

Lab Validation Report

DataCore SANsymphony and DataCore Hyper-converged Virtual SAN

Application-adaptive Data Infrastructure Software

By Tony Palmer, Senior Lab Analyst and Jack Poller, Lab Analyst

January 2016

Contents

Introduction	3
Background	3
DataCore Software	4
ESG Lab Validation	5
Getting Started	5
Hyper-converged Data Infrastructure	9
Performance	13
High Availability	18
ESG Lab Validation Highlights	22
Issues to Consider	22
The Bigger Truth	23
Appendix	24

ESG Lab Reports

The goal of ESG Lab reports is to educate IT professionals about emerging technologies and products in the storage, data management and information security industries. 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 DataCore.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. Copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482.0188.

Introduction

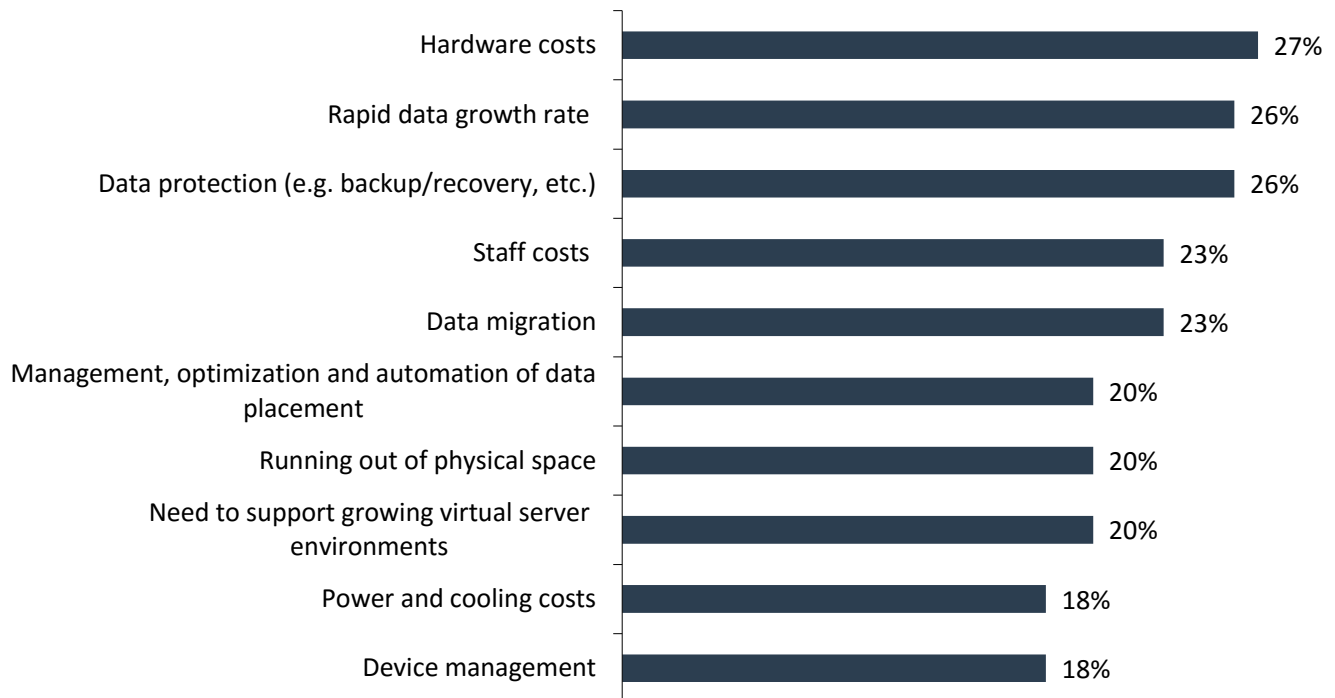
ESG Lab evaluated [DataCore](#) SANsymphony and DataCore Hyper-converged Virtual SAN software with a focus on assessing the value of infrastructure-wide storage virtualization and hyper-converged solutions for the data center. Testing focused on centralized management, ease of use, capacity efficiency, data mobility, and performance. SANsymphony's enhanced high availability and disaster recovery capabilities are also examined.

Background

ESG recently conducted a survey of 373 IT professionals and respondents were asked to identify what they would consider to be their biggest challenges with respect to their storage environment. As one might expect from a captive audience of server virtualization users, there was significant focus on data growth (cited by 26% of respondents), data protection (26%), staff costs (23%), and data migration (23%). Perhaps of greatest interest and significance, however, is that hardware costs (27%) was the most-cited storage challenge, as seen in Figure 1.¹

Figure 1. Top Ten Biggest Storage Challenges

In general, what would you say are your organization's biggest challenges in terms of its storage environment? (Percent of respondents, N=373, top ten shown)



Source: Enterprise Strategy Group, 2016.

The data center as we know it is in the midst of a significant transformation; the data center of today is increasingly virtualized and next-generation data centers will incorporate more public, private, and hybrid cloud-based applications. Virtualization and cloud bring server consolidation, workload mobility, self-provisioning, management, multi-tenancy, and the ability to rapidly scale up and down. Traditional data center storage solutions, customarily designed for static workloads tied to physical servers, are challenged to provide agility, security, data protection, and performance without costly and proprietary hardware.

¹ Source: ESG Research Report, [2015 Data Storage Market Trends](#), October 2015.

DataCore Software

DataCore has been providing software-based storage virtualization solutions for well over a decade with more than 10,000 customers and over 30,000 licenses deployed. DataCore's mission is to enhance the value of the storage hardware users prefer to own by maximizing the performance, availability, and utilization of those resources. Their SANsymphony Software-defined Storage Platform and DataCore Hyper-converged Virtual SAN products can be used independently or combined in several different ways, some of which we explore below.


DataCore SANsymphony and DataCore Hyper-converged Virtual SAN Software

SANsymphony works across all popular brands and models of disk, flash, and hybrid storage arrays, providing an integrated and consistent set of provisioning, data protection, and performance acceleration functions on an infrastructure-wide basis.

SANsymphony is designed as a flexible software platform that runs on virtual or physical industry-standard x64 servers. The product is agnostic with regard to the underlying storage hardware and can effectively virtualize whatever storage is on a user's floor, whether direct-attached, SAN-connected or in the cloud. SANsymphony supports a wide array of storage devices and host connectivity options, including iSCSI, Fibre Channel, and Fibre Channel over Ethernet (FCoE) determined by the HBAs or NICs installed in the DataCore nodes.

DataCore software can be installed on dedicated servers to manage very large multi-petabyte, physical storage pools, or can be configured as a hyper-converged solution, residing in the same physical server hosting virtual machines for applications, as shown in Figure 2. In Microsoft Hyper-V environments, the software can be installed on the root partition, without having to live inside a virtual machine.

Figure 2. DataCore SANsymphony and DataCore Hyper-converged Virtual SAN



The SANsymphony and Hyper-converged Virtual SAN solutions are built with business- and mission-critical applications in mind. The software is designed to enable users to provision, share, reconfigure, migrate, replicate, expand, and upgrade storage without downtime or performance impact.

Both of DataCore's solutions leverage storage virtualization and parallel I/O technologies to maximize data infrastructure flexibility, productivity, and cost savings. They offer users a rich set of enterprise-class functionality, with workflow-oriented wizards to ease administration and management. Parallel I/O, high-speed caching, auto-tiering, random write accelerator, and quality of service (QoS) features are dedicated to enhancing storage performance. Synchronous mirroring, asynchronous replication, in-pool mirroring, snapshots, and continuous data protection (CDP) provide data availability. Online migration and upgrades, thin provisioning, storage pooling, and data deduplication and compression increase efficiency. All of these features can span heterogeneous storage with the DataCore software running between the storage hardware and the application hosts.

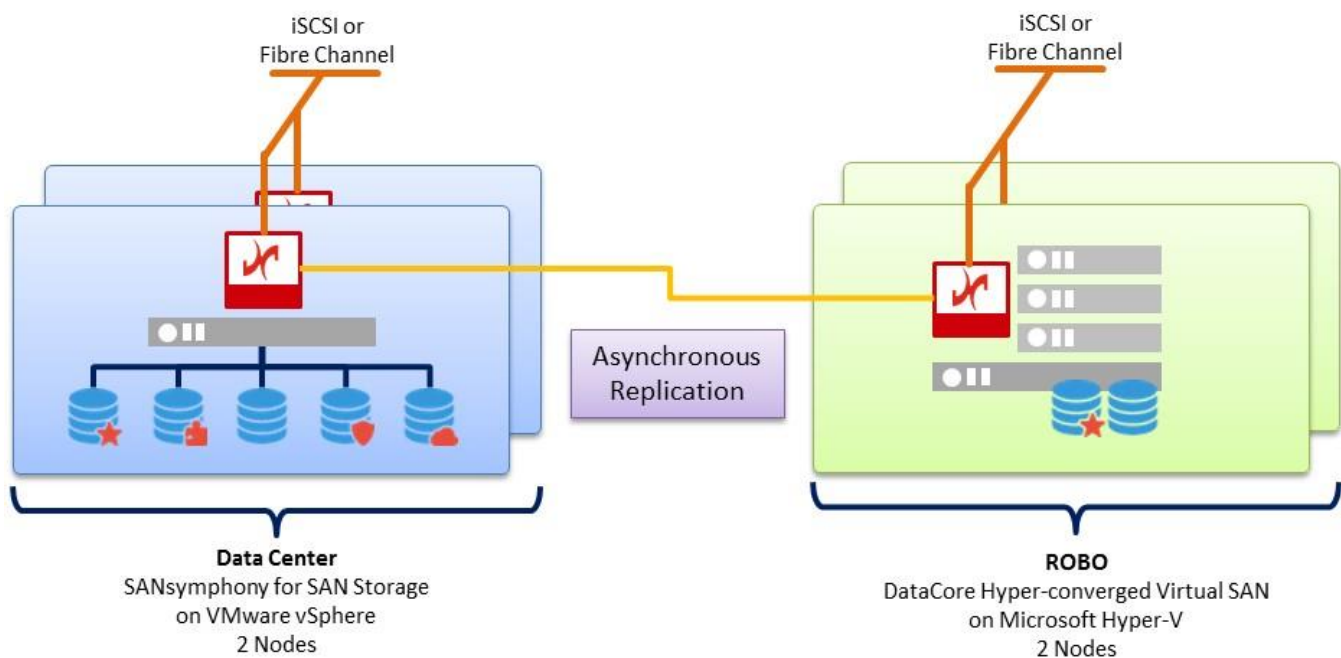
ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of DataCore SANsymphony and Hyper-converged Virtual SAN solutions at DataCore's facility in Fort Lauderdale, Florida. Testing was designed to validate the flexibility and ease of management in a heterogeneous, highly virtualized environment. ESG also examined DataCore software's performance, efficiency, and availability in fully virtualized and hyper-converged configurations.

Getting Started

Figure 3 illustrates the test bed used by ESG Lab for this validation report. It represents a common configuration using SANsymphony at the primary data center connected to DataCore Hyper-converged Virtual SAN at remote office/branch office (ROBO) sites. Two physical servers were utilized, with VMware vSphere 6.0 installed on one and Microsoft Windows Server 2012 R2 with Hyper-V on the other. Two virtual machines on the vSphere server were configured as SANsymphony storage virtualization nodes, while DataCore Hyper-converged Virtual SAN was installed on the root partition of two virtual Hyper-V nodes.

Figure 3. The ESG Lab Test Bed



The DataCore management console (which can be run local or remotely) provided a centralized management interface that was used to perform the configuration, management, and monitoring tasks across both environments. Any console command may also be issued from PowerShell scripts to coordinate workflows.

Self-Provisioning from vSphere with VVols

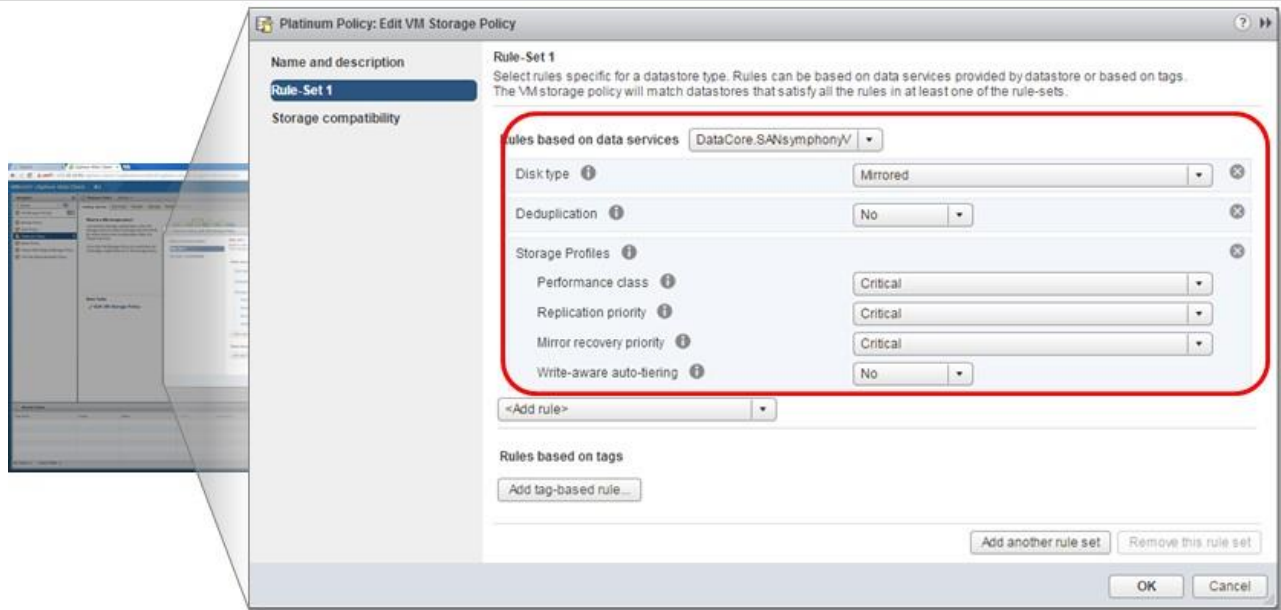
First ESG Lab tested self-provisioning with VMware Virtual Volumes (VVols). VMware designed VVols to be an integration and management framework for external storage with the goal of providing control at the VM-level to simplify storage operations by centering them on the VM rather than on the physical infrastructure.

Put another way, rather than creating individual volumes and datastores to be assigned to groups of virtual machines, a storage administrator can simply create storage containers with specific properties and advertise them for use. Then the vSphere administrator simply selects the appropriate container when provisioning a new virtual machine. The vSphere storage (VASA) API connects the new VM to the corresponding storage container supplied by DataCore. In this scenario we set up Platinum Gold, Silver, and Bronze policies to define different levels of service.

ESG Lab Testing

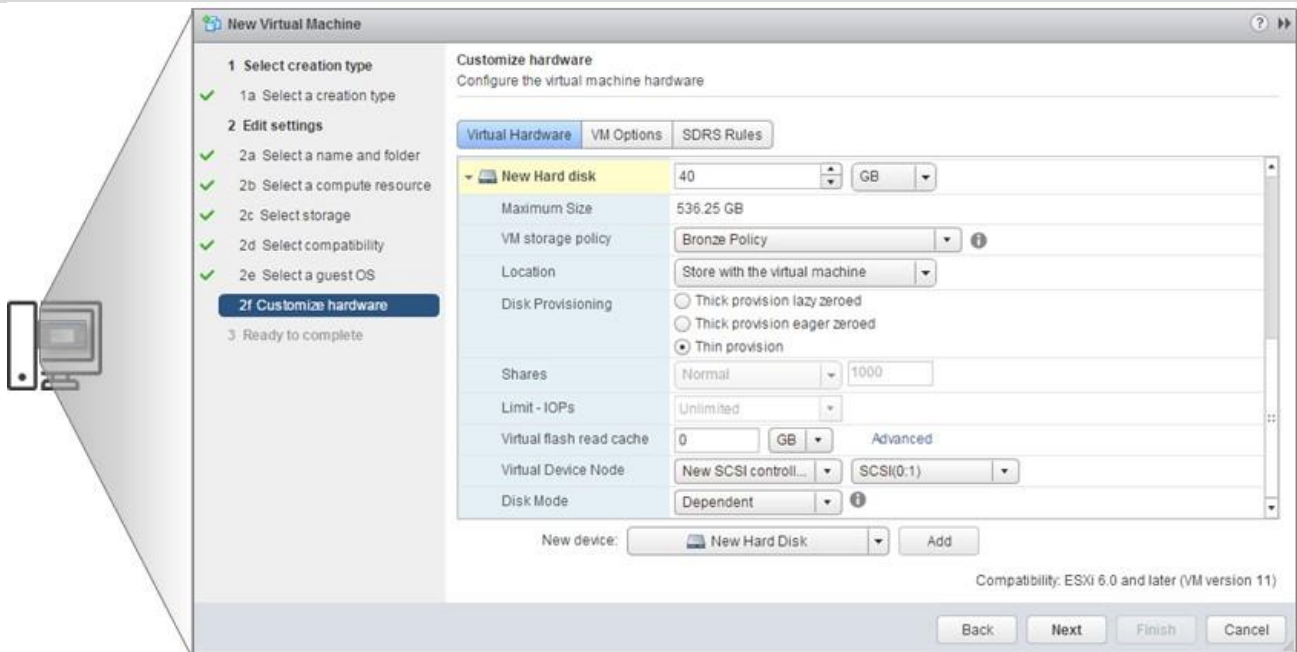
ESG Lab initialized the environment by defining a few storage policies from the vSphere client, as shown in Figure 4. When first setting up a policy, the administrator selects rules for the datastore type desired. The options presented are based on the published capabilities of DataCore's virtual storage pools. For the Platinum policy, ESG selected **mirrored** (synchronously across nodes) for the disk type, **no deduplication**, and **critical** for the highest performance, replication, and mirror recovery priority.

Figure 4. Setting Up a Storage Policy from vSphere



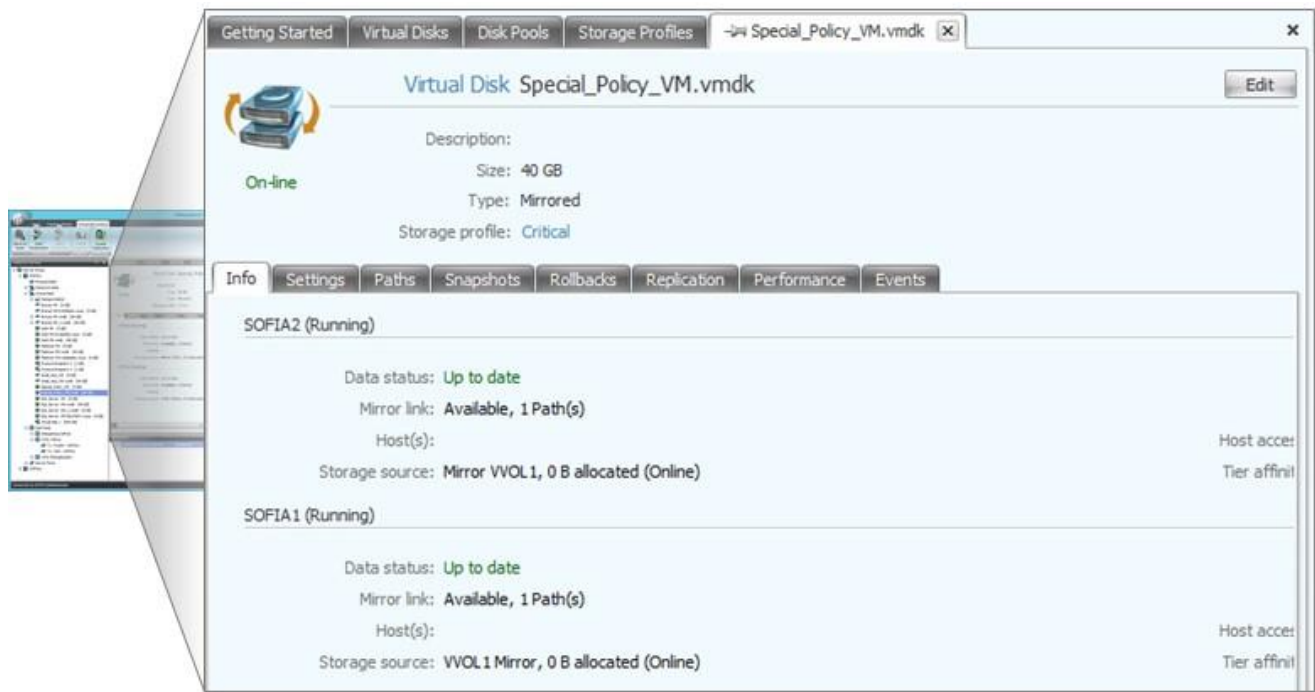
A Bronze storage policy was also set up with a disk type of **Single** (only one copy on one node), and **Deduplication** enabled. With those one-time steps completed, vSphere administrators can provision storage for themselves without having to know anything about the DataCore software. Using the standard VMware interface, ESG Lab created one virtual machine called *Special_Policy_VM* using the Platinum storage policy, and another virtual machine called *Small_test_VM* using the Bronze storage policy, as shown in Figure 5.

Figure 5. Creating a VM Using the Bronze Storage Policy



To understand how the storage is being allocated to those VMs, ESG Lab examined the corresponding VMDK files using the DataCore management console. Details for the *Special_Policy_VM* virtual machine, created with the Platinum storage policy are shown in Figure 6.

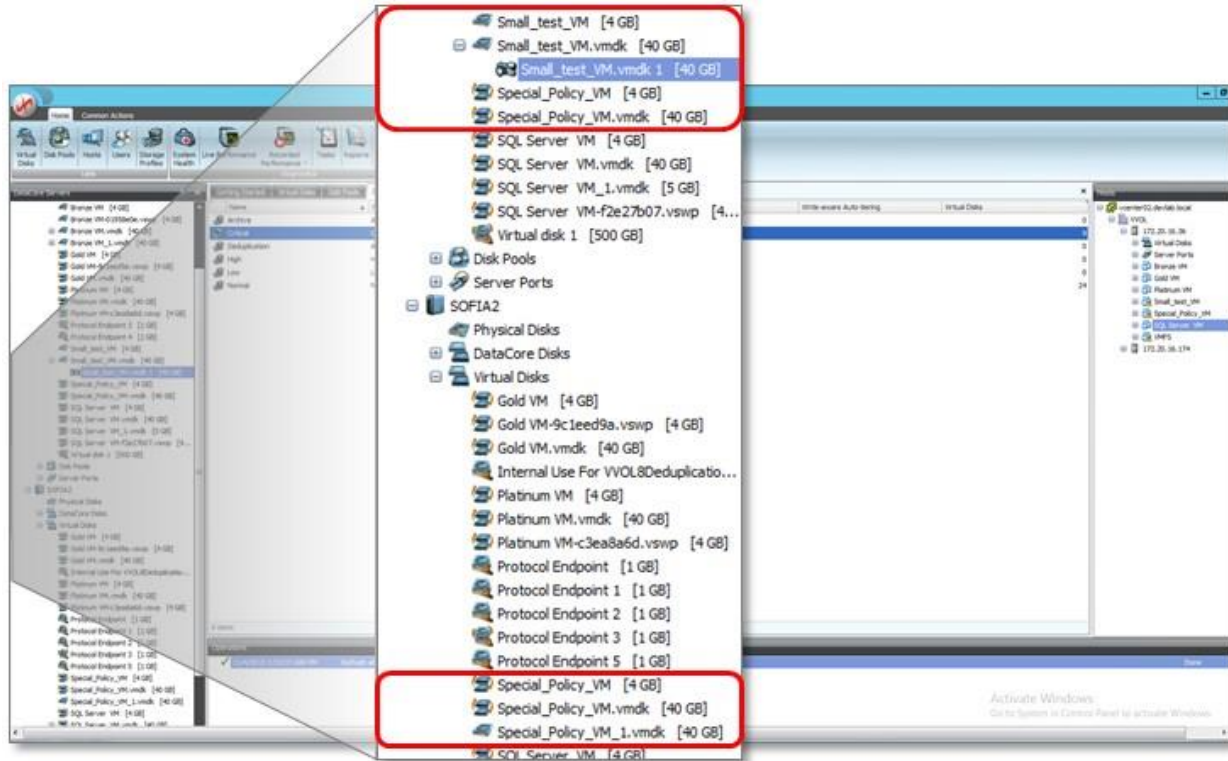
Figure 6. Managing VMDK File Storage from the DataCore Console



Finally, ESG Lab looked for the VMDK files for both virtual machines using the DataCore management console.

Figure 7 shows that the VMDK files for *Special_Policy_VM*, created with the Platinum storage policy, are mirrored across both SANsymphony nodes while the VMDK files for *Small_test_VM*, created with the non-mirrored Bronze storage policy, reside only on one node, as expected.

Figure 7. VMDKs Mirrored and Non-mirrored Across DataCore nodes



Why This Matters

As virtualized IT environments continue to grow and evolve, the ability to quickly provision and easily manage capacity is essential if organizations are to provide cost-effective services to the business. Forty seven percent of organizations surveyed by ESG cited this need for greater agility to better align with the needs of the business as driving consideration of software-defined storage solutions.²

DataCore SANsymphony VVols integration made provisioning storage a seamless, integrated part of VM creation, empowering VMware administrators to provision storage for virtual machines using predefined storage policies that define performance, availability, and locality of data without ever having to touch the back-end storage. ESG Lab created two storage profiles and two virtual machines based on those profiles in minutes using native VMware tools.

SANsymphony automatically took care of the entire behind-the-scenes configuration, which is the responsibility of specialized storage administrators in traditional storage environments. ESG Lab was impressed with the speed, simplicity, and completeness of the DataCore integration.


It's worthwhile to note here that although much of the storage in data centers today does not support VVols, and can never be retrofitted to do so, SANsymphony extends VVols self-provisioning to those resources and any new devices customers may choose in the future.

² Source: ESG Research Report, [2015 Data Storage Market Trends](#), October 2015.

Hyper-converged Data Infrastructure

DataCore Virtual SAN software is a hyper-converged solution that combines storage, compute, and virtualization into a compact industry-standard x64 server footprint. By combining the hardware resources from each node into a shared-storage pool, DataCore can deliver simplified operations, improved agility, and greater flexibility for virtualized workloads. The complete set of advanced features is available for VMware and Microsoft environments as well as other hypervisors, and each node can be configured with any combination of SSDs for workloads requiring high performance and low latency, and HDDs for lower cost, higher-capacity storage. Server RAM acts as high-speed cache, which, together with DataCore Parallel I/O further accelerates performance for latency-sensitive applications.

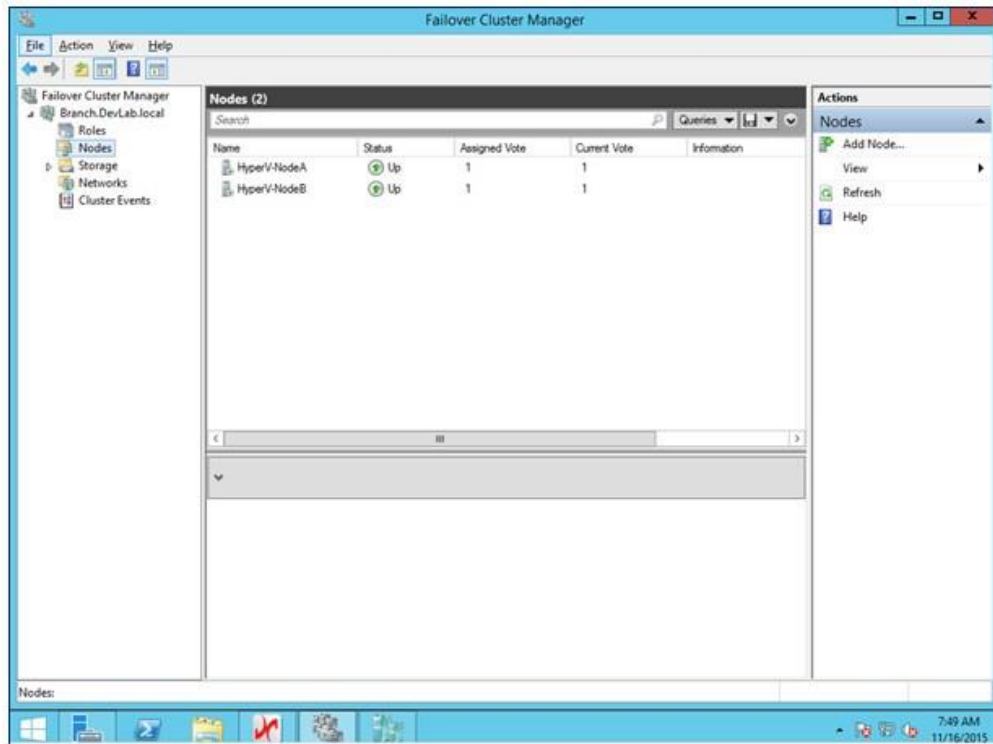
Figure 8. DataCore Hyper-converged Virtual SAN



ESG Lab's goal in these tests was to demonstrate how DataCore Virtual SAN can be used to provide a highly available hyper-converged environment hosting applications and their storage for organizations' data centers and remote office/branch office (ROBO) environments.

Figure 9 shows the test environment—a two node Microsoft Hyper-V Failover Cluster with DataCore Virtual SAN software running on the root partition of each server. From the Windows Failover Cluster console the nodes appear as HyperV-NodeA and HyperV-NodeB. The DataCore software provides active-active copies of the internal storage between the two nodes, yet the cluster sees them as a single shared storage resource that passes cluster qualification tests.

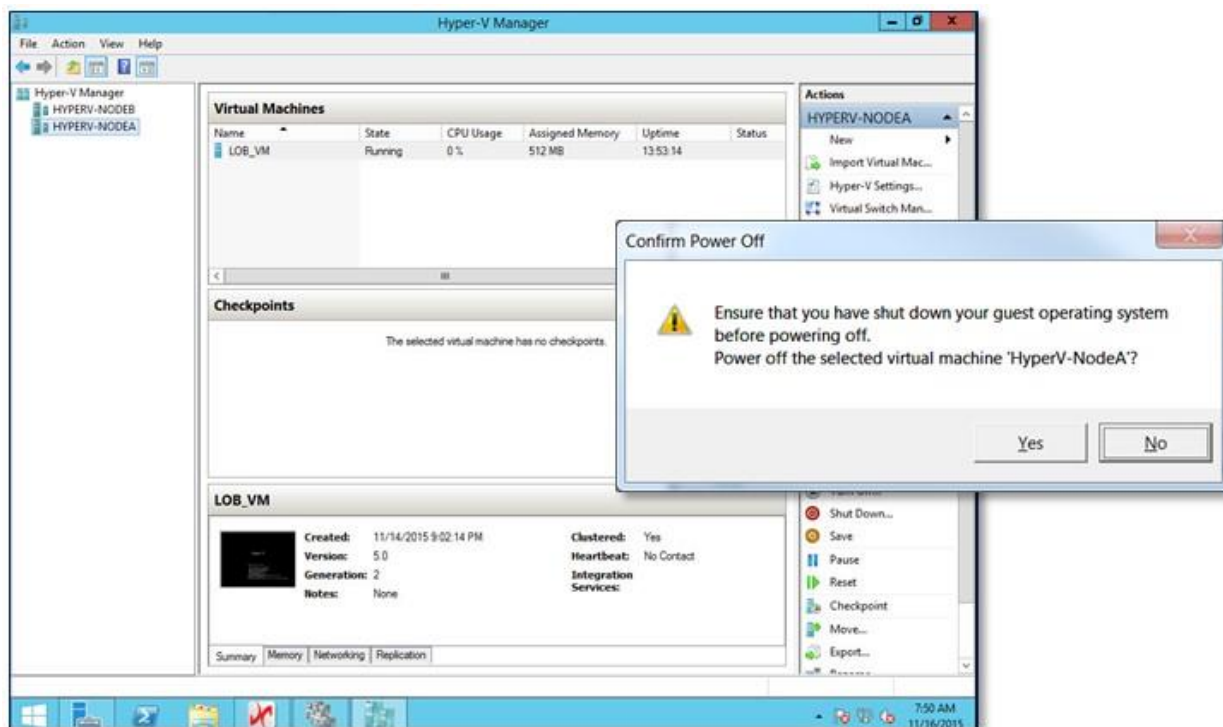
Figure 9. DataCore in a Hyper-converged Failover Cluster using Hyper-V



ESG Lab Testing

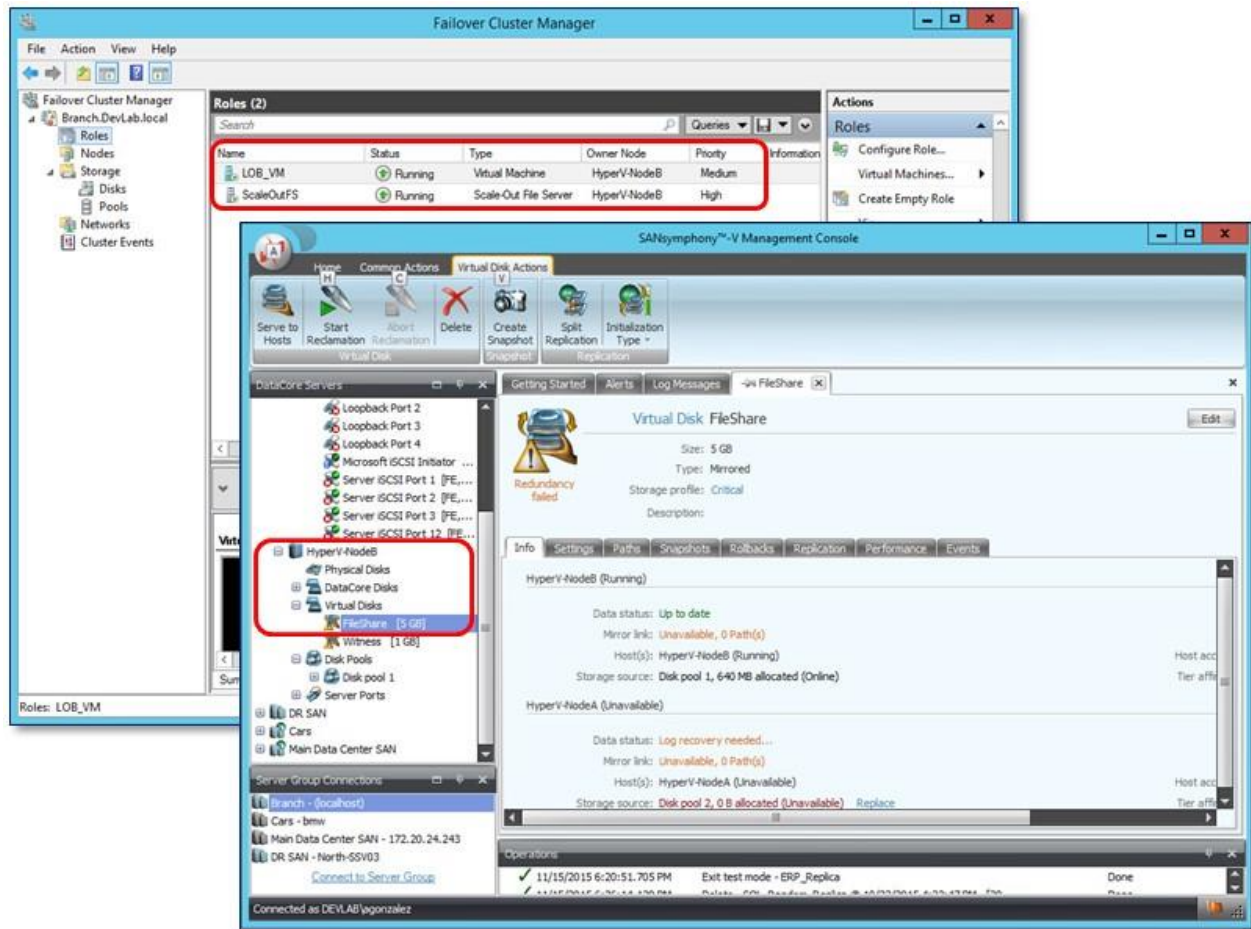
First, *HyperV-NodeA*, with the *LOB_VM* virtual machine running in it, was powered off. The machine was powered down without shutting down the OS to simulate an abrupt failure as seen in Figure 10.

Figure 10. Powering Off a Node



Failover of the virtual machine from *HyperV-NodeA* to *HyperV-NodeB* was fast and automatic. As seen in Figure 11, the *LOB_VM* virtual machine was up and running on *HyperV-NodeB* using the synchronously-mirrored copy of the VM.

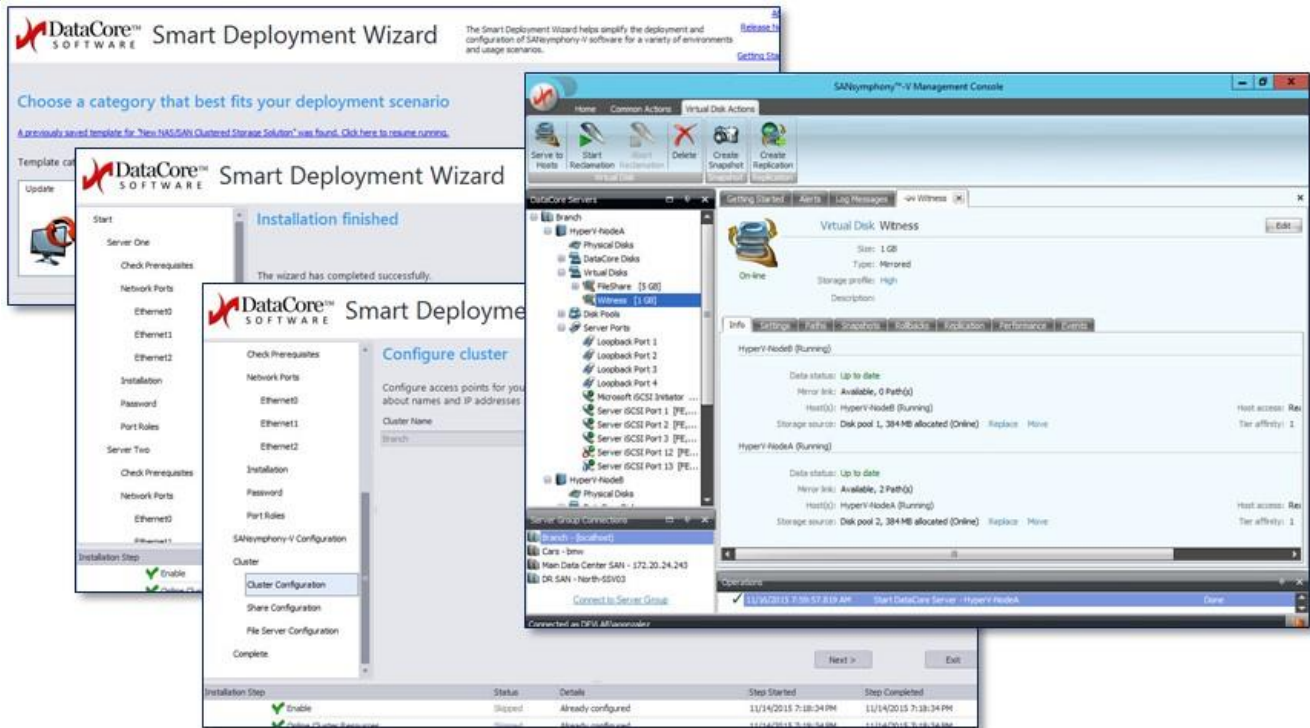
Figure 11. Storage and VMs Failed Over to the Surviving Node



Continuing with the failed server scenario, we then ran the DataCore Smart Deployment Wizard to simulate bringing a new server into the Failover Cluster. We selected **Clustered Virtual Machines** to configure a hyper-converged environment.

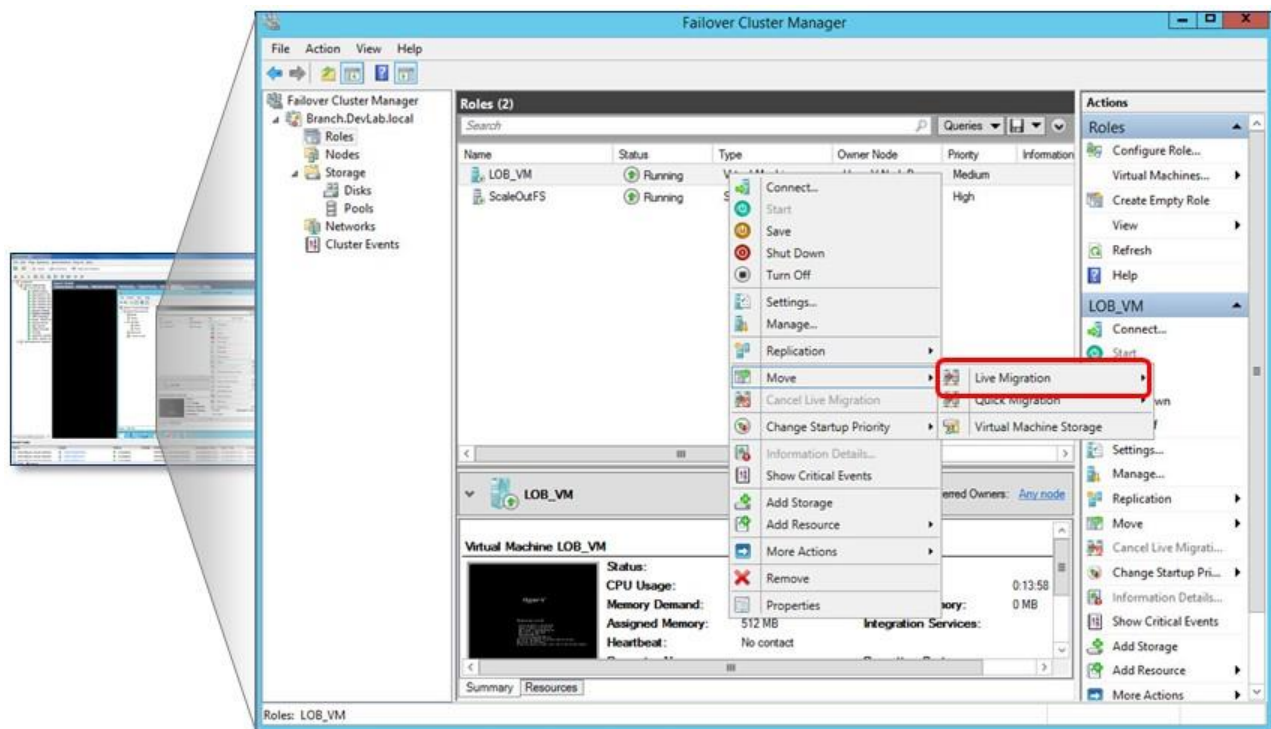
Figure 12 shows the progression of the smart deployment wizard through the steps to install DataCore Virtual SAN software on a new server and add it to the Hyper-V cluster. The wizard also offers to set up a fully-redundant Scale-out File Server for a combined NAS/SAN solution.

Figure 12. The Smart Deployment Wizard



Installation of the Hyper-V Failover Cluster configuration was automated using the wizard to assist the user from start to finish. ESG Lab used live migration to move the VMs to the new *HyperV-NodeA*, as seen in Figure 13.

Figure 13. Moving the Running Virtual Machine Back to HyperV-NodeA



The entire process, from start to finish, including installing the software, connecting the server to the cluster, and fully resynchronizing the data, was accomplished with a handful of mouse clicks and no outages in less than 30 minutes.

Why This Matters

Organizations have embraced hyper-converged technology for a number of reasons, including a faster and simplified deployment process, improved services and support, and converged management of traditionally disparate resources. With such a long list of benefits, ESG asked respondents what the primary reason was for deploying or considering deploying a hyper-converged solution, and the top response was improved total cost of ownership.³ Combine this finding with the fact that return on investment is a top reported consideration for justifying new IT investments in general, and it makes sense that hyper-converged solutions are growing in popularity due to the mix of simplicity and cost savings that they offer.⁴

ESG Lab found that virtualizing infrastructure with DataCore Hyper-Converged Virtual SAN was intuitive and straightforward. A pair of Windows 2012 servers running Hyper-V and DataCore Virtual SAN provided a highly available platform for multiple simulated applications. When a server failure was simulated with a hard power off, virtual machines and storage failed over to the surviving node immediately and automatically. A new server was installed and added to the Failover Cluster—seamlessly and painlessly—while applications remained online and available.

Performance

DataCore leverages cost-effective server RAM, advanced caching techniques, load balancing, automated storage tiering, random write acceleration, and QoS in combination with their parallel I/O technology to enhance system performance.

Caching—DataCore cache resides between the operating system on the hosts and the physical storage, and applies to all storage devices configured throughout a user's SAN. Any application server connected to SANsymphony benefits from cache acceleration regardless of the physical location of the cache. The cache is implemented using commodity server RAM, which is typically less expensive than dedicated storage array cache, yet far faster than flash or SSD-based caches, providing cost efficiency and excellent price-performance.

Load Balancing—Load balancing overcomes short-term bottlenecks that may develop when the queue to a given disk channel is full, or when one channel fails or is taken offline. Hosts can take advantage of all communication channels available between application servers (hosts) and SANsymphony because SANsymphony presents the virtual disks to all front-end channels simultaneously. SANsymphony automatic disk pool balancing ensures an even spread of all data blocks across the physical disks in the pool. In addition to improving performance, disk pool balancing avoids hot spots when modifying the number of disks in the pool.

Automated Storage Tiering—SANsymphony's automated storage tiering dynamically moves data between different disk types and spans flash, disk, and cloud technologies to meet space, performance, and cost requirements. SANsymphony monitors I/O behavior to determine frequency of use, and automatically promotes the most frequently used data blocks to the fastest tier and demotes the least frequently used data blocks to the slowest, lower cost tier. Storage profiles align the process with service level objectives (SLAs).

Random Write Accelerator—Writes to random storage locations present multiple performance challenges to storage systems. For small blocks of data, a read is required before the write to correctly calculate the data protection codes (RAID parity). When writing to traditional spinning disks, physical head movement and rotational delays incur performance penalties. Conversely, while SSDs don't have to deal with physical parts, each write of a data block is preceded by an erase operation. DataCore overcomes these random write challenges and accelerates performance by storing the random writes in sequential order in new locations.

Quality of Service—DataCore provides QoS controls to ensure that high-priority workloads competing for access to storage can meet their service level agreements (SLAs) with predictable I/O performance. QoS controls regulate the

³ Source: ESG Research Report, [Trends in Private Cloud Infrastructure](#), April 2014.

⁴ Source: ESG Research Report, [2015 IT Spending Intentions Survey](#), February 2015.

resources consumed by workloads of lower priority. QoS controls can be applied to both individual hosts and host groups, and can simultaneously regulate the data transfer rate (MB/sec) and I/O rate (IOPS).

Parallel I/O Software and Performance

Parallel I/O—DataCore’s software-defined storage and Parallel I/O software technology are designed to adaptively harness available multi-core processors to optimize and schedule I/O processing across many different cores simultaneously. It actively senses I/O load being generated by multiple VMs concurrently and dynamically assigns CPU cores as needed to process the I/O load. This enables DataCore to take full advantage of modern multi-core server technologies to eliminate I/O bottlenecks, speed up application performance, and drive greater workload and virtual machine density per server.

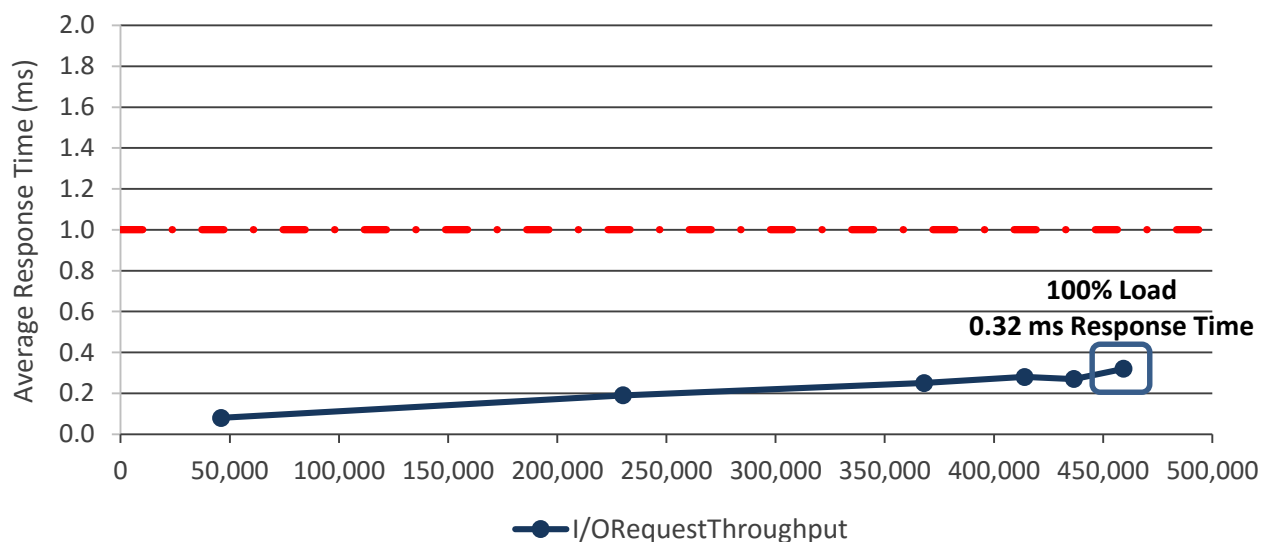
SPC-1 Results

ESG Lab reviewed DataCore’s published results of the SPC-1 application-level industry-standard benchmark suite maintained by the Storage Performance Council. SPC-1 testing generates a series of workloads designed to emulate the typical functions of transaction-oriented, real-world database applications. Transaction-oriented applications are generally characterized by largely random I/O and generate both queries (reads) and updates (writes). Examples of those types of applications include OLTP, database operations, and mail server implementations. SPC results can be roughly mapped by users into easily understood metrics. For a credit card database system, for instance, it might be the number of credit card authorizations that can be executed per second.

It is important to note that the SPC-1 benchmark consists of over 60% writes, a mix of random and sequential I/O, and a variety of block sizes. As such, the results should not be compared with marketing performance numbers consisting of 100% random reads with a homogeneous block size.

DataCore has published an excellent result of 459,290 SPC-1 I/O requests per second at 100% load with an average response time of only 0.32 milliseconds in a hyper-converged configuration.⁵ Figure 14 shows a response time/throughput curve, which visually represents the performance of the system under test as load is increased. A long, flat curve indicates better performance, as this denotes that response time stays low as IOPS increase. The 0.32 millisecond result at 100% load is the fastest response time ever reported on SPC-1 at the time of this writing, and showcases the power of parallel I/O software to significantly reduce the time it takes for applications to access, store, and update their data.

Figure 14. DataCore SPC-1 Results



⁵ http://www.storageperformance.org/results/benchmark_results_spc1_active#a00164

Table 1 summarizes the published results.

Table 1. DataCore SPC-1 Benchmark Results Summary

SPC-1 Load	SPC-1 IOPS	Response Time (ms)	Read Response Time (ms)	Write Response Time (ms)	\$/SPC-1 IOPS
10%	46,003.40	0.08	0.14	0.04	
50%	230,023.79	0.19	0.37	0.08	
80%	367,984.28	0.25	0.47	0.11	
90%	413,959.19	0.28	0.51	0.13	
95%	436,570.42	0.27	0.49	0.13	
100%	459,290.87	0.32	0.58	0.16	\$0.08

At 459,290 SPC-1 IOPS, DataCore SANsymphony 10.0 is currently ranked 11th overall in SPC-1 performance results. DataCore achieved this result with a solution cost of just \$38,400 compared to a range of \$488,000 to more than \$3 million for the top-ten. With high performance and low cost, SANsymphony set a new SPC-1 price-performance record, with a cost in dollars per SPC-1 IOPS of just \$0.08. At full load, SANsymphony was responding at 0.32ms, nearly two orders of magnitude below the 30ms response time threshold set by the Storage Performance Council (SPC) and less than a third of the 1ms threshold considered to be the standard for all-flash systems.

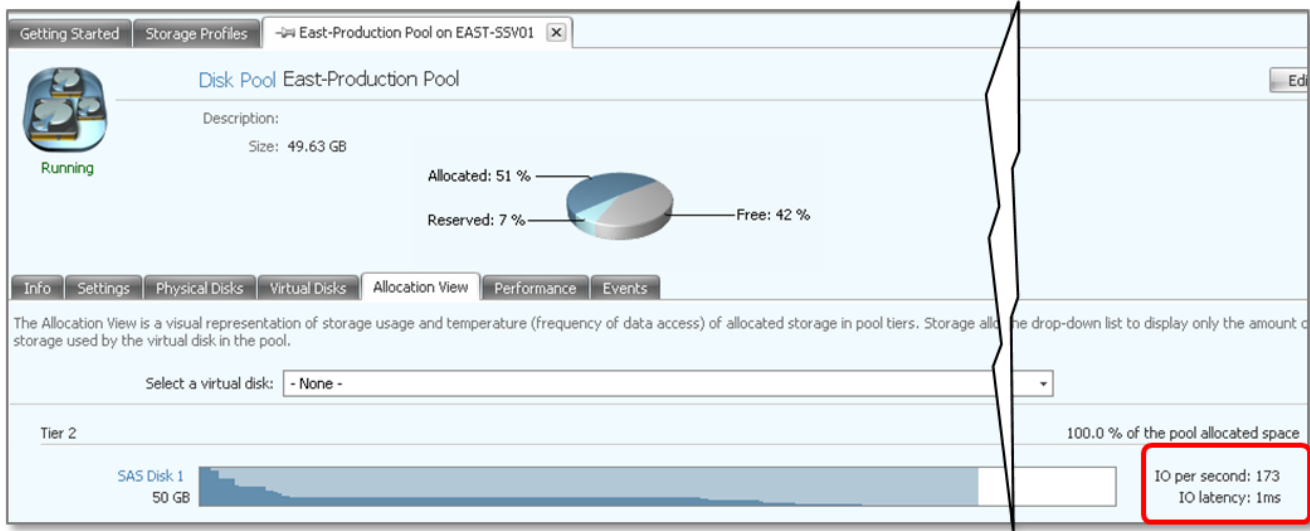
The DataCore SPC-1 result audited by the SPC and peer reviewed by competitors proves its suitability for response time sensitive applications and demonstrates the headroom available to scale up and scale out to much larger configurations and capacities.

The tested DataCore configuration is a hyper-converged solution; the benchmark workload was running on the same system that was performing storage operations, rather than on separate servers. This is a very important consideration. In contrast, the cost of the servers to drive the workload on external storage systems is not included in their dollar per SPC-1 IOPS calculation. In the real world this will translate to an even better price-performance benefit from the DataCore solution.

ESG Lab Testing

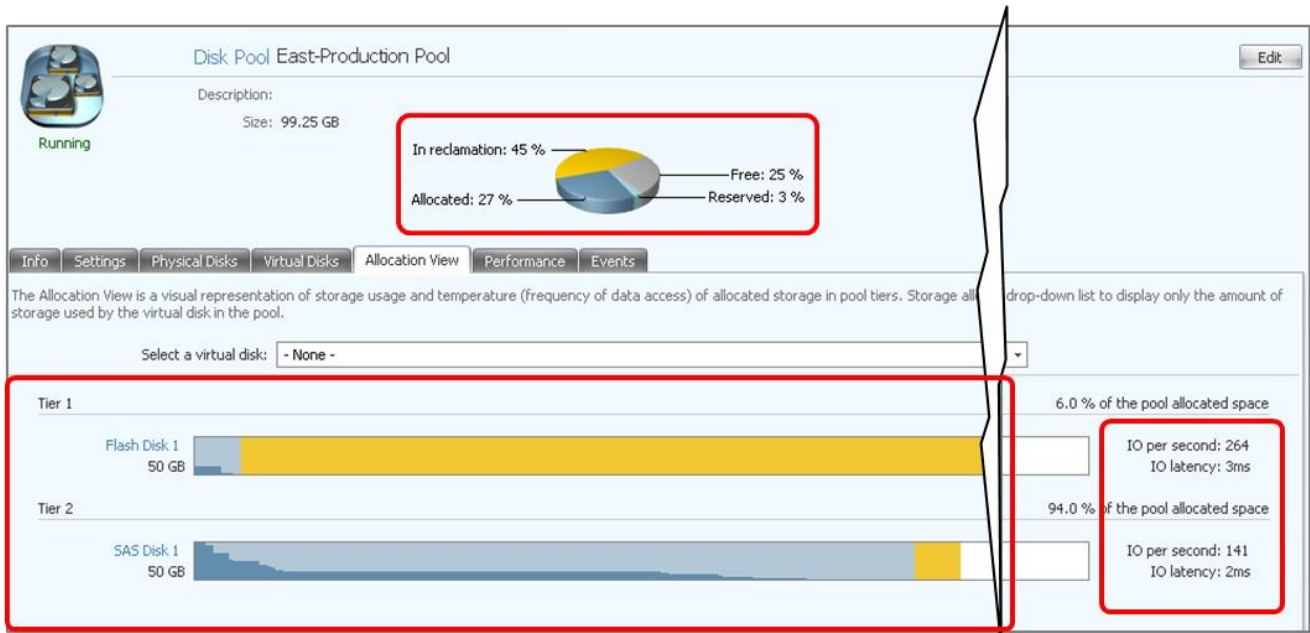
ESG Lab also validated the value of SANsymphony automated storage tiering. All performance testing was completed using the industry-standard Iometer test tool to simulate a typical OLTP workload. First, we measured the performance of SANsymphony on a non-tiered storage pool using a SAS disk. The SANsymphony management console, as shown in Figure 15, reported that the system was able to sustain 173 IOPS with 1 millisecond response times.

Figure 15. Auto Tiering Configuration—Non-tiered Storage Pool



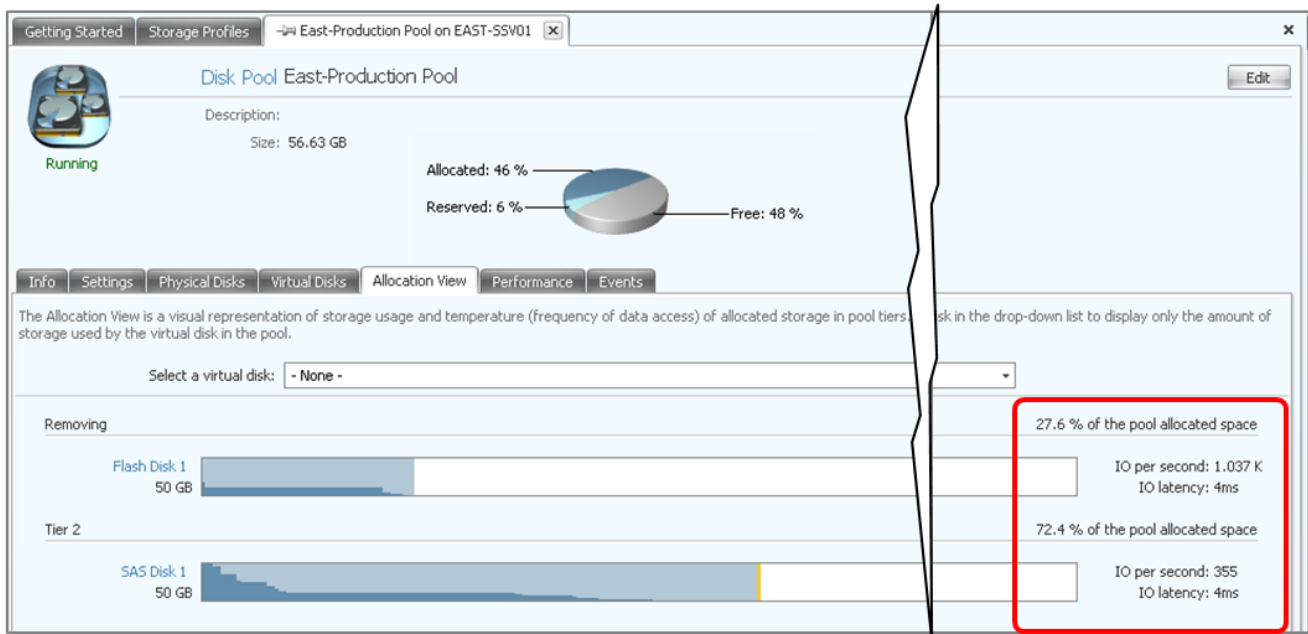
Next, we added an SSD to the storage pool. SANsymphony immediately started the process of tiering, moving hot, or frequently accessed, data blocks from the SAS disk to the higher performance flash disk. This can be seen in both the pie chart and bar graph presented on the DataCore console, as shown in Figure 16. The yellow indicates data blocks being zeroed out or reclaimed. During the rebalancing process, SANsymphony improved performance, and reported 264 IOPS and 2ms response for the flash disk and 141 IOPS and 3 ms response for the SAS disk

Figure 16. Auto Tiering Configuration—After Adding a Flash Tier



After completing the rebalancing process, the SANsymphony solution dedicated all resources to performing storage operations for the application. The effect of storage tiering was immediately visible, with SANsymphony reporting a total of 1,392 IOPS—a 700% improvement over the non-tiered storage pool, as shown in Figure 17. Moving hot data to the higher performance flash disk reduced the pressure on the queues for the SAS disk, which delivered 355 IOPS, a 100% improvement over the non-tiered storage pool configuration. Note that this was a shared lab environment not intended to show the software's highest performance.

Figure 17. Auto Tiering Configuration—Rebalance Complete



Why This Matters

Storage scalability and performance are significant challenges for the highly virtualized modern data center. With pools of virtualized databases, application servers, and desktops all relying on a shared storage infrastructure, storage performance and response time become critical considerations.

Through careful examination of SPC-1 results combined with hands-on testing of multiple applications and databases, ESG Lab has verified that DataCore SANsymphony can be deployed in the modern data center in a dedicated SANsymphony configuration or as a Hyper-converged Virtual SAN solution to cost-effectively provide extremely high-performance storage with consistently low response times. SANsymphony's automated storage tiering provided additional benefits, both improving overall IOPS by 700% and, by reducing load on lower performance tiers, improving SAS disk IOPS by 100%.

High Availability

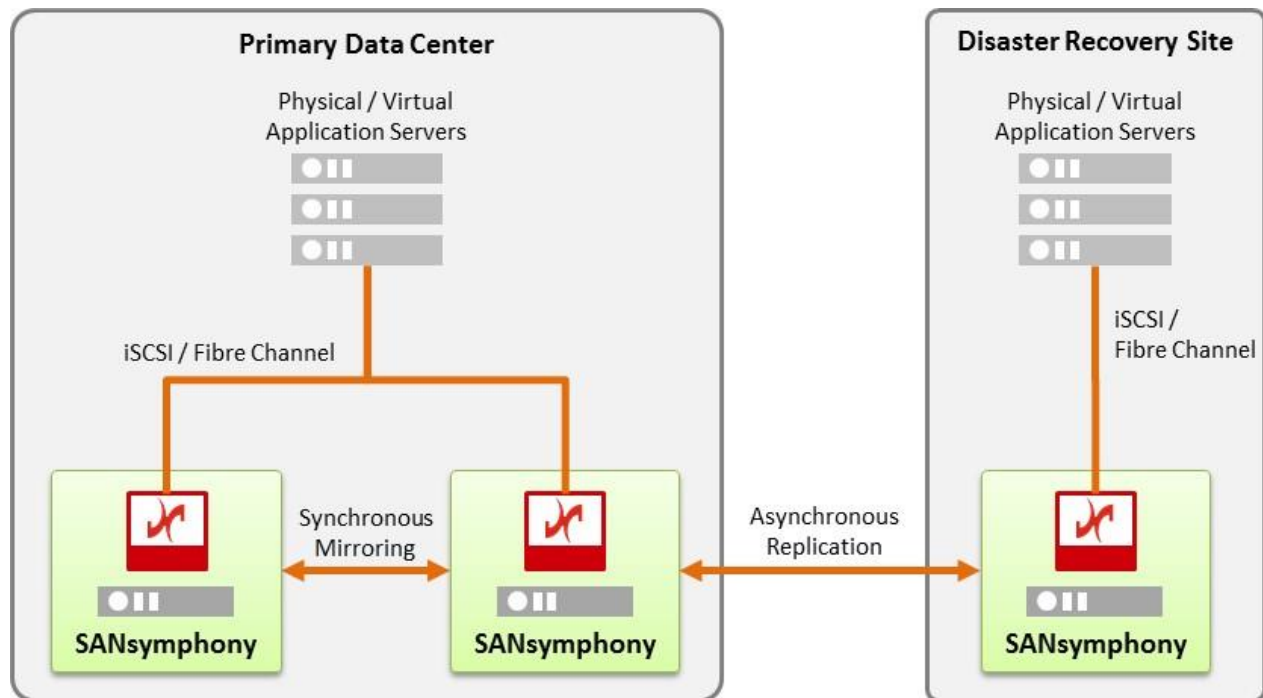
DataCore SANsymphony and DataCore Hyper-converged Virtual SAN offer multiple, complementary technologies to ensure that organizations' data remains accessible despite hardware faults, human error, and environmental disruptions. SANsymphony offers N+1 redundancy to eliminate single points of failure in a data center and across multiple data centers in a metropolitan area. Asynchronous remote replication along with Advanced Site Recovery (ASR) automate and simplify failover and failback of active workloads in the event of a regional disaster or a planned site switchover. ASR is integrated with VMware Site Recovery Manager (SRM).

This level of high availability is provided via synchronous mirroring between nodes and multiple I/O paths from hosts to nodes. DataCore recommends that customers place redundant nodes with their respective storage pools in separate rooms, ideally in separate buildings on a campus where a water leak or air conditioning problem, for example, can only disturb one of the nodes while the other transparently absorbs its load. Larger customers often operate distributed data centers split between hot sites several miles apart, with zero-touch failover and failback.

Data layer protection is provided via snapshots (full clones or copy-on-first write differentials) as well as Continuous Data Protection (CDP). When a volume is protected with CDP, SANsymphony logs and time stamps all activity to the virtual disk so that users can create a rollback at any point in time within the rollback window.

To guard against regional disasters, DataCore offers asynchronous remote replication over conventional LANs and WANs using industry-standard TCP/IP protocols. SANsymphony automatically compresses the replication stream to reduce bandwidth requirements, allowing customers to use lower cost links with narrower bandwidth. A configuration using both synchronous mirroring and asynchronous remote replication is shown in Figure 18.

Figure 18. Highly Available Data Infrastructure



ESG Lab Testing

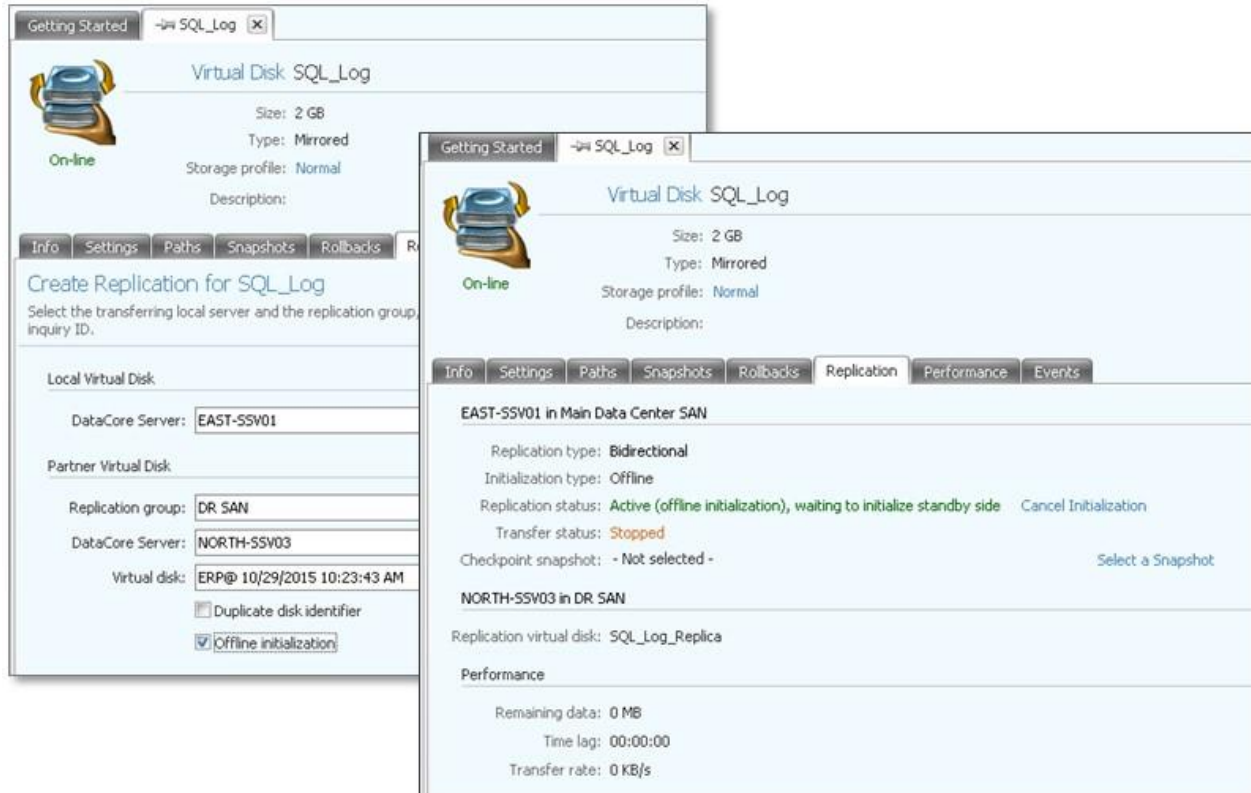
In a previous report ESG Lab tested synchronous mirroring and asynchronous replication using a simulated Exchange server.⁶ ESG Lab validated that DataCore SANsymphony provides an array of data protection capabilities that can cost-effectively satisfy the most stringent business continuity and disaster recovery requirements.

⁶ Source: ESG Lab Validation Report, [DataCore SANsymphony-V](#), April 2011.

Synchronous mirroring within metropolitan areas, snapshots, CDP, and asynchronous remote replication to distant disaster recovery sites can be used without being dependent on any specific model or brand of storage device.

DataCore asynchronous replication can initialize a volume for replication offline, preparing the data set before installation at the remote site.

Figure 19. Offline Initialization for Asynchronous Replication



Offline initialization is configured at the time of replication creation, by simply selecting a checkbox. ESG Lab configured a volume for offline initialization, as shown in Figure 19. Next, ESG Lab enabled Continuous Data Protection for the volume *SQL_Random*, as shown in Figure 20.

Figure 20. Enabling Continuous Data Protection (CDP)

Getting Started - SQL_Random [X]

Virtual Disk: SQL_Random

Size: 30 GB
Type: Mirrored
Storage profile: Critical
Description:

Info Settings Paths Snapshots Rollbacks Replication Performance Events

☒ Continuous Data Protection Enabled

DataCore Server: EAST-SSV01

Disk pool: East-Production Pool

Maximum history log size: 20 GB

System Health Thresholds

Minimum retention period warning threshold: 0 hours

Maximum history log size warning threshold: 100 % of the maximum history log size

System Health Thresholds

Allocated space thresholds:

Attention: 100 %
Warning: 101 %

I/O latency thresholds:

Attention: 5,000 ms
Warning: 10,000 ms

Storage latency thresholds:

Attention: 5,000 ms
Warning: 10,000 ms

Warning: 101 % Warning: 10,000 ms Warning: 10,000 ms

A batch file was run in a continuous loop, creating a series of small text files, one per second to simulate continuous writes to a log file.

Figure 21. Creating a Rollback

Getting Started - SQL_Random [X]

Virtual Disk: SQL_Random

Size: 30 GB
Type: Mirrored
Storage profile: Critical
Description:

Info Settings Paths Snapshots Rollbacks Replication Performance Events

Create Rollback for SQL_Random

Select a restore point and pool for the rollback. Check the duplicate disk identifier option if both source virtual disk and rollback will be served to hosts requiring the same SCSI inquiry ID. Choose to preserve or automatically expire the rollback when the history log goes past the restore point.

Name: SQL_Random @ 11/16/2015 2:20:26 PM

Description:

Restore point: 11/16/2015 2:20:26 PM

DataCore Server: EAST-SSV01

Rollback type: ☐ Expire rollback ☒ Persistent rollback

The rollback will be accessible until it is manually deleted or split. When the history log advances past the restore point the rollback source will be affected; a mirrored virtual disk will have compromised mirror redundancy, a single virtual disk will become inaccessible. The rollback type is not editable after creation.

Storage profile: Normal

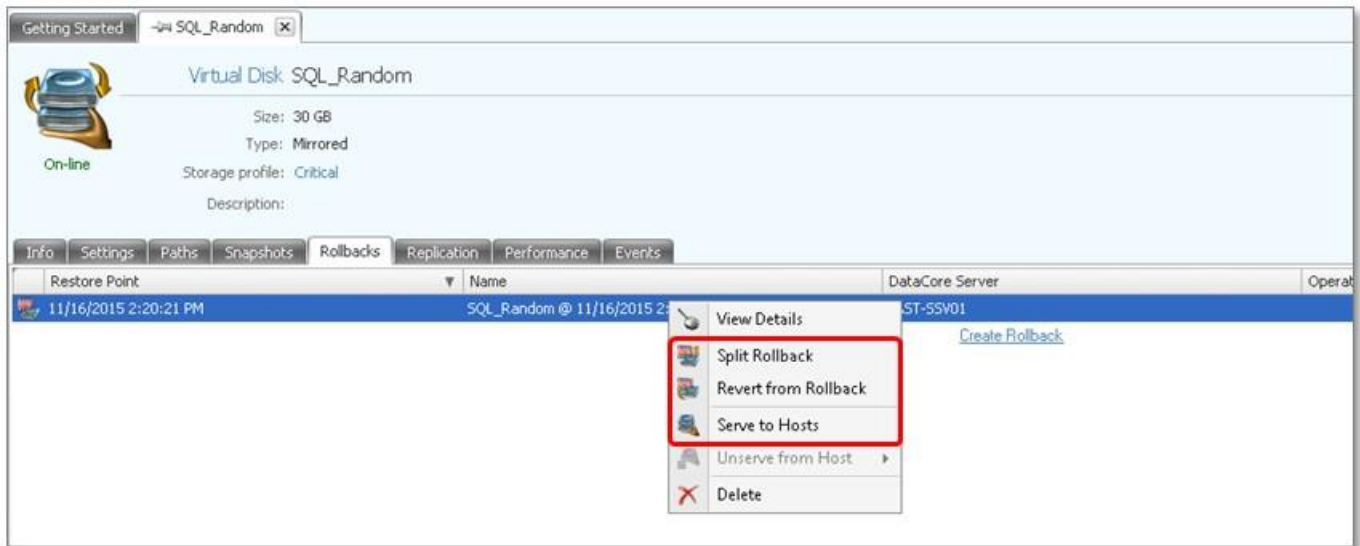
☐ Duplicate disk identifier

Pool Name	Size	Status	SAU Size	Available Storage	Single Virtual Disk Count	Mirrored Virtual Disk C...	Dual Virtual Disk Count	Oversubscription	Usage
East-Bulk Storage Pool	49.63 GB	Running	128 MB	49.50 GB	1	0	0	0.8	1%
East-Production Pool	99.25 GB	Running	128 MB	70.50 GB	3	6	0	261.75 GB	29%

Create Disk Pool

Next, ESG Lab simulated an accidental file deletion event and stopped the batch file. To recover, a rollback point was created, at 2:20:26 PM, as seen in Figure 21.

Figure 22. Utilizing a Rollback



Finally, ESG Lab accessed the rollback for recovery. Once created, rollbacks can be served directly to hosts, split to create a point-in-time clone, or the source volume can be reverted to the point in time selected, as shown in Figure 22. ESG Lab un-served the volume from the server, and reverted to the rollback. After serving the volume back to the server, the text files created with the batch file were confirmed to have been restored back to the 2:20:26 PM point in time of the rollback.

Why This Matters

Data growth and the rapid proliferation of virtualized applications are increasing the cost and complexity of storing, securing, and protecting business-critical information assets, and IT organizations running mission-critical applications need to guard against service interruptions. An interruption could be something unanticipated, such as a hardware failure or human error, but more often, routine equipment upgrades, firmware updates, and hardware refreshes can require equipment to be taken out of service. In highly virtualized, consolidated environments these disruptions will cause major outages with significant impact.

An “always available” storage solution with management tools that makes it easy to deploy and centrally manage a multi-site storage deployment reduces time, cost, and risk.

ESG Lab validated that DataCore SANsymphony and Hyper-converged Virtual SAN provide an array of data protection capabilities that can cost-effectively satisfy the most stringent business continuity and disaster recovery requirements. Synchronous mirroring within metropolitan areas, snapshots, CDP, and asynchronous remote replication to distant disaster recovery sites can be used without being dependent on any specific model or brand of storage device. For example, customers could take advantage of a hyper-converged system at a remote site to establish a contingency site for larger data centers while reducing the total cost of disaster avoidance.

ESG Lab Validation Highlights

- ☑ DataCore SANsymphony VVols integration made provisioning storage a seamless, integrated part of VM creation using predefined storage policies matching performance, availability, and locality of data needs without ever having to touch the back-end storage, using native VMware tools.
- ☑ ESG Lab was impressed with the speed, simplicity, and completeness of the DataCore integration.
- ☑ ESG Lab found that virtualizing infrastructure with DataCore Virtual SAN was intuitive and straightforward and DataCore Virtual SAN provided a highly available hyper-converged platform for multiple simulated applications.
- ☑ DataCore has demonstrated the power of its parallel I/O software by having the Storage Performance Council audit and publish an excellent result of 459,290 SPC-1 I/O requests per second at 100% load with an average response time of only 0.32 milliseconds and a record price-performance cost of \$0.08 per SPC-1 IOPS in a hyper-converged solution.⁷
- ☑ SANsymphony's automated storage tiering provided additional benefits, both improving overall IOPS by 700% and, by reducing load on lower performance tiers, improving SAS disk IOPS by 100%.
- ☑ DataCore Continuous Data Protection (CDP) was easy to configure and use, enabling rollback to a specific point in time without having to create multiple snapshots.

Issues to Consider

- ☑ The test results presented in this report are based on applications and benchmarks deployed in a controlled environment with industry-standard testing tools. Due to the many variables in each production data center environment, capacity planning and testing in your own environment are recommended.
- ☑ While ESG Lab feels that DataCore is an agile, mature organization with a robust and clearly differentiable solution, DataCore should keep a sharp eye on new hyper-converged and software-defined storage players that threaten to integrate similar functionality into their solutions.

⁷ http://www.storageperformance.org/results/benchmark_results_spc1_active#a00164

The Bigger Truth

Traditional storage systems were designed to support many physical servers and a relatively static infrastructure. In the virtual world, where application workloads are detached from physical systems and rapidly migrate across the environment, the traditional infrastructure complicates every aspect of the storage environment, and drives the urgency to virtualize storage. Virtualization is one of the few IT tools with a genuine ability to significantly address the challenge of unabated demand for a limited supply of the resources that deliver IT services and the tools to manage them.

Virtualization was a “nice-to-have” five or ten years ago, but is now rapidly becoming a hard requirement for many production environments. IT has no choice but to virtualize, as the experience and success of server virtualization has transitioned from initially supporting test and development to now being a software infrastructure for data center architectures. Storage virtualization *has* to be a part of that virtual IT infrastructure.

SANsymphony 10 is DataCore’s newest release based on more than a decade of experience, and DataCore has thousands of users that will attest to its ability to deliver both quality and business value. The company now finds itself with incredibly relevant capabilities that truly matter to users. ESG Lab tested DataCore SANsymphony and DataCore Hyper-converged Virtual SAN. We found the software easy to implement and manage, virtualizing any storage infrastructure with enterprise-class features and functionality while enhancing performance. DataCore SANsymphony kept data available and online through both planned and unplanned outages flawlessly.

ESG Lab was especially impressed with DataCore’s parallel I/O technology and its ability to deliver enterprise-class performance running on low-cost commodity hardware. DataCore’s recently published SPC-1 results set a record of \$0.08 per SPC-1 IOPS, besting its closest competitor by 300%. SANsymphony’s 459,290 IOPS rank it 11th overall in SPC-1 IOPS results, while the 0.32 millisecond response time at 100% load is the fastest response time ever reported on SPC-1 at the time of this writing. Even more impressive, the results were obtained using a single, hyper-converged system that combined all storage, compute, and network infrastructure in one server.

The quest for realistic and affordable options to deal with the challenges inherent in IT virtualization and consolidation is daunting. Storage administrators are increasingly being retired and replaced or re-invented as network and virtualization administrators; storage management in a modern IT environment has to be simple and practical as well as functional because users will eventually be compelled to virtualize everything.

DataCore SANsymphony and DataCore Hyper-converged Virtual SAN proved to be robust, flexible, and responsive and can deliver major value in terms of utilization, economics, improved response times, high availability (HA), and easy administration. ESG Lab firmly believes that it would benefit any organization considering or implementing an IT virtualization project to take a long look at DataCore software.

Appendix

Table 2. ESG Lab Test Bed

Data Infrastructure Software	Version
DataCore SANsymphony DataCore Hyper-converged Virtual SAN	R10
Hardware	OS & Hypervisor
2x Dell PowerEdge R720 Server 2.2 GHz Xeon Processor E5-2660 V2 384GB RAM, 2x 200GB SAS SSD (RAID-1) [Lab environment shared with other workloads]	Windows Server 2012 R2 with Hyper-V VMware vSphere 6.0
Other Test Software	Version
WANem (WAN emulator) 1.544 Mbps, 10 ms Latency	2.2
Workload Generation	Workload Characterizations
Iometer 2006.07.27	SQL Server 8KB I/O 100% random, 67% read



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