

EMC ViPR Controller

Version 3.6

New Features and Changes

302-003-909

01

Copyright © 2013-2017 Dell Inc. or its subsidiaries. All rights reserved.

Published May 2017

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

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS-IS.” DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

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

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

CHAPTER 1

New Features

- Overview4
- Service order enhancements4
- Node recovery on vApp for offline node5
- Diagutils6
- ViPR Controller backup and restore enhancements8
- ViPR Controller log file collection enhancements9
- Log retention policy enhancement9
- Storage driver auto installation9
- Multiple concurrent actions on export groups are permitted10
- Export group path adjustment11
- Improved maintenance of export group inventory12
- Ensuring consistent HLU assignments for cluster exports12
- Changes made to the Remove Block Volume service14
- Switch affinity added to port allocation rules16
- SRDF Metro on VMAX supports adding new storage while solution remains active17
- Improvements to discovery processing18
- HTTP protocol can be used when AllowUnencrypted flag is disabled for WinRM
Windows discovery19
- SHA-256 support19
- Vblock enhancements19
- Isilon file enhancements20
- File Protection Policy templates22

Overview

This document summarizes the new features and changes provided with the ViPR Controller 3.6 release.

ViPR Controller REST API updates are documented in the *ViPR Controller REST API Reference*, available as a zip file from the [ViPR Controller Product Documentation Index](#).

ViPR Controller UI and CLI updates are highlighted with the new feature descriptions in this document. For more details, see the following:

- For the ViPR Controller UI, see the ViPR Controller Online Help, or the *ViPR Controller User Interface Virtual Data Center Configuration Guide*.
- For the ViPR Controller CLI, see the *ViPR Controller CLI Reference Guide*

All documents can be accessed from the [ViPR Controller Product Documentation Index](#).

Service order enhancements

This release includes improved order querying features and new features for managing order life cycles.

Order querying is improved with the following new features and changes:

- A Date Range filter is added to the order query functionality on the **Catalog > All Orders** and **Catalog > My Orders** pages of the ViPR Controller User Interface.
- To avoid out-of-memory errors, the displayed result of an order query is limited to 6000 orders. An informational message informs the user if the actual results contain more than 6000 orders.
- Order query processing efficiency was improved by changes to the underlying database schema.

Order life-cycle management is improved with the following new features and changes:

- **Download** and **Delete** buttons are added to the **Catalog > All Orders** page in the ViPR Controller User Interface. Using these features, administrators with Tenant Administrator role can download information about selected orders to a local folder, and then delete orders from the database.
- A check box column was added to the table on the **Catalog > All Orders** page to provide a way to select the orders to download or delete.
- A select-all checkbox in the column header selects all orders that satisfy the most recent query. For example, if a query returns 9,000 orders, only 6,000 orders can display on the UI, but the select-all checkbox applies to all 9,000 for the purposes of using the **Download** or **Delete** features.
- Administrators can use the **Delete** feature as an effective database maintenance tool. The feature deletes the selected orders and all objects and logs related to them from the database. The expected workflow is to download orders (for reference purposes) before deleting them.
- If a download or delete job is running, an informational message about the job appears at the top of the page. The informational message includes the job id and describes how many orders are included in the job, how long the job has been

running, and the expected remaining time until completion. Only one such job can run at a time. The **Download** and **Delete** buttons are disabled for the duration of the job run.

- If an order does not meet the following conditions, its deletion fails and it is not deleted.
 - The order status must be one of the following: SUCCESS, PARTIAL_SUCCESS, ERROR, CANCELLED, REJECTED.

Note

Orders associated with scheduled events have a status of SCHEDULED and can not be deleted.

- The order must be older than 30 days.
- No more than 20,000 order deletions can be processed at a time. If a delete job contains more than that, the first 20,000 orders are deleted initially, and the job continues 5 days later to delete the next 20,000 orders. This processing continues until all of the requested orders are deleted. If a job contains 50,000 orders, the job takes more than 10 days to complete. The job continuation occurs automatically without any user interaction. The 5-day delay is an internally set database configuration parameter

Node recovery on vApp for offline node

Perform a node recovery from the ViPR Controller UI if a node (or minority nodes) is not able to rejoin the cluster when it's been down for more than five days.

Some nodes in a customer environment can be down for more than five days for unknown reasons. The nodes are then rejected from rejoining the cluster after being brought up in order to avoid introducing stale data. In the past, a manual recovery script was provided to support engineers.

You can now do a database service recovery from the ViPR Controller UI if a node (or minority nodes) is not able to rejoin the cluster when it's been down for more than five days.

- A System Administrator can know the node status and can drill down to get detailed information, and then go to the recovery page as necessary.
- A System Administrator can recover an offline node that cannot re-join the cluster after being offline or down too long. Both 2+1 and 3+2 environments are supported. Geo/Disaster Recovery environments are not supported.
- An alert notification is periodically sent via email to the System Administrator if a node is down for more than one day.
- Backward compatibility is maintained; there is no impact to other platforms (non-vApp, etc.).

The following ViPR Controller UI Help pages have been updated:

- **System > Node Recovery**
- **Dashboards > Health**

Diagutils

Diagutils is a built-in ViPR Controller utility program that enables you to easily download critical troubleshooting information, including a database dump, logs, ZK properties, ViPR system properties, etc.

Command

```
# /opt/storageos/bin/diagutils
```

Note

In order to execute the `diagutils` utility, the user should enable `root` console login and permit `root` `ssh` access, which are disabled by default.

Note

For a `diagutils` command log, refer to `/tmp/diagutils.out`.

Syntax

```
diagutils
<-all|-quick>
<-min_cfs|-all_cfs|-zk|-backup|-logs|-properties|-health>
[-ftp <ftp/ftps server url> -u <user name> -p <password>]
[-conf <config file>]
```

Options

-all

Includes the output gathered by options: `-backup` (with default backup name), `-zk`, `-logs`, `-properties`, `-health`, and `-all_cfs`.

Equivalent to

```
-backup -zk -logs -properties -health -all_cfs
```

`-ftp` and `-conf` are the only two other options allowed in conjunction with `-all`.

Note

`diagutils` collects a large amount of data when the `-all` option is used, and will take longer to run.

-quick

Includes the output gathered by options: `-zk`, `-logs`, `-properties`, `-health`, and `-min_cfs`.

Equivalent to

```
-zk -logs -properties -health -min_cfs
```

`-ftp` and `-conf` are the only two other options allowed in conjunction with `-quick`.

-min_cfs

Collect a minimum set of column families via the output of `dbutils list` and/or `cqlsh`

The default `cfs` list includes:

BlockConsistencyGroup, BlockMirror, BlockSnapshot, Cluster, ExportGroup, ExportMask, FCZoneReference, Host, Initiator, Network, NetworkSystem, ProtectionSet, ProtectionSystem, StorageProvider, StoragePool, StoragePort, StorageSystem, Vcenter, VirtualArray, VirtualDataCenter, VirtualPool, Volume.

-all_cfs

Collect all column families via the output of `dbutils list` and/or `cqlsh`

Note

`diagutils` collects a large amount of data when the `-all_cfs` option is used, and will take longer to run.

-zk

Collect `zk` jobs and queues through `zkutils`.

-backup [backup name]

Creates a new ViPR system backup/dump of DB and ZK through `bkutils`, which can be restored later.

If `[backup name]` is not specified, the timestamp will be used instead. If `[backup name]` already exists, the utility won't create a new backup, but will copy the existing backup into the archive.

-logs

Collect all system logs (`/var/log/messages`), and ViPR logs, including the rotated ones and orders in the last 30 days.

Note

Collecting orders is possible only with ViPR Controller 3.6 and above.

Note

ViPR Controller keeps a maximum of 30 days of logs. `diagutils` will take longer to run if there is a large amount of logs.

-properties

Collect system properties (version, node count, node names, etc.).

-health

Collect system health information (for example, node and service status, etc.), performance data of local node from top output.

-ftp <ftp/ftps server url> -u <user name> -p <password>

If specified, the output will be transferred to the external ftp/ftps server and removed from local storage after the transfer.

Note

It is recommended to always transfer the output to an ftp/ftps server, so as to retain space on ViPR nodes.

-conf <config file>

If specified, diagutils will use the settings in a configuration file (`/opt/storagesos/conf/diagutils-sample.conf`) instead of the default settings.

Examples

```
diagutils -all -ftp ftp://10.247.101.11/logs -u usera -p xxx
```

```
diagutils -quick
```

```
diagutils -min_cfs -zk -logs
```

Output

diagutils creates a compressed zip archive of several formatted output text files. When it completes running, it provides the name and location of the output file generated.

For example: "The output file is: `/data/diagutils-data/diagutils-20161104011601.zip`"

Exclusions and Limitations

- When majority nodes are down, and the database and ZK are not accessible, diagutils can only collect logs on live nodes and system properties.
- diagutils collects information only from the ViPR Controller system/cluster
- diagutils generates a warning if the `/data` directory is 50% full, and it will exit if the `/data` directory reaches 80% full.

ViPR Controller backup and restore enhancements

The backup status is now shown in the ViPR Controller Dashboard. Use this information to identify success, warnings, or failure and to configure your backup schedule.

This information displays with links to the **System > Data Backup and Restore** configuration screens:

- Last successful backup: The date and time of the last successful backup, and whether the backup was manual or scheduled.
- Last manual backup: The date and time of the last manual backup. "NA" indicates that there has not been a manual backup.
- Last scheduled backup: The date and time of the last scheduled backup. "NA" indicates that there has not been a scheduled backup.
- Next scheduled backup: The date and time of the next scheduled backup.
- Last upload status: The date and time of the last backup upload.

Refer to the ViPR Controller **Dashboards > Overview** UI Help page for additional information.

ViPR Controller log file collection enhancements

The log file collection process has been improved.

From the **System > Logs** page you can access log messages associated with ViPR Controller services and system events (alerts). The download button enables you to download a zip file containing the logs that correspond to the current filter setting. The new features allow you to:

- Configure the log retention time from the **System > General Configuration > Log Retention** page. The default setting is now 30 days. It can be adjusted to fewer if desired.
- Download log files when you have system administrator privileges. The system administrator can download order logs for a given tenant from a subpage after selecting the **Download** button.
- Collect order logs for all tenants when you have provider tenant administration privileges. In the past, log files could only be collected for the provider tenant.
- Include platform (operating system) messages when downloading or collecting the default logs files.

Note

This feature is not available from the user interface. Use the `diagutils` utility to accomplish this.

Log retention policy enhancement

You can now dynamically configure the number the number of days for log retention.

In the past, ViPR Controller logs (`dbsvc`, `controllersvc`, etc) were kept for seven days. Logs older than seven days were cleaned up automatically. This sometimes caused problems for troubleshooting and customer support.

You can now dynamically configure the log retention days from the ViPR Controller GUI.

- The default log retention days is 30.
- The permitted configurable range is from 7 to 30 days.

The following ViPR Controller UI Help page has been updated:

- **System > General Configuration > Log**

Storage driver auto installation

You may install storage drivers for use with storage systems or storage providers from the EMC ViPR Controller user interface.

The storage driver auto installation operation ensures the storage driver persists during concurrent restarting of all controller services. And it supports installation, uninstallation, and upgrade of the driver jar files.

After a driver install, uninstall, or upgrade operation, controller services are restarted concurrently to take the operation into effect. No reboot occurs. During the driver

installation process, the user interface is always available and the system is stable. The only adverse effect is that ongoing orders may fail.

Storage driver auto installation

Use the **Physical > Storage Drivers** page to install, upgrade, or delete storage drivers for storage systems used by ViPR Controller.

- Only active sites can process driver install, uninstall, or upgrade requests.
- All sites must be in a stable state. New driver operations are not allowed if any system services (`sys svc`) are offline. The `sys svc(s)` are needed to download the new driver file from the coordinator node.
- Standby site status should be Standby_Synced, not paused or degraded.
- All driver operations must be active or in use and not in the process of installing, uninstalling, or upgrading.
- There cannot be ongoing orders inside the controller service because controller services have to restart as part of the auto installation process.

The **Storage Drivers** page lists the storage driver and the following attributes.

Table 1 Storage Driver options and attributes

Column name	Description
Driver name	The name of the storage driver must be 50 characters or less. Only letters, digits, dashes, and underlines are allowed.
Version	The version of the storage driver.
Supported storage system	The storage driver may support both a storage system and a storage provider. When this happens, the Supported Storage Systems column has two items separated by a comma.
Type	The type of storage system. For example, File or Block.
Default Non-SSL Port	Set the default non-SSL port here.
Default SSL Port	Set the default SSL port here.
Status	In-use or Ready. When set to In-Use, the driver may be upgraded, but not deleted. When set to Ready, the driver may be upgraded or deleted. If a driver is In-use and you want to delete it, you must first delete the storage systems that use the driver. Then the status will change to Ready.
Actions	Upgrade or Delete

Multiple concurrent actions on export groups are permitted

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

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

In previous releases, only one actionable event per export group could run at a time. With this release, multiple actionable events for the same export group can be approved and run concurrently.

This change is supported by changes to the underlying API. The update operation for ExportGroup is now thread-safe, permitting multiple concurrent update calls for an export group.

Export group path adjustment

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

Configuration requirements and information

This operation requires the Tenant Administrator role in ViPR Controller.

ViPR Controller UI

The following ViPR Controller UI pages, and options provide the functionality to use the feature:

ViPR Controller UI Pages	Description
Service Catalog > Block Storage Services > Export Path Adjustment	<p>Allows you to adjust the export paths for a storage system within a host or cluster export group (including VPLEX with exports from both VPLEX clusters).</p> <p>Use this service to retain all the existing paths and add new paths. Or choose to recalculate the best paths based on current port usage data (including reducing the number of paths).</p> <p>This order updates the masking view and storage view, reconfigures zoning, and initiates a host rescan for discoverable hosts.</p> <p>You may suspend the operation and preview the impact of adding, removing, or changing paths. You then either resume or roll back the operation. If you choose the rollback option, no paths will be removed. However, any new paths that have been provisioned as a result of the order will remain in place.</p> <p>View the changes made to the paths by looking at Resources > Volumes.</p>

ViPR Controller CLI

Use the `path_adjustment_preview` and `path_adjustment` options with the ViPR Controller CLI `viprcli exportgroup` command to add, change, or delete export group port path assignments. Here are examples of the command syntax. The *ViPR Controller CLI Reference Guide* provides more information.

```
./viprcli exportgroup path_adjustment_preview -n losathost.emc.com -pr rrprj1 -va testva -ss 999334388821 -minp 1 -maxp 4 -ppi 4 -useex
```

In this example, `rrproj1` is the project. `testva` is the virtual array. A minimum of one initiator is attached to the host and it can have up to four paths (`ppi`) connected to storage system, `999334388821`. Existing paths will be used if possible because `useex` is included in the command.

```
./viprcli exportgroup path_adjustment -n
"testty11_20170315085001826" -pr proj1 -va va -ss 000196801612 -
minp 2 -maxp 4 -ppi 2 -useex -wait
```

In this example, `proj1` is the project. `va` is the virtual array. A minimum of two initiators are attached to the host and up to four paths (`ppi`) may be connected to storage system, `000196801612`. Existing paths will be used if possible because `useex` is specified in the command. The absence of `verbose` means this will return only the short description.

The presence of the `-wait` option means the provisioning process will be suspended until you have a chance to review the impact of the path adjustments. When you are ready to accept the results, you have to issue a command to resume processing of the path adjustment. You can resume activity with the `viprcli task` command and the Task ID.

Improved maintenance of export group inventory

In a continuing effort to provide more self-healing capabilities natively within ViPR Controller, stale export group inventory data is updated so that volumes, initiators, hosts, or cluster fields in the database are reset to "null."

Prior to the 3.6 version of ViPR Controller, users were allowed to create or leave behind empty export groups or export groups with compute resources and initiators tied to them. This is still allowed up to the point of removing Host(s) or initiators. After that, however, the ViPR Controller database is updated with current export group information.

Stale data could result from

- failure to follow through on an API execution path
- removal of devices without removing the export group itself
- prior failures, including failures of rollback

Ensuring consistent HLU assignments for cluster exports

If you choose the default (automatically assign) values when provisioning shared volumes to a cluster, ViPR Controller applies logic to ensure consistent unique HLU assignments to these volumes. This ensures the same host logical unit (HLU) will be generated for all hosts available in the cluster. (This logic does not apply for standalone exclusive host exports using the automatically assign value.) This feature proactively prevents a possible LUN violation error to be thrown from the array.

In the past, for shared or cluster exports, there was a chance that an underlying storage system could assign duplicate HLUs to some hosts in the cluster, which could cause problems. It is important to have uniform or consistent presentation of storage area network LUNs across all managed VMware ESX servers in a cluster environment in order to prevent data corruption. In the latest version of ViPR Controller, when the HLU is set to `-1` (default), ViPR Controller adjusts the HLU assignments so that no duplicate HLUs are issued.

Note

The logic supporting consistent HLU assignments occurs only for specific storage systems. The maximum allowed HLU value for each storage system is configurable. You may change the default maximum value through CLI or REST API calls. Follow the recommendations of the storage vendors when reconfiguring the HLU upper limit. (Review requirements for HBAs, types of hosts, storage arrays, or virtualization considerations). ViPR Controller cannot enforce upper HLU limits in the code. The Support Matrix and the e-Lab Navigator provide more information.

Storage system	Parameter	Default HLU maximum value
VMAX	controller_vmax_max_allowed_HLU	4096
XtremIO	controller_xtremio_max_allowed_HLU	2048
Unity	controller_unity_max_allowed_HLU	4096
VPLEX	controller_vplex_max_allowed_HLU	4096

Manual adjustment of the upper HLU limit

Use the ViPR Controller REST API or the command line interface (CLI) to adjust the upper HLU limit. Ensure you review the vendor-specific recommendations for your storage system before changing the limit. ViPR Controller cannot enforce changes in HLU upper limits.

Example of adjustment using the REST API

The maximum HLU keys are based on storage system type.

Storage system	Parameter	Default HLU maximum value
VMAX	controller_vmax_max_allowed_HLU	4096
XtremIO	controller_xtremio_max_allowed_HLU	2048
Unity	controller_unity_max_allowed_HLU	4096
VPLEX	controller_vplex_max_allowed_HLU	4096

In this example, the default value for maximum HLU for XtremIO is 2048. If you want to change it to 4096, use the REST API call to update the value:

```
URI: https://<<ipaddress>>:4443/config/properties
Method: PUT
Content-type: application/xml
Content:
```

```
<property_update>
  <properties>
    <entry>
      <key>controller_xtremio_max_allowed_HLU</key>
      <value>4096</value>
    </entry>
  </properties>
</property_update>
```

Example 1 Example of system configuration change using viprcli command

```
viprcli system set-properties -pn
controller_xtremio_max_allowed_HLU -pvf <file contains the new
value>
```

Changes made to the Remove Block Volume service

The service called **Remove Block Volumes**, in earlier releases, has been renamed **Unexport and Remove Block Volumes**. A new service, **Remove Block Volumes** has been added.

Use **Remove Block Volumes** to remove unexported block volumes or consistency groups.

Use **Unexport and Remove Block Volumes** to remove block volumes or consistency groups and related exports.

Unexport and remove block volumes

Deletes the block volume and orchestrates other changes (such as removing snapshots or full copies) if needed to successfully complete the deletion of the volume.

Before you begin

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

Procedure

1. Select **User > Service Catalog > Block Storage Services > Unexport and Remove Block Volumes**.
2. Select the **Project** from which to delete the volumes.
3. Select the block **Volumes** to be deleted.
4. Choose the **Deletion Type**.
 - A Full deletion removes the volume from ViPR Controller and from the underlying array. You can no longer use the volume. Use this choice when you want to free up more space.
 - An Inventory-only deletion removes the volume from ViPR Controller management, but does not remove the volume on the array. If you want to use the volume again, you must first ingest it back into ViPR Controller.

The inventory-only deletion type removes from ViPR Controller:

- volumes that have been exported to hosts or clusters
- volumes that may have been internally exported to VPLEX or RecoverPoint
- RP target or journal volumes
- volumes with full copies or snapshots
- volumes that have been partially ingested (for SRDF or RP)

Use this choice if you ingested a volume into the wrong virtual pool and want to back out your changes.

5. Click **Order**.

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

Remove block volumes

Remove unexported block volumes or consistency groups. If there are snapshots or other objects associated with the volumes or consistency groups, you must manually remove them before this service can succeed.

Before you begin

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

Procedure

1. Select **User > Service Catalog > Block Storage Services > Remove Block Volumes**.
2. Select the **Project** from which to delete the volumes.
3. Select the block **Volumes** to be deleted.

This service lists unexported volumes only. If you want to remove block volumes that have been exported (for example, that have snapshots or copies still associated with the volumes), use the **Unexport and Remove Block Volumes** service.

4. Choose the **Deletion Type**.
 - A Full deletion removes the volume from ViPR Controller and from the underlying array. You can no longer use the volume. Use this choice when you want to free up more space.
 - An Inventory-only deletion removes the volume from ViPR Controller management, but does not remove the volume on the array. If you want to use the volume again, you must first ingest it back into ViPR Controller.

The inventory-only deletion type removes from ViPR Controller:

- RP target or journal volumes
- Unexported volumes that have been partially ingested (for SRDF or RP)

Use this choice if you ingested a volume into the wrong virtual pool and want to back out your changes.

5. Click **Order**.

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

Switch affinity added to port allocation rules

During volume exports, ViPR Controller selects the switch storage ports. Switch affinity is a new feature added to the storage port selection rules that selects ports on the same network and the same physical switch as the host initiator, if possible. This feature can reduce traffic on the storage network, such as inter-switch link (ISL) traffic.

Set up switch affinity

Switch affinity considers the location of the initiator port and gives preference to storage ports that are local on the same physical switch as the initiator port. By allocating local ports whenever possible, inter-switch link (ISL) traffic is reduced.

When the **Switch Affinity Enabled** parameter value is **True**, the switch affinity logic is considered in the list of rules for port allocations. The default state is **True**.

Note that the switch affinity logic is preceded by other rules for port allocation. Even when **True**, switch affinity might be overridden by other rules.

Switch affinity is supported for:

- VMAX
- VNX
- XtremIO
- Unity
- VPLEX
- VPLEX Backend—(export backend volumes to VPLEX), the VPLEX backend ports are considered the initiators.

To set the value of the **Switch Affinity Enabled** parameter:

Procedure

1. Select **Physical > Controller Configs > Port Allocation**.
2. Select **Switch Affinity Enabled** from the drop-down list.
3. Click **Add**.
4. For **Scope Type**, select **Global**.
5. For **Scope Value**, select **Default**.
6. For **Value**, select **True** to enable switch affinity logic or **False** to disable it.

Results

After a volume export, messages in the `controllersvc.log` file show the switch affinity results, as follows:

- initiators with switch affinity
- initiators with partial switch affinity—Some of the allocated storage ports are connected to the same switch as the initiators for the compute resources, and others are not.
- initiators without switch affinity

Summary of port allocation rules

ViPR Controller allocates storage ports for exported volumes from hosts and clusters using rules that consider the type of host initiator and various user-defined settings.

Port allocation evaluates available ports using the following rules, in the order listed. Rules 1 through 5 do not apply to the first port in a request, but are applied to the subsequent (redundant) storage port allocations in the same export.

Table 2 Summary of port allocation rules

Number	Rule	Explanation
1	DirectorRule17	Skipped for first port. Applies to VMAX only. Disabled for VMAX3. Select additional storage ports so that the director index, when added to the index of the first selected storage port, adds up to 17.
2	Different director type	Skipped for first port. Applies to VPLEX only. Select additional storage ports with a director type (A or B) different from the type used by the first selected storage port.
3	Different engine	Skipped for first port. Applies to VMAX, VPLEX, XtremIO, and HDS. From the available storage ports after the second rule, select storage ports whose engine is different from the previously selected storage port.
4	Different director	Skipped for first port. Applies to all array types, except XtremIO. From the available storage ports after the third rule, select storage ports using a different director from the previously selected storage port director.
5	Different CPU	Skipped for first port. Applies to VMAX only. From the available storage ports after the fourth rule, select storage ports using a different CPU from the previously selected storage port CPU.
6	Switch affinity (if set to true)	Applies to the first port and subsequent selections, if switch affinity is enabled. Applies to all supported host initiators and VPLEX. From the available storage ports after the fifth rule, select storage ports that connect to the same switch as the initiator.
7	Different switch	Applies to VMAX, VNX, XtremIO, Unity, and VPLEX. When exporting from clusters, there might be multiple initiators connected to different switches. From the available storage ports after the sixth rule, select storage ports that connect to a different switch from the previously selected storage port. If the sixth rule does not find any ports, use the results of the fifth rule, and select storage ports that connect to a different switch from the previously selected storage port.
8	Metrics-based port selection	Algorithms determine ports to avoid based on collected metrics.

SRDF Metro on VMAX supports adding new storage while solution remains active

In previous releases, adding new RDF pairs into an active group without suspending the existing pairs was not possible in an SRDF metro environment. With this

enhancement, you no longer need to suspend and disrupt the high availability of existing SRDF metro pairs to add more storage.

When SRDF pairs need to be added to a non-empty ViPR Controller project (SRDF Group), two scenarios are possible:

Create new SRDF pairs. The new RDF volume pairs will not have any data on them.

In this case, use the ViPR Controller Create Volume operation to add the new volume pairs. If the associated virtual pool contains Asynchronous, Synchronous, Active (METRO) SRDF protection, the Create Volume operation adds the new volume pairs without suspending other active pairs using that SRDF Group.

Note

The Create Volume operation, in the case of a non-empty SRDF Active (METRO) group, uses the `format` option (`createpair -format`) when adding the new pairs. The `format` option removes any existing data on the new pairs. ViPR-Controller does not check for the existence of data on the new pairs. In the case of an empty SRDF Metro group, the Create Volume operation does not use the `format` option.

Convert to SRDF protection. The RDF pairs created will have existing data on the volumes. This existing data needs to be preserved.

In this case, use the ViPR Controller Change Virtual Pool operation on existing volumes to add SRDF protection. As part of the Change Virtual Pool operation, ViPR Controller suspends all the pairs belonging to the SRDF Group before proceeding with the required task. This operation preserves data on the existing volumes being upgraded to SRDF protection. This is true for SRDF Active (METRO) protection as well.

Improvements to discovery processing

Improvements were made in the underlying processing of discovery operations and associated tasks.

These improvements result in more efficient processing, fewer hung or stopped tasks, fewer instances of overloaded queues and permanent locks that need to be manually released.

No changes were made to any interface or to the default discovery scan interval settings. No changes to your existing procedures are required, and no actions are required to accommodate or enable these changes.

Administrators might notice the following differences in discovery operations and associated tasks.

- The job scheduler does not schedule a job if an identical job is already running or scheduled. Previously, if scanning or discoveries took longer than a configured scan interval setting, the scanning jobs would overlap and queues would become overloaded.
- The time between scans is calculated using the time that the previous job was picked from the queue. Previously, the time was calculated using the time that the previous job was placed on the queue, potentially causing actual runs to occur more frequently than the configured scan interval setting.

- The scanning operation was split into several operations. Each interface type has its own operation with unique operation IDs. For example, one operation scans only SMI-S providers, another operation scans all Vplex providers, and so on. Previously, a problem accessing any provider caused the entire discover process to hang.
- Architectural changes were made to reduce the likelihood of stranded pending tasks, especially for failed operations.

HTTP protocol can be used when AllowUnencrypted flag is disabled for WinRM Windows discovery

When discovering Windows hosts where encrypted traffic is running under the HTTP protocol, you no longer have to adjust the WinRM setting for the AllowUnencrypted flag or install HTTPS certificates on all Windows hosts.

This information has been added to the **Physical > Hosts > Add a Host** online help topic.

SHA-256 support

A `strong_ciphers` configurable parameter has been added to ViPR Controller to support use of SHA-256.

Use either `viprcli` or API commands to set the parameter. You will need to create a property values file with a comma-separated list of the encryption algorithms you want ViPR Controller to support. For example,

```
TLS_RSA_WITH_AES_128_CBC_SHA256,TLS_RSA_WITH_AES_128_CBC_SHA
```

In this `viprcli` example, the file, `propvalue`, holds the list of new values to be set:

```
./viprcli system set-properties -propertyname strong_ciphers
-propertyvaluefile propvalue
```

Note

Both AES-128 and SHA-256 encryption algorithms are supported by default.

Vblock enhancements

The operations that affect Vblock catalog services have been enhanced to prevent data unavailability incidents.

Enhancements include:

- Displaying the UCS service profile and UCS blade mapping on the **Physical > Hosts** page.
- Improved handling of mixed clusters in Vblock catalog services.
 - Allows the "remove host from cluster" operation only if the host is part of the Vblock

- Allows "decommission cluster" operation only if all hosts are in the Vblock
- Use **Clusters > Edit hosts** to manage operations for non-Vblock hosts
- Detection of changes made to UCS Service Profile, for example, when the profile is moved to a different blade.
 - Manual changes in the UCS are followed by a UCS rediscovery in ViPR Controller. After rediscovery, the compute element association can be verified on the **Physical > Hosts** page.
 - ViPR Controller detects dissociation of service profile for a host from compute element or blade
 - ViPR Controller detects movement of UCS service profile for a host from one compute element or blade to another
 - Prevents the provisioning or decommissioning of operations on a UCS that fail discovery. Provides information to resolve conflict so rediscovery can be successful.
- Providing new validations during UCS discovery.
 - Detects the UUID of service profile changes including duplicate UUIDs found during discovery or provisioning
 - Generates error messages and log files indicating which objects are in conflict.
- Providing reliable rollback for Vblock catalog services.
- Providing the ability to set boot volume HLU as 0.
 - The boot volume HLU field is added to these services:
 - Provision cluster
 - Add Host(s) to cluster
 - Provision bare metal cluster
 - Add bare metal host to cluster
 - Boot volume HLU value defaults to 0 but you can override
 - The valid range of values are 0 to 255. Use 0, which is the default and the best practice, or provide an HLU which will not result in a LUN violation. If the HLU value results in a LUN violation, then the order fails. You may correct the HLU and resubmit.

Note

Entering -1 to automatically assign the HLU to a boot volume is not supported.

- If the HLU value is already in use or will cause a LUN violation on a VMAX array, the order fails and rolls back. You can then resubmit the order with a different HLU.
- Allowing you to decommission the boot volume on VPLEX . Prior to version 3.6, a boot volume provisioned on VPLEX could not be decommissioned when deactivating a host. Now the boot volume is decommissioned when the host is deactivated.

Isilon file enhancements

This release offers the following Isilon file enhancements:

- [Root squash for Isilon NFS Exports](#)

- [Replicate Isilon file system configurations during failover/failback](#)
- [Use dedicated SCZs for SyncIQ replication workload](#)

Root squash for Isilon NFS Exports

ViPR Controller now supports "root squash" for Isilon NFS exports. Root Squash is available only to the user who is logged in to ViPR Controller.

ViPR Controller creates NFS exports with `root` defined as the root user. This is equivalent to a `No_Root_Squash` condition:

- Remote `root` users are treated as local `root`. Requests from `root` are not mapped to the `anonymous` user and group ID.
- `root` on the client machine has the same level of access to the files on the system as `root` on the server. This can have serious security implications and is not best practice.

By default, ViPR Controller squashes `root` to user `nobody`. This is equivalent to a `Root_Squash` condition:

- All requests from the `root` user are translated or mapped to the `anonymous` user.
- By default, any file request made by the `root` on the client machine is treated as if it is made by user `nobody` on the server.

You can change root mapping from `nobody` to the ViPR Controller login user to get root mapping to the login user during an NFS export.

The following ViPR Controller UI Help pages have been updated:

- **Catalog > View Catalog > File Storage Services > Create File System and NFS Export**

Replicate Isilon file system configurations during failover/failback

Isilon file system configurations, consisting of NFS exports, export rules, NFS ACLs, and quota configurations, are replicated during Failover and Failback operations.

- During a Failover operation, the source file system's NFS export, export rules, and NFS ACLs are replicated to the target file system.

Note

Quota directories and snapshots are not replicated during a Failover operation. Sub-directory exports are not replicated if data synchronization has not taken place.

During the first Failover, all configurations are replicated to the target file system, but during subsequent Failovers only the changes are replicated.

- During a Failback operation, the target file system's NFS export, export rules, NFS ACLs and quota configurations are replicated to the source file system. Only the changes are replicated.

The following EMC ViPR Controller Guide has been updated:

- "System Disaster Recovery, Backup and Restore Guide", Chapter 2, "System Disaster Recovery."

Use dedicated SmartConnect Zones for SyncIQ replication workload

For Isilon systems, you can now specify a dedicated SmartConnect Zone for data access.

It is recommended that there be a dedicated SmartConnect Zone for data access, per SyncIQ Policy, rather than sharing the SmartConnect Zone between policies. After a port is registered as a Disaster Recovery port, the dedicated port is picked up for SyncIQ activities.

The following ViPR Controller UI Help page has been updated:

- **Physical > Storage Systems > Storage Ports**

The following EMC ViPR Controller Guide has been updated:

- "User Interface Virtual Data Center Configuration Guide", Chapter 2, "Adding and Configuring Physical Assets."

File Protection Policy templates

Isilon policy management provides the ability to define file protection policies at different directory levels.

Isilon policy management provides:

- The ability to define different policies (Replication and Snapshot) at different levels (file system, vPool, or project).
- Simplified policy management that make it easier for both IT administrators and Isilon users to manage Isilon policies.
- File protection policies that are multi-tenant aware and have controlled access.
- The ability to validate policies with similar snapshot/replication parameters and provide recommendations.
- Seamless policy based fileshare provisioning.

Note

Scheduled policies from earlier ViPR Controller versions are migrated to File Protection Policy Templates during an upgrade. They can be found under **Virtual > File Protection Policy Templates** and have the following **Description**: "Policy created from virtual pool <name> replication while system upgrade."

Note

With the introduction of File Protection Policy Templates, the following services became redundant and were deprecated:

- **Catalog > View Catalog > File Protection Services > Create Replication Copy**
- **Catalog > View Catalog > File Protection Services > Remove Replication Copies**

These services will not be part of a fresh installation. However, if you are performing an upgrade installation, you will need to remove these services by going into:

- **Catalog > Edit Catalog > File Protection Services > Create Replication Copy**
 - **Catalog > Edit Catalog > File Protection Services > Remove Replication Copies**
-

The following new ViPR Controller UI Help pages have been created:

- **Virtual > File Protection Policy Templates**
- **Virtual > File Protection Policy Templates > Add > Create File Protection Policy Template**
- **Virtual > File Protection Policy Templates > Assign Policy**

The following ViPR Controller UI Help pages have been updated:

- **Virtual > File Virtual Pools > Add/Edit**

The following ViPR Controller UI Help pages have been deprecated:

- **Tenants > Schedule Policy**
- **Catalog > View Catalog > File Protection Services > Create Replication Copy**
- **Catalog > View Catalog > File Protection Services > Remove Replication Copies**

The following EMC ViPR Controller Guides have been updated:

- "User Interface Virtual Data Center Configuration Guide", Chapter 5, "Creating and Configuring Virtual Assets."
- "Service Catalog Reference Guide", Chapter 4, "ViPR Controller File Storage Services."

The following ViPR Controller CLI commands have been added:

- `vipercli filepolicy assign`
- `vipercli filepolicy create`
- `vipercli filepolicy delete`
- `vipercli filepolicy list`
- `vipercli filepolicy show`
- `vipercli filepolicy unassign`
- `vipercli filepolicy update`

File Protection Policy Templates

Use the **Virtual > File Protection Policy Templates** page to view, create, edit, and delete file protection policies for Isilon. You can create file snapshot policies and file replication policies and apply them at the Virtual Pool, Project, or File System level. Policies are multi-tenant aware and have controlled access.

The ViPR Controller roles required for file protection policy management are as follows:

Table 3 File Protection Policy role requirements

Function	Role requirements
Create/edit policy template	System Administrator, System Monitor
Create vPool	System Administrator, System Monitor
Policy assignment at vPool level	System Administrator

Table 3 File Protection Policy role requirements (continued)

Function	Role requirements
Policy assignment at Project level	System Administrator AND System Monitor
Policy assignment at File System level	System Administrator AND System Monitor

The **File Protection Policy Templates** page lists the file protection policies and the following attributes:

Table 4 File Protection Policy attributes

Column name	Description
Selection column	Select one or more rows and click Delete to delete the file protection policies. Select one or more rows and click Assign Policy to assign the policies to one or more virtual pools, projects, or file systems. Select one or more rows and click Unassign Policy to unassign the policies.
Name	The name of the file protection policy.
Type	The file protection policy type: File Snapshot Policy, File Replication Policy.
Applied At	The level at which the policy is applied: Virtual Pool, Project, File System.
vPool(s)	The Virtual Pools to which the policy is assigned.
Project(s)	The projects to which the policy is assigned.
Priority	The policy priority: High, Normal.
Description	A description of the policy.

Note

Scheduled policies from earlier ViPR Controller versions are migrated to File Protection Policy Templates during an upgrade. They can be found under **Virtual > File Protection Policy Templates** and have the following **Description**: "Policy created from virtual pool <name> replication while system upgrade."

Creating a file protection policy template

Click **Add** on the **Virtual > File Protection Policy Templates** page to open the **Create File Protection Policy Template** page and create a file protection policy. You can create a file snapshot policy or a file replication policy and apply it at the Virtual Pool, Project, or File System level.

The ViPR Controller roles required for file protection policy management are as follows:

Table 5 File Protection Policy role requirements

Function	Role requirements
Create/edit policy template	System Administrator, System Monitor
Create vPool	System Administrator, System Monitor
Policy assignment at vPool level	System Administrator
Policy assignment at Project level	System Administrator AND System Monitor
Policy assignment at File System level	System Administrator AND System Monitor

Procedure

1. Specify information for the following options:

Option	Description
Type	Select: File Snapshot Policy, File Replication Policy.
Name	Specify a name for the file protection policy.
Description	Specify a description for the file protection policy.
Snapshot Name Pattern	File Snapshot Policy only. The default pattern is: {Cluster}_{vNas}_{VPool}_{Policy_TemplateName}_{%Y-%m-%d-%H-%M}, and is not editable.
Replication Type	File Replication Policy only. Select: Remote, Local.
Copy Type	File Replication Policy only. The default is Asynchronous.
Priority	File Replication Policy only. Select High, Normal. When a policy needs precedence over other policies, select High.
Worker Threads	File Replication Policy only. Select a value from 3 to 10. The default is 3. Increase the number of worker threads if there is a large amount of data to be replicated.
Run every	Specify a Frequency and a value. For Frequency, select Minutes, Hours, Days, Weeks, Months. For Weeks, specify Day of Week. For Months, specify Day of Month.
Run/Start at	Specify the time to start the run in HH:MM format.
Snapshot Expiration	File Snapshot Policy only. Select Never expires, Snapshot expires. For Snapshot expires, specify Hours, Days, Weeks, Months, and a value.
Apply Policies at	Select Virtual Pool, Project, File System.

2. Click **Save** to create the new file protection policy.

Assigning a file protection policy template

Click **Assign Policy** on the **Virtual > File Protection Policy Templates** page to open the **Assign Policy** page for the selected policy. You can assign the policy at the Virtual Pool or Project level.

The ViPR Controller roles required for file protection policy management are as follows:

Table 6 File Protection Policy role requirements

Function	Role requirements
Create/edit policy template	System Administrator, System Monitor
Create vPool	System Administrator, System Monitor
Policy assignment at vPool level	System Administrator
Policy assignment at Project level	System Administrator AND System Monitor
Policy assignment at File System level	System Administrator AND System Monitor

Procedure

1. Specify information for the following options:

Option	Description
Policy	The name of the policy selected in the File Protection Policy Templates page.
Apply Policies at	The policy level specified in the File Protection Policy Templates page: Virtual Pool, Project, File System. Not editable.
Virtual Pool	Specify a virtual pool for the Project or Virtual Pool policy. A file snapshot policy can be applied to multiple vPools; a file replication policy can be applied to one vPool.
Projects	For Project policy only. Specify one or more projects for the Project policy.
Source Virtual Array	For Replication policy only. Specify the source virtual array for the Replication policy.
Target Virtual Array	For Replication policy only. Specify the target virtual array for the Replication policy.

2. Click **Save** to assign the file protection policy.

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

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>).