Lab Validation Report

EMC ScaleIO

Transforming Commodity Hardware into Simple, Scalable, High-performance, Shared Storage

By Aviv Kaufman, Lab Analyst and Brian Garrett, Vice President ESG Lab

April 2014

Contents

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by EMC.

Introduction

This ESG Lab Validation documents the hands-on testing of ScaleIO technology with a goal of demonstrating how ScaleIO can turn commodity hardware into a powerful and scalable converged storage and compute cluster. The testing focused on validating the performance scalability, elasticity, and reliability of a converged ScaleIO cluster in an Oracle Real Application Cluster (RAC) environment.

Background

No matter how big or small your IT organization is, the challenge is the same: Do more with less. With budgets flattening and end-users continuing to demand services that require greater performance and agility from the supporting IT infrastructure, companies must find innovative ways to create scalable solutions. With capital and operational expenses constantly under scrutiny, IT admins must be conscious of the impact that their technology choices have on the bottom-line and be ready to defend their choices. ESG research echoes this sentiment, with total cost of ownership (TCO) as respondents' most-cited important criterion for selecting a storage vendor or solution.[1] Also of great importance to these organizations were ease of management and ease of implementation. The top five most-cited responses are shown in Figure 1.

Figure 1. Top Five Most Important Criteria in Selecting a Storage Vendor or Solution

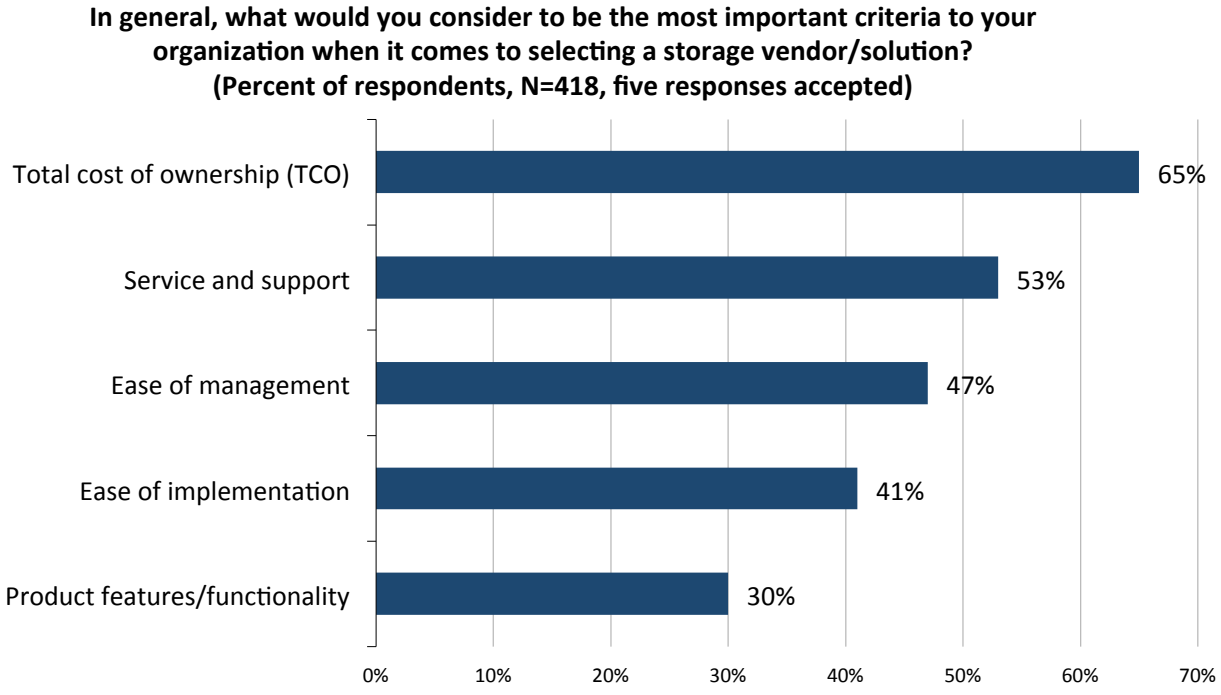

Source: Enterprise Strategy Group, 2014.

Traditional IT infrastructures consisting of silos of dedicated servers, networking, and storage are extremely complex and expensive to purchase and maintain. To limit capital and operational expenses, many organizations have turned to cloud-based solutions such as infrastructure-as-a-service (IaaS). ESG research indicates that 65% of enterprise organizations currently use or plan to use IaaS, with another 20% indicating that they are perhaps interested in using it later.[2] Many organizations leverage the cost effectiveness of these solutions as a temporary solution as they struggle to design an in-house infrastructure that can deliver the same level of cost effectiveness, flexibility, and scalability to their end-users while also satisfying the demands of today's high-performance applications.

[1] Source: ESG Research Report, *2012 Storage Market Survey*, November 2012.
[2] Source: ESG Research Report, *2014 Public Cloud Computing Trends*, March 2014.

The Solution: EMC ScaleIO

EMC ScaleIO is a flexible software-only solution that leverages host-based internal storage to create a scalable virtual SAN with performance that is comparable with or better than an external SAN, at a fraction of the cost and complexity.

ScaleIO works by installing lightweight software components on application hosts. Application hosts contribute internal disks and any other direct attached storage (DAS) resources to the ScaleIO cluster by installing the ScaleIO Server software (SDS) and can be presented volumes from the ScaleIO cluster by leveraging the ScaleIO Client (SDC). These components can run alongside other applications on any server (physical, virtual, or cloud) and can leverage any type of storage media (disk drives, flash drives, PCIe flash cards, or cloud storage).

ScaleIO can be deployed as storage only or can greatly simplify the infrastructure by converging the storage, compute, and networking resources into a single scalable building block that is easy to manage. Capacity and performance of all available resources are aggregated and made available to every participating ScaleIO server and application. Storage tiers can be created with media types and drive types matching the ideal performance or capacity characteristics to best suit the application needs.

Storage and compute resources can be added to or removed from the ScaleIO cluster as needed, with no downtime and minimal impact to application performance. The self-healing, auto-balancing capability of the ScaleIO cluster ensures that data is automatically rebuilt and rebalanced across resources when components are added, removed, or failed. Because every server and local storage device in the cluster is used in parallel to process I/O operations and protect data, system performance scales linearly as additional servers and storage devices are added to the configuration.

Figure 2. ScaleIO: Turning Commodity Hardware into a Converged, Elastic, Scalable Infrastructure

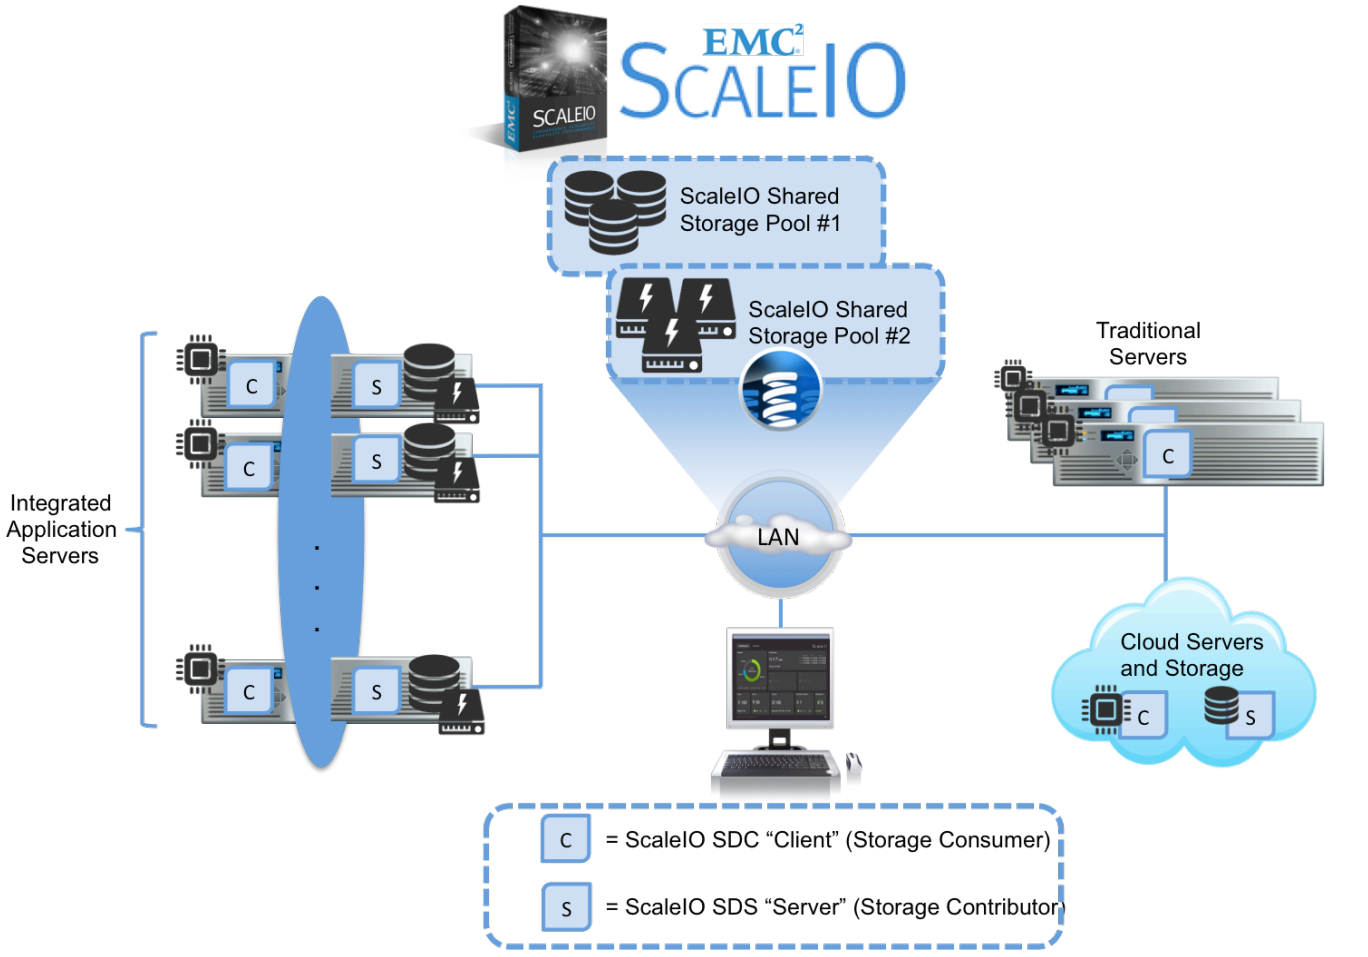

ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of ScaleIO at EMC's facilities, in Hopkinton, Massachusetts. Testing was designed to demonstrate how ScaleIO can be used to converge commodity hardware into a scalable, elastic, high-performance server SAN for use in an Oracle RAC environment using industry-standard tools and methodologies. ESG Lab also tested the reliability and self-healing capability of the ScaleIO cluster.

Getting Started: Convergence

To get started, ESG Lab examined the hardware components and reviewed the steps that were necessary to turn the 17 rack units (17U) of commodity hardware into a high-performance converged Oracle RAC database deployment. The hardware included eight identically configured Cisco UCS 240 M3 Servers and a Mellanox 40Gb QDR Infiniband switch. Each server contained a dual 40Gb HCA for Ethernet connectivity, a 10krpm SAS drive, a PCIe SLC Flash card, and six 2.5" cMLC flash drives.

ESG Lab Testing

To begin, Oracle Linux (OL) v.6.4 was installed on the 10krpm SAS drive of each server and the ScaleIO files were copied to one of the servers. A simple menu driven configuration tool was used to easily generate a script to configure a ScaleIO cluster for first time use by setting node IP addresses, creating a protection domain, and adding the internal disk devices for all eight nodes. Two tiered storage pools were defined for performance testing: The first consisted entirely of the PCIe SLC Flash cards, and the second was made up of the cMLC flash drives. Executing the automated script installed ScaleIO software on all eight nodes and created the ScaleIO cluster. The entire process took only a few minutes to complete and required no manual intervention.

The process installed and configured the three required ScaleIO software components that are used to create the virtual SAN and converge the compute and storage infrastructure on the servers. The ScaleIO Data Server (SDS) was installed on all eight servers and was used to present the PCIe SLC Flash card and six cMLC SSDs to the ScaleIO storage pool. The SDS can contribute any type of local storage to the pool including disks, disk partitions, or even files. The ScaleIO Data Client (SDC) was also installed on all eight servers and was used to provide access to the presented ScaleIO volumes by communicating with all other nodes. Finally, the Metadata Manager (MDM) was installed on only three of the servers and functioned as a quorum that handles cluster-wide mapping information and decision making for all system functions. The functionality of the SDC and SDS software components is shown in Figure 3.

Figure 3. ScaleIO SDC Client and SDS Server Functionality

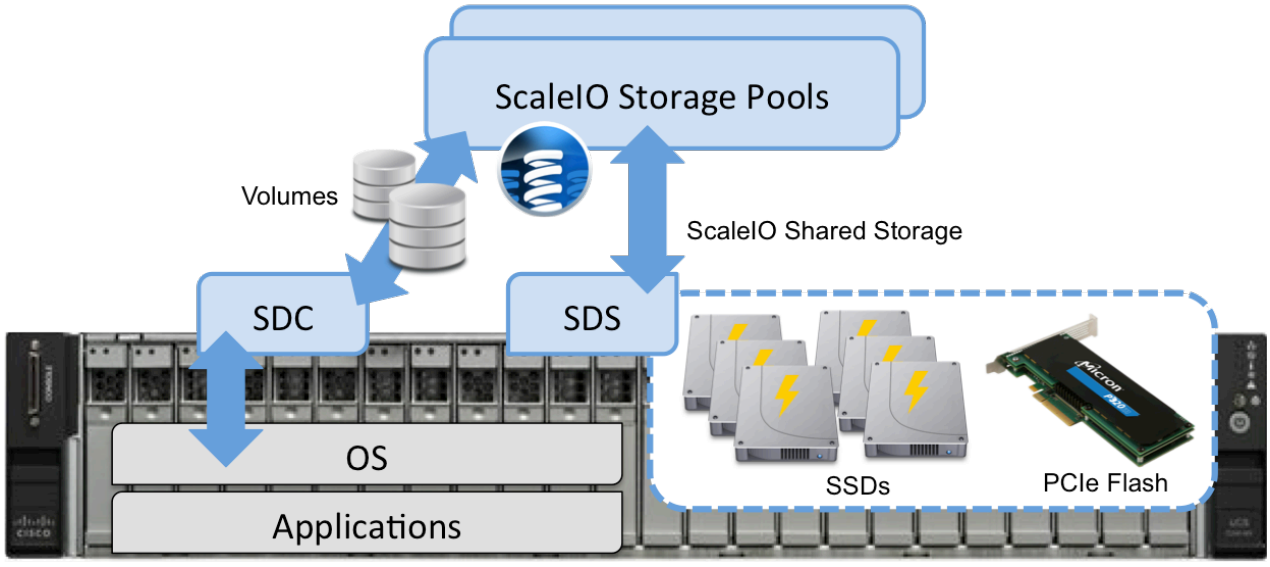

Once the configuration script was completed, a web browser was used to monitor the configuration by pointing at the virtual IP address of the cluster. The monitoring interface was extremely simple to use and well designed. Using the monitor interface, it was possible to see the entire converged infrastructure in a single, modern, intuitive dashboard view. The dashboard included real-time cluster capacity and performance statistics, system status and alerts, as well as client, server, and volume information. ESG Lab was able to see all eight clients (SDC) and eight servers (SDS) that made up the cluster and all of the capacity available to the storage pools. Figure 4 shows the ScaleIO Graphical Monitoring Interface.

Figure 4. ScaleIO Graphical Monitoring Interface

Next, ESG Lab created three volumes and mapped them to the SDC on all eight nodes by issuing several very simple command line interface (CLI) commands. A single 5TB volume was created on the "tier1" SLC PCIe storage pool and a single 40TB volume was created on the "tier2" cMLC SSD storage pool to store the Oracle RAC database tables. Oracle RAC was then installed and configured on all eight nodes and a large database was configured and populated on each of the ScaleIO volumes for use in testing. Figure 5 shows the resulting test bed: a high-performance, converged infrastructure, eight node Oracle RAC environment made up completely of industry-standard hardware components and ScaleIO software.

Figure 5. Converged ScaleIO-powered Oracle RAC Test Bed

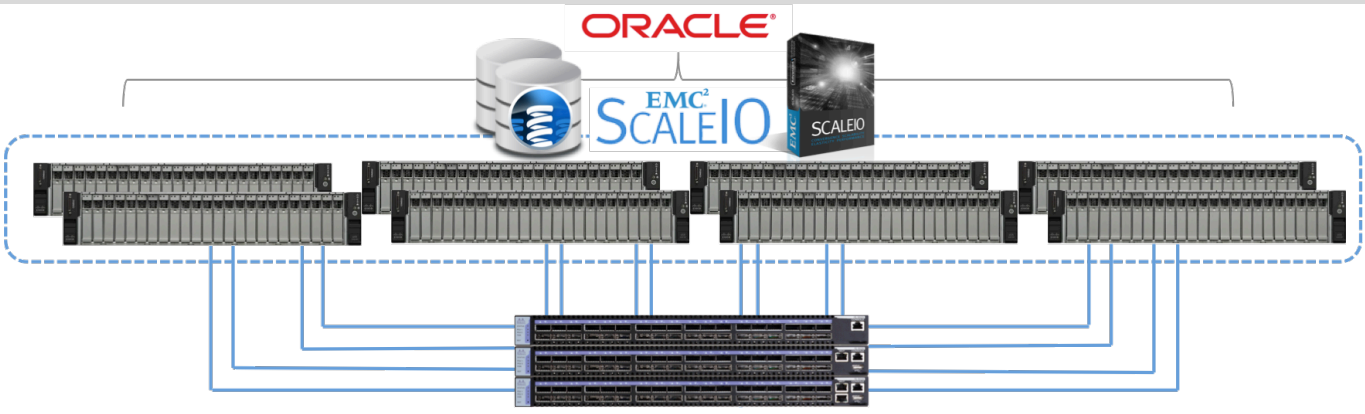

Hardware (Each of 8 Nodes):
Cisco UCS 2U C240 M3 w/ 24 x SFF Bays
2 x Intel 2.90 GHz E5-2690 (8 Cores)
128GB DDR3-1600 Memory
1 x 300GB 6Gb SAS 10Krpm SFF HDD (OS)
1 x 700GB Micron P320 SLC PCIe Flash Card
6 x 960GB Micron M500 cMLC SSD
1 x Dual Port Mellanox 40Gb HCA

Software (Each of 8 Nodes):
Oracle Linux 6.4 OS
Oracle 12c RAC w/ ASM External
ScaleIO 1.2 SDS
ScaleIO 1.2 SDC

ScaleIO Storage:
Pool #1: 5.1TB (SLC PCIe)
Pool #2: 41.9TB (cMLC SSD)

It should be noted that while ESG chose to test the ScaleIO configuration in a single layer converged configuration (database servers and storage sharing the same hardware), ScaleIO can also be configured as a two-layer SAN solution, connected to an existing Oracle deployment by simply installing the SDC on the Oracle servers. ScaleIO also gives the end-user the flexibility to make this configuration choice non-disruptively at a later date, starting with a converged solution, and non-disruptively reconfiguring to standalone servers with a dedicated ScaleIO storage solution later in the hardware lifecycle.

Why This Matters

Organizations are under constant scrutiny to do more with less. While technological demands on an IT data center are ever increasing, budgets are often flat. Organizations must look for innovative ways to solve business problems by limiting both capital purchases and operational expenditures. Many IT solutions are not only expensive to purchase, but are also costly to operate and maintain. Traditional solutions require multiple interfaces and dedicated administrators for servers, SAN, networking, and storage. ESG research indicates that 62% of IT organizations' 2014 budgets will be spent managing existing infrastructure.[3]

Using ScaleIO, ESG Lab was able to transform eight industry-standard servers and a single switch into a powerful, converged architecture capable of supporting a high-performance, highly available, Oracle RAC environment. This simple, elastic, ScaleIO cluster can be built on demand, leveraging cost-effective commodity hardware or repurposing underutilized servers with internal storage. A ScaleIO cluster is simple to create and can be managed easily by a single administrator, saving valuable time and minimizing operational expenditures.

[3] Source: ESG Research Report, *2014 IT Spending Intentions Survey*, February 2014.

Performance and Scalability

The driving factor behind the design of ScaleIO was to create a software-only product that leverages commodity hardware to create shared storage that can scale from a minimum of three to thousands of nodes. ScaleIO was designed to run efficiently, using minimal computing resources, while distributing volumes access and metadata across all nodes and never relying on a "master" node to route access. This design enables ScaleIO to scale in both performance and capacity. ESG Lab validated the performance and scalability of the system by testing an eight-node system in an Oracle RAC environment and auditing the results of internal scalability tests.

ESG Lab Eight-node Oracle RAC Testing

ESG Lab first looked at the performance that ScaleIO can deliver in a simulated real-world populated database in an Oracle RAC environment. Testing focused on measuring the throughput and latency of Oracle SQL-driven I/O and the bandwidth achieved with parallel table scans. ESG Lab utilized the widely adopted and publicly available Silly Little Oracle Benchmark kit (SLOB) to efficiently generate realistic system-wide, random, single block, application-independent SQL queries. The SLOB benchmark exercised the ScaleIO storage by stressing the physical I/O layer of Oracle through SGA-buffered random I/O without being limited to a specific load-generating application.

Eight Cisco UCS C240 servers were configured, each running Oracle Linux (OL 6.4) with Oracle Database 12c Grid Infrastructure. The clusterware disk objects were stored in an ASM disk group consisting of a ScaleIO volume created from a pool of the single 300GB SAS HDD in each server. Each server also contained a dual port 40Gb NIC for network connectivity, a single 700GB PCIe SLC Flash card, and six 2.5" 960GB cMLC SATA flash drives. Both the ScaleIO 1.2 SDS (server) and ScaleIO 1.2 SDC (client) drivers were installed on each server. The PCIe flash storage and the cMLC flash drives were used to create two storage pools. All servers were then presented a ScaleIO volume from each storage pool for the Oracle tables and another volume for the Oracle logs.

ESG Lab first tested the eight-node performance of the Oracle RAC simulated database when performing SQL reads. The SLOB benchmark was configured to generate a total of 400 sessions that were generating SQL read queries to the database. The ScaleIO graphical performance monitor and Oracle Enterprise Manager (OEM) were used to monitor the performance of the ScaleIO cluster during the run, and the number was validated upon completion by auditing the Oracle Automated Workload Repository (AWR) report. ESG Lab verified that the eight-node cluster was able to perform over **830,000 random 8KB Oracle database I/Os per second**. It should be noted that this number does not represent the maximum achievable performance with ScaleIO, and ESG Lab expects the performance to increase as more nodes are added to the cluster. Figure 6 shows the results of the test as seen by the ScaleIO GUI, Oracle AWR report, and OEM performance report.

Figure 6. Eight-node Oracle RAC Environment Performance: 100% 8KB Read SLOB SQL Query Generated IOPS

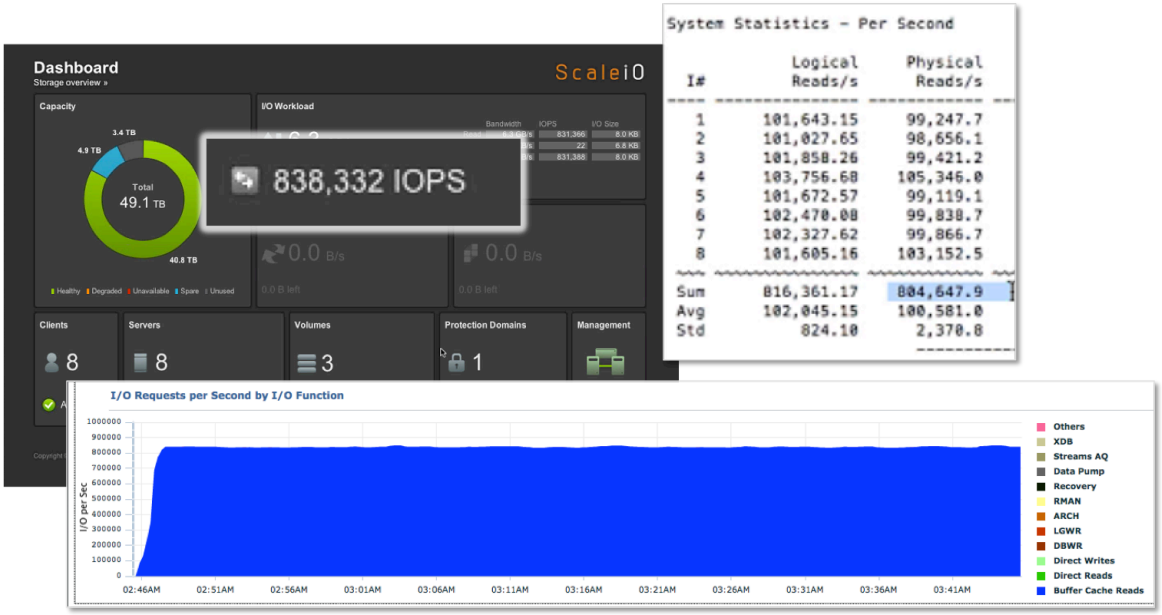

Next, ESG Lab tested the performance of the eight-node Oracle RAC configuration using a workload that is more representative of a typical production OLTP environment performing a mix of database queries and database updates. The SLOB configuration was reconfigured to produce a mixed workload consisting of 75% SELECT (reads) and 25% UPDATE (read/modify/write) that resulted in a disk activity of roughly 84% random reads and 16% random writes. The workload was run over a period of one hour while OEM and the ScaleIO monitoring interface were used to monitor and ensure consistent performance through the duration of the test. The eight-node environment produced a total of more than **565,000 IOPS**, as shown in Figure 7.

Figure 7. Eight-node Oracle RAC Environment Performance: 75% Query/25% Update SLOB SQL Generated IOPS

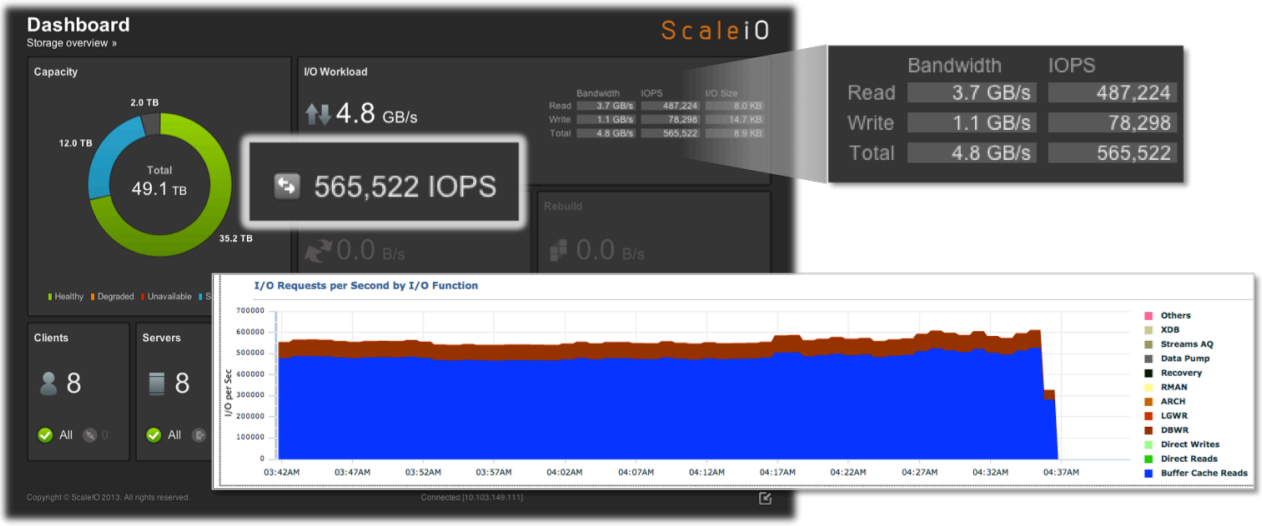

Finally, ESG Lab validated that the ScaleIO system can deliver the high throughput demanded by Oracle admins when performing database table scans. For this test, a script was created to generate three concurrent parallel SQL queries against the tables in the Oracle RAC database. OEM and the ScaleIO monitoring tool were used to monitor the performance of the table scans. Because the performance of table scans can vary through the run, the test was repeated several times while incrementally increasing the number of concurrent scans. The eight-node environment consistently delivered Oracle Parallel Query table-scan rates of over **21 GB/sec** as seen in Figure 8.

Figure 8. Eight-node Oracle RAC Environment Performance: OLAP-like Parallel Table Scans

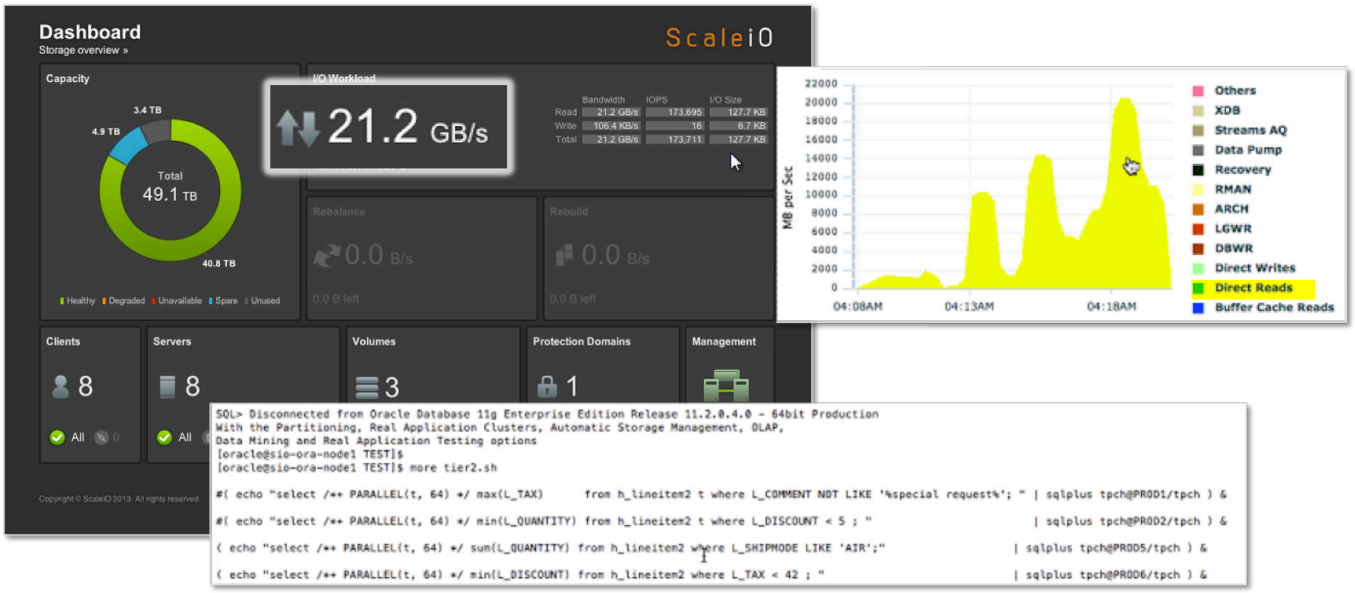

ESG Lab Audited FIO Results

To demonstrate scalability, ESG Lab audited the results of internal tests performed at the ScaleIO QA engineering facilities in Israel. The tests were conducted on a 53-node ScaleIO cluster made up of Cisco UCS C240M3 servers, each with a single 700GB PCIe SLC flash card. Using the industry-standard FIO load generator tool to measure random performance, the configuration was able to achieve a truly impressive **8.5 million read IO/sec**, and **3.7 million write IO/sec**. Equally impressive was the **114 GB/sec** achieved by the cluster when performing 64KB random reads. But most remarkable of all to ESG Lab was the fact that these numbers still do not represent the maximum performance of a ScaleIO cluster. Figure 9 shows a summary of the audited test results.

Figure 9. 53-node ScaleIO QA Environment Audited Scaling Test Results

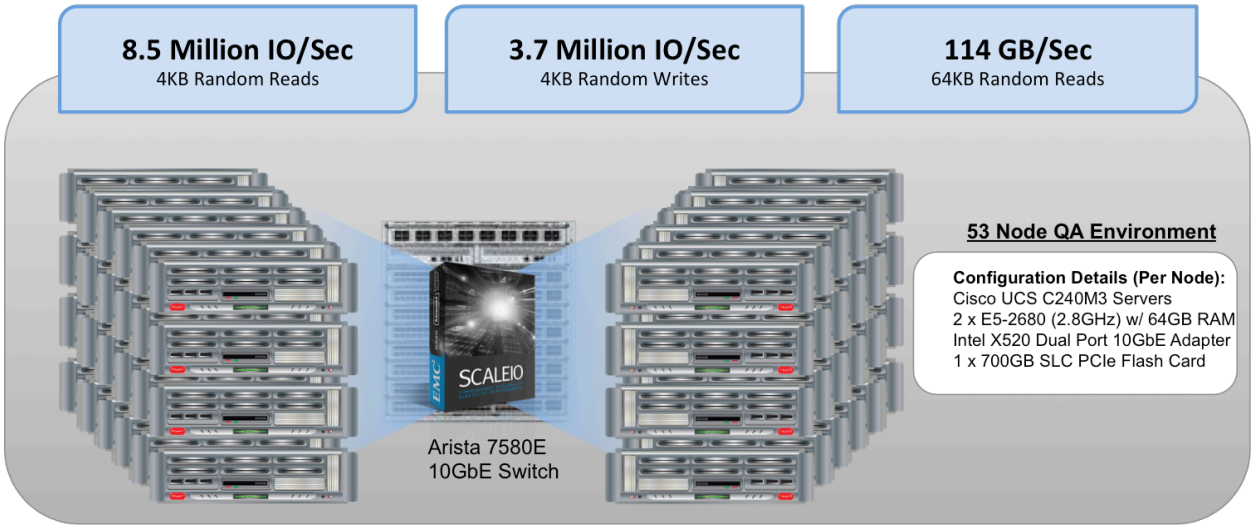

What the Numbers Mean

- The eight-node Oracle RAC ScaleIO cluster was able to perform over 800K query SLOB IOPS and 565K mixed query and database update SLOB IOPS, proving that it is capable of handling some of the most demanding transactional database workloads.
- The ScaleIO powered eight-node Oracle RAC cluster was also able to perform over 21 GB/sec while performing parallel table scans of the database, demonstrating bandwidth performance that is on par with some of the more expensive storage solutions on the market today.
- In addition to high IOPS, ESG noted that the response times seen in the tests were low as well: under 1.2 milliseconds for the mixed query and update testing, and under 0.7 milliseconds for the read-only tests.
- A 53-node cluster performed over 8.5 million IOPS and 114 GB/sec, demonstrating that a ScaleIO cluster can be scaled to meet the demands of nearly any application.

Why This Matters

Database performance is critical to business operations. A higher performing database can produce quicker sales and faster access to processed information, ultimately resulting in higher production and increased revenue for the organization. But deploying a high-performance database is expensive. The cost of acquisition to maintain silos of servers, SAN, and storage, as well as the dedicated expertise to manage these complex resources can add up very fast, and unpredictable business demands make it difficult to plan for future growth.

Using ScaleIO and cost-effective commodity hardware, ESG Lab was able to create a high-performance Oracle RAC database that was able to perform over 830K random SLOB SQL database reads per second, and scan three tables at a combined rate of over 21 GB/sec. This configuration was simple to manage, scaled linearly, and achieved equal or better performance at a hardware cost that ESG estimated to be about half of the hardware cost of an equivalently performing Oracle RAC converged solution on the market today**.**

Elasticity and Reliability

EMC ScaleIO is designed to be elastic: Workloads can share the same set of resources without impacting one another and hardware components can be added or removed as needed, non-disruptively, with little or no impact to existing applications. The self-healing capability of the cluster ensures that data is balanced and protected at all times. ESG Lab validated the elasticity and reliability of the ScaleIO cluster by simultaneously running Oracle OLTP queries and table scans, and by adding, removing, and failing nodes non-disruptively while performing database queries.

ESG Lab Testing

First, ESG Lab tested the ability of the ScaleIO cluster to add and remove nodes to the cluster to meet the ever-changing demands of the business. A case was simulated in which an organization had decided to repurpose two of the eight nodes in the Oracle RAC ScaleIO cluster for a period of time due to reduced demands on the database and higher priority needs for the hardware in another ScaleIO cluster. The previously tested mixed query and update SLOB workload was used to measure the performance of the configuration before the removal of the nodes, after the removal of the nodes, and after replacing the nodes back into the ScaleIO cluster.

After measuring the performance of the eight-node Oracle RAC configuration, two nodes were removed from the ScaleIO cluster. The array automatically rebuilt and rebalanced the data contained on the nodes being removed to the remaining six nodes. The measured performance dropped from roughly 565,000 SLOB IOPS to roughly 437,000 SLOB IOPS, which is a near perfect linear decrease of 23%. After adding the two nodes back into the configuration and waiting for the data to be automatically redistributed, the performance was measured once again and had returned back to its original eight-node level of roughly 565,000 SLOB IOPS.

Figure 10. Linear Performance Scaling when Shrinking and Growing the ScaleIO Cluster

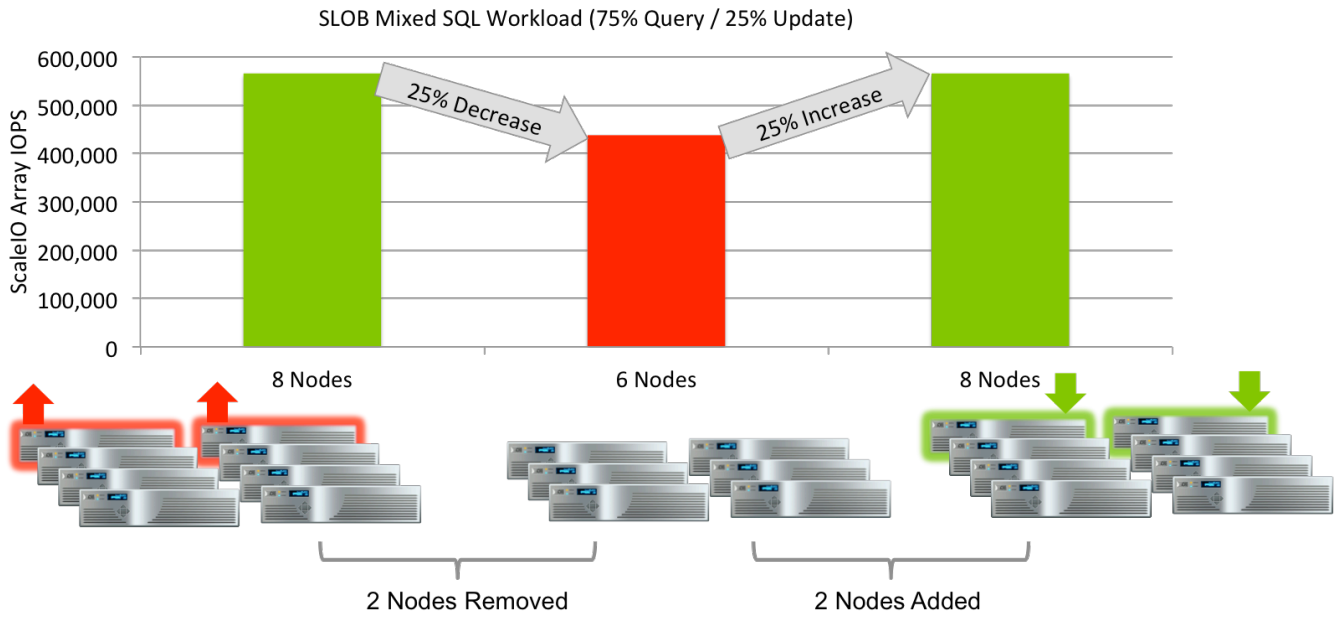

Next, ESG Lab tested the flexibility of ScaleIO to handle mixed workload environments by performing a high number of random small block mixed workload OLTP-like database queries to a simulated development database, while simultaneously performing throughput-intensive Oracle Parallel Query table scans to a simulated development database. The SLOB benchmark was used to generate a mix of database queries and database updates to the "production" database on the "tier1" storage pool that resulted in 84% reads and 16% writes on the ScaleIO cluster. While the first database was performing over 350K IOPS, three Oracle table scans were generated in parallel against the "development" database on the "tier 2" storage pool.

ESG Lab noted that there was no measureable drop in performance for either application as the ScaleIO cluster was able to simultaneously deliver 359K mixed workload OLTP IOPS while also performing table scans to a second database at a rate of 16.2 GB/sec. In traditional systems, mixing workloads with these characteristics is difficult and requires manual intervention and tradeoffs to fine-tune the system to best meet the demands of the applications. Because ScaleIO was built for scale and performs operations in parallel across all servers, it was not necessary to fine-tune the system to handle multiple workloads. Figure 10 shows the results of the multiple workload testing.

Figure 11. Eight-node Oracle RAC Environment: Oracle Parallel Queries with Simultaneous OLTP Mixed Workload

Next, ESG Lab validated the self-healing capability of the ScaleIO cluster. As is the case with most scale-out architectures, the impact of losing a node will be greater with fewer nodes in the system because each node in the system represents a percentage of the total system performance. ESG Lab performed the ScaleIO node failure testing in a "worst-case" scenario using only three nodes (the minimum configuration required to maintain a tie-breaking quorum). The testing was completed on an Oracle database lab development cluster made up of Supermicro servers with a single internal PCIe Flash card running Red Hat Enterprise Linux 6.3.

The SLOB mixed query and update workload (84% reads) was run against a single volume generating close to 120K IOPS on the ScaleIO cluster. After running for several minutes, one of the three servers was issued a hard reboot using a Linux kernel command, immediately removing it from the cluster and forcing an expected drop in performance. The cluster immediately began rebuilding the data to ensure that all data was protected on the remaining two nodes, during which time the rebooted node powered back on and reentered the configuration. After only five minutes, all of the data had been rebuilt and the cluster began to rebalance the data across the three nodes. In just four minutes, the data had been rebalanced across all three nodes and the database was back to the original level of performance.

The synthetic worst-case scenario test showed an expected performance impact as one-third of both the disk resources and load generation was removed. It bears mention that the ESG Lab workloads were meant to push the components in the configuration close to their limits. The impact of this test will vary depending on a number of factors, including number and type of nodes and drives as well as the intensity of the workload. In practice, for most real-world ScaleIO deployments, there may be little to no noticeable impact to the performance of the database at all when rebalancing and rebuilding data in the cluster. The results of the automated self-healing test are shown in Figure 12.

Figure 12. ESG Lab Validation of ScaleIO's Automated Self-healing Capability

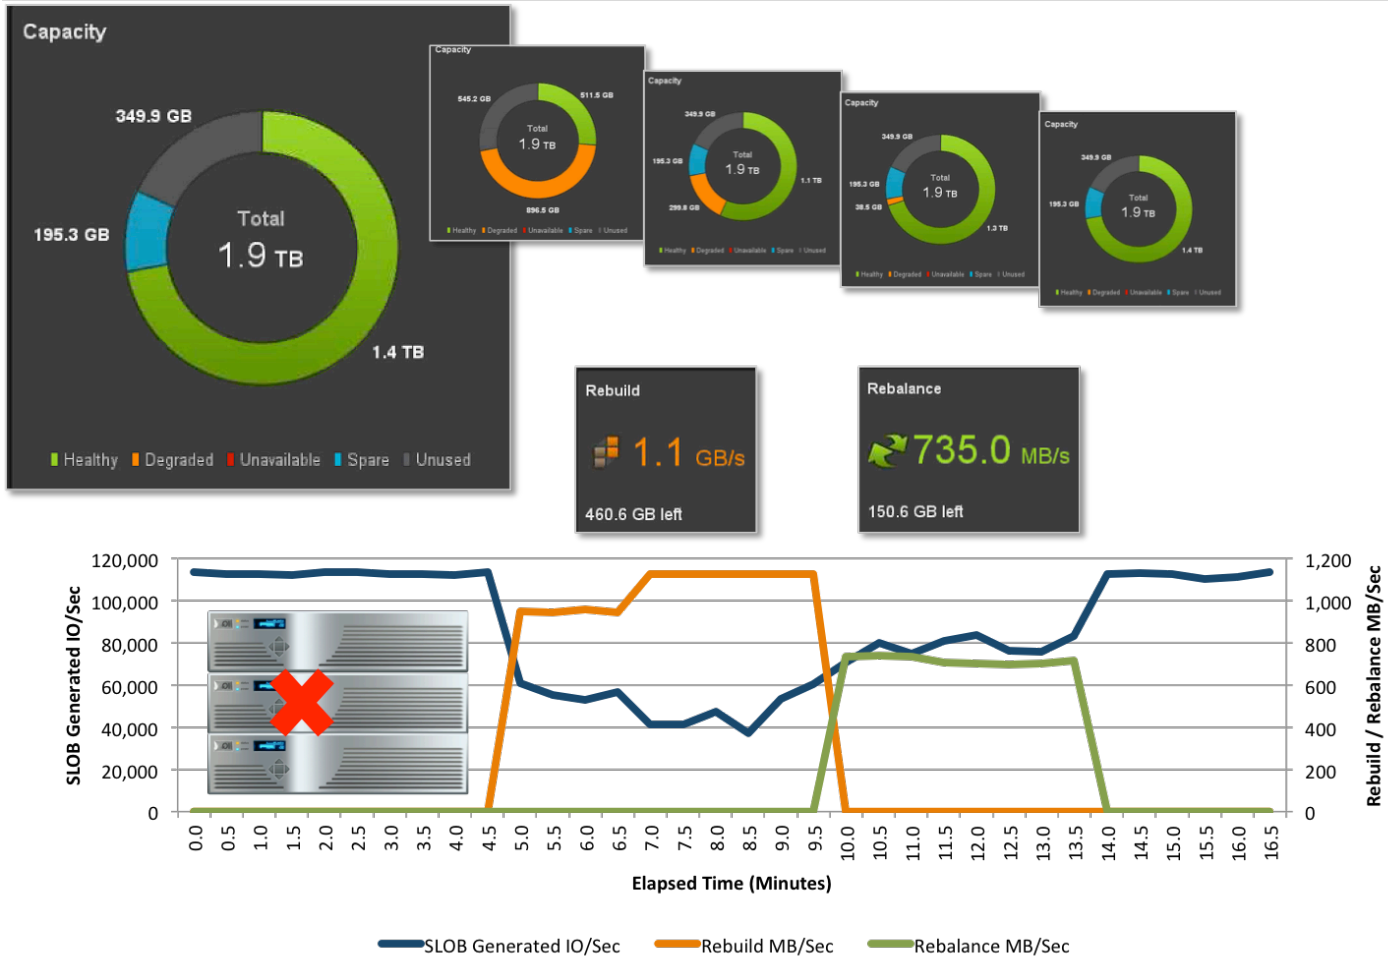

Why This Matters

To meet the needs of an unpredictable mobile workforce expecting on-demand access to information from a variety of sources, many organizations have leveraged the elasticity offered by cloud-based infrastructure-as-service (IaaS) solutions. ESG research shows that roughly two out of every three organizations currently use, or have plans to use, an IaaS cloud service.[4] While many organizations are happy using cloud services for non-mission-critical applications suited to the public cloud, IT departments are hard pressed to offer the same level of elasticity, reliability, and cost effectiveness for internal private cloud-based applications.

ESG Lab was able to grow and shrink a ScaleIO cluster non-disruptively in a matter of minutes, offering predictable linear performance to an Oracle RAC cluster and simultaneously satisfying the dissimilar performance requirements of databases performing over 360K OLTP IOPS and 16.3 GB/sec table scans. The self-balancing, self-healing functionality allows the ScaleIO cluster to grow or shrink non-disruptively as needed to meet the current demands of the business. ScaleIO can help transform an IT organization from a cost center into a service center, saving an organization hundreds of thousands of dollars by enabling IT to easily and cost-effectively deploy simple-to-manage, elastic, high-performance IaaS for internal operations.

[4] Source: ESG Research Report, *2014 Public Cloud Computing Trends*, March 2014.

ESG Lab Validation Highlights

- ☑ ESG Lab was able to use ScaleIO to easily create a high-performance converged infrastructure eight-node Oracle RAC database using off-the-shelf commodity hardware while eliminating the need for an expensive and complex SAN.
- ☑ The performance of a converged infrastructure Oracle RAC database was tested with the freely available SLOB benchmark and performed as good or better than other premade database appliances who's hardware can cost at least twice as much as the commodity hardware deployed in the ScaleIO solution.
- ☑ The converged infrastructure Oracle RAC performance tests showed that the ScaleIO solution was capable of delivering low sub-millisecond response times in addition to high IOPS and throughput.
- ☑ The results of internal scalability testing were audited by ESG and showed that a 53-node ScaleIO QA deployment was capable of over 8.5 million IOPS and 114 GB/sec, and this performance was still limited only by the amount of hardware available.
- ☑ ESG Lab validated the elasticity and reliability of the ScaleIO cluster by validating the ability to self-heal by adding, removing, and failing nodes while continuing to run database activity non-disruptively and delivering linear, predictable performance.

Issues to Consider

- ☑ The majority of the testing in this report was conducted using Cisco servers and flash media to push the systems to higher levels of performance for demonstration purposes. It should be noted that organizations have the flexibility to choose the makeup of CPU and storage technologies that provide the best balance between cost, capacity, and performance to meet their requirements.
- ☑ While performance ultimately will vary based on choice of servers, storage media, network congestion, and workload characteristics, ScaleIO offers the same level of scalability and functionality for all installations.
- ☑ The testing in this report focused solely on physical server installations. Although outside the scope of this report, ScaleIO can also be deployed in virtual environments, and can even be deployed in the cloud.
- ☑ The graphical user interface of ScaleIO at this point is for monitoring purposes only. All management tasks are completed with simple and intuitive CLI commands. While this allows for agility in deployment and easy third-party integration, ESG Lab looks forward to the ability to perform common management tasks in the user interface such as volume creation, server mapping, and snapshots.
- ☑ By design, ScaleIO's data protection is limited at this point to meshed, two-copy mirroring only. ESG Lab would like to see the option to protect with more copies of the data, or other data protection schemes or capacity efficiencies that offer more flexibility in capacity and DR planning.
- ☑ While ESG Lab was particularly impressed with the performance and scalability of ScaleIO, for customers who want to deploy a hyper converged infrastructure, it will be interesting to see if and how the product evolves and approaches enterprise class availability and features while maintaining a balance between simplicity and functionality.

The Bigger Truth

Organizations are under constant scrutiny to solve problems that are drastically increasing in complexity and performance requirements with budgets that remain flat. Moore's Law only partially satisfies the requirement, so IT decision makers must also look to innovative solutions that meet their needs while also minimizing the total cost of ownership (TCO) and maximizing the return on investment (ROI). The latest ESG research indicates that ROI and reduction in operational expenditures are the two most important considerations for justifying IT investments over the next 12 months. ESG research also identifies server, storage, and network infrastructure as the top three most reported priorities over this same 12-month span. [5]

ScaleIO software leverages the freedom of choice and cost effectiveness of commodity hardware to converge new or existing compute, SAN, and storage resources into a single, scalable, building block to help solve IT's most challenging problems. ScaleIO can be deployed on any server whether physical, virtual, or cloud-based to create a server SAN of shared, mixed media, storage that is accessible to all other servers.

ESG Lab validated the convergence, performance, elasticity, and reliability of a ScaleIO solution in an eight-node, high-performance Oracle RAC environment. This deployment, built on commodity hardware, achieved over 830K random SLOB SQL query IOPS and performed over 21 GB/sec when scanning Oracle tables. ScaleIO's self-healing capability allowed ESG Lab to easily grow, shrink, and fail components in the cluster non-disruptively. ESG Lab was very impressed with the performance, flexibility, and simplicity of ScaleIO. While this testing focused on an Oracle RAC deployment, very few applications would not benefit from this simple, scalable, converged infrastructure.

A product named ScaleIO had better scale. ESG Lab audited the results of an internal scaling test of a 53-node cluster that achieved over eight million IOPS and 114 GB/sec. ESG was even made aware of other internal testing that showed linear scalability in a cloud deployment of 995 nodes. ScaleIO did not only claim this extreme level of scale in theory, but also most impressively, demonstrated it in practice.

ScaleIO does not only offer the technical agility to meet demand, perhaps more importantly, it also offers the financial agility to meet demand cost effectively. Because ScaleIO is a software-only solution, organizations can take advantage of the most cost-effective commodity hardware that best suits their application requirements at the time they need it. Capacity planning is greatly simplified as the cluster can be easily grown when needed, even mixing hardware characteristics or repurposing capitalized hardware, to arrive at the perfect balance of price, capacity, and performance. ScaleIO software is licensed in one-time, 1TB raw capacity increments.

The management of ScaleIO is extremely simple and can be performed by a single administrator rather than maintaining dedicated, compute, SAN, and storage experts. This, combined with the simple maintenance enabled by the self-healing properties of the cluster, can save an organization hundreds of thousands of dollars in operational expenses.

By testing ScaleIO in an Oracle RAC configuration, ESG Lab demonstrated the ability to perform in an environment where performance and elasticity are top concerns. It should be noted that this is by no means the only environment that is well suited to a ScaleIO deployment. The flexible configurations offered by ScaleIO also make it a great solution for a wide variety of applications ranging from VDI and server consolidation to internally developed scale-out web applications, and any other application that favors a simple, elastic, scale-out architecture.

ScaleIO was an Israeli startup that had come out of stealth mode for only a few months before being acquired by EMC in June of 2013. EMC has had a history of identifying and acquiring disruptive technologies at the right time and in ScaleIO, EMC believes it has found yet another. If your organization is struggling to deploy technologies that are flexible enough to meet unpredictable application demands while also minimizing costs, then ESG Lab strongly suggests that you make a small investment to license and try ScaleIO today. You may just end up finding your solution for tomorrow.

[5] Source: ESG Research Report, *2014 IT Spending Intentions Survey*, February 2014.

Enterprise Strategy Group | **Getting to the bigger truth.**

20 Asylum Street | Milford, MA 01757 | Tel: 508.482.0188 Fax: 508.482.0218 | www.esg-global.com