

Dell EMC SC Series storage and Oracle OLTP

Abstract

Testing delivered through this paper shows that Dell EMC SC Series all-flash optimized arrays are able to deliver the requirements of high-demand OLTP systems.

July 2017

Revisions

Date	Description
October 2015	Initial release
July 2017	New format. Consolidated information and made SC series storage agnostic. Added performance numbers for SC5020

Acknowledgements

This reference architecture was produced by Mark Tomczik, Matt Reichow, and Ken Cheng.

The information in this publication is provided “as is.” Dell Inc. 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 software described in this publication requires an applicable software license.

© 2012–2017 Dell Inc. or its subsidiaries. All Rights Reserved. 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. [8/1/2017] [Reference Architecture] [2045-M-RA-O]

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

Table of contents

Revisions.....	2
Acknowledgements.....	2
Table of contents	3
Executive summary.....	4
Audience	4
Prerequisites	4
1 Solution architecture.....	5
1.1 Servers and configuration.....	7
1.2 Ethernet network configuration.....	7
1.3 Ethernet network switch and adapter	7
1.4 Fibre channel switch and adapters.....	8
1.5 SAN and zoning.....	9
1.6 Multipathing	9
1.7 Dell EMC SC 5020 configuration.....	9
2 Testing methodology and monitoring tools	11
2.1 I/O profiling	11
2.2 Test results	11
3 Conclusion.....	18
A Dell EMC SC4020 reference architecture	20
A.1 Servers and configuration.....	20
A.2 Linux configuration	21
A.3 Network configuration.....	22
A.4 Network switches and adapters.....	22
A.5 Fiber channel switches and adapters.....	23
A.6 Storage specifications and connectivity	24
A.7 SAN and zoning.....	25
A.8 Oracle architecture	25
A.9 Testing methodology and monitoring tools.....	25
A.10 ORION IO profiling and test results	25
A.11 Conclusion	27
B Technical support and resources	28
B.1 Additional resources	28

Executive summary

Transaction oriented applications (such as financial transactions, order entry, and retail sales) perform the routine transactions that support these businesses. They require large amounts of storage and the ability to manage a vast number of transactions with sub-second response times.

Performance bottlenecks are not easily mitigated because there are many different components of the infrastructure stack that can contribute to the sub-optimal performance metrics. One component that is normally looked at first is the storage system and the delivered IOPS. Traditional and legacy storage is typically hard disk drive (HDD) based and cannot deliver the requirements of high-demand online transaction processing (OLTP) systems. Methods like disk overprovisioning, short stroking and manual performance tuning are commonly used to meet and troubleshoot performance requirements and issues. Performance tuning usually requires the greatest effort and cost because of application changes and a significant amount of testing before the changes are released into production. As a result, application changes are typically looked at last.

Dell EMC™ SC Series Storage offers flash technology that simplifies performance tuning. An all-flash SC Series array provides a storage medium with extreme performance characteristics. Unlike other mid-range all-flash solutions, SC Series all-flash arrays can be configured to distinguish multiple types of solid-state disks (SSDs) and use the SSDs to their fullest ability. SC Series all-flash arrays with this configuration are referred to as all-flash optimized arrays and are more cost effective than pure HDD configurations. These arrays can also be configured as hybrid arrays containing a mixture of SSDs and HDDs, or configured with only HDDs. Any of the configurations are quick to deploy, provide good value, have expansion potential and are economically priced for SMBs.

The reference architecture in this paper highlights the Dell EMC SC5020, but can be used with any SC series array for Oracle OLTP environments. Results of several performance tests against the storage sub-system are also included. As expected, the results show exceptional performance of the SC5020 which was able to sustain and deliver approximately 2ms latency and over 255K IOPS in a 70/30 random read/write mix.

Appendix A contains results from a similar reference architecture with the SC4020.

Audience

This document was written for DBAs, system administrators and storage administrators seeking a reference architecture for using Oracle OLTP on an SC series array.

Prerequisites

Readers should have formal training or advanced working knowledge of the following:

- RAID, Fibre Channel, multipathing, serial-attached SCSI (SAS), and IP networking administration
- Operating and configuring Dell EMC SC Series storage arrays
- Installing, configuring, and administering an OS
- General IP networking concepts
- Dell EMC server architecture and administration
- General understanding of SAN technologies
- Operating Dell EMC Storage Manager (DSM)
- Installing, configuring, and administering Oracle 11g and 12c; ASM, RAC, or single instance database
- Read and understand the content of the best practice guide for SC series storage.

1 Solution architecture

The reference architecture and configuration in this section supports an Oracle OLTP environment. It is provided as a guide rather than an explicit set of absolutes. Successful implementations are based on many factors. For a detailed discussion, see the [Dell EMC SC series arrays and Oracle](#) best practice document. For a list of implementation requirements, refer to documentation of components within the target environment, for example:

- Oracle Linux installation documentation
- Dell Storage Center Administration guide
- Brocade switch documentation
- Cisco® switch documentation
- Oracle documentation
- Dell EMC SC series storage Linux best practices
- Dell EMC SC series storage and Oracle best practices

The components used in the lab environment were:

Hardware components

- Dell PowerEdge™ R730
- Cisco Ethernet switch
- Brocade® Fabric channel switch
- Dell EMC SC5020

Software components

- Dell SCOS 7.2.10.101.01
- Dell Enterprise Manager 2017 r1, build 17.1.1.350
- Microsoft® Windows Server® 2012 R2 Enterprise, 64-bit
- IOmeter 1.1.0




Figure 1 Reference architecture

1.1 Servers and configuration

One two-socket 2U Dell PowerEdge R730 rack server was used in the environment. It was used to:

- Host Dell EMS DSM to simplify the management of the storage environment and provided statistics on the underlying storage
- Simulate the synthetic I/O loads using IOMeter.

Table 1 Dell Enterprise Manager and IOMeter server

Component	Description
Model	PowerEdge R730
CPU	Intel® Xeon® E5-2660 v3 @ 2.60GHz, 2 processors, 10 cores ea
Memory	128 GB RAM (2 x 64GB DDR4, 2133 MHz DIMMs)
OS	Microsoft® Windows Server® 2012 R2 Enterprise 64-bit
BIOS version	2.4.3
HBAs	Four QLogic® QLE2662 16Gb 2 port FC adapters
HBA firmware	9.2.3.20
Boot	SAN boot
Integrated NIC1	Intel 4P X520/4P I350 rNDC 10g
IPv6	Not enabled
IPv4	Enabled

1.2 Ethernet network configuration

A single management/public network was configured in the environment and used for normal public communication as well as between the R730 PowerEdge server and the SC5020. Version 4 internet protocol (ipv4) and default network ports were configured.

1.3 Ethernet network switch and adapter

The public network utilized a Dell PowerConnect™ 5548 switch. The R730 server had an integrated Intel 4P X520/4P I350 rNDC Ethernet adapter. Two of the four ports were 10GbE and two were 1 GbE ports. Only one of the two 10GbE ports were used with full duplex and RX/TX flow control.

Table 2 Server network adapter characteristics

Component	Description
Model	Intel 4P X520/4P I350 rNDC
Firmware version	15.0.28

1.4 Fibre channel switch and adapters

A single Brocade 300 switch with 24-16 Gb F-ports was deployed in the environment to host the I/O load. In an environment that requires high availability, deploy two fabric switches to mitigate fabric switch failure in a single fabric switch environment. For information on deploying two fabric switches, see [Dell EMC SC series arrays and Oracle](#) best practices.

The R730 server had four QLogic QLE2662 16Gb 2 port Fibre Channel adapters.

Table 3 Server HBA characteristics

Component	Description
Model	QLE2662
Description	PCI-Express Dual Channel 16Gb Fibre Channel HBA
Driver version	STOR Miniport 9.2.3.20
Firmware version	8.03.01
Optrom_bios_version	3.44
Optrom_efi_version	6.12
Optrom_fcode_version	N/A
Optrom_fw_version	8.03.01

Table 4 QLogic settings

Component	Description
Frame size	2048
Loop reset delay	5
Adapter hard loop ID	Disabled
Hard loop ID	0
Spinup delay	Disabled
Connection options	1
Fibre Channel tape support	Enabled
Data rate	2
LUNs per target	256
Enable LIP reset	Yes
Enable LIP full login	Yes
Enable target reset	Yes
Login retry count	60
Port down retry count	5
Link down timeout	30
Operation mode	0

Component	Description
Interrupt delay timer	0
EV Controller Order	Disabled
Execution Throttle	256

1.5 SAN and zoning

The SC series array was configured with two controllers that had multiple front-end ports zoned across two separate VSANS as shown in Figure 1. This two controller configuration with multiple front-end ports is recommended to avoid a single point of failure. Virtual port mode was also configured in the array to improve increased connectivity and improved redundancy.

Eight soft zones were configured in the environment.

In an environment that requires high availability, deploy two fabric switches and place each VSAN on a separate dedicated switch.

1.6 Multipathing

Multipathing was configured to ensure front-end redundancy and eliminate single points of failure of I/O paths.

1.7 Dell EMC SC 5020 configuration

The 3U, two controller, SC5020 was configured in FC virtual port mode with 15 write-intensive SSDs. One SC420 enclosure was configured with 15 additional write-intensive SSDs and cabled to the SC5020.

Table 5 SC5020 array configuration

Component	Description
SCOS	7.2.10.101.01
Processors	Intel Xeon Processor E5-2630 v3, 2.4 GHz, 8 cores, 8 MB cache
Memory	42.35 GB per array
Controllers	Two per array. Each controller had a 16 Gbps quad-port HBA
Battery-backed controller write cache	1 GB per array (512 MB per controller)
Expansion enclosures	One SC 420 enclosure for 12 Gb SAS
Internal Storage	Head: 15 x 2x5" write intensive 800 GB SSDs; enclosure: 15 x 2x5" write intensive 800 GB SSDs in tier 1 (29 + 1 spare)
Front-end connectivity	Eight 16 Gb FC ports
Back-end connectivity	Four 12 SAS ports
NAS connectivity	None-configured

DSM was used to perform the following configuration:

- A single logical server object was created for the physical server. All four HBAs in the physical Dell R730 Power Edge server were added to the server object.
- Only one disk pool, the default disk folder named **Assigned**, was created to ensure one virtual pool of storage was created from the 30 SSDs (29 active drives and one spare drive).
- Redundancy level and datapage size were left at the default values of **redundant** and **512 KB** respectfully.
- A single tier of storage, Tier 1, was defined. RAID levels within the tier 1 were left at the default values: **RAID 10** and **RAID 5.-9**.
- The default standard profile was used for tier 1.
- Snapshot profiles were not configured in the environment.
- Read and write cache were managed globally, with write cache disabled and read cache enabled.
- Eight 1.121 TB volumes were created for testing. Point-to-point mapping between each FC port and each volume was used. The eight volumes were balanced between both controllers. For example, 4 volumes were running on the lead controller while the remaining 4 volumes were running on the peer controller. Data reduction with deduplication and compression was not selected. In an environment that requires high availability, it is suggested that volumes be mapped to multiple or all available FC ports.

2 Testing methodology and monitoring tools

A number of IOMeter tests were run in the test environment and performance metrics were gathered from each one. The goal of the tests was to measure SC5020 storage subsystem I/O performance when simulating synthetic OLTP workloads.

Tests were broken down into types of I/O loads. Within each test, queue depths and block sizes varied in a geometric sequence from 2 to 128 (queue depth) and 512 bytes to 2048 Kb (block sizes).

2.1 I/O profiling

The objective of the I/O profiling was to measure the performance characteristics (IOPs, latency and throughput) of the SC5020 and the PowerEdge R730 while varying the queue depth and block size. IOMeter test scenarios were based on:

- Random 70/30 read/write mix
- Random 80/20 read/write mix
- Random 50/50 read/write mix
- Random read
- Random write
- Sequential read
- Sequential write

2.2 Test results

For brevity, a subset of the test types and results (IOPS, latency, and throughput) are presented. Emphasis is on 70/30 random read/write mix and sequential read tests for block sizes 4k, 8k, 16k, 32k and 64k.

With respect to throughput of the 70/30 random read/write tests, queue depths and block sizes were increased until a maximum throughput level was reached. Figure 2 shows a chart of the test data collected and illustrates that when queue depth of each server FC HBA port is set to 32, optimal throughput can be achieved. With queue depths greater than 32, additional throughput potential is possible; however the throughput rate of return starts to decline until throughput saturation is somewhat realized when queue depth of each server FC HBA port is set to 128.


Random read/write 70/30, MBs, RAID10

Figure 2 Throughput with varying block sizes and queue depths for 70/30 random read/write tests

The 100 percent sequential read test produced similar findings. Figure 3 shows that when each server FC HBA port is set to 32, optimal throughput rates are realized for block sizes 4k, 8K, 16k, and 32K. For these block sizes, marginal additional throughput is possible with greater queue depths. When block size is 64K, optimal throughput is reached when queue depth is set to 16 for each server FC HBA ports, and marginal throughput increases are possible with queue depths greater than 16.


Sequential read, MBs, RAID10

Figure 3 Throughput with varying block sizes and queue depths for 100 percent read tests

Figures 3, 4, and 5 compare throughput test results of 80/20, 70/30, and 50/50 random read/write tests. One interesting conclusion is that throughput is greater for 64 K blocks than for 128KB blocks when queue depth is set to 32 or greater.

Diminishing throughput rates: random read/write 50/50, MBs, RAID10

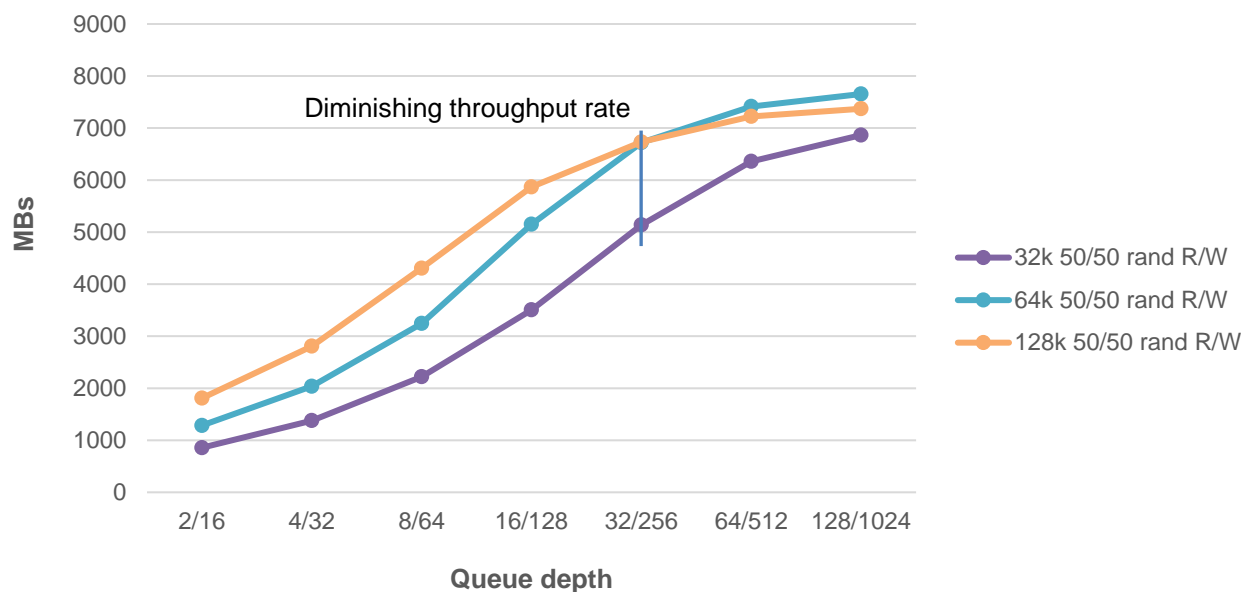


Figure 4 Diminishing throughput with 50/50 read/write mix and 64K vs 128K block size

Diminishing throughput rates: random read/write 70/30, MBs, RAID10




Figure 5 Diminishing throughput with 70/30 read/write mix and 64K vs 128K block size

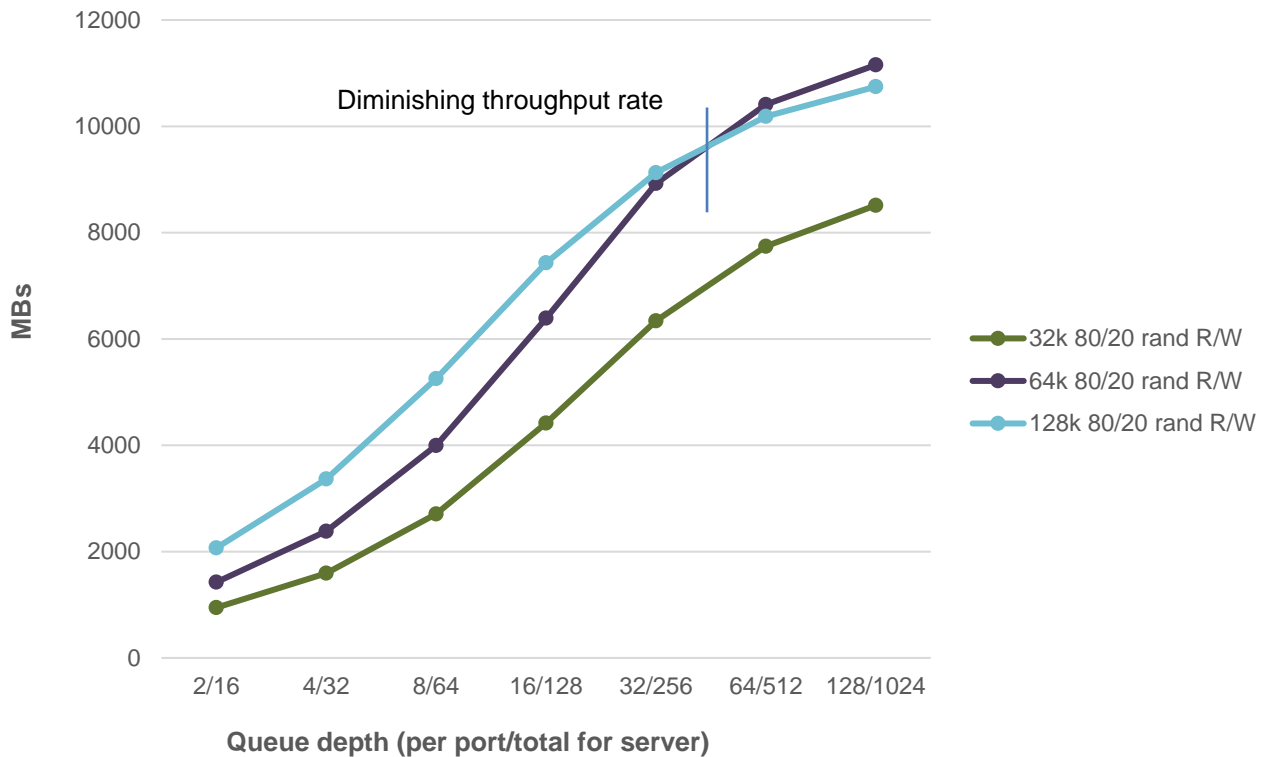
Diminishing throughput rates: random read/write 80/20, MBs, RAID10

Figure 6 Diminishing throughput with 80/20 read/write mix and 64K vs 128K block size

This decrease in throughput is not caused by the SC5020 array, but by how the server FC HBA manages I/O.

Summary of throughput tests:

- As expected, throughput increases as queue depths and block sizes increase
- Start with a queue depth of 32 as a best practice
- Block sizes of 4K, 8K, 16K, and 32K provide excellent throughput

In addition to excellent throughput rates, IOPS experienced the same point of diminishing returns when queue depth exceeds 32 per server FC HBA port.

When using queue depths of 128 and lower, the SC5020 is able to return high IOPS with 4K, 8K, 16K, and 32K blocks in a 70/30 random read/write mix. While 8K blocks and a queue depth of 128 per server HBA port returned 256K IOPS, 4K, 16K and 32K blocks returned 256K, 251K, 243K IOPS respectfully with the same queue depth setting.

For IOPS, increasing the queue depth from 16 to 32 for 8K blocks generates an increase in IOPS and throughput of approximately 50K and approximately 400MB respectfully. While an increase in queue depth from 32 to 64 for 8K blocks increase IOPS and throughput by approximately 30K and approximately 200MB respectfully.

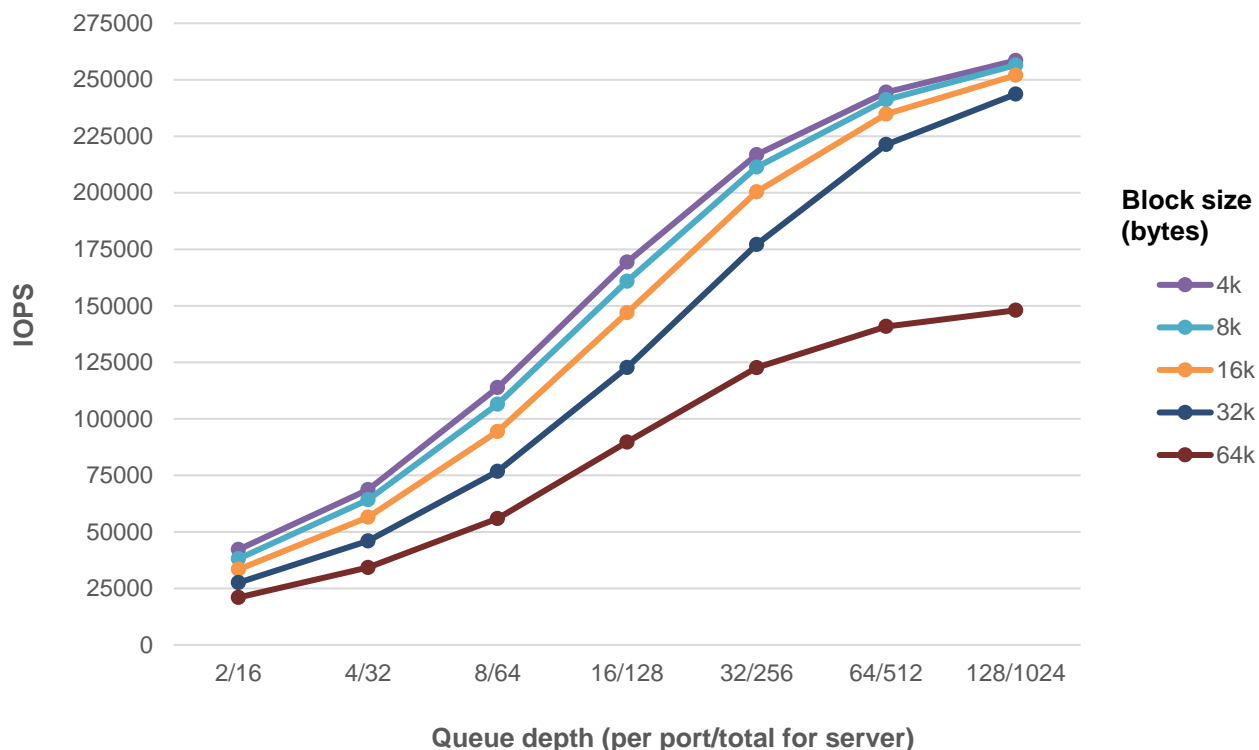
Random read/write 70/30, IOPS, RAID 10

Figure 7 Random read/write 70/30 IOPS test

Similar IOPS resulted with the 100 percent sequential read tests as seen in Figure 8. When using queue depths of 128 and lower per server HBA port, the SC5020 array is able to return exceptional IOPS with 4K, 8K, 16K, and 32K blocks. With 64K blocks and queue depth of 128, IOPS are about 47 percent less than with 32K blocks, but still very attractive at 200K.


Sequential read, IOPS, RAID10

Figure 8 IOPS with sequential 100 percent sequential read

If the server HBA ports have a queue depth of 32 and block size is 32K or less, IOPS still improve, but the change is somewhat the same. This can be seen in Figure 8 by the convergences of IOPS for block sizes less than 64 and server HBA port queue depths 32/256 and greater. Since Oracle OLTP databases generally have database block sizes of 8K, or 16K, and Oracle redo logs should use 4k blocks, if the server HBA port queue depths are set to 32/256, the database should provide good performance because the storage sub-system will return approximately 350K IOPS for OLTP processing.

SC5020 arrays are also able to generate excellent average latencies with high IOPS. As seen in Figure 9, at a queue depth 32/256, latencies are below 1.5 milliseconds for block sizes under 64k in a 70/30 random read/write mix. With a 64K block, the SC5020 array yields approximately 2 milliseconds under the same synthetic OLTP load and queue depth value.

70/30 random read/write, average latency, RAID10

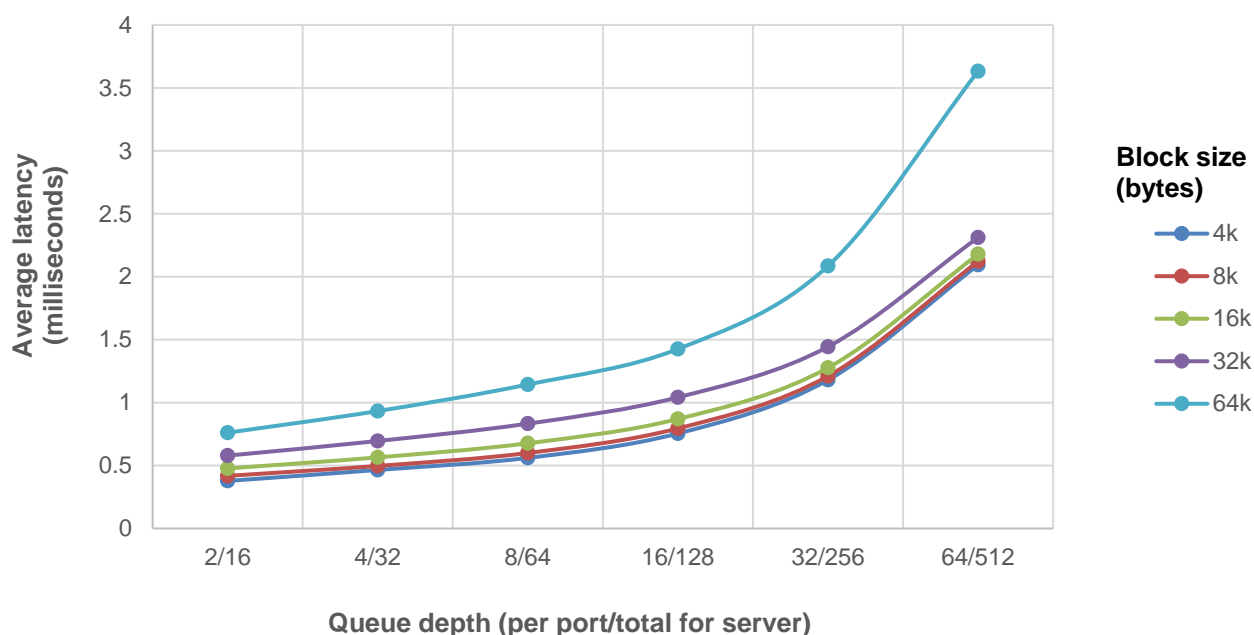


Figure 9 Average latencies with 70/30 random read/write load.

In a 100 percent sequential read test and queue depth set to 32/256, latencies are under 1.5 milliseconds for block sizes under 128k. When block size is under 64k, latencies are under 0.8 milliseconds as shown in figure 10.




Figure 10 Average latencies with 100 percent sequential read load

3 Conclusion

This paper presents how a single SC5020 all-flash array can be deployed for a cost-effective storage solution for an Oracle OLTP environment. Test results for a synthetic OLTP 70/30 random read/write load show that the SC5020 can deliver exceptionally low latencies, high IOPS and high throughput rates with a variety of block sizes and queue depths settings with room to spare. For example, Figure 11 shows that the SC5020 array delivers approximately 211K IOPS with an average latency of approximately 1.2 milliseconds when server HBA port queue depths are set to 32 and database block size is 8K. If queue depth changes to 64 with 8k blocks, an SC5020 will deliver approximately 241K IOPS and an average latency of approximately 2.1 milliseconds.

Random read/write 70/30 - IOPS vs latency, RAID 10

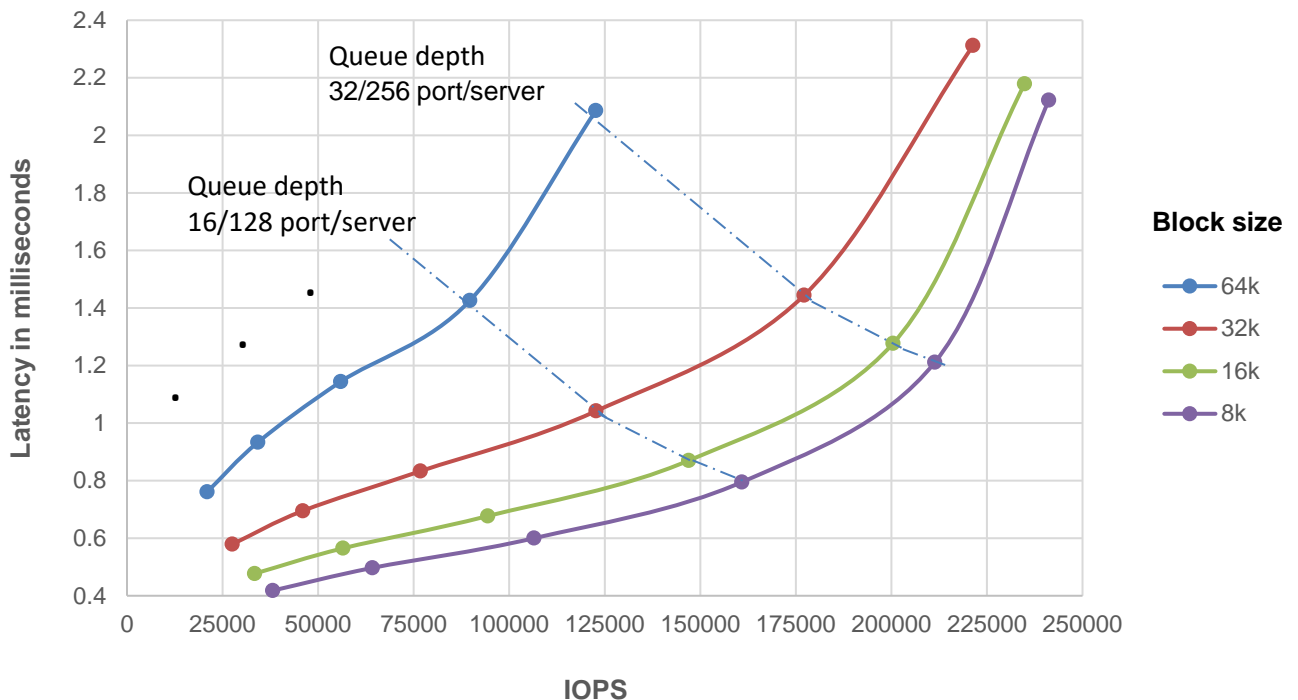


Figure 11 70/30 random read/write, latency and IOPS chart

Test results for a synthetic sequential read test show similar performance characteristics: exceptionally low average latencies, high IOPS and high throughput rates with a variety of block sizes and queue depths settings with room to spare. Figure 12 shows over 362K IOPS with an average latency of approximately 1.4 milliseconds when queue depths are set to 64/512 and database block size is 8K. If queue depth changes to 32/256, an SC5020 will deliver approximately 352K IOPS and an average latency of approximately 0.7 milliseconds.

Sequential read - IOPS vs latency, RAID 10




Figure 12 100 percent sequential read, latency and IOPS chart

In summary, this reference architecture shows that the SC5020 can be used to deploy a well performing Oracle OLTP database. With 8K blocks and queue depth 64/512, SC5020 delivers approximately 350,000 IOPS with average latencies under 1.5 milliseconds, which are attractive for mission critical OLTP Oracle applications.

A Dell EMC SC4020 reference architecture

This appendix lists information regarding a similar reference architecture for the SC4020 and Oracle OLTP.

Hardware components

- Dell PowerEdge R730xd
- Dell PowerEdge R710
- Dell PowerEdge R620 (Dell Enterprise Manager)
- Cisco Ethernet switch
- Brocade Fabric channel switch
- Dell SC4020

Software components

- Oracle Unbreakable Enterprise Linux Kernel 6.5
- Oracle 11g (11.2.0.4) Grid Clusterware
- Oracle 11g (11.2.0.4) RDBMS single instance standalone
- Dell SC4020 Storage Center OS (SCOS) 6.5.10
- Dell Enterprise Manager 2015 r1, build 15.3.1.110
- Microsoft Windows Server 2008 R2 Enterprise SP1, 64-bit

A.1 Servers and configuration

Three servers were used in the environment: two for the deployment of Oracle 11g, where each server hosted an 11.2.0.4 RDBMS installation, and 11.2.0.4 Grid Infrastructure. The third server hosted Dell Enterprise. Dell Enterprise Manager was not a required component of the infrastructure, but was used to simplify the management of the storage environment and provide statistics on the underlying storage.

Table 6 Dell Enterprise Manager application server - optional

Component	Description
Model	Dell PowerEdge R620
System Revision	1
CPUs	Intel Xeon E5-2670 v2 @ 2.50 GHz, 2 processors, 10 cores/ea
Memory	129 GB RAM (8x16GB DDR-3, 1600 MHz DIMMs)
OS	Microsoft Windows Server 2008 R2 Enterprise SP1, 64-bit
BIOS version	2.4.3
Firmware Version (iDRAC)	1.66.65 (Build 07)
iDRAC 7 NIC	Dedicated
Lifecycle Controller Firmware	1.4.2.12
IDSDM firmware	N/A
HBAs	QLogic QLE2562 8Gb 2Port FC adapter
HBA firmware	3.21.04

Component	Description
Boot	SAN boot
Integrated NIC 1	Intel 2P X540/2P I350 rNDC
IPv6 Enabled	No
IPv4 Enabled	Yes

Table 7 Oracle 11gR2 Grid Infrastructure, ASM

Component	Description
Model	Dell PowerEdge R730xd
System Revision	1
CPUs	Intel(R) Xeon(R) CPU E5-2697 v3 @ 2.60GHz, 2 processors, 14 cores/ea
Memory	256 GB RAM (16x16GB DDR-4, 2133 MHz DIMMs)
Swap	64GB
OS	3.8.13-16.2.1.el6uek.x86_64
BIOS version	1.2.10
Firmware Version (iDRAC)	2.10.10.10
iDRAC 8 NIC	Dedicated
Lifecycle Controller Firmware	2.10.10.10
IDSDM firmware	N/A
HBAs	QLogic QLE2532 8Gb 2Port FC adapter
HBA firmware	03.11.18
Boot	SAN boot
Integrated NIC 1	Intel® 2P X520/2P I350 rNDC
IPv6 Enabled	No
IPv4 Enabled	Yes

A.2 Linux configuration

Linux 3.8.13-16.2.1.el6uek.x86_64 was installed and configured per best practices

A.3 Network configuration

A single network was configured in the environment:

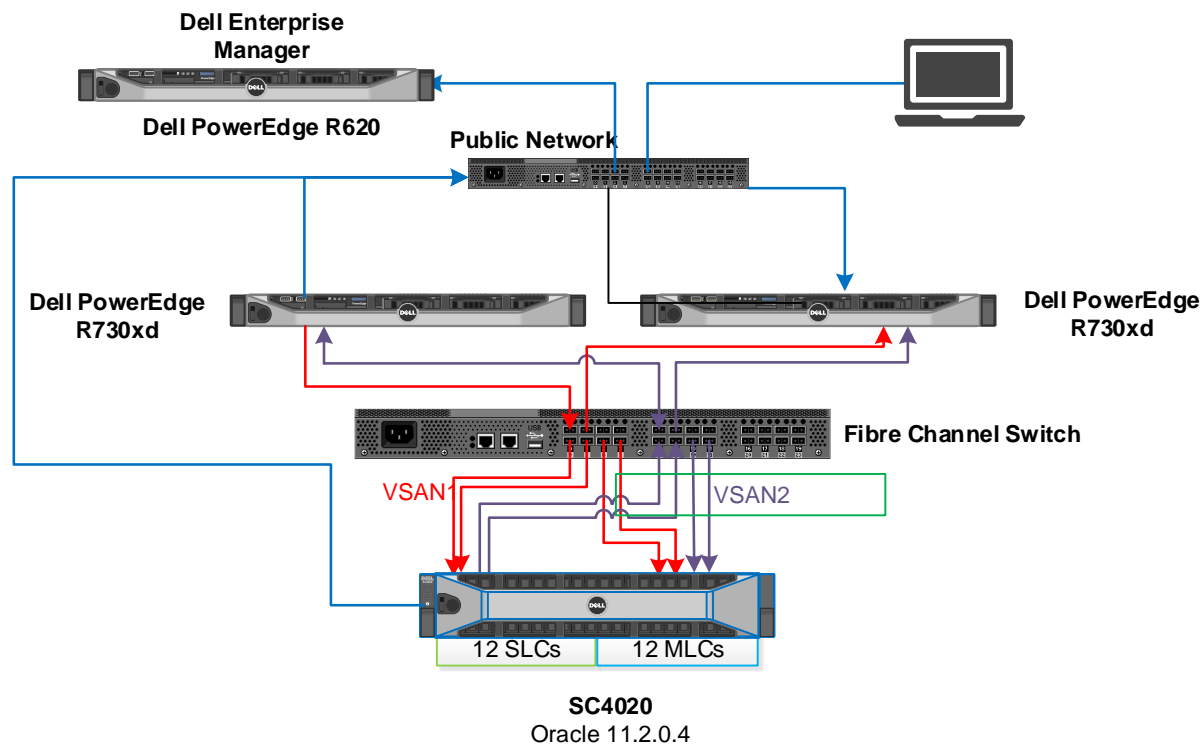


Figure 13 Network architecture

The public network was used for communication between all servers, Dell SCv4020, and users. It was also used by other applications accessing the environment, as well as Dell EM, and users administrating and managing the test environment. Version 4 internet protocol (ipv4) and default network ports were configured and all firewalls and SELinux were disabled in the environment.

A.4 Network switches and adapters

The public network was connected to a shared Powerconnect 5548 switch.

Each database server’s integrated NIC 1 is an Intel 2P X520/2P I350 rNDC 4 port adapter. Two of the four ports are 10 GbE and the other two were 1 GbE ports. Only one of the 2 1 GbE ports were used; network interface em3 with full duplex and RX/TX flow control.

Table 8 Server network adapter characteristics

Characteristic	Value
Model	Intel(R) 2P X520/2P I350 rNDC
Firmware version	15.0.28

A.5 Fiber channel switches and adapters

A Brocade 300 switch, designed for small to medium solutions, with 24 8Gb ports was deployed in the environment to host the Grid voting and RAC database files. In an environment that requires high availability, it might be prudent to deploy two fabric switches to mitigate fabric switch failure to in a single fabric switch deployment.

A Cisco MDS 9134 (DS-C9134-K9) multilayer fabric switch, designed for midrange to large enterprise environments, was deployed in the environment to host the Grid and Oracle home installations, as well as a san from boot volume. It provided line-rates of 4Gb and 10Gb ports. In an environment that requires high availability, it might be prudent to deploy two fabric switches to mitigate fabric switch failure to in a single fabric switch deployment.

Each RAC server had two QLogic ISP2532-based 8Gb Fibre Channel to PCI Express (rev 02) HBAs.

Table 9 Server HBA characteristics

Name	Value
model	QLE2562
description	PCI-Express Dual Channel 8Gb Fibre Channel HBA
driver version	8.04.00.11.39.0-k
firmware version	5.08.00 (90d5)
optrom_bios_version	3.24
optrom_efi_version	6.03
optrom_fcode_version	3.19
optrom_fw_version	7.01.00 32896

Table 10 QLogic settings

Name	Value
Frame Size	2048
Loop Reset Delay	5
Adapter Hard loop ID	Disabled
Hard Loop ID	0
Spinup Delay	Disabled
Connection Options	1
Fibre Channel Tape Support	Enabled
Data Rate	2
Luns per target	256
Enable LIP reset	Yes
Enable LIP Full Login	Yes

Name	Value
Enable Target Reset	Yes
Login Retry Count	60
Port Down Retry Count	60
Link Down Timeout	30
Operation Mode	0
Interrupt Delay Timer	0
EV Controller Order	Disabled
Execution Throttle	256

Table 11 Brocade switch characteristics

Name	Value
model	300
Port speed and type	24 x 8Gb F-Port

A.6 Storage specifications and connectivity

The 2U, two controllers, Dell EMC SC4020 was configured in FC virtual port mode with 24 SSDs (12 read-intensive MLCs and 12 write-intensive SLCs). No expansion enclosures were used.

Table 12 Dell EMC SC4020 configuration and characteristics

Name	Value
SCOS	6.5.10
Processors	Intel Xeon Processor E3-1265L v2, 2.500 GHz, 4 cores, 8MB cache
Memory	32GB per array (16GB per controller)
Controllers	2 per array. Each controller has 4 x 8Gb ports.
Battery-backed controller write cache	1GB per array (512MB per controller)
Expansion enclosures	None
Internal Storage	24 x 2.5" SSDs ((11 + 1 spare) x 372.61 GB MLCs in tier 1, (11 + 1 spare) x 1.46 TB SLCs in tier 2)
Front-end connectivity	8 x 8Gb FC ports
Back-end connectivity	4 x 6Gb SAS ports
NAS connectivity	None configured

A.7 SAN and zoning

An SC2040 was configured with two controllers with multiple front-end ports zoned across two separate VSANS as shown in Figure 3 Network architecture. This two controller configuration with multiple front-end ports is recommended to avoid a single point of failure should a controller fail. Virtual port mode was also configured in Storage Center to improve increased connectivity, and improved redundancy.

For additional information on virtual port mode, see the SC4020 Storage System Owner's Manual.

A.8 Oracle architecture

Oracle 11g GI, ASM, and RDBMS single instance standalone were used in the environment. No database was created as Oracle ORION was used to capture performance metrics. ORION is Oracle's performance measuring tool and resides in ORACLE_HOME/bin, and is part of the RDBMS installation.

Table 13 Oracle architecture

Component	Version
Grid Infrastructure	11.2.0.4
RDBMS	11.2.0.4
ASMLib	2.0.4-1.el6.x86-64

For the 11.2.0.4 standalone RDBMS environment, ORION was configured one each database server in file <test-name>.lun to use 4 x 200GB LUNs on each server,

```
/dev/multipath/lun1
/dev/multipath/lun2
/dev/multipath/lun3
/dev/multipath/lun4
```

A.9 Testing methodology and monitoring tools

A number of different ORION tests using different parameters were run in the environment. Performance metrics were gathered from each ORION test. The Orion scenarios tested were based on 70-30% OLTP read/write mix, with a varying queue depth.

During each run, ORION gathered and created a number of files containing trace data and performance metrics:

A.10 ORION IO profiling and test results

ORION was used to simulate OLTP transactions. The goal was to measure maximum IPS capability of the all-flash configuration. The tests used a combination of the below parameters.

Table 14 ORION test parameters

Parameter	Purpose
-run oltp	Direct ORION to simulate OLTP transactions
-run advanced	
-duration 1800	Length of a test run in seconds
-num_disks 11	Number of physical disks in the array that should be used in the simulation.
-size_small 8, 16	Size of the IO block in kbytes
-write 30	Percentage of writes performed during the simulation.
-verbose	Directive to generate more information
-matrix point	Orion test consisting of multiple data points
-type rand	Randomly distributed IOs
-num_small 32, 64, 128, 256	Number of outstanding small IOs

IOPS 8K vs 16K random 70/30 RW mix

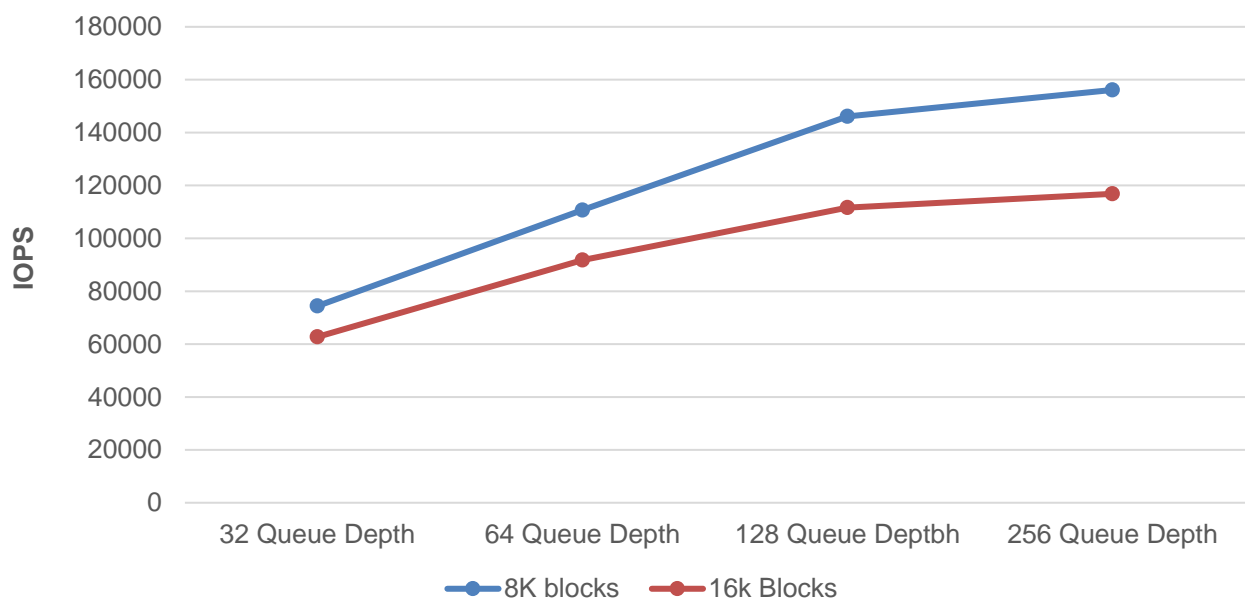


Figure 14 IOPS 8K vs 16K with varying queue depth




Figure 15 Average latencies with 8K and 16K blocks and varying queue depths

A.11 Conclusion

This SC4020 reference architecture demonstrated deploying an Oracle OLTP environment using two single instance standalone RDBMS installations in a 11.2.0.4 Grid Infrastructure deployment on two R730xd database servers. The architecture was placed under a series of Orion tests to simulate an OLTP workload and record the performance.

- The SSD based solution presented in this section is a cost effective alternative to deploying a HDD based infrastructure to accomplish the same. The reference architecture discussed in this paper is capable delivering:
- Approximately 20 TB raw storage (11 x 372.61 GB SSDs tier 1, and 11 X 1.46 TB SSDs tier 2), up about 50% more storage than tradition HDD based configuration (11 x 279 GB 15K) + (11 x 951.51 GB 7k)
- 156000 IOPs with 8K block IO, 70/30 read write mix, with queue depth 256, and 11 write intensive SSDs
- Sub second read and write latency

B Technical support and resources

[Dell.com/Support](#), [EMC Online Support](#), and [E-LAB NAVIGATOR](#) are focused on meeting customer needs with proven services and support.

[Dell TechCenter](#) is an online technical community where IT professionals have access to numerous resources for Dell software, hardware and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell Storage platforms.

B.1 Additional resources

Referenced or recommended Dell EMC publications:

- [Dell Storage Center with Red Hat Enterprise Linux \(RHEL\) 6x Best Practices](#)
- [Dell Storage Center with Red Hat Enterprise Linux \(RHEL\) 7x Best Practices](#)
- [Flash-optimized Data Progression](#)
- [Using Oracle Database 10g's Automatic Storage Management with EMC Storage Technology](#)
- [PowerPath Family, version 6.x, CLI and System Messages Reference](#)
- [Dell SC Series All-Flash Arrays and Oracle OLAP Processing](#)
- [Dell Storage Center OS 7.0 Data Reduction with Deduplication and Compression Solution Guide](#) on Dell EMC TechCenter

Referenced or recommended Red Hat publications:

- [Red Hat Enterprise Linux 6 Storage Administration Guide](#)
- [Red Hat Enterprise Linux 6 DM Multipath, DM Multipath Configuration and Administration](#)
- [Deploying Oracle Database 12c on RHEL 7 - Recommended Practices](#)
- [Deploying Oracle Database 12c on RHEL6 - Recommended Practices](#)
- [Deploying Oracle Database 11g on RHEL 6 - Recommended Practices](#)
- [Deploying Oracle RAC 11g R2 Database on RHEL 6 - Recommended Practices](#)
- [Redhat customer portal support note 222473](#)

Other recommended publications:

- [Standard RAID levels](#)

Other recommended publications: (Oracle manuals):

- [Installing Oracle ASMLib](#)
- [Configuring Oracle ASMLib on Multipath Disks](#)
- [Configuring Oracle ASMLib on Multipath Disks](#)
- Oracle Database – Filesystem & I/O Type Supportability on Oracle Linux 6 ()
- [Using Oracle Database 10g's Automatic Storage Management with EMC Storage Technology.](#)
- [Oracle Database Storage Administrator's Guide 11g Release 1 \(11.1\)](#)

- [Oracle Automatic Storage Management, Administrator's Guide, 11gR2](#)
- [Oracle Automatic Storage Management, Administrator's Guide, 12cR1](#)
- [Oracle Database Concepts, 12c Release 1](#)
- [Oracle Database, Installation Guide, 11g Release 2 \(11.2\) for Linux, E47689-10](#)
- [Oracle Database Administrator's Reference 12c Release 1 \(12.1\) for Linux and UNIX-Based Operating Systems E10638-14](#)
- [Oracle Database Administrator's Guide 12c Release 1](#)
- [Oracle® Grid Infrastructure Installation Guide 11g Release 2 for Linux](#)
- [Oracle® Grid Infrastructure Installation Guide 12c Release 1 for Linux](#)
- [Oracle Database Data Warehousing Guide 12c Release 1 \(12.1\)](#)
- [Oracle Database Data Warehousing Guide 11g Release 1 \(11.1\)](#)
- [Oracle Linux Administrator's Guide for Release 6](#)
- [Oracle Linux Administrator's Guide for Release 7](#)

Other recommended publications: My Oracle Support Doc IDs:

- Oracle Support Doc 1077784.1: Can I create an 11.2 disk over the 2 TB limit"
- [Lun Size And Performance Impact With Asm \(Doc ID 373242.1\)](#)
- Oracle Support Doc 1601759.1: Oracle Linux 5 - Filesystem & I/O Type Supportability
- Oracle Support Doc 1487957.1: ORA-1578 ORA-353 ORA-19599 Corrupt blocks with zeros when filesystemio_options=SETALL on ext4 file system using Linux ()
- [ORA-15040, ORA-15066, ORA-15042 when ASM disk is not present in all nodes of a Rac Cluster. Adding a disk to the Diskgroup fails. \(Doc ID 399500.1\)](#)
- [How To Setup Partitioned Linux Block Devices Using UDEV \(Non-ASMLIB\) And Assign Them To ASM? \(Doc ID 1528148.1\)](#)
- [ASMFD \(ASM Filter Driver\) Support on OS Platforms \(Certification Matrix\). \(Doc ID 2034681.1\)](#)
- [How to Install ASM Filter Driver in a Linux Environment Without Having Previously Installed ASMLIB \(Doc ID 2060259.1\)](#)
- [How To Setup Partitioned Linux Block Devices Using UDEV \(Non-ASMLIB\) And Assign Them To ASM? \(Doc ID 1528148.1\)](#)
- [How To Setup ASM on Linux Using ASMLIB Disks, Raw Devices, Block Devices or UDEV Devices? \(Doc ID 580153.1\)](#)
- Oracle Linux 6 - ASM Instances Fail with 4K Sector Size LUN (Doc ID 2211975.1)
- Supporting 4K Sector Disks [Video] (Doc ID 1133713.1)
- Supporting ASM on 4K/4096 Sector Size (SECTOR_SIZE) Disks(Doc ID 1630790.1)
- New Block Size Feature for Oracle ASM and oracleasm-support (Doc ID 1530578.1)
- Alert: After SAN Firmware Upgrade, ASM Diskgroups (Using ASMLIB) Cannot Be Mounted Due To ORA-15085: ASM disk "" has inconsistent sector size. (Doc ID 1500460.1)
- Supported and Recommended File Systems on Linux (Doc ID 236826.1)