OPTIMIZING EXCHANGE SERVER IN A TIERED STORAGE ENVIRONMENT

WHITE PAPER | NOVEMBER 2006

compellent

compellent

EXECUTIVE SUMMARY

Microsoft ® Exchange Server is a disk-intensive application that requires high speed storage to deliver low latency transactional response time. When configuring storage for Exchange Server, administrators typically focus on transactional performance requirements and store all Exchange Server-related data on expensive high-performance disks. As with most enterprise applications, however, the majority of Exchange Server data is inactive and does not require the fastest spindles available to deliver high write performance.

The ability to automatically classify and migrate inactive portions of an Exchange Server database at the block-level onto lower-cost drives with Compellent's Storage Center™ is causing many organizations to rethink their Exchange Server storage platform. Using Compellent's Automated Tiered Storage solution, administrators can easily place Exchange Server databases on multiple tiers of storage, reducing the overall system cost while maintaining a high level of performance.

Automated Tiered Storage: Ideal for Exchange Server

Storage Center is the industry's only SAN with Automated Tiered Storage. Utilizing intelligence about blocks of data, gathered by Compellent's Dynamic Block Architecture™, Storage Center can automatically classify and migrate data onto the optimal tier of storage. Compellent's Dynamic Block Architecture records and tracks specific information about every block of data that provides the system intelligence on how that block is being used. Information about the blocks is gathered without system overhead but can be extensive, including time written, the type of disk drive used, the type of data stored, RAID level, and more.

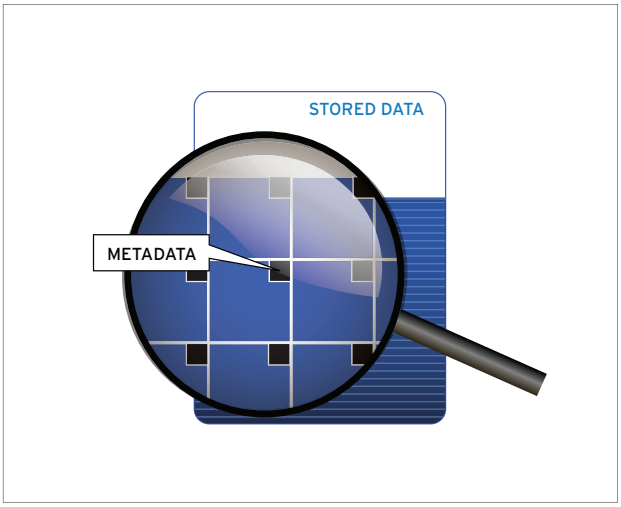

Figure 1: *Metadata provides Storage Center information on frequency of use, enabling blocks to be placed on the optimum tier of storage.*

This block-level intelligence is changing traditional views of how storage is managed, and how Exchange Server storage platforms are implemented. All of this metadata, or "data about the data," can be utilized to create a cost-effective, yet high-performance Exchange Server environment. Storage Center's Data Progression™ software uses this metadata to identify inactive blocks in an Exchange Server configuration and automatically migrate those blocks off an expensive storage tier to a lower cost tier of storage.

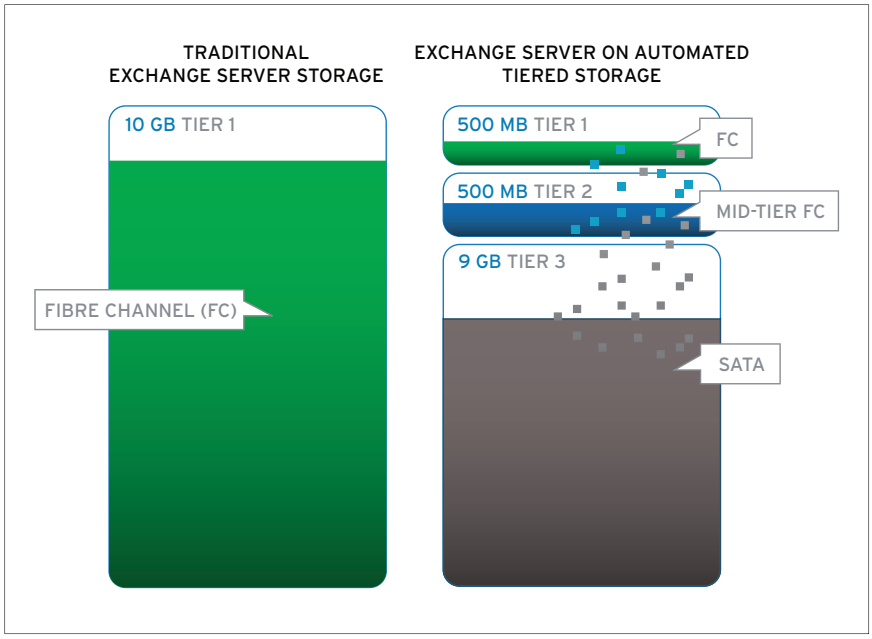

Figure 2: *10 GB in an Automated Tiered Storage environment can be split across three tiers. The 500 MB that stores the most recent data would be on Tier 1, while the majority of data including previous calendar items, older sent mail and contacts would be stored on Tier 3.*

Data Progression is ideally suited for use in an Exchange Server environment, resulting in several tangible benefits, including:

Increased User Load. Leveraging a second or even third tier of storage for inactive data can allow administrators to increase the number of recommended users without adding additional high-performance spindles. With Data Progression, increased user loads are a scalable benefit. As an Exchange Server configuration gets larger and user mailbox sizes increase, the amount of mail that ages and is then automatically migrated to a lower tier increases. More users with more mail can be managed with slower, cheaper spindles on a second tier instead of constantly expanding by purchasing faster more expensive spindles.

In tests using QuickStart ILM, an Automated Tiered Storage bundle designed for small to medium size enterprises, the recommended users for Exchange Server increased by more than 100 simply by moving infrequently accessed blocks of data to a lower tier. This increase was accomplished without the need for additional storage processing power, representing a significant advantage for a smaller Exchange Server environment.

Unlimited Snapshot Storage. Compellent automatically stores snapshots, called Replays on the lowest tier of storage. The lower tier of storage can be used to store an unlimited number of Replays without any impact on Exchange Server performance.

Cost-Effective Email Archive. For companies that are required to keep copies of email for several years to meet regulatory requirements, a lower-cost tier of storage provides an ideal online archive.

Performance Sizing Accuracy. Tiered storage allows companies to build an application tier with focused spindles to ensure appropriate performance without wasted storage capacity.

Automated Tiered Storage provides the foundation for a powerful Exchange Server solution, but is only one facet of an ideal storage platform for Exchange Server. Compellent's suite of software features like virtualization, continuous snapshots and thin provisioning seamlessly integrate with Exchange Server to enhance functionality.

Advanced Virtualization Increases Exchange Server Performance

A key factor in the performance of an Exchange Server solution running on Storage Center is Compellent's advanced virtualization and caching features. Compellent virtualizes at the disk-level within Storage Center, accelerating data access for Exchange Server by spreading read/write operations across all the disk drives in the SAN so multiple requests can be processed in parallel. Full disk parallelization increases I/Os and reduces hot spots in Exchange Server environments without necessitating performance tuning. Compellent's advanced approach to virtualization enables a single volume to span the multiple tiers of storage created with Data Progression.

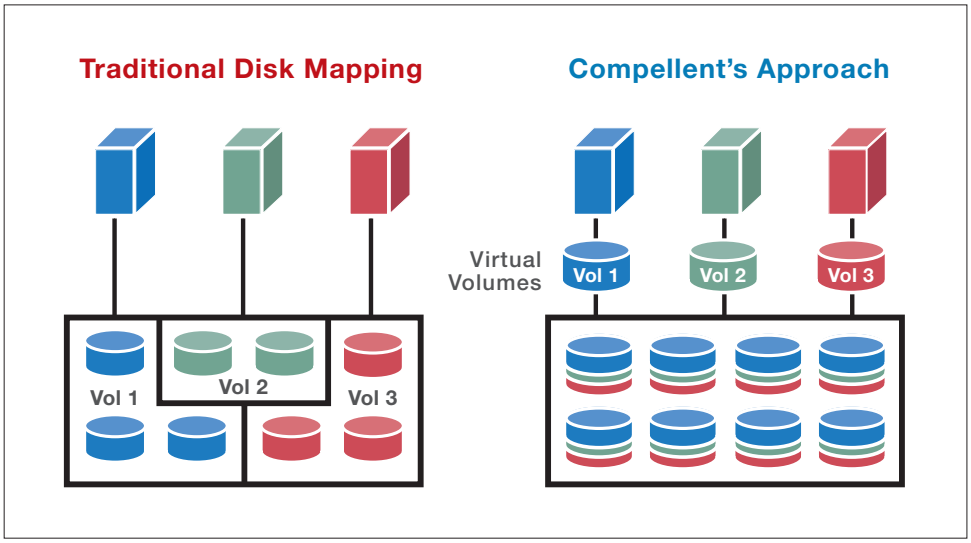

Figure 3: *Compellent's virtualization stripes data across all drives in a volume to increase performance and simplify configuration.*

Moss & Associates, the largest construction company in South Florida, uses Data Progression to automatically migrate inactive portions of users large .pst Exchange Server files to a lower-cost tier of storage. After installing Storage Center, Moss loaded all their data onto Tier 1 Fibre Channel and enabled Data Progression. Within 5-7 days the inactive data, which comprised nearly 80 percent of Moss's stored data, automatically migrated down to lower-cost SATA drives. Moving this inactive data off high-performance drives freed up Tier 1 spindles, increasing Tier 1 performance by 60 percent.

This advanced virtualization also simplifies configuring Exchange Server log data and database volumes, and optimizes performance for both. When configuring Exchange Server, it is important to keep log data on a volume optimized for sequential write performance and use caching intelligently for that volume to reduce write latency. At the same time, the database volume must be able to randomly read and write data efficiently while also reducing latency inherent in physical disk systems. By utilizing a single, virtual pool of storage, administrators can easily provide these divergent features to specific volumes within the same disk pool.

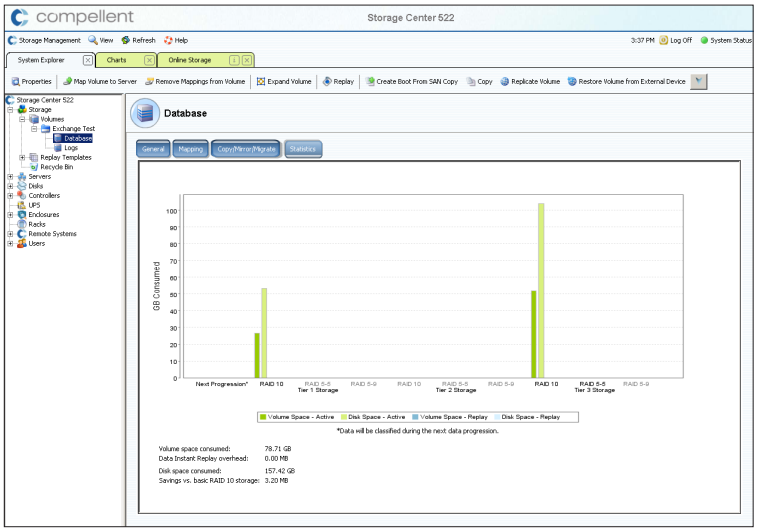

Figure 4: *Compellent's intuitive interface allows administrators to easily manage and view database LUN configurations. This screenshot shows volume space, disk space and RAID configuration using Automated Tiered Storage with data spanning Tier 1 and Tier 2.*

With Compellent, administrators do not have to be as concerned about isolating log I/O from database I/O because Storage Center uses data blocks, not individual RAID devices, as the core method of data access. This allows Compellent to better leverage all of the spindles within a specific configuration. Using a virtualized pool of resources also creates more inherent redundancy than typical 'dedicated spindle' systems and can reduce the complexity of managing multi-terabyte advanced storage systems. In addition, Compellent utilizes efficient data placement to enable faster seek times for Exchange Server.

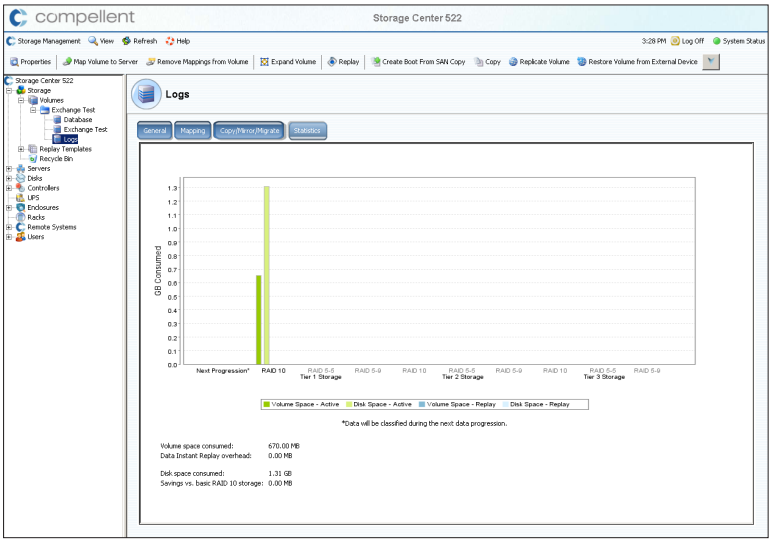

Figure 5: *This screenshot shows an Exchange Server Log LUN configuration in an Automated Tiered Storage Environment. Log files are placed on Tier 1 storage and utilize RAID 10 for performance.*

Continuous Snapshots Protect Exchange Server Data

Storage Center Data Instant Replay™ allows snapshots, or Replays, to be created in seconds on a production volume. Compellent's unique architecture allows the creation of unlimited Replays that consume minimal storage space. Replays are created without an initial clone and contain only changed data rather than the allocated but unused storage typical in other systems. This space-efficient design enables administrators to create and store Replays of Exchange Server volumes without consuming excess storage capacity or negatively impacting performance.

With Data Instant Replay, a large number of users can be supported within a single storage group with fast recovery capabilities. Using Data Instant Replay, administrators create Replays of a volume and restore a single mailbox from that copy within minutes instead of hours. Because Data Instant Replay stores delta changes instead of complete duplicate copies of a volume, it is possible to store a nearly unlimited set of Replays for records retention or to meet business Service Level Agreements (SLAs).

Although Replays only consume a small amount of storage space, every Replay is a readable and writeable volume that can be instantly mapped to any server. Each snapshot can be mounted to a backup server for backing up to tape or for mailbox recovery as needed. Mounting of the volume and checking for database integrity takes only a few minutes and does not affect the production Exchange Server. This allows for tape backups to occur as needed without the limitations of backup windows. In addition, Data Instant Replay can be configured to work with a second Storage Center system to allow for remote replication of data. This enables disk-only solutions to be used in production environments.

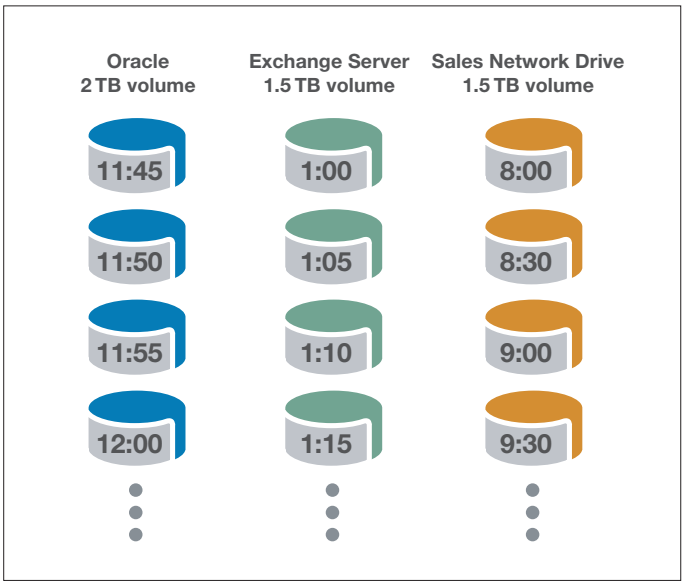

Figure 6: *Data Instant Replay delivers instant recovery from any point in time via unlimited space-efficient copies that consume minimal storage space.*

Replay Manager, combined with Data Instant Replay, ensures reliable backup of Exchange Server data. Using Replay Manager, customers avoid the difficulties caused by skipped files or data missing during backups due to files that are open or in use. Replay Manager provides the ability to take automated non-disruptive Replays of Windows Server volumes without taking Exchange Server offline.

Replay Manager seamlessly integrates with Microsoft's Volume Shadow Copy Service (VSS) to perform an open file Replay of Windows Server application data at a given point in time. Exchange Server can be running while backups occur – without the need for a specific backup window. When VSS receives a Replay request from the Replay Manager, VSS communicates with the Windows Server application to pause new transactions, finish all current transactions, and flush all the cached data to disk. VSS then communicates with the Replay Manager to initiate a Replay. Once a Replay has been created, VSS signals the Windows Server application to resume transactions to the disk.

Replay Manager, in combination with Data Instant Replay, ensures data remains consistent even when Windows Server applications are running during a Replay. Together they deliver the following features:

- Consistent, online backup and recovery of Exchange Server without volume clones or downtime
- Integrated with Microsoft's Volume Shadow Copy Service (VSS) and Virtual Device Interface (VDI) technologies
- Unlimited number of Replays at any time interval
- Multi-volume Replay support for VSS-enabled applications including Exchange Server
- Support for Microsoft Windows Clustering

Thin Provisioning Maximizes Exchange Server Utilization

In Exchange Server environments using traditional SAN platforms, up to 60 percent of disk space ends up allocated but unused. Compellent's Thin Provisioning, called Dynamic Capacity™, delivers the highest storage utilization rate possible by eliminating allocated but unused capacity. Dynamic Capacity completely separates storage allocation from utilization, enabling administrators to create Exchange Server volumes for future growth upfront yet only consume actual physical capacity when data is written by Exchange Server. In fast growing mid-range companies the ability to allocate storage and assign capacity for future growth without having to purchase the physical storage can represent a significant cost and administrative time savings.

Best Practices for Implementing Exchange Server

Many variables can impact Exchange Server performance. When implementing or troubleshooting Exchange Server on Compellent, administrators can utilize the following best practices:

Log Volume. Compellent's advanced virtualization removes the need to isolate spindles for performance. A single storage group can be used with a separate log file volume. Because of Compellent's disk virtualization technology, the same disk folder can be used for both volumes. This allows administrators to take advantage of RAID 10 performance and redundancy across all available spindles in the disk pool instead of limiting performance to specific spindles. This capability can increase performance of the log volume because Compellent leverages all of the spindles on Tier 1 storage by spreading the RAID 10 device across all available spindles.

Isolating Log I/O. Isolating log and database volumes to specific spindles is not necessary on Storage Center. Virtualization manages device redundancy and increases the ability to recover from disk failures. System administrators do not have to be as concerned about isolating log I/O from database I/O because Compellent Storage Center uses data blocks, not individual RAID devices, as the core method of data access. Instead, Compellent systems better leverage all of the spindles within a specific configuration.

Volume Sizes. Volume sizes are dependent on business SLAs. Using Microsoft VSS can allow for very large databases to be mounted to a recovery storage group in a short amount of time.

Disk Management. Disks are best used as a group – placing disks in one disk folder instead of managing them as separate groups is both easier to manage and more efficient on Compellent Storage Center. Unlike other systems there is not a penalty or limit to the number of disks that can be grouped into a virtual LUN.

Cache Settings. These settings will allow most efficient use of cache to support Exchange Server:

- Log volume: Write cache only
- Database volume: Read and Write cache enabled

Partitioning. In Compellent's internal system testing, the database and log volumes were partitioned using the Windows Diskpart command line utility as follows:

CREATE PARTITION PRIMARY ALIGN = 64

Internal testing using this offset has resulted in minimal performance differences but because of the I/O characteristics of Exchange Server it may prove beneficial in certain circumstances, therefore Compellent recommends this setting for Exchange Server volumes.

Troubleshooting. When troubleshooting performance the key items to review are disk latency, IOPS and data transfer rates. Administrators can follow these troubleshooting measures:

Latency

- Use Compellent Enterprise Manager to review server and volume latency. Look for a 20ms or less average latency on both read and write operations.
- Review the Microsoft System Monitor and look at the physical disk counter for "Avg. disk sec/read" and "Avg. disk sec/write". These should also be 20ms or less for the Exchange Server database and log volume.
- If the above indicators show longer latency times then consider adding spindles to the solution or removing load from the SAN when possible. Because of Compellent's modular design it is possible to easily add another disk enclosure rather than migrate to an entirely new system.

IOPS

- Leverage Compellent Enterprise Manager to monitor total IOPS used by the Exchange Server volumes as well as any other systems connected to the SAN. If the IOPS is increasing ensure that you have the correct spindle types to provide the needed IOPS.
- Microsoft System Monitor should be leveraged to review the "disk transfers/sec" counter for the physical disk object. Check these results against the Compellent Enterprise Manager numbers.

Data Transfer Rates

- Compellent Enterprise Manager provides information on disk drive transfer rates as well. Monitor drive transfer rates to optimize performance.
- Again, review the Microsoft System Monitor for the physical disk counter "Disk Bytes/Sec" and compare the results to Compellent Enterprise Manager.

For small to medium size businesses interested in implementing Exchange Server on Automated Tiered Storage, Compellent's QuickStart ILM Bundle is an ideal platform. A complete hardware and software solution out of the box, QuickStart ILM includes the following software components:

DATA PROGRESSION
Automates classification and migration of Exchange Server data and snapshots between tiers.

DATA INSTANT REPLAY
Allows Exchange Server administrators to mount any Replay to any server in less than ten seconds for quick recovery from viruses or human errors.

DYNAMIC CAPACITY
Allows Exchange Server administrators to size volumes upfront for growth without having the physical capacity in the SAN.

MANAGEMENT
One common, intuitive interface manages QuickStart ILM, reducing licensing and support costs, minimizing training and eliminating long hours of administration.

General Troubleshooting

- Review the Storage Center management interface for warning logs on a regular basis. Equipment problems, connectivity issues and other items can be diagnosed by looking at the Storage Center event log and system status indicators. If a failure of a component occurs, Storage Center's easy to use interface will highlight the component with a red or yellow warning indicator.
- Finally, Compellent CoPilot™ support is available 24x7 to assist if issues occur.

Cut Exchange Server Costs With Automated Tiered Storage

For most companies, an Exchange Server database represents perhaps the most significant amount of stored and inactive data. Within 5-7 days after receiving an email, most users no longer access attachments and calendar items. In fact for many, Exchange Server becomes a kind of online repository, and users don't have or take the time to delete inactive items. Because of this, the benefits of Automated Tiered Storage can be particularly significant in an Exchange Server environment.

For some companies, just migrating calendar data can result in significant savings. For many companies, using Automated Tiered Storage results in as much as a 74 percent drop in hardware costs alone. And by eliminating manual data classification and movement, Compellent's Data Progression cuts storage administration time in half, allowing companies to achieve the benefits of a tiered storage environment without the traditional costs or complexities.

COMPELLENT

12982 Valley View Road
Eden Prairie, MN 55344
Tel: 877-715-3300
Fax: 952-294-3333

www.compellent.com

Copyright © 2006 Compellent. All Rights Reserved. Compellent, Data Instant Replay, Dynamic Block Architecture, CoPilot, and Storage Center are trademarks or registered trademarks of Compellent. WPEXC111706