

# FOUR WAYS TO LOWER THE COST OF REPLICATION

How an Ultra-Efficient, Virtualized Storage Platform Brings  
Disaster Recovery within Reach for Any Organization

## EXECUTIVE OVERVIEW

Within the last decade, replication technologies have come to play a major role in the data protection paradigm. At first, replication processes required large, duplicate systems for off-site recovery, with a price tag that shut out all but the largest of enterprise customers. But more recent innovations have helped to dramatically cut cost and complexity, making disaster recovery (DR) a viable solution for organizations of all sizes. Today, even small IT groups with limited resources can create a cost-effective plan for both near-line and off-site replication over heterogeneous storage platforms.

Whether you are currently using tape backup and want to create a faster, more reliable data protection framework, or need to increase efficiencies within your current DR infrastructure, Compellent offers a range of technologies that bring the cost of replication within reach. This white paper outlines four key ways to drive down replication costs while delivering enterprise-class performance and manageability with Compellent Fluid Data storage:

- 1. Storage Virtualization** – abstracts and aggregates physical drives into a shared pool of storage, ensuring data availability and rapid recovery for non-stop business environments.
- 2. Thin Replication** – utilizes space-efficient, continuous snapshots to protect data, yielding substantial savings in storage hardware, bandwidth and management.
- 3. Portable Volume** – dramatically reduces the timeframe for initial replication synchronization by using pre-configured external hard drives instead of high-speed data links or duplicate arrays.
- 4. Scalable Serial Attached SCSI (SAS)** – allows for the combination of high-performance and high-capacity drives in a single enclosure, providing remote site configuration flexibility and reducing costs while enabling administrators to leverage enterprise functionality like automated tiering.

### 1. Increase Availability and Recoverability with Virtualization

Storage virtualization is the backbone of any cost-effective plan for data protection and business continuity. With Compellent Fluid Data storage, all the physical disk drives in the system are abstracted and aggregated into a single virtual storage

pool presented to servers. A single command can be used to create and provision a logical unit number (LUN) from any part of the pool and map it to a physical host, simplifying data management and optimizing storage resources. The remote replication of virtual volumes means that mission-critical data is fully protected and quickly recoverable, even in the event of a site-wide disaster.

Because the Compellent SAN virtualizes storage at the disk level, read/write operations are spread across all the drives in the storage pool and multiple requests are processed in parallel for optimum performance. Virtual volumes can span multiple heterogeneous arrays, and administrators can allocate storage based on usage, reducing both the likelihood of underutilized drives and the once-common issue of inadequate space. System administrators can respond immediately to changing storage requirements and ease management by adding capacity to virtual volumes when pre-set thresholds are reached.

A virtualized storage environment also enables administrators to quickly and easily assign and manage virtual volumes from a centralized user interface. As needs change, more space can be provisioned on the fly in a matter of a few mouse clicks without taking the system offline. Given today's budget constraints, the ability to employ less IT staff to manage larger amounts of data brings welcome increases in return on investment (ROI) and a parallel reduction in total cost of ownership (TCO).

In addition, storage virtualization ensures data availability and recoverability by safeguarding against all possible points of failure along the SAN. Compellent's clustered controllers operate in unison, but connect to drives and enclosures independently, providing higher availability than traditional implementations. Dual paths from servers to disk drives, along with automatic controller failover, provide high resilience. And non-disruptive software upgrades, performed without requiring the shut-down of applications, ensure continuous data access.

Virtual Ports deliver yet another level of hardware virtualization by enabling multiple virtual nodes to share a single physical port. This not only enhances the performance of existing I/O hardware – increasing port capacity, disk bandwidth and I/O connectivity – but enhances failover capabilities. Virtual Ports also eliminate the need for dedicated reserve hardware ports,

increasing flexibility and mobility. With support for both FC and iSCSI, this increased flexibility reduces replication expenditures for both entry-level and enterprise systems, delivering end-to-end virtualization with the greatest possible efficiency.




Figure 1: Storage virtualization optimizes storage utilization by abstracting and aggregating physical disk resources into a single virtual pool of storage presented to servers simply as capacity.

## 2. Right-Size Storage Assets with Thin Replication

The challenge of implementing a full-scale replication solution can seem daunting. IT managers must evaluate their risk tolerance and performance requirements, and map them to available hardware budgets and staffing for long-term planning. Compellent's Thin Replication technology, called Remote Instant Replay, provides enhanced data protection while minimizing cost, time and effort.

Unlike conventional solutions, Thin Replication uses continuous, space-efficient Replays (snapshots) to protect enterprise data. Once an initial Replay of a virtual volume is taken, the Compellent SAN only needs to capture incremental changes in data. This eliminates the need to create full-volume clones at the replication target site, minimizing initial hardware investments. The SAN can maintain an unlimited number of Replays in increments defined by the user, and the Replays can easily be mapped to any server for near-instant recovery.

During ongoing replication, the Compellent SAN intelligently transfers only the changed blocks of data, rather than entire volumes, optimizing bandwidth utilization and the use of remote site resources. System administrators can closely estimate

bandwidth requirements and configure replication processes based on actual Replay size using an intuitive interface. Transfer rates can be customized based on link speed, time of day and replication priority. In addition, Replays can be used to test new applications and service packs without disruption, efficiently support boot-from-SAN operations and virtually eliminate backup windows on production systems.

Because Thin Replication does not require primary and remote locations to have identical configurations, users have the flexibility to choose lower-cost drives for remote sites. To further maximize efficiencies, Fluid Data storage features a built-in Fibre Channel (FC)-to-iSCSI converter, enabling users to replicate natively over IP. While most SANs require expensive FC-to-iSCSI conversion, Compellent's native iSCSI interface allows users to leverage their existing infrastructure for additional savings in hardware and staff training. Whether using traditional FC or cost-effective IP-based transfer links, administrators can replicate data from high-performance SSD, FC or SAS drives at the primary site to high-capacity SAS or SATA drives at a remote site.

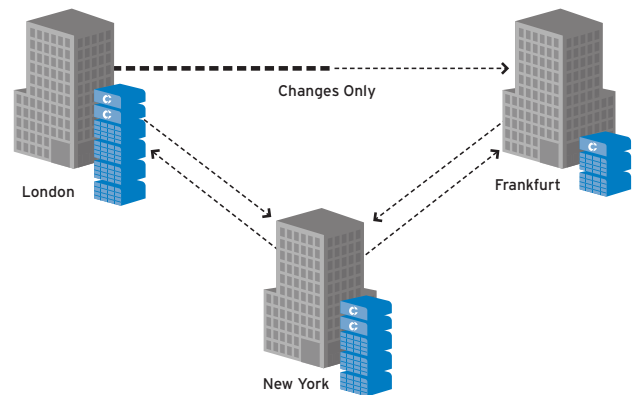


Figure 2: Thin Replication captures unlimited, space-efficient snapshots of incremental changes in enterprise data for robust site-to-site data protection without the need for full-volume clones.

## 3. Jumpstart Replication with Portable Volume

For organizations that do not have a dedicated fiber link or other high-speed connection, the initial replication synchronization of data between primary and remote sites can present a significant time and bandwidth challenge. Depending on the size of the volumes being replicated and the bandwidth type, the initial process of creating identical data sets at a remote site can take weeks or months – an unacceptable timeframe for many organizations.

Compellent Portable Volume dramatically shrinks the synchronization timeframe to days or hours, regardless of the type of data link between sites. It jumpstarts replication by enabling users to prime portable, high-capacity hard drives with the volume(s) base and Replay information necessary for the initial synch-up from the source array. The drives are then physically transported to the remote site, where they are connected to a target array. The Compellent SAN auto-detects the Portable Volume and immediately begins the creation of identical replication volumes on the target array.

Businesses that have historically required rapid synchronization of critical data previously would have had to send a duplicate storage enclosure to their remote locations or invest thousands of dollars in a separate, dedicated high-speed data link. The Portable Volume kit, which consists of two 2 TB enterprise-class external hard drives, two USB cables and a ruggedized case, presents a far more cost-effective option for creating mirror-image data stores within hours or days. This elegantly simple solution enables administrators to implement remote replication without the bottlenecks caused by low-bandwidth links or large data stores. In addition, system administrators can replicate data more quickly with Portable Volume than with traditional replication running as a background task over a shared link.




Figure 3: Portable Volume jumpstarts replication by using external enterprise-class hard drives with USB connectivity for initial synchronization in lieu of duplicate arrays or high-speed data links.

For small to midsize data centers with limited bandwidth, and for organizations with stringent data availability requirements, Portable Volume leverages the common USB 2.0 interface to

### SCALABLE SAS: THE BEST OF BOTH WORLDS

Compellent Scalable SAS provides a cost-effective, all-in-one storage solution for replication target sites. System administrators can establish a remote site with just one controller and one SAS enclosure, then scale up over time to accommodate growth.

One controller supports as many as 32 enclosures, each with 12 or 24 drive bays, and users can scale from 6 to 384 SAS drives in a single system. In addition to being highly scalable, unlike other SAS solutions on the market, Compellent's SAS bundles allow users to mix and match performance-oriented drives and capacity-oriented drives in the same enclosure. The result is an optimal balance of performance and value - a cost-effective, stand-alone SAN with the capability to support enterprise-class functionality such as Automated Tiered Storage right "in the box."

deliver rapid replication startups without sacrificing existing FC ports needed for front- or backend connectivity. The data on Portable Volume drives is readable and visible only to the Compellent SAN and can be encrypted using standard AES 128-bit technology for maximum data-at-rest protection during transit. After synchronization is completed at the remote site, the hard drives can be used to initiate replications at other target sites.

### 4. Use SAS Strategically at Remote Sites

Typically, performance-oriented FC drives and high-capacity SATA drives are housed in separate enclosures. This often results in enclosures that are only half-populated at remote sites, consuming valuable rack space and driving up the cost per usable GB for replication targets.

With Compellent Scalable SAS, system administrators can mix high-performance SAS drives and high-capacity SAS drives in the same enclosure. The result is an optimum combination of speed and value with a minimal footprint. High-performance SAS drives are comparable in price to FC drives, but SAS enclosures and interconnects are less costly than their FC counterparts. In addition, since Compellent SAS drives are fully interoperable with FC and SATA drives, all three drive types can be mixed in the same system for increased remote replication flexibility.

Configuring remote sites with a blend of SAS drives also enables system administrators to take advantage of enterprise SAN

functionality like Automated Tiered Storage without FC or SATA drives. Fluid Data storage classifies the two drive types into tiers based on capacities and speeds, and uses block-level intelligence to automatically migrate data bi-directionally between the two tiers based on changing access patterns. This brings a whole new level of efficiency to the utilization of mixed drives within a single array. By decreasing reliance on costly drives for replication targets and reducing the number of drives required, companies can slash overall storage costs by as much as 80 percent.

For storage administrators who are already performing remote replication, shrinking the data center footprint with fewer enclosures can bring cost/benefit ratios in line with tightening IT budgets. For those looking to implement offsite replication for the first time, Compellent's support for SAS makes remote replication affordable while delivering an unprecedented range of enterprise-class features. Either way, an investment in SAS delivers significant savings in TCO and ensures that current investments are protected as the industry moves away from relying on high-cost FC drives.




Figure 4: With Scalable SAS, high-capacity and high-performance drives can be combined in a single enclosure for cost-effective remote site configurations capable of enterprise functionality.

## CONCLUSION

Utilizing Compellent's innovative replication technologies, it becomes fully feasible for data centers of all sizes to implement a business continuity plan that is resilient, robust and sharply focused on providing high IT impact on a budget. It all begins with a dynamic, block-level architecture, which provides the granularity required to create a virtualized pool of storage resources. Storage virtualization ensures optimal resource utilization, enabling system administrators to replicate data efficiently and recover from system failures or natural disasters almost instantly. Initial replication synchronization can be performed in a matter of days instead of weeks or months – all without expensive, high-speed connections or the need to ship duplicate arrays from site to site. And Scalable SAS provides a cost-effective yet high-performance replication target solution with a minimal footprint.

### Compellent Technologies

7625 Smetana Lane | Eden Prairie, MN 55344

877 715 3300 tel | 952 294 3333 fax | [www.compellent.com](http://www.compellent.com)