

EMC SCALEIO

High-Performance Converged Infrastructure for Oracle Database Deployments

ABSTRACT

Industry standard servers and storage have evolved into dependable, high-performance platforms suitable to handle a high percentage of today's Oracle Database deployments. However, the applications depending on Oracle Database require scale-out and availability attributes that generally surpass the capability of a single, simple commodity storage server. EMC ScaleIO™ is software that enables one to build a converged infrastructure of compute and shared-disk storage using simple industry-standard servers. A ScaleIO server SAN is complementary to Oracle Database performance and availability requirements. This paper introduces EMC ScaleIO™ and offers a case study of the power and flexibility of ScaleIO converged infrastructure in an Oracle Database deployment.

EMC WHITE PAPER

EMC²

TABLE OF CONTENTS

- ABSTRACT 1**
- INTRODUCTION..... 3**
- EMC SCALEIO PLATFORM SUITABILITY FOR ORACLE DATABASE 3**
 - What Is a Suitable Platform for Oracle Database? 3
 - Server SAN--An Alternative Choice for Oracle Database Storage and Server Provisioning 3
 - Converged Storage and Compute 3
 - Linear Scalability..... 4
 - Architectural Components 4
 - Elastic Storage 4
 - Oracle Automatic Storage Management with Elastic Storage 4
- CASE STUDY 5**
 - Scope..... 5
 - Platform Suitability Criteria..... 5
 - Test System Description 5
 - Shared Storage Configuration for Oracle Database 6
 - Performance Analysis 7
 - Concurrent Table Scan Test Results 8
 - OLTP-Style Test Results 9
 - Oracle Database 12c Performance Characterization..... 11
- SUMMARY..... 12**

INTRODUCTION

EMC ScaleIO™ is software that creates a server-based SAN from simple industry-standard storage servers to deliver elastic and scalable performance and capacity on demand. EMC ScaleIO™ transforms application servers' local disks into an enterprise-class, scale-out virtual SAN. Convergence realized through ScaleIO spares customers the complexity and costs associated with a three-layer architecture of servers, fabric, and traditional external storage. While providing robust, elastic shared block storage and offering high performance, ScaleIO provides a complete, on-premises storage software solution that offers the same low-cost storage economics as the in-house software of the largest cloud players while meeting the demands of large enterprises.

Flexibility and elasticity are at the core of EMC ScaleIO™, which offers customers the freedom to use any type of server hardware. EMC ScaleIO™ manages itself, automatically rebalances capacity when servers are added or removed, and self-heals to automatically recover from server or disk failures. With EMC ScaleIO™, a virtual SAN can be yours—without the headache, complexity, and cost of traditional external storage.

EMC SCALEIO PLATFORM SUITABILITY FOR ORACLE DATABASE

What Is a Suitable Platform for Oracle Database?

Platform requirements for Oracle are well understood and Oracle practitioners have come to appreciate the attention to detail that must be given when selecting a particular architecture for an Oracle database deployment. For example, some of the Oracle databases in a given IT deployment might be Standard Edition, while others are Enterprise Edition. Further, some deployments might require shared disk for either Real Application Clusters or failover high-availability. Historically, shared disk storage equated to either SAN or NAS. At times, IT engineers have had to make difficult decisions about whether to co-mingle databases of widely varying requirements. For example, does it make sense to burden a central storage array—optimized for shared disk—with databases that have no shared-disk requirements? What about varying use-cases?

Consider, for example, the sometimes-difficult choice of comingling both production and development databases just because they both have shared-disk requirements? If comingling is not an option, what is the alternative? Is deploying another high-end central storage array the right answer? These are the difficult questions IT managers are faced with today—and IT managers are seeking viable options.

Choice is a necessary ingredient for success in today's complex IT shop with critical Oracle Database deployments.

Server SAN – An Alternative Choice for Oracle Database Storage and Server Provisioning

As mentioned above, platform requirements for Oracle are well understood. In fact, most of the industry-standard storage servers on the market today are quite sufficient for Oracle database. Indeed, all of the major hardware vendors offer small form factor (e.g., 2U) storage servers that often support as many as 24 2.5" hard drives that can range in capability from 7200 RPM SATA to 10,000 RPM SAS and even Solid State Disks. Moreover, these storage servers have ample PCI Gen3 slots to support PCI Flash cards for extreme low-latency I/O. However, these are not just *storage servers*, these are very powerful **servers**.

Indeed, in the first quarter of 2014, 2U storage servers are available with 2 Intel Xeon E5-2600v2 processors with as many as 12 processor cores per CPU. Moreover, every core in these modern processors is able to process more TPC-C transactions per minute than the largest, most expensive refrigerator-sized servers of the recent past. So if all the performance ingredients are in place, what's missing? Why aren't more non-critical, development or test databases deployed on these incredibly powerful storage servers? The answer is quite simple: enterprise Oracle IT shops still need SAN technology—and will for the fathomable future. They need SAN—or at least external storage—technology because even databases with no shared disk requirement cannot suffer unavailability of data when a server becomes unavailable. This all leads to an important question; just what is a SAN anyway?

From a pure functionality perspective a SAN serves one main purpose where Oracle Database is concerned: to present network-attached, coherent block storage. That's all good and true, but for some use-cases, one could imagine how nice it would be to have all the aggregate storage in a collection of industry-standard storage servers "look" like a SAN—a server SAN, if you will. EMC ScaleIO™ does just this.

Converged Storage and Compute

ScaleIO converges storage and compute resources into a single-layer architecture, aggregating capacity and performance and simplifying management. With ScaleIO, storage is just another application running alongside other applications, and each server is a building block for the global storage and compute cluster. Converging storage and compute simplifies the architecture and reduces cost without compromising on any benefit of external storage. ScaleIO enables the IT administrator to individually manage the entire data center stack, improving operational effectiveness and lowering operational costs.

Linear Scalability

ScaleIO is designed to linearly scale from three to thousands of nodes. Unlike most traditional storage systems, as the number of storage devices grows, so do throughput and IOPS. The scalability of performance is linear with regard to the growth of the deployment. Whenever the need arises, additional storage and compute resources (i.e., additional servers and drives) can be added modularly. Storage and compute resources grow together so the balance between them is maintained. Storage growth is therefore always automatically aligned with application needs.

Architectural Components

ScaleIO consists of two major functional components: ScaleIO Data Client (SDC) and ScaleIO Data Server (SDS). The SDC is a block device driver that exposes ScaleIO shared block volumes to applications. The SDC runs locally on any server that requires access to the block storage volumes in the cluster. The local application processes issue I/O requests and the SDC fulfills it regardless of where the particular blocks reside. The SDS is a software component installed on each server that contributes local storage to the overall ScaleIO storage pool. The SDS serves incoming read and write requests from any of the SDCs within the cluster. The SDC possesses full knowledge of the data locations throughout the cluster and always directs I/O requests to their correct destination SDS, whether on the same server or any other server. This way, I/O operations are never routed through a central routing point, preventing bottlenecks.

Elastic Storage

With ScaleIO, storage capacity and compute can be increased or decreased whenever the need arises. The system automatically rebalances data “on the fly” with no downtime. Additions and removals can be done in small or large increments. No capacity planning or complex reconfiguration due to interoperability constraints is required, which reduces complexity and cost. The ScaleIO system reconfigures itself as the underlying resources change; data is rearranged and spread evenly on the servers to optimize performance and enhance resilience. All of this happens automatically without operator intervention. Figure 1 depicts adding 3 disks to a ScaleIO pool after which a ScaleIO rebalance occurs. This functionality is critical because it ensures that adding capacity does not create storage imbalances—the IOPS are spread across all storage devices.




Figure 1: Adding Storage to ScaleIO Results in Data Rebalancing. Adding Storage Adds I/O Bandwidth.

Oracle Automatic Storage Management with Elastic Storage

ScaleIO storage is both elastic and protected, therefore Oracle Database administrators can create simple Automatic Storage Management (ASM) disk groups, with EXTERNAL REDUNDANCY, consisting of a single, expandable ASM disk. ScaleIO naturally balances volume contents across newly added storage. To that end, growing an ASM disk group in the ScaleIO environment is as simple as adding space to the ScaleIO volume and then utilizing the ASM command to resize the disk group. Increasing the size of an ASM disk group in this manner is non-disruptive from both the ScaleIO and ASM perspective. After adding storage to the ScaleIO volume, the administrator simply executes the ALTER DISKGROUP RESIZE command. The administrator doesn't need to specify the size when executing the disk group resize option because Oracle Database will simply execute an operating system call to detect the new size of the block device (the ScaleIO volume). Single-disk ASM disk groups can grow through resize

operations until the disk is 16TB. Once single-disk ASM disk group reaches 16TB in size, the administrator must add another ASM disk to the disk group. Disks added to ASM disk groups should come from separate ScaleIO pools. The administrator simply creates another volume (which SDC presents as a block device at the operating system layer) from a separate pool and adds it to the ASM disk group as a new ASM disk. Depending on the size of disk being added the admin can choose whether or not to allow ASM to perform a rebalance. ScaleIO does not perform rebalance operations across pools so redundant rebalancing won't occur if the administrator decides to rebalance an ASM disk group after adding a ScaleIO volume from a different pool to an ASM disk group.

Snapshots

ScaleIO offers volume snapshot technology. Snapshots in ScaleIO are:

- Nearly immediate to create
- Thinly provisioned
- Writable
- Allow for associating groups of ScaleIO volumes into consistency groups.

Storage snapshot functionality is critical in today's complex Oracle environments. Snapshots enable administrators to test patching and updates against the full version of production data in isolation. Snapshots are also very helpful in offloading the processing overhead of backups to other hosts. Datacenter operations, such as development and test activity, are also greatly improved with storage snapshots.

CASE STUDY

Scope

ScaleIO is enabling software free of fixed boundaries and presumptions of underlying hardware. In other words, the performance attributes of any given ScaleIO deployment is only a reflection of the components used to build the solution. The goal in this Case Study was neither to demonstrate limits of ScaleIO nor to suggest a best practice. Instead, EMC engineers aimed to provide an example of how simple, inexpensive and small ScaleIO configurations can provide Oracle Database a platform with very powerful performance attributes and robust data protection.

This Case Study is focused on a hyper-converged deployment of Oracle Database on ScaleIO. In other words, each host plays the role of Oracle Database server, ScaleIO data client and ScaleIO data server.

Platform Suitability Criteria

Suitable platforms for Oracle Database come in all shapes and sizes but there is a very fixed list of absolutes that the platform must deliver:

- **CPU.** The hosts must have a suitable ratio of processor bandwidth to I/O capability. I/O driven by SQL processing is also relatively CPU intensive.
- **I/O Throughput.** The platform cannot possess data-flow bottlenecks.
- **I/O Latency.** The IT industry has entered the age of low-latency commercial computing. High I/O service times for random I/O are simply not tolerated.

To assess suitability of ScaleIO with these three performance-pillars for Oracle Database in mind, EMC engineers engaged to test ScaleIO with Oracle Real Application Clusters. The testing focused on IOPS-capability and I/O latency, with OLTP-style workloads, as well as table scan throughput with data warehouse-style workloads.

Test System Description

To model a ScaleIO deployment optimized for Oracle Database, EMC engineers built an eight-node Oracle Linux 6.4 test configuration with components depicted in Figure 2:


<ul style="list-style-type: none">• 8 x 2U Cisco C240 M3 Server, each with:<ul style="list-style-type: none">○ 2 x E5-2690 Intel Xeon processors.○ 1 x 2.5" SATA drive for the Operating System○ 1 x 2.5" SAS 10K RPM drive○ 1 x EMC XtremSF SLC 700 PCIe Flash card○ 6 x 960GB 2.5" SATA cMLC SSD drives—expandable to 23 SSDs per host in this configuration○ 1 x Dual port 40 Gb QDR Infiniband IB HCA• Mellanox Infiniband switch	<h1>17U</h1>	
---	--------------	---

Figure 2: Test Configuration Description.

Oracle Database 12c Grid Infrastructure was the chosen release of clusterware for Real Application Clusters since it supports testing both Oracle Database 11g Release 2 (11.2.0.4) and Oracle Database 12c instances concurrently.

In preparation for the Oracle testing, EMC Engineers configured a single ScaleIO Protection Domain and three pools with the following and attributes:

1. Pool: "spindle". This pool is a collection of each the 2.5" 10K RPM SAS drive in each host.
2. Pool: "pcie". This pool consists of each XtremSF SLC PCIe card from each host.
3. Pool: "ssd". This pool is comprised of the 6 cMLC drives from each of the hosts

Figure 3 shows a screenshot of the ScaleIO Dashboard displaying information about the pools described above.

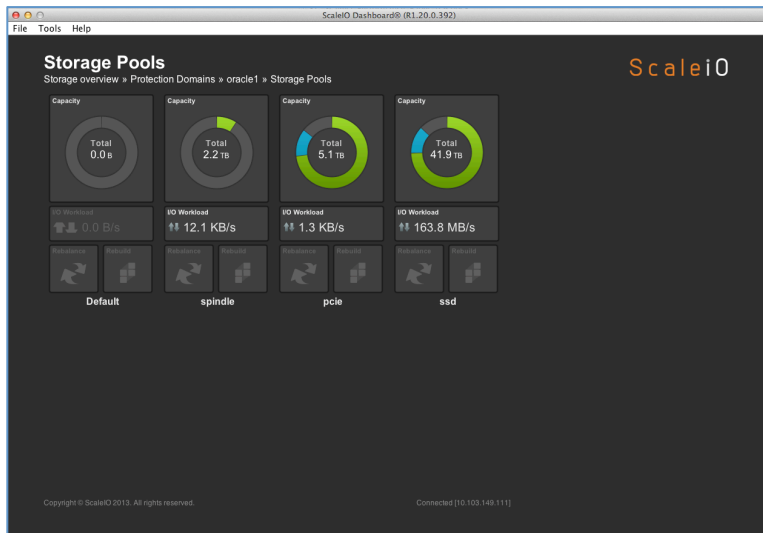


Figure 3: ScaleIO Dashboard Depicting ScaleIO Storage Pool Monitoring.

Shared Storage Configuration for Oracle Database

In order to test Oracle RAC one must provision shared storage via either ASM, cluster file system or NFS mounts. EMC engineers chose to test with ASM from the 12c Grid Infrastructure distribution.

To prepare for importing the ScaleIO elastic storage into ASM, the ScaleIO SDC agent presented a volume from each of the three pools mentioned above. The presentation of ScaleIO disk space appears to the administrator as a simple block device. For illustrative purposes please refer to Figure 4, which shows:

- A listing of the three block devices (one with a small partition) that map to the volumes presented by ScaleIO SDC.
- A listing of the udev(7) rules file associating a consistently named block device with the underlying ScaleIO SDC device.
- A listing of the three block device entries in the Linux /dev directory that map to the ASM disks.

```
$
$
$ grep scini[abc] /proc/partitions
250      0 1996488704 scinia
250     16 16777216000 scinib
250     32 109051904  scinic
250     33  52429808  scinic1
$
$ more /etc/udev/rules.d/99-asm-devices.rules
KERNEL=="scini*[0-9]", SUBSYSTEM=="block", PROGRAM="/bin/scaleio/ecs/drv_cfg --query_block_device_id --block_device /dev/%k", RESULT=="637b242a1247e909-29
92d83000000000", NAME="asm-grid", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="scini*", SUBSYSTEM=="block", PROGRAM="/bin/scaleio/ecs/drv_cfg --query_block_device_id --block_device /dev/%k", RESULT=="637b242a1247e909-2992d8f
b00000002", NAME="flash_ssd", OWNER="oracle", GROUP="dba", MODE="0660"
KERNEL=="scini*", SUBSYSTEM=="block", PROGRAM="/bin/scaleio/ecs/drv_cfg --query_block_device_id --block_device /dev/%k", RESULT=="637b242a1247e909-2992d8f
800000001", NAME="flash_pcie", OWNER="oracle", GROUP="dba", MODE="0660"
$
$ ls -l /dev/asm-grid /dev/flash_ssd /dev/flash_pcie
brw-rw---- 1 oracle dba 250, 33 Apr 17 06:29 /dev/asm-grid
brw-rw---- 1 oracle dba 250,  0 Apr 17 06:29 /dev/flash_pcie
brw-rw---- 1 oracle dba 250, 16 Apr 17 06:29 /dev/flash_ssd
$
$
$
```

Figure 4: Linux Operating System. Mapping ScaleIO Volumes as SDC Block Devices. Volumes Are Given Persistent Names With udev(7).

After the ScaleIO storage was provisioned, and the aforementioned task of configuring udev(7) rules were complete, EMC engineers moved on to import the disks into ASM and create disk groups. This process is as simple as invoking the Oracle ASM Configuration Assistant (ASMCA) and setting the ASM disk_string parameter to "/dev/flash*". Please refer to Figure 5, which shows:

- A small disk group called DATA, which was used for the Oracle Database 12c Grid Infrastructure objects. This disk group used partition 1 from the /dev/scinic SDC block device (see Figure 4).
- An approximately 15TB disk group called TIER2FLASH. Based on its size we can see it maps to the /dev/scinib ScaleIO SDC block device shown in Figure 4. This disk group is comprised solely of space from the 6 mCMLC SSD drives in each of the 8 hosts.
- A disk group of roughly 1.9TB called TIER1FLASH. This disk space is comprised of the aggregate of all PCIe Flash card capacity.

Notice how the disk groups were created with ASM EXTERNAL redundancy (Figure 5). This is because ScaleIO protects data with mesh-mirroring. Mesh-mirroring chunks data and creates redundant copies ultimately providing two copies of each chunk. Data will be protected as long as the stored copies of the same chunk are stored on separate physical servers. Mesh-mirroring also allows for all nodes to be involved with the rebuild process during a disk failure.

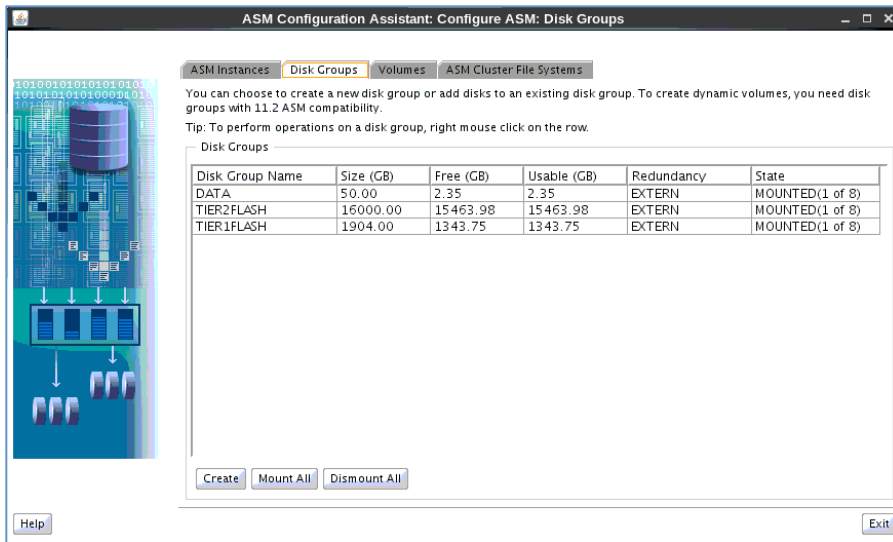


Figure 5: Automatic Storage Management Configuration Assistant. ASM Disks from ScaleIO Volumes.

After configuring ASM, EMC engineers proceeded to load application data in preparation for a series of performance tests.

Performance Analysis

EMC engineers structured their assessment of ScaleIO suitability for Oracle Database based roughly on the following test steps:

- Data Warehousing Model
 - Load a 300GB TPC-H Lineitem table with synthetic data generated by the DBGEN utility.
 - Conduct a series of table scans—with increasing levels of concurrency—using the Parallel Query Option of Oracle Database 11g.
 - Measure storage throughput
- OLTP Model
 - Load 1TB of test data with the Silly Little Oracle Benchmark kit (SLOB¹). SLOB guarantees completely random access to data and is a very simple, scalable and freely available tool kit anyone can use to test storage suitability for Oracle Database. SLOB allows users to test varying read/write ratios as well as selectable transaction logging weight.

¹ For more information on SLOB please refer to <https://community.emc.com/docs/DOC-33564>

² The commonly misunderstood db file sequential read wait event is a single-block, random read. See

Concurrent Table Scan Test Results

In order to assess ScaleIO disk throughput with Oracle Parallel Query Option, EMC engineers conducted a series of concurrent queries. Each query scanned the Lineitem table with full table scans—but varying selectivity predicates and aggregation. Figure 6 shows a screen capture of Enterprise Manager during the ramp-up testing. As Figure 6 shows the testing reached peak performance of slightly more than 20 GB/s before the suite concluded.

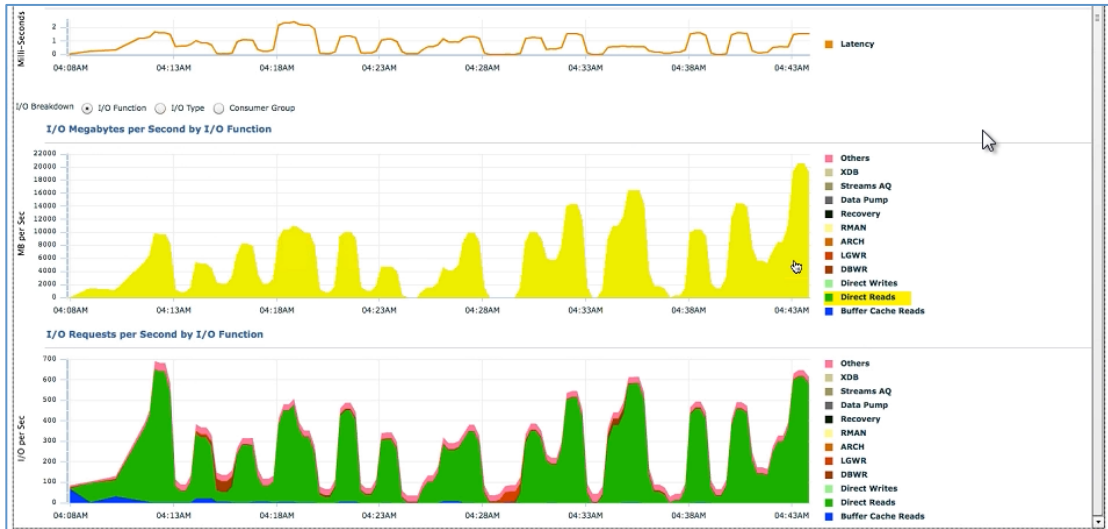


Figure 6: Oracle Enterprise Manager. Monitoring Parallel Query Scan Performance. Ad-hoc Queries.

Figure 7 shows a screen capture of the ScaleIO dashboard during the peak concurrency testing. The dashboard confirms the test configuration peaked at 21.2 GB/s derived from 173,711 physical reads from storage. The dashboard also shows that the I/O transfer sizes as 128KB, which multiplied by 173,711 works out to 21.2 GB/s. EMC engineers set the Oracle Database initialization parameter `db_file_multiblock_read_count` to 16 in order to establish the table scan read size to 128KB.

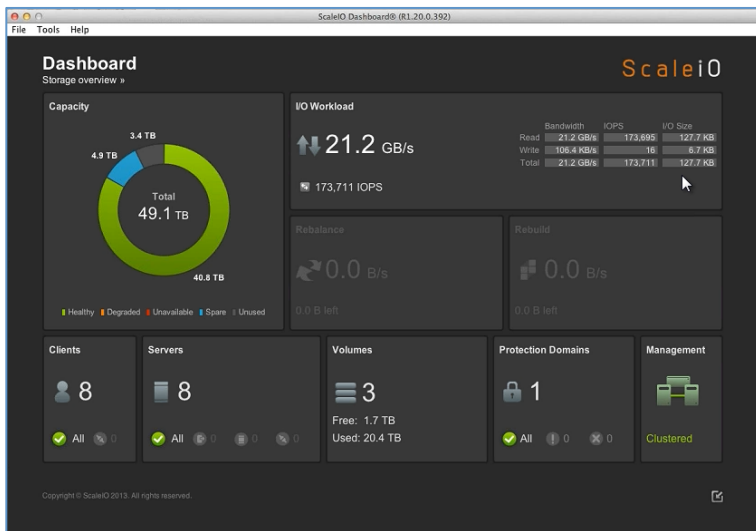


Figure 7: ScaleIO Dashboard. Monitoring Parallel Query Scan Performance of 21GB per second.

OLTP-Style Test Results

Read Only Maximum I/O Rate

After completing the table scan performance testing, EMC engineers switched attention to OLTP-style workload analysis. Task number one was to ascertain the maximum random I/O capability of the test configuration. EMC engineers conducted a series of tests to determine the number of directly connected Oracle database sessions necessary to generate maximum random I/O. This tests suite was entirely read-centric and it was determined that 72 sessions connected to each of the 8 database instances generated peak IOPS. Once these sizing parameters were established, EMC engineers conducted an hour-long test of SLOB with 576 directly connected, zero think-time sessions. Figure 8 shows the Enterprise Manager screen capture depicting the fact that the configuration held the I/O rate above 800,000 IOPS for an entire hour.

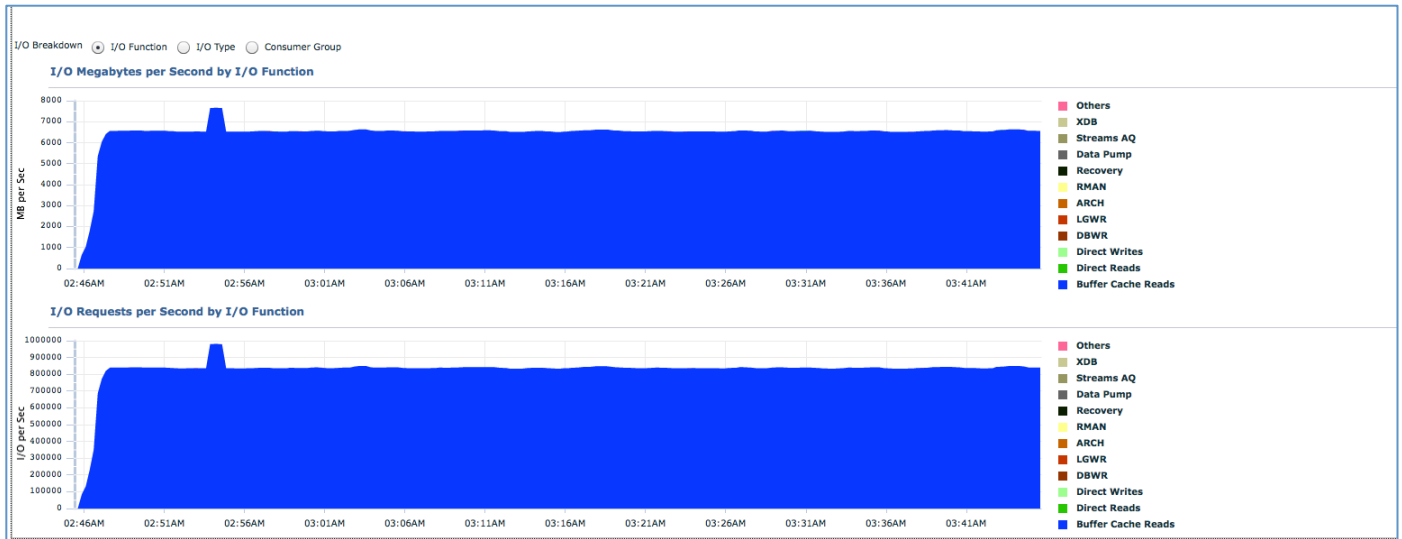


Figure 8: Enterprise Manager. Hour-long Monitoring of Random IOPS Testing.

During the hour long test EMC engineers captured several screen shots of the ScaleIO dashboard. Figure 9 shows one such graphic. As the graphic shows, the I/O sizes were 8KB and the workload consisted entirely random reads. Figure 9 also shows a moment at which the current I/O rate was 893,116 read IOPS.

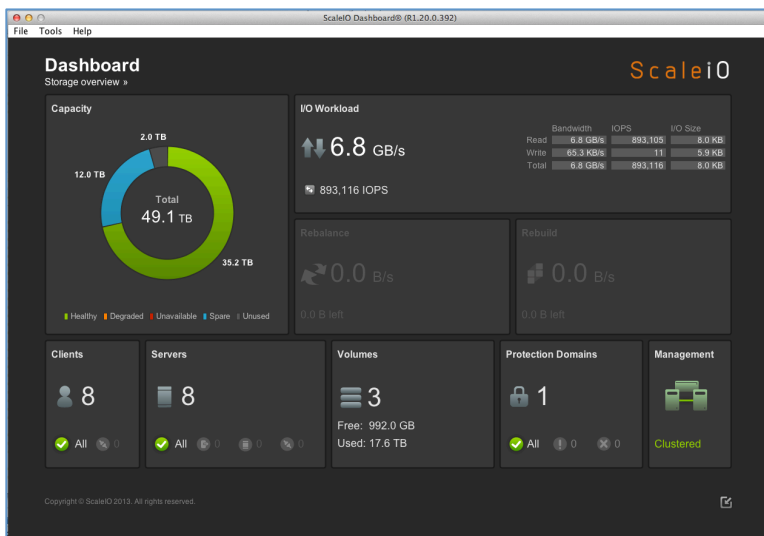


Figure 9: ScaleIO Dashboard. Monitoring Oracle Database Random I/O Rate of 893,116 IOPS.

Read Only I/O Service Times

Today's enterprise IT applications need more than IOPS capacity. Latency is critical too. Figure 10 shows a snippet of the one of the Automatic Workload Repository (AWR) reports generated during the hour-long read-only performance test. Figure 10 also shows that the top wait event was for the *db file sequential read*² I/O event and that the average service time measured by the Oracle Database wait interface was 690us.

#	Wait		Event		Wait Time			Summary Avg Wait Time (ms)				
	Class	Event	Waits	%Timeouts	Total(s)	Avg(ms)	%DB time	Avg	Min	Max	Std Dev	Cnt
*	User I/O	db file sequential read	225,455,791	0.00	155,136.26	0.69	89.72	0.69	0.68	0.69	0.00	8
		DB CPU			16,066.78		9.29					8
	Cluster	gc cr grant 2-way	37,645,271	0.00	6,646.40	0.18	3.84	0.17	0.14	0.25	0.04	8
	Cluster	gc cr grant congested	1,014,230	0.00	431.08	0.43	0.25	0.39	0.25	0.52	0.10	8
	Concurrency	library cache lock	2,432	0.00	206.55	84.93	0.12	45.39	2.65	88.13	60.44	8
	Concurrency	library cache: mutex X	19,741	0.00	89.24	4.52	0.05	2.06	0.92	7.44	2.18	8
	Cluster	gc buffer busy acquire	1,773	0.00	48.15	27.16	0.03	26.12	17.09	38.16	7.21	8
	Configuration	enq: HW - contention	159	0.00	32.01	201.32	0.02	174.21	0.79	308.36	113.29	8
	Cluster	gc remaster	1,320	0.00	31.62	23.95	0.02	24.26	20.59	28.51	2.53	8
	Other	gcs drm freeze in enter server mode	1,401	0.00	25.99	18.55	0.02	20.55	12.55	34.17	8.58	8

Figure 10: Automatic Workload Repository Report Showing Average Random Single-Block I/O Latency of 690 Microseconds.

Mixed I/O OLTP-Style Testing

After completing the assessment of the ScaleIO test configuration with the read-only workload, EMC engineers changed the SLOB parameters governing the SQL SELECT-to-UPDATE ratio to 25% UPDATE and 75% SELECT. It's necessary to point out that a SQL blend of 25% UPDATE alongside 75% SELECT does not map to a physical I/O blend of 75/25. Database products like Oracle defer writes to data files until a checkpoint is forced. With all this in mind, please direct your attention to Figure 11. Figure 11 is a screen capture of the ScaleIO dashboard showing an overall I/O rate of 565,522 IOPS where—at that moment—the ratio of reads to writes was roughly 85:15 (487,224 read IOPS and 78,298 write IOPS).

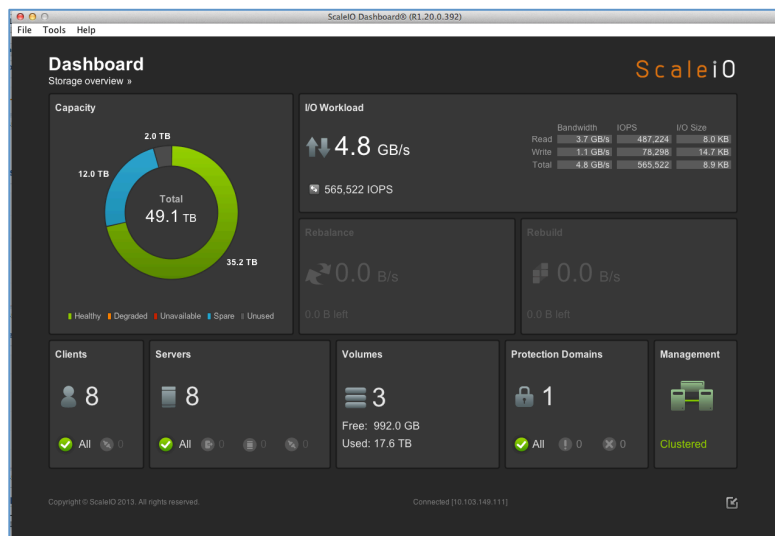


Figure 11: ScaleIO Dashboard. Monitoring 75:25 OLTP-Style Workload Achieving 565,522 IOPS.

To complement the hour-long testing with the 100% SQL SELECT workload, EMC engineers executed another hour-long test with the 75:25 (SELECT:UPDATE) workload. Figure 12 shows a screenshot of Enterprise Manager after the hour-long test completed.

The graphic shows that the workload held steady—free of pathologies. More importantly, however, is how Enterprise Manager shows that the random reads enjoyed service times in the 600-microsecond range

² The commonly misunderstood db file sequential read wait event is a single-block, random read. See http://docs.oracle.com/cd/E16655_01/server.121/e15857/pfgrf_instance_tune.htm#TGDBA94485



Figure 12: Enterprise Manager. Monitoring Hour-Long Sustained 75:25 OLTP-Style Workload.

Oracle Database 12c Performance Characterization

After finishing the Oracle Database 11g testing, the EMC engineers created a new database with Oracle Database 12c using four of the hosts in the test configuration. The database was loaded with a 300GB scale TPC-H Lineitem table populated with the DBGEN utility.

After executing a series of queries the EMC engineers took a screenshot of the Oracle Database 12c Enterprise Manager Express performance hub tab. As Figure 13 shows, EM Express proved that the smaller Oracle Database 12c set of instances (4 in total) were able to scan storage at 11 GB/s in parallel.

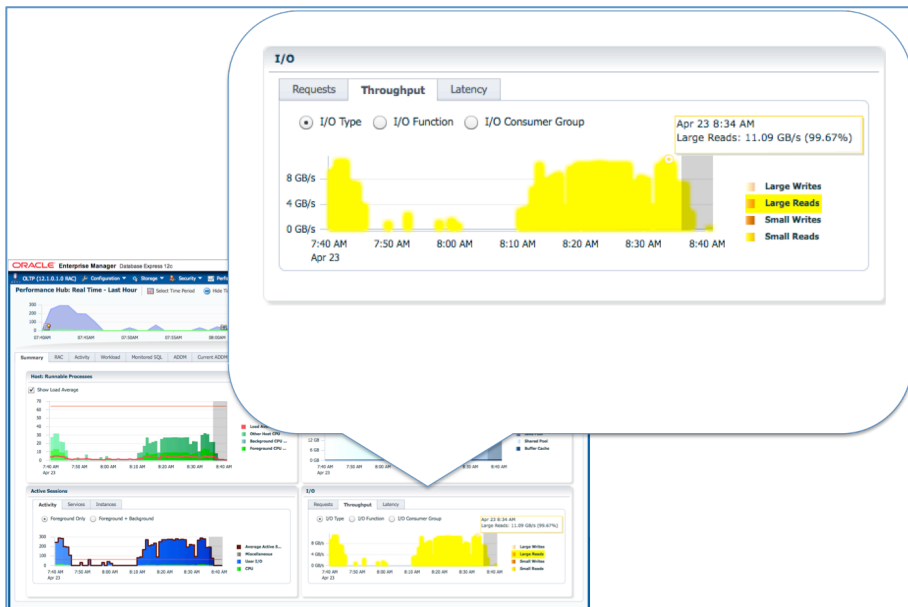


Figure 13: Oracle Database 12c Enterprise Manager Express. Monitoring 4-Node Parallel Query Performance.

For further elaboration on the Oracle Database 12c table scan testing please refer to Figure 14 which shows a highlighted screenshot of an interactive SQL*Plus session where EMC engineers proved the following:

- The Oracle Database release was 12.1.0.1.0
- A query against the gv\$instance view shows there were four instances of this database online
- The table scan took 25.01 seconds
- Calculating the difference between the before and after global cumulative physical reads proved 293 GB had been scanned at a rate of 11.7 GB/s.

```

$
$ sqlplus tpch/tpch @test_lineitem_scan.sql

SQL*Plus: Release 12.1.0.1.0 Production on Wed Apr 23 09:57:53 2014

Copyright (c) 1982, 2013, Oracle. All rights reserved.

Last Successful login time: Wed Apr 23 2014 09:26:13 -07:00

Connected to:
Oracle Database 12c Enterprise Edition Release 12.1.0.1.0 - 64bit Production
With the Partitioning, Real Application Clusters, Automatic Storage Management, OLAP,
Advanced Analytics and Real Application Testing options

SQL> SELECT COUNT(*) NUMBER_OF_RAC_INSTANCES_ONLINE FROM gv$instance;

NUMBER_OF_RAC_INSTANCES_ONLINE
-----
                                4

SQL> SELECT sum(VALUE) / 1024/1024/1024 PREAD_GB FROM gv$sysstat
 2 WHERE STATISTIC# = (SELECT STATISTIC# FROM v$statname WHERE NAME = 'physical read total bytes');

PREAD_GB
-----
 63380.0

SQL> SET TIMING ON
SQL> SELECT avg(L_DISCOUNT) FROM h_lineitem WHERE L_TAX < 7 and L_QUANTITY > 9 ;

AVG(L_DISCOUNT)
-----
      .049993254

Elapsed: 00:00:25.01
SQL> SET TIMING OFF
SQL> SELECT sum(VALUE) / 1024/1024/1024 PREAD_GB FROM gv$sysstat
 2 WHERE STATISTIC# = (SELECT STATISTIC# FROM v$statname WHERE NAME = 'physical read total bytes');

PREAD_GB
-----
 63673.4

SQL> select (63673.4 - 63380.0) / 25.01 GB_SCAN_RATE_PER_SECOND from dual;

GB_SCAN_RATE_PER_SECOND
-----
          11.7313075

SQL> █

```

Figure 14: Screen Capture of Interactive Session. EMC Engineers Measure 11.7GB/s Scan Rate with 4 RAC Nodes.

SUMMARY

EMC ScaleIO architecture and features enable administrators to build exceedingly powerful converged infrastructure platforms that are well suited to Oracle Database. The data protection, fault-resilience and elasticity features of ScaleIO complement the scale-out nature of Oracle Real Application Clusters. Oracle Database editions ranging from Standard Edition to Enterprise Edition with Real Application Clusters benefit from the highly parallel shared-disk infrastructure which assures high bandwidth, low-latency IOPS and high availability. IT architects now have the technology solution that enables them to transform low cost, powerful industry standard servers into highly-available, highly-performing converged infrastructure optimized for Oracle Database.

CONTACT US

To learn more about how EMC products, services, and solutions can help solve your business and IT challenges, [contact](#) your local representative or authorized reseller—or visit us at www.EMC.com.

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.

EMC2, EMC, the EMC logo, and the RSA logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. 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. Published in the USA 8/13 White Paper H12249