

EMC ViPR SRM 3.5.1 Performance and Scalability Guidelines

Abstract

The performance and scalability guidelines outlined here are based on tests run on EMC ViPR SRM configured in a production-like environment. Both live and simulated configuration items were used to measure the performance and scalability of ViPR SRM.

September 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.

P/N 302-001-060 Rev 01

Table of Contents

Overview	4
Audience.....	4
Purpose	4
Infrastructure planning guidelines	4
Deployment Planner tool	4
Guideline P1- Assumptions to determine the configuration size	5
Recommended deployment type based on inventory size and metric count.....	5
Resource description and specification of deployment types.....	7
Guideline P2- Increase default disk space	8
Guideline P3- Deploy ViPR SRM on ESX host with right SPEC score	8
Guideline P4- Plan datastore read/write latency for application deployment	8
Guideline P5- Maintain network latency	8
Guideline P6- Plan for EMC VMAX Collector configuration	9
Guideline P7- Plan for VNX Collector configuration	9
Guideline P8 – Plan for Generic Host Collector.....	10
Guideline P9– Plan for Load-Balancer-Arbitrator and Load-Balancer-Connector Configuration	10
Deployment and configuration guidelines	10
Guideline D1- Consider CPU utilization of EMC Host Interface resources on production host	10
Guideline D2- Consider additional hardware requirements for enabling LUN and Path	
metrics of EMC PowerPath	11
Guideline D3- Compliance policy configuration guidelines	12
ESX/ESXi hardware specifications used for testing	12
Finding the SPEC score	13
Appendix A – Metric count and percentage change for ViPR SRM SolutionPacks.....	15
Metric count	15
Percentage change in collection	22
What has changed in 3.5.1 Guidelines?	23

Overview

The guidelines for an optimized performance and scalability are related to planning the hardware and virtual environment, disk space, network latency, and deployment and configuration.

- Planning guidelines help determine the right ESX hardware and the number of virtual machines needed for a specific data center environment.
- Disk space guidelines help prepare and configure datastore capacity before ViPR SRM deployment.
- Network latency guidelines help identify the proximity and placement of ViPR SRM virtual machines with respect to the configuration items.
- Deployment and configuration guidelines assist with deployment and post-deployment.

ViPR SRM is fully deployable on non-ESX hosts. EMC ViPR SRM Performance and Scalability Guidelines does not address non-ESX host deployment scenarios.

EMC ViPR SRM Performance and Scalability Guidelines do not address all of the SolutionPacks contained in ViPR SRM.

Audience

This article is for ViPR SRM installers, administrators or anyone who manages the ViPR SRM application.

Purpose

After reading this article, you will have guidelines to follow to ensure you can configure ViPR SRM for optimal performance based on the scale of your environment.

Infrastructure planning guidelines

Use these guidelines to plan the infrastructure such as the right ESX hardware, number of virtual machines needed, and other details described in the tables to ensure an optimum performance and scalability.

Deployment Planner tool

The ViPR SRM Deployment Planner tool can be used in conjunction with these guidelines to help scope specific environment requirements. Using known data about the environment, the tool will estimate ViPR SRM hardware needs, provide metric

counts (with growth factor), and include pointers to performance and scalability guidelines specific to the environment.

Access the Deployment Planner tool on the Downloads for ViPR SRM page at:

https://support.emc.com/downloads/34247_ViPR-SRM

Guideline P1- Assumptions to determine the configuration size

The size of a particular ViPR SRM configuration depends on multiple factors such as count and size of various configuration items, arrival rate of alerts, and so on. Since these variations could be significant in real-time datacenters, certain assumptions have been made to arrive at few examples of deployment configuration. The major assumptions are:

- The ratio between EMC VNX Block and EMC VMAX arrays is 60% to 40%
- The ratio between Brocade FC and Cisco MDS/Nexus switches is 50% to 50%
- All NAS arrays are of EMC VNX File array type only
- 50% of all the ESX and non-ESX hosts may contain EMC PowerPath, collecting default PowerPath metrics
- All these ViPR SRM SolutionPacks are planned for the deployment:
 - SolutionPack for Brocade FC Switch
 - SolutionPack for Cisco MDS/Nexus
 - SolutionPack for EMC VMAX
 - SolutionPack for EMC VNX
 - SolutionPack for VMware vCenter
 - SolutionPack for Physical Hosts
- Compliance rules are scheduled during period of low user activity
- Alerts processing at the rate of 33,000 alerts per day

The ViPR SRM deployment can be scaled-out by deploying additional Backend and Collector VMs on top of the standard deployment.

Deployment configuration based on inventory size and metric count.

The recommended ViPR SRM deployment configuration types at different scale levels, configuration item count, and metric count are described here.

Recommended deployment type based on inventory size and metric count

Inventory	Deployment Configuration
-----------	--------------------------

Scale Level	Configuration Items	Configuration Item Count	Estimated Metric Count (in Millions) ²	Deployment Type	VM Count ¹
Minimal	Storage Arrays/LUNs	3/32,300	1	All-in-one	1
	NAS/File Systems	3/3,460			
	Hosts/Host Devices	40/2,500			
	Switches/Ports	7/780			
	ESXs/VMs	26/520			
0	Storage Arrays/LUNs	14/132,000	4.5	Standard Deployment:	
	NAS/File Systems	14/14,643		SRM 4-VM vApp	4
	Hosts/Host Devices	330/21,120		Additional Backend	0
	Switches/Ports	29/3,300		Additional Collector	5
	ESXs/VMs	110/2,200			
1	Storage Arrays/LUNs	29/264,000	8.5	Scale-out Deployment:	
	NAS/File Systems	29/29,286		SRM 4-VM vApp	4
	Hosts/Host Devices	660/42,240		Additional Backend	2
	Switches/Ports	57/6,600		Additional Collector	13
	ESXs/VMs	220/4,400			
2	Storage Arrays/LUNs	43/396,000	13	Scale-out Deployment:	
	NAS/File Systems	43/43,930		SRM 4-VM vApp	4
	Hosts/Host Devices	990/63,360		Additional Backend	2
	Switches/Ports	86/9,900		Additional Collector	19
	ESXs/VMs	330/6,600			
3	Storage Arrays/LUNs	57/528,000	17	Scale-out Deployment:	
	NAS/File Systems	57/58,573		SRM 4-VM vApp	4

¹ The number of Collectors represented is approximate requirement for the number of configuration items specified. To account for factors such as geographic distributions of data centers, your environment may require additional Collector Virtual Machines.

² Estimated Metric Count: indicates the approximate number of metrics that ViPR SRM database(s) will need to store.

Inventory				Deployment Configuration	
Scale Level	Configuration Items	Configuration Item Count	Estimated Metric Count (in Millions) ²	Deployment Type	VM Count ¹
	Hosts/Host Devices	1,320/84,480		Additional Backend	3
	Switches/Ports	114/13,200		Additional Collector	26
	ESXs/VMs	440/8,800			
4	Storage Arrays/LUNs	72/660,000	22	Scale-out Deployment:	
	NAS/File Systems	72/73,216		SRM 4-VM vApp	4
	Hosts/Host Devices	1,650/105,600		Additional Backend	4
	Switches/Ports	143/16,500		Additional Collector	31
	ESXs/VMs	550/11,000			

To configure Additional Backend and Additional Collector VMs refer to the article [“Using the 1 VM vApp to Scale Out EMC ViPR SRM”](#).

Resource description and specification of deployment types

Deployment Type	Purpose of VM	vCPUs	Memory (GB)	Disk Space(GB)	VM Count
All-in-one	All components in a single VM	4	32	600	1
Standard Deployment (SRM 4-VM vApp)	Frontend	4	16	120	1
	Backend-Primary	4	24	600	1
	Backend-Secondary	4	24	1,200	1
	Collector	4	16	120	1
Scale-out Deployment (Requires Standard Deployment)	Additional Backend	4	24	1,200	1
	Additional Collector	4	8	120	1

Notes:

- The disk space specification in this table assumes default retention policies.
- The CPU, disk space, and memory requirements do not change for the mode of installation, so binary and vApp installations require the same resources.
- The Frontend is tested for 10 concurrent users

Guideline P2- Increase default disk space

After deploying ViPR SRM standard deployment (SRM 4-VM vApp), the disk space must be increased for both primary and secondary backends. The default disk space on the primary and secondary backends is 132 GB. Use the above table to estimate the disk space requirements for your environment.

The default disk space on Frontend is 120 GB. Increase the disk space on the Frontend as per the following table.

Scale level	Disk Space (GB)
0	120
4	320

Guideline P3- Deploy ViPR SRM on ESX host with right SPEC score

All recommendations assume there is no resource contention with other VMs that reside on the same ESX server of data store. Do not over allocate virtual resources as to cause resource contention.

Use only ESX servers with the SPEC rate of SPECint_Rate_base2006 of 350 or higher to deploy ViPR SRM.

Use Finding the SPEC score to determine the SPEC score for the ESX host.

Guideline P4- Plan datastore read/write latency for application deployment

To ensure the optimum performance of ViPR SRM, the average read/write latency for datastores of Fibre Channel disks should be 10 ms or below and of SATA II disks should be 14 ms or below. The disk read/write latency exceeding the specified values degrades the end user response time. The read/write latency metrics can be viewed within the VMware vCenter Performance view or by using ESXTOP.

The average IO requests on the Datastore in the test environment is 20 requests per second.

Guideline P5- Maintain network latency

While deploying ViPR SRM, the network latency must stay within the limits for the specified components as mentioned in the table.

Components	Maximum Network Latency
Between Frontend and Backend	5 ms
Between Backend and Collector	300 ms

Between VMAX Collector host and Solutions Enabler or Unisphere host	20 ms ²
Between VNX Collector host and VNX provider host	20 ms (2)
Between RSC (Host) Collector and Host	20 ms (2)
Between VMware Collector and VCenter	20 ms (2)
Between SNMP collector and FC Switch	100 ms

Guideline P6- Plan for EMC VMAX Collector configuration

The maximum tested scalability for a ViPR SRM VMAX Collector collecting performance and capacity statistics from a remote EMC Solutions Enabler provider and Unisphere for VMAX is 65,000 VMAX LUNs.

The following must be noted when deploying a VMAX Collector:

- Six gatekeepers are required for each VMAX array whether the provider is dedicated for ViPR SRM or shared with other management software.
- Do not exceed 20 ms of latency between VMAX Collector host and Solutions Enabler provider or Unisphere for VMAX host.
- The VMAX Collector with default heap size can manage a maximum of 30,000 VMAX LUNs.
- The VMAX Collector heap size should be increased to 4 GB when collecting performance and capacity statistics for more than 30,000 VMAX LUNs.
- Recommended VMAX polling intervals (seconds)

Unisphere Performance	7200
SMI Performance	1800
Topology³	10800

Guideline P7- Plan for VNX Collector configuration

The maximum tested scalability for a ViPR SRM VNX Collector collecting performance and capacity statistics is 15,000 devices.

The following must be noted when deploying a VNX Collector:

- The VNX collector with default heap size can manage a maximum of 5,000 VNX LUNs
- The VNX collector heap size should be increased to 4 GB when collecting performance and capacity statistics for more than 5,000 VNX LUNs.

² If latency between the provider and the collector is more than 20 ms, consider increasing the polling interval or decreasing the number of configuration items per collector.

³ VMAX Topology polling interval can be modified by editing the configuration file “xmlcollector-vmx-topo.xml” under the folder “/opt/APG/Collecting/XML-Collector/emc-vmx/conf” on the VMAX Collector Host. Modify the polling interval to <polling-interval>108000</polling-interval>.

- The recommended VNX Polling intervals (seconds)

Performance	1200
Topology	3600

Guideline P8 – Plan for Generic Host Collector

The maximum tested scalability for a SRM Generic Host Collector collecting performance and capacity statistics without PowerPath is 1000 hosts.

Guideline P9– Plan for Load-Balancer-Arbiter and Load-Balancer-Connector Configuration

Use the below table to increase the heap size and MaxPermSize for Load-Balancer-Arbiter on the Primary Backend Host and Load-Balancer-Connector on the Collector Hosts

	Heap Size (GB)	MaxPermSize (MB)
Load-Balancer-Arbiter	8	256
Load-Balancer-Connector	4	256

Deployment and configuration guidelines

Follow these guidelines while deploying and configuring ViPR SRM for an optimum performance and scalability of CPU utilization, JVM heap size, CPU, and memory resources.

Guideline D1- Consider CPU utilization of EMC Host Interface resources on production host

The EMC host interface resources on a production host is minimal. The CPU consumption is nearly 0% when the Solution Pack for EMC Host Interface Collector is not polling it.

The table shows the resource utilization of the EMC Host interface when Solution Pack for EMC Host Interface Collector is polling on a Windows host with 64 devices and 4 paths and the following specifications:

- 4x3.0 GHz processors with 2 GB of RAM
- Windows 2008 Enterprise Edition SP1

Host interface footprint

Default Polling Type	Frequency	Duration (mm:ss)	Average CPU Utilization	Required Disk Space
Performance Collection	Every 5 minutes	< 1 second	9%	150 MB
Topology and Capacity Collection	Every 15 minutes	5:30	9%	150 MB

The maximum tested scalability for a SRM Host Interface is a maximum of 40 hosts.

Guideline D2- Consider additional hardware requirements for enabling LUN and Path metrics of EMC PowerPath

The PowerPath LUN (device) and Path metrics are turned-off by default. Before enabling these collections, refer the following table.

Assumptions:

- (a) 50% of PowerPath enabled Hosts/ ESX Servers have only LUN and Path metrics enabled.
- (b) 10% of PowerPath enabled Hosts / ESX Servers have all PowerPath metrics (LUN, Path and Range Bound) enabled.

Host Type	Avg. # Devices per Host	Avg. # I/O Paths per Host	# Metrics collected	Total Metrics Collected per host	# Hosts @ Scale level 0	Total Metrics collected @ Scale level 0
Physical Host (LUN and Path metrics enabled)	64	4	38	9,728	165	1,605,120
Physical Host (LUN, Path and Range-bound metrics enabled)	64	4	77	19,712	17	335,104
VMWare ESX Server (LUN and Path metrics enabled)	256	4	38	38,912	55	2,140,160
VMWare ESX Server (LUN, Path and Range-bound metrics enabled)	256	4	77	78,848	6	473,088

It is recommended to use additional Collector for all Hosts that require PowerPath metrics collection. Plan for additional backend based on Scale level and the total metric count.

Guideline D3- Compliance policy configuration guidelines

Both the compliance rule-validation and database maintenance script execution consume considerable CPU and memory resources. The nightly indexing of properties store starts at 8:00 P.M. (UTC) and takes a maximum of 4 hours.

It is best to schedule the compliance rule-validation during periods of low user activity and not during nightly database maintenance script execution. Do not schedule all compliance policies to start simultaneously, because some policy validation may consume considerable CPU and memory resources, which causes delay in user response times. Schedule the policies to run at least 30 minutes apart.

Example: Scheduling policy run

Consider the nightly database maintenance is complete by midnight, then schedule

- Array configuration policy to run at 12:00 A.M.
- EMC Support Matrix policy to run at 12:30 A.M.
- Path Management policy to run at 1:00 A.M. and so on.

The following section details configuration of Path based rules for resource optimization.

The rules that are considered resource intensive are:

1. I/O Path Redundancy
2. I/O Path Redundancy for VPLEX Frontend
3. Missing Path

Tips to configure Path based rules:

1. Configure Path based policies for Hosts that are considered critical by creating Scope specific to the host or group of hosts
 - a. Use the below table as an example for “Create Scope” dialog

Name	LINBG Hosts
Criteria	Devtype='Host' & device='LINGB%'

2. Multiple Scopes can be created for various groups like Hypervisors, AIX and so on.

ESX/ESXi hardware specifications used for testing

The table describes the hardware specification of the ESX/ESXi servers used for testing.

Component	Specification
SPECint_Rate_base2006 score	583
Manufacturer	Cisco
Model	C260-BASE-2646
Processor type	Intel® Xeon® CPU E7-2870 @ 2.40 GHz
License	VMware vSphere 5.1 Enterprise
NIC type	1 Gigabit, Full Duplex NIC
Processor socket	2
Cores per socket	10
Total number of cores	20
Hyper threading	Active
Memory	512 GB
Number of NICs	9

The drive and RAID configurations were provisioned to the ViPR SRM virtual appliance.

Finding the SPEC score

If you are planning to compare your host performance with the equipment being used in the lab to measure the performance and scalability of ViPR SRM, information provided on the SPEC website (www.spec.org) helps.

There are two distinct CPU benchmark results available on this web site. Hosts using old technology (for example, single-core hosts) are benchmarked using SPEC CPU2000. This benchmark was retired in February 2007 and a new benchmark (SPEC CPU2006) is used by most hardware vendors.

1. Navigate to the SPEC website by typing www.spec.org in the web browser.
2. Select Result > CPU2006 > Search CPU2006 Results from the top navigation bar.
3. Select SPECint2006 Rates for Available Configurations.
4. Select Advanced for Search Form Request and click Go!.
5. Fill out the form with appropriate information. For example, Hardware Vendor, System, # Cores, # Chips, # Cores Per Chip.

In SPEC nomenclature, the term for Processor is Chip.

6. Click Fetch Results to submit the search. The value listed in the Baseline column for the host is the value that is compared to the minimum specified in this document.

Appendix A – Metric count and percentage change for ViPR SRM SolutionPacks

The following table can be used to determine the overall metrics for an environment. For example a single EMC VMAX array with 16000 LUNS, 444 Disks, 8000 Thin Devices will produce 232000 metrics. For each LUN, 5 metrics are collected, 11 metrics for a disk and 5 metrics for each Thin Device are collected. An overall 5% is factored for the rest of the metrics.

Metric count

Solution Pack	Category	Metric in 3.5.1	Typical Example
EMC VMAX	Unique metrics	144	
	Total metric count = Sum (Scaling factor * metric count)	VMAX Array * 645 ^a	1 * 645
		Disks * 11	444 * 11
		Devices * 21	16,000 * 21
		Directors * 45	16 * 45
		Pools * 9	40 * 9
	Total metrics		342,609
VNX-Block	Unique metrics	149	
	Total metric count = Sum (Scaling factor * metric count)	VNX Array * 587 ^a	1 * 587
		Disks * 18	213 * 18
		LUNs * 35	5,048 * 35
		Pools * 27	7 * 27
		Initiator * 1	5,058 * 1
	Total metrics		186,348
VNX-File	Unique metrics	97	
	Total metric count =Sum (Scaling factor * metric count)	VNX File * 1560 ^a	1 * 1560
		File System * 14	1000 * 14
		Storage Pool * 4	18 * 4
		Disks * 2	100 * 2
		Data Movers * 47	35 * 47
		Check Point *10	210 * 10
	Total metrics		19,577
EMC XtremIO	Unique metrics	76	
	Total metric count = Sum (Scaling factor * metric count)	Array * 852 ^a	1 * 852
		Volumes * 7	8,192 * 7

		SSDs * 7	100 * 7
		Snapshots * 2	8,192 * 2
		VolumeFolders * 6	8,192 * 5
		Initiators * 5	512 * 5
	Total metrics		119,500
EMC VPLEX	Unique metrics	107	
	Total metric count = Sum (Scaling factor * metric count)	VPLEX System * 339 ^a	1 * 339
		Device * 4	964 * 4
		Virtual Disk * 8	1,346 * 8
		Virtual Volume * 12	611 * 12
		Extent * 3	1,568 * 3
		StorageView * 2	240 * 2
		DistributedComponent * 3	202 * 3
		ConsistencyGroup * 1	1146 * 1
	Total metrics		29,231
EMC Isilon	Unique metrics	95	
	Total metric count = Sum (Scaling factor * metric count)	Isilon System * 221 ^a	1 * 221
		NfsExport * 1	4,930 * 1
		Quota * 4	1,000 * 4
		Disk * 13	25 * 13
		Snapshot * 1	346 * 1
	Total metrics		9822
EMC Centera	Unique metrics	42	
	Total metric count = Sum (Scaling factor * metric count)	Node * 5	64 * 5
		Storage Pool * 5	6 * 5
		Centera System * 32	1 * 32
	Total metrics		382
Brocade FC Switch - Discovery through SNMP, Fabric & zoning discoveries through SMI-S	Unique metrics	86	
	Total metric count = Sum (Scaling factor * metric count)	Switch * 7 ^a	1 * 7
		Port & Interface * 75	32 * 75
		Zoneset * 1	1
		Zone * 1	1
		ZoneMember * 2	652 * 2
	Total metrics		3,713
Brocade FC Switch	Unique metrics	35	

SMI-S discovery	Total metric count = Sum (Scaling factor * metric count)	Switch * 23 ^a	1 * 23
		Port * 23	32 * 23
		Zoneset * 1	1
		Zone * 1	1
		ZoneMember * 2	652 * 2
	Total metrics		2,065
Cisco MDS/Nexus Switch	Unique metrics	91	
	Total metric count =Sum (Scaling factor * metric count)	Switch * 849 ^a	1 * 849
		Port & Interface * 103	32 * 103
		Zoneset * 1	1
		Zone * 1	1
		ZoneMember * 2	652 * 2
		VSAN * 3	15 * 3
	Total metrics		5496
EMC Host Interface	Unique metrics	39	
	Total metric count =Sum (Scaling factor * metric count)	File System Count * 4	7*4
		Disk count * 7	32*7
		Volume Group * 4	2*4
		Hosts * 54 ^a	1 * 54
	Total metrics		314
Physical Hosts Without PowerPath	Unique metrics	39	
	Total metric count =Sum (Scaling factor * metric count)	Hosts * 54 ^a	1* 54
		File Systems* 4	7*4
		Disks* 10	32*10
		Volume Group * 4	2*4
		Path-Availability * 1	60 * 1
	Total metrics		470
Physical Hosts With PowerPath	Unique metrics	75	
	Total metric count = Sum (Scaling factor * metric count)	Hosts * 54 ^a	1 * 54
		File System * 4	16 * 4
		Disk count * 10	32 * 10
		Volume Group * 4	4 * 4
		Path * 9	60 * 9
		Remote LUN * 14	32 * 14
	Total metrics		1,442

VMware vCenter Without PowerPath	Unique metrics	102	
	Total metric count = Sum (Scaling factor * metric count)	ESX * 99	5 * 99
		VM * 44	200 * 44
		RDMs/Datastores * 15	500 * 15
	Total metrics		16,795
VMware vCenter With PowerPath	Unique metrics	164	
	Total metric count = Sum(Scaling factor * metric count)	ESX * 99	5 * 99
		VM * 44	200 * 44
		Disk * 15	500 * 15
		Path = Disk * 6	500 * 6
	Total metrics		19,795
EMC Data Domain	Unique metrics	75	
	Total metric count = Sum (Scaling factor * metric count)	DataDomainSystem * 20 ^a	1 * 20
		FileSystems * 14	1 * 14
		Disks * 15	12 * 15
		NodeInterfaces * 26	25 * 26
	Total metrics		864
EMC Atmos	Unique metrics	158	
	Total metric count = Sum (Scaling factor * metric count)	RMGSSite * 2	6 * 2
		Node * 12	59 * 12
		Disk * 5	480 * 5
		Filesystem * 4	20 * 4
		Subtenant * 42	50 * 42
		User * 42	1,000 * 42
			Others 11 %
	Total metrics		47,300
Microsoft Hyper-V	Unique metrics	55	
	Total metric count = Sum (Scaling factor * metric count)	FS*15	16,000 * 14
		HBA*11	40*11
		Snap*1	100*1
		System*11	10*11
		Processor*4	80*4
		Memory*13	10*13
	Total metrics		225,100
IBM-XIV	Unique metrics	135	

	Total metric count = Sum (Scaling factor * metric count)	LUN * 10	3,000 * 10
		Disk * 2	180 * 2
		Storage Pool * 16	100 * 16
		XIV Array * 107 ^a	1 * 107
	Total metrics		84197
NetApp (7-Mode)	Unique metrics	123	
	Total metric count = Sum (Scaling factor * metric count)	Disk * 4	264 * 4
		Aggregate * 23	12 * 23
		Volumes * 27	4096 * 27
		Luns * 11	1,024 * 11
		NetApp Arrays * 756 ^a	1 * 756
		CIFS * 19	150 * 19
		NFS * 53	90 * 53
	Total metrics		131564
NetApp Cluster-Mode with 2 nodes	Unique metrics	167	
	Total metric count = Sum (Scaling factor * metric count)	Disk * 4	384 * 4
		Aggregate * 24	192 * 24
		Volumes * 20	400 * 20
		Luns * 5	1638 * 5
		NetApp Arrays * 756 ^a	1 * 756
		CIFS * 13	225 * 13
		NFS * 103	135 * 103
	Total metrics		39920
IBM-LPARS	Unique metrics	42	
	Total metric count = Sum (Scaling factor * metric count)	Devices * 12	18 * 5
		LPARs * 21	21 * 6
		CPUs * 9	84 * 9
	Total metrics		972
IBM-SVC	Unique metrics	196	
	Total metric count = Sum (Scaling factor * metric count)	MDisk * 16	1024 * 16
		SVC Array * 180 ^a	1 * 180
		MDiskGroup * 8	128 * 8
		VDisk * 49	2048 * 49
		MirrorExtent * 2	82 * 2
		Port * 37	8 * 37

		Host * 3	8 * 3
	Total metrics		118424
EMC ViPR	Unique metrics	183	
	Total metric count = Sum (Scaling factor * metric count)	Array * 72	24 * 72
		ECS * 238	2 * 238
		File Share * 10	600 * 10
		Host * 4	600 * 4
		LUN * 19	20000 * 19
		Pool * 28	144 * 28
		VDC * 1650 ^c	2 * 1650
	Total metrics		397936
Hitachi Device Manager	Unique metrics	98	
	Total metric count = Sum (Scaling factor * metric count)	HDS Array * 280 ^a	1 * 280
		Disk * 2	444 * 2
		Volumes * 12	5048 * 12
		Storage Pool * 14	10 * 14
		ParityGroup * 5	10 * 5
		Masking View * 2	500 * 2
	Total metrics		62934
EMC ScaleIO	Unique metrics	136	
	Total metric count = Sum (Scaling factor * metric count)	ScaleIOSystem * 76 ^a	1 * 76
		SDSController * 15	256 * 15
		Drive * 15	128 * 15
		ProtectionDomain * 15	64 * 15
		LUN * 1	8192 * 1
		StoragePool * 14	256 * 14
	Total metrics		18572
Cisco UCS	Unique metrics	95	
	Total metric count = Sum (Scaling factor * metric count)	UCSInstance * 53	1 * 53
		BladeChassis * 6	16 * 6
		FabricInterConnect * 13	26 * 13
		FabricInterConnect – FCPort * 15	6 * 15
		FabricInterConnect – EthernetPort * 28	20 * 28
		BladeServer * 10	1 * 10
		BladeServer-NIC * 22	4 * 22

		BladeServer-Port * 1	2 * 1
	Total metrics		1237
HP StorageWorks	Unique metrics	98	
	Total metric count = Sum (Scaling factor * metric count)	Array * 38 ^a	1 * 38
		Disk * 2	444 * 2
		LUN * 7	5048 * 7
		MaskingView * 2	500 * 2
		RaidGroup * 5	10 * 5
		StoragePool * 14	10*14
	Total metrics		37452
HP 3PAR	Unique metrics	106	
	Total metric count = Sum (Scaling factor * metric count)	Array * 217 ^a	1 * 123
		Disk * 18	180 * 18
		VirtualVolume * 22	3000 * 12
		StoragePool * 10	20 * 13
	Total metrics		39623
HP EVA	Unique Metrics	199	
	Total metric count = Sum (Scaling factor * metric count)	Array * 130 ^a	1 * 130
		Disk * 16	32 * 16
		Disk Group * 16	10 * 16
		Storage Pool * 16	10 * 16
		Virtual Disk * 21	32 * 21
	Total Metrics		1634
IBM-DS (DS-3k-4k-5k)	Unique metrics	55	
	Total metric count = Sum (Scaling factor * metric count)	Array * 69 ^a	1 * 69
		Controller * 11	2 * 11
		Disk * 11	112 * 8
		LUN * 20	2000 * 20
		Storage Pool * 7	20 * 7
	Total metrics		41127
IBM-DS (DS-6k-8k)	Unique metrics	65	
	Total metric count = Sum (Scaling factor * metric count)	Array * 64 ^a	1 * 64
		Disk * 2	1500 * 2
		LUN * 7	16000 * 7
		Storage Pool * 4	20 * 4
	Total metrics		115144

DPA	Unique Metrics	64	
	Total metric count = Sum (Scaling factor * metric count)	DataDomainSystem * 13 ^a	1*13
		DataDomainNodeInterface * 12	25*12
		RecoverPointSystem * 19	1*19
		RecoverPointConsistencyGroup * 14	15 * 14
		RecoverPointReplicationSet * 3 (Depends upon the volume)	2000 * 3
	Total Metrics		6542
RecoverPoint	Unique Metrics	31	
	Total metric count = Sum (Scaling factor * metric count)	RecoverPointSystem * 101 ^a	1 * 101
		Cluster * 3	2 * 3
		Volume * 1	765 * 1
		ReplicationSet *1	284 * 1
		ConsistencyGroup * 15	40 * 15
	Total Metrics		1756

^a Overall metrics other than scaling factors

^b ZoneMemIdentifier is the metric that also used for Passive Host discovery. The count of this metric increases in proportion to the number of LUNs masked to each Passive host.

^c Each VDC is assumed to contain 150 projects and 10 tenants. There are 10 metrics per project and 12 per tenant and 30 per VDC. (VDC = (10*150 + 12*10 + 30 = 1650))

Percentage change in collection

The following table shows the change in the number of metrics collected in each release vs. the number of metrics collected in version 3.0. The change in metrics collections is depicted in percentages.

For example, in case of NetApp-7 Mode, 3% is the increase over the number of metrics collected in version 3.0.2 vs 3.0; and 60% increase in version 3.5 vs 3.0.

Solution Pack	ViPR SRM Version					Comments
	3.0	3.0.1	3.0.2	3.5	3.5.1	
VMAX	-	89%	88%	87%	87%	The wide increase is a result of LUN level metrics collected using SMI-S
VNX-Block	-	-2%	1%	56%	57%	

NetApp – 7 Mode	-	3%	3%	60%	-60%	Reason for reduction in metrics - Disabled Snapshot and Volume performance metrics collection in 3.5.1
NetApp – C Mode	-	1%	1%	1%	-30%	Reason for reduction in metrics - Disabled Snapshot and Volume performance metrics collection in 3.5.1
Physical Hosts	-	-	-	-	20%	Path metric is introduced for compliance IO redundancy rule validation. The increase will be observed on non-PowerPath host configurations. In case of hosts where PowerPath is configured, refer to “Guideline D2”

What has changed in 3.5.1 Guidelines?

- Guideline D3 - Modified compliance policy configuration guidelines
- Appendix A - Updated the “Metric count and scaling factor” factor table with latest data
- Appendix A – Updated “Percentage change in collection” table with 3.5.1 data.