

Lab Validation Report

Microsoft Windows Server 2012

Storage and Networking Analysis

By Mike Leone, ESG Lab Engineer

December 2012

Contents

Introduction	3
Background	3
Introducing Windows Server 2012	4
ESG Lab Validation	5
Performance with Storage Spaces and the SMB Protocol	5
Storage Efficiency with Deduplication	11
Volume Agility with Chkdsk	15
Data Transfer Speed Efficiency with ODX	17
ESG Lab Validation Highlights	20
Issues to Consider	20
The Bigger Truth	21
Appendix	22

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 Microsoft.

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

This ESG Lab Validation documents the results of expert, independent, hands-on testing of the improved storage and networking performance, efficiency, and scalability of Windows Server 2012. Testing included a look at Storage Spaces and the SMB 3.0 protocol with direct-attached storage to highlight cost-efficiency with industry-standard hardware. Also tested were improvements to Chkdsk performance and functionality, new deduplication capabilities, and the new ODX feature, all of which highlight the flexibility, efficiency, and agility that Windows Server 2012 has to offer organizations of any size.

Background

IT managers within organizations of all sizes are struggling with the cost, complexity, and performance challenges associated with managing growing volumes of digital data. These challenges are particularly vexing for IT managers within midmarket organizations (100-999 employees) due to a shortage of manpower and dedicated storage expertise compared to their peers within enterprise-class organizations (1,000 employees or more).

More than half of IT managers within midmarket companies surveyed by ESG (60%) reported managing 25TB or more of data¹ and they're struggling to keep up with data growth. Nearly half of the respondents to a separate ESG research survey (46%) reported that data is growing at an annual rate of 20% or more.² For those organizations, data volume is doubling every two to five years.

Faced with these challenges, it's no surprise that the top storage challenges reported by IT managers would include rapid growth and management of unstructured data, hardware costs (CAPEX), and poor performance, as shown in Figure 1.³ It should also be noted that the need to support growing virtual server environments and the management, automation, and placement of data made the list of top ten challenges as well.

Figure 1. Top Ten Storage Environment Challenges



Source: Enterprise Strategy Group, 2012.

¹ Source: ESG Research Brief, [Disk-based Storage Capacity Trends](#), September 2012.

² Source: ESG Research Report, [Trends in Data Protection Modernization](#), August 2012.

³ Source: ESG Research Report, *2012 Storage Market Survey*, due to be published October 2012.

Introducing Windows Server 2012

Windows Server 2012 provides powerful new technologies that help transform an existing IT infrastructure into a modern datacenter that effectively addresses every available opportunity with agility and efficiency. New and improved features provide high performance and scalability that create a dynamic, multitenant infrastructure which securely scales workloads to meet even the most complex SLAs, while offering the flexibility to adapt to ever-changing business needs. With new enterprise-class storage features and functions delivered on cost-effective, industry standard hardware, Customers can expect capabilities such as:

- *Storage Spaces* - Provide virtualized storage pools and virtual hard disks that are configured with specific provisioning and allocation attributes, and can leverage failover clustering for high availability and resiliency.
- *Enhanced SMB protocol* - Enables file server support for applications needing high performance and availability. SMB Direct provides RDMA hardware support for high performance storage capabilities without high cost. SMB Multichannel facilitates network bandwidth aggregation and fault tolerance for multiple paths between SMB clients and file servers.
- *Data deduplication* – State-of-the-art data chunking and compression for individual servers as well as branch office extension.
- *Thin provisioning and trim* – Native support for thin provisioning of storage plus reclamation when storage is no longer needed.

These new storage solutions deliver continuous application uptime while maximizing performance, scalability, and choice with industry leading partners through benefits such as:

- *Resilient File System (ReFS)* – New file system designed to maximize data availability and online operations.
- *NTFS* - Improvements in data integrity and availability.
- *Cluster Shared Volume V2 (CSV2)* – Shared file storage solution for consolidated storage and application clusters in non-Hyper-V environments. Storage space integration with failover clustering also provides continuously available service deployments.
- *Offloaded Data Transfer (ODX)* – Offloads data movement from the server, enabling fast and easy migration of large data sets or entire VMs between storage devices without impacting server or network activities and maximizing returns on your storage array investments.
- *Live storage migration* – Supporting high availability while moving virtual hard disks between hosts.
- *Support for VMware and NFS 4.1* – Enabling VMware virtual machines and other third-party platforms to run with NFS while gaining Windows continuous availability.

These technologies enable enterprise IT to simplify services and management, while quickly supporting process and application deployment, scaling, and integration. Along with improved availability and access to applications, Windows Server 2012 simultaneously reduces costs, capital investment, and risk.

ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of Windows Server 2012 at Microsoft's EEC facility in Redmond, Washington. Testing was designed to demonstrate the performance, efficiency, and scalability of Server 2012 using industry standard tools and methodologies. This includes the testing of new Chkdsk, deduplication, and ODX functionality, as well as the utilization of storage spaces and the SMB 3.0 protocol for a complete, cost-effective solution.

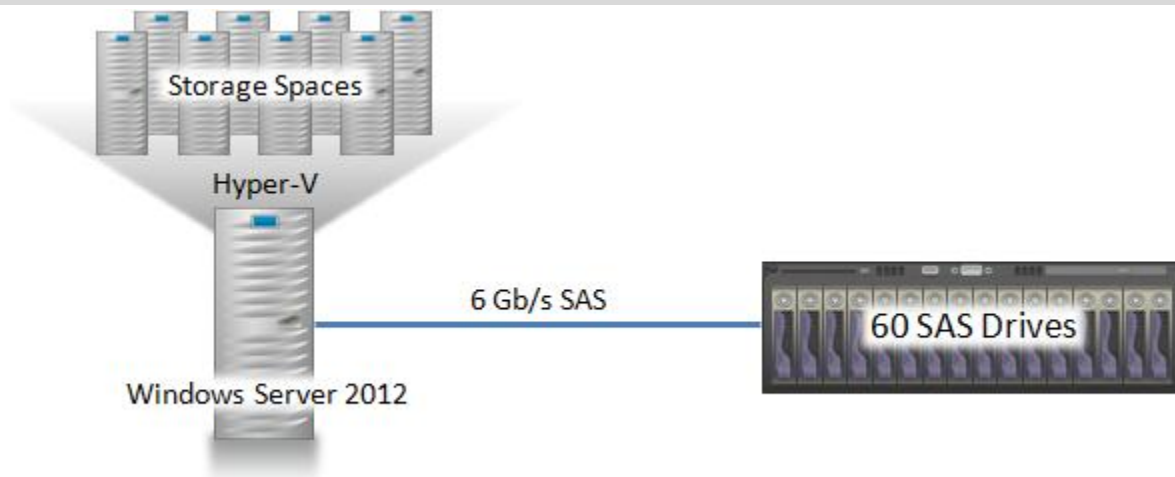
Performance with Storage Spaces and the SMB Protocol

The Storage Spaces feature of Windows Server 2012 makes enterprise-level features such as scalability and high-availability affordable by virtualizing commodity hardware. Organizations can consolidate physical and virtual storage capacity into pools that are easy to scale, and deliver high performance and availability. For private clouds and service providers, it enables a service catalog of virtual disks that can be created and managed securely in multi-tenant configurations. With enhancements to the storage message block protocol, Server 2012 brings enterprise levels of performance as well as bandwidth aggregation and fault tolerance. Organizations can gain enterprise-class functionality while saving money and reducing complexity with easy-to-manage file shares. Solving this key challenge can enable expanded virtualization for mission-critical applications. With new remote direct memory access (RDMA) support, one computer can now access another computer's memory without operating system involvement, leading to extremely efficient low-latency networking with very high throughput potential.

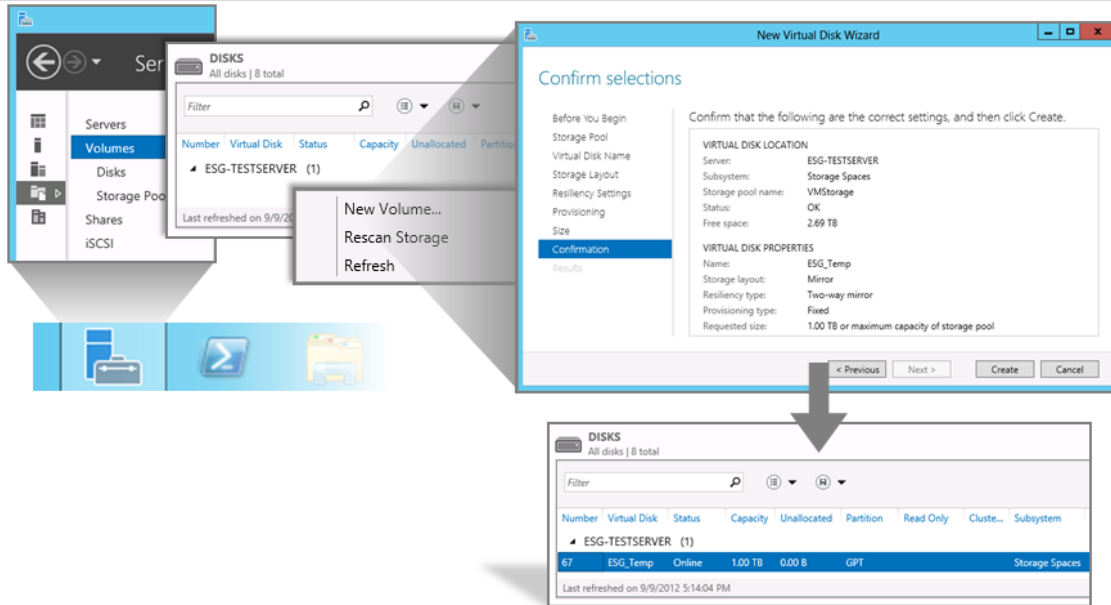
ESG Lab Testing

ESG Lab tested Storage Spaces with a goal of showing the ease of configuration, performance scalability, and cost-effectiveness. Figure 2 shows the hardware configuration used during testing. ESG Lab used a Dell R910 as the host machine running Windows Server 2012. Eight Hyper-V enabled guest VMs also running Server 2012 were configured with four vCPUs and 12GB of RAM. The host machine was connected via 6Gb SAS to a RAID Inc. JBOD with 60 600GB 15K SAS drives.

Figure 2. ESG Lab Test Bed with Storage Spaces



ESG Lab configured the storage using the new, easy-to-follow wizards offered through Server 2012. The process is highlighted in Figure 3. ESG Lab selected the Server Manager in the taskbar, browsed to the Volumes section, and created a new volume. The wizard asked for information regarding which storage pool, the desired name of the volume, storage layout type, resiliency settings, provisioning type, and size. After selecting everything, a confirmation screen is displayed before the volume is created. For the testing outlined in this paper, ESG Lab configured the storage layout to be mirrored (as opposed to simple or parity), the resiliency setting to be two-way mirrored (as opposed to three-way mirrored), and the provisioning to be fixed (as opposed to thin). The volume was then created and properly assigned to a specific VM.

Figure 3. Configuring Virtual Volumes with Storage Spaces

ESG Lab used an OLTP online brokerage application to simulate the activity of thousands of Microsoft SQL 2012 users with a goal of demonstrating the performance and scalability of storage spaces. Each of the eight virtual machines was assigned a volume consisting of a 2TB database and a 90GB log. The workload used during ESG Lab testing was designed to emulate the database activity of users in a typical online brokerage firm as they generated trades, performed account inquiries, and did market research. The workload was composed of ten transaction types with a defined ratio of execution. Four of the transactions performed database updates, the rest were read-only. A 3,000 customer database was configured within each virtual machine with a database scaling goal of 24,000 customers. The workload generated a high level of I/O activity with small access sizes and spent a lot of execution time at the operating system kernel level. These characteristics, combined with a large, cache-resident working set, created a workload that was well suited for evaluating the performance efficiency and scalability of storage spaces.

The average response time for the ten transaction types was monitored as the scale of the databases and number of concurrent users increased within the virtual machines. Perfmon was also used to monitor the number of disk transfers/sec (IOPS) and SQL batch requests/sec. The results are summarized in Figure 4 and Table 1.


Figure 4. SQL Server Workload Scalability with Storage Spaces

Table 1. SQL Server Workload Scalability with Storage Spaces

Virtual Machines	Database Scale (Total Customers)	SQL Server Batch Requests/Sec (Total)	Transactions/Sec (Total)	Average Transaction Response (Seconds)
1	3,000	937	500	.01
2	6,000	1,834	1,192	.01
4	12,000	3,569	1,867	.02
6	18,000	5,869	2,983	.02
8	24,000	7,161	3,964	.03

What the Numbers Mean

- SQL Server batch requests increased from 937 to 7,161 as the number of virtual machines running the OLTP workload increased from one to eight. To put this into perspective, Microsoft documentation indicates that “over 1,000 batch requests per second indicate a very busy SQL Server.”⁴ Of course, this is a relative number that depends on the power of the hardware used to deploy SQL Server applications.
- Each VM averaged just fewer than 1,000 SQL Server batch requests per second in each test scenario.
- Performance was predictable and scaled linearly, as measured in increments of two VMs.
- Response times were measured at the database transaction level for a workload composed of a mix of read-only and read-write transactions. Ten transaction types generated between two and 17 table accesses each, with each table access requiring one or more disk accesses. The average transaction response times remained low, falling between .01 and .03 seconds.

ESG Lab audited the results of a Microsoft Windows Server 2012 launch demo as another proof point of the enterprise-class performance and scalability of Storage Spaces. Microsoft configured three Intel Romley based servers in a cluster, each with two Xeon E5-2670 processors, 48GB of RAM, and four LSI 9205-8e SAS HBAs. The servers were connected to a pair of 24-bay JBODs with 12 6Gb/sec 4-lane SAS cables. With each of the JBODs containing 24 SSDs, a three Server 2012 cluster sustained 2.3 million IOPS with a 4KB random read workload.

Why This Matters


User expectations of enterprise data centers and hosted service providers are running higher than ever. Today’s infrastructure managers must deliver high availability of applications and data, and be able to scale quickly and easily while maintaining high performance and meeting strict budget requirements. These features are generally associated with more complex, high-end storage implementations, and with costs on the mind of most organizations, it can be difficult to afford to swap out their current storage to gain these features. For service providers, profitability depends on finding a way to deliver these features while keeping costs down.

ESG Lab confirmed that the performance efficiency and scalability of Storage Spaces can be used to consolidate tier-1 application workloads with confidence. Provisioning of storage was made simple with intuitive configuration wizards to help meet storage resiliency and performance expectations. A single server hosting a virtualized SQL Server 2012 infrastructure deployed within eight Hyper-V virtual machines supported the tier-1 application needs of database instances serving up to 24,000 simulated online brokerage customers. Response times remained manageably low, while SQL Server batch requests and disk transfers scaled linearly as virtual machines and simulated brokerage users were added. ESG Lab was most impressed with the fact that these results were achieved with a cost-effective industry standard server and a direct-attached disk array that eliminated the cost of SAN adapters, switches, and storage systems.

⁴ Source: SQL Team Forum, [Memory Bottleneck Analysis](#), April 2006.

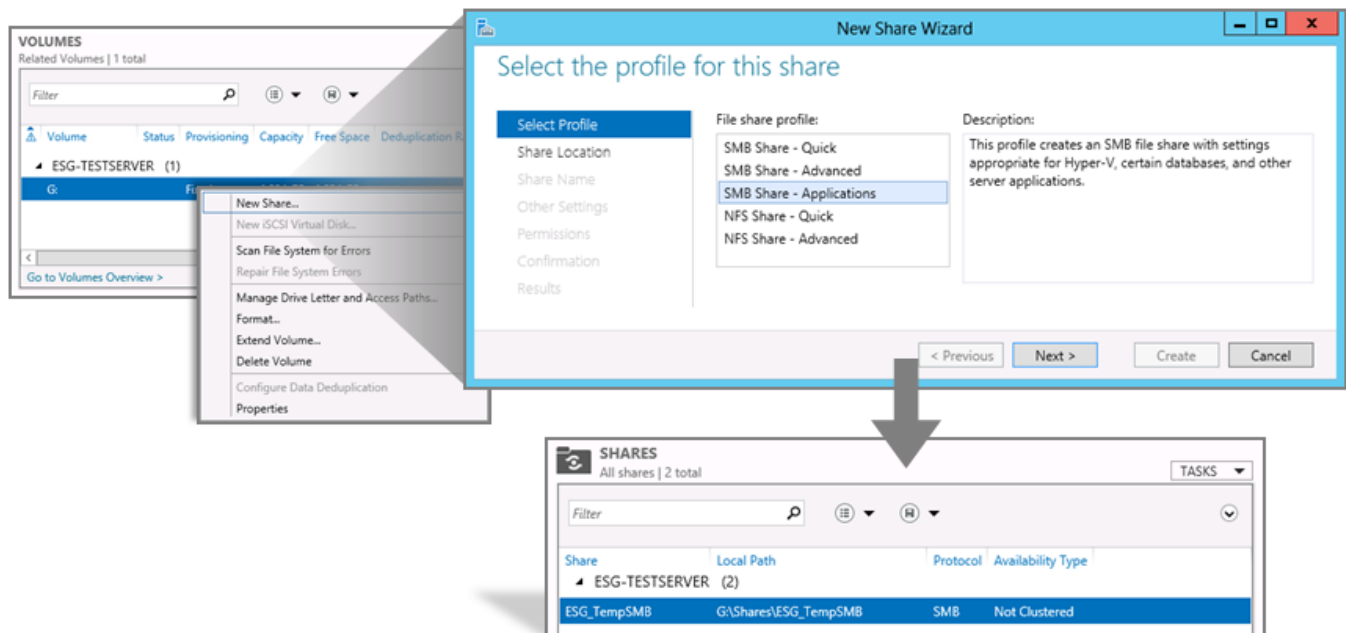
ESG Lab tested the improved SMB 3.0 protocol with a goal of showing the ease of configuration, performance scalability, efficiency, and cost-effectiveness. Figure 5 shows the hardware configuration used during testing. ESG Lab used a Dell R910 as the host machine running Windows Server 2012. The host machine was connected to a file server also running Server 2012 with a 10GbE LAN. The file server was connected via 6Gb SAS to a RAID Inc. JBOD with 60 600GB 15K SAS drives. Eight Hyper-V enabled VMs running Server 2012 were configured with four vCPUs and 12GB of RAM and presented to the main server under test through the SMB 3.0 protocol. ESG Lab also configured RDMA enabled Chelsio network adapters between the main server under test and the file server.

Figure 5. ESG Lab Test Bed with the SMB Protocol



Configuring SMB shares was made easy through another user-friendly wizard. Figure 6 shows the SMB share setup process. By selecting New Share from the Volumes section of Server Manager, the wizard guided ESG Lab through the setup. For this particular test case, ESG Lab selected the SMB share offered specifically for a database application workload. After providing information in regards to the share location, name, and access permissions, a new share was created and presented to the main test server.

Figure 6. Configuring SMB Shares



ESG Lab used the same OLTP workload application to simulate the activity of Microsoft SQL Server users with a goal of demonstrating the performance, scalability, and efficiency of the SMB protocol, the Microsoft Hyper-V hypervisor, and the Microsoft SQL Server database engine on cost-effective commodity hardware. A 3,000 customer database was configured within each of the eight SQL Server virtual machines with a goal of achieving linear scalability for the number of transactions per second as the number of consolidated SQL Server VMs increased from one to eight. The transactions per second and average response time were monitored as the number of customers and VMs increased. The results are summarized in Figure 7 and Table 2.

Figure 7. SQL Server Workload Scalability with the SMB 3.0 Protocol




Table 2. SQL Server Workload Scalability with the SMB 3.0 Protocol

Virtual Machines	Database Scale (Total Customers)	SQL Server Batch Requests/Sec (Total)	Transactions/Sec (Total)	Average Transaction Response (Seconds)
1	3,000	892	499	0.01
2	6,000	1,798	1,123	0.01
4	12,000	3,454	1,989	0.02
6	18,000	5,783	3,027	0.03
8	24,000	6,903	3,885	0.05

What the Numbers Mean

- Each of the eight virtual machines was populated with a SQL Server database supporting 3,000 brokerage customers.
- The virtual machines were configured with four virtual CPU cores and 12GB of RAM with a goal of balancing the utilization of all the available physical server resources during the peak eight-VM test.
- The total number of SQL Server batch requests across all eight VMs increased from 892 to 6,903.
- Transactions per second increased from 499 to 3,885 as the number of virtual machines running on a single physical server increased from one to eight.
- ESG Lab recorded manageably low average transaction response times between 0.01 and 0.05 seconds, with a goal of avoiding end-user-perceived slow response times for transactions taking longer than three seconds.
- Network utilization remained remarkably low, with utilization reaching a maximum of 4.2%.

ESG Lab noted that for both the Storage Spaces test scenario and the SMB test scenario, the same hardware was used including: server type, storage type, VM configuration, storage configuration, and OLTP workload. This was done with a goal of showing a high-level comparison between the two test cases. Table 3 shows a comparison between the total number of IOPS for Storage Spaces (local) and SMB (remote). ESG Lab witnessed impressive comparison results, which show the remote SMB IOPS to be between 94-98% of the local Storage Spaces IOPS. Microsoft guidance suggests that customers could see an even greater benefit from the SMB protocol and in some cases exceed Storage Spaces performance. This adds to the already cost-effective, ease-of-use value of a file-server hosting multiple VMs that serve various applications.

Table 3. SQL Server Workload Scalability with the SMB 3.0 Protocol

Virtual Machines	Storage Spaces Local IOPS (Total)	SMB Remote IOPS (Total)	Remote/Local
1	900	850	94.4%
2	1,750	1,700	97.1%
4	3,500	3,350	95.7%
6	5,850	5,600	95.7%
8	7,000	6,850	97.9%

ESG Lab audited the results of a Microsoft benchmark with a goal of quantifying the additional performance benefits that can be gained by leveraging the combination of RDMA support and the enhanced SMB 3.0 protocol in Windows Server 2012. The hardware used for the testing was two Intel Romley based servers, a single port from a Mellanox network interface, InfiniBand FDR with 54Gbps, and persistent storage with four FusionIO ioDrive 2 cards. Two workloads were run to show IOPS and throughput. An 8KB read workload achieved a high rate of 343,388 IOPS, while a 512KB read workload sustained an impressive throughput rate of 5.792Gb/sec.

Why This Matters

With less tolerance for poor application performance, it's no surprise that many organizations have resisted moving tier-1 applications to virtual servers for fear that workload aggregation will slow performance. Prior to Windows Server 2012, remote storage options for Hyper-V were limited to costly, hard to provision SAN solutions for Hyper-V guests or featureless, less expensive storage options. The enhanced performance and efficiency of the SMB 3.0 protocol in Windows Server 2012 allows organizations to save money while getting great performance on industry standard hardware.

ESG Lab validated that the improvements to the SMB protocol allow organizations to deploy tier-1 applications like SQL Server 2012 without a need for concern. Deploying an extremely cost-effective solution that leverages the new benefits of the SMB 3.0 protocol, consolidation through Hyper-V, and easy provisioning of direct-attached storage with 15K SAS drives, more than adequate performance was achieved from an enterprise SQL Server 2012 application workload. The number of transactions/sec scaled linearly from 499 to 3,885 while response times remained low as the number of VMs scaled from one to eight. Finally, very little bandwidth was consumed, giving more freedom for network configurations within IT infrastructures. With the added benefits of hardware offload with SMB direct and multi channel, organizations can achieve similar returns as if using MPIO in a SAN, while transparent failovers provide the much needed high availability of many enterprise-class storage deployments.

Storage Efficiency with Deduplication

Windows Server 2012 now includes data deduplication as a standard feature at no extra cost. This new feature is a highly scalable and efficient solution that allows businesses to store more data in less physical space. With the added capabilities of capacity optimization, non-intrusive performance scalability, and reliability from the perspective of data integrity, Server 2012 allows users to store and access data as efficiently as possible. Deduplication is included with Windows Server 2012 at no extra cost, and is scalable and non-intrusive. Using a combination of chunking and compression technologies, Windows Server 2012 applies state-of-the-art deduplication to each server, freeing up storage capacity.

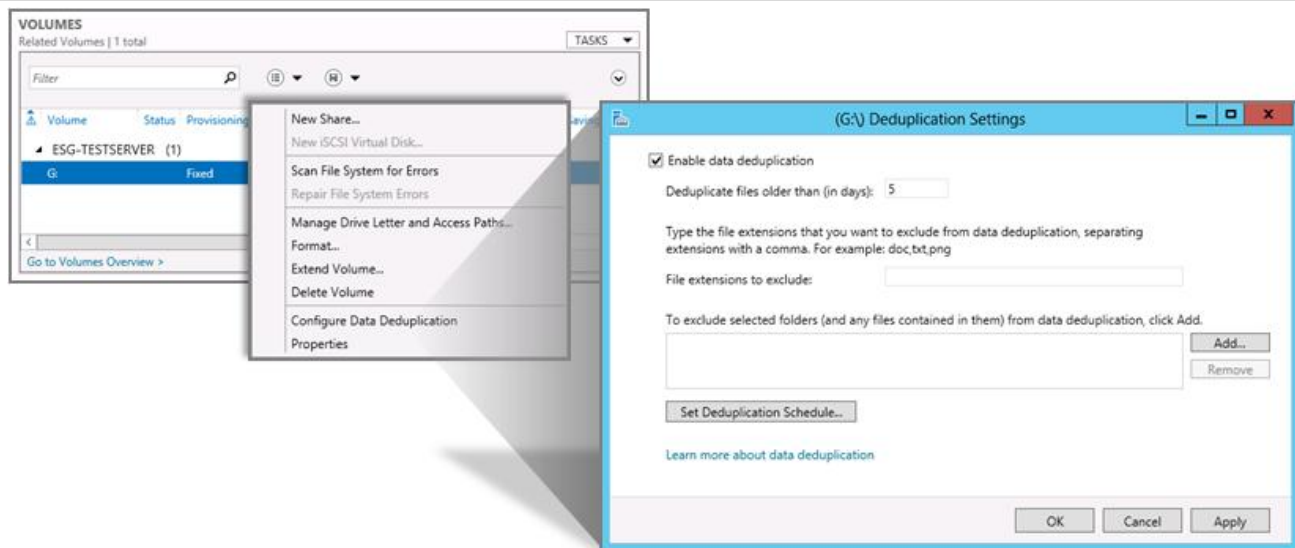
ESG Lab Testing

ESG Lab tested the new Server 2012 deduplication capabilities with a focus on capacity savings and potential deduplication overhead. Two different data sets were tested:

- **File Server / Home Directory** – The volume size was 900GB and consisted of mostly small files such as Microsoft Office documents and PDFs. This data set is considered to be moderately deduplicatable due to the size and variability of file content.
- **Operating System Images** – The volume size was 3TB and consisted of mostly large files such as pre-configured VHDs containing Server 2008, Server 2012, and SQL Server 2012. This data set is considered to be highly deduplicatable due to the combination of size and the amount of common data.

ESG Lab's first objective was to understand the ease of configuring deduplication on a particular volume. Setting up deduplication was made simple through a single window that was accessed from Server Manager. The short process is shown in Figure 8. After installing the deduplication role on Server 2012, ESG Lab selected the volume in which to deduplicate and accessed the Deduplication Settings window. The window allowed for either a quick and easy setup, or offered the ability to create a deduplication schedule for more complex storage environments.


Figure 8. Configuring Deduplication



Next, ESG Lab used Powershell commands to initiate deduplication jobs on the two volumes specified above.⁵ The command was issued to one volume at a time. There are two deduplication modes: throughput mode and background mode. Throughput mode is meant to consume more server resources (CPU and memory) and complete faster. Background mode is meant to consume less server resources and complete slower. With a goal of showing capacity savings, ESG Lab used throughput mode to quickly deduplicate each volume. The capacity savings are shown in Figure 9.

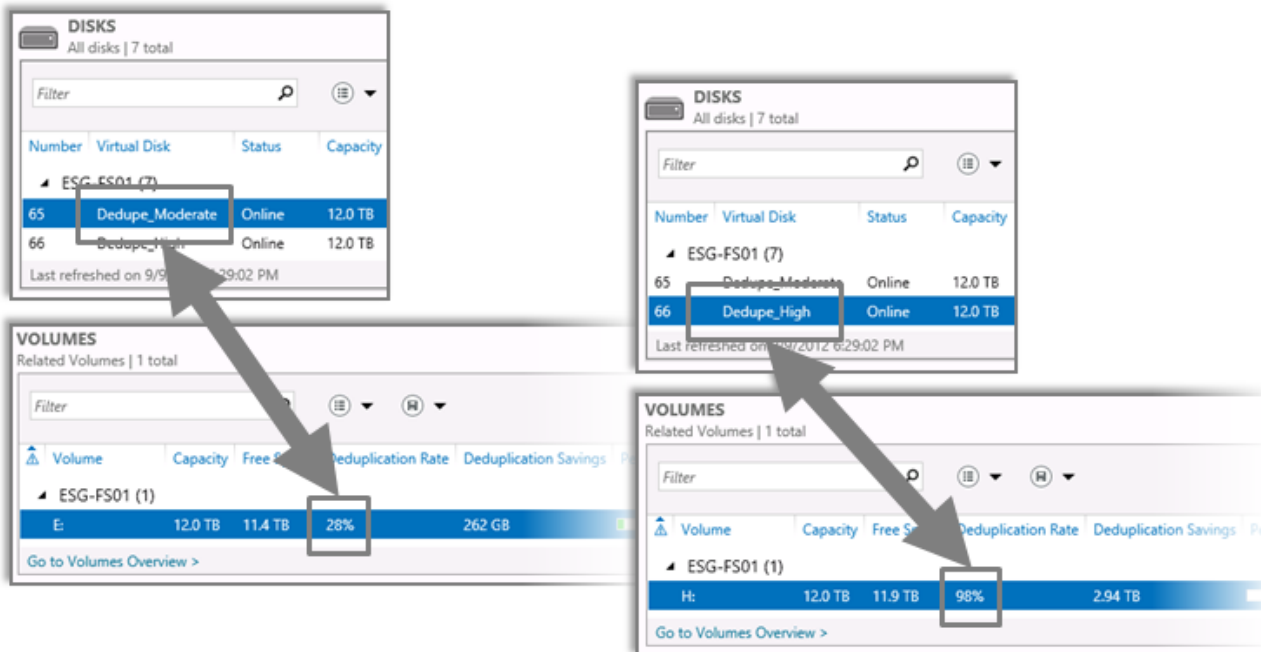
⁵ Listed in appendix.

Figure 9. Server 2012 Deduplication Capacity Savings




The File Server volume decreased in consumed capacity from 900GB to 648GB, yielding a savings of 28%. Microsoft guidelines suggest a typical user could see between 25-60% capacity savings with similar datasets. The volume consisting of OS VHDs decreased in consumed capacity from 3TB to 57GB, yielding a savings of 98%. ESG Lab viewed the deduplication rate through Server Manager, as shown in Figure 10.

Figure 10. Deduplication Rate through Server Manager



Impressed with the capacity savings, ESG Lab next tested the overhead of rehydrating files in two different ways: opening small files and copying large files. Using a Powershell script, ESG Lab selected four different file counts consisting of Microsoft Office files. Each test opened the selected number of files from a volume where the files were not deduplicated and then from a volume where they were deduplicated. Processor utilization and elapsed time were monitored and the results are shown in Figure 11 and Table 4.


Figure 11. Deduplication Overhead – Opening Multiple, Small Files*Table 4. Deduplication Overhead – Opening Multiple, Small Files*

Number of Files	Total Time to Open All Files Before Deduplication (Seconds)	Total Time to Open All Files After Deduplication (Seconds)
1	1.312	1.321
22	7.171	7.702
83	27.765	28.171
125	35.638	36.799

What the Numbers Mean

- In all cases, the time taken to open a group of files after deduplication took no longer than 3% compared to the same files before they were deduplicated.
- Users should see little to no overhead when opening smaller, deduplicated Office files.
- Processor utilization remained relatively low while rehydration of the office files occurred, serving as another layer of transparency to users concerned about the impact of rehydrating to other server workloads.

The second way ESG Lab tested the potential overhead of rehydrating deduplicated files was by copying large VHDs from one volume to another volume by using the xcopy command. A scenario in which this would fit would be having a library of VHDs and copying one of them to provision a new VM on a Hyper-V host. Reasonable rehydration rates with minimal resource usage are crucial to any business looking to benefit from Server 2012's freely available deduplication technology. ESG Lab monitored CPU utilization and elapsed time to characterize the overhead of rehydrating larger, more fragmented files. The results are shown in Figure 12 and Table 5.

Figure 12. Deduplication Overhead – Copying Large Files*Table 5. Deduplication Overhead – Opening Multiple, Small Files*

Number of VHDs	Consumed Capacity (GBs)	Total Time to Copy VHD Before Deduplication (Seconds)	Total Time to Copy VHD After Deduplication (Seconds)
1	10.7	19.5	25.3
1	16.5	29.8	39.7
6	81.8	164	209.9

What the Numbers Mean

- ESG Lab copied VHD files contained within the volume that was deduplicated at a rate of 98%.
- In all cases, the time taken to copy deduplicated VHDs from one volume to another volume took around 30% longer than the VHD(s) prior to deduplication. This was true regardless of the number of VHDs.
- Due to the fragmentation of the highly deduplicated VHD files, CPU utilization increased modestly by 20% when rehydrating those files.

Why This Matters

Data growth continues to challenge IT organizations. ESG research proves this point; respondents to ESG's annual IT spending surveys for the past three years have listed managing data growth as one of the top five IT priorities. The costs of storing and managing duplicate file data can stress capital and operational budgets unnecessarily. Deduplication appliances are available, but they add another point of management to the IT administrator's task list, another line item to the purchase order, and additional CAPEX and OPEX costs.

ESG Lab validated the deduplication capabilities of Server 2012 and feel that the combination of capacity savings and rehydration rates offer a great option to any size business looking to consolidate data in a single cost-effective solution. Deduplication was easy to configure and monitor through Server 2012's Server Manager interface. ESG Lab witnessed capacity savings of up to 98% with a highly deduplicatable volume containing OS images. Combined with the capacity savings offered by Server 2012, the transparency of rehydration to users was impressive. For smaller, deduplicated files, the overhead to rehydrate when opening an Office file went unnoticed. For larger files, rehydration took about 30% longer, but with a 98% deduplication rate, ESG Lab feels that 30% overhead is more than reasonable.

Volume Agility with Chkdsk

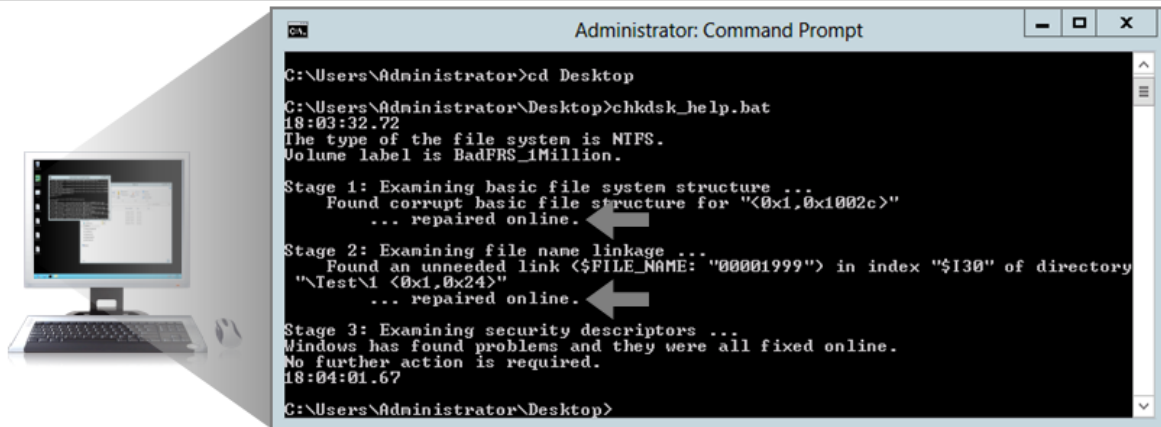
Chkdsk is a command used to check a disk volume's file system and metadata for logical or physical errors. Windows Server 2012 brings additional self-healing capabilities and better scanning techniques to any sized IT environment. These enable Chkdsk to find and repair errors, even for very large volumes, in seconds to prevent file system corruption. Organizations can remain online and productive, and IT no longer must manage volume size around file repair downtime.

ESG Lab Testing

ESG Lab tested the improvements to Chkdsk, with a goal of showing Server 2012's ability to fix and repair volumes quickly and efficiently when compared to Server 2008. Using a Hyper-V enabled server and direct-attached storage, ESG Lab created a 2012 and 2008 VM with four vCPUs and 16GB of RAM. Each VM was presented with nine virtual volumes with varying degrees of size and complexity. The three sizes consisted of volumes with 100,000 files, 1,000,000 files, and 10,000,000 files. The three types of corruptions tested were a bad file record segment, 1,000,000 orphan files, and a corrupt index entry.

The errors were corrected using two methods. The first method used a Chkdsk scan to detect and fix simple errors. Figure 13 shows a command prompt with the output of a Server 2012 Chkdsk scan command. The command repaired a volume consisting of 1,000,000 files that contained a bad file record segment corruption. The ability of Server 2012 to keep the volume online during the repair is highlighted with arrows.


Figure 13. Server 2012 Chkdsk Volume Availability



The second repair method ESG Lab tested was the new functionality of Chkdsk called spotfixing. Spotfixing allows volumes to be scanned while remaining online and then gives the user the ability to take the volume offline to fix the errors when time permits. ESG Lab tested the efficiency of spotfixing with the 1,000,000 orphan file corruptions. The volume repair in both Server 2008 and Server 2012 were monitored with a focus on duration of the repair to complete, as well as volume availability. The spotfixing results are shown in Table 6 and Figure 14.

Table 6. Server 2012 Chkdsk Spotfixing Results

Operating System	Drive Volume Status	File Count	Time to Repair
Server 2008	Offline for Scan and Repair	100,000	Over 24 hours
		1,000,000	Over 24 hours
		10,000,000	Over 24 hours
Server 2012	Online for Scan, Offline for Repair	100,000	10 minutes 11 seconds
		1,000,000	12 minutes 6 seconds
		10,000,000	16 minutes 48 seconds

Figure 14. Server 2012 Chkdsk Spotfixing Results

What the Numbers Mean

- ESG Lab repaired a complex corruption (1,000,000 orphan files) through Server 2012's new spotfixing functionality and compared results to Server 2008.
- Using Server 2012 the volume remained online during the scan and only went offline to repair the error. Server 2008 took the volume offline for the entire duration of the process.
- The time to repair a complex error was significantly shorter using spotfixing. In Server 2008 a repair could take upwards of multiple days, as opposed to Server 2012, which took just under 17 minutes to scan and repair a volume consisting of 10,000,000 files.

Why This Matters

Scanning file systems and repairing errors is a key feature of an operating system. Unfortunately, the process can result in hours of downtime for large volumes, even if only a few errors are present, because the downtime required is proportional to the size of the total volume, not the number of errors. As a result, business operations are interrupted, and this can cause a productivity drain across your organization. Many of today's stringent corporate SLAs demand application and availability around the clock to accommodate global enterprises. And for service providers, high availability of customer operations is critical to business success. IT often manipulates volume size to combat this problem, creating smaller volumes than they prefer just to minimize Chkdsk downtime.

ESG Lab confirmed that the new efficiency and functionality of Server 2012's Chkdsk is a major improvement over its Server 2008 predecessor. The ability to quickly repair basic corruptions while volumes remained entirely online was impressive to ESG Lab. Also, the time taken to repair a 1,000,000 orphan file complex corruption using the new feature of spotfixing was extremely faster. Spotfixing reduced the repair time from days to minutes and with partial drive availability throughout the process, productivity from a user's perspective would be only minimally impacted.


Data Transfer Speed Efficiency with ODX

ODX is a new data transfer technology from Microsoft that advances the movement of storage data from within the storage itself as opposed to over a network. Using storage resources for data movement relieves server and network resources from a heavy burden and reduces latency. Whether initiated from the GUI or command line, Server 2012 file transfers, provisioning, and virtual machine migrations are quickly, automatically, and transparently offloaded to the storage arrays.

ESG Lab Testing

The process for which data is transferred using new Server 2012 ODX capabilities directly correlates to the capacity of the the ODX-capable storage solution. This differs from the traditional form of data transfers, which can potentially consume significant network and server resources. ODX transfers improve data center capacity and scale by removing host-based data transfer load from servers and networks. Figure 15 shows the data path for a Non-ODX and ODX data transfer for a VM live migration.

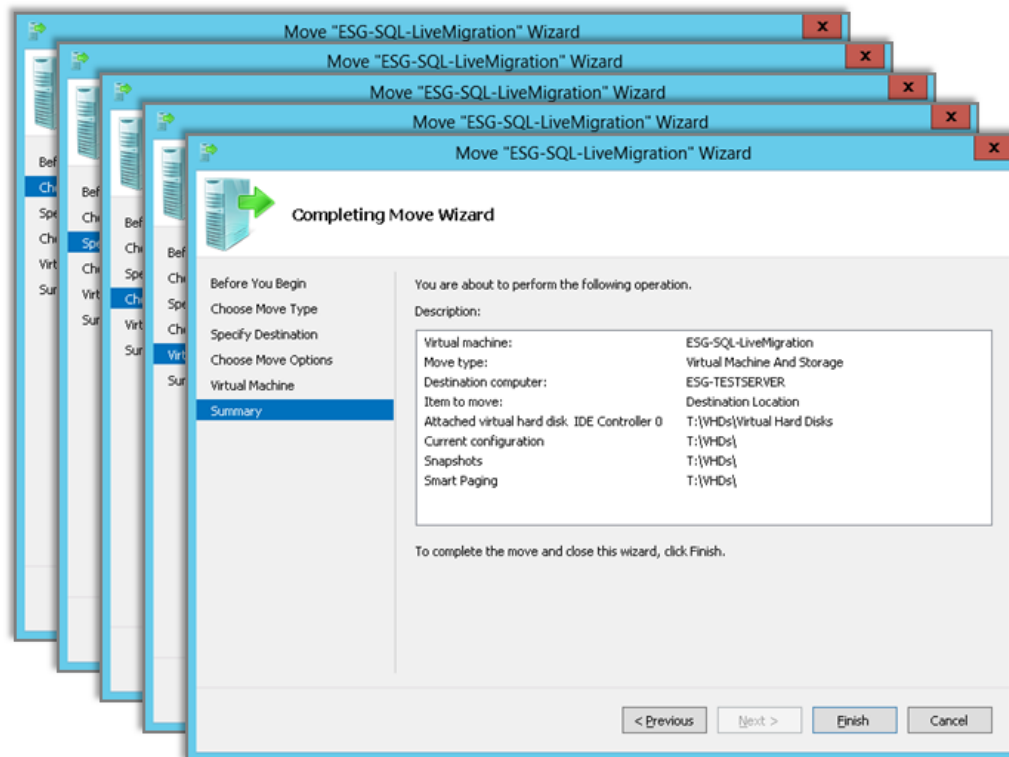
Figure 15. Data Movement with Non-ODX and ODX Data Transfers



The red dotted line is the data path for a Non-ODX data transfer. The transfer begins at the storage level, travels through the host server consuming host-server resources, over a LAN consuming network resources, through the destination-server, consuming destination-server resources, and finally lands back in the storage. The green arrow is the data path for an ODX data transfer. Information is shared across the LAN to initiate the data transfer, but the information is a single, small-sized token that represents the data about to be shared from the storage device. The storage device then performs the data movement internally according to the token. Though the above scenario shows what a VM live migration would look like, this scenario is true for all large data transfers, including database and video files.


ESG Lab tested the efficiency and functionality of an ODX data transfer. Two servers were connected to an ODX-compliant Dell EqualLogic storage array. The storage array consisted of 12 600GB SAS drives. A single RAID5 pool was created with two volumes, one of which contained a 75GB VM, and the other was empty. Using an intuitive wizard, ESG Lab configured a VM live migration from one server to another within a SAN. The wizard is shown in Figure 16. ESG Lab specified the type of move, which server was going to receive the data, move options, and destination virtual machine options. Below, the summary screen highlights all options selected by the Lab.

Figure 16. ODX Data Transfer Wizard



ESG Lab transferred a VM using the traditional non-ODX method and the new ODX method. The Lab monitored network utilization and elapsed time for the transfer to complete in both test cases. The results are shown in Figure 17.

Figure 17. ODX Data Transfer Results



What the Numbers Mean

- ESG Lab witnessed striking improvements using the ODX data transfer.
- The ODX transfer took approximately six and half minutes for the VM to completely migrate to the other server and the average network bandwidth consumption was around 64Kb/sec.
- Using the Non-ODX method, moving the 75GB VM over the network took approximately 52 minutes and consumed 4Mb/sec of network bandwidth.
- The ODX method completed eight times faster than the non-ODX method while consuming virtually no server CPU or network resources.

Why This Matters

The demand for high-speed data transfers has never been higher. With virtual deployments being designed to consolidate IT infrastructure resources and increase agility with VM migrations, a push for system and storage optimization is on the rise. Organizations continue to bog down network bandwidth by transferring large files and VM images through client-server networks. A need for a more efficient data transfer mechanism is at the top of the list; one that consumes less resources and reacts quicker to the needs of the business. ESG Lab validated that ODX technology, which was introduced in Windows Server 2012, can be used to move large files and virtual machines quickly and directly between storage arrays while completely removing server and network resource consumption. ESG Lab specifically noted the transparency of the ODX file transfer, whether performed through a basic drag-and-drop from one volume to another, the command line, or easy-to-use Server 2012 wizards. The time it took ESG Lab to transfer a 75GB Hyper-V VM using the new ODX capabilities of Server 2012 was remarkable, completing eight times faster than the traditional non-ODX method.

ESG Lab Validation Highlights

- ☑ Configuring and managing storage spaces and SMB shares was made easy through Server Manager's intuitive interfaces and helpful configuration wizards.
- ☑ Using cost-effective, industry standard direct-attached storage, Hyper-V VMs, and SQL Server 2012, the performance of Storage Spaces and SMB shares exceeded ESG Lab's expectations. Scaling from one to eight VMs, both test cases yielded predictable and linearly scalable OLTP performance in regards to transactions/sec, SQL Server batch requests, disk transfers/sec, and average transaction response times.
- ☑ The capacity savings witnessed by ESG Lab were exceptional, with a highly deduplicatable volume of OS VHDs going from 3TB to 56GB, saving 98% of the original consumed capacity.
- ☑ The overhead of rehydrating deduplicated files was more than acceptable, especially when combined with the potential capacity savings. When opening a collection of 125 Microsoft Office documents before and after deduplication, a barely noticeable 3% difference was measured.
- ☑ Server 2012 Chkdsk improvements allowed volumes to remain online while scanning and correcting basic drive errors.
- ☑ The new spotfixing Chkdsk feature reduced the amount of time needed to correct complex drive errors, such as 1,000,000 orphan files, from days to minutes.
- ☑ The live migration of a 75GB VM using the ODX data transfer completed eight times faster than the traditional LAN method while consuming virtually no network resources.

Issues to Consider

- ☑ Default server BIOS, operating system, and application settings were used during ESG Lab testing. As expected, after any testing of this magnitude, analysis of the results indicates that tuning would most likely yield slightly higher absolute results. Given that the goal of this report was not to generate a big number, ESG Lab is confident that the results presented in this report meet the objective of demonstrating new functionality and capabilities of Server 2012, as well as performance and scalability benefits of tier-1 applications and workloads.
- ☑ Though Storage Spaces can meet the ranging needs and expectations from entry-class to enterprise-class SLAs, it's still a new storage solution. Storage Spaces does not yet replace every feature an enterprise-class SAN has to offer. A prime example is that Windows Server 2012 does not yet support in-box replication.
- ☑ ESG Lab witnessed an eight times performance increase using Server 2012's new ODX technology, which proved to be a drive bottleneck. Through detailed analysis, this result is directly proportional to how many data drives are available at the storage array level. With that being said, ESG Lab feels performance levels of ODX could reach as high as ten to twelve times faster than traditional large file and VM transfers over a LAN.

The Bigger Truth

Storage requirements are becoming harder and harder to manage—enterprise applications, virtualization, management of all data types—the list goes on. With strict budget limitations, IT managers are being pressured to do more with less. To address the storage-complexity increases, the high costs, and the performance expectations of most organizations, they need a better approach providing simplified storage management, greater storage efficiency, enterprise-level features all at a low cost, now more than ever.

Virtual server technologies have been proven to deliver significant business value, reducing both OPEX and CAPEX and keeping virtualization at the top of the IT priority list. In addition, organizations are leveraging virtualization to gain mobility and agility, in many cases deploying private clouds and taking advantage of fast and easy provisioning from public cloud providers. In fact, a recent ESG survey shows that 49% of respondents are looking to add additional Hypervisors to their infrastructures in the next 12-18 months. Of those respondents, 81% said that Microsoft Hyper-V was being looked at as a strong alternative to their existing virtualization solution.⁶

But while customers continue to enjoy efficiency improvements, they want to move past the initial consolidation gains and virtualize more mission-critical applications. However, the workload aggregation of virtual server environments can create an “I/O blender” effect, with multiple types and sizes of workloads overloading disk spindles, stressing networks, and consuming server processing resources. These challenges can affect application performance, often preventing expansion of virtual environments and limiting its benefits.

ESG Lab testing demonstrated just how groundbreaking the release of Windows Server 2012 is for Microsoft. The new Server Manager interface offered a refreshing way to configure and manage storage with easy-to-use wizards, and intuitive navigation. The new storage spaces and improved SMB protocol provided flexibility to organizations with limited budgets looking to leverage inexpensive, commodity hardware, while still meeting strict performance requirements. ESG Lab proved that tier-1 applications, such as SQL Server 2012, can not only meet performance expectations, but exceed them, in a consolidated virtual server environment powered by Windows Server 2012.

With data growing faster than disk drive prices are dropping, capacity optimization is as important as ever. ESG Lab tested the newly offered deduplication technology that comes at no extra cost with Server 2012 and was impressed with the great capacity savings and minimal performance overhead. Deploying large volumes on commodity hardware can be done confidently with improved NTFS availability through Chkdsk enhancements. ESG Lab verified that volumes remained online for longer periods of time while common data corruptions were repaired. For more complex errors, ESG Lab confirmed that repairs of volumes with a large number of files, which used to take days to complete, now take minutes. The volume remained online while scanned for errors, and was only taken offline for the quick repair. The new Windows Server 2012 ODX capabilities simplify private cloud deployments by quickly provisioning and transferring large VMs. ESG Lab transferred a 75GB VM from one server to another eight times faster than the non-ODX method while consuming virtually no server or network resources.

Windows Server 2012 was designed to help IT managers optimize a cloud-optimized infrastructure today and in the future. It is the type of foundation required to support the needs and expectations of today’s user, who demands maximum information access with fast performance and minimal interruption, regardless of their location or endpoint device. The new and enhanced storage features bring enterprise levels of performance, scalability, availability, reliability, and efficiency in an affordable package. Microsoft is helping make virtualization and cloud computing easier and faster to deploy, leveling the playing field for enterprises and service providers of all kinds.

The storage efficiency, agility, and transparency of new and improved features make the high-performing, cost-effective Windows Server 2012 a no brainer for small and large businesses alike.

⁶ Source: This data comes from a custom research project conducted by ESG on behalf of Microsoft on the topic of virtualization and private cloud trends in May 2012.

Appendix

Table 7. ESG Lab Test Bed

Storage Spaces Testing	
Physical Server	Type: Dell PowerEdge R910 Processor: 32 core 2.55GHz Intel Xeon X7555 Memory: 524GB Operating System: Windows Server 2012 OS Build: 9200
Virtual Machines	Virtual Machine Count: 8 vCPUs: 4 RAM: 12GB Operating System: Windows Server 2012 OS Build: 9200 Database: SQL Server 2012
Direct-Attached Storage	Type: RAID Inc. 4U 60-Bay EBOD Drives: 60 600GB 15K SAS Volume Count: 8 (one for each VM) Volume Size: 2 TB (including OS, DB, and DB Logs)
SMB Testing	
Physical Server	Type: Dell PowerEdge R910 Processor: 32 core 2.55GHz Intel Xeon X7560 Memory: 524GB Operating System: Windows Server 2012 OS Build: 9200
File Server	Type: Dell PowerEdge R810 Processor: 16 core 2.00 GHz Intel Xeon X6550 Memory: 256GB Operating System: Windows Server 2012 OS Build: 9200
Virtual Machine	Virtual Machine Count: 8 vCPUs: 4 RAM: 12GB Operating System: Windows Server 2012 OS Build: 9200 Database: SQL Server 2012
Network	Speed: 10 GbE
Direct-Attached Storage	Type: RAID Inc. 4U 60-Bay EBOD Drives: 60 600GB 15K SAS Volume Count: 8 (one for each VM) Volume Size: 2 TB (including OS, DB, and DB Logs)
Deduplication Testing	
Physical Server	Type: HP DL380p G8 Processor: 16 core 2.6GHz Intel Xeon E5-2670 Memory: 196GB Operating System: Windows Server 2012 OS Build: 9200

Direct-Attached Storage	Type: RAID Inc. 4U 60-Bay EBOD Drives: 60 3TB 7.2K Near Line SAS Volume Count: 4 <ul style="list-style-type: none"> - 2 for original data - 2 for deduplicated data Volume Size: 2TB
Chkdsk Testing	
Physical Server	Type: HP DL380p G8 Processor: 16 core 2.6GHz Intel Xeon E5-2670 Memory: 196GB Operating System: Windows Server 2012 OS Build: 9200
Virtual Machine	Virtual Machine Count: 2 vCPUs: 4 RAM: 12GB Operating System 1: Windows Server 2012 Operating System 2: Windows Server 2008
Direct-Attached Storage	Type: RAID Inc. 4U 60-Bay EBOD Drives: 60 3TB 7.2K Near Line SAS Volume Count: 18 (9 for each VM) <ul style="list-style-type: none"> - Bad file record segment – 100,000 files - Bad file record segment – 1,000,000 files - Bad file record segment – 10,000,000 files - Corrupt Index Entry – 100,000 files - Corrupt Index Entry – 1,000,000 files - Corrupt Index Entry – 10,000,000 files - 1,000,000 orphan files – 100,000 files - 1,000,000 orphan files – 1,000,000 files - 1,000,000 orphan files – 10,000,000 files Volume Size: 2TB
ODX Testing	
Physical Servers	Type: 2 Dell PowerEdge R910 Processor: 32 core 2.55GHz Intel Xeon X7555 Memory: 524GB Operating System: Windows Server 2012 OS Build: 9200
Network	10 GbE
SAN Storage	Type: Dell EqualLogic PS Series Storage Array Drive Count: 12 600GB 15K SAS drives Volumes: 2 LUNS within RAID5 storage pool VM Size : 75GB

Table 8. ESG Lab Deduplication Powershell Commands

Powershell Command	Description
<i>Enable-DedupVolume <volume>:</i>	This command enables deduplication on a specified volume
<i>Set-DedupVolume -<volume>: -MinimumFileAgeDays 0</i>	By default, deduplication optimizes files that are modified at least 5 days in the past. This command allows for the optimization of all files regardless of creation and modification timestamp.
<i>Set-DedupSchedule * -Enabled:\$false</i>	This command ensures that the dataset in the volume does not get deduped prior to measurement.
<i>Start-DedupJob -<volume>: -Type Optimization</i>	This command runs deduplication in throughput mode.
<i>Start-DedupJob -<volume>: -Type Optimization -Memory 25 -StopWhenSystemBusy</i>	This command runs deduplication in background mode.



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