



CATE Storage and Backup Engineering
CATE CitiSystems Engineered Systems
Citigroup Architecture and Technology Engineering

ScaleIO v1.31.1277

CSBE-15007-RN

Release Notes

1.0

©Citigroup Inc. 2015, All Rights Reserved.

This document contains material confidential to Citigroup. The contents are protected by specific employee and contractor agreements regarding confidentiality. Reproduction for distribution outside of Citigroup is prohibited.

1 Version Control

Version	Date	Author	Reviewed	Comments
0.1	1-Feb-2015	Scott Torgerson	Chris Clayton, Chris Howarth, Bob Small, Eliot Wilson, Ludwig Derodel	Peer
0.5	2-June-2015	Scott Torgerson, Chris Clayton, Chris Howarth	Global Stakeholders	RFT
1.0	22-June-2015	Scott Torgerson, Chris Clayton, Chris Howarth	Global Stakeholders	Release to Production

2 Legal Notice

Information contained herein is for Internal Use and may be used only for business purposes authorized by Citigroup.

CITI, CITIBANK, CITICORP, and CITIGROUP are trademarks and service marks of Citicorp and are used and registered throughout the world. Citibank, Citicorp, Citigroup and their subsidiaries also claim rights in certain other trademarks and service marks contained in these pages.

3 Contact

Please address any comments or questions regarding the information contained in this document via email to the [*CATE Global Engineering Support – ScaleIO](#) distribution list.

4 Latest Versions

It is recommended that the latest version of the document be used from the CSBE website at the following URL:
<http://catecollaboration.citigroup.net/domains/platstor/gdse/products/default.aspx>.
Newer versions may include certification on additional platforms, or comments on issues discovered after deployment.

5 Table Of Contents

1	VERSION CONTROL	2
2	LEGAL NOTICE.....	3
3	CONTACT.....	3
4	LATEST VERSIONS	3
5	TABLE OF CONTENTS	4
6	TABLE OF FIGURES.....	6
7	INTRODUCTION.....	7
8	APPLICABLE ENVIRONMENTS.....	7
9	COMPONENT CERTIFICATION MATRIX.....	8
9.1	CERTIFIED COMPONENT AND FUNCTIONALITY MATRIX	8
10	KNOWN ISSUES.....	9
10.1	THIN PROVISIONING ON PHYSICAL SERVERS	9
10.2	EMC SRM.....	9
10.3	NO CHARGEBACK-RELATED DATA IN SRM	10
10.4	DATA MOBILITY	10
10.5	MAINTENANCE MODE	11
10.6	SILENT DATA CORRUPTION.....	11
10.7	REBUILD/REBALANCE THROTTLES	11
10.8	RECOVERY FROM MDM LOSS.....	12
10.9	NO VASA SUPPORT	12
10.10	REBUILDS DO NOT COMPLETE	12
10.11	SDC BALANCING ALGORITHMS.....	13
11	APPENDICES	15
11.1	APPENDIX A: TEST ENVIRONMENTS	15
11.1.1	Warren LAB Cluster	15
11.1.2	Grid LAB Cluster	15
11.1.3	Maas (Mongo as a Service) POC Environment.....	16
11.1.4	Raas (Oracle RAC as a Service) POC Environment	17
11.1.5	Georgetown Tornado++ Cluster.....	17
11.2	APPENDIX B: TEST PLANS.....	19
11.3	APPENDIX C: TEST STEPS AND RESULTS.....	21
11.3.1	Test #1 – ScaleIO Cluster Components Installation.....	21
11.3.2	Test #2 – ScaleIO Cluster Configuration.....	21
11.3.3	Test #3 – SDC Installation/Configuration (VMWare).....	21
11.3.4	Test #4 – SDC Installation/Configuration (Physical Servers).....	22
11.3.5	Test #5 – Specific ScaleIO Security Settings.....	22
11.3.6	Test #6 – Citi Signed SSL Certificates	23
11.3.7	Test #7 – RAM Read Cache Settings.....	23

11.3.8	Test #8 – Spare Capacity Settings	24
11.3.9	Test #9 – Pool/Disk type Throttling	24
11.3.10	Test #10 – Thick Provisioning	25
11.3.11	Test #11 – Thin Provisioning (physical servers only)	26
11.3.12	Test #12 – Standby MDM Failover Testing.....	26
11.3.13	Test #13 – Various Cluster Failure Testing.....	27
11.3.14	Test #14 – Event Notification (Tivoli/NetCool)	27
11.3.15	Test #15 – Log Reconciliation (SEMs/ArcSight)	28
11.3.16	Test #16 – Automated SBCC (HPNA).....	29
11.3.17	Test #17 – EMC SRM.....	29
11.3.18	Test #18 – Obtaining ScaleIO permanent License	30
11.3.19	Test #19 – Disk Failure Testing.....	30
11.3.20	Test #20 – Network Failure Testing (Cluster and SDC).....	31
11.3.21	Test #21 – Network Testing (Various)	32
11.3.22	Test #22 – Durham, NC Performance Testing.....	33
11.3.23	Test #23 – Tornado++ Performance Testing	33

6 Table of Figures

Figure 1 - Warren LAB Cluster Test Environment.....	15
Figure 2 - Grid LAB Cluster Test Environment.....	16
Figure 3 - Maas Cluster Test Environment	16
Figure 4 - Raas Cluster Test Environment.....	17
Figure 5 - Tornado++ Cluster Test Environment.....	18

7 Introduction

ScaleIO is a software-only solution that uses existing server's local disks and IP LAN to create a virtual SAN that has all the benefits of shared storage (but at a lower cost). ScaleIO utilizes the existing local internal storage of many commodity servers and turns it into internal shared block storage volumes. These ScaleIO block volumes are then provisioned over 10Gb IP network to clients.

The lightweight ScaleIO software components are installed on the application servers and communicate via a standard IP LAN to handle the application I/O requests sent to ScaleIO block volumes. An efficient decentralized block I/O flow, combined with a distributed, sliced volume layout, results in a massively parallel I/O system that can scale up to hundreds of nodes.

ScaleIO is designed and implemented with enterprise-grade resilience. Furthermore, the software features an efficient distributed self-healing process that overcomes media and node failures, without requiring administrator involvement.

Dynamic and elastic, ScaleIO enables administrators to add or remove nodes and capacity on-the-fly. The software immediately responds to the changes, rebalancing the storage distribution and achieving a layout that optimally suits the new configuration.

Because ScaleIO is hardware agnostic, the software works efficiently with various types of disks, including: magnetic (HDD) and solid-state disks (SSD), flash PCI Express (PCIe) cards, networks, and hosts.

8 Applicable Environments

This certification ONLY applies to ScaleIO DEV and UAT environments running on Tornado++ hardware. POC (Proof Of Concept) or Pilot environments not running on Tornado++ hardware will only be supported working in conjunction with CATE. Any POC/Pilot environments should first contact CATE at [*CATE GLOBAL Engineering Support – ScaleIO](#) prior to starting the installation of the environment.

9 Component Certification Matrix

9.1 Certified Component and Functionality Matrix

Component Matrix	
ScaleIO ¹	1.31.1277
ScaleIO Dashboard	1.31.1277
REST Gateway	1.31.1277
ScaleIO Dashboard Supported O.S. versions ²	Windows 2008 R2 (SOE Build) Windows 2012 (SOE Build) Windows 2012 R2 (SOE Build)
MDM	RHEL 6.X (SOE Build)
Tie-Breaker	RHEL 6.X (SOE Build)
SDS	RHEL 6.X (SOE Build)
SDC Supported O.S. versions ³	RHEL 6.X (SOE Build) Windows 2008 R2/2012 (SOE Build) ESXi 5.5 SP1 (SOE Build)

Note: This project only covers ScaleIO v1.31.1277 which is slated for DEV/UAT and POC uses only. ScaleIO v1.32 will be GA 29th May 2015 which will be the targeted release for ScaleIO production use, after the CATE certification completes.

Before installation or configuration of any ScaleIO Cluster it should be validated by the implementer that all component versions are listed in the above certified matrix. Any software versions not listed in the above matrix will not be supported.

Complete best practices, installation, and configuration instructions and details for all of the components listed in the above table can be found in the [EMC ScaleIO R1.31.1277 Solutions Guide](#).

¹ All servers that makeup the ScaleIO cluster (MDM, SDS, REST Gateway) will only be certified on Linux SOE built servers.

² Due to the Java requirements, the GUI will only be certified for SOE built windows servers.

³ Currently there is no ScaleIO support for Solaris and AIX.

10 Known Issues

During the certification testing of the EMC ScaleIO project, the following issues were found and recorded. An ongoing list of all Bugs and Requests for Enhancements for EMC ScaleIO are being tracked and recorded in [JIRA](#) tracking system.

10.1 Thin Provisioning on Physical Servers

Impact – High

Affects – Any cluster that has thinly provisioned LUNs when a pool fills up.

Description

When LUNs are thinly provisioned they do as expected during normal use in that free storage is only taken from the pool when the thinly provisioned LUN uses space. However, when the pools that contain thin LUNs were tested against the pool reaching the hard limit, there were unexpected results. When the pool should have filled and locked the thin LUNs as read-only, the pool started using protected space in the pool to fulfill the actual LUN write request. Since a higher priority was given to the thin LUN, it was able to finish the write using protected space making the cluster show as unprotected.

Workaround

There is no current workaround for this issue. As a result, CATE will not certify thin provisioned LUNs for physical servers in this initial release of ScaleIO.

Further Information

This issue was reported to EMC and they were able to duplicate the issue. EMC has stated that this issue will be resolved in the 1.32 release due to GA in May 2015. CATE will re-evaluate and test this issue in that release and attempt to certify this functionality for that release.

10.2 EMC SRM

Impact – Medium

Affects – Performance and capacity reporting of ScaleIO with SRM

Description

Performance and capacity reports required in SRM for ScaleIO management are either not available or where available, some are inaccurate.

Workaround

Use SIO Dashboard to get pool level metrics for capacity planning purposes.

Further Information

EMC are planning to meet Citi's reporting requirement over the next few releases of SRM focusing on:

1. Fixing any algorithm issues in existing reports that Citi detects.
2. Reporting on data already in SRM database for which there were no reports
3. Pulling additional data (e.g. volume performance metrics) into SRM database and reporting on that
4. Adding new performance metrics (e.g. client level metrics) into the SIO REST API to allow them to be consumed by SRM

10.3 No chargeback-related data in SRM

Impact – Medium

Affects – Chargeback

Description

No volume or client-related data in SRM that can be used for chargeback.

Workaround

Chargeback for virtual environments will need to be undertaken using VMware tools (as today). No chargeback is possible for physical hosts.

Further Information

Volume related data will start to appear in SRM from 3.6.3 onwards.

10.4 Data Mobility

Impact – Medium

Affects – Movement of data between pools (tiers) or protection domains.

Description

It is not possible to move data at the ScaleIO level between pools (tiers) or protection domains.

Workaround

All data migration between pools or protection domains will need to be undertaken using VMware tools (e.g. storage VMotion).

Further Information

The functionality is on EMC's roadmap for post R2.0.

10.5 Maintenance mode

Impact – Medium

Affects – Removal of SDS nodes for maintenance tasks

Description

Need to ability to migrate data from an SDS prior to shut down without completely removing SDS from cluster

Workaround

For short outages (reboots) CATE is recommending not removing the SDS but instead using a hardened DRL to ensure the data is degraded for the minimum period possible.

Further Information

Native maintenance mode is coming in ScaleIO R2.0.

10.6 Silent data corruption

Impact – Medium

Affects – Data integrity

Description

ScaleIO is unable to detect and correct copies of data that have suffered from silent data corruption.

Workaround

None for R1.311.

Further Information

R1.32 will introduce the ability to detect (but not correct) data corruption. Correction will not come until post-2.0 and certain server hardware enhancements are available.

10.7 Rebuild/rebalance throttles

Impact – Medium

Affects – Client response time during rebuild & rebalance operations.

Description

The R1.311 IO priority controls do not function correctly in R1.311.

Workaround

None for R1.311.

Further Information

R1.32 will fix this bug.

10.8 Recovery from MDM loss

Impact – Medium

Affects – Cluster availability

Description

Cluster only has a maximum of 2 MDMs. There is exposure if one is down.

Workaround

Use the `mdm_failover` script provided by ScaleIO.

Further Information

R2.0 will have 3 MDMs minimum.

10.9 No VASA support

Impact – Medium

Affects – VMware operations

Description

VMware cannot query the ScaleIO storage to determine its capabilities.

Workaround

Use a sensible pool naming convention to identify pool capabilities.

Further Information

Stretch goal for EMC for R2.0

10.10 Rebuilds do not complete

Impact – High

Affects – Data availability

Description

Rebuilds do not complete even though there is still spare capacity available. This is believed (at the time of writing) to be due to the use of fault sets.

Workaround

Leave additional unused capacity to allow rebuild to complete. Values to be determined by trial and error (already done and documented for T++ pools).

Further Information

Waiting on EMC for new Best Practice on how to determine the correct amount of spare capacity to assign when fault sets are in use.

10.11 SDC balancing algorithms

Impact – Low

Affects – Optimal use of resources

Description

SDC & SDS traffic is not always evenly balanced across all paths

Workaround

None.

Further Information

EMC will analyze ScaleIO interface balancing algorithms to determine if changes are needed to improve performance and utilization of multiple connections to storage. EMC has confirmed this is a bug and plans a fix in v2.0.

References

Document	Location
CATE Scale-out commodity block storage position paper	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/Eval/Scale_Out_Block_storage_position_paper_V1.0.pdf
ScaleIO Product Page	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/techstdshmpgs/emcscaleIO.aspx
ScaleIO Standard Build Page	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/EMC%20ScaleIO/ScaleIO_csb.aspx
ScaleIO v1.31.1277 GNAST Solutions Guide.	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/EMC%20ScaleIO/GNAST%20Design%20and%20Configuration%20Guide%20-%20EMC%20ScaleIO%20R1.31.pdf
EMC ScaleIO v1.31.1277 Users Guide	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/platforms/EMC%20ScaleIO/EMC%20ScaleIO%20Users%20Guide%20v1.31.pdf
ScaleIO Security Model	http://ewiki.nam.nsroot.net/w/images/f/f6/ScaleIO_Security_Model_071014.pdf
VMWare VSI Storage	https://catecollaboration.citigroup.net/domains/platstor/gdse/products/Designs/CSBE-11079-DD.pdf

11 Appendices

11.1 Appendix A: Test Environments

Multiple test environments were used when testing ScaleIO performance and functionality. This was in part due to the lack of a large scale test environment when most functional tests were completed. As the certification progressed more realistic test environments became available and tests were performed again where needed. At the very end of this initial certification, the first Tornado++ environment was made available so that CATE could test on it for approximately one week.

The following sections will show the test environments used for testing ScaleIO as part of this certification.

11.1.1 Warren LAB Cluster

The Warren LAB cluster is a small (4) node cluster running on HP DL360-G7 servers. Each of the SDS nodes only contains (3) SAS drives that are in the cluster. This environment was used for most of the security testing that was completed. This was the Cluster used for all CIVA testing.




Figure 1 - Warren LAB Cluster Test Environment

11.1.2 Grid LAB Cluster

The Grid LAB cluster is a small (5) node cluster running on IBM HS23 blade servers. Each of the SDS nodes only contains (1) SAS drive that are in the cluster. This test cluster also did not have a second NIC for creating a private

network for the ScaleIO cluster traffic. As a result, this environment was only used for functional testing.




Figure 2 - Grid LAB Cluster Test Environment

11.1.3 Maas (Mongo as a Service) POC Environment

The Maas Cluster was a POC/Pilot environment that contained (4) nodes in the cluster. However, the nodes in the cluster much closer resembled production nodes in a Tornado++ hardware environment. Each of the nodes in the cluster had (20) SAS drives and (2) SSD drives. Although the nodes in this cluster resembled production nodes, the network did not match a production configuration. As a result, many of the tests in this environment revolved around the functionality of multiple Tiers and SSD pools etc.




Figure 3 - Maas Cluster Test Environment

11.1.4 Raas (Oracle RAC as a Service) POC Environment

The Raas Cluster was a POC/Pilot environment that contained (4) nodes in the cluster. Like the Maas cluster, these nodes also resembled a Tornado++ production node. Each of the nodes in the cluster had (20) SAS drives and (3) SSD drives. This test cluster also had the added benefit of having a planned production type network configuration also. Most every kind of test that could be run on this cluster was run as time permitted since it closely resembled production.





Figure 4 - Raas Cluster Test Environment

11.1.5 Georgetown Tornado++ Cluster

The Tornado++ cluster was a real production configuration as described in the [ScaleIO Environments on Tornado++ Best Practices document](#). CATE was given this environment for one week to test on prior to it being turned over to the Operations team. This environment was used to do any performance and network testing that could be completed. Many functional tests were also completed on this environment by setting up the environment by using the CATE Solutions Guides for ScaleIO.



“SDC” MDM_IP=**10.123.32.11,10.123.32.75,10.123.32.91,10.123.96.11,10.123.96.75,10.123.96.91**

“CLUSTER” MDM_IP=**10.123.128.11,10.123.128.75,10.123.128.91,10.123.160.11,10.123.160.75,10.123.160.91**

Management IP=**10.123.32.11,10.123.32.75,10.123.32.91,10.123.96.11,10.123.96.75,10.123.96.91**

Figure 5 - Tornado++ Cluster Test Environment

Note: The above diagram of the test environment only shows the high level components of the cluster. A complete layout of a Tornado++ cluster can be found in the [ScaleIO Environments on Tornado++ Best Practices document](#)

11.2 Appendix B: Test Plans

Test	Test Environments Used	Test
<u>1</u>	Tornado++, Maas, Raas, Warren LAB, Grid LAB	ScaleIO Cluster components Installation (MDM, SDS, Tie-Breaker, REST Gateway)
<u>2</u>	Tornado++, Maas, Raas, Warren LAB, Grid LAB	ScaleIO Cluster Configuration
<u>3</u>	Warren LAB VMWare 5.5 Cluster pointing at Tornado++, Maas, Raas, Warren LAB, and Grid LAB clusters	SDC Installation/Configuration (VMWare)
<u>4</u>	Windows and Linux servers in Warren Lab pointing at Tornado++, Maas, Raas, Warren LAB, and Grid LAB clusters	SDC Installation/Configuration (physical)
<u>5</u>	Tornado++, Maas, Raas, Warren LAB, Grid LAB	Security Settings
<u>6</u>	Tornado++, Maas, Raas, Warren LAB, Grid LAB	Citi Signed SSL Certificates
<u>7</u>	Tornado++, Maas, Raas, Warren LAB	RAM Read Cache Settings
<u>8</u>	Tornado++, Maas, Raas, Warren LAB, Grid LAB	Spare Capacity Settings
<u>9</u>	Tornado++, Maas, Raas, Warren LAB, Grid LAB	Pool/Disk type Throttling
<u>10</u>	Tornado++, Maas, Raas, Warren LAB, Grid LAB	Thick Provisioning (virtual and physical environments)
<u>11</u>	Warren LAB, Grid LAB	Thin Provisioning (physical servers only)
<u>12</u>	Maas, Raas, Warren LAB, Grid LAB	Standby MDM Failover Testing
<u>13</u>	Maas, Raas, Warren LAB, Grid LAB	Various Cluster Failure Testing
<u>14</u>	Maas, Raas, Warren LAB, Grid LAB	Event Notification (Tivoli/NetCool)
<u>15</u>	Maas, Raas, Warren LAB, Grid LAB	Log Reconciliation (SEMs/ArcSight)
<u>16</u>	Maas, Raas, Warren LAB, Grid LAB	Automated SBCC (HPNA)
<u>17</u>	Maas, Raas, Warren LAB, Grid LAB	EMC SRM
<u>18</u>	Maas, Raas, Warren LAB, Grid LAB	Obtaining ScaleIO permanent License

<u>19</u>	Maas, Raas, Warren LAB, Grid LAB	Disk Failure Testing
<u>20</u>	Tornado++	Network Failure Testing (Cluster and SDC)
<u>21</u>	Tornado++	Network Testing (Various)
<u>22</u>	EMC Durham NC LAB environment	Durham Performance testing
<u>23</u>	Tornado++	Georgetown Tornado++ cluster performance testing.

11.3 Appendix C: Test Steps and Results

The following sections will show the test results of the tests that were run as part of this project certification. It is important to note that some tests that were run using steps in the Solutions Guide will refer to that document for the steps.

11.3.1 Test #1 – ScaleIO Cluster Components Installation

Item	Test	Expected Results / Explanation	Actual Results	Pass/ Fail	Notes
1	Install all components of the ScaleIO Cluster (MDM, Tie-Breaker, SDS, REST Gateway)	All ScaleIO components should install without errors	As Expected	Pass ✓	Use documented procedure in ScaleIO Cluster Installation section of the Solutions Guide .

11.3.2 Test #2 – ScaleIO Cluster Configuration

Item	Test	Expected Results / Explanation	Actual Results	Pass/ Fail	Notes
1	Configure the ScaleIO Cluster	Cluster should configure without error and be able to return basic commands such as <code>scli -query_cluster_all</code>	As Expected	Pass ✓	Use documented procedure in ScaleIO Cluster Configuration section of the Solutions Guide .

Note: All installation and configuration commands used in Test #1 and #2 were taken from the documented steps in the Solutions Guide. The same commands were used by operational partners to build all POC environments.

11.3.3 Test #3 – SDC Installation/Configuration (VMWare)

Item	Test	Expected Results / Explanation	Actual Results	Pass/ Fail	Notes
1	Install and configure a VMWare SDC. This will allow VMWare Clusters to see ScaleIO storage.	At this point a ScaleIO LUN should be able to be provisioned to the VMWare cluster. ScaleIO Cluster should also show the newly added VMWare SDC as Connected in the following command:	As Expected	Pass ✓	Use documented procedure in VMWare SDC Installation section of the Solutions Guide .

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
		scli -query_all_sdc			

11.3.4 Test #4 – SDC Installation/Configuration (Physical Servers)

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Install and configure a Windows physical server SDC	At this point a ScaleIO LUN should be able to be provisioned to the Windows server. Cluster should also show the newly added VMWare SDC as Connected in the following command: scli -query_all_sdc	As Expected	Pass ✓	Use documented procedure in <u>Windows SDC Installation</u> section of the Solutions Guide .
2	Install and configure a Linux physical server SDC	At this point a ScaleIO LUN should be able to be provisioned to the Linux server. Cluster should also show the newly added VMWare SDC as Connected in the following command: scli -query_all_sdc	As Expected	Pass ✓	Use documented procedure in <u>Linux SDC Installation</u> section of the Solutions Guide .

11.3.5 Test #5 – Specific ScaleIO Security Settings

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Set the remote read-only state to all remote access to ScaleIO cluster is read-only.	ScaleIO CLI commands should only be allowed to run on primary MDM. The following command will also show the remote read-only as enabled:	As Expected	Pass ✓	Use documented procedure in the <u>Global Standard (Security) - Configuration</u> section of the Solutions Guide .

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
		<code>scli --query_remote_read_only_limit_state</code>			
2	Setup restricted SDC mode.	The following command will show SDC nodes as not approved that were added prior to this feature being enabled: <pre>scli --query_all_sdc</pre> The following command will show SDC restricted mode as enabled. <pre>scli --query_restricted_sdc_mode</pre>	As Expected	Pass ✓	Use documented procedure in Obtaining and Installing Citi Trusted Certificates section of the Solutions Guide .

11.3.6 Test #6 – Citi Signed SSL Certificates

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Install and configure Citi Trusted certificate	The following command will return a session key: <pre>curl -X GET -H 'Accept: application/json;version=1.0' -u admin https://<FQDN>/api/login --cacert <CA public certificates file></pre>	As Expected	Pass ✓	Use documented procedure in Obtaining and Installing Citi Trusted Certificates section of the Solutions Guide .

11.3.7 Test #7 – RAM Read Cache Settings

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Calculate and configure RAM Read Cache on SDS nodes.	The following command will show that all SDS RAM Read cache sizes were able to be set at	As Expected	Pass ✓	Use documented procedure in RAM Read Cache section of the Solutions Guide . The procedure is in the

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
		65% of the total memory on each SDS: <pre>scli --query_properties --object_type SDS -- properties RMCACHE_SIZE_IN_KB -- all_objects</pre>			Planning Steps Checklist section of the Solutions Guide.

11.3.8 Test #8 – Spare Capacity Settings

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Calculate and configure spare capacity.	The following command should show that all pools have been set to the proper spare setting: <pre>scli --query_properties --object_type STORAGE_POOL -- properties SPARE_PERCENT -- all_objects</pre>	As Expected	Pass ✓	Use documented procedure in Spare Capacity section of the Solutions Guide . The procedure is in the Planning Steps Checklist section of the Solutions Guide.

11.3.9 Test #9 – Pool/Disk type Throttling

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Volume QoS. Limit the IOPS that a volume can have.	Volume performance will be limited to IOPS throttle set. <pre># scli -- set_sdc_volume_limits --mdm_ip 10.40.99.173 - -volume_name cac01 -- sdc_ip 10.40.99.171 -- limit_iops 100</pre> Reest to unlimited with <pre># scli -- set_sdc_volume_limits --mdm_ip 10.40.99.173 - -volume_name cac01 --</pre>	As Expected	Pass ✓	SIO Dashboard was to confirm that the throttle was being adhered to.

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
		<code>sdcli --limit-bandwidth 0</code>			
2	Volume QoS. Limit the bandwidth that a volume can have.	Volume performance will be limited to bandwidth throttle set. <pre># sdcli -- set_sdc_volume_limits --mdm_ip 10.40.99.173 - -volume_name cac01 -- sdcli 10.40.99.171 -- limit-bandwidth 10</pre> Reest to unlimited with <pre># sdcli -- set_sdc_volume_limits --mdm_ip 10.40.99.173 - -volume_name cac01 -- sdcli 10.40.99.171 -- limit-bandwidth 0</pre>	As Expected	Pass ✓	SIO Dashboard was to confirm that the throttle was being adhered to.

11.3.10 Test #10 – Thick Provisioning

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Create a Thick volumes and provision it to a test SDC	LUN should show up to SDC and be mountable.	As Expected	Pass ✓	Use documented procedure in Provisioning LUNs for Physical Environments and Provisioning LUNs for Virtual Environments section of the Solutions Guide .
2	Validate that LUN can be used RAW or with filesystem.	LUN should accept I/O in either configuration.	As Expected	Pass ✓	Once raw device showed up to O.S., then created a filesystem on it.
3	Run reads and writes to the LUN	If a copy is used, a file should copy and have same checksum. If a tool was used to generate I/O it should	As Expected	Pass ✓	First ran iorate tool against the raw device. Then used a filesystem and made sure files could be copied to it.

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
		completed without errors.			

11.3.11 Test #11 – Thin Provisioning (physical servers only)

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Create a Thin volumes and provision it to a test SDC		As Expected	Pass ✓	Created thin LUN scott01 as 8GB and mapped to SDC gdsexiv02. Used capacity of the LUN showed as a few KB.
2	Run reads and writes to the LUN	Used space on the LUN should increase	As Expected	Pass ✓	After a few minutes of I/O the actual used space on the LUN was up to 6GB
3	Delete the LUN and fill up all but 1GB of the entire cluster with thick LUNs. Then re-create the thin LUN with a size of 8GB. On the SDC create a filesystem on the LUN. Then start copying a big file (8GB or bigger) to the filesystem.	Cluster should go to 100% and put the SDC thin LUN as read-only. This should happen very quick (minutes only)	Not Expected	Failed ✗	Even though the cluster only had 1GB of free space, it did not mark the filesystem read-only. Instead the cluster allowed the whole file to be copied by using protected space. So although the file copied, the cluster was in an unprotected mode. See Known Issues Section about thin provisioning .

11.3.12 Test #12 – Standby MDM Failover Testing

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Configure standby MDM	The node that is configured as the standby MDM should have the following process show up as running. <code>/usr/bin/python /opt/emc/scaleio/mdm_failover/bin/mdm_failover.py --log_dir /opt/emc/scaleio/mdm_fa</code>	As Expected	Pass ✓	Use documented procedure in Configuring automated MDM/TB replacement section of the Solutions Guide .

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
		<code>ilover/bin/./logs --conf_file /opt/emc/scaleio/mdm_fa ilover/bin/./cfg/conf.txt</code>			
2	Set the failover time variables low so that rebooting an MDM will cause the standby to take over.	Set the following parameters: <code>--retry_timeout=10</code> <code>--retry_count=6</code>	As Expected	Pass ✓	
3	Reboot the primary MDM.	Standby MDM should invoke. The following command should show the standby MDM as the Secondary IP: <code>scli --query_cluster</code>	As Expected	Pass ✓	

11.3.13 Test #13 – Various Cluster Failure Testing

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Primary MDM Reboot	Secondary MDM takes over as primary	As Expected	Pass ✓	
2	SDS Reboot	Cluster self-heals, reprotecting data impacted by loss of SDS. Alert raised.	As Expected	Pass ✓	
3	Tie-Breaker Reboot	Cluster operations continue without impact. Alert raised.	As Expected	Pass ✓	
4	Loss of a disk	Cluster self-heals, reprotecting data impacted by loss of disk. Alert raised.	As Expected	Pass ✓	See also test 12.3.19

11.3.14 Test #14 – Event Notification (Tivoli/NetCool)

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Configure SNMP alerts to go to Tivoli servers.	The following parameter should be set in the	As Expected	Pass ✓	Use documented procedure in <u>Tivoli Event Notification</u>

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
		/opt/emc/scaleio/gateway/webapps/ROOT/WEB-INF/classes/gatewayUser.properties file to show both SNMP trap receivers. snmp.traps_receiver_ip			(SNMP/NetCool) under ScaleIO Cluster Configuration section of the Solutions Guide .
2	Make an event happen that will cause an alert.	And event should be shown on the Tivoli/NetCool console.	As Expected	Pass ✓	To create the event, the Cluster was taken out of cluster mode using the following command: scli --switch_to_single_mode

11.3.15 Test #15 – Log Reconciliation (SEMs/ArcSight)

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Configure remote_syslog on ScaleIO cluster	The following command should show that the remote_syslog is pointed to the SEMs/ArcSight servers: scli --query_remote_syslog	As Expected	Pass ✓	Use documented procedure in Log Reconciliation (SEMs/ArcSight) under ScaleIO Cluster Configuration section of the Solutions Guide .
2	Create some change configuration events and non-change configuration events on the ScaleIO Cluster	SEMs/ArcSight server should get the events in their syslog server	As Expected	Pass ✓	For change configuration events I created and removed a volume. For non-change configuration events I did several different types of query commands.
3	Review SEMs reconciliation report	The SEMs/ArcSight report should only contain the change configuration events.	As Expected	Pass ✓	

11.3.16 Test #16 – Automated SBCC (HPNA)

Item	Test	Expected Results / Explanation	Actual Results	Pass/ Fail	Notes
1	Configure HPNA on ScaleIO Cluster		As Expected	Pass ✓	Used documented procedure in Automated SBCC (HPNA) under ScaleIO Cluster Configuration section of the Solutions Guide .
2	Create a parameter on the ScaleIO Cluster that goes against the HPNA policies and rules	Cluster should allow these <u>non</u> -best practices settings to be set.	As Expected	Pass ✓	For the EMC ScaleIO – Misc Policy the remote_read-only_limit_stated was disabled. For the EMC ScaleIO – Pools Policy the default storage pool was created.
3	Run the Policy and Rules on the HPNA server and validate that the problems are found.	When the HPNA rules run, it should throw an exception for the two different policies.	As Expected	Pass ✓	

11.3.17 Test #17 – EMC SRM

Item	Test	Expected Results / Explanation	Actual Results	Pass/ Fail	Notes
1	Ensure that SRM data can be set up in a resilient manner	SRM can collect data when MDM fails over or when primary gateway is unavailable.	As Expected for MDM test.	Partial Pass ✓	SRM can now deal with MDM failover. Also SRM permits set up of a second REST gateway to use, but there may be issues with using the token from the primary with the secondary gateway. Further testing will be undertaken.
2	Ensure ScaleIO data is collected by SRM	Data appears in SRM	As Expected	Pass ✓	
3	Ensure ScaleIO data collected by SRM is valid	Metrics in SRM match metrics in SIO Dashboard	Not Expected	Failed ✗	Logged with EMC via ticket SR# 70835146
4	Ensure SRM has all the reports Citi needs	Check reports match requirements doc	Only a subset of	Failed	EMC are working to add missing reports in the next

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
			required reports are currently available.	✗	few releases of SRM

11.3.18 Test #18 – Obtaining ScaleIO permanent License

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Validate that a permanent license is needed and then get a permanent license	Prior to getting the license the following command should show the number of days left in the temporary license. <code>scli --query_license</code>	As Expected	Pass ✓	Use documented procedure in <u>Obtaining a permanent license</u> section of the <u>Solutions Guide</u> .
2	Load the permanent license on the ScaleIO Cluster	After the permanent license is loaded the usage time left should show unlimited.	As Expected	Pass ✓	

11.3.19 Test #19 – Disk Failure Testing

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Pull a disk drive out of one of the SDS nodes in the ScaleIO Cluster.	The drive for that slot/path should show as errored in the following command: <code>scli -query_sds -sds_name [name of SDS]</code> The drive should also show up as failed in the HP Smart Array Controller CLI commands similar to this: <code>hpssacli</code>	As Expected	Pass ✓	Use documented procedure in <u>Handling a disk failure in SDS node</u> section of the <u>Solutions Guide</u> .

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
		<code>ctrl sot=2 show config</code>			
2	Flash the drive and have the hardware engineer put a new drive in place.	LED for the simulated bad drive in this test should light up.	As Expected	Pass ✓	
3	Remove and then Add the device in the effected SDS.	The following commands were used (command options not shown) <code>scli -- remove_sds_device scli --add_sds_device</code>	As Expected	Pass ✓	

11.3.20 Test #20 – Network Failure Testing (Cluster and SDC)

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Impact to storage traffic caused by component failure	No or minimal impact	As Expected	Pass ✓	The impact to storage traffic during most component failures was minimal with the exception of a leaf switch losing more than 50% of its uplinks. In this smaller lab environment we simulated this with the failure of 1 of 2 spine switches to introduce contention on the uplinks.
2	Linux NIC failure	Even load balancing	See notes	Failed ✗	Native Linux network interface failure handling does not adjust host routing tables as necessary for resilient use of multiple independent network interfaces. Workaround has been to provide a Citi software package to manage routing and NIC configuration (sds-route

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
					and carrier-mon packages)
3	SDC balancing algorithms	Load is evenly balanced	See notes	Failed ✘	EMC will analyze ScaleIO interface balancing algorithms to determine if changes are needed to improve performance and utilization of multiple connections to storage. EMC has confirm this is a bug and plans a fix in v1.32
4	SDC recovery when one path loses	No impact of access to storage	As Expected	Pass ✓	

11.3.21 Test #21 – Network Testing (Various)

Item	Test	Expected Results / Explanation	Actual Results	Pass/Fail	Notes
1	Fail single network adapter on SDS private network	No disruption to ScaleIO traffic as routed via secondary interface	As Expected	Pass ✓	
2	Fail single network on SDS public network to ensure that critical O.S. and GSF components such as Kerberos, LDAP, Tectia, Powerbroker continue to function	All critical components failed over to using the secondary interface on the public network	As Expected	Pass ✓	Some GSF products required a minor modification to ensure failures are handled correctly
3	Verify that the carrier-mon-package correctly shuts down a logical interface if the carrier is down and restores once carrier is up again	Logical interface shuts down or starts as expected	As Expected	Pass ✓	
4	Shut down primary public network connection to verify sdsroute package	Secondary adapter used as default route used when primary adapter fails	Expected	✓	

Item	Test	Expected Results / Explanation	Actual Results	Pass/ Fail	Notes
	correctly fails over default route				

11.3.22 Test #22 – Durham, NC Performance Testing

Item	Test	Expected Results / Explanation	Actual Results	Pass/ Fail	Notes
1	CATE engineers went to the EMC Durham NC LAB to run various performance tests on a larger scale environment	The expected results for these tests varied depending on the test that was being run.	As Expected	Pass ✓	All testing completed in the Durham LAB are recorded in the Durham Perf.pdf document.

11.3.23 Test #23 – Tornado++ Performance Testing

Item	Test	Expected Results / Explanation	Actual Results	Pass/ Fail	Notes
1	CATE engineers were given the Georgetown Tornado++ ScaleIO cluster for one week to do testing. Various performance tests were run.	These tests were run to fine tune previously calculated performance expectations.	As Expected	Pass ✓	Although the test was successful in determining if previous calculation were correct, a few non-performance related test issues were discovered. See Rebuilds do no complete known issue.