LAB VALIDATION REPORT

Compellent Storage Center v4.0

Sophisticated Storage Simplicity

By Brian Garrett
With Tony Palmer

February, 2008

Table of 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 Compellent.

Introduction

Organizations of all sizes are struggling with the growth, cost and complexity of online storage. New software running on a recently upgraded industry standard server platform extends the sophisticated storage simplicity of the Compellent Storage Area Network (SAN) storage solution. This ESG Lab validation highlights the ease of use and cost-effectiveness of Compellent Storage Center version 4.0 with a focus on valuable capabilities including Fast Track, Data Progression, Dynamic Capacity, Data Instant Replay and Thin Remote Replication.

Background

Storage professionals within large organizations have realized the benefits of SAN-attached storage solutions for more than a decade. Using Fibre Channel (FC) SAN technology to implement a shared storage infrastructure, data center managers have increased utilization, centralized management and improved reliability. The benefits of a centralized SAN infrastructure have been extended to the broader market in recent years due to the availability of lower cost FC solutions and more recently due to iSCSI SAN solutions that leverage commodity Ethernet technology.

The priorities shown in Figure 1 reflect the challenges that SAN administrators are faced with when trying to meet growing online storage requirements.[1] These results were collected from planned iSCSI adopters during a recent ESG survey, which indicates that iSCSI has entered the early mainstream. Purchasing priorities for these organizations correlate well with surveys and conversations that ESG has had with SAN storage administrators within companies of all sizes. Taken together, the priorities are clear: organizations want to reduce the cost and complexity of deploying shared SAN storage.

FIGURE 1. STORAGE PURCHASING PRIORITIES

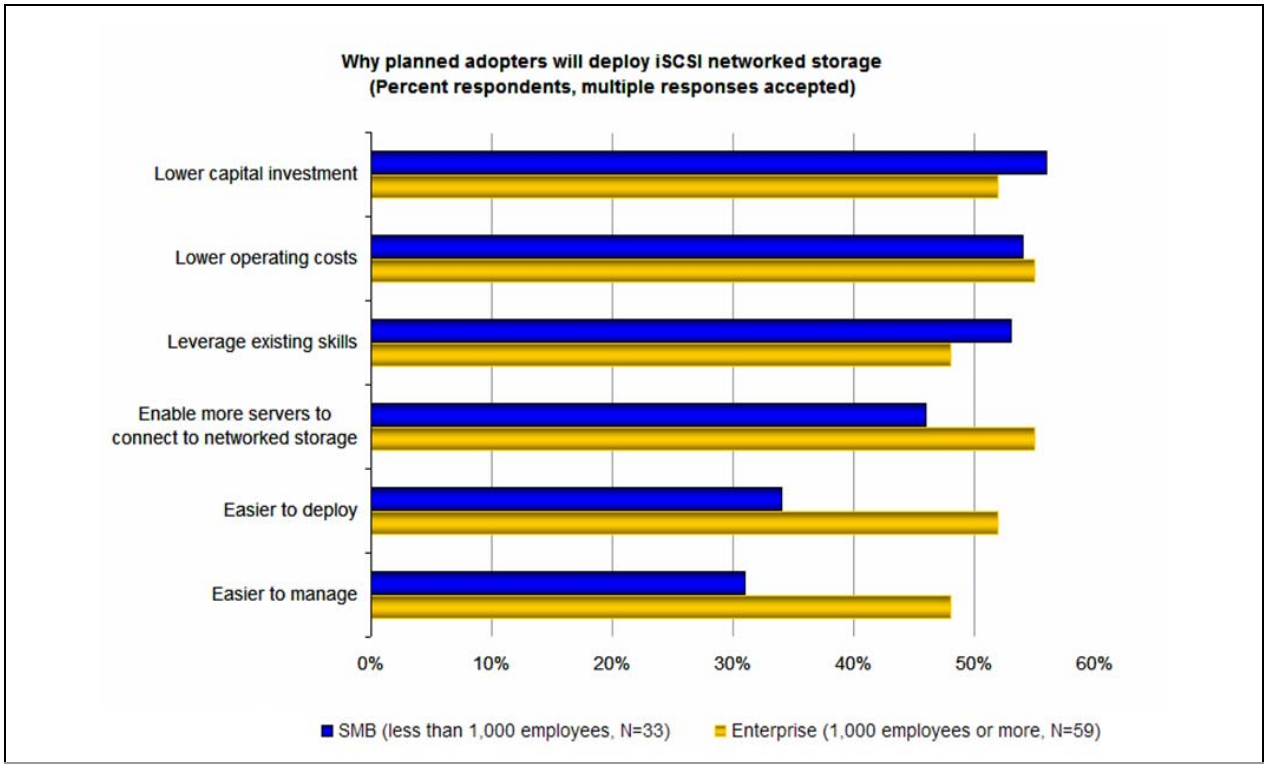

[1] ESG Research Report: *iSCSI Enters the Mainstream,* January 2006

Server virtualization is driving a new wave of networked storage adoption. Using server virtualization technology to consolidate servers saves money while reducing power, cooling and space demands in the data center. While the same benefits can be realized with a consolidated storage strategy (less to buy, manage and cool), the combination of server virtualization and SAN-attached storage enables server mobility and increased availability. As a result, organizations that have already invested in Fibre Channel SANs are retooling their data centers to combine server virtualization with SAN-enabled storage consolidation. Organizations new to SAN technology are considering iSCSI SANs to compliment their server consolidation strategy. Regardless of the technology, storage networks are being adopted in increasing numbers in order to reduce costs, increase flexibility and enhance reliability.

The latest frontier for SAN innovation is simplicity and ease of management. The first wave of enterprise-class Fibre Channel SANS that came into the market were notoriously hard to configure and manage. As SANs moved into the broader market and iSCSI SAN solutions became available in recent years, easier deployment and management have become key concerns. Regardless of the size of the organization, SAN solutions that are easier to manage enable administrators to keep up with growing capacity requirements.

SAN-attached storage solutions have evolved over the years to support a number of advanced storage services. Sophisticated recoverability services—including snapshots and remote replication—are being used to better meet compliance and service level agreements. Tiered storage and the ability to migrate application data to the right level of performance, cost and protection over time are being used to reduce the cost of capacity. Advanced provisioning techniques are being used to reduce the cost of allocated, but unused, storage. While each of these sophisticated storage services is valuable, they all would be considerably more valuable if they were easy to configure and manage.

Compellent

ESG has been working with Compellent since it was founded in 2002 by a team of industry pioneers with decades of storage experience. ESG noticed then that Compellent uses an innovative approach to develop, sell and support SAN storage solutions—and that hasn't changed. First, and quite unique in the storage industry, Compellent derives 100% of its sales through partners and value added resellers. Choosing not to compete with their partners using a direct sales force, Compellent works closely with resellers who know and understand the needs of their customers. Product features and capabilities are prioritized based on feedback from an active customer council. Resellers and dozens of early adopters clearly indicated that simplicity and ease of use were important. When the product was first launched, Compellent let dozens of resellers and early customers tell the story based on their experiences.

Compellent Technologies is publicly traded and has been deployed in more than 740 worldwide organizations. Its success is based on a growing network of international channel partners and an award winning product line.

Storage Center, version 4.0

Storage Center is a standards-based, modular solution that connects to servers through a storage network (iSCSI, FC, NAS). Multiple storage tiers are supported to meet a wide variety of scalability, performance and protection objectives (FC, SATA). Sophisticated storage services are delivered with the following customer objectives in mind:

- Purchase less storage capacity
- Manage more without adding staff
- Mix and match technologies
- Optimize performance
- Reduce power and cooling demands
- Enable rapid and reliable recovery
- Simplify capacity planning

FIGURE 2. COMPELLENT STORAGE CENTER, VERSION 4.0

Storage Center is a combination of software accessed from a web browser, a pair of clustered controllers and a number of drive enclosures as shown in FIgure 2. Version 4.0 Storage Center software is offered in conjunction with a 3rd generation controller platform with an upgraded multi-core processor, increased memory and a faster bus. Version 4.0 adds a number of valuable enhancements including a virtual tier of storage called Fast Track. This valuable new capability takes advantage of the faster performance that can be achieved by placing data towards the outside edges of a hard drive. This report presents the results of ESG Lab's experience with Storage Center version 4.0, including an in-depth examination of the cost-effective performance benefits of Fast Track.

ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of Storage Center version 4.0 at Compellent corporate headquarters in Eden Prairie, Minnesota. Testing was designed to demonstrate the flexibility, power and simplicity of Compellent SAN solutions with a focus on features including Fast Track, Data Progression, Dynamic Capacity, Data Instant Replay and Thin Remote Replication.

Ease of Implementation and Management

The ESG Lab Validation began with an evaluation of ease of initial deployment, configuration and ongoing management of the test bed shown in Figure 3. Windows and Linux servers were connected to Compellent Storage Center through FC and iSCSI storage area networks. A pair of clustered Compellent controllers were attached to drive enclosures containing a mix of drive types including high performance FC drives, mid-tier FC drives and bulk capacity SATA drives. The primary storage system running Storage Center version 4.0 was connected through a simulated WAN to a second storage system running Storage Center version 3.6.[2]

FIGURE 3. THE ESG LAB TEST BED

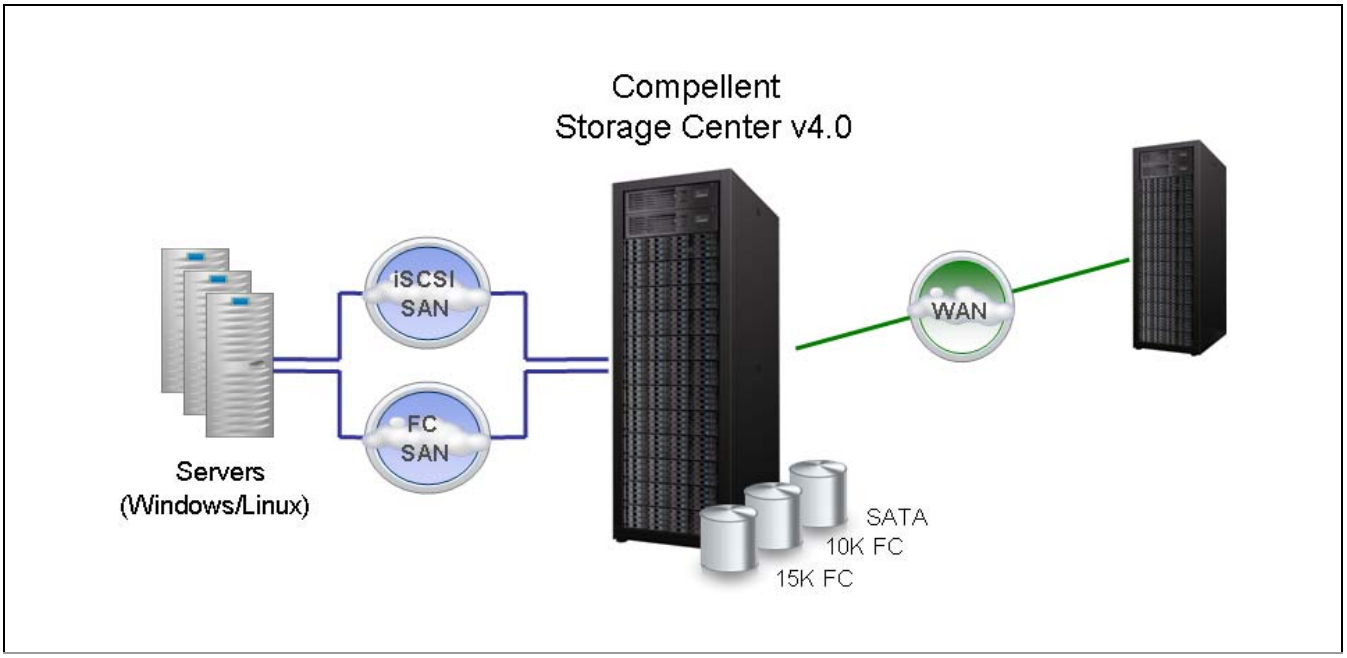

ESG Lab Testing

A configuration from scratch on a pre-wired test bed began with the assignment of IP addresses. The web-based Storage Center GUI, as shown in Figure 4, was used for the balance of the configuration. Seven intuitive wizard-driven steps later, a terabyte volume was being exercised on a Windows server through an iSCSI interface. Twenty clicks and six minutes after beginning the configuration, files had been copied from a local C: drive on a Windows server to a SAN attached F: drive under Storage Center management.

[2] Configuration details can be found in the Appendix.

FIGURE 4. STORAGE CENTER VERSION 4.0

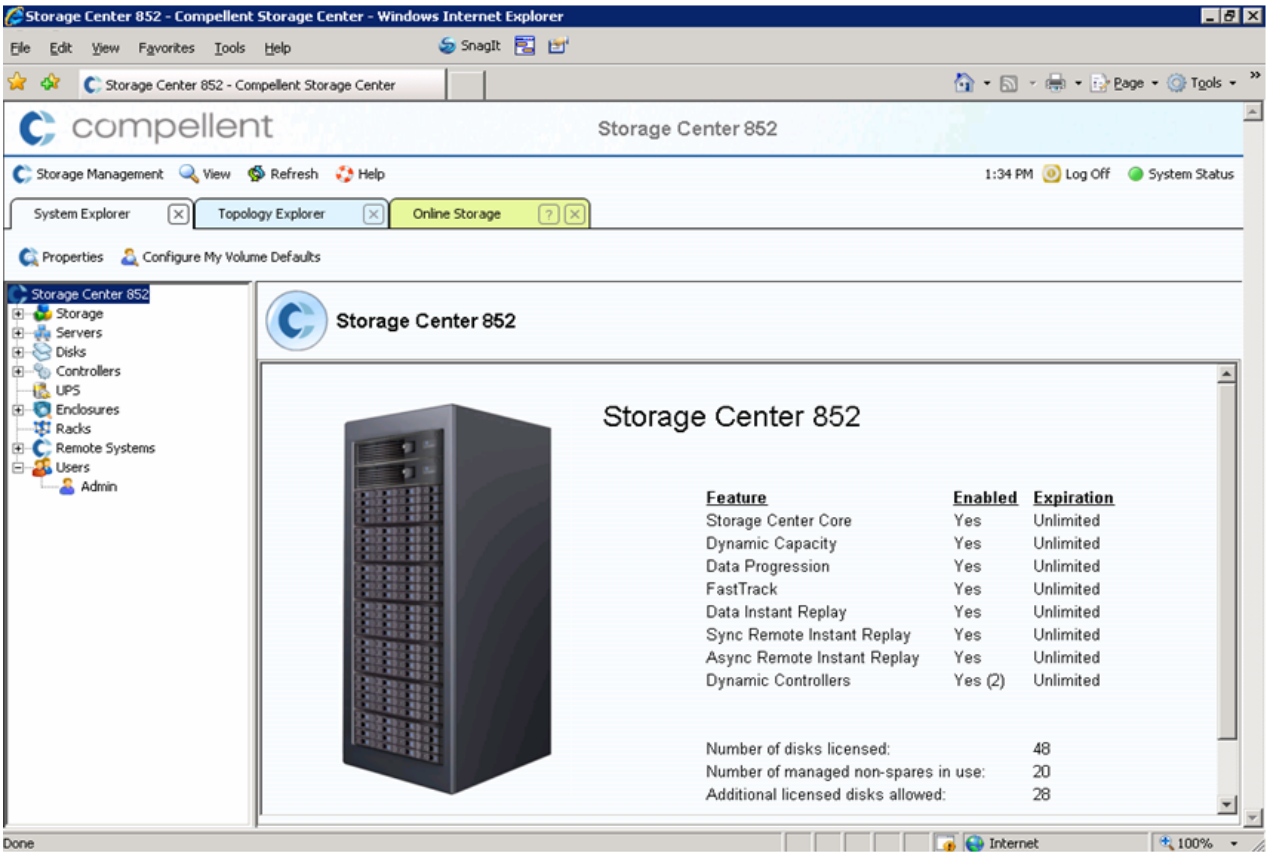

Initial configuration was followed by a tour of the Storage Center version 4.0 user interface. Navigation using the menu tree on the left was easy and intuitive. Excellent reporting is built into the base product at no additional cost. Optional modules include threshold notification and chargeback. A rich set of intuitive wizard-based capabilities were exercised, including the licensed data movement capabilities shown in Figure 4 and presented later in this report, Dynamic Capacity, Data Progression, Fast Track, Data Instant Replay and Remote Instant Replay.

Why This Matters

Organizations of all sizes are struggling with the growth of online storage capacity—especially small to medium sized organizations that can't afford a full-time storage administrator. Storage area networks enable consolidation and cost savings, but have historically been hard to configure and manage. ESG Research indicates that iSCSI technology has entered the early mainstream as a viable and cost-effective alternative to FC for creating a SAN. Compellent supports both (iSCSI and FC) in a single platform that is easy to configure and manage. ESG Lab's "time to first access" of six minutes and 20 mouse clicks is a good indicator of the simplicity and ease of use of Storage Center.

Fast Track and Data Progression

Compellent uses an innovative approach to store, manage and migrate data between tiers of storage. Unlike traditional architectures that manage at the volume or file level, Compellent manages pools of storage inside the volume at the block level. Attributes are tracked at the block level including time written, time accessed, frequency of access and RAID level. This metadata, or data about data, is used to automatically migrate blocks to the optimum tier of storage based on user-defined policies. This capability, which Compellent refers to as Data Progression, has been enhanced in Storage Center version 4.0 to support movement to and from a new tier of storage.

Disk drive vendors use a technology called zone bit recording to pack more information on the outer tracks of a hard drive. The Fast Track towards the outer edge of the drive contains more information than the slow tracks towards the middle of the drive as shown in Figure 5. More data packed on the outer edge of the drive enables more information to be accessed during each revolution of the drive.

Data Progression is used to define policies at the application volume level so that frequently accessed data is retained on high-performance media while infrequently accessed data is stored on lower-cost media. For example, Data Progression can not only be used to migrate frequently accessed application data from slower inner tracks to Fast Track, it can also be used to migrate infrequently accessed blocks from high performing drives (e.g., FC RAID-10) to a more affordable tier of storage for long term retention (e.g., SATA). In other words, with Storage Center version 4.0, Compellent has extended its ability to move data *between* tiers and now can be used to move data *within* a tier.

FIGURE 5. BLOCK-LEVEL AUTOMATED TIERED STORAGE

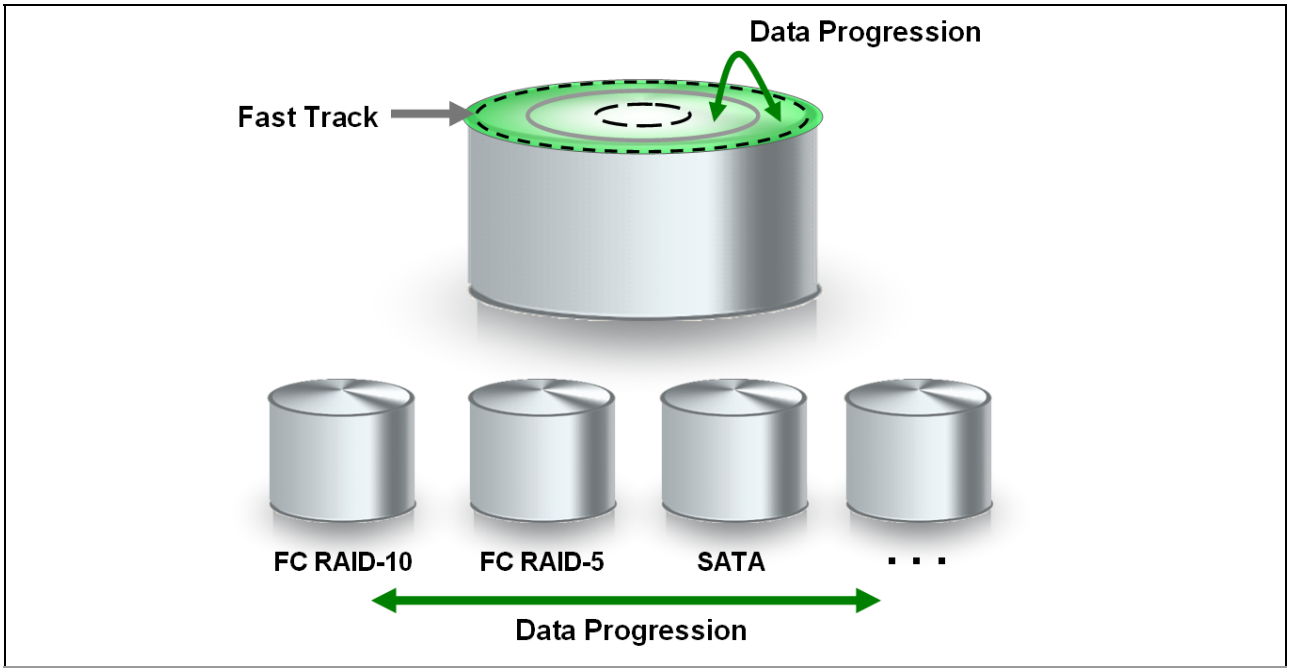

ESG Lab Testing

ESG Lab configured Fast Track and Data Progression for a test volume using the dialogue box shown in Figure 6. In this example, the test volume was configured to use multiple tiers of storage:

- Fast Track on 10K RAID-10 FC drives
- Standard tracks on 10K RAID-10 FC drives
- Replay Data (a.k.a. snapshots) on RAID 5 Tier 3 drives

This single user interface is all that was needed to enable automated block-level migration between fast and standard tracks. Storage Center kept track of usage and automatically migrated frequently accessed test volume data to the outer tracks. Adding another tier of storage for this volume was simply performed by clicking another checkbox. The Advanced Data Progression Settings link can be used to modify the default guidelines for automated migrations.

FIGURE 6. CONFIGURING DATA PROGRESSION POLICIES FOR A NEW VOLUME

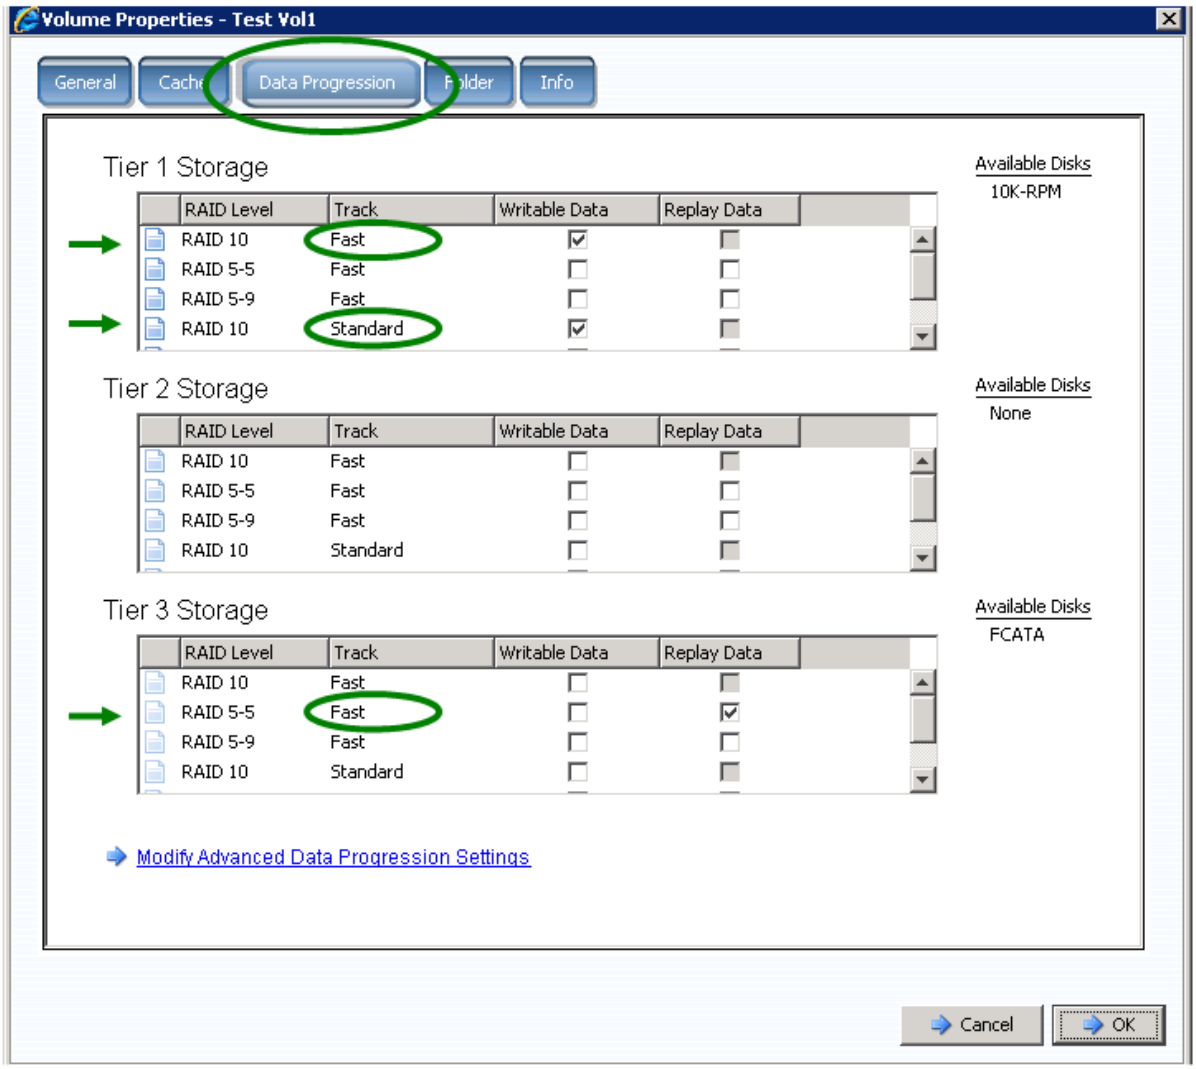

The next phase of testing was focused on quantifying the performance benefits that can be achieved with Fast Track. The drive specifications for a FC drive[3] and an audit of Compellent drive-level performance test results revealed that a volume entirely on fast outer tracks can perform approximately 25% better than a standard track volume for a simulated database workload. These results correlate well with the 23% to 25% benefit measured by researchers at the University of Southern California for a file system workload with all data on outer tracks of drives.[4] It's important to note that both of these results represent the maximum attainable result if all data were on the fast outer tracks of drives and none of the inner 'standard' tracks were used. In the real world, users want

[3] http://www.seagate.com/docs/pdf/datasheet/disc/ds_cheetah_15k_5.pdf
[4] Van Meter, Rodney. *Observing the Effects of Multi-Zone Disks*, http://www.usenix.org/publications/library/proceedings/ana97/full_papers/vanmeter/vanmeter/zcav.html

to utilize as much of the capacity of the drive as possible. This is the fundamental value proposition of using Fast Track with Data Progression; users get the performance of the faster outer tracks for the data that needs it while being able to use the rest of the drive for less performance critical data, automatically.

With the drive-level results in hand, an experiment was designed to quantify the performance boost that can be expected for real-world applications in real world configurations using Fast Track and Data Progression working together.

Microsoft Exchange e-mail traffic was evaluated based on the tools and guidelines specified by the Microsoft Exchange Solution Reviewed Program (ESRP). Microsoft ESRP testing uses a workload generator (Jetstress) that simulates the disk traffic created by a number of simulated e-mail users. Two configurations were tested as shown in Figure 7. The first was configured with the Exchange database file, located on a Compellent volume residing on standard (inner) tracks (the shaded area). The second test was designed to simulate the effect of automated Data Progression with Fast Track having moved the most active Exchange data to the faster outer tracks over time. Table 1 shows the configuration and results of this test.[5]

FIGURE 7. QUANTIFIYING THE FAST TRACK PERFORMANCE ADVANTAGE

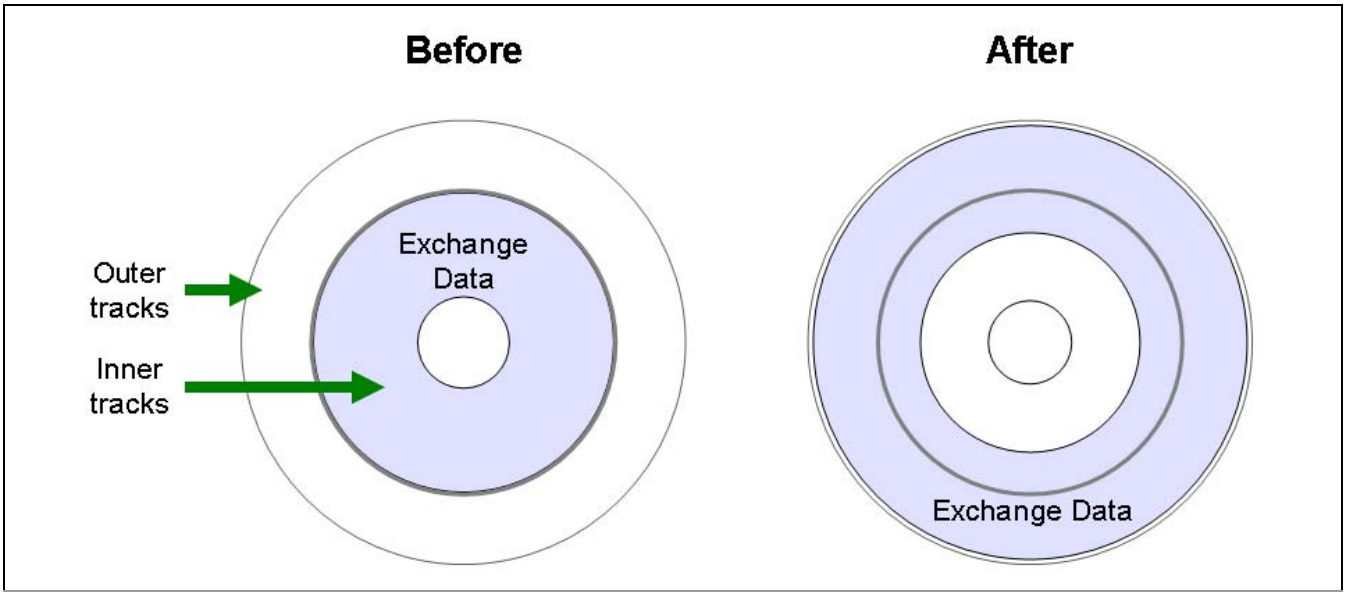

TABLE 1. FAST TRACK AND DATA PROGRESSION PERFORMANCE ANALYSIS

Configuration	Mailboxes	Mailbox Quota	IOPS	Response
Standard	1,500	250 GB	578	8 ms
Fast/Progression	1,500	250 GB	636	7 ms
			10% more	12.5% faster

ESG Lab observed a performance increase of 10% Exchange IOs per second (IOPS) and a 12.5% shorter response time. This increase was due to the most active parts of the Exchange database moving from the slower standard tracks to the faster outer tracks.

[5] Both configurations were tested using the same drive and Exchange configurations: 16 RAID-10 15 K FC drives, 1,500 mailboxes, 0.32 IOPs per mailbox, a 250 MB mailbox quota, 4 storage groups and 4 databases.

What the Numbers Mean

- The response times recorded are well below Microsoft's guideline of 20 milliseconds. Response times in excess of 20 milliseconds lead to users waiting to send and retrieve e-mails. Shorter response times increase productivity.
- Automated block-level movement of frequently accessed e-mail data to outer tracks by Fast Track software improved response time by 12.5% with no other tuning in a Microsoft Exchange simulation audited by ESG Lab.
- In the same simulation, the number of IOPS increased by 10% without having to increase the number of drives. Better performance from the same number of drives reduces the cost of meeting performance service level agreements.

Why This Matters

Mechanically spinning hard drives are the slowest link in the compute chain. Choosing the right drive technology to cost effectively meet application performance and capacity requirements can be extremely difficult with traditional storage system architectures—especially when trying to balance the cost effectiveness of the latest high capacity drives vs. a high number of drives working in parallel. Deploying extra drives to meet performance requirements increases cost, decreases utilization and increases the demand for power, cooling and space in the data center.

ESG Lab has confirmed that moving an application to Fast Track yields a performance benefit of up to 25%, which is similar to the performance *and cost* difference between 10K RPM and 15K RPM FC drives. Combining Fast Track with automated Data Progression makes applications perform faster over time (12.5% was measured during ESG Lab testing with Exchange workloads).

Fast Track and Data Progression deliver on the promise of information lifecycle management without the cost and complexity of migration software running on servers. With application data automatically migrating to the optimal tier of storage at the block level, organizations can purchase less storage capacity and decrease the load on the data center.

Dynamic Capacity

Dynamic Capacity is often referred to as thin provisioning in the industry. Compellent was one of the first in the industry to provide this valuable capability, which was architected into Storage Center from day one. Dynamic Capacity is used to present a virtual pool of shared capacity that is larger than the actual amount of physical storage available. It enables system administrators to deliver capacity on demand from a common pool of storage.

The difference between traditional provisioning and Dynamic Capacity is shown in Figure 8. In this example a server that requires 60 GB today is projected to grow to 500 GB over the next couple years. Traditional methods would require the full 500 GB be allocated today to avoid lengthy downtime for volume expansion or migration. Dynamic Capacity allows a 500 GB volume to be presented to the server with only 60 GB of physical capacity actually consumed. If and when the amount of capacity exceeds a pre-defined threshold, additional physical capacity can be added online with no impact to users and applications.

FIGURE 8. THE DYNAMIC CAPACITY ADVANTAGE

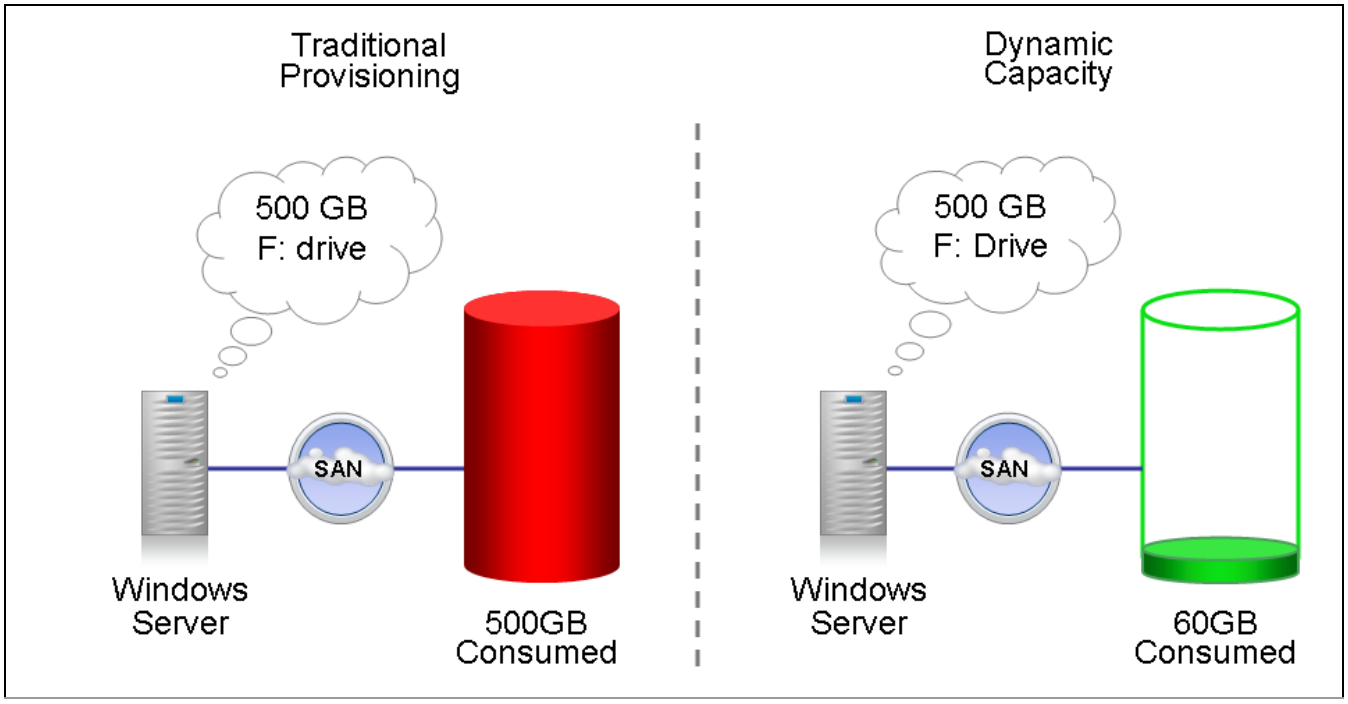

ESG Lab Testing

The first volume created by ESG Lab was a 500 GB iSCSI-attached Windows volume. The Compellent Storage Center wizard that was used to define the capacity of a new volume clearly stated the goals of Dynamic Capacity with the following instructions:

> Please enter the desired size of the volume to be created.
> The size will be allocated on an "as needed" basis from storage.

The tools and alerting capabilities for managing an over-provisioned pool of storage were examined after two more volumes had been created: a 1 TB FC-attached Windows volume and a 1 TB iSCSI-attached Linux volume. One of the screens and reports examined is shown in Figure 9.

FIGURE 9. DYNAMIC CAPACITY STORAGE SAVINGS

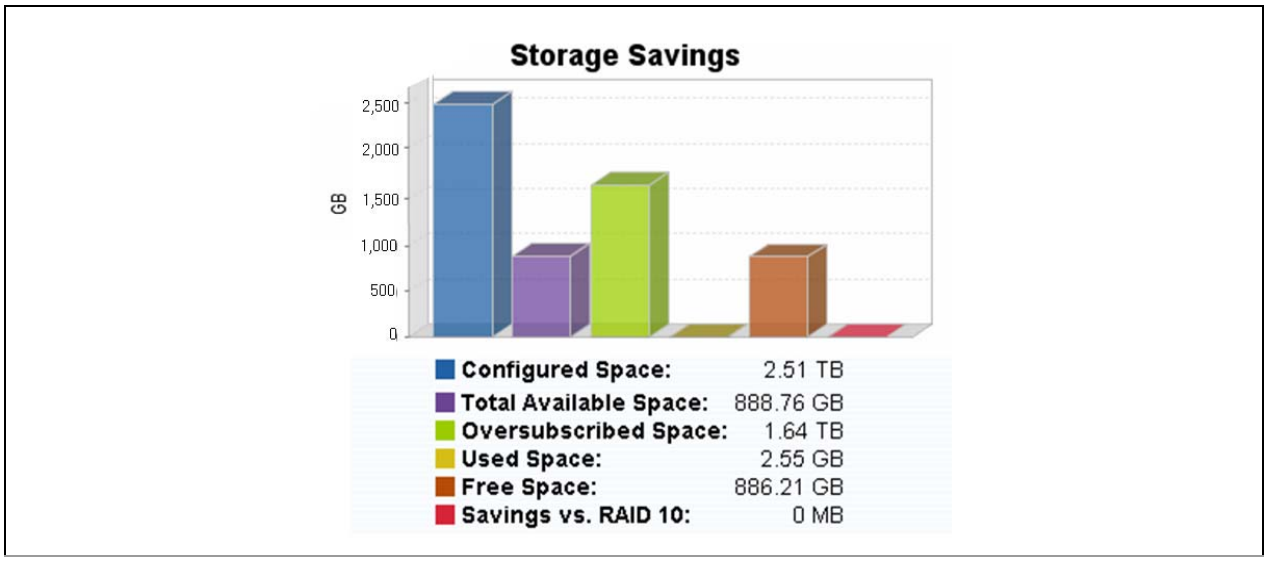

Note that 2.51 TB of space has been configured using only 888.76 GB of total available disk space. This yields 1.64 TB of oversubscribed space. Oversubscribed space appears to applications as if it is configured and available, yet requires no capacity until needed. The used space of 2.55 GB represents the actual capacity used at this point in the validation. In this case, 2.55 GB of file data was copied from a local hard drive to a SAN-attached pool of storage.

ESG Lab confirmed that Compellent has the reporting, alerting and Phone Home capabilities needed to ensure that physical storage capacity can be added non-disruptively before applications run out of actual physical capacity. Compared to a number of thin provisioning implementations that ESG Lab has examined recently, ESG Lab feels that Compellent's years of market experience dealing with the operating system and application-specific nuances of thin provisioning puts Dynamic Capacity in a class of solutions that ESG refers to as Thin Provisioning 2.0.

Why This Matters

ESG Research indicates that traditional storage provisioning can be complex, costly and wasteful. A majority of storage managers report that they have stranded and unused capacity and 80% feel that provisioning is a time and resource drain on their organization.[6] Underutilized storage capacity affects the bottom line in a number of ways including increased acquisition and management costs; accompanied by increased power, cooling and space demands in the data center.

A survey of early adopters of thin provisioning technologies, including a number of Compellent Dynamic Capacity customers, indicates that thin provisioning can be used to eliminate stranded storage, reduce capacity costs and simplify provisioning and capacity planning. Like the majority of early adopters of thin provisioning technology who report that they've had no issues or problems when implementing thin provisioning, ESG Lab found that Dynamic Capacity is easy to configure and manage.

[6] ESG Report: *Thin Provisioning Analysis and Research,* N = 30 in-depth interviews, November 2006

Data Instant Replay and Thin Replication

Data Instant Replay is used to provide disk-based recovery from deleted, corrupted or overwritten data. This technology, often referred to as snapshot or continuous data protection (CDP), provides space-efficient versions of application volumes that are frozen at a point in time. Unlike traditional snapshot technologies that use "copy on write" technology which consumes disk space and time, Compellent uses metadata and pointers to create Replays which are space efficient, fast and easy to manage.

Remote Instant Replay leverages similar technology to restore application data at a remote site after a site failure or disaster. After an initial copy of data has been sent to a remote site, only the blocks of data that have changed since the last Replay are sent to the remote site. The resulting WAN-efficient disaster recovery approach is referred to as Thin Replication.

FIGURE 10. SPACE AND BANDWIDTH EFFICIENT RECOVERY OPTIONS

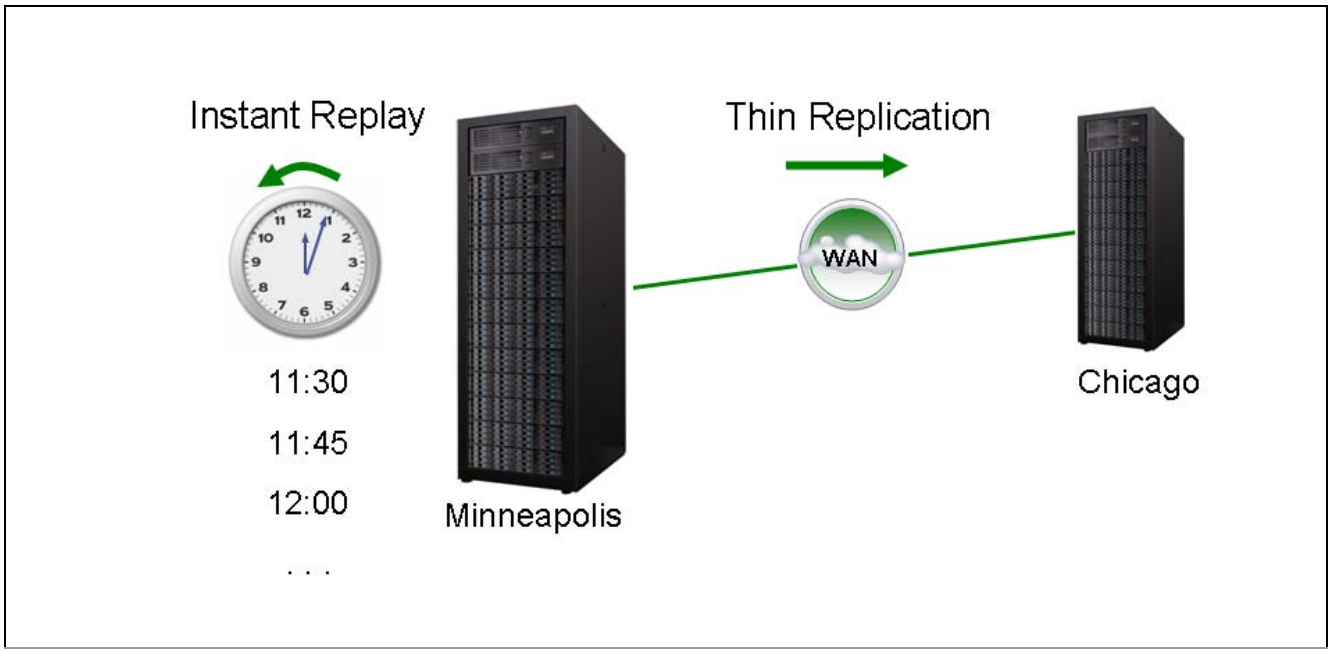

ESG Lab Tested

Data Instant Replay was tested using the Storage Center GUI and a FC attached volume accessed as drive F: from a Windows server. Files were copied from a local C: drive to the SAN-attached F: drive before the first Replay was configured. Replay configuration was performed using an intuitive Wizard-based Outlook-style interface from the web-based Storage Center GUI. No scripting was needed. Replay schedules and retention policies were easy to configure using a pre-defined Replay template. For example, a *Schedule of Daily every 12 hours between 12:05 AM and 6:00 PM* and a *Replay Lifetime of 5 days* were used during ESG Lab testing. Once a template had been created, policies could be applied to multiple volumes. It was noted that Replay policies can be defined at any time, even after a volume is in use.

A file was deleted to simulate a common user error. A point-and-click interface on the Storage Center GUI was used to initiate an Instant Replay. The Instant Replay from just before the simulated user error was mapped to another Windows server as an R: drive. The "lost" file was verified by inspection on the R: drive. It took only six mouse clicks and six seconds to complete the Instant Recovery and restore the "lost" file.

A number of valuable Compellent Data Instant Replay capabilities were noted:

- Read-only and writeable Replays are supported.
- Multiple Replay schedules and retention policies are supported within a single system.
- There is no limit on the number of Replays (as long as capacity is available for Replay data).
- Volume Shadow Services (VSS) integration provides reliable recovery of Microsoft applications including Microsoft Exchange and SQL Server.

Remote Instant Replay and Thin Replication
Remote Instant Replay was configured from the Compellent Enterprise Manager GUI. The Enterprise Manager GUI provides a single management interface for multiple Storage Center systems. In other words, Enterprise Manager is a "manager of managers."

Replication was configured between a system running the latest version of Compellent Storage Center (version 4.0) to a system running on a previous version of Storage Center hardware and software (version 3.6). A pair of Ethernet links between the two systems was used to emulate a pair of T1 WAN links between a primary data center in Minneapolis, Minnesota to a remote site in Chicago, Illinois. A Windows file system volume was configured for replication using the intuitive Enterprise Manager GUI. No scripting was needed. Remote Instant Replay was running six clicks and four minutes after beginning the configuration.

FIGURE 11. THIN REPLICATION AND BANDWIDTH THROTTLING

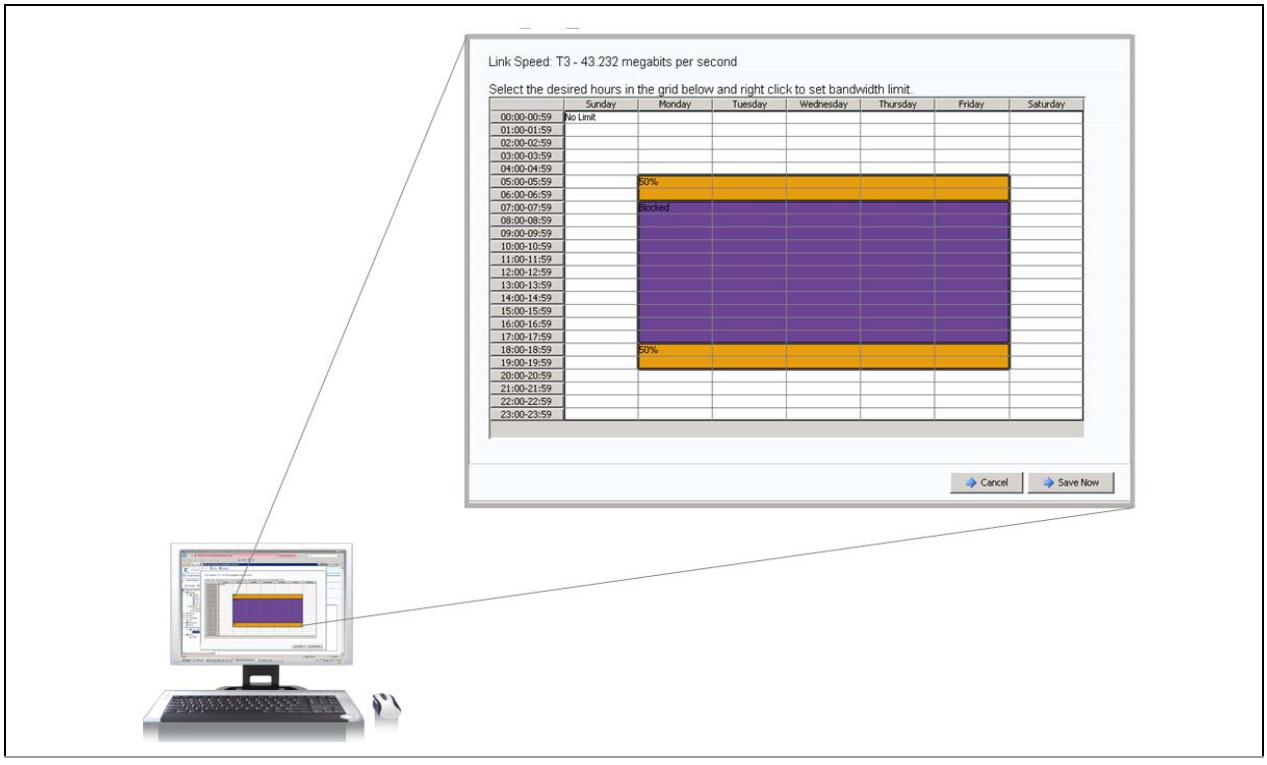

As seen in Figure 11, ESG Lab set bandwidth throttling by clicking and dragging across a grid made up of days of the week and hours of the day. Once the days/hours are selected, ESG Lab right-clicked on the selection to set the bandwidth throttle percentage. In Figure 11, the white area indicates no bandwidth limits, the orange area represents a 50% throttle and the purple area is throttled down 100%, completely blocking access during business hours.

A number of valuable Remote Instant Replay capabilities were noted:

- Replication over native IP networks reduces the cost and complexity of disaster recovery.
- Synchronous and asynchronous replication meets a wide range of disaster recovery objectives.
- A "Replicate to Simulation" option is supported to estimate WAN bandwidth consumption.
- After the initial replication, only changed blocks are replicated (Thin Replication)
- Configuring WAN bandwidth consumption using Quality of Service (QoS) settings is intuitive
- The click-and-drag user interface for throttling WAN bandwidth during peak business hours is extremely intuitive (Figure 11)

Why This Matters

Organizations are susceptible to a wide range of data hazards including viruses, human error, hardware failures, power outages and natural disasters. The downtime associated with data hazards can lead to lost revenue and lost productivity. As a matter of fact, recent ESG Research indicates that 63% of organizations report that they can't tolerate more than four hours of downtime for their most critical applications before experiencing significant revenue loss or other adverse business impact.[7]

Snapshot and remote replication technologies have become available in recent years as a quick and reliable alternative to recovery from a tape backup. Traditional snapshot and replication technologies can be difficult to configure and often require highly skilled administrators to perform complicated scripting.

ESG Lab has confirmed that Compellent Data Instant Replay technology is easy to configure using point-and-click wizards and templates. The block-based virtualization Compellent architecture reduces the storage capacity, WAN bandwidth and time required to create, access and manage an unlimited number of Instant Replays.

[7] ESG Research: *Data Protection Survey*, N = 398, 2007

ESG Lab Validation Highlights

- ☑ ESG Lab configured a pre-wired iSCSI and FC attached Storage Center configuration from scratch in 20 clicks and six minutes. Configuring and managing a pool of storage using the Storage Center web-based GUI was easy and intuitive.
- ☑ Fast Track, a feature new in Storage Center version 4.0, uses the tracks on the outside of drives to provide a performance boost, depending on application and workload.
- ☑ The combination of Fast Track and automated block-based Data Progression improved Microsoft Exchange database response time in simulations by 12.5% and yielded 10% more transactions with the same number of drives.
- ☑ Data Progression delivered the value of Information Lifecycle Management (ILM) without the cost and complexity of host-based data migration software. Data was automatically migrated at the block level based on policies that dictate the optimal cost and performance of application data over time.
- ☑ Dynamic Capacity was used to present a 1.5 TB of capacity to three servers using only 888 GB of actual disk capacity. The reporting, monitoring and alerting capabilities of the Storage Center GUI were reviewed to ensure that it's easy to plan for an online capacity upgrade.
- ☑ Data Instant Replay was used to restore a "lost file" using a space-efficient disk-based recovery image in six mouse clicks and six minutes
- ☑ Remote Instant Replay and Thin Replication were configured using the Enterprise Manager GUI in six mouse clicks and four minutes. A "Replicate to Simulation" option was noted for planning WAN bandwidth usage. Easy to configure bandwidth throttling and QoS capabilities were noted.
- ☑ A number of enhancements included in the Storage Center version 4.0 release were reviewed:
  - o 3rd generation hardware with faster processors, more memory and bandwidth
  - o Forwards and backwards compatibility for maximum investment protection:
    - Version 4.0 software runs on previous generation hardware
    - Version 4.0 can be replicated to systems running previous software versions
  - o Jumbo frames and VLAN tagging is supported for optimized iSCSI operation
  - o Variable page sizes are supported for optimized application performance
  - o A space recovery utility that runs on Windows servers reclaims deleted fragments on Dynamic Provisioned Windows volumes.
  - o Thin Import – an underutilized full provisioned volume on an existing SAN can be converted to a thin provisioned volume as it is imported into Compellent.

Issues to Consider

- ☑ Configuring Compellent feels different than configuring a traditional RAID system. In addition to placing drives in RAID groups, there's a vital process of creating pools and data progression policies. While the Data Progression user interfaces at the core of this process become familiar after a bit of practice, they might benefit from a bit of simplification or a new wizard.
- ☑ The Compellent hardware upgrade path tracks the latest advances in server, drive and interface technology. The 3rd generation hardware platform announced in conjunction with Storage Center version 4.0 sports an upgraded multi-core processor, increased memory and a faster bus. It's also "SAS-ready." While SAS drive support has not been announced, this is an important consideration for the future. SCSI drives have evolved into SAS drives and are emerging as a price/performance alternative to FC drives.

ESG Lab's View

Consolidation and savings have become a priority within IT organizations in recent years. A renewed focus on bottom line business results and service level agreements has led an increasing number of organizations to adopt server and storage consolidation strategies. Green initiatives and power and cooling concerns are also a problem. More and more, we are seeing buying decisions that begin with the statement, "How can we get more done with less?" Filling in the blanks for a storage solution, it's clear that IT managers want solutions that deliver more capacity, more performance, more flexibility and more resilience at less cost, less complexity and less overhead.

ESG Lab is impressed with the success of Compellent's channel-focused business model and the latest version of its flagship product, Storage Center version 4.0. ESG Lab testing revealed that Storage Center is rock solid and extremely easy to configure and manage. That's impressive given the innovative block-based virtualization and sophisticated storage services at the heart of the Storage Center architecture. Fast Track and automated Data Progression are a great example of the sophisticated simplicity provided in Storage Center version 4.0. ESG Lab testing confirmed that Fast Track provides a performance boost of approximately 25% using drives that Compellent customers already own. That's roughly equivalent to the performance and cost difference between 10 K and 15 K FC drives. Combining Fast Track with automated Data Progression yields a capability that ESG believes is unique in the industry—automated block-based migration of application data to the faster tracks on the outside of a drive, which makes applications run faster over time.

In conclusion, let's compare the evolution of storage infrastructure within an organization to the haphazard evolution of the housing market within a community in New England. To meet the varying needs of the community, the storage infrastructure has evolved to contain a mix of houses, apartments and dorms. High-end mission critical applications like e-mail and databases live in gated communities and mansions (expensive enterprise-class disk arrays with the fastest Fibre Channel drives), second-tier applications like file sharing live in apartments (mid-tier arrays with a mix of FC and SATA drives) and near-line applications live in dormitories (DAS, JBOD and bulk SATA arrays).

In contrast, the Compellent approach looks more like a well-planned community built around a state-of–the-art condominium complex. Sharing a common pool of infrastructure built from affordable commodity components, Storage Center reduces the cost of acquisition, maintenance, power, cooling and space. Residents are free to choose their own upgrades and control their own utility bills. Freed from cost and complexity of managing the outside infrastructure, residents can take advantage of customizable inside options (e.g., Fast Track and Instant Replay policies). The end result is a well planned storage community built around a highly virtualized storage complex, with a rich set of system and data management features on a reliable high performance platform that enables IT managers to cost effectively provide best-in-class storage services to the business community.

Appendix

Table 1. TEST CONFIGURATION DETAILS

Storage	
Compellent Storage Center (local)	Version 4.0, 4.4TB raw capacity, 2.1TB 10K FC drives, 3.3TB 7,200 RPM FATA drives
Compellent Storage Center (remote)	Version 3.6, 3.4TB raw capacity, 3.42TBD 7,200 RPM SATA drives
Servers	
Windows server #1	3.4GHz, 2GB RAM, Server 2003 O/S, Qlogic FC adapter, Qlogic FC driver
Windows server #2	3.4 GHz, 2GB RAM, Server 2003 O/S, Intel GigE interfaces, Microsoft iSCSI driver
Linux server #1	3.4GHz, 2GB GB RAM, Red Hat O/S, Intel GigE interfaces, Red Hat iSCSI driver
Switches	
Gigabit Ethernet	Extreme Networks Summit24, 24 ports
Fibre Channel	McData 4700, 24 ports @ 4 Gbps

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

www.enterprisestrategygroup.com