0f:9d:19 (Fabric B)
	9G1	50:00:09:73:00:0f:9d:a1 (Fabric B)

Existing Volume

Name: 28D1:TestVolume

ViPR Virtual Array

VA1 includes:

- Brocade switches - Brocade.170 (Fabric A) and Brocade.171 (Fabric B)
- 2 HBA ports and 4 storage ports.
- Automatic selected for SAN zoning.

ViPR Virtual Pool

VP3 includes:

- VA1
- 1 minimum path
- 4 maximum paths
- 2 paths per initiator
- 1 FAST pool from VMAX999

Result

As a result of this particular ViPR provisioning operation, a ViPR error is generated.

Remediation

To complete this provisioning operation:

1. Cascaded Storage Groups must be created. This is an offline event.

In a brownfield configuration like this, SYMAPI does not support adding a FAST LUN to a Storage Group without a Cascaded Storage Group. Therefore, ViPR would not support this either. The ViPR error looks like this:

Create Volume and Datastore

Create block volume and VMware Datastore.

Order Number:	36	Datastore Name:	DS0224B3
Date Submitted:	Feb 24th 2014, 12:17:31 pm	vCenter:	vCenter:22
Submitted By:	root	Datacenter:	CSE 176 Lab
Status:	✖ Error Occurred Processing Order	Storage Type:	Exclusive
Execution Time:	8 minutes	ESX Host/Cluster:	10.247.188.13
Precheck Steps:	7 ■ ■ ■ ■ ■ ■ ■	Virtual Array:	VSA1
Execution Steps:	3 ■ ■ ■	Virtual Pool:	VSP3
Rollback Steps:	1 ■	Project:	Human Resources
		Name:	volume0224B3
		Consistency Group:	
		Size (GB):	2
		HLU:	-1
		Re-Submit:	Order

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 ESXServer1_SG was found on the array with non-FAST volumes in it. Adding FAST volumes to this Storage Group is not permissible.

Use Unisphere to create a new Storage Group, example: ParentESXServer1_SG. In the **Create Storage Group – Welcome** dialogue box, select **Cascaded Storage Group** in the **Storage Group Type**.

1 Create Storage Group - Welcome

* Storage Group Name:

Storage Group Type:

Volumes Type:

Select **ESXServer1_SG** from the list of available Storage Groups. ESXServer1_SG will become a Child to the new Storage Group, ParentESXServer1_SG.

2 Create Storage Group

Select one or more storage groups to be children in the cascaded group

Storage Group	Volumes	Cap (GB)
d956f8db363153ec_CSE_TestEnv6_DATA	5	1202.44
d956f8db363153ec_lppa220_iss_emc_com_BOOT	1	20
ESXServer1_SG	1	4
FC15K_3RAID5-thinDeviceGroup	0	0
FC_Gold_BS	4	4
FLASH_3RAID5-thinDeviceGroup	0	0
Gold-thinDeviceGroup	0	0
Gold_FC-thinDeviceGroup	0	0
Graham_Cluster_2_1386674791982_Graham_Cluster_2_SG_1	1	1
Graham_Cluster_2_1386789264029_Graham_Cluster_2_SG_1	1	1
Iglah014	5	1202.44
Iglah134_4784	6	0.06
Iglah134_4785	6	0.06
Iglah135_20b12e	6	0.06

[< Back](#) [Next >](#) [Finish](#) [Cancel](#) [Help](#)

The application will have to be taken offline in order to move the new Parent Storage Group to the new MV. Follow these steps:

1. Delete the old MV, ESXServer1_MV
2. Create a new MV, ESXServer1_MV2, that will contain the same Initiator Group and Port Group as the original MV. Add the new Parent Storage Group.

Create Masking View

* Masking View Name: ESXServer1_MV2

* Initiator Group: ESXServer1

* Port Group: ESXServer1_PG

* Storage Group: ParentESXServer1_SG

Set dynamic LUNs

OK Cancel Help

Now that Cascaded Storage Groups are available, retry the ViPR provisioning operation. It will succeed now that ViPR can create a Child Storage Group for the FAST LUN.

The new Storage Group structure looks like this in Unisphere.

000195700999 > Hosts > Masking Views > ESXServer1_MV2 > Storage Group > ParentESXServer1_SG > Cascaded Storage Groups

Storage Groups

Name	Parent	Child	Child SGs	FAST Policy	Capacity (GB)	Volumes
ESXServer1_SG		●	0	N/A	4	1
ParentESXServer1_SG_SG_Bronze		●	0	Bronze	2	1

Brownfield Use Case 5: Provisioning a ViPR LUN to a Host with an Existing Masking View - Virtual Pool has More Paths than the Existing Port Group

In this brownfield scenario, a SAN administrator has an existing MV with one or more LUNs provisioned to the host. This existing MV was built with a Port Group with 2 VMAX FA ports. The ViPR Virtual Pool has a four storage port requirement.

As previously discussed, ViPR will attempt to reuse a MV when possible. However, ViPR will not modify the Port Group within the existing MV. A ViPR provisioning operation is primarily based on the attributes of the Virtual Pool. It specifies the number of paths per initiator. The number of paths in the existing Port Group may differ from the Virtual Pool. ViPR will neither increase nor decrease the number of paths in an existing Port Group. For example, if there are 2 VMAX FA ports in the existing Port Group and the Virtual Pool requires four VMAX FA ports, then the number of VMAX FA ports remains at 2.

Since the number of storage ports remains static, the user must be aware of this before continuing with a ViPR provisioning operation in a brownfield environment. In this scenario, the provisioning operation will complete successfully. However, the user may not be aware that they did not get the number of paths that they expected.

Existing Masking View

Name: ESXServerA_MV

Existing Storage Group

Parent Name: ESXServer1_SG

Child Name: ESXServer1_SG_C

Existing Initiator Group

Initiators: 10:00:00:00:c9:77:48:01
10:00:00:00:c9:77:46:d5

Existing Port Group

Name: ESXServer1_PG2ports

VMAX FA Ports selected: 9G0 50:00:09:73:00:0f:9d:a0 (Fabric A)
9G1 50:00:09:73:00:0f:9d:a1 (Fabric B)

Existing Volume

Name: 28D1:TestVolume

ViPR Virtual Array

VA1 includes:

- Brocade switches - Brocade.170 (Fabric A) and Brocade.171 (Fabric B)
- 2 HBA ports and 4 storage ports.
- Automatic selected for SAN zoning.

ViPR Virtual Pool


VP2 includes:

- VA1
- 1 minimum path
- 4 maximum paths
- 2 paths per initiator

Result

As a result of this ViPR provisioning operation, a volume is successfully created and provisioned.

Note: The Port Group is unmodified.

 **Create Volume and Datastore**
Create block volume and VMware Datastore.

Order Number: 7

Date Submitted: Mar 4th 2014, 2:01:39 pm

Submitted By: root

Status: Order Successfully Fulfilled

Execution Time: 9 minutes

Precheck Steps:

Execution Steps:

Datastore Name: DSE0304

vCenter: vCenter.22

Datacenter: CSE 176 Lab

Storage Type: Exclusive

ESX Host/Cluster: 10.247.188.13

Virtual Array: VSA1

Virtual Pool: VSP2

Project: Human Resources

Name: volumeE0304

Consistency Group:

Size (GB): 6

HLU: -1

In ViPR Virtual Pool VP2, the user would expect 4 paths since there are 2 HBAs in the host and the paths per initiator is 2.

Minimum Paths: *

Minimum number of total paths from the host to storage array

Maximum Paths: *

Maximum number of total paths from host to storage array

Paths Per Initiator: *

Number of paths per host initiator

The Port Group, ESXServer1_PG2ports, has 2 target ports, 9G0 and 9G1. The Port Group remains the same before and after ViPR completes the operation.

000195700999 > Hosts > Masking Views > ESXServer1_MV > Port Group > ESXServer1_PG2ports > Ports			
Ports			
Dir:Port	Port Groups	Masking Views	
FA-9G:0	5	5	
FA-9G:1	5	5	

To satisfy the requirements of continued provisioning operations using VP2, the SAN administrator should perform the following actions in Unisphere or SYMCLI:

1. Manually add FA ports to the Port Group.
2. Zone the initiators to the FA ports.

000195700999 > Hosts > Masking Views > ESXServer1_MV > Port Group > ESXServer1_PG2ports > Ports

Ports			
Dir:Port	Port Groups	Masking Views	
FA-7E:0	5	5	
FA-7E:1	5	5	
FA-9G:0	5	5	
FA-9G:1	5	5	

By adding FA ports, the path count for all MVs using this Port Group will be impacted. This can adversely impact the number of possible LUNs and paths for an operating system. As an example, the ESX host system fails to report a physical path to storage. The maximum number of supported paths (1024) could have been reached.

In recommended brownfield configurations prior to ViPR's introduction, administrators create a separate Port Group for each MV, it allows for that port group to be modified for an individual host rather than modifying a *global* Port Group which affects multiple hosts.

Brownfield Use Case 6: Provisioning a ViPR LUN to a Host with an Existing Masking View - Virtual Pool has Fewer Paths than the Existing Port Group

In this brownfield scenario, a SAN administrator has an existing MV with one or more LUNs provisioned to the host. This existing MV was built with a Port Group with 4 VMAX FA ports. The ViPR Virtual Pool has a 2 storage port requirement. This is reverse of the scenario in the previous section.

ViPR will attempt to reuse a MV when possible. However, ViPR will not modify the Port Group within the existing MV. A ViPR provisioning operation is primarily based on the attributes of the Virtual Pool. It specifies the number of paths per initiator. The number of paths in the existing Port Group may differ from the Virtual Pool. ViPR will neither increase nor decrease the number of paths in an existing Port Group.

Since the number of storage ports remains static, the user must be aware of this before continuing with a ViPR provisioning operation in a brownfield environment. In this scenario, the provisioning operation will complete successfully. However, the user may not be aware that they did not get the number of paths that they expected.

Existing Masking View

Name: ESXServerA_MV

Existing Storage Group

Parent Name: ESXServer1_SG

Child Name: ESXServer1_SG_C

Existing Initiator Group

Initiators: 10:00:00:00:c9:77:48:01
10:00:00:00:c9:77:46:d5

Existing Port Group

Name: ESXServer1_PG

VMAX FA ports selected	7E0	50:00:09:73:00:0f:9d:18 (Fabric A)
	9G0	50:00:09:73:00:0f:9d:a0 (Fabric A)
	7E1	50:00:09:73:00:0f:9d:19 (Fabric B)
	9G1	50:00:09:73:00:0f:9d:a1 (Fabric B)

Existing Volume

Name: 28D1:TestVolume

ViPR Virtual Array

VA1 includes:

- Brocade switches - Brocade.170 (Fabric A) and Brocade.171 (Fabric B)
- 2 HBA ports and 4 storage ports.

- Automatic selected for SAN zoning.

ViPR Virtual Pool


VP2 includes:

- VA1
- 1 minimum path
- 4 maximum paths
- 1 paths per initiator

Result



As a result of this ViPR provisioning operation, a volume is successfully created and provisioned.

The Port Group is unmodified.



Create Volume and Datastore

Create block volume and VMware Datastore.

Order Number:	9	Datastore Name:	DSA0305
Date Submitted:	Mar 5th 2014, 8:17:13 am	vCenter:	vCenter.22
Submitted By:	root	Datacenter:	CSE 176 Lab
Status:	✓ Order Successfully Fulfilled	Storage Type:	Exclusive
Execution Time:	11 minutes	ESX Host/Cluster:	10.247.188.13
Precheck Steps:	7 	Virtual Array:	VSA1
Execution Steps:	8 	Virtual Pool:	VSP1
		Project:	Human Resources
		Name:	volumeA0305
		Consistency Group:	
		Size (GB):	2
		HLU:	-1

In ViPR Virtual Pool VP1, the user would expect 2 paths since there are 2 HBAs in the host and the paths per initiator is 1.

Minimum Paths: *
Minimum number of total paths from the host to storage array

Maximum Paths: *
Maximum number of total paths from host to storage array

Paths Per Initiator: *
Number of paths per host initiator

The Port Group, ESXServer1_PG, has 4 target ports, 9G0, 9G1, 7E0, and 7E1. The Port Group remains the same before and after ViPR completes the operation.

000195700999 > Hosts > Masking Views > ESXServer1_MV > Port Group > ESXServer1_PG2ports > Ports			
Ports			
Dir:Port	Port Groups	Masking Views	
FA-7E:0	5	5	
FA-7E:1	5	5	
FA-9G:0	5	5	
FA-9G:1	5	5	

To satisfy the requirements of continued provisioning operations using VP1, the SAN administrator should perform the following actions:

1. Manually remove FA ports from the Port Group.
- or
2. Remove 1 zone from Fabric A and B.

By removing FA ports, the path count for all MVs using this Port Group will be impacted. By removing zones, the path count for a particular host will be impacted. This can adversely impact the number paths for existing LUNs to any operating system using this Port Group. With paths removed, I/O performance will be directly impacted.

In recommended brownfield configurations prior to ViPR's introduction, administrators create a separate Port Group for each MV, it allows for that port group to be modified for an individual host rather than modifying a *global* Port Group which affects multiple hosts.

Brownfield Use Case 7: Bypassing an Existing Masking View

As discussed in the previous sections, ViPR leverages existing MVs when possible. In some cases, ViPR can seamlessly provision into the brownfield environment. In other cases, modifications to the existing MV may be required. However, another option is available.

By changing the name of an existing MV, ViPR will ignore it as if it does not exist. This effectively looks like a greenfield environment to ViPR because it doesn't see a MV for a specific host that it can use. ViPR will create an entirely new MV.

Add `no_vipr` or `NO_VIPR` to the MV name. In the example below, `no_vipr` was prepended to the `ESXServer1_MV` MV name. However, it may be added anywhere in the name.

000195700999 > Hosts > Masking Views > ESXServer1_MV

Details : Masking View : ESXServer1_MV

Properties	
Name	no_vipr_ESXServer1_MV
Initiator Group	ESXServer1
Number of initiators	2
Port Group	ESXServer1_PG_7E0_7E1
Number of ports	2
Storage Group	ESXServer1_CSG
Number of volumes	2

Once complete, the SMI-S Provider used by ViPR to communicate with the Symmetrix must be refreshed. Otherwise, ViPR won't know of the changes made through Unisphere. Now, ViPR will ignore the MV named `no_vipr_ESXServer1_MV` when provisioning to the same host.

Brownfield Use Case 7: Using Fibre Channel Switch Zoning to Limit Paths

Zoning can be used to limit the number paths to a LUN. In some configurations, Initiator Groups and Port Groups may have more initiator and target ports than what the administrator actually wants to initially use. However, the zones may be used to expand or limit host to storage connectivity. The number of paths to a LUN can easily be changed by modifying existing zonesets.

For example, an administrator creates a MV with 4 ports in the Port Group. The Initiator Group has 2 HBA ports. However, the administrator zoned 2 of those ports, one on fabric A and one on fabric B. In the ViPR provisioning operation, it will make use of the existing MV. ViPR reads the existing MV. It does not read the existing zones. New Fibre Channel switch zones are created by ViPR to match the existing Port Group. Therefore, the customer's original intent to limit the ports via zoning will be overridden by the new ViPR zones.

In the example, the existing Port Group has 4 target ports, 9G0, 9G1, 7E0 and 7E1.

Existing Masking View

Name: ESXServerA_MV

Existing Storage Group

Parent Name: ESXServer1_SG

Child Name: ESXServer1_SG_C

Existing Initiator Group

Initiators: 10:00:00:00:c9:77:48:01
10:00:00:00:c9:77:46:d5

Existing Port Group

Name: ESXServer1_PG

VMAX FA ports selected

7E0	50:00:09:73:00:0f:9d:18 (Fabric A)
9G0	50:00:09:73:00:0f:9d:a0 (Fabric A)
7E1	50:00:09:73:00:0f:9d:19 (Fabric B)
9G1	50:00:09:73:00:0f:9d:a1 (Fabric B)

Existing Volume

Name: 28D1:TestVolume

In the zoning example below, Fabrics A and B only contain one FA port each. In this configuration, the SAN administrator has limited the number of paths to the volume by restricting the zoning rather than removing FA ports from the Port Group.

Existing FC Switch Zoning

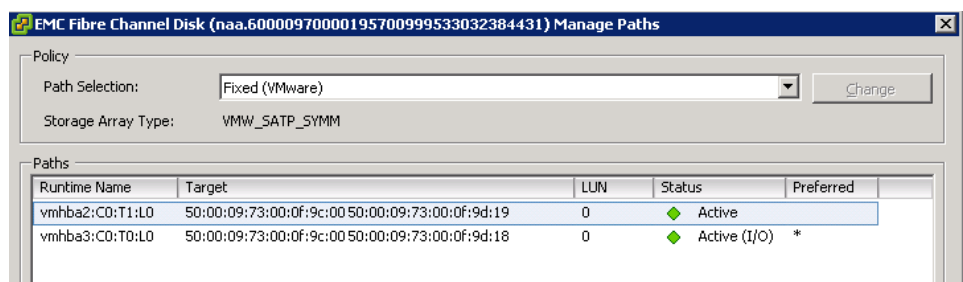
Fabric A

zone: ESXServer_46d5_7E0
10:00:00:00:c9:77:46:d5
50:00:09:73:00:0f:9d:18 7E0

Fabric B

zone: ESXServer_4801_7E1
10:00:00:00:c9:77:48:01
50:00:09:73:00:0f:9d:19 7E1

Looking at the Manage Paths window in vCenter for volume 28D1, there are 2 paths to the LUN.



Result

As a result of this ViPR provisioning operation, two things occur.

1. Volume is successfully created
2. The volume is provisioned into the existing MV.

Notes

Port Group is unmodified.

New zones are created that include ALL ports in the VMAX Port Group.
Paths are increased to the existing volume.

ViPR identifies the initiators and targets within the existing MV. Even though zones already exist, ViPR does not read the existing zoneset. It creates ViPR zones in each fabric that takes into account the HBA ports in the Initiator Group, VMAX FA ports in the Port Group and the requirement for pathing in the ViPR Virtual Pool.

The Virtual Pool for this provisioning operation requires 2 paths per initiator. Since this Initiator Group in the existing MV has 2 HBA ports, a total of 4 paths will be created.

Minimum Paths: *
Minimum number of total paths from the host to storage array

Maximum Paths: *
Maximum number of total paths from host to storage array

Paths Per Initiator: *
Number of paths per host initiator

ViPR creates the following zones in Fabrics A and B. In comparison, note the new ViPR zones with the additional FA ports.

Fabric A

zone: SDS_lppa013_lss_emc_com_0000C97746D5_0999_FA7E0
10:00:00:00:c9:77:46:d5
50:00:09:73:00:0f:9d:18

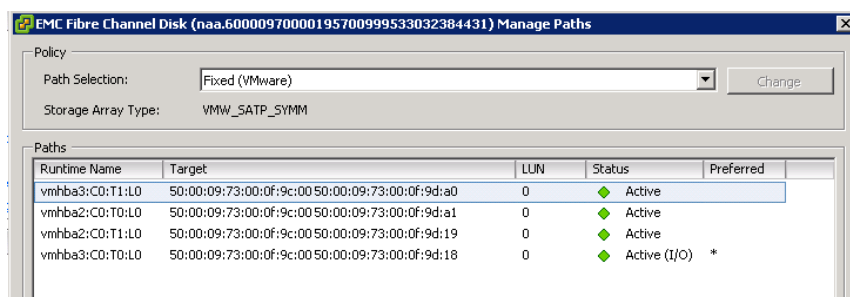
zone: SDS_lppa013_lss_emc_com_0000C97746D5_0999_FA9G0
10:00:00:00:c9:77:46:d5
50:00:09:73:00:0f:9d:a0

Fabric B

zone: SDS_lppa013_lss_emc_com_0000C9774801_0999_FA7E1
10:00:00:00:c9:77:48:01
50:00:09:73:00:0f:9d:19

zone: SDS_lppa013_lss_emc_com_0000C9774801_0999_FA9G1
10:00:00:00:c9:77:48:01
50:00:09:73:00:0f:9d:a1

After the provisioning process, looking at the **Manage Paths** window in vCenter for volume 28D1, there are now 4 paths to the LUN.



Conclusion

In conclusion, EMC ViPR provides SAN administrators a new management option for provisioning block, file and object storage. A ViPR provisioning operation configures storage, SAN, and hosts parameters. With EMC VMAX, the MV includes the Initiator Group, Storage Group, and Port Group. In a brownfield configuration, ViPR is provisioning to a host that already has a MV associated with it. Because EMC ViPR is architected to follow EMC best practices for block SAN provisioning, there are cases where ViPR will integrate seamlessly. When best practices are not followed, it may be necessary to modify the existing MV. This paper covered both cases and will provide the administrator the necessary tools to ensure EMC ViPR can successfully provision EMC VMAX storage.