

FLUID DATA STORAGE: DRIVING FLEXIBILITY IN THE DATA CENTER

Eight Must Have Technologies for the IT Director.

EXECUTIVE OVERVIEW

Today's dynamic business landscape is no place for complex, inefficient storage systems with dead-end growth paths. The generation of data continues to explode, and ready access to that data is more critical than ever. Yet most enterprise storage solutions fall short, underutilizing data center resources, wasting IT staff time and leaving organizations at risk of unnecessary downtime.

All of these factors quickly eat into the bottom line. IT departments are forced to purchase, house, power and cool more hardware than needed. Administrators waste considerable time and effort managing complicated storage systems. And backup and recovery is slow and unreliable, leading to gaps in productivity and backlogs in business transactions. Quickly the costs associated with managing enterprise data can become overwhelming.

To compound matters, organizations continue to find themselves locked into proprietary storage platforms that can't be easily expanded or upgraded. This lack of flexibility and scalability leads to early obsolescence and costly forklift upgrades. Ultimately, administrators are left to coerce rigid storage systems to address core business needs, only to repeat the cycle with each inevitable system overhaul.

With Compellent, enterprise data is actively, intelligently managed throughout its lifecycle so you can constantly adapt to changing business conditions. Together, a Fluid Data architecture, storage virtualization, intelligent software and modular hardware bring unprecedented efficiency, simplicity and security to enterprise storage. Data center resources are utilized to their full potential. Built-in intelligence and automation ensure data is available when and where it's needed. And an open, persistent hardware platform scales in line with business needs to protect storage investments over the long term.

This paper discusses eight Compellent technologies that combine to create new levels of efficiency and flexibility while cutting cost now and in the future:

- 1. Fluid Data Architecture** – Storage is managed at the most granular level with built-in system intelligence to enable the dynamic flow of enterprise data.
- 2. Storage Virtualization** – Storage is virtualized at the disk level to create a flexible pool of storage resources shared by all servers all the time.
- 3. Thin Provisioning** – Allocation is completely separated from utilization so any size volume can be created at any time, yet capacity is only consumed when data is written.
- 4. Automated Tiered Storage** – Data dynamically cascades from tier to tier according to actual usage, freeing up high-performance drives for mission-critical applications.
- 5. Space-efficient Snapshots** – Continuous snapshots only capture changes in data for real-time protection with instant recovery to any point in time.
- 6. Thin Replication** – Data is replicated between local and remote sites using space-efficient snapshots and native IP connectivity, eliminating the need for high-speed data links or identical system configurations.
- 7. Unified Storage Resource Management** – All storage resources are managed through a single, point-and-click interface providing a complete view of the entire storage environment.
- 8. Open, Agile Hardware Platform** – Storage is designed for persistence, not obsolescence, leveraging a single, modular hardware platform coupled with technology independence.

1. A Fluid Data Architecture Enables the Flow of Enterprise Data

The underlying problem with most storage systems is that they manage data in complete volumes, restricting system flexibility and performance from the outset. To realize storage efficiencies and reduce data center costs, you need to manage data at a more granular level – within the volume itself.

Compellent storage is enabled by a dynamic Fluid Data architecture that actively, intelligently manages data at the block level. Specific information about each block is tracked and recorded to provide real-time system intelligence for dynamically storing, migrating and recovering data. This metadata, or data about the data, is gathered using minimal system overhead but can be extensive. Examples include the type of data stored, disk drive used, RAID level, time written, frequency of access and more.

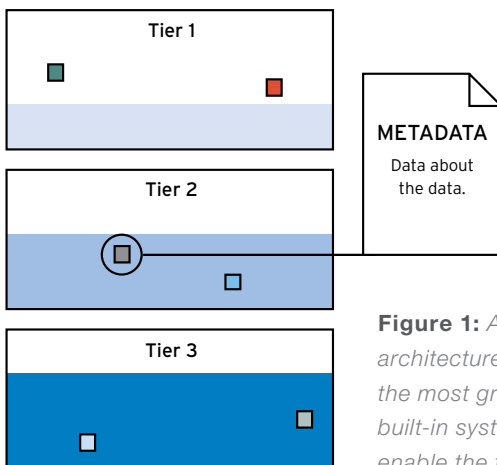


Figure 1: A Fluid Data architecture manages data at the most granular level with built-in system intelligence to enable the flow of data.

Empowered by this intelligence, built-in storage automation optimizes the provisioning, placement and protection of data throughout its lifecycle. Business applications are implemented faster, information to make decisions is always available, new technologies are instantly deployed and data is continuously protected.

2. Storage Virtualization Creates a Flexible Pool of Storage for All Servers

Managing data at the block level also enables Compellent to virtualize storage at the disk level. Storage virtualization not only consolidates resources and reduces disk costs, but dramatically increases system flexibility. Administrators no longer need to allocate particular drives to specific servers. Instead, an




Figure 2: Storage is virtualized at the disk level to create a flexible pool of resources shared by all servers all the time.

advanced storage virtualization engine creates a shared pool of storage spanning all the drives in the system. Storage is presented to servers simply as capacity, regardless of disk type, RAID level or server connectivity. In short, all storage resources are available to all servers all the time.

With Compellent, virtual storage volumes can be provisioned and mapped to servers in seconds. The system even expands volumes automatically when an application requires more storage. Administrators don't need to waste time with complicated capacity planning, or even take the system offline for modifications. This ability to change, shift or dynamically scale storage on demand minimizes disk expenditures, saves staff time and ensures continuous availability.

Since read/write operations are spread across all drives, multiple I/O requests can be processed in parallel. This eliminates the "hot spots" associated with conventional storage systems while expediting data access and boosting application and OS performance. As more capacity is added to the pool, data is automatically restriped across all available drives for hands-free load balancing and performance tuning.

Storage virtualization amplifies the benefits of server virtualization. In fact, users can quickly create hundreds of virtual volumes to support any virtual server platform and optimize the placement of virtual applications without wasting time, money or disk space.

3. Thin Provisioning Fully Optimizes the Utilization of Disk Capacity

Compellent significantly reduces the cost of storage by enabling you to purchase and manage fewer disk drives now and in the future. With other storage systems, physical disk capacity is pre-allocated when the volume is created. Administrators estimate how much capacity may be required for a given application and

allocate “extra” space to accommodate growth. If the volume created is 500 GB, all 500 GB are set aside for that application. No other applications can use any of the pre-allocated disk space, and none of it can be reclaimed later if actual utilization doesn’t coincide with staff estimates. In most cases, only a fraction of the pre-allocated capacity is ever actually used, resulting in the accumulation of purchased but “stranded” storage.

Such inefficient disk utilization inflates capital expenditures, operating expenditures and, ultimately, your total cost of ownership (TCO). Administrators are forced to buy more capacity than needed upfront, when the price per GB is sure to fall. Over time, as capacity is consumed (or stranded), even more capacity must be purchased, further expanding the data center footprint. And all of this storage must be provisioned manually, a time-consuming process that often requires downtime. In the end, regardless of how much data is truly stored, all of these disks require continuous power and cooling.

Compellent Thin Provisioning software, called Dynamic Capacity, completely separates allocation from utilization, eliminating pre-allocated but unused capacity. Administrators can provision any size virtual volume upfront yet only consume physical capacity when data is actually written to disk. That means you purchase the data you need to store your data today, then continue saving by expanding system on demand, adding the right capacity at the right time as your business needs change. In most cases, organizations can regain 40 to 60 percent of disk space that would have been lost to pre-allocation. You can even reclaim capacity from volumes provisioned with legacy systems using Thin Import.




Figure 3: Thin Provisioning separates allocation from utilization so capacity is only consumed when data is written to disk.

4. Automated Tiered Storage Dynamically Classifies and Migrates Data

To continue containing costs throughout the lifecycle of enterprise data, Compellent leverages an innovative data movement engine that integrates intelligent tiering with advanced virtualization. Traditionally, information lifecycle management has been a tedious and manual process with no end. Data is continually classified and reclassified based on its “importance” to the organization, an approach rife with political implications.




Figure 4: Automated Tiered Storage dynamically classifies and migrates data to the optimum tier based on frequency of access.

Politics aside, manually moving data between high-performance drives and more cost-effective, capacity-oriented drives is complicated and time consuming. Add-on migration engines can help, but increase overall software costs and waste valuable staff time for systems integration. Either way, each volume must be moved in its entirety, although some of that data is probably more frequently accessed. That means administrators have to continuously fine-tune data placement over time. Still, the alternative – retaining all enterprise data on high-performance drives – is not only costly in terms of disk expenditures, but wastes energy and squanders limited data center space.

With Compellent, since data is actively, intelligently managed at the block level, manually moving data between tiers is a thing of the past. Using unique Automated Tiered Storage software, known as Data Progression, Compellent dynamically classifies and migrates data to the optimum tier based on actual usage. The “importance” of data becomes a matter of fact according to actual usage, and data placement organically remains in tune with changing business needs.

Data is written to high-performance Solid State (SDD) or Fibre Channel (FC) drives on Tier 1 storage. Then, as the frequency of access declines, the less active blocks of data migrate to FC or SAS drives on Tier 2 storage. Over time, completely inactive data moves to high-capacity SAS or SATA drives on Tier 3 storage. To further free up high-performance drives for mission-critical applications, snapshots automatically flow to cost-effective, energy-saving drives. Meanwhile, the most active data is dynamically placed on the outer tracks of each drive for increased performance. Since most enterprise data is inactive, on average organizations can reduce disk expenditures by as much as 80 percent with Automated Tiered Storage.

Regardless of the current tier, with Compellent storage, enterprise data always remains readily available. Even once-inactive data is promoted to a higher tier if it becomes regularly accessed again. Administrators simply customize the tiering algorithm according to specific organizational needs, or use default settings established based on current industry best practices. In fact, with Automated Tiered Storage, your staff can spend half as much time managing data than with other systems.

5. Snapshots Provide Continuous Data Protection and Instant Recovery

Protecting enterprise data from server failures, viruses, power outages, human error and other inevitable circumstances is critical to continued business success. Today's organizations simply can't remain competitive without the ability to quickly and accurately recover "lost" or inadvertently deleted data.

Using conventional backup technologies is slow, unreliable and consumes excess storage space. Creating continuous snapshots protects data more effectively, but most snapshots still require full mirror images and subsequent full-volume clones. RAID overhead only adds to the problem. Such a heavy burden on disk resources significantly hinders system performance, and full-volume clones delay recovery. Also, because most systems provision disk space so inefficiently, the number of snapshots that can be taken is limited – usually 10 per volume. Limited snapshots mean fewer time-interval recovery points, putting mission-critical applications at risk of extended downtime.

With Compellent, snapshots are space efficient, continuous and unlimited. Data Instant Replay software doesn't require a




Figure 5: Space-efficient snapshots capture incremental changes in data for real-time protection with instant recovery to any point in time.

full mirror image or subsequent full-volume clones. In fact, only changes in data since the last snapshot, or Replay, need to be captured, dramatically reducing the amount of disk space required. The integrity of enterprise application data spanning multiple virtual volumes is protected using Consistency Groups. And Compellent storage doesn't restrict the number of Replays that can be taken. That means changes in data can be captured continuously in short intervals for the utmost level of security using minimal capacity.

Since Replays are so space efficient, recovery takes less than 10 seconds and volumes can be mapped to any server without disruption. To simplify the process, administrators set up Replay schedules using an intuitive point-and-click interface. Replays can also be used to test new applications and service packs on actual data without risk, efficiently support boot-from-SAN operations and virtually eliminate backup windows on production systems.

6. Thin Replication Protects Multi-Site Data with Configuration Flexibility

Due to the high costs and complexity traditionally associated with off-site backup strategies, remote replication continues to be a low priority for many organizations. That's because replicating data between locations typically requires identical site configurations and expensive, high-speed data links. And even then replication can be slow and unreliable, especially when the process requires the transfer of full-volume clones.

Compellent Thin Replication technology, Remote Instant Replay, provides affordable, verifiable and simplified multi-site data protection leveraging space-efficient snapshots called Replays. Following initial site synchronization, only incremental changes in data need to be replicated. This cuts hardware costs, reduces bandwidth requirements and significantly expedites recovery. Compellent storage can maintain an unlimited number of

Replays without affecting system performance, and the Replays can easily be mapped to any server for instant recovery to any point in time.

Since Compellent storage intelligently transfers only changed blocks of data, bandwidth utilization is optimized, obviating the need for high-speed connections. Built-in bandwidth simulation and shaping further streamline the process. Transfer rates can be customized based on link speed, time of day and replication priority.

Another advantage of Thin Replication is that it doesn't require identical system configurations at each site. That gives you the flexibility to choose lower-cost SAS or SATA drives for remote sites. To further maximize efficiencies, Compellent storage features a built-in Fibre Channel-to-iSCSI converter, enabling administrators to replicate data natively over existing IP networks. Users can also expedite initial replication synchronization using Portable Volume, secure enterprise-class external hard drives preconfigured for use with Compellent storage.

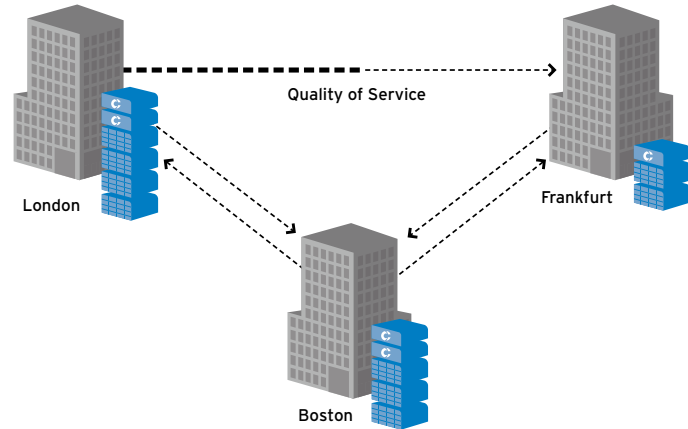


Figure 6: Thin Replication protects data between local and remote sites using space-efficient snapshots and native IP connectivity.

7. A Unified User Interface Simplifies Storage Resource Management

Managing conventional storage systems is complex and time consuming. In fact, IT staff expenditures often exceed the cost of the storage platform itself. Administrators need to monitor always-changing capacity requirements, manually migrate data from tier to tier, configure backup sequences and more. And in most cases, all these tasks must be performed using multiple, standalone user interfaces.



Figure 7: An intuitive, point-and-click user interface simplifies storage resource management through a single pane of glass.

Compellent storage is designed to help administrators manage more data in less time. This is largely because of the built-in efficiency and intelligent automation of Compellent storage. It is also because Compellent storage features an intuitive, point-and-click interface that provides a complete view of the entire storage environment through a single pane of glass.

With Compellent, there is no need for specialized skill sets or ongoing systems training. Wizards guide users through system setup and application configuration, making even advanced operations simple and straightforward. Since common, repetitive (and often time-consuming) tasks are fully automated, users can focus on other important projects. Storage consumption and usage trends are automatically monitored and displayed, eliminating the need for manual capacity planning. And a unique Phone Home feature provides automated real-time alerts and notifications for remote diagnostics and monitoring. Compellent storage even generates executive summaries, cost-savings calculations and utilization chargeback reports with just a few clicks.

8. An Open, Agile Hardware Platform Scales on Demand Without Limit

Unlike systems that require you to rip and replace hardware as your business needs change, Compellent supports the continual adoption of new technologies on a single, modular platform. This is a major departure from the status quo, which imposes artificial restrictions on everything from drive type to server interconnect. Simply put, most storage systems are designed for early obsolescence, forcing organizations into costly forklift upgrades.

Compellent storage is designed for persistence. Administrators can scale from entry-level to enterprise on the same platform (from 2 TB to 1,000 TB in the same system) incrementally in line with business needs. They can also deploy any combination of FC and iSCSI server interfaces, as well as mix and match SSD, FC, SAS and SATA drives in the same system. Even SAS drives with varying capacities and rotational speeds can be used in the same enclosures. You can seamlessly adopt new and emerging technologies on the fly as they become available. And upgrades, port modifications and other changes in configuration can be implemented on demand without disruption.




Figure 8: A single, modular hardware platform scales from entry-level to enterprise with technology independence.

Compellent storage also features fully redundant hardware and advanced failover capabilities. Clustered controllers, each with redundant fans and power supplies, operate in unison for optimum system performance. Yet each controller connects to enclosures and drives independently to ensure no single point of failure. Controller port virtualization and dual paths from servers to disk drives also enhance availability, with built-in multi-path I/O failover eliminating the need for custom software.

With Compellent, you have the freedom to choose the technologies that support your IT infrastructure today and readily adapt to change tomorrow – all without disruption. And you don't have to repurchase your base software license when you upgrade controllers to integrate new technologies.

CONCLUSION

Organizations no longer have to be confined by the rigid boundaries and perpetual expense of conventional storage systems. By actively managing data at a more granular level, Compellent provides dynamic, virtualized storage that readily adapts to constant change. The utilization of storage resources is optimized to cut storage costs now and in the future. Built-in intelligence and automation simplify and streamline data management, saving staff time and minimizing human error. Advanced, space-efficient snapshot technology protects against downtime for continuous availability without squandering storage resources. And an open, agile hardware platform safeguards storage investments over the long term by remaining technology independent and scaling without limit. By leveraging Fluid Data in the data center, organizations can significantly lower the cost of managing enterprise data throughout its lifecycle – and readily adapt to constant change.