



LAB VALIDATION REPORT

IBM SAN Volume Controller Optimizing Storage Utilization and Availability

By Brian Garrett

August, 2008

Table of Contents

Table of Contents	i
Introduction	1
<i>Background</i>	1
ESG Lab Validation	4
<i>Getting Started</i>	5
<i>Space Efficient Virtual Disk</i>	9
<i>Data Migration</i>	12
<i>Scalability and Performance</i>	14
<i>Flash Copy</i>	18
<i>Metro and Global Mirror</i>	20
<i>Virtual Disk Mirror</i>	22
<i>Advanced Manageability</i>	23
<i>Fault Tolerance</i>	25
ESG Lab Validation Highlights	26
Issues to Consider	26
ESG Lab's View	27
Appendix	28

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 IBM.

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.

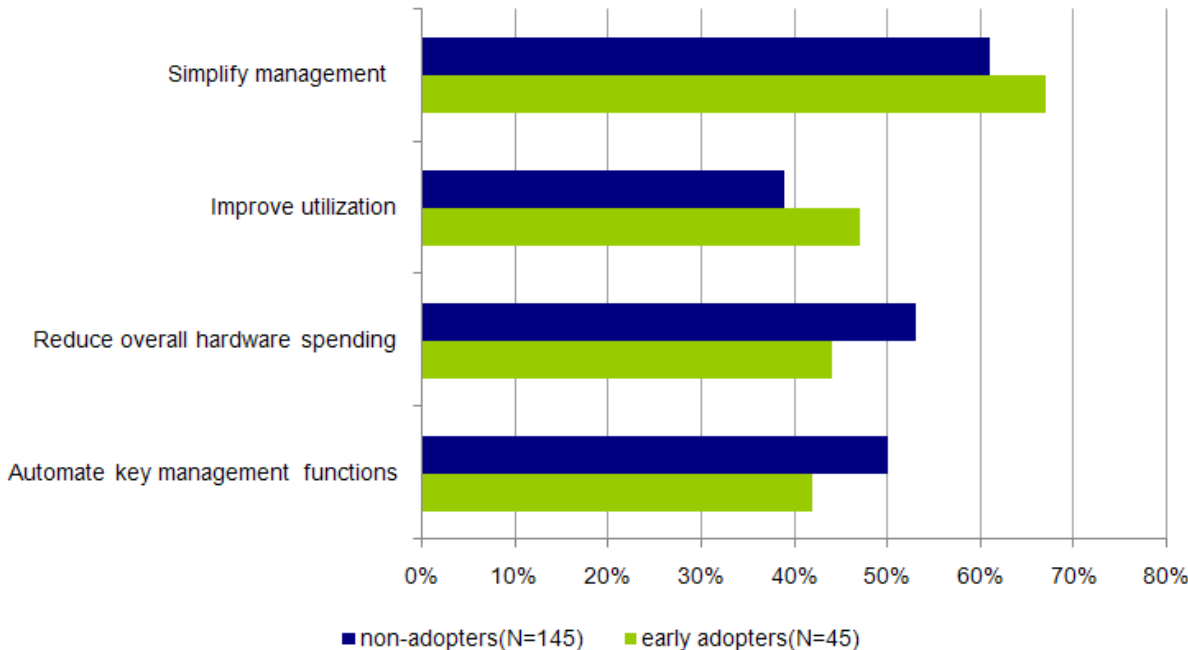
Introduction

IBM SAN Volume Controller (SVC) is a network-based storage virtualization solution with over 13,000 appliances shipped running in more than 4,300 systems worldwide. This report documents the results of ESG Lab testing of IBM SVC performed in 2006 with a focus on the fundamentals of network based storage virtualization including ease of use, non-disruptive virtualization, data mobility and copy services. The enhanced storage utilization and availability capabilities of SVC 4.3 as validated by ESG Lab in 2008 are also presented.

Background

ESG surveyed 210 IT decision-makers who evaluate and purchase storage networking technologies to see who is considering an intelligent storage network solution and why. For the purposes of this study, ESG defined an intelligent storage network solution as one “that provides storage services like storage virtualization, volume management, data migration, and data replication that execute directly on network-resident devices such as SAN switches or purpose-built appliances as opposed to hosts or storage subsystems.”¹ As shown in Figure 1, simplified management, improved utilization and reduced costs topped the list of potential benefits that non-adopters were most interested in. A strikingly similar list of realized benefits was reported by early adopters. As a matter of fact, for the top two benefits on both lists, the actual benefits achieved by early adopters exceeded the expected benefits of planned adopters.

FIGURE 1. EXPECTED VS. ACTUAL BENEFITS OF STORAGE VIRTUALIZATION

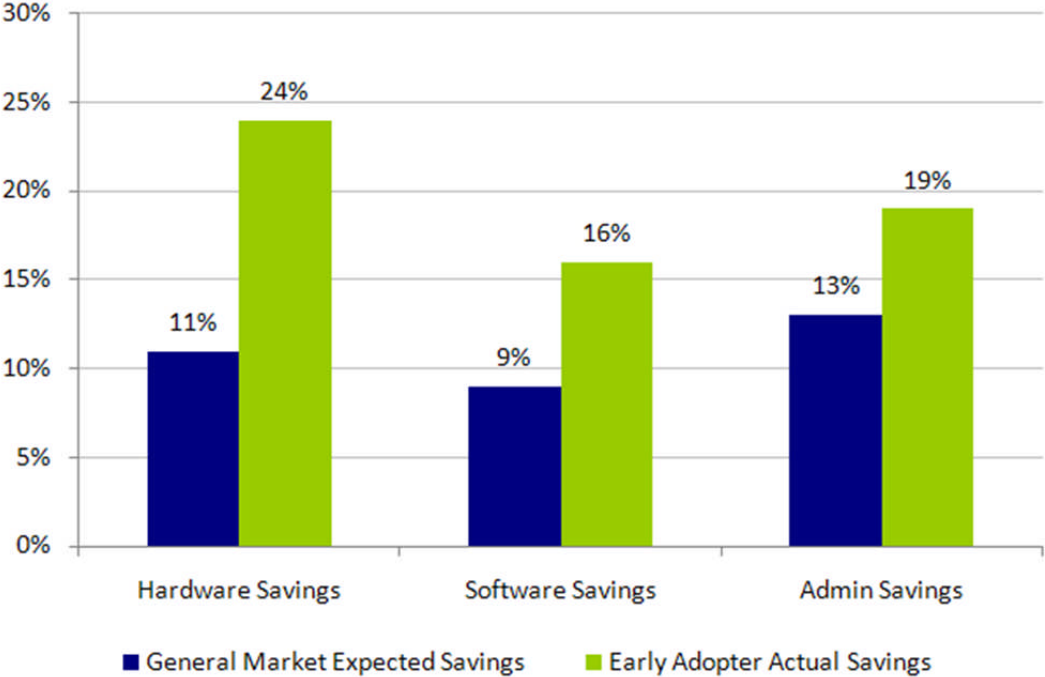


¹ Source: ESG Research Report, *The Future of Network-Based Storage Intelligence*, September 2004, 145 non-adopters, 45 early adopters

The benefits shown above should look strikingly familiar to IT professionals who have experience with server virtualization. Like server virtualization, storage virtualization enables consolidation, which reduces complexity and cost. Storage and server virtualization are also similar in their ability to provide increased availability, fault tolerance and mobility. Given the similarities, it's not surprising that, in a more recent study, ESG learned that an increasing number of organizations are deploying server and storage virtualization together. Twenty-four percent have already deployed storage virtualization in conjunction with server virtualization and 33% plan on doing so within the next 24 months.²

One of the most significant findings—if not the most significant—is the substantial economic benefit of deploying an intelligent storage network solution. For example, 79% of early adopters with a large SAN believe they have reduced their annual storage hardware spending to some degree, with these users reporting a mean average annual savings of 19.7%. As shown in Figure 2, storage hardware savings were the most significant due to the ability to reclaim and re-use existing storage (the least expensive storage is the storage you already have). Consolidating storage software that previously ran on multiple servers or storage systems onto a centrally managed infrastructure also helped reduce storage software spending and storage administration costs.

FIGURE 2. EARLY ADOPTER SAVINGS



² Source: ESG Research Report, *The Impact of Server Virtualization on Storage*, December 2007, N=332

The IBM SAN Volume Controller

SVC is an appliance-based storage virtualization solution that is deployed in a storage area network. Up to four pairs of appliances, deployed in a fault tolerant cluster and managed as a single solution from a centralized console, can be deployed in a data center to provide:

- Optimized capacity utilization and availability
- A centralized platform for the management of SAN-attached storage capacity
- Heterogeneous interoperability between a wide variety of hosts, operating systems and storage systems
- Online data migration services
- Remote replication and disaster avoidance services
- Space efficient provisioning and copy services

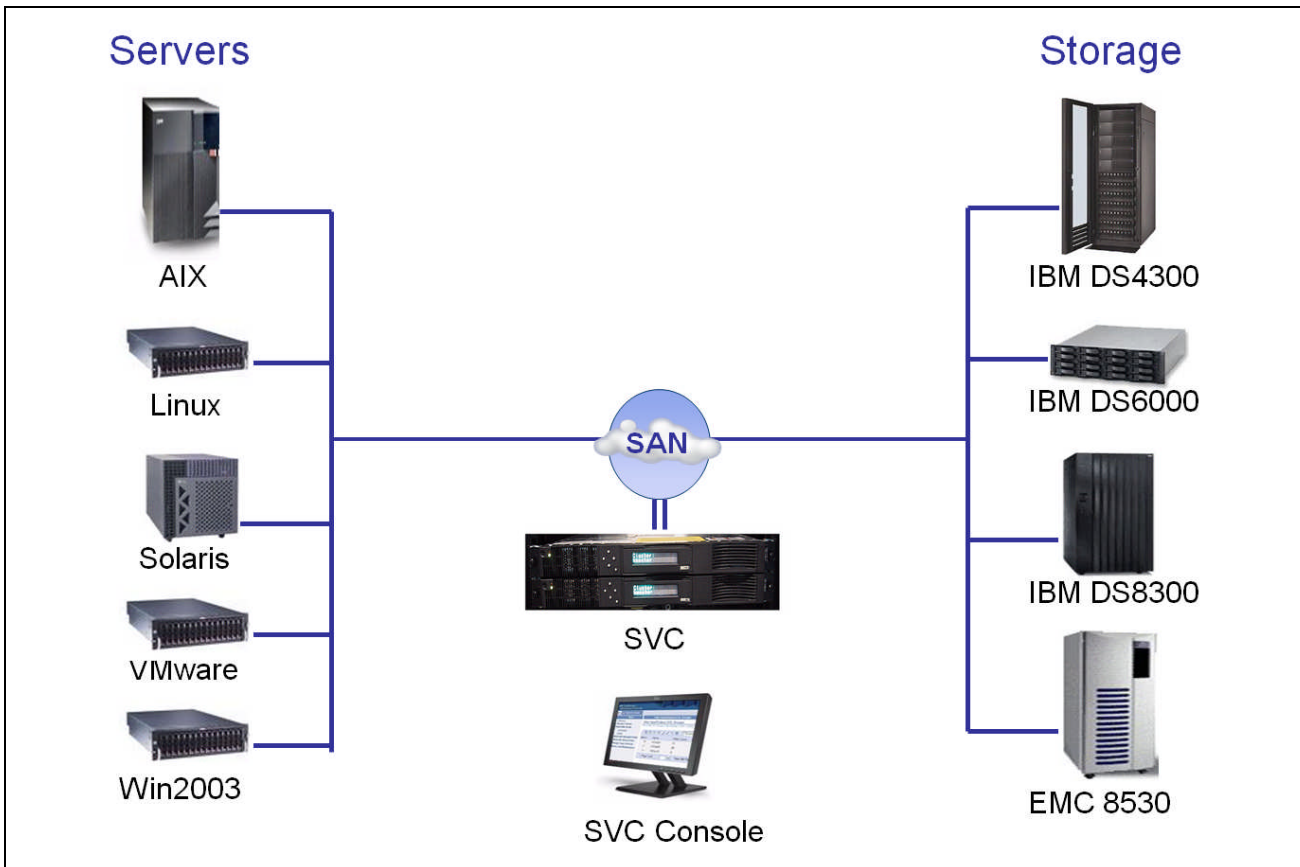
SVC may be used to provide these functions to enhance existing installed storage or as part of a new storage deployment.

The goal of ESG Lab testing of the IBM SVC was to validate the fundamentals of network-based storage virtualization including ease of use and non-disruptive virtualization, data mobility and copy services. The performance impact of SVC virtualization services and the capacity and availability enhancements introduced in SVC version 4.3 were also explored. ESG Lab found that the IBM SVC is a rock-solid, feature-rich platform that delivers on the promise of network based storage virtualization by reducing the complexity and cost of managing SAN attached storage.

ESG Lab Validation

ESG Lab validation of the SVC virtualization platform was performed at an IBM facility in Gaithersburg, Maryland. The test bed used during the validation is shown in Figure 3 and is documented in the Appendix. Separate test beds used to test Metro and Global Mirroring are presented later in this report.

FIGURE 3. THE ESG LAB TEST BED



SVC was used to provide storage virtualization services on each of the operating systems and storage systems shown here during the ESG Lab Validation. Host Bus Adapters from QLogics and MDS switches from Cisco were used for SAN server connectivity and switching. The IBM SDD multi-path driver, which is provided at no extra charge with SVC, was used for a fault tolerant connection between servers and the SAN fabric for all of the operating systems tested, except for VMware, where the native multi-path driver built into VMware ESX was used.

Since SVC was first announced in 2003, IBM has been committed to expanding the interoperability of SVC. The breadth and depth of the test bed used during the ESG Lab Validation represents a fraction of the servers, operating systems, host bus adapters, switches and storage systems supported by SVC. A complete list of the hardware, firmware and driver levels supported by IBM is available on the IBM support website.³

³ <http://www.ibm.com/systems/storage/software/virtualization/svc/interop.html>

Getting Started

SVC appliances are logically configured in-band between servers and SAN attached storage. Physically introducing SVC into an existing SAN environment is performed in two steps:

1. Run FC cables between the SVC appliances and a SAN switch.
2. Add the SVC to the server and storage to the Fibre Channel zones that you would like to manage.

Once cabled, the SVC console is used to configure virtual storage in three easy steps:

1. Run a discovery process (or wait for the next automatic discovery to occur).
2. Configure a logical unit (LUN) presented by a storage array as a managed disk (MDisk).
3. Create a virtual disk (VDisk) to be presented to a server using one or more MDisks.

SVC can be configured to work with volumes that are already in use by an existing application. After the three steps above have been performed and the server has performed a re-scan or reboot, applications can be restarted. When used in this manner, SVC operates in a mode referred to as Image Mode. Image Mode sets up a one-to-one mapping of a physical MDisk to a host-accessible VDisk. Because the VDisk has exactly the same mapping as the underlying MDisk, any data already on the disk is still accessible when migrated to an SVC environment. Image Mode VDisk support is a powerful concept that can be used for initial testing of SVC connectivity. A storage administrator can optionally de-configure or “undo” an Image Mode VDisk during initial SVC testing.

After an Image Mode VDisk has been configured and tested, it is typically migrated to a MDisk Group (MDG) to gain the full benefits of SVC virtualization and copy services. An MDisk Group is a collection of managed disks presented by one or more storage arrays. MDisk Groups can be configured in two modes: striped or concatenated. Host-accessible VDIs, which have been built using capacity from an MDisk Group, can be expanded, moved and copied while applications remain online and available.

ESG Lab Testing

ESG Lab validation testing began by moving through a three step process that is typically performed by customers getting started with SVC. As shown in Figure 4, a LUN residing within an IBM disk array and accessed from a Windows server as drive letter F: was tested in three modes:

1. Un-managed

This is the state before SVC was introduced. A LUN within the disk array is being accessed directly by a SAN connected Windows server as driver letter F:.

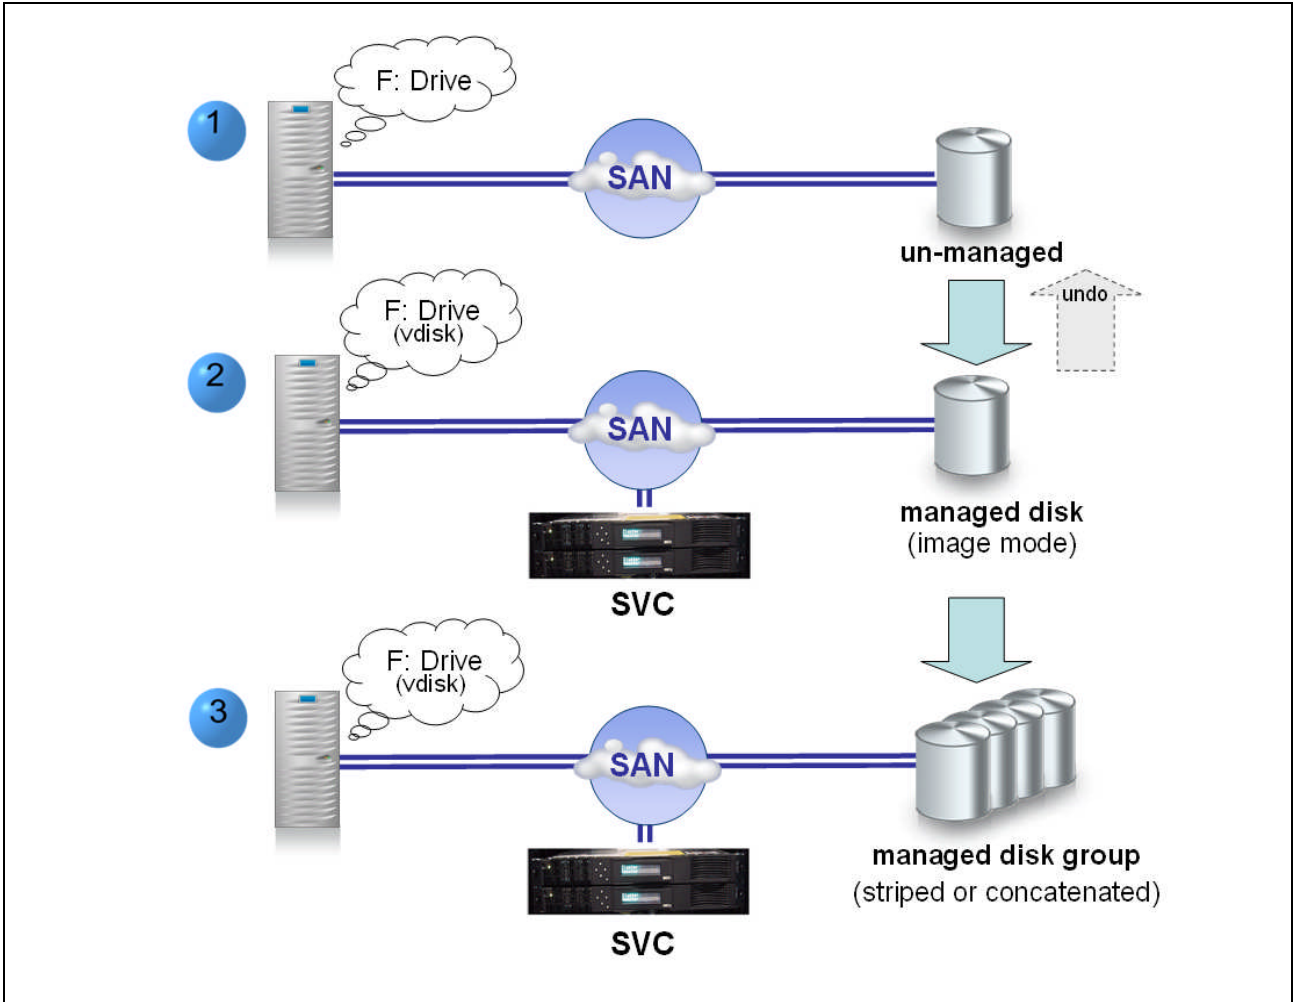
2. Managed Image Mode VDisk

This is the state after SVC was first cabled into the SAN and configured. The LUN is now being managed by SVC and accessed by the Windows server as an Image Mode VDisk. Data written to the F: drive before SVC was introduced remains accessible and unchanged.

3. A built over a Managed Disk Group

This is the more typical usage of an SVC managed disk group. In this case, the F: drive being used by the Windows server is composed of capacity from four managed disks within a storage array. This can now be migrated, copied within an array (FlashCopy), and replicated between arrays (Metro or Global Mirror), and expanded—all while applications continue to access the F: drive without interruption.

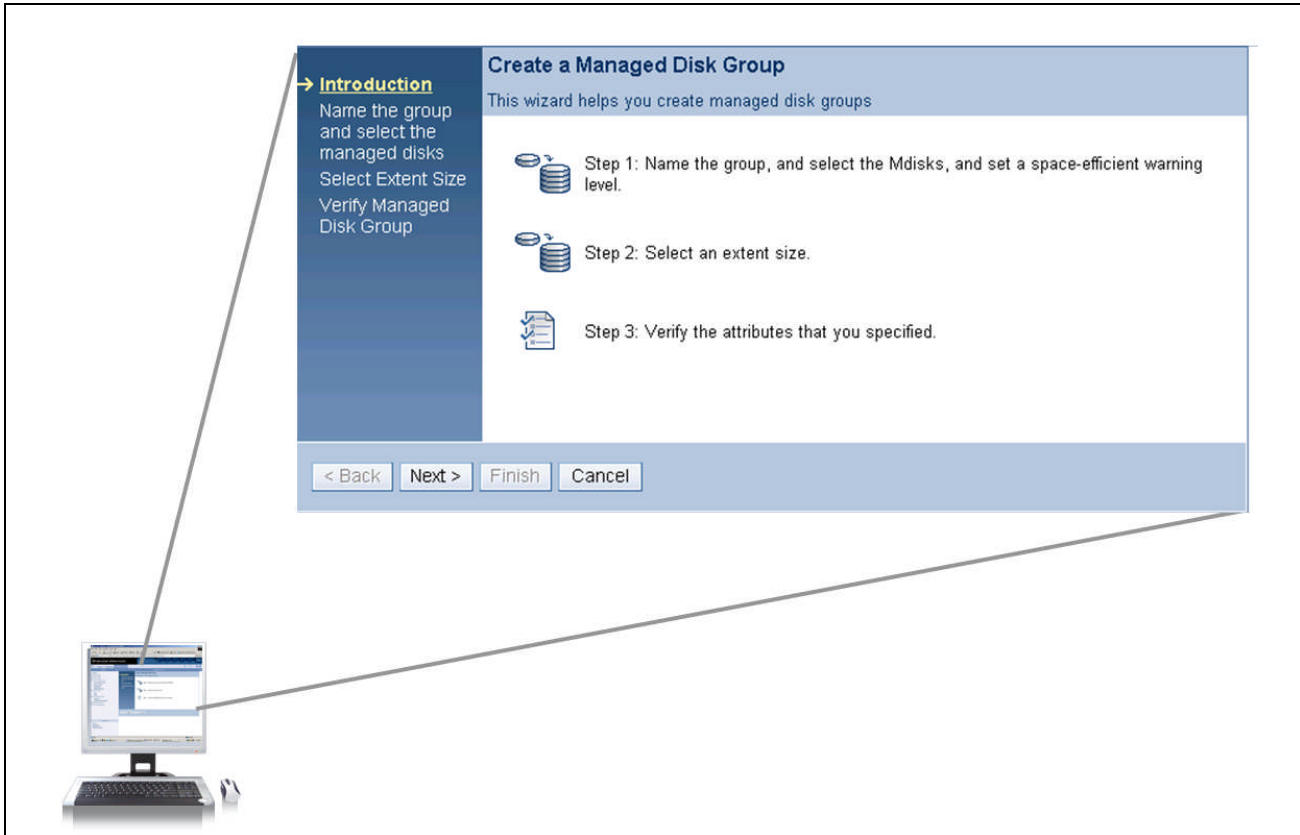
FIGURE 4. GETTING STARTED WITH SVC VIRTUALIZATION



ESG Lab testing began with an un-managed LUN residing on an IBM storage array being used by an NTFS file system on a Windows 2003 cluster. The NTFS file system was accessed as drive letter F (as shown above in step 1). The contents of the program file directory on the C: drive were copied to the F drive before it was configured as an Image Mode VDisk using SVC (as shown above in step 2). After a Windows disk administrator utility rescan, the drive was again accessed as drive letter F: and the copy of the program files directory was inspected. From a user standpoint, there was no indication that anything had changed at all.

The SVC Validation continued with a quick tour of the basic capabilities of the SVC. ESG Lab found the SVC management console to be intuitive and straightforward. Commonly performed tasks are wizard-driven as shown in the screen shot in Figure 5, which was used to create a managed disk group.

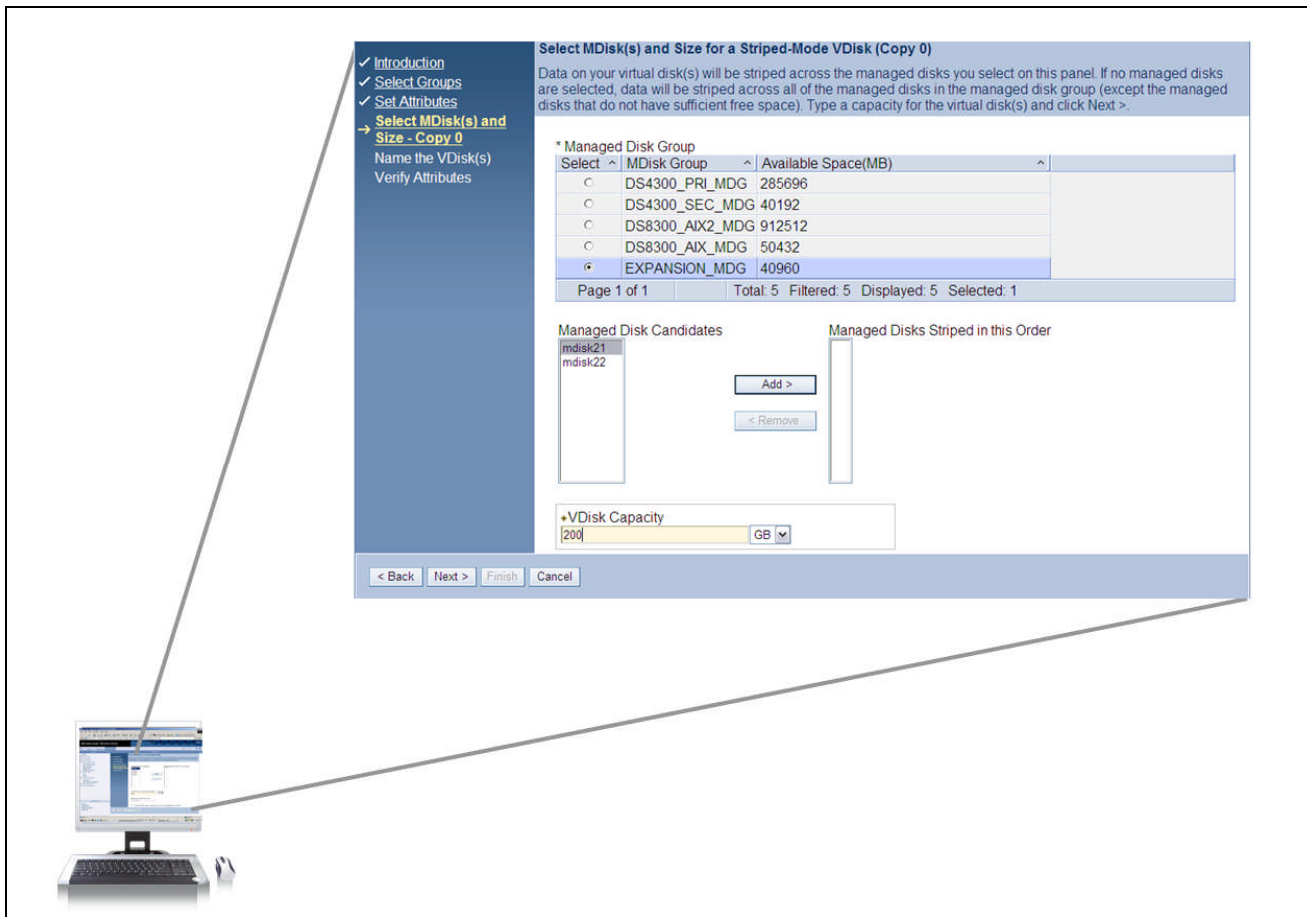
FIGURE 5. WIZARD-DRIVEN EASE OF USE



Managed disk groups are used to create a pool of storage from which virtual disks are carved. Managed disk groups can be striped or concatenated. A 200 MB striped mode managed disk was configured using the intuitive wizard shown in Figure 6.

ESG was impressed by how easy it was to get up and running with SVC. The console application has a very consistent look and feel, includes wizards for commonly performed tasks, and has excellent online help. With virtually no supervision, the following tasks were completed in just under an hour using multiple storage arrays from multiple vendors, connected to a variety of operating systems, created managed disk groups (striped and concatenated), created a VDisk, expanded a VDisk, deleted a VDisk, migrated a VDisk, created a flash copy of a VDisk, monitored task progress and viewed logs.

FIGURE 6. CREATING A MANAGED DISK GROUP



The IBM System Storage Multipath Subsystem Device Driver (SDD) was pre-installed on the Windows server for this phase of testing. SDD is a multipath driver that is installed on servers to provide a fault tolerant path to shared network attached storage. IBM provides SDD at no additional charge for SVC customers, and is committed to the support of emerging native multi-path drivers built into operating systems (e.g., the native VMware multipath driver used later during ESG Lab testing). For customers with many servers, this can significantly reduce the total cost of ownership of a shared fault-tolerant SAN storage solution, especially when compared to competitive virtualization solutions that mandate the use of for-fee multipath software. As a matter of fact, the savings in multi-path software licensing alone may in some cases provide cost justification for the initial purchase of SVC.

Why This Matters

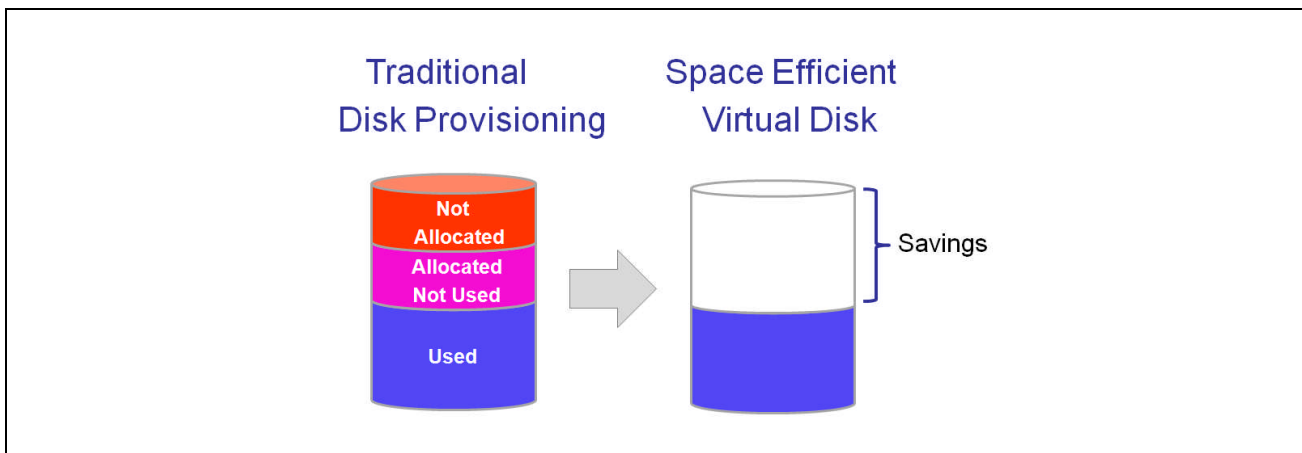
ESG Lab found that getting started with SVC virtualization was intuitive and straightforward. An existing volume being used by an application can be changed into an SVC Image Mode VDisk (and back again if desired) in a matter of minutes. Once under SVC control, volumes can be expanded, migrated, copied and replicated to a remote site—all while applications remain online and available.

Space Efficient Virtual Disk

Space Efficient Virtual Disk, a feature introduced in SVC version 4.3, can be used to present a virtual pool of shared capacity that is bigger than the actual amount of physical storage available. This valuable capability, often referred to in the industry as thin provisioning, was designed to maximize capacity utilization and simplify storage capacity allocation.

When provisioning storage using traditional methods, system administrators allocate a fixed amount of capacity that is dedicated to an application. To avoid the potentially devastating impact of an application running out of storage capacity, storage administrators typically allocated more capacity than was needed to accommodate growth. As shown in Figure 7, on top of this allocated, but unused pool of capacity, a second pool of capacity that has not been allocated is typically set aside as a reserve for applications that have exceed their pre-allocated capacity. Provisioning storage capacity from a centrally managed pool of just-in-time capacity, the IBM SVC Space-Efficient Virtual Disks function simplifies the cumbersome task of storage provisioning while providing significant capacity savings.

FIGURE 7. OPTIMIZING DISK CAPACITY UTILIZATION WITH SPACE EFFICIENT VIRTUAL DISK



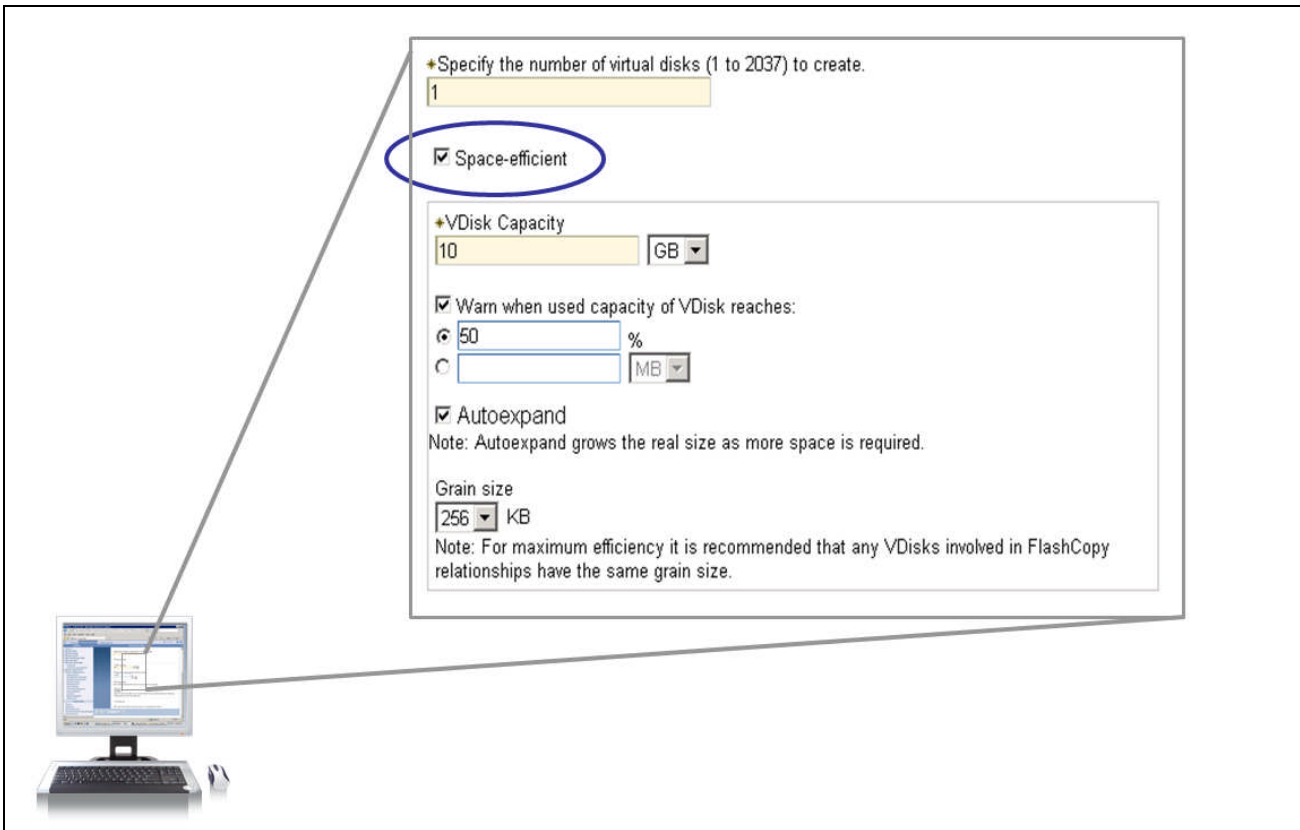
As an example of what happens during traditional storage provisioning, consider a 500 GB volume that has been allocated to an application with only 100 GB of actual data being created. In this example, the other 400 GB has no data stored on it. That allocated and unused capacity belongs to that application—no other application can use it. This means that 400 GB of the allocated 500 GB is wasted storage, which means that it also wastes money and energy. And even though all of the storage capacity may eventually be used, it may take years to do so, which is a waste of storage capacity and money.

With IBM SVC space-efficient virtual disks, the storage provisioning process starts out the same. The storage administrator configures a virtual disk with 500 GB of capacity to be presented to a server. Later, when only 100 GB of actual data has been written, the unused 400 GB is still available for other applications. This approach allows the application to grow transparently while simultaneously ensuring that capacity is not wasted. In other words, SVC delivers just in time storage capacity. The application thinks it has 500 GB of storage, but SVC allocates capacity as needed. The rest of the capacity stays in a shared managed disk pool. Administrators can set thresholds to be alerted when capacity needs to be added.

ESG Lab Testing

ESG Lab configured a space-efficient virtual disk for use on a Windows 2003 server. ESG Lab noted that this valuable new capability, introduced in SVC version 4.3, was added as a checkbox to the virtual disk creation wizard as shown in Figure 8. After selecting the space-efficient checkbox, a dialogue box was used to define the size of the volume presented to the Windows server (10 GB), enable a warning when capacity utilization reached a defined threshold (50%), and enable automatic capacity expansion. ESG Lab noted that capacity threshold warnings can be set not only at the virtual disk level, but also for a managed disk pool shared by a number of virtual disks.

FIGURE 8. CONFIGURING A SPACE EFFICIENT VDISK

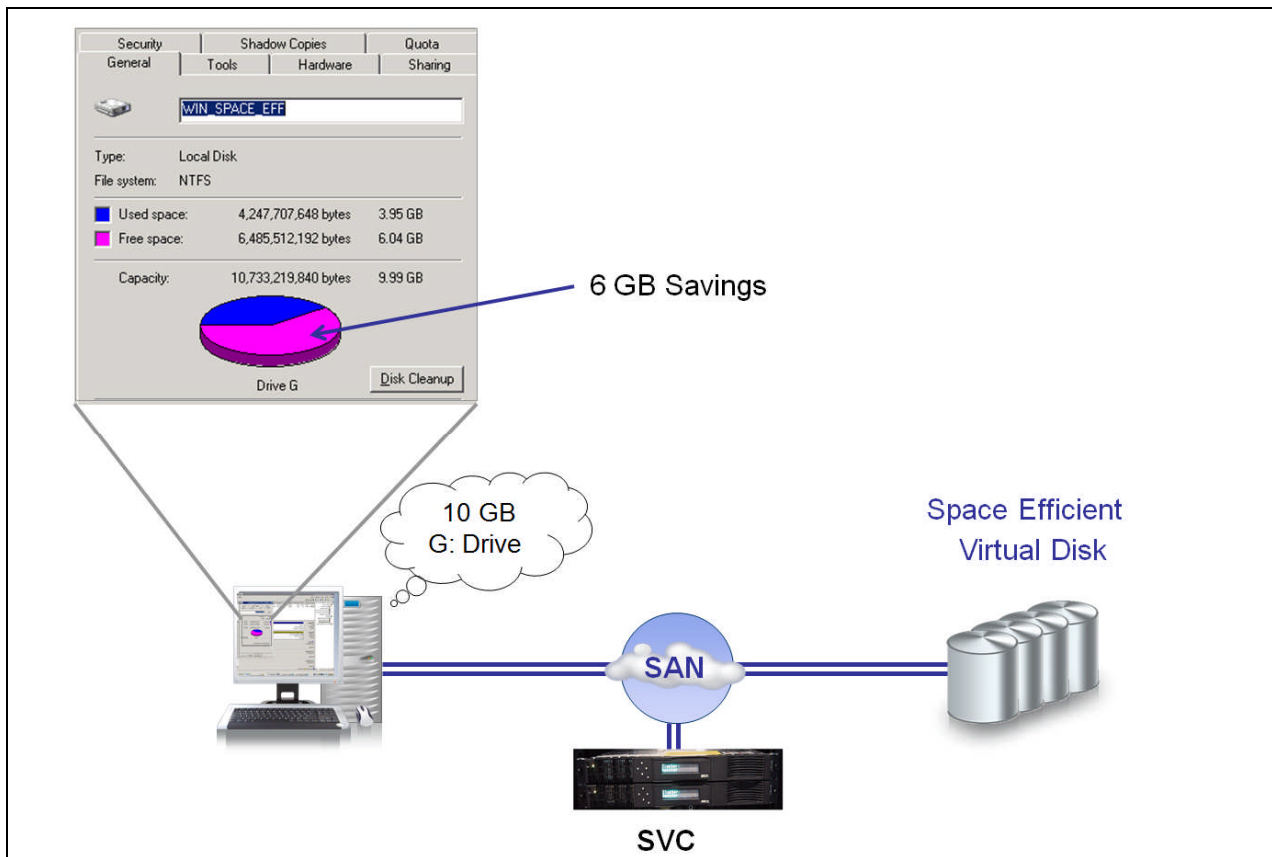


The 10 GB space efficient disk was presented to the server and discovered using the Windows disk administrator utility. An NTFS file system was created and presented as a G: drive. The Microsoft Windows Explorer utility was used to monitor used and free capacity as files were copied from the C: drive on the Windows server to the space efficient G: drive.

The SVC management console was used to monitor the actual physical capacity used. Automatic expansion was observed as physical capacity was allocated as needed. When the used capacity as reported by the Windows Explorer utility exceeded 50%, an SVC threshold warning appeared in the SVC management console log as expected.

At the point in time shown in Figure 9, a pair of SAN-attached IBM SVC appliances was presenting a 10 GB G: drive using only 4 GB of actual disk capacity. With traditional provisioning methods, the 6 GB shown as free space by the Windows Explorer utility would be stranded and wasted. In contrast, a quick look at the IBM SVC console confirmed that only 4 GB of capacity had been consumed. In other words, 6 GB of capacity savings was saved as SVC delivered just-in-time capacity from a shared pool of managed disk.

FIGURE 9. OPTIMIZING DISK CAPACITY UTILIZATION WITH SPACE EFFICIENT VIRTUAL DISK



Why This Matters

ESG conducted a telephone survey of enterprise-class storage administrators with a focus on the limitations of traditional storage provisioning methods.⁴ Fifty-four percent were aware that they had stranded and unused storage capacity due to inefficient provisioning methods. Fifty-five percent reported that between 31 and 50 percent of purchased capacity was stranded and unused. Eighty percent felt that storage provisioning was a significant time and resource drain.

Thin provisioning simplifies the cumbersome task of storage provisioning and improves capacity utilization. ESG believes that thin provisioning is one of the most useful functions of a storage system and yet few storage systems currently support this capability. Space-efficient virtual disks provide the benefits of thin provisioning for all SVC-managed disk arrays even if they do not support thin provisioning themselves.

With a wizard-driven configuration process and advanced management capabilities including alert thresholds at the virtual disk and pool level, ESG Lab testing has confirmed that SVC space-efficient virtual disks can be used to reduce the cost and complexity of storage provisioning while providing significant capacity savings.

⁴ Source: ESG Report, *Thin Provisioning*, April 2006, N=20

Data Migration

Data migration is the movement of data that is currently in use by an application to a new set of storage devices, often on a different storage array. Data migrations occur most often when an organization buys new storage. They also occur to provide more capacity or performance for an application to meet the needs of new applications or to move information assets between tiers of storage with different performance and cost attributes according to the needs of the business.

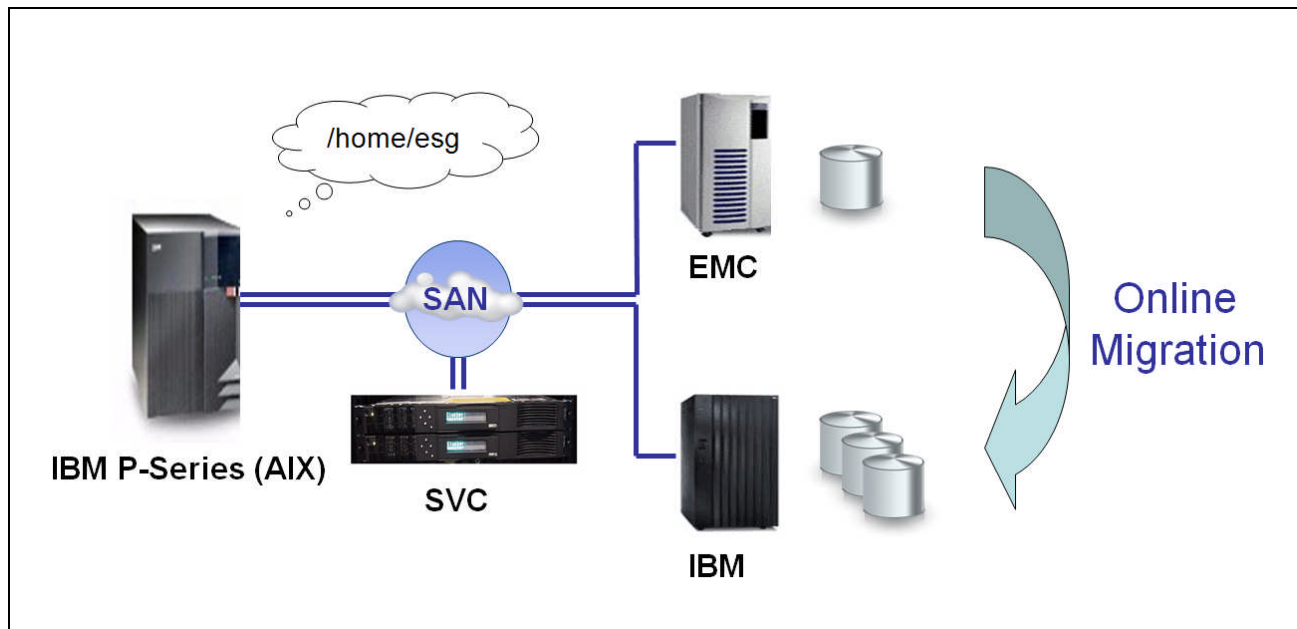
Traditional data migration methods are complicated and disruptive. Applications must typically be shut down for the duration of the migration. Data is moved from one system to another using backup tapes or software, which moves the data over a LAN. Regardless of how the data is moved, traditional data migrations cause extensive downtime, consume significant resources and are prone to error.

The SVC platform performs data migrations without causing application downtime. Migrations can be performed within and between disk arrays from multiple vendors. Instead of consuming server, backup, or LAN resources to move the data, SVC appliances are used as the engine which drives data migrations over a high speed SAN. SVC migrations can be started and monitored from the SVC console or scripted using the SVC CLI. To minimize any potential performance impact on running applications, the priority/speed of SVC data migrations is user selectable.

ESG Lab Testing

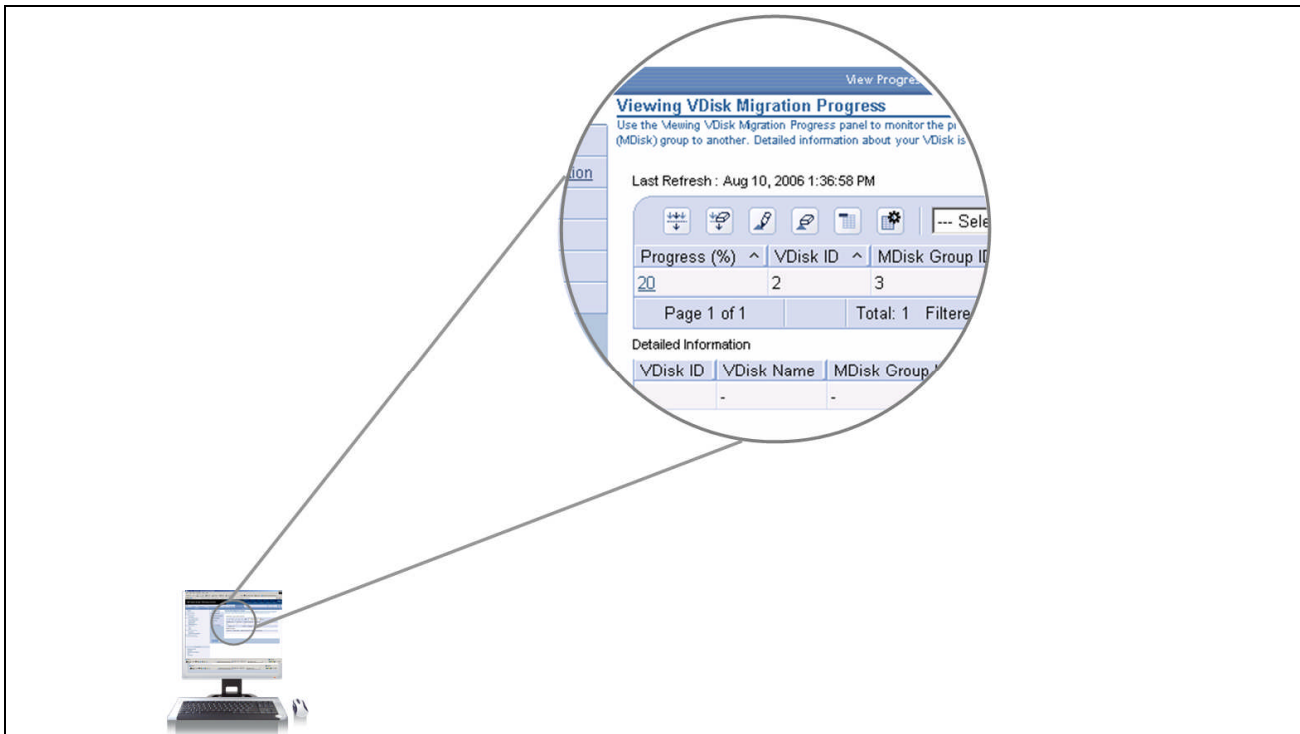
ESG Lab performed a migration from an EMC Symmetrix to an IBM DS8300 disk array, as shown in Figure 10. The volume was being used as a file system (/home/esg) on a server running the IBM AIX operating system. A script was started on the AIX server to exercise the volume during the migration. The script created a file every two seconds and ran for the duration of the migration. No performance impact was noticed during the migration.

FIGURE 10. AN ONLINE MIGRATION FROM AN EMC SYMMETRIX TO A IBM DS8300



The migration was initiated using an intuitive wizard from the SVC console. Migration status was monitored from the SVC console, as shown in the screen shot depicted in Figure 11 that was taken when the migration was 20% complete. A 6 GB volume took approximately one minute to migrate, yielding an effective single volume migration rate of 100 MB/sec.

FIGURE 11. MONITORING MIGRATION STATUS



ESG Lab spoke with a customer who is a "...huge fan of SVC. With our four node SVC cluster, we were able to migrate over 50 TB of data in less than four months. The first couple migrations were performed during maintenance windows, but we didn't have to shut down applications. Now that we are comfortable with the system, migrations are routinely performed outside of maintenance windows with applications up and running." In addition, the customer said, "SVC is now a core part of our infrastructure. We first embraced virtualization in our server infrastructure with VMware. Now that we've virtualized our storage infrastructure with SVC, we can quickly deploy new applications without worrying about over provisioning. If an application developer comes to me with a request for a hundred user application that grows to a thousand users overnight, I can move the application to a new server or storage pool on demand."

Why This Matters

Data migrations using traditional methods cause application downtime and can have significant business impact. ESG has found that most customers can tolerate very little downtime without impacting revenue. ESG research indicates that the majority of respondents could not tolerate more than one to four hours of downtime before experiencing a significant impact to their businesses (23%). Nearly as many said that they could not tolerate more than an hour (22%) and many couldn't tolerate any downtime at all (14%).³ ESG Lab has verified that SVC performs data migrations with no application downtime.

Scalability and Performance


When considering the performance impact of a virtualization layer inserted between applications and SAN attached storage, the first issue to consider is the delay caused by such a layer during normal processing. ESG Lab's experience with a variety of storage virtualization solutions has shown that while these delays deserve scrutiny, they tend to be so small in comparison to the length of a storage operation that they won't be noticed. The potential performance benefits that can be gained using caching and drive striping within a virtualization layer and the performance scalability of the virtualization layer are just as important, if not more so.

ESG Lab Latency Analysis

ESG Lab audited the results of IBM tests, which measured the latency inserted by SVC during a 4 KB sequential read from disk. A single logical drive was tested. The elapsed time needed to read 6 GB of data was measured before and after the insertion of SVC between the server and a storage array. The sequential pre-stage capability of the SVC was disabled to negate the performance benefits of caching. The latency added by SVC to each I/O operation was 60 microseconds.

To put a 60 microsecond delay into perspective, consider the anatomy of a 4 KB disk read request as shown in Figure 12. The timeline for the I/O starts with a read/write command being sent to the drive from an application. This is followed by long, mechanical access times waiting for the drive to move the actuator (seek) and spin the actuator over the data requested (latency). Next, the data is transferred from the drive to the server and a status handshake is performed to terminate the request. The timeline shown here uses the average seek and latency numbers for a popular high-performing Fibre Channel drive and command and status overhead measured using a bus analyzer.⁵ Note that the access time associated with the mechanical disk drive (seek + latency) is clearly responsible for the majority of the "wait time."

FIGURE 12. THE ANATOMY OF A DISK OPERATION



A 60 microsecond delay is very small in comparison to the drive wait time, which is typically measured in milliseconds. As a matter of fact, it is so small that it can't be drawn to scale on the diagram. If it could be drawn, the bulk of the delay would be inserted at the very beginning of the I/O on the far left. To put the delay into perspective, consider a typical disk I/O that takes 5 milliseconds. Sixty microseconds represents only 1.2% of five milliseconds. In other words, at the application level where performance delays matter most, a microscopic delay of this magnitude cannot be detected by users. Moreover, this impact only occurs during the minority of I/O requests that are reads serviced directly from disk instead of SVC cache.

⁵ The data transfer, command and status overheads had to be increased so that you could even notice them on the graph.

ESG Lab Caching and Striping Analysis

The micro-level latency impact of SVC is dwarfed by the macro-level performance benefits of SVC caching and striping. Experienced storage professionals are well aware of the performance benefits of caching. SVC uses advanced algorithms, similar to those in caching disk arrays, to perform writes at memory speed, pre-stage sequentially accessed read data and deliver frequently accessed reads from memory instead of disk.

SVC has eight gigabytes of cache per node (and thus a minimum of 16 GB per cluster), which is larger than many midrange disk arrays. This large centralized cache provides dramatic performance benefits. Cache hits, which use memory to perform data transfers in microseconds, are dramatically faster than cache misses, which take milliseconds because they are serviced from mechanically spinning disks. ESG Lab analyzed the cache pre-stage benefits for a 64 KB sequential read. When compared to a read miss with cache disabled, ESG Lab noted that SVC caching improved performance by 25%. The test bed and methodology for this result was the same as that presented above during the ESG Lab Latency Analysis.

Experienced storage professionals should also be aware of the performance advantages that can be realized by striping application data over multiple disk drives. The more drives per application volume, the more work that can be done in parallel. This is a simple matter of physics. A mechanically spinning disk drive can only find and transfer one request at a time. For applications like databases where multiple users tend to access data on the same application volume, striping over many physical drives can provide dramatic performance benefits. Striping has traditionally been implemented in host based volume manager software and within RAID groups inside storage arrays. One of the key benefits of storage virtualization is the ability to deliver this performance capability as a centrally managed service. SVC can be used to stripe application volumes over as many drives as desired. Drives can be striped within and between arrays. Drives can be added to an existing volume to increase performance on the fly.

Why This Matters

SAN attached storage is often deployed for performance sensitive applications including databases and e-mail. Customers therefore need to be assured that a virtualization layer inserted between applications and SAN attached storage doesn't negatively impact application performance. ESG has confirmed that in the worst case of a cache read miss, the SVC adds a very short delay (60 microseconds), which is invisible to users and applications. In the more typical case, the caching and wide striping capabilities of SVC actually improves application performance.

ESG Lab Performance Scalability Analysis

The modular clustered design of SVC was originally conceived at the IBM Almaden Research Center with infinite scalability in mind. Industry standard servers are clustered together to work as a single fault tolerant system. Servers are added to the cluster in pairs to increase the aggregate performance of the system. As of the writing of this report, up to eight nodes per SVC cluster are supported.

The performance scalability of SVC is evident in published Storage Performance Council (SPC) results. SPC is a vendor-neutral standards body designed to be a source of comparative storage subsystem performance information.⁶ There are two application workloads supported by SPC. SPC-1 simulates an interactive database application. SPC-2 simulates a bandwidth-intensive application such as those that process large files, large database queries or video on demand. Results are shown in the following tables.

TABLE 1. TOP THREE SPC-1 PERFORMANCE RESULTS

Rank	Product	SPC-1 IOPS	\$/IOPS	Capacity (GB)
1	IBM SVC 4.2	272,505	\$12.05	24,434
2	HP StorageWorks XP24000	200,246	\$17.96	26,000
3	Hitachi Data Systems USP V	200,246	\$17.61	26,000

TABLE 2. TOP THREE SPC-2 PERFORMANCE RESULTS

Rank	Product	SPC-2 MBPS	\$/MBPS	Capacity (GB)
1	IBM SVC 4.2	7,084	\$463.66	101,155
2	IBM SVC 4.1	4,544	\$400.78	51,264
3	IBM SVC 3.1	3,518	\$563.93	20,615

What the Numbers Mean

- The IBM SVC has been used to record the best disk-based SPC-1 and SPC-2 results in the industry.⁷
- Each of the SVC tests were run on an eight node cluster in front of a large number of fast Fibre Channel drives within a number of modular IBM DS4000 disk arrays.
- An IBM SVC 4.2 cluster recorded an amazingly high throughput rate of 7.084 GB/sec (7,084 MB/sec).
- A helpful price/performance metric based on published list prices is included in all SPC results. Note that IBM delivered the industry leading disk-based SPC-1 performance result at a cost effective price of \$12.06 per database I/O operation per second.
- ESG Lab is convinced that an SVC configured in front of a pool of modular storage arrays can be used to deliver aggregate performance that is on par with or above that of a single enterprise-class storage controller.

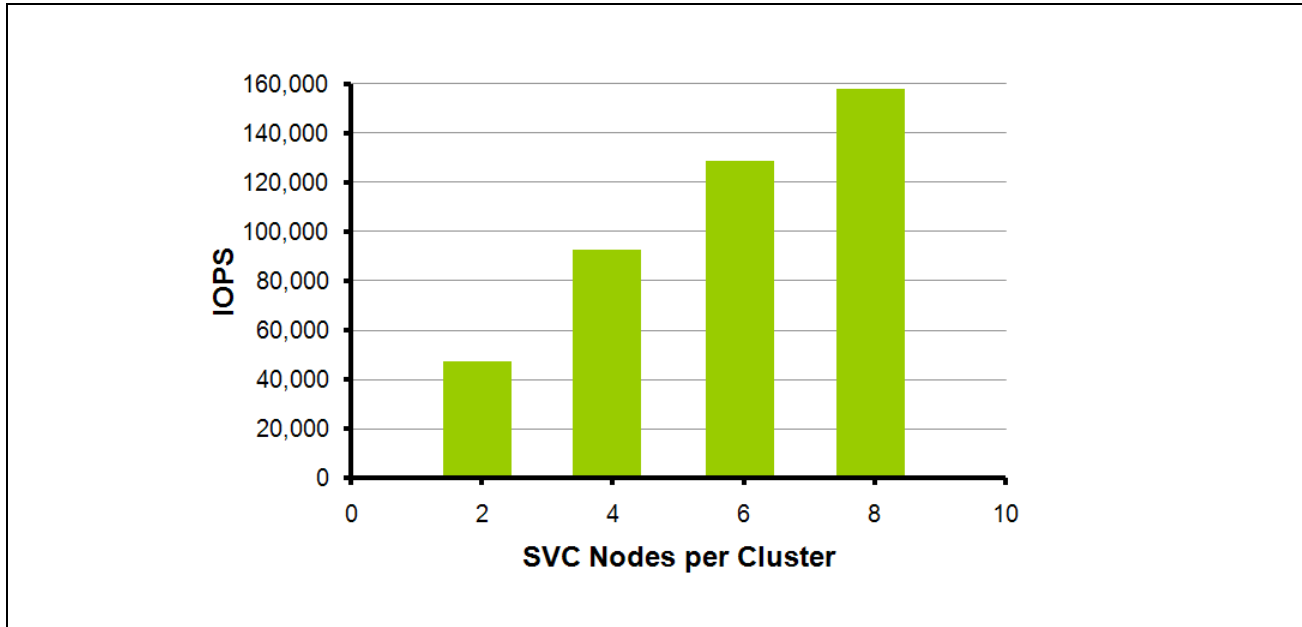
For a deeper look into SVC performance scalability, ESG Lab examined performance results for a two, four, six and eight node SVC cluster running SVC version 3.1. As shown in Figure 13, performance scaled in a near-linear fashion and topped out at nearly 160,000 I/Os per second. Note that the eight node result here is roughly the same as the 155,519 SPC-1 IOPS result published using the same version of SVC (3.1).

⁶ SPC methodology and results can be found at <http://www.storageperformance.org/>

⁷ The top SPC-1 result as of the writing of this report is held by Texas Memory systems for a solid state disk array with an ASU capacity of 137 GB compared to tens of thousands of GB for the disk-based results shown in Table 1). EMC has chosen not to publish SPC results.

SVC performance has continuously improved since the SVC version 3.1 results shown in Figure 13 were collected as evidenced in the SPC-1 results shown in Table 1. As a result, ESG is confident that the scale of the graph shown in Figure 13 would be significantly higher for SVC version 4.3. ESG Lab is also confident that performance would scale in a near-linear fashion as nodes are added to an SVC 4.3 cluster.

FIGURE 13. SVC PERFORMANCE SCALABILITY (DATABASE WORKLOAD)



The latency, caching, striping and scaling effects presented above are an interesting academic exercise, but what matters more is real-world application performance as reported by SVC users. ESG Lab spoke with an end-user who has no issues at all with SVC performance. As a matter of fact, the customer stated that he was able to obtain better overall performance by moving some applications from a Tier 1 storage array to a more cost-effective Tier 2 storage array due to the newfound ability to dynamically add drives to a striped managed disk group to meet performance needs on demand. With well over 10,000 customers running IBM SVC in production, it is clear that SVC is meeting the performance needs of a wide variety of users.

Why This Matters

The benefits of storage virtualization increase dramatically as the size of the shared storage pool grows. With many applications running in parallel to a common pool of shared SAN attached storage, customers can be assured that an SVC virtualization layer inserted between applications and their storage can scale to meet peak performance demands.

ESG has confirmed that performance scales in a predictable linear fashion as the number of nodes in an SVC cluster grows. As a matter of fact, ESG has confirmed that the IBM SVC has been used to record industry leading SPC-1 and SPC-2 performance ratings with excellent results of 271,505 IOPS and 7.1 GB/sec, respectively.⁸

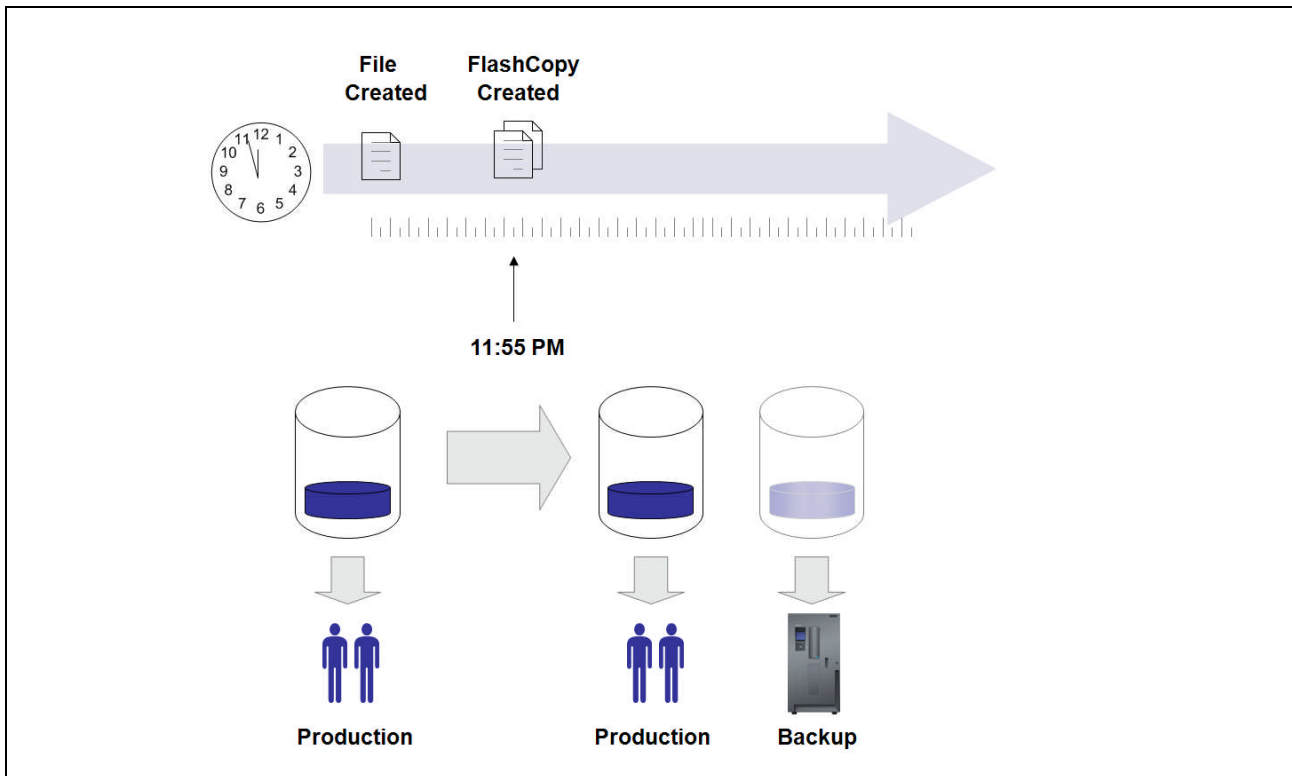
⁸ As shown in Tables 1 and 2. See footnote 7.

Flash Copy

The FlashCopy utility has been embraced by IBM customers for years. A FlashCopy replica is a point-in-time copy of an application volume that is created nearly instantaneously. These copies are instantly available and can be created while applications are up and running. FlashCopy provides a powerful means for quickly making images of application data for backups, application testing and data movement. Multiple volumes involved in an application (e.g., a database table space volume and its log volume) can be placed in a consistency group before being copied. Because FlashCopy runs at the block level, it is necessary to flush cache and OS buffers before executing a FlashCopy in order to ensure consistency at the application level.

One of the uses of FlashCopy is to decrease the time required to obtain a backup of a production database. Without FlashCopy, the application must be shut down to obtain a consistent usable backup of the database application. For a very large database, it can take hours, or possibly days, for the backup to complete. The database typically remains unusable during the entire duration of the backup. In contrast, FlashCopy shrinks the time that the database is unavailable from hours to minutes. As shown in Figure 14, this is accomplished by making a virtual, consistent copy of the volume that is frozen at a point in time. A space-efficient variant of this technology, called Space-Efficient FlashCopy, was introduced in SVC version 4.3.

FIGURE 14. USING A SPACE EFFICIENT FLASHCOPY FOR BACKUP

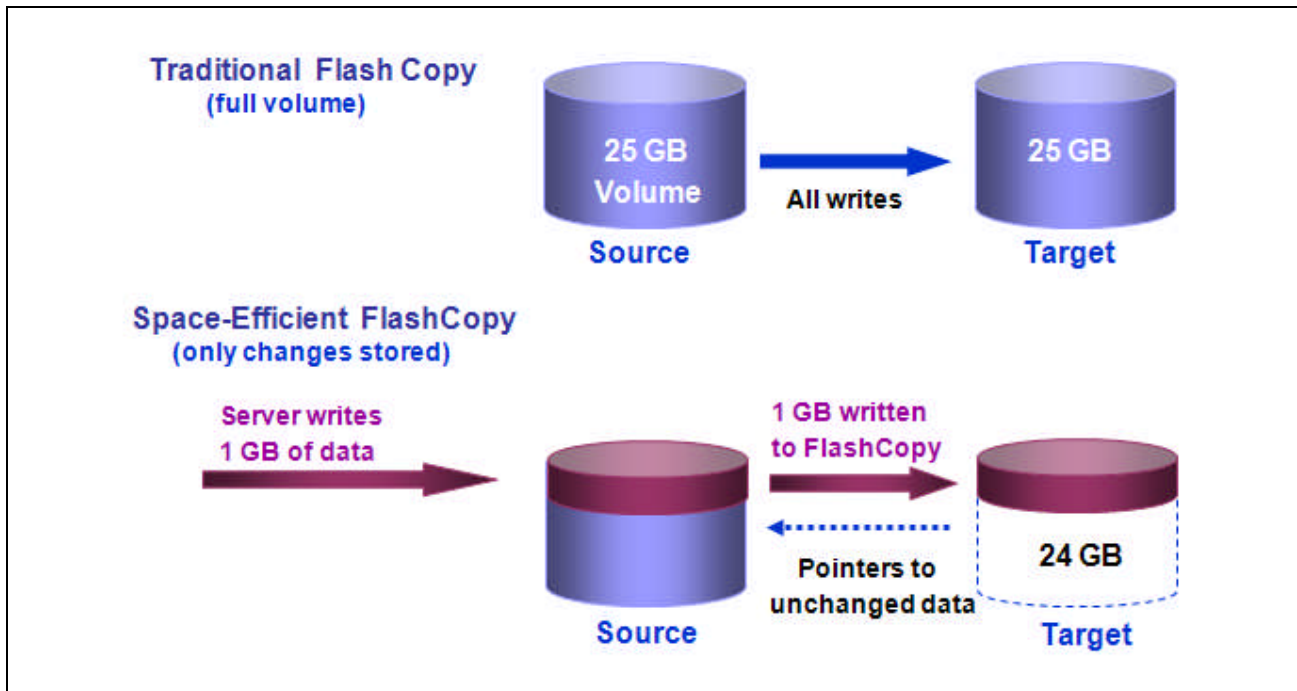


ESG Lab Testing

ESG Lab used SVC console to create a space-efficient FlashCopy replica of a 25 GB VDisk. The traditional and space efficient methods introduced in SVC version 4.3 were tested.

As shown in Figure 15, a traditional Flash Copy replica allocates space for a full copy. In this example, 25 GB of DS8300 capacity was reserved for a copy of a VDisk accessed by an IBM System p server running the AIX operating system. The traditional FlashCopy was used to simulate an application test and development workflow. Write-intensive test and development workflows are excellent candidates for a traditional flash copy.

FIGURE 15. FLASH COPY IN ACTION



A utility that endlessly created a file per second was started on the primary volume. After the FlashCopy replica had been created, the target VDisk was mapped to another logical partition on the server and accessed. No interruption to application access was noticed as the FlashCopy replica was created. As expected, the target VDisk was frozen at the exact point in time when the FlashCopy command had been executed.

The experiment was repeated using the Space-Efficient FlashCopy feature introduced in SVC version 4.3. Space-Efficient FlashCopy replicas are ideally suited for a read-intensive backup job. As shown in Figure 15, aside from a bit of space needed to maintain pointers to un-changed data, the only space consumed was a record of writes completed since the FlashCopy command was executed. In this example, after 1 GB of data of new data had been written by the test utility, 24 GB of target VDisk capacity was saved compared to a traditional FlashCopy.

ESG Lab spoke with a customer who has found some interesting and powerful uses for FlashCopy. FlashCopy is used as an insurance policy when backups fail. When a backup failure is noticed in the morning, the volumes that failed last night's backup are copied. In this manner, the storage administrator can perform restores using the disk-based FlashCopy replica if needed while the backup problem is being fixed. This customer also uses FlashCopy before major system and application upgrades. If something goes wrong during the upgrade, the FlashCopy replica provides a disk-based point-in-time image from just before the upgrade for quick recovery.

Why This Matters

Creating consistent copies of application data while applications are up and running has been a problem that data managers have been struggling with for years. Copies are needed for a variety of reasons including backup and recovery and application testing. Disk array vendors have been providing point-in-time functionality to address this problem in recent years. Instead of running PiT functions on many storage arrays, ESG Lab has confirmed that the FlashCopy function of SVC can be used as a centralized platform for making full volume and space-efficient point in time copies.

Metro and Global Mirror

The world in which we live in comes with harsh realities that can threaten our businesses. Threats to the data center include flood and fire, as well as other natural disasters that can have a devastating impact. Additionally, events such as major power outages and even acts of terrorism occur throughout the world today. In today's business environment, organizations must also protect themselves against litigation and corporate investigations, and comply with federal regulations. A growing number of IT professionals use mirroring as a way to protect their data and recover from such disasters. The Metro and Global Mirror functions of SVC provide a centralized utility for remote mirroring between two or more storage systems from one or more vendors.

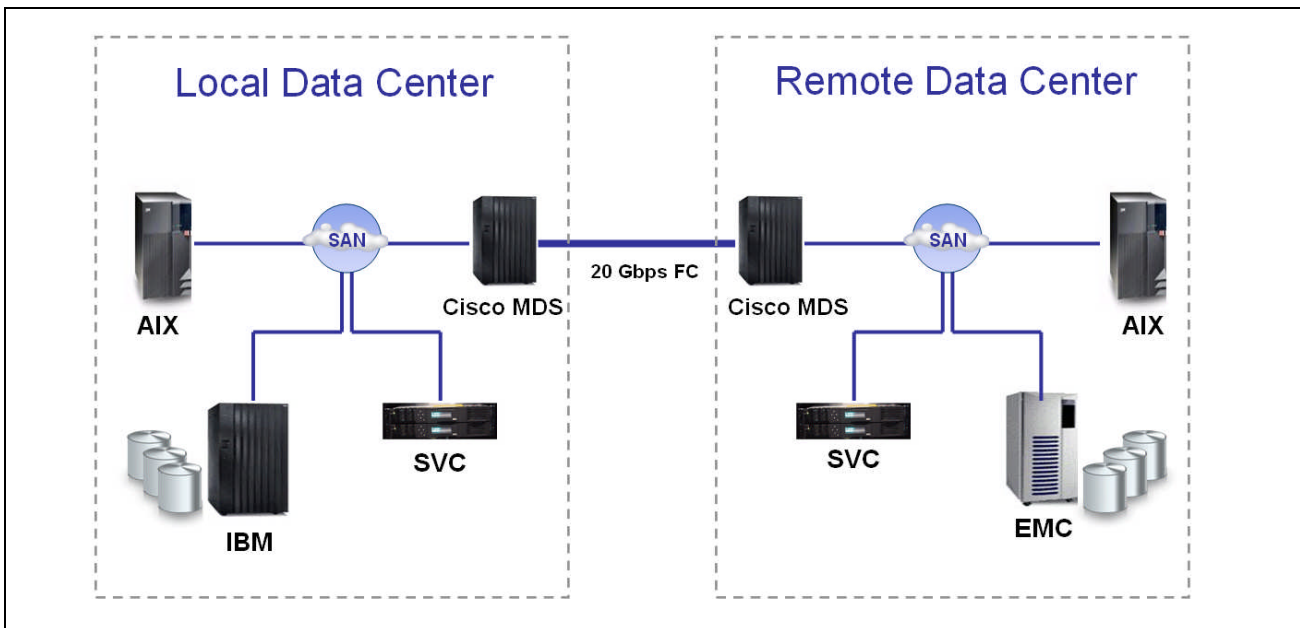
Metro and Global Mirror are used to make a remote copy of application data within one SVC-managed cluster to another cluster for disaster avoidance. Typically, the copies are made between data centers. If the primary data center fails, then the application can continue using a copy at the remote site.

The Metro Mirror function provides a synchronous real-time copy of the data and is typically used over relatively short distances (e.g., within a campus or a metropolitan region). The Global mirror function provides an asynchronous copy of application data and is typically deployed over longer distances (e.g., between two geographic regions). A Metro Mirror copy is always up to date, complete and instantly usable. A Global Mirror copy runs slightly behind the primary data center image and provides the ability to restart an application using a consistent time-ordered image of application data—minus the last few transactions that may have been caught in the synchronization process between data centers.

ESG Lab Testing

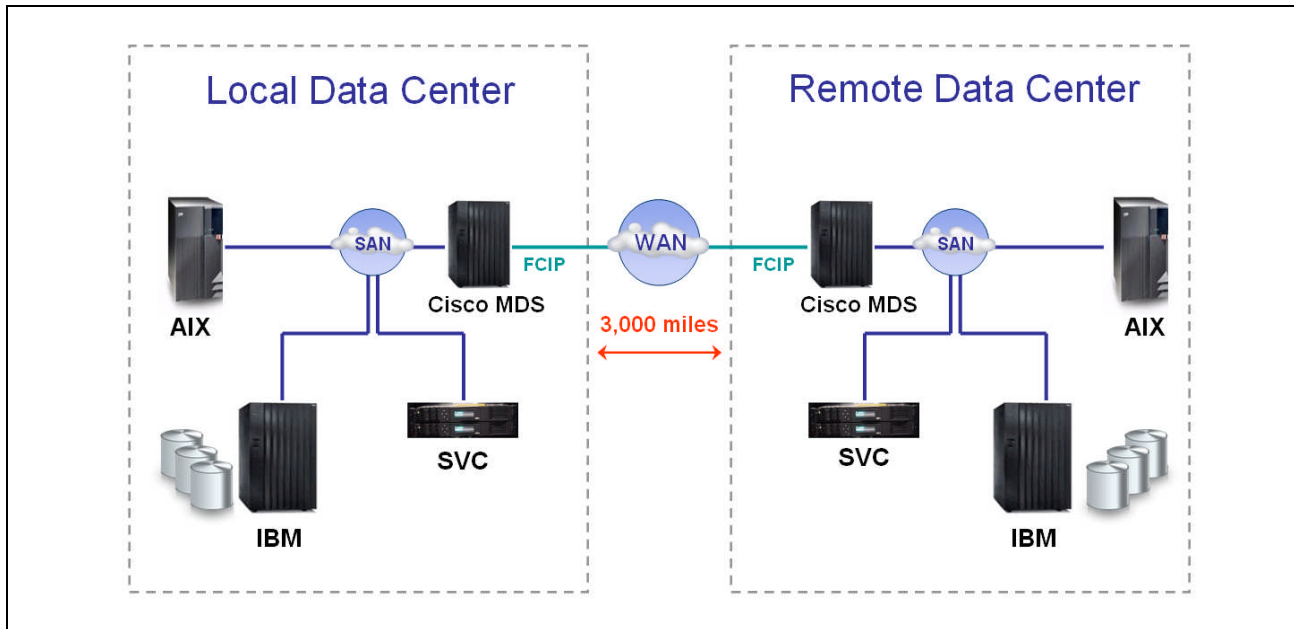
The configuration used for Metro Mirror testing is shown in Figure 16. Each of the SVC clusters being mirrored was located within the same building at IBM, but on different floors. Servers running the AIX operating system were used to access data at the local and remote sites. Data residing on an IBM storage array at the local site was mirrored to an EMC storage array at the remote site. A 20 Gbps FC trunk configured between Cisco MDS storage switches at each site was used as a high-speed data replication pathway between the sites.

FIGURE 16. THE METRO MIRROR TEST CONFIGURATION



The SVC console was used to configure and manage Metro Mirror during both a planned and an un-planned outage. The unplanned outage was simulated by breaking the FC connection between the SVC at the primary site and its storage array. A script, which created time-stamped files in a loop, was used to verify that the mirroring process had no impact on application performance and to verify that each and every file written on the local SVC cluster was copied and readable on the remote SVC cluster.

FIGURE 17. THE GLOBAL MIRROR TEST CONFIGURATION



A similar test bed was used for Global Mirror testing, as shown in Figure 17. Instead of the high speed FC link used between the two sites during the Metro Mirror validation, an Emperix WAN simulator was used to simulate a DS3 link between two data centers located 3,000 miles apart. An FCIP connection within the Cisco MDS switches at each site was used to connect to the simulated WAN. The SVC console was used to monitor the synchronization process. Files created before the mirroring operation had begun were validated by inspection after the SVC console indicated that the mirroring operation had ended with a state of "Consistent Synchronized."

Why This Matters

Business critical applications with non-stop availability requirements need to be able to survive a data center failure due to potential disasters including power failure, flood, fire and acts of terrorism. Traditionally, the remote mirroring capability of storage arrays had been used to recover quickly from such disasters. Array-based replication solutions differ in the way that they are configured and used, and most require the same storage systems on each end of the wire.

ESG Lab confirmed that SVC Metro and Global Mirror provide a single solution which can be used to implement disaster tolerant applications running on different servers, operating systems and storage systems. As an option, SVC customers can deploy a smaller more affordable storage system, or an old storage system they already own, at the remote site to significantly reduce the total cost of disaster avoidance.

Virtual Disk Mirror

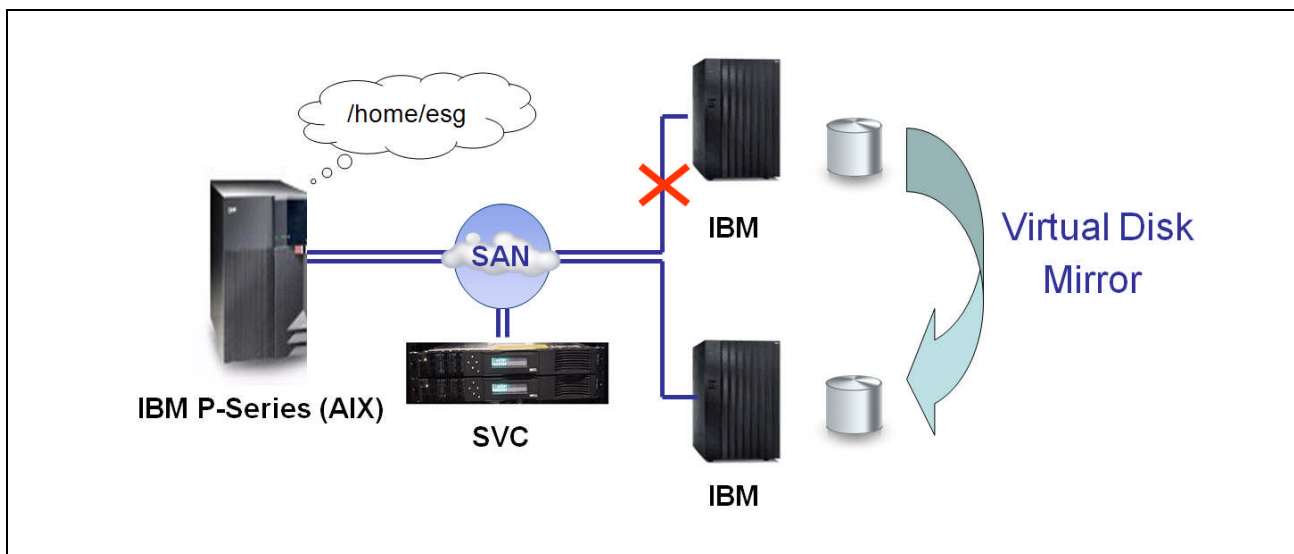
Virtual Disk mirror support, introduced in SVC version 4.3, is used to replicate application data between two disk arrays in the same data center. Virtual Disk Mirroring is used to survive a catastrophic disk array or SAN fabric failure. This feature is ideally suited for mission critical applications where an application outage would have a devastating impact on the business.

ESG Lab Testing

The test bed used to validate SVC version 4.3 Virtual Disk Mirroring is shown in Figure 18. A home directory on an IBM System p server running AIX version 4.3 was mirrored across two enterprise-class IBM DS8300 disk arrays. A test utility which performed update operations at a rate of four I/Os per second was used to simulate the write traffic emanating from a mission critical online transaction database application. The test utility was started before the SVC disk mirror operation was executed.

A host-based AIX utility was used to monitor response time during the mirror synchronization and a simulated disk array failure. An average virtual disk response time of 2 milliseconds was maintained throughout the mirror synchronization operation. After the synchronization was complete, a disk array failure was simulated by disabling the SAN fabric zone connecting one of the disk arrays. The test application continued without error during the failover. Response time increased to 16 milliseconds for approximately two seconds during the failover.

FIGURE 18. TESTING SVC 4.3 VIRTUAL DISK MIRRORING



Why This Matters

Mirroring between enterprise-class disk arrays has been embraced by organizations with extreme high availability needs for more than a decade. SVC provides a centralized, cost effective alternative with a common user interface and command line interface—regardless of the vendor, type or model of disk array being mirrored. ESG Lab confirmed that Virtual Disk Mirroring, introduced in SVC version 4.3, can be used to transparently survive a catastrophic disk array failure with no perceivable application performance impact.

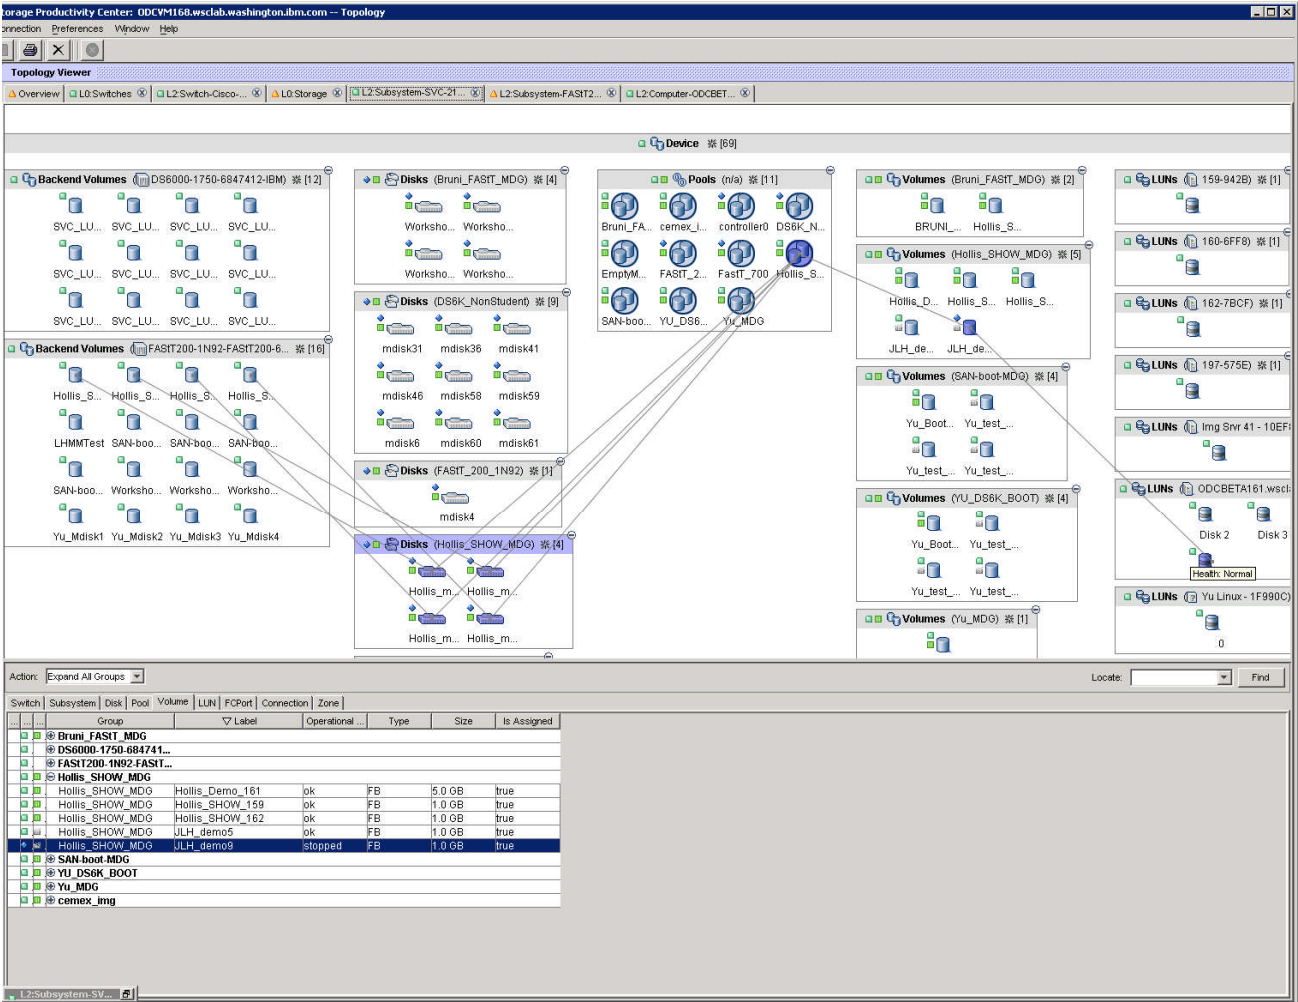
Advanced Manageability

IBM TotalStorage Productivity Center software simplifies and automates the management of storage infrastructure through a centralized, common suite of tools for the management of storage devices, storage networks and data. The Productivity Center software suite for storage is comprised of four licensed components, each of which is fully integrated with SVC: Productivity Center for Data, Productivity Center for Fabric, Productivity Center for Disk and Productivity Center for Replication. Productivity Center for Data uses server resident agents to track file system and application storage capacity utilization—a valuable storage planning metric that cannot be provided by block-based storage and virtualization solutions.

ESG Lab Testing

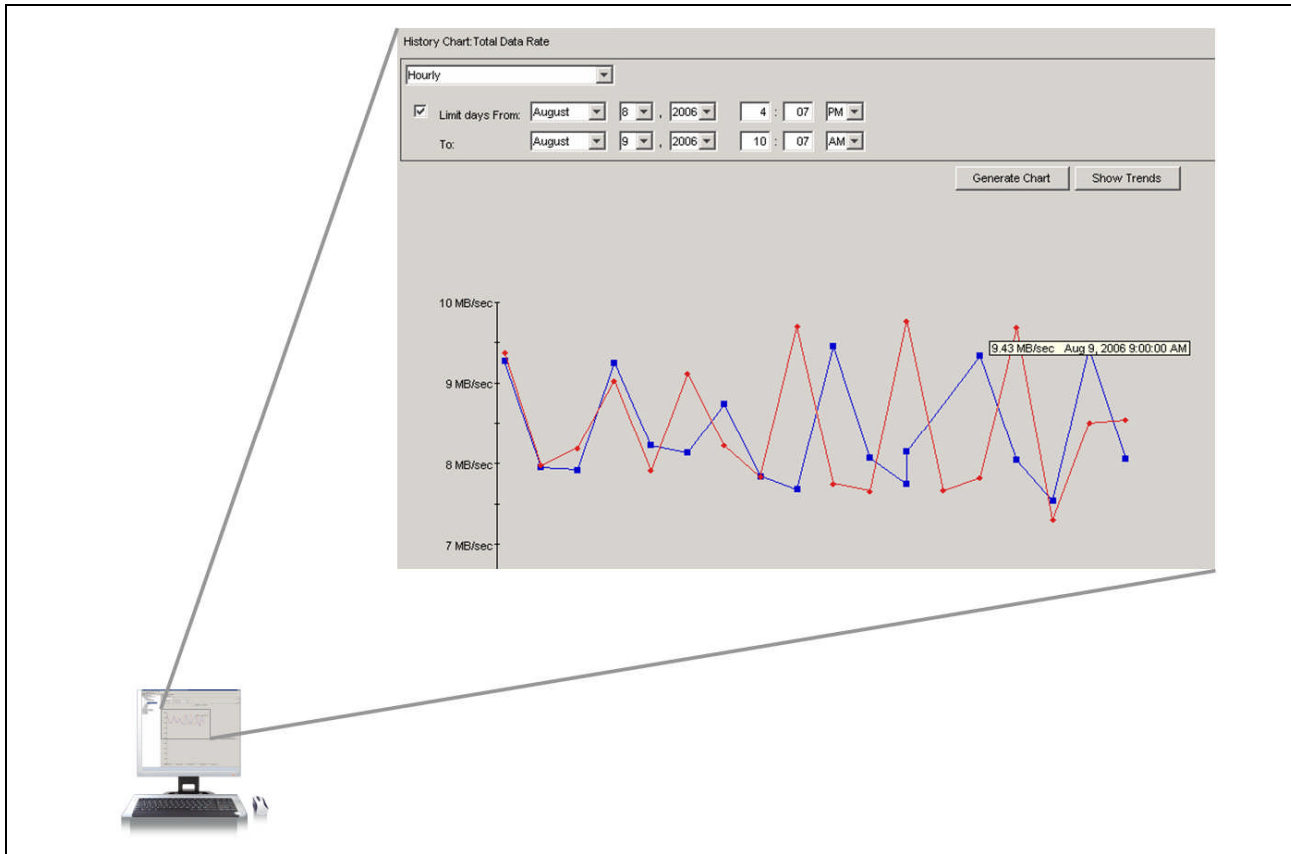
ESG used Productivity Center to examine an end-to-end topology view of an SVC-managed storage environment, as shown in Figure 19. Reading from left to right, ESG noted that Productivity Center presents a clear and easy to read picture of the relationship between **Backend Volumes** (four physical LUNs within an IBM FASTT200 array), managed disks (**Disks**), managed disk groups (**Pools**) and VDisks (**Volumes**) presented to servers (**LUNs**). ESG Lab also noted that Productivity Center graphically presents the health of volumes for quick and easy problem isolation.

FIGURE 19. AN END-TO-END TOPOLOGY VIEW OF SVC MANAGED STORAGE



ESG Lab used Productivity Center to examine the historical performance of SVC-managed storage as shown in Figure 20.

FIGURE 20. HISTORICAL PERFORMANCE DATA FOR SVC MANAGED STORAGE



ESG Lab spoke with a customer who has deployed the entire Productivity Center storage suite. Besides the administrator productivity savings that have been realized, Productivity Center has become a valued tool for the justification of storage infrastructure upgrades.

Why This Matters

An end-to-end view of storage infrastructure, and the data that resides within that infrastructure, is needed to predict and avoid space and performance problems before they arise, optimize storage resource utilization, identify and retire orphan data, forecast future growth and enhance the productivity of storage administrators. ESG Lab confirmed that IBM TotalStorage Productivity Center, working in concert with SVC, provides a centralized and automated interface for the management of a heterogeneous storage network so that administrators can make better and timelier decisions.

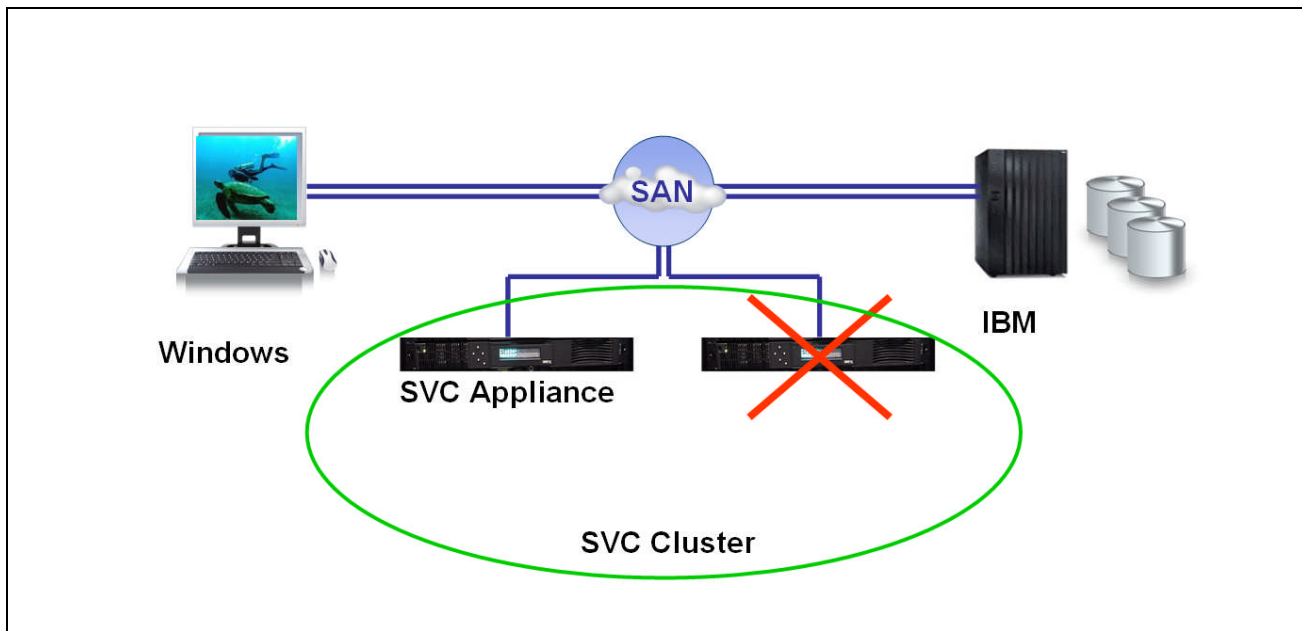
Fault Tolerance

SVC appliances are configured in fault tolerant pairs, with each appliance providing failover services for another appliance in the cluster—much like a buddy system. Up to eight appliances, configured in fault tolerant pairs, are arranged in a cluster and managed as a single system.

ESG Lab Testing

ESG Lab killed the power on one of the SVC appliances in a dual-node cluster as a Windows server played a video residing on an SVC managed volume, as shown in Figure 21.⁹ The failure was noted on the SVC console as the video played without interruption.

FIGURE 21. ESG LAB FAULT TOLERANCE TESTING



Why This Matters

Because SVC appliances are logically configured in-band between servers and storage, transparent fault tolerance is a crucial requirement when deploying applications that rely on SVC-managed storage. ESG Lab verified that SVC appliance failover is transparent and automatic.

⁹ The SVC console was used to determine which SVC appliance should be powered down before the error was injected. In other words, ESG Lab made sure that the SVC appliance serving the video was the one powered off.

ESG Lab Validation Highlights

- ☑ ESG Lab found that SVC can be easily introduced into (and removed from, if desired) an existing application environment.
- ☑ ESG Lab confirmed that the SVC console has a very consistent look and feel and includes wizards for commonly performed tasks.
- ☑ With virtually no supervision, ESG Lab completed a number of valuable storage management tasks in under an hour using multiple storage arrays from multiple vendors connected to a variety of operating systems. Managed disk groups were created, VDisks were expanded and deleted, a VDisk was migrated, a VDisk was copied, tasks were monitored and logs were viewed—all while applications remained online and available.
- ☑ ESG Lab confirmed that the SVC Space-Efficient Virtual Disks function can be used to reduce the cost and complexity of storage provisioning while providing significant capacity savings. Threshold monitoring capabilities at the virtual disk and managed disk pool levels were exercised.
- ☑ ESG Lab testing has confirmed that SVC was used to perform online migrations within and between storage arrays while applications were running.
- ☑ A copy of a volume was mounted and accessed from another server using FlashCopy. The copied volume was available immediately, even as the copy ran in the background on an SVC appliance. The capacity saving capabilities of the Space-Efficient FlashCopy feature introduced in SVC version 4.3 were verified.
- ☑ Metro Mirroring was used to replicate SVC-managed storage between two data centers within the same campus and Global Mirroring was used to replicate between two data centers connected via a simulated DS3 WAN link of 3,000 miles.
- ☑ Virtual Disk Mirroring was used to replicate a virtual disk in real time between two IBM DS8300 disk arrays. A test application ran without interruption as a catastrophic disk array failure we emulated.
- ☑ TPC was used for an end-to-end topology view of an SVC-managed storage environment and a view of the historical performance of SVC-managed storage.
- ☑ An SVC appliance managing the storage being used to view a video from a Windows server was powered off to verify the fault tolerance of an in-band SVC solution.

Issues to Consider

- ☑ *One to Many Mirroring:* Metro and/or Global Mirroring from one data center to multiple remote data centers was not supported when this report was published. In particular, the ability to mirror a production volume to a nearby data center synchronously while simultaneously replicating to a third data center in another geography asynchronously would be powerful. This capability could be used to ensure the very highest level of disaster recovery services for the most mission-critical applications.
- ☑ *iSCSI Support:* IP-based SANS, which use the iSCSI protocol to leverage commodity Ethernet technology, have emerged in recent years as a viable, cost-effective alternative to FC—especially for the growing number of small to medium sized businesses adopting server virtualization technology and iSCSI for their first SAN. Given the complementary benefits of deploying server virtualization with storage virtualization, ESG believes that SVC support for iSCSI attached disk arrays would be a valuable SVC enhancement.

ESG Lab's View

Processes are largely broken in IT today due to the fact that data growth has continued unabated. Innovation and automation are needed to reign in the cost and complexity associated with unrelenting growth and ever-changing business requirements. Virtualization technologies are needed to simplify, consolidate and automate routine IT functions.

Innovative virtualization technologies are being deployed by a growing number of IT organizations. A recent ESG survey of 602 IT professionals indicates that server virtualization is leading the charge. Thirty-three percent of respondents report that server virtualization is in "wide-spread production use."¹⁰ Server virtualization, like storage virtualization, is powerful game-changing technology that can be used to consolidate and simplify complex IT infrastructure. While the cost savings of consolidation drove the first wave of server virtualization adoption, the second wave is being driven by the mobility, flexibility and enhanced fault tolerance that can be achieved with a networked virtualized IT infrastructure.

Like the early adopters of server virtualization, early adopters of network-based storage virtualization report that management complexity has been simplified and costs have been reduced. Given the synergistic benefits, it is not surprising that ESG research indicates that a growing number of forward-looking organizations are deploying storage and server virtualization together. Twenty-four percent have deployed both already and 33% plan on doing so within the next 24 months.

With more than 13,000 appliances shipped since the introduction of SVC in 2003, IBM is the leading storage virtualization vendor in the market. ESG has confirmed that SVC simplifies storage provisioning as it provides a centralized platform for valuable data services including online data migration, copy services, mirroring and remote replication. For those who might be concerned about the potential performance impact of a storage virtualization appliance, ESG Lab is pleased to report that SVC not only minimizes performance overhead to an extent that can't be detected by applications and end-users, it can actually improve performance using field-proven caching techniques. For those who might be concerned about enterprise-class fault tolerance of SVC, ESG Lab has confirmed that pairs of SVC appliances use the buddy system to ride through a catastrophic hardware failure with ease.

ESG Lab recently tested the enhanced utilization and availability capabilities of SVC version 4.3—the most powerful being the Space-Efficient Virtual Disk function. Presenting virtual disk capacity to servers as SVC allocates physical capacity on a just in time basis, space-efficient virtual disks eliminate the cost of unused and unallocated storage as they simplify the task of assigning storage capacity to applications. SVC provides a centrally managed platform for deploying this valuable emerging technology which is generally referred to as Thin Provisioning. ESG Lab was impressed by advanced SVC capabilities which ESG likes to refer to as Thin Provisioning 2.0. Thin Provisioning 2.0 capabilities verified by ESG Lab include wizard-driven configurability, programmable thresholds and alerting at the volume and pool level, optional automatic volume extension and the ability to define the capacity associated with each unit of allocated capacity.

The bottom line with SVC is its ability to reduce complexity and costs. One SVC user with whom ESG spoke said that he "bet the farm" on SVC storage virtualization and that he is thrilled with the results. With one less storage administrator, he is managing three times more storage capacity. Using a tiered storage approach that he refers to as tier "one and a half," he's reduced his storage capital equipment budget by \$1M over 18 months by moving applications from tier one to tier two storage.

SVC is a mature, interoperable and rock-solid storage virtualization solution. For customers shopping for their next storage infrastructure upgrade, especially those considering IBM storage, SVC is a no-brainer. Try it for the migrations that you know you need to get done, and ESG believes you'll like it for your newfound ability to turn an out-of-control heterogeneous storage infrastructure into a centrally managed on-demand storage utility.

¹⁰ Source: ESG Research Report, *IT and Service Management Process and Automation Priorities*, March 2008

Appendix

TABLE 3. TEST CONFIGURATION

Server/OS	HBA/driver	Multi-path	Switch	Storage
Win2K	QLogic 2312/ 8.2.0.10 (bios 1.29)	IBM SDD 1.6.0.7	Cisco MDS	IBM DS6000 EMC SYMM IBM ESS F20
VMware ESX 2.5.3 Build 22981	Qlogic 231x/2340 / qla2300_607.o (firmware 3.03.06)	Native VMware	Cisco MDS	IBM DS6000 EMC SYMM ESS F20
Win2003 cluster	QLogic 2340/9.1.2.11 (bios 1.43)	IBM SDD 1.6.1.0	Cisco MDS	IBM DS4300 IBM DS6000 IBM DS8300
Win2K	QLogic 2312/ 8.2.0.10 (bios 1.29)	IBM SDD 1.6.0.7	Cisco MDS	IBM DS4300 IBM DS6000 IBM DS8300
Linux Red Hat Enterprise AS 4 Kernel 2.6.9-22	QLogic 2300 /8.01.00B5 (Bios 1.17)	IBM SDD 1.6.1.0	Cisco MDS	IBM DS4300 IBM DS6000 IBM DS8300
AIX	5759/ 5.3.0.31	IBM SDD 1.6.1.0	Cisco MDS	IBM DS4300 IBM DS6000 IBM DS8300
SUN450 Solaris9	QLogic	IBM SDD 1.6.1.0	Cisco MDS	IBM DS6000 EMC SYMM IBM ESS F20
AIX	5759/ 5.3.0.31	IBM SDD 1.6.1.0	Cisco MDS	IBM DS6000 EMC SYMM IBM ESS F20



20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com