

THE ARCHITECTURAL ADVANTAGES OF COMPELLENT AUTOMATED TIERED STORAGE

How Integration, Granularity and Real-time Intelligence Differentiate One
Sub-LUN Tiering Solution

EXECUTIVE OVERVIEW

With the recent profusion of tiered storage solutions into the enterprise storage market, organizations are increasingly adopting tiering strategies to accommodate rapid data growth while controlling costs. Almost all major offerings leverage multiple types of drives, with varying capacities and performance levels, and attempt to match data with the optimal drive type based on storage profiles. Typically, mission-critical data is stored on high-performance drives, while “less important” data is moved to high-capacity drives.

Yet most of those solutions fail to maximize the potential benefits of tiered storage. Poor integration of tiering into the storage platform; a lack of granular, real-time system intelligence; labor-intensive, inefficient data movement schemes; and limited scalability keep many solutions from meeting real-world business needs. Applications are tacked onto existing storage infrastructures. Data is migrated as large “all or nothing” pages, often according to outdated information. The tiering profiles are limited in scope, generally requiring manual intervention. Data must be reformatted from a disk group to a pooled configuration. And solution licensing is restricted to a particular hardware platform. Such limitations completely offset the cost- and time-saving benefits of truly virtualized storage.

Unlike the many newcomers to the storage tiering market, Compellent® engineered its Fluid Data™ architecture from the ground up to include Automated Tiered Storage. This approach enables organizations to move data dynamically, intelligently and efficiently among multiple storage tiers and RAID levels. All write transactions and frequently accessed data are placed on high-performance drives with performance-optimized RAID levels, while less frequently accessed data cascades to more cost-effective drives and/or lower-overhead RAID levels.

Since Compellent Automated Tiered Storage, called Data Progression™, is built right into the virtualized storage platform, it does not require additional hardware or server-side agents to operate. Automated Tiered Storage also integrates seamlessly with the full range of other Compellent technologies to maximize storage performance, efficiency, flexibility, reliability and manageability. In addition, because Automated Tiered Storage is not tied to just one storage system model or platform,

organizations can take advantage of the functionality without downtime or a costly forklift upgrade—and continue to benefit as they scale the solution in line with changing business conditions.

To realize the full potential of Automated Tiered Storage, organizations need a solution that:

1. **Is built into a virtualized storage platform, not added as a separate agent**
2. **Manages and migrates data as highly granular 512KB, 2MB or 4MB blocks**
3. **Utilizes in-flight use characteristics for ongoing, real-time intelligence**
4. **Provides hands-free management through policy-based automation**
5. **Integrates seamlessly with snapshots, delivering high write and read performance**
6. **Leverages other enterprise functionality to maximize tiering benefits**
7. **Offers platform independence and scales on demand without limits**

1. Built-in virtualization provides foundation for automated tiering

A virtualized—or pooled—storage environment provides an ideal foundation for tiered storage. With a truly virtualized environment, there are no constraints on where data can reside, as data is not confined to conventional disk groups. Data can be moved without limitation to a tier with a particular drive type or performance level, or to a particular RAID level within a storage tier, depending on an organization’s needs.

Unlike static storage architectures, Compellent Fluid Data storage was architected to provide truly virtualized storage that spans all disks in the storage environment. Compellent virtualizes storage at the disk level, creating a dynamic, shared pool of storage resources available to all servers, all the time. All storage types and RAID levels are pooled by default. Administrators can keep the entire array virtualized as a single pool of high-performance storage—the most common approach—as well as configure particular subsets of the array as individual pools, if desired, according to enterprise needs.




Figure 1: A Fluid Data architecture enables true storage virtualization, providing a single pool of storage resources that spans all disks and RAID levels.

Other storage architectures fall short of true storage virtualization. Those solutions pool only a subset of the storage environment by stretching a volume, or LUN, across a particular set of drives. Administrators are required to predefine the ratio for each storage tier (or drive type) within each pool. The pool is confined to that set of drives, leaving little flexibility for optimizing storage later.

To take advantage of sub-LUN tiering with other solutions, administrators must provision virtual volumes within a distinct pool of disks upfront or convert existing volumes to virtual volumes, which are then pooled across specific disks in the storage environment. That conversion process creates

downtime as volumes are reformatted and temporary volumes are created, and it requires twice the existing storage capacity for the migration to the pooled configuration.

With Compellent, sub-LUN tiering is not confined to a distinct set of drives with predefined storage ratios for each tier.

2. Granular data management optimizes migration

A Fluid Data architecture is what enables true storage virtualization and facilitates a highly granular approach to storage tiering. Compellent divides data into pages—preformatted collections of allocated and unallocated disk blocks. The pages are 2 MB by default, though administrators have the option to manage storage as 512 KB or 4 MB pages, and details about each block are captured in action. This system intelligence is collected at the 512 byte level—the smallest addressable data block available. These tiny blocks are what make up the pages.

This granular approach optimizes tiering. Using real-time system intelligence, the Compellent SAN identifies very small data blocks that are eligible for movement based on frequency of access and then moves just the small pages that include those blocks to the appropriate tier or RAID level, creating a finely tuned tiered storage environment. Using small pages also increases the efficiency of data movement since it is more efficient to read, write and migrate small pages than larger ones.

Static storage architectures are limited to much larger blocks of data, ranging from 16 MB to 1 GB. Even if only a very small quantity of data in the block changes, making it eligible for migration, the system has to move the entire, larger block. The efficiency of the Compellent architecture can move and place data with a greater level of precision within the tiered environment.

A Fluid Data architecture allows data to be migrated in small 512 KB, 2 MB or 4 MB blocks, not large 16 MB to 1 GB blocks like other tiering solutions.

3. Real-time intelligence facilitates precise data movement

Compellent maintains continual awareness about small data blocks and captures real-time use characteristics about each block. Collected transparently in flight, these use characteristics include information on when blocks were created, which drives

hold the blocks, the associated virtual volume, how frequently the blocks are accessed or changed, and whether the blocks represent actual data or virtual pointers to data. This information provides the intelligence to determine whether and when blocks of data should be moved from one storage tier, or one RAID level, to another. Because the Compellent architecture operates at such a granular level, this block-level intelligence requires negligible system overhead to ensure data always stays in tune with application needs.




Figure 2: Fluid Data storage continuously collects use characteristics to ensure that data migrations are based on real-time information.

When it is time to migrate data (typically once per day), pages eligible for movement are automatically transitioned to either a new storage tier or a new RAID level within the existing tier. Moving data to a new storage tier (from one type of drive to another) frees up capacity in the high-performance tier. Moving data to a new RAID level in the same tier helps sustain read performance while reclaiming capacity. Typical storage architectures limit which blocks can migrate for the optimum balance of performance and cost because they can move data between tiers, but they do not offer more than one RAID level within each tier.

With the Compellent Fluid Data architecture, there are no predefined restrictions on how much data can reside on a certain type of storage. Other storage architectures require pre-determined tiering allocations—for example, an administrator might have to set a rule that only 10 percent of a volume can reside on solid-state drives while the rest must reside on Fibre Channel or SATA drives. Static storage solutions lack the flexibility organizations need to optimize their environment as needs change, and as a result, administrators could easily spend more time managing performance and capacity problems.

Most storage architectures do not include a real-time intelligence approach. Therefore, many tiered storage solutions have been produced by tacking software onto an existing storage platform. They employ a new application—or “agent”—to determine whether data should be moved. The agent, which requires a distinct server, is run periodically or on demand. It assesses the state of data, identifying, for example, “old” data and then moving that data to a lower tier of storage. Because the agent collects data only periodically, the information it collects is often outdated by the time the migration process is ready to run. In addition, data is moved in large blocks, although much of the data constituting those blocks may not be eligible for migration.

Use characteristics about each block are captured in flight to ensure precise data movement using the most current information.

4. True automation simplifies tiering administration

Compellent Automated Tiered Storage provides fully automated, hands-free sub-LUN tiering. Administrators use policy-based profiles to drive placement and movement of data. Users can choose predefined storage profiles, assembled by Compellent through years of automated tiering experience, or define their own. Profiles can be applied to a single LUN, a group of LUNs or multiple LUNs within a group, fine-tuned by application if desired. The profiles can specify not only the storage tiers to be utilized, but also the various disk types, rotational speeds and/or RAID levels within each tier. Profile settings can be changed at any time without downtime or disruption to the production environment.




Figure 3: A built-in migration engine provides hands-free management through policy-based tiering profiles for each volume or group of volumes.

Migration occurs automatically at a set time defined by the user, or on demand, while the system is still online. The migration process runs in the background and does not affect data availability or application performance. There is no need to bring down an application, pause I/O or wait for a minimum I/O requirement. If a read request comes in to a page that is being moved, the request is satisfied from the original placement of the page. The page is then moved after the read is complete. If a write request comes in, it will not interfere with the migration process, as new information is always written to tier 1, RAID 10—and therefore is not eligible for migration. Overwriting a block of protected information is also satisfied on tier 1, RAID 10, so moving a block of data receiving writes simply will not occur. There is never a situation when application I/O is denied—application requests always receive priority.

Administrators can easily expand LUNs on the fly or add capacity online to accommodate changes without system disruption or downtime. When drives are added to any tier, the system automatically restripes all data across appropriate RAID levels during the next migration sequence within a 24-hour window.

Other solutions do not provide the same flexibility. Since volumes are not easily expanded, administrators might need to apply compression as LUNs fill up. Administrators would then need to apply decompression if data becomes active again, which may delay data access. Even if compression is used, LUNs will ultimately run out of room and require data migration to a new LUN. And since read and write requests hit the same blocks, new writes will incur the I/O and latency performance penalties inherent to lower tier, high-capacity disks with RAID 5 or 6 protection.

Profiles specify not only the tiers to be used for each volume, but also the disk types, rotational speeds and RAID levels within each tier.

5. Snapshot integration delivers write and read performance

Compellent Data Instant Replay™ technology plays a key role in delivering outstanding performance in an Automated Tiered Storage environment. Data Instant Replay protects data without wasting capacity. Administrators can create space-efficient snapshots (Replays) of changes in data for continuous protection from loss or corruption. Replays use minimal storage space

because only data written to the volume since the previous Replay was taken is captured on the array.

In the Compellent architecture, new data is written by default to tier 1, RAID 10 storage to provide the best write performance. Replays move to a lower storage tier with RAID 5 or 6 protection during the next migration cycle within a 24-hour window. And over time, according to the tiering profile, infrequently accessed blocks of data move to a lower storage tier and RAID level, or to a different RAID level within the same tier. Moving this read-only data from RAID 10 to RAID 5 within the same tier enables administrators to maintain the same read performance. Whether data is moved to a new tier or RAID level, the migration frees up space on the higher tier or RAID level.

When new data needs to be written to an existing block that has since been converted to read-only and migrated to a lower tier, those writes are redirected to the tier 1, RAID 10 storage. A new writable block is automatically allocated to provide the highest transaction performance. Virtual pointers utilize the use characteristics of those blocks to maintain data continuity.




Figure 4: All data is written to tier 1, yet snapshots move to the lower tier available within 24 hours for the highest write and read performance possible. Virtual pointers retain continuity between all associated blocks.

For many organizations, data exhibits only a 3 to 10 percent rate of change. Consequently, most data is read-only data that can be migrated to lower tiers. Organizations often can reclaim 20 to 40 percent of tier 1 capacity simply by migrating read-only data.

Tiered storage solutions within traditional storage architectures cannot deliver the same level of write performance. In those

solutions, data is written to a particular block and kept in that block. If the block is migrated to tier 3 and a new write comes in for that volume, the write will occur on tier 3. With Compellent there is no write penalty for migrating data to lower tiers.

Data Instant Replay also helps sustain outstanding performance even as snapshots are taken. With static architectures, preparing storage volumes requires administrators to carve out storage space for the volume plus a distinct cache for volume-specific snapshot data. When new data is written, the storage controller determines whether that volume is protected by snapshots. If the answer is yes, the system reads the existing active production data, writes changes to the snapshot cache and then writes over the volume block with the new data. That sequence requires numerous I/O processes. In addition, other systems store snapshots inside the original volume alongside the active production data. Consequently, there is a finite amount of space available for the snapshots—only a certain number of snapshots can be taken.

The Compellent system does not require administrators to reserve space for snapshot data or write data. Furthermore, Compellent snapshots do not reside within the volume with production data—they sit outside the production volume in pages within the general virtual storage pool. With the Compellent system, administrators can take frequent snapshots while conserving capacity and maintaining excellent performance.

Users do not need to reserve space for snapshots or new writes, and there is no write penalty for migrating data to lower tiers.

6. Integration with other enterprise features maximizes tiering benefits

Because Compellent was engineered with automated sub-LUN tiering in mind, other critical storage features are tightly integrated with Data Progression. Like Data Instant Replay, these technologies help to enable—and maximize the benefits of—Automated Tiered Storage, delivering the efficiency, performance and ease of management that organizations require.

Thin Provisioning

Compellent Thin Provisioning, called Dynamic Capacity™, helps optimize disk utilization within the virtualized storage environment. Thin Provisioning helps create the flexible, efficient

virtualized environment with the pooled storage resources required for effective tiering.

With Thin Provisioning, since even free blocks are preformatted during the provisioning process yet no zeroes are written to set aside volume space, administrators can expand or shrink volumes on demand without being bound to RAID set capacity or performance limitations. Other solutions require administrators to predefine RAID sets ahead of time and later format additional space to accommodate LUN expansion. To make changes with those solutions, administrators have to free up additional storage capacity and migrate data.

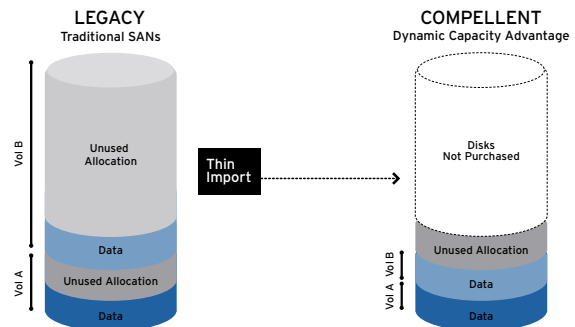


Figure 5: *Compellent Thin Provisioning helps to create the flexible pool of storage required for efficient and effective tiering.*

Some other vendors offer a version of Thin Provisioning, but the capability applies only to storage contained within a newly created virtualized storage pool. These vendors often recommend against using Thin Provisioning to support enterprise database or e-mail environments because activating Thin Provisioning could degrade write performance, and produce latency, as the controller prepares the storage environment. With Compellent, Thin Provisioning is a core component of the virtualized storage platform and is always ready to go—using it has no performance impact.

Fast Track

Compellent Fast Track augments Automated Tiered Storage by delivering optimal placement of data on each individual disk. Fast Track uses the intelligence continuously collected by the Compellent system to identify the most active, frequently accessed blocks of data on each spindle. It then places those data blocks together on the outer tracks of each drive. Keeping the head focused on that one area of the disk, where the active data resides, delivers better performance than if the head were forced to move all over the disk.

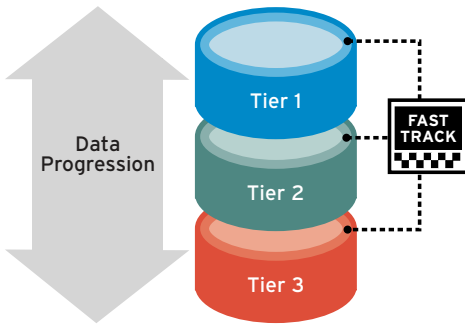


Figure 6: *Compellent Fast Track groups the most active data blocks and places them together on the outer tracks of each drive to improve performance.*

Fast Track works on every spindle-based drive type within the system. Data movement happens automatically, on a daily basis, without any manual intervention. With Fast Track, organizations optimize performance of each platter based on actual usage patterns, reducing seek times for the most active blocks of data. Fast Track also helps reduce long-term storage costs by maximizing performance on cost-effective drives.

Enterprise Manager

Compellent Enterprise Manager helps to simplify administration of Automated Tiered Storage. The storage management interface provides a comprehensive view of the entire Compellent storage environment, including real-time tiering statistics. Administrators can use the interface to monitor system performance and capacity, such as the percentage of data in each tier and RAID level by LUN. Or users can opt to simply receive regular tiering reports via e-mail.

With Enterprise Manager, administrators can interact with the system to whatever degree is desired. Following initial setup, no staff interaction is required, as all tiering functionality is fully automated to provide hands-free management. However, administrators who prefer more direct involvement can take advantage of full configuration access and control—from augmenting the storage profiles and policies that drive the tiering algorithm for each volume to changing Replay schedules, all of which can be performed without disruption.

Solutions that add tiering functionality to an existing storage platform do not offer integrated management. Administrators are required to use a separate tool to manage the tiering process. Integration with other virtualization features is limited.

Administrators can expand or shrink volumes on demand without being bound to RAID set capacity or performance limitations.

7. Platform independence facilitates scalability

Compellent Automated Tiered Storage functionality is not tied to any particular storage system model or platform. Organizations can scale up and out without limits, downtime or the need to alter the licensing and service agreements. Incorporating Automated Tiered Storage functionality into an existing Compellent system is as easy as activating that particular built-in software module.

With other solutions, the tiering license might be tied to a particular storage system model. To take advantage of tiering with those solutions, organizations might need to update to a new platform. If a forklift upgrade is required, it could cause downtime for data migration. Administrators would have to migrate all existing volumes from the current RAID group configuration to a pooled configuration for automated tiering to work. To complete the process, administrators would need twice the capacity as the original volume during migration to the new platform.

Making that upgrade would also change the software license and service agreements. These agreements are typically tied to the storage controller. As soon as an organization upgrades a controller, it needs to upgrade the license and service agreements. The same is true if the organization moves from one platform to another. Businesses can avoid these issues with Compellent Fluid Data storage.

Sub-LUN tiering is not tied to a particular model or platform, so taking advantage of it does not require a costly upgrade.

CONCLUSION

As organizations continue to implement sub-LUN tiering solutions to cost-effectively address explosive data growth, IT management must understand the architectural limitations inherent to most solutions. Without a truly virtualized storage platform engineered specifically with Automated Tiered Storage in mind, even more sophisticated, policy-based tiering solutions fall short. Such solutions lack the integration, granularity, intelligence and automation needed to ensure data remains in tune with application needs.

With its Fluid Data architecture, Compellent brings proven efficiency, automation and scalability to sub-LUN tiering. Automated Tiered Storage has been an integral feature of Compellent storage since 2005, and Compellent has been enhancing the functionality ever since based on pragmatic knowledge of how to utilize it with particular applications and storage configurations for optimum performance and capacity utilization.

Continuously aware of each granular block of data spanning the entire storage system, Compellent Automated Tiered Storage enables organizations to strike a constant balance between performance and cost. Administrators do not need to replace or add hardware, reformat existing volumes or install server-side agents to take advantage of dynamic tiering. The functionality is built right into the storage architecture—and not only integrates with, but leverages a full range of enterprise features to drive efficiency, performance and manageability.

With Compellent, data is not only managed based on actual use between storage tiers, but also RAID levels within each tier and together on the fastest tracks of every drive. Sophisticated tiering profiles provide hands-free management, while allowing full administrative control as needed depending on business needs. And organizations can scale the solution on demand to better meet those business needs.