Lab Validation Report

Dell Compellent Storage Center 5.4

Fluid Networked Storage

By Jason Buffington with Brian Garrett

November 2011

Contents

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 Dell.

Introduction

Storage Center from Dell Compellent was designed to provide *efficient, agile, and resilient* networked storage with a goal of reducing storage cost and complexity while meeting the growing needs of the business. This report documents hands-on testing of Storage Center version 5.4 with a focus on the capacity, recoverability, and performance efficiency of the Compellent Dynamic Block Architecture. Testing covered Compellent's approach to multi-tiered storage, referred to as Data Progression, as well as its deep integration and manageability in VMware virtual server environments and automated migration of virtual machines between sites with the recently-added Live Volume feature. One of the key goals of this report is to highlight the capabilities of Compellent Data Progression as compared to emerging sub-LUN tiering solutions from other leading storage vendors.

Background

IT departments continue to struggle with the costs of storage, which includes not only the cost of acquisition but also the ongoing operational expenses associated with storage management. Growing demand for increased capacity and the need to provide assured availability and flexible recovery scenarios are challenges as well. As shown in Figure 1, a recent ESG survey indicates that rising costs, rapid data growth, and data protection are at the top of the list of challenges associated with meeting the storage needs of applications.[1]

Figure 1. Storage Challenges Associated with Supporting Applications

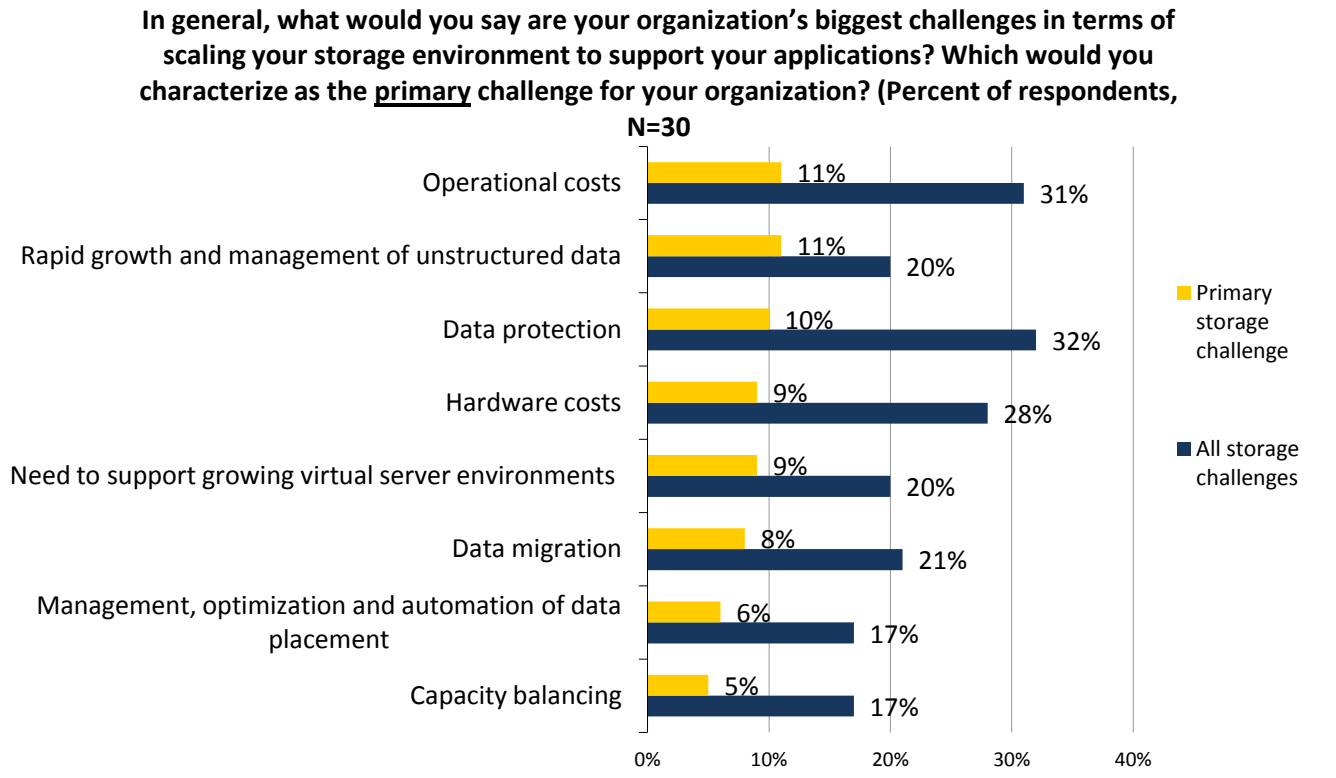

Source: Enterprise Strategy Group, 2010.

Keeping up with the growing storage needs of virtual server environments and the cost and complexity of data migrations are critical concerns as well. And last, but not least, IT managers are looking for products and tools which can automate the management, optimization, and capacity balancing of storage resources associated with ever-growing information assets.

[1] Source: ESG Research Report, *Scale-out Storage Market Trends*, December 2010.

A storage solution with an architecture that dynamically automates data placement and protection at the block level is needed to address each of these challenges as it stores and protects application data using the right tier of storage, at the right time, at the right price.

Storage Center 5.4

Storage Center is a SAN attached storage solution that uses a combination of intelligent software, a pair of clustered storage controllers built from industry-standard server hardware, host interface controllers (e.g., Fibre Channel, iSCSI), and one or more drive enclosures to store and protect an organization's data. Drive enclosures can be populated with multiple tiers of storage to meet a variety of price and performance requirements (e.g., SAS and SATA), with the ability to "mix and match" storage technologies and interfaces in a modular fashion.

Along with a fourth generation controller boasting a 25% IO performance increase and more controller memory and PCI-e slots, Storage Center 5.4 also offers support for 6 Gbps SAS drives, 2.5" SAS enclosures for higher density, and support for the latest host connectivity options (10 GigE iSCSI and FCoE).

Compellent, which was recently acquired by Dell, was founded in 2002 by a visionary team of storage industry veterans. Storage Center was architected from the ground up with a powerful page-based architecture that automates the movement and protection of data at the block level data using capacity residing on different tiers of storage. As shown in Figure 2, policies are used to automatically use higher performing "tier-1" storage for write data and frequently accessed read data. Dell's *Data Progression* functionality automatically migrates infrequently accessed blocks of data to a more cost effective tier of near-line storage. Replays (a.k.a., snapshots) are automatically placed on lower tier(s) of storage for cost effective recovery.

Figure 2. Dell's Dynamic Block Architecture Enables "Fluid Data"

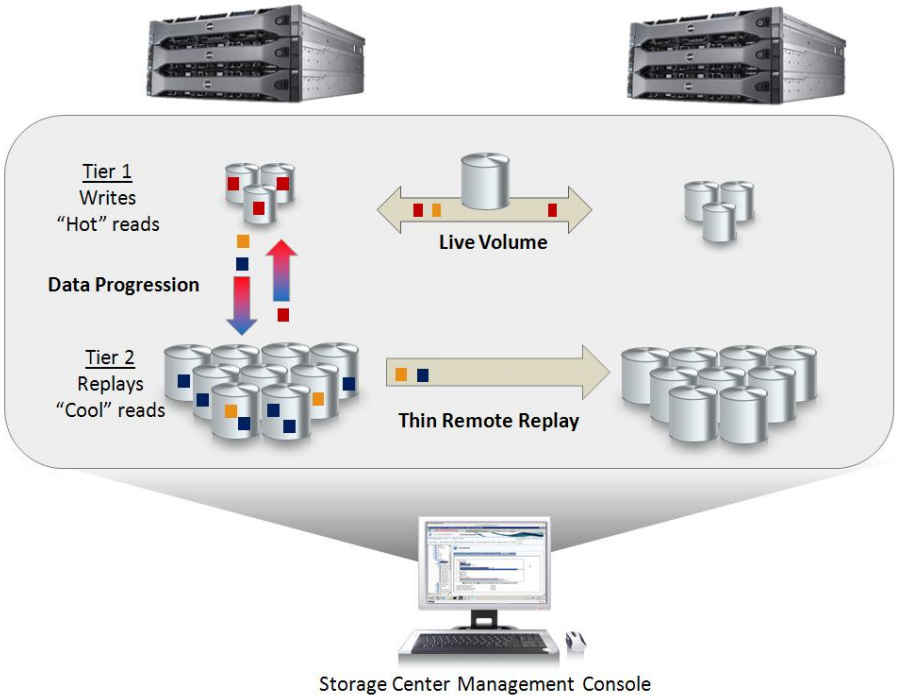

Since Compellent first began shipping systems with its innovative Data Progression feature in 2003, a number of leading storage vendors have begun to adopt a similar approach with a technology that has since become known as sub-LUN tiering. At that time, the Dynamic Block Architecture was also used to reduce the cost of disk capacity via a thin provisioning feature which Dell refers to as Dynamic Capacity. Once again, Compellent was one of the first to introduce what may be better known by many as thin provisioning. The Dynamic Block Architecture also enables a number of powerful capabilities including WAN efficient replication of recovery data between sites (Remote

Replays) and Live Volume which create a single view of application data residing on a pair of Compellent Storage Center solutions deployed within the same—or different—data centers.

ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of Storage Center version 5.4 at a Dell facility in Eden Prairie, Minnesota. Testing was intended to validate the Dynamic Block Architecture's delivery of *efficient, agile, and resilient* storage services. ESG Lab looked at several aspects of the Storage Center solution, focusing on:

- Capacity efficiency with *Dynamic Capacity*
- Performance acceleration with *Data Progression*
- Efficient disk-based recoverability with local and remote snapshots (*Remote Instant Replay*)
- Storage and server virtualization with VMware and Microsoft Hyper-V integration
- Advanced agility and recoverability with *Live Volume*

Getting Started

The test bed used during the ESG Lab Validation is shown in Figure 3. A Microsoft Windows physical server and a VMware vSphere virtual server were connected via 8 Gbps Fibre Channel to a Storage Center solution with a mix of SSD, FC, and SAS drives. The test bed specifics can be found in Table 1 of this report's appendix, with nearly 80% being large but lower-performing "tier 3" drives and the rest being tier-2 or tier-1 to quantify the performance benefit after tiering up to faster storage with Data Progression and the efficiency of snapping down to lower cost storage during Replay testing. A similar Storage Center configuration located in a simulated remote site was used during remote Replay and Live Volume testing.

Figure 3. ESG Lab Test Bed

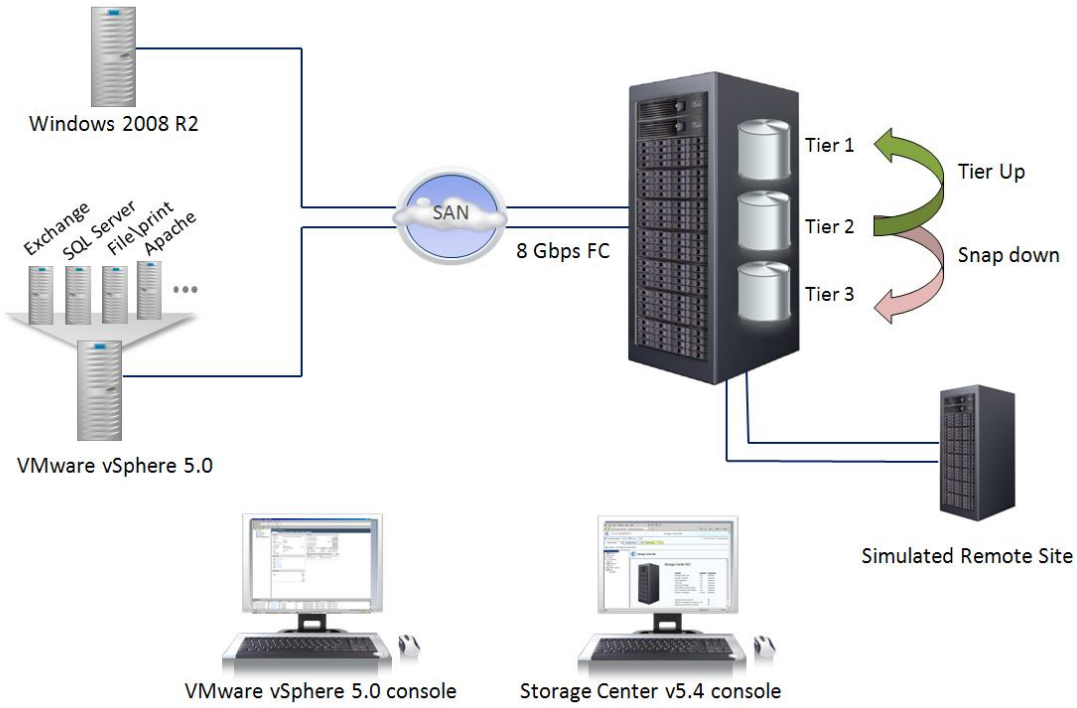

Capacity

Compellent's page-based Dynamic Architecture is used to implement the Dynamic Capacity feature generally known in the industry as thin provisioning. Dynamic Capacity is used to quickly and efficiently allocate storage on demand without consuming blocks that would otherwise be wasted before a server actually used them. As opposed to legacy storage systems that have recently been upgraded to support thin provisioning, Dynamic Capacity was built into the Compellent architecture from the beginning and is the default behavior for storage provisioning.

Figure 4 illustrates the difference between Compellent Dynamic Capacity and traditional provisioning. In this example, a new 250 GB volume has been created and presented to a server. In both cases, the server sees 250 GB of available capacity. With traditional provisioning, that capacity is exclusively dedicated to the server, regardless of when or how much the server actually consumes. On the right, Dynamic Capacity presents the same capacity to the server, but consumes only the blocks that have been written by applications. In this example, a small fraction of the 250 GB in usable capacity has been used. ESG research and conversations with a number of early adopters of thin provisioning indicates that customers have reduced the cost of storage capacity by a factor of 50% or more compared to traditional provisioning method.

Figure 4. Thin Provisioning

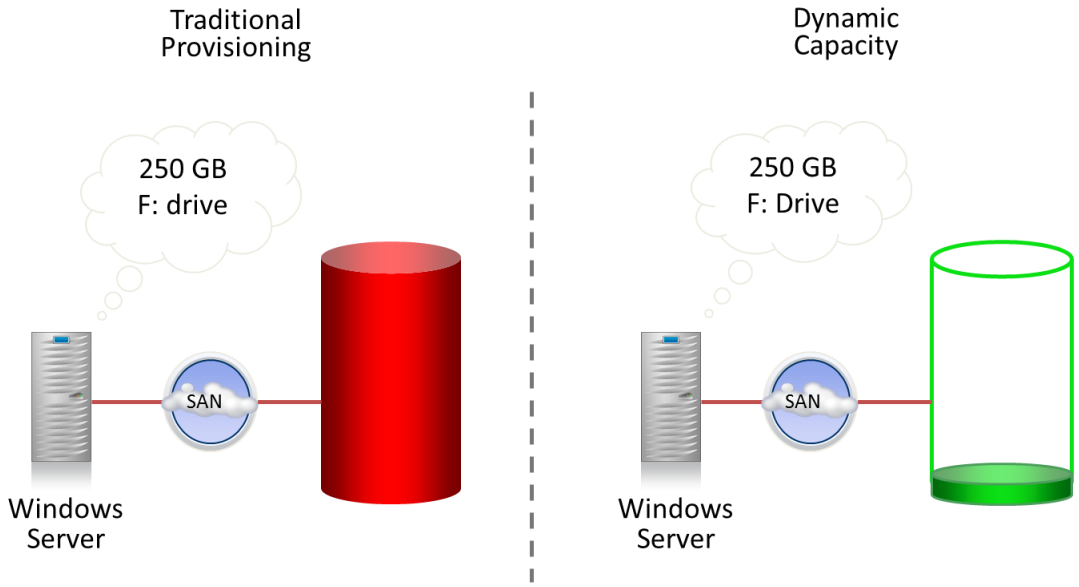

ESG Lab Testing

ESG Lab tested Dynamic Capacity using a simple deployment of file server storage. A 249 GB volume was created using the web-based Storage Center management console and attached to a server running the Windows Server 2008 R2 operating system. A file system was created and a drive letter was assigned with the Microsoft Disk Administrator utility. A Windows Explorer properties view of the drive is shown in Figure 5.

Figure 5. Thin Provisioning: 249 GB volume, as Seen by Windows Explorer

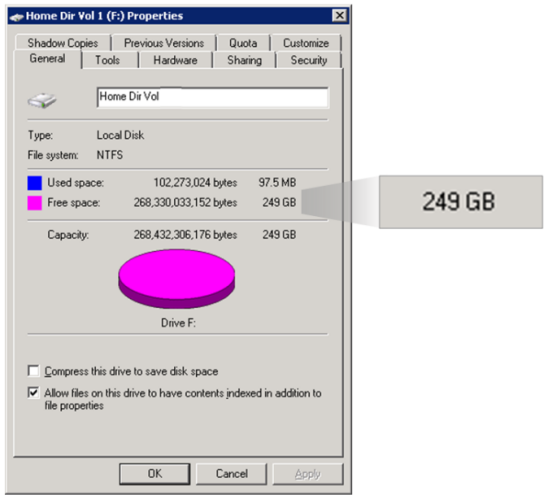

Figure 6 shows that the actual disk space consumed for the 249 GB volume is only 138 MB.

Figure 6. Thin Provisioning: 250 GB volume, as Seen by Storage Center

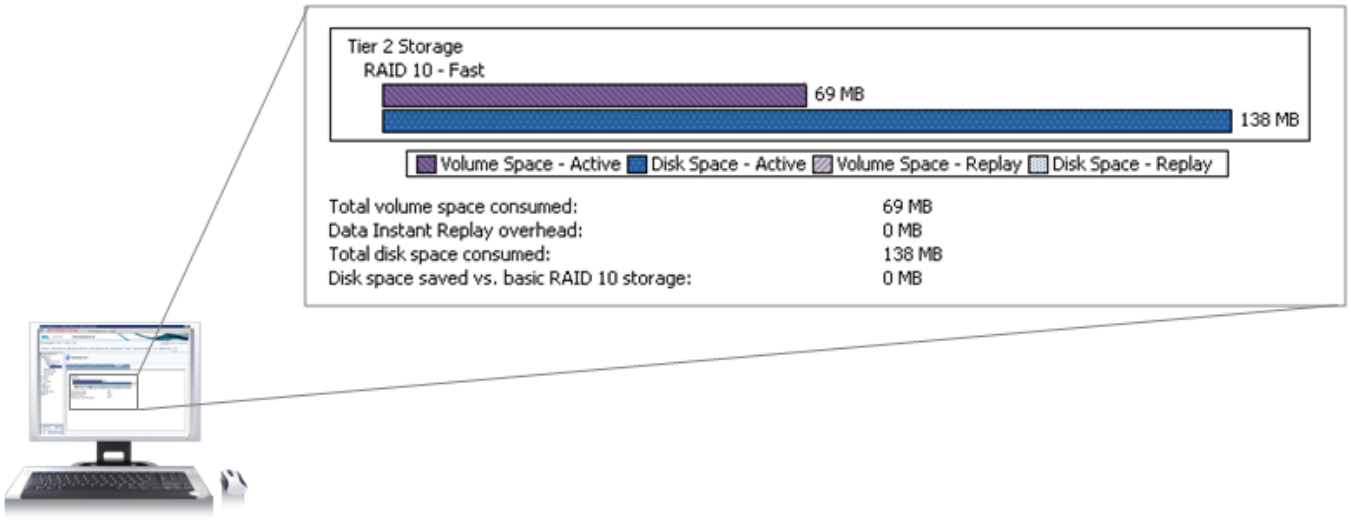

Why This Matters

Managing data growth is a key IT priority as reflected by ESG's *2011 IT Spending Intentions Survey,* where it ranked third in the list of top priorities behind increased use of server virtualization and major application upgrades. As much as data is growing, consumption tends to grow at a higher pace in most environments as storage capacity is pre-allocated before it is consumed.

ESG Lab has confirmed that Dynamic Capacity, built into the Dynamic Block Architecture of Compellent Storage Center since its beginnings, can be used to simply and effectively reduce the cost of storage capacity. Compared to a number of legacy storage solutions that have recently added thin provisioning as a feature, Compellent Dynamic Capacity is a default provisioning method that's extremely easy to configure and manage. Even today, while thin provisioning may require manual activation and sometimes with workload cautions from other vendors, it is the default for Compellent solutions.

Performance

Data Progression leverages the Dynamic Block Architecture to move individual blocks of data to different tiers of storage based on predefined or custom policies. Data Progression is managed at the page level, with each page being composed of one or more blocks of data (the default page size is 2 MB). These pages can then be managed for performance and other goals by intelligently placing them on the most appropriate type of storage (e.g., SSD, SAS, SATA), RAID-levels (e.g., RAID-10, RAID-5), and even track-placement on a given spindle (Fast Track).

The Compellent Dynamic Block Architecture stores write data on a pool of blocks residing on the fastest performing tier of storage. As illustrated in the example shown in Figure 7, writes are stored on a pool of storage protected with a high performance RAID-10 algorithm. Reads are serviced from that same group of drives until the Data Progression algorithm kicks in (the default is once per day at 7 PM).

Data Progression migrates infrequently used blocks to a more cost effective tier of storage based on policies managed by an administrator. In this example, data progression migrates first to RAID-5 storage in the tier-1 pool (note that this novel approach provides a cost effective way to avoid the performance penalty associated with RAID-5 writes). Later, the automated Data Progression algorithm moves infrequently accessed data to a cost effective second tier of storage. If and when data becomes "hot" again, Data Progression reverses the process and migrates the blocks back to the higher performing tier.

Figure 7. Data Progression

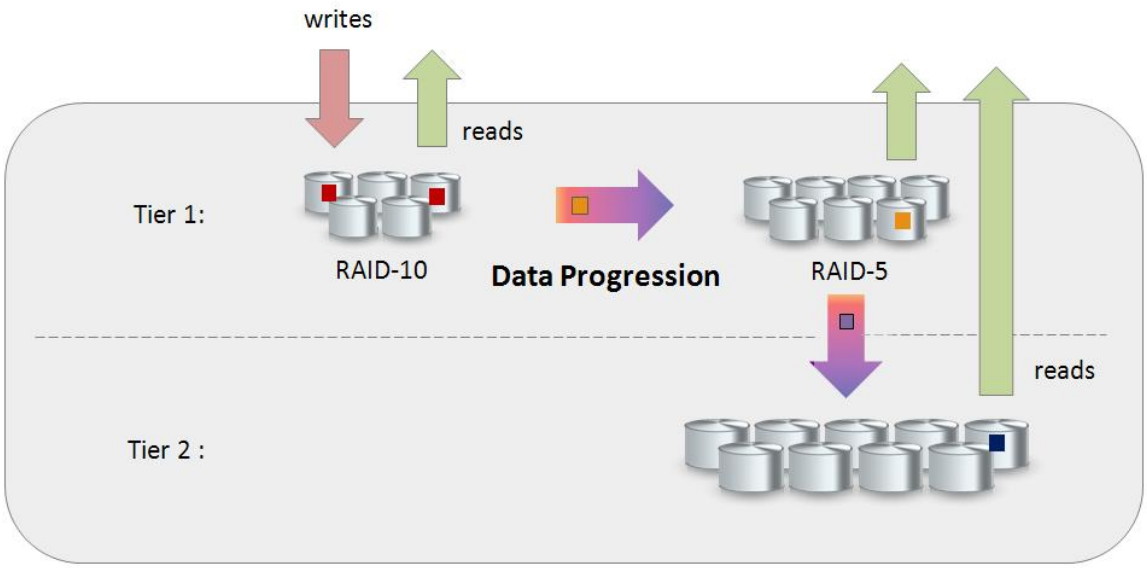

ESG Lab Testing

ESG Lab application workloads and the industry standard open source IOmeter utility were used to quantify the performance benefit that can be achieved with Data Progression.[2] An online transaction processing (OLTP) database workload that simulates the activity of a multi-user Microsoft SQL Server application was tested. A virtual desktop infrastructure (VDI) workload that simulates a VMware View virtual desktop infrastructure using the linked clone method to share operating system "gold" images was also tested.

Data Progression testing was performed with an 8 Gbps Fibre Channel attached Compellent Storage solution that was configured with three tiers of storage:

- Tier 1: Four 200 GB SSD devices
- Tier 2: Forty-three 150 GB 15K RPM SAS drives
- Tier 3: Eleven 2 TB 7.2K RPM SATA drives

Application workloads were tested after Data Progression had run with a goal of measuring the performance boost that can be achieved with the fully automated page-level tiering algorithms built into Compellent Storage Center version 5.4. The results are summarized in Figure 8 and Figure 9.

Figure 8. Data Progression Response Time Analysis

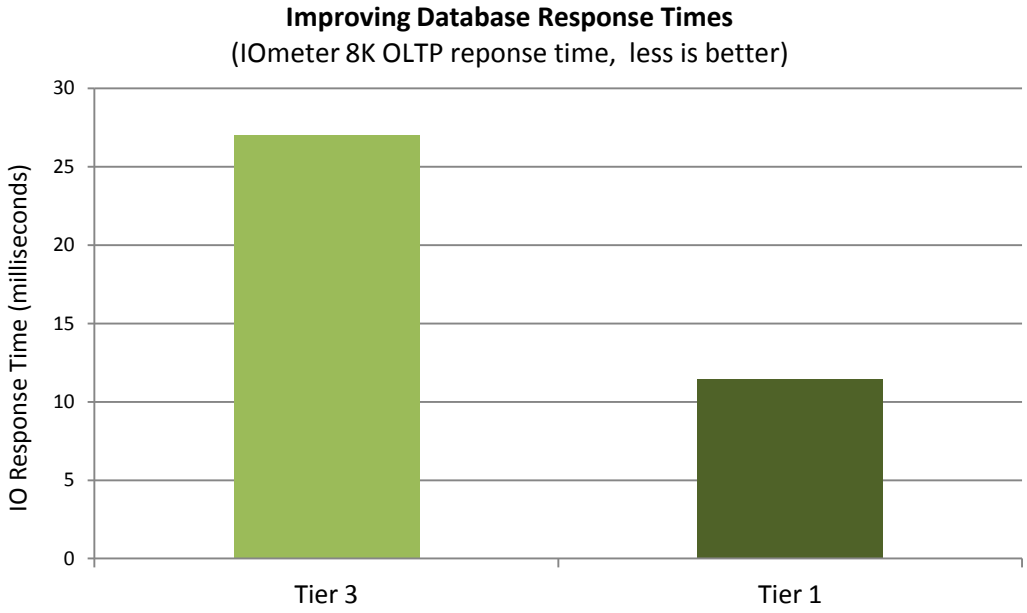

ESG Lab observed what Dell referred to as common guidance in that only a relatively small amount of the higher-performing storage was added to the environment. In conversation, ESG learned that most Compellent customers solve their capacity issues by purchasing only additional lower tier drives. Similarly, those same customers solve their performance issues by buying a smaller high-performance drives for a new or augmented top tier. This is a welcome contrast to the legacy approach of having to relocate entire data sets between storage arrays or tiers as workload performance needs change over time.

[2] http://www.enterprisestrategygroup.com/using-esg-lab-workloads/

Figure 9. Data Progression Aggregate Workload Analysis

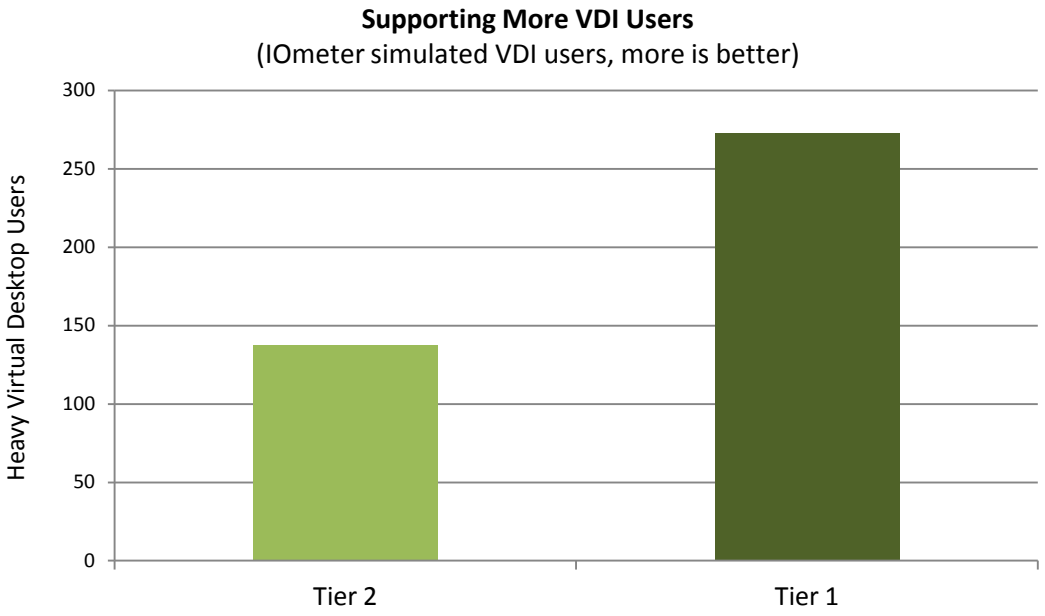

What the Numbers Mean

- The response times for the database application were 60% faster when moved from tier-3 SATA devices to tier-1 SSD devices.
- Faster response times at the IO level, which are shown as less in Figure 8, are magnified at the application level where each database transaction typically requires multiple IO requests.
- Moving the virtual desktop workload from tier-2 SAS drives to tier-1 SSD drives increased the number of IOs per second by 98%.
- ESG Lab measured IO activity of 20 IOs per second (IOPs) for an ESG analyst over the course of an eight hour work day. Using that value as representative of a typical knowledge worker, a recorded IO rate of 5,453 IOPS can support up to 272 heavy VDI users—98% more than was achieved with tier-2 SAS drives.

As notable as the numbers are, the experience was equally noteworthy. After configuring the IO tests with the initial storage, ESG Lab found it very easy to enable a higher-performing tier of storage. As the Compellent algorithms took effect, ESG was able to observe the increased performance as the faster storage began being utilized.

Why This Matters

Companies continuously face challenges in cost-effectively meeting the capacity and performance requirements of applications, especially those with strict performance requirements. Failure to meet these requirements can result in lost productivity and costly loss of services, but over-provisioning to avoid performance problems is also a waste of money.

ESG Lab verified that Compellent Dynamic Block Architecture simplifies storage tiering as it maximizes performance and reduces costs. After setting a data progression policy, performance was automatically optimized using pools of blocks that were allocated at the right time, for the right data, at the right price.

Recoverability

The Compellent Dynamic Block Architecture is used to create disk-based point-in-time Replays (a.k.a., snapshots). Administrators use Replays to perform non-disruptive backups of production data, as seen in Figure 10, as well as for application testing, development, and fast recovery of deleted files or from corrupted data. In this example, a file created at 11:55 AM was recovered from disk using a Replay Volume created at 12:10. From an end-user standpoint, the Replay was used to dial back in time and recover the file using a point-in-time image created before the corruption.

Figure 10. Using a Compellent Replay to Recover a Corrupted File

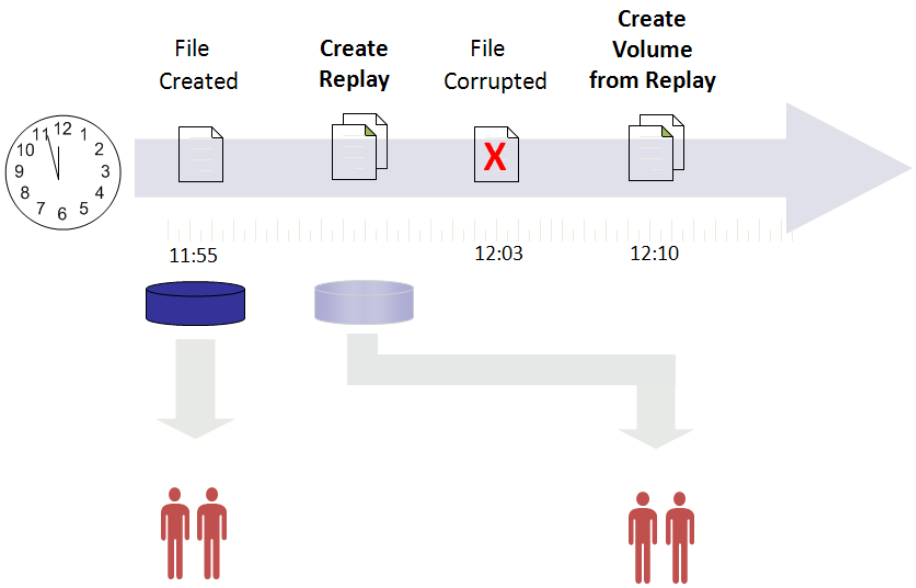

Compellent's Dynamic Block Architecture not only enables blocks to seamlessly move between media tiers for optimized performance and cost, but is also the enabling technology for local and remote Replays. The Dynamic Block Architecture enables capabilities that are often not supported in snapshot implements within legacy storage solutions:

1. Replay data is stored on a lower cost tier of storage compared to the pool where the primary data resides. For example, a volume residing on high speed SAS drives can be protected with Replay capacity that resides on denser, more affordable SATA drives.
2. Replay capacity is delivered on demand from a pool of available blocks. Replay capacity does not have to be reserved up front.
3. Remote Replays send only the latest changes over the network to a remote Storage Center. This reduces WAN bandwidth requirements as it creates a cost effective remote mirror for disaster recovery.

ESG Lab Testing

ESG Lab used a Compellent Replay to recover from simulated data corruption. Specifically, ESG Lab created a volume from a replay so that the production server would continue as is, while a mountable version of the volume from a previous point in time was used for the recovery.

Because a Replay is simply a different list of pointers to blocks within the same Dynamic Block Architecture, a previous point in time consumes almost no space on its own, as seen in Figure 11 where both the production volume and the replay volume are visible in the left-hand tree, but the Replay Volume only consumes only 78 MB of actual disk space.

Figure 11. Replay Volume from the Perspective of Storage Center

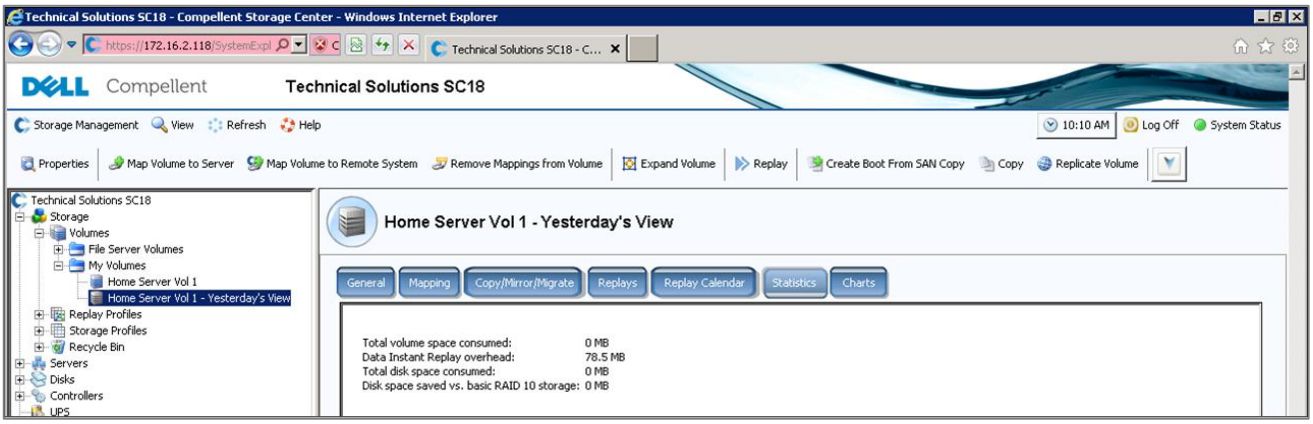

Figure 12 shows the same Replay volume as a separately mountable and usable resource on a Windows Server.

Figure 12. Replay Volume from the Perspective of Windows Server

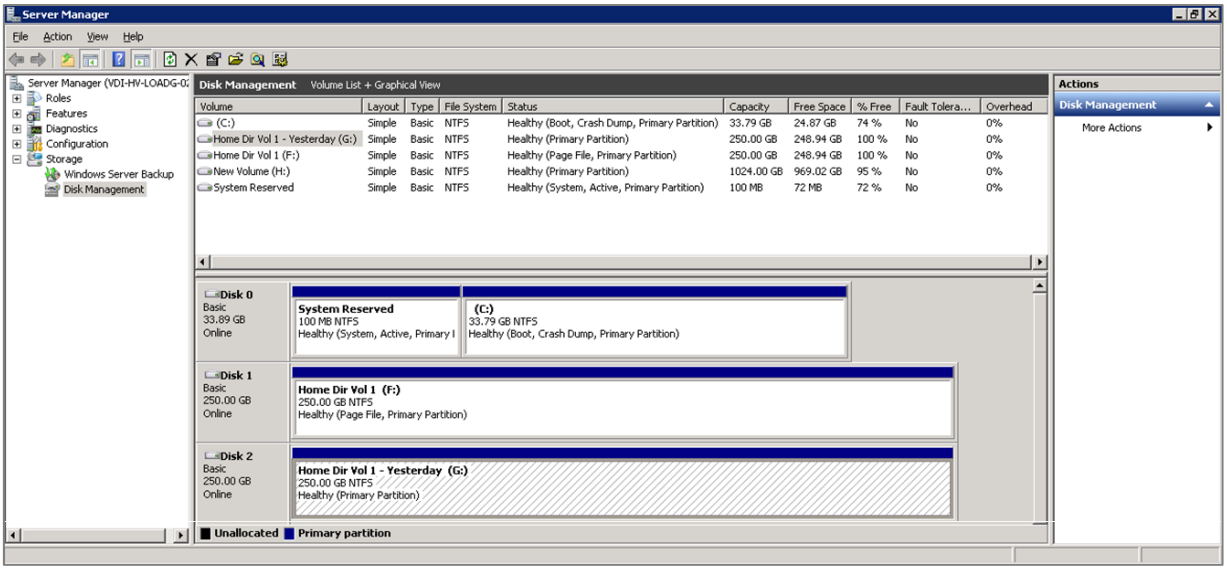

The key here is to recognize that the Replay volume is in fact "just pointers," but Windows couldn't tell the difference. While Figure 11 shows the replay volume made up of only 78 MB of overhead, Windows sees an entirely separate 250 GB volume. For this test, ESG Lab was able to easily create a Replay volume and then mount it to the original production server so that a previous version of a file could be accessed. The same capability can be used to create new server volumes from a master-image scenario.

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, application testing, and master-image deployments.

ESG Lab confirmed that the Dynamic Block Architecture delivers space-efficient Replays for quick and easy recovery of accidentally deleted or corrupt files. ESG Lab was particularly impressed by the ability to store Replay updates on a cost effective lower tier of storage, the WAN-efficient replication of Replays to a remote site, and the fact that Replay capacity is delivered on demand without reservations.

Virtualization

According to recent ESG research,[3] 70% of organizations are using more than one server virtualization platform. Just as IT has managed and maintained a heterogeneous operating system environment, evidence also suggests that multiple hypervisors will be deployed. IT is heterogeneous by nature, so this makes complete sense. All hypervisors essentially do the same job, but licensing fees for proprietary products are driving enterprises to adopt a second choice. The increasing maturity of Citrix XenServer and Microsoft Windows Server 2008 servers running Hyper-V are grabbing enterprise attention, but VMware is not necessarily being displaced where it exists.

A natural pairing to Compellent's virtualization of storage is to combine Storage Center with server virtualization platforms, including both VMware and Microsoft Hyper-V. This pairing is done from two perspectives:

- A hypervisor-aware perspective from within Storage Center
- Compellent-embedded information within server virtualization management tools

ESG Lab Testing

Figure 13 shows the multiple Storage Centers used during testing and the range of operating systems and virtual server hypervisors connected to them. In the expanded view on the right side, one can see ESXi and Windows Server 2008 R2 (Hyper-V) servers, each with multiple storage LUNs attached to each hypervisor.

Figure 13. Compellent Enterprise Manager Showing Storage Center Connected to Multiple Hypervisors

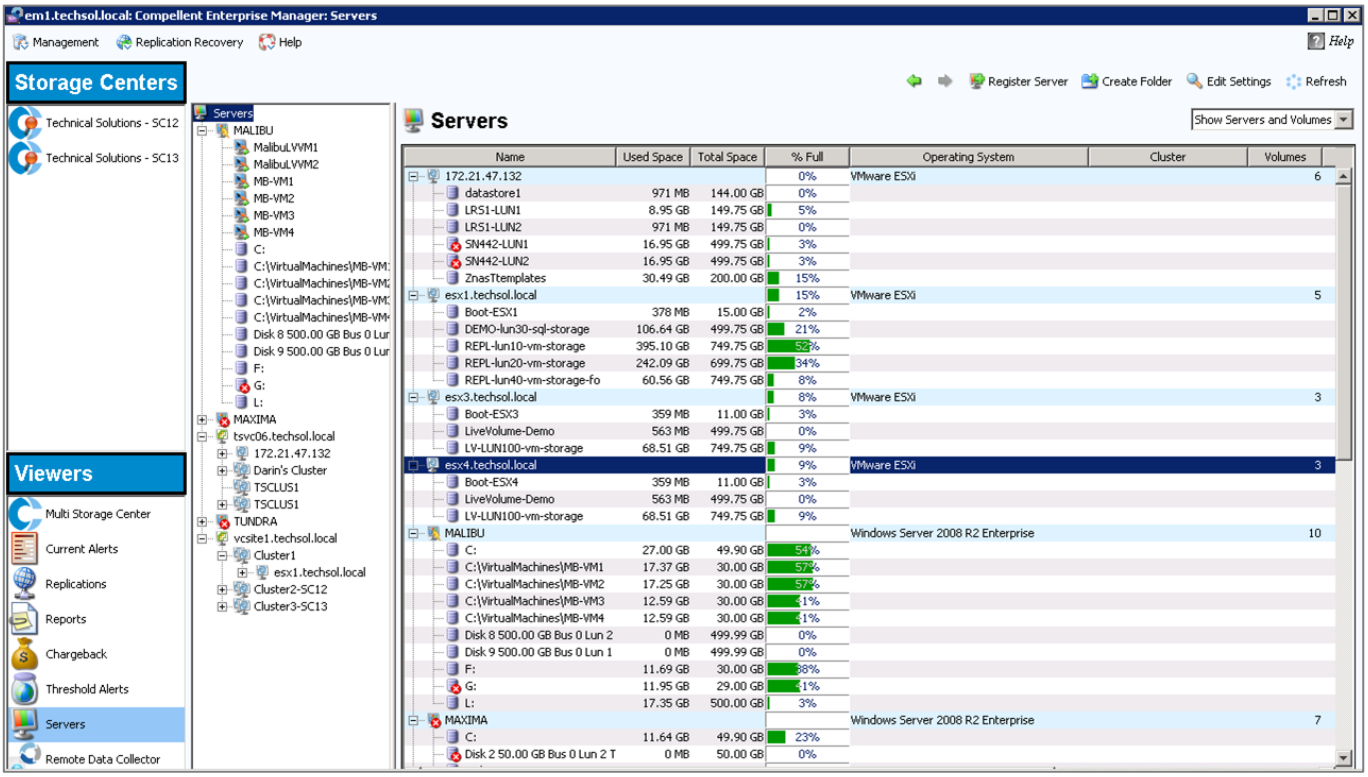

ESG Lab tested Compellent's integration with VMware ESX and vSphere 5.0 by monitoring and managing several typical storage functions from within the vCenter interface. It should be noted that similar integration was observed through the Microsoft PowerShell command language that provides scripting control for Hyper-V environments; however, the integration was best observed in the graphical vSphere UI.

ESG Lab observed many of the same provisioning activities and interfaces done earlier with physical servers but using the vSphere Client interfaces. Specifically, Figure 14 shows similar data to what was observed in Figure 6, related to a particular volume's consumption of storage across multiple tiers within the Storage Center.

[3] Source: ESG Research Report, *The Evolution of Server Virtualization*, November 2010.

What is most notable about the integration is the "Compellent View" tab in the VMware vSphere management console as shown in Figure 14. In this example, the Compellent attributes of the storage associated with a virtualized server are shown. Several common monitoring and management tasks were executed from this VMware management console.

Figure 14. Storage Center Integrated Interfaces within vSphere

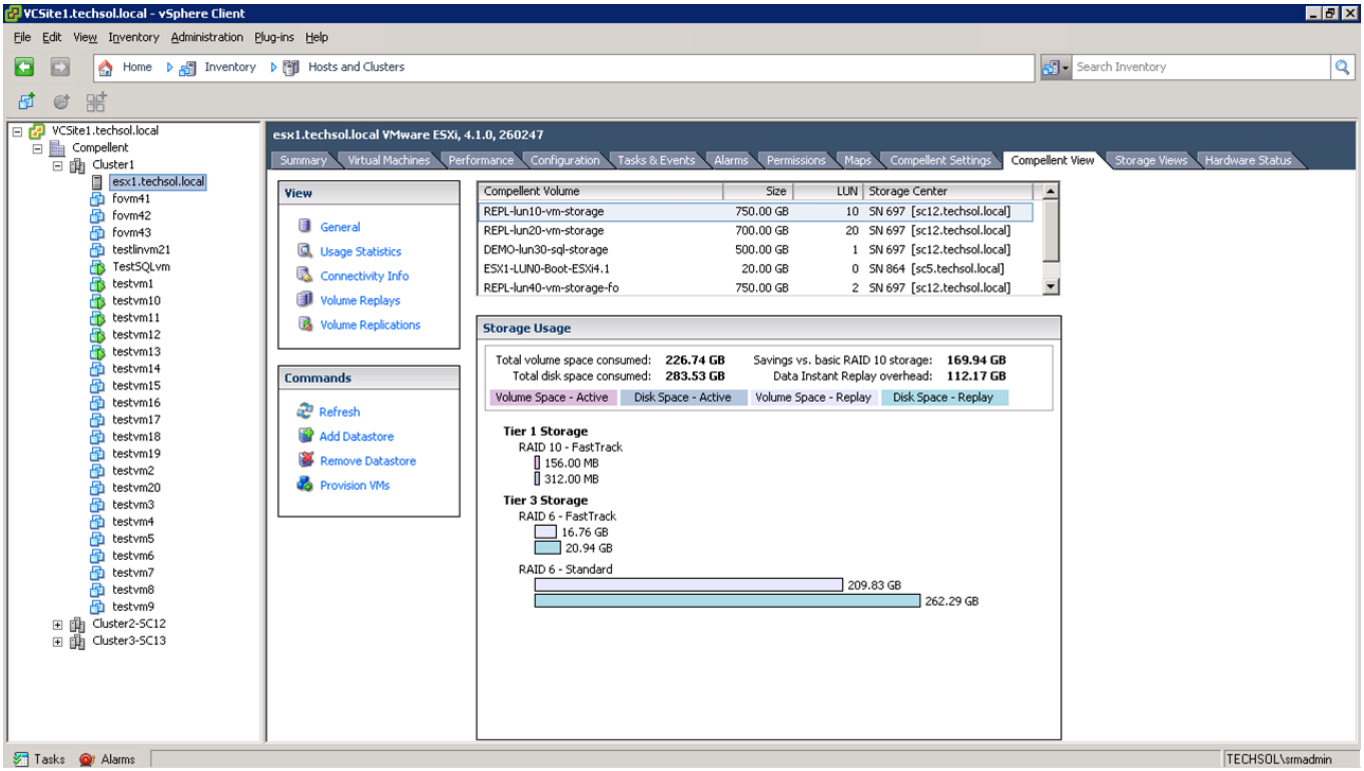

To further test the integration, ESG Lab attempted to manage the ESX storage from its normal view within the vSphere Configuration Tab, as seen in Figure 15.

Figure 15. Storage Center Actions within vSphere Right Clicks

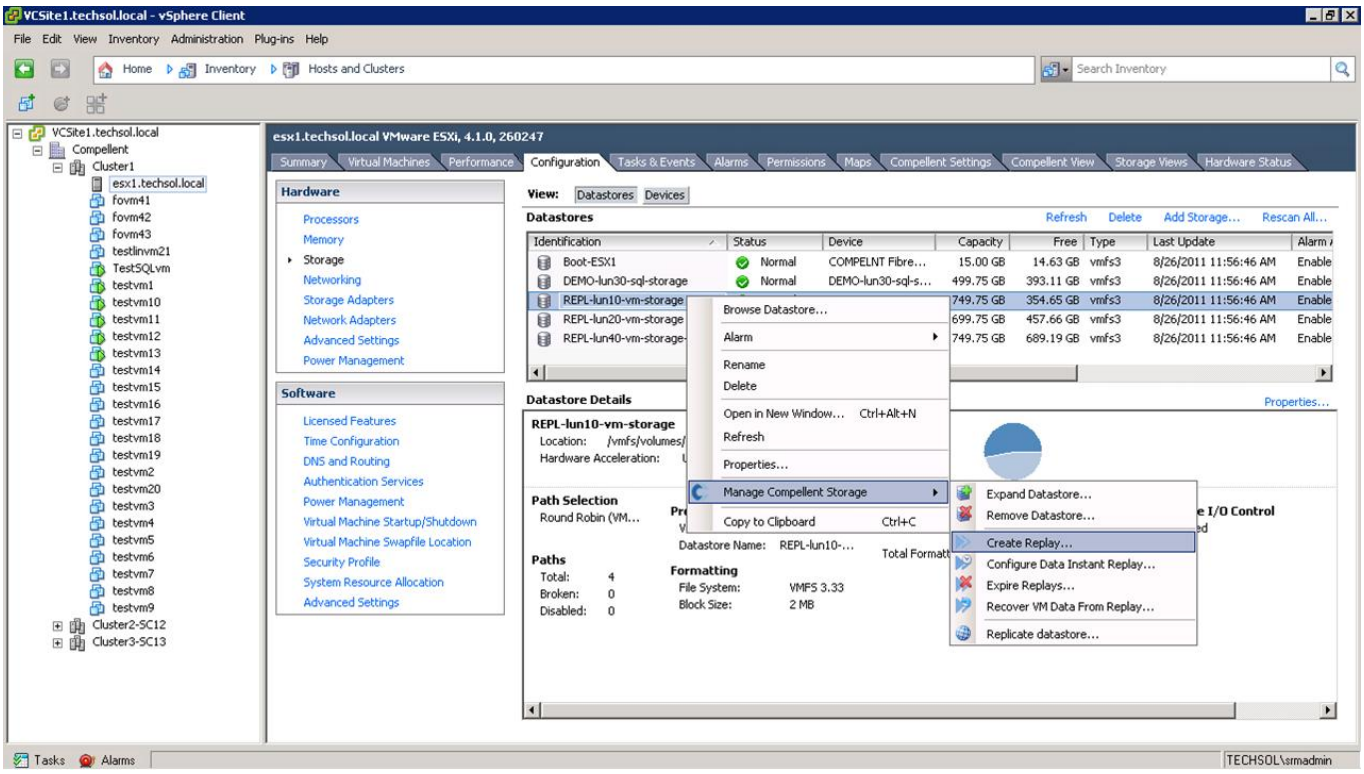

Specifically, ESG Lab tested the ability to invoke a "Replay," commonly referred to as a snapshot, from within vSphere. To accomplish this, ESG Lab selected an ESX cluster in the left-pane of vSphere. After selecting storage in the hardware selection pane, a right click on a volume revealed the Compellent sub-menu of commands shown in Figure 15.

ESG Lab anticipates that VMware administrators will be highly impressed with and empowered by the clean integration of Compellent storage with ESX hosts. With only minimal orientation to the Compellent terms and concepts, some very advanced storage/virtualization activities were accomplishable without ever leaving vSphere.

Why This Matters

ESG research indicates that server virtualization has been the top IT priority for the past three years. The need to support multiple hypervisors is becoming a key requirement as IT environments continue consolidating not just their server infrastructures, but also their storage. And while compatibility or awareness of the hypervisor is a positive for virtualized storage solutions, what is more important is the reverse: integration from the hypervisor's perspective toward its storage.

ESG Lab confirmed that the Compellent Storage Center management console is server virtualization-aware. ESG Lab has also confirmed that plug-ins for VMware and integration with Microsoft Hyper-V can be used to manage a consolidated virtual server and Compellent-enabled virtual storage infrastructure from a single pane of glass. By virtualizing the storage and then integrating it into the server management tools, IT environments can manage their computing fabric from a unified interface without having to bounce between management interfaces, saving money and accelerating time to value.

Data Agility and Mobility

The same Dynamic Block Architecture that enables blocks to seamlessly move between media tiers for optimized performance and recovery can also be used for online migration and synchronization of blocks between Storage Centers. This powerful capability simplifies the migration of virtual machines between sites and the migration of data between Compellent Storage Centers. The virtual machine migration capability is depicted in Figure 16.

In this example, servers located in two different data centers have been configured as a virtual server cluster. The virtual server platform is being used to move an application running in a virtual machine from one data center to the other. While the virtual machine is moved using industry standard virtual server technology (e.g., VMware VMotion, Microsoft Live Migration), Compellent Live Volume replicates changed blocks over the WAN and creates a single consistent view of data regardless of which data center is being used for production.

Figure 16. Live Volume Enabling Online Virtual Machine Migration

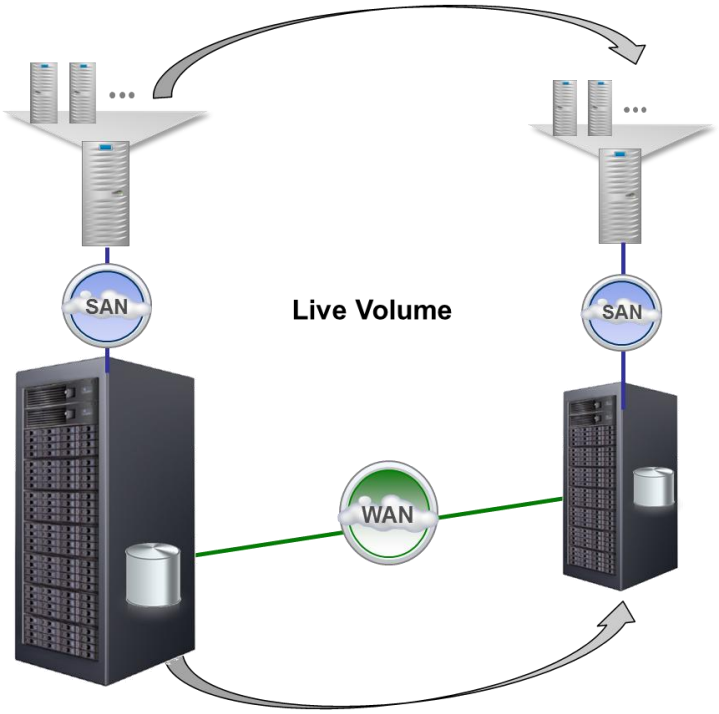

This powerful capability is an efficient, automated, and cost effective alternative to methods that have become available over the past few years. For example, VMware storage VMotion provides similar functionality, but requires too much server processing power and network bandwidth to make it practical for use between data centers. This can also be achieved with expensive clustering hardware that provides a view of storage between data centers deployed at each site.

ESG Lab Testing

ESG Lab tested the Live Volume feature by first creating an asynchronously replicated volume between two Storage Centers connected across a campus environment in a VMware environment.

Shown in Figure 17, ESG Lab selected a volume that was storing virtual machines from one of the hypervisors and was able to configure replication by simply right clicking on the primary volume and selecting which Storage Center would hold the alternate copy.

Figure 17. Configuring Live Volume Replication

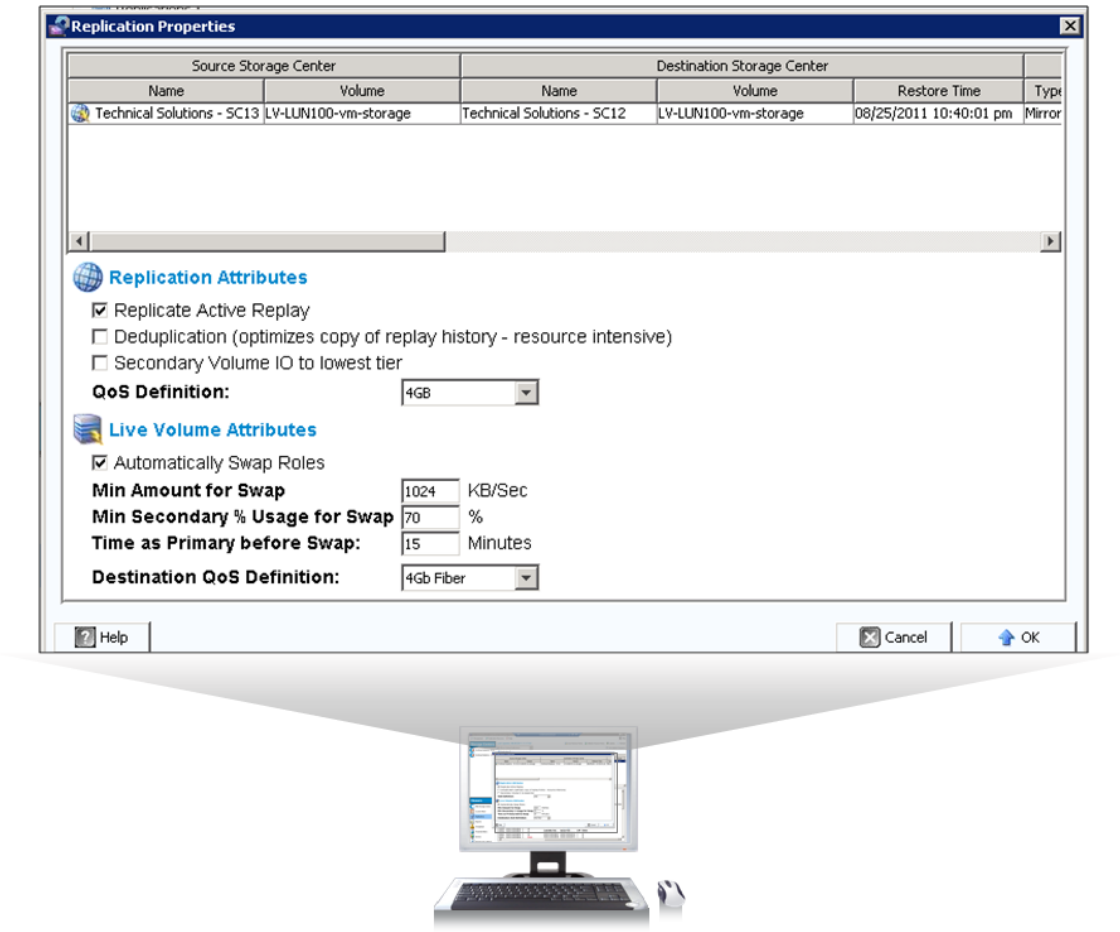

The most notable option in Figure 17 is the ability to automatically swap roles between the copies. This is the key benefit of the Live Volume feature in that, regardless of which host may be initially driving IO, the volume is adaptive. If the IO were to move from a host on one Storage Center to the other—for example, as part of a server-level recovery or migration scenario—the Storage Centers would mutually identify that the IO was now being generated from the other side and the direction of the replication traffic would transparently reverse. ESG Lab tested this by moving workloads between production hosts on either side of the Storage Center pair and observing the IO shifts between the sides and the directional shift of replication.

Figure 18. Live Volume Actions within Storage Center

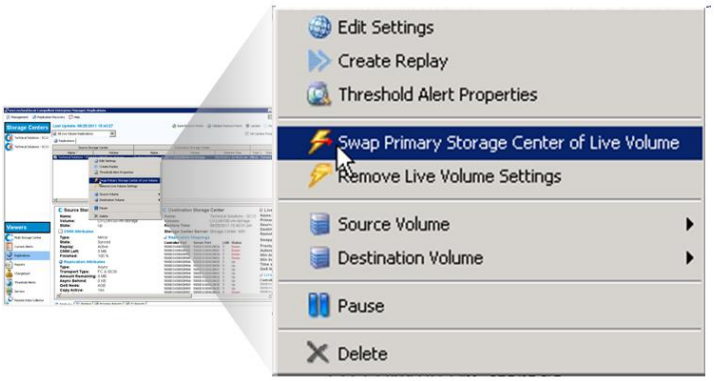

With Live Volume configured and replicating automatically, moving a running virtual machine to another data center was simply performed using the VMware vSphere client. Using the standard user interface for VMware VMotion, a virtual machine was dragged and dropped from one virtual server to the other. In the screenshot shown in Figure 19, the migrate option accessed from a right click on a virtual machine was used to launch the migration wizard. The virtual machine migration completed in less than five minutes. A heavy IO workload started before the VMotion was used to notice a momentary dip in performance, which lasted about three seconds. The virtual machine remained up and available throughout the test.

Figure 19. VMotion Between Data Centers with Live Volume

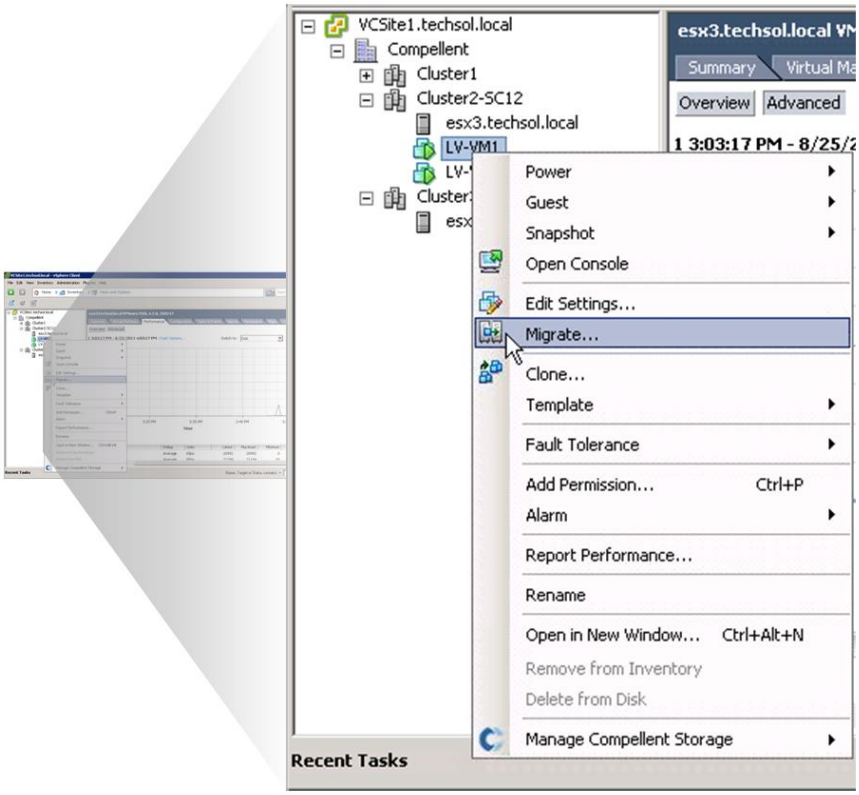

As seen in Figure 19, planned downtime or other re-hosting events was tested by ESG Lab with a few easy mouse clicks within vSphere that then leveraged Compellent's Live Volume capability for a successful migration.

Why This Matters

IT organizations need to move applications from one data center to another for a variety of reasons, including data center consolidations, planned maintenance, disaster avoidance, and load balancing. Traditional methods typically require planned downtime. If the application must remain online during the move, a costly and complex combination of host clustering software and disk array remote mirroring is needed.

ESG Lab has confirmed that Compellent Live Volume can be used to quickly and easily move a running virtual machine between data centers. After configuring Live Volume to maintain a single of view of virtual machine data that's replicated between sites automatically, the VMotion process looks and feels exactly as if Live Volume were not there. For administrators familiar with the drag and drop VMotion interface, there is nothing new to learn.

ESG Lab Validation Highlights

Along with the various "Why this Matters" conclusions throughout the report, the following is a summary of the actual hands-on tests and their outcomes:

- ☑ ESG Lab validated the performance and flexible cost-benefit of the Dynamic Block Architecture that underlies the Storage Center platform.
- ☑ ESG Lab confirmed that the right data blocks are placed on the most appropriate layer of storage based on relative demand and price/performance. In fact, Compellent anecdotally states that many customers often do not purchase any incremental tier-1 storage following the initial purchase because their capacity growth issues are addressed by the lower storage—not withstanding new workload requirements or performance innovations such as SSD as a new top tier.
- ☑ ESG Lab confirmed the ease of use and transparent operation of Compellent's multi-tier storage solution.
- ☑ ESG Lab observed impressive integration between Compellent's own management tools and those of hypervisors and server platforms as well as third-party monitoring, which collectively adds to Compellent's value because the capabilities become attainable from existing management interfaces.
- ☑ ESG Lab confirmed the Replay capabilities of Storage Center to yield site-level resilience, as well as easy volume-level data recoveries.
- ☑ Live volume was used to perform an online VMotion of a VMware virtual machine between data centers. The VMotion operation was easy to configure and monitor from the VMware vSphere console. The virtual machine remained up and available during the migration.

Issues to Consider

While ESG Lab found many commendable aspects of Dell's Storage Center, the following are a summary of considerations for future development by the Compellent team and/or factors to consider by prospective customers:

- ☑ While the integration of Storage Center with VMware vCenter was impressive and logical, parity for other hypervisors is not yet in place. With the anticipated release of Microsoft's next generation of management tools in System Center 2012, Compellent has an opportunity to provide similar integration into not only virtualization management (VMM 2012), but also systems management (OpsMgr 2012) which would likely provide parity for Citrix Xen configurations.
- ☑ While Live Volume can be used for online migration of storage between Compellent Storage Center solutions, it does not support data migrations between heterogeneous disk arrays from different vendors. If a downtime window cannot be tolerated, a traditional host-based data migration utility or a network resident data mover can be used to migrate storage from another product to a Dell Compellent Storage Center.
- ☑ While the VMware vSphere plug-in provides access to typical Storage Center management tasks, the ability to perform a Live Volume swap operation from a VMware management console is not yet supported. Dell advised ESG that this is planned for an upcoming release.

The Bigger Truth

One of the greatest challenges IT faces today is managing unabated data growth in dynamic, virtualized environments. 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 1,602 IT professionals indicates that server virtualization is leading the charge. Seventy-four percent of respondents report that they are actively using server virtualization.[4] Server virtualization, like storage virtualization, is powerful game-changing technology that can be used to consolidate and simplify complex IT infrastructure. Like the early adopters of server virtualization, forward-looking IT managers are turning to highly virtualized storage solutions to reduce costs, increase efficiency, and enhance the recoverability and agility of a consolidated storage infrastructure. ESG has confirmed that Compellent Storage Center version 5.4 is a powerful virtual storage solution that's built on top of a pool-based Dynamic Block Architecture.

ESG Lab tested several sought-after storage capabilities, including thin provisioning, remote replication, rapid disk-based recovery, virtual machine migration between sites, and deep integration with leading virtual server platforms. In all cases, the answer to "Why is this different from a traditional SAN solution?" came back to Compellent's Dynamic Block Architecture. The Dynamic Block Architecture is the DNA that enables simple and capacity efficient deployment of new volumes on physical and virtual servers. Dynamic Block Architecture also makes it easy to move applications and virtual machines between sites from a storage perspective, and to quickly recover data with capacity and network efficiency in mind.

ESG Lab was impressed with the flexibility and power of the Dynamic Block Architecture when first testing Storage Center in 2008.[5] Valuable capabilities that are now known in the industry as thin provisioning and sub-LUN tiering were built into the architecture and have been proven in the field by thousands of customers since the first product was shipped in 2005. As legacy storage architectures evolve to support these valuable capabilities, the flexibility and power of the foundational Dynamic Block Architecture is turning into a competitive advantage for Dell. Features like Live Volume, which provides cost-effective clustered access to a virtualized pool of storage over distance, would be very difficult to implement without a page-based architecture.

While advanced features like Live Volume are powerful, the bottom line with the Dynamic Block Architecture is its ability to reduce the overall cost of purchasing and maintaining storage. Writes are automatically directed to the highest performing tier. Data Progression moves infrequently used data to a more cost effective tier and Replays are stored on the most cost effective tier. Due to these core capabilities, the Dell team indicates that many customers are using affordable SATA drives for most of their capacity with a couple of fast SAS drives for an automatic performance boost. That's significantly more affordable than a traditional disk array full of expensive FC drives.

ESG Lab commends the Compellent team for the forethought needed to envision the value of the Dynamic Block Architecture back in 2002. This architecture has created a sound foundation that makes it possible for Dell to cost effectively deliver features and value that established storage vendors are struggling to match. ESG Lab congratulates Dell for another smart acquisition that will continue to add value to its data center portfolio. At the end of the day, ESG Lab finds that the goal of developing *efficient, agile, and resilient* storage has been met with the latest release of Compellent Storage Center. IT managers shopping for their next storage infrastructure upgrade—especially those that have embraced the benefits of server virtualization—should seriously consider the benefit of virtualizing their storage infrastructure with the growing family of Fluid Data solutions from Dell.

[4] Source: ESG Research Report, *The Evolution of Server Virtualization*, November 2010.
[5] A list of previously published Compellent-focused ESG publications can be found in the Appendix.

Appendix

***Table 1**. Primary ESG Lab Test Bed*

Server	
Operating System	Windows Server 2008 R2 SP1
Hardware	Dell R710 - 1x Single 2.67 GHz Quad-Core CPU, 32 GB 800 MHz RAM
Compellent Storage Center version 5.4	
Tier 1	4x 200 GB SSD
Tier 2	43x 150 GB 15K RPM SAS
Tier 3	11x 2 TB 7.2K RPM SATA
SAN	
Fibre Channel	8 Gbps
Performance Characterization Tool	
IOmeter	Version 2008.06.18
Performance Characterization Workloads	
OLTP	8 KB block size 70% random reads 30% random write
VDI	16 KB block size 70% random reads 30% random write

Additional ESG Coverage of Compellent

ESG has covered (Dell) Compellent in previous reports, including:

- ESG Market Landscape Report, *Replication Technologies for Business Continuity*, August 2011
- ESG Market Landscape Report, *Storage Tiering*, July 2011
- ESG In the News, *Dell nears Acquisition Deal for Compellent*, December 2010
- ESG Technology Brief, *Compellent – Extremely Efficient Storage*, April 2009
- ESG Lab Validation, *Storage Center 4.0*, February 2008

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

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