DATA INSTANT REPLAY

FEATURE BRIEF | AUGUST 2006

compellent

compellent

DATA INSTANT REPLAY OVERVIEW

The Data Instant Replay™ feature of the Compellent® Storage Center™ storage area network (SAN) creates space-efficient snapshots – copies of volumes, which Compellent calls Replays – to provide extremely fast recovery from any type of threat to data. Creating or recovering from a Replay is nearly instantaneous, typically accomplished in 10 seconds or less. Besides protecting data, Replays can help administrators test new applications and service packs, efficiently support server boot-from-SAN operations, and virtually eliminate backup windows on production systems.

With Data Instant Replay, Compellent's Storage Center SAN provides a safe, reliable, and easily managed solution for preserving the integrity of data, enabling rapid recovery from any data loss. Not only can the system maintain an unlimited number of Replays that can be mapped to any server, but Replays also require significantly less storage space than traditional snapshots. As a result, Data Instant Replay delivers the flexibility that IT administrators need to perform their most critical functions efficiently and effectively.

This feature brief describes how Data Instant Replay works, some of its common uses, and the benefits you can derive from using the feature.

Unlimited Number of Snapshots Improves Recoverability

Compellent's Dynamic Block Architecture™ is an innovative foundation that provides the basis for numerous enhancements to traditional data management. Dynamic Block Architecture divides the storage used by data into 2 MB pages and maintains information about the blocks called metadata. This metadata includes information on when the blocks were created, what disk drives hold the blocks, frequency of access, and whether the blocks represent actual data or Replays. With this intelligence about the blocks of data, many of the restrictions of traditional SAN snapshots are eliminated.

For example, Data Instant Replay supports an unlimited number of Replays, providing data protection at many points in time so recovery can be targeted as close as possible to when the failure occurred. The more frequent the Replays, the more your data is protected against loss, whether due to equipment failure, virus attacks, or human error.

Many storage experts consider 8 snapshots per volume and 10 to 20 total snapshots for a storage system to be the limits of most SANs. In fact, many storage systems were never designed to support an unlimited number of snapshots. Most traditional systems use a technique referred to as copy-on-write, which creates a unique copy of the volume with every new write. This technique consumes valuable disk space which severely limits the number of possible snapshots. Using Data Instant Replay, however, IT staff can take snapshots every 10 or 15 minutes, selectable on a volume-by-volume basis, and save them for any length of time. This provides significantly improved recoverability compared to what is possible with most legacy SANs.

Efficient Use of Storage Enables Frequent Snapshots

The unlimited Replays provided by Data Instant Replay consume significantly less storage space than snapshots from other SAN systems. This is because most traditional SAN systems require a full copy (a clone) of a snapshot in order to use it—so a 1 TB volume would typically require a copy of all of that data before creating a new volume. Not only does this consume storage space, but it is time consuming and may impact the performance of the original volume.

Compellent Replays do not require cloning. Storage Center simply maps the Replay as a logical unit number (LUN) to be used by the server. Negligible storage is consumed for mapping, and Replays can be mapped any number of times as read/write volumes. In fact, the size of the Replay grows only to accommodate the newly written data. This efficient use of storage, compared to other SANs, means many more Replays can reside in a given amount of physical storage.

Intuitive, Point-and-Click Interface Speeds Replay Management

Compellent's Replay management is designed for efficiency and ease of use, so administrators can adopt a "set it and forget it" philosophy. Replay operations are managed with a simple administrative tool for defining schedules that give the best possible protection within specific storage space requirements. This tool, which does not require scripting, is similar in appearance to a Microsoft® Outlook® calendar, helping make it intuitive to use. Scheduling is easy and flexible. Using point-and-click templates, administrators can rapidly establish many different types of schedules, such as the following:

- Create a Replay every 10 minutes and save each one for 4 hours
- Create a Replay once a day at noon and save it for a week
- Create a Replay on the first day of the month and save it for six months

Furthermore, the flexibility does not end with schedules. Once a scheduling rule is created, it can be applied to multiple volumes. And rules can be added at any time, even after the volume is in use.

Restoring Data is Quick and Simple

Restoring an accidentally deleted or corrupted file is a simple task using Data Instant Replay. Here is an example of how it works. Figure 1 shows multiple Replays of a volume taken at successive points in time, starting at T_0. There are two Replays, taken at T_0 and T_1, as well as the online volume, shown at T_2. The older Replays are all read-only, while the current version of the volume is read/write.

Now, a user calls the IT administrator to say a file has been deleted by mistake, sometime after T_1. To retrieve the file, the administrator first maps the T_1 Replay to either the same server that accesses the original volume or to a different server. From that server, the administrator simply locates the original volume and file and copies the requested file to the current volume, quickly resolving the issue for the user.

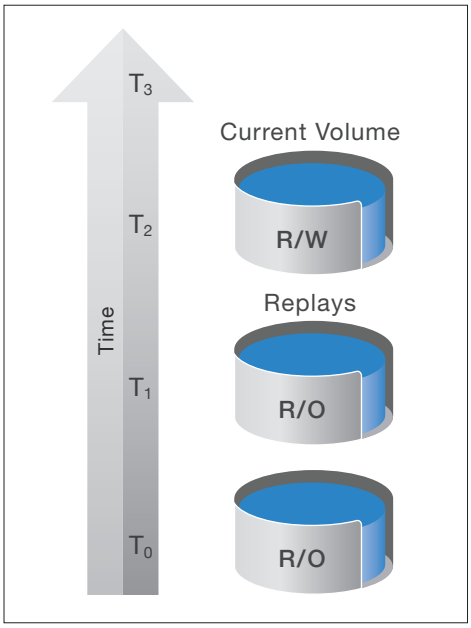

Figure 1: *Example of multiple Replays taken at successive points in time.*

Virtualized Replay Mapping Minimizes Space Requirements

Storage Center keeps track of virtualized data in 2 MB pages. This size makes efficient use of storage without incurring a large amount of housekeeping overhead. A larger page size can waste storage space. On the other hand, a much smaller page size incurs software overhead to keep track of all items being virtualized, which can slow performance.

Figure 2 shows how the Compellent Storage Center creates Replays without consuming additional space except to accommodate new data that is written to the volume. The rows represent Replays taken at different time intervals from the same volume. They also represent mount points that can be mapped to a server as a logical unit number (LUN). Each column in the figure represents a page. The number of pages required to represent a Replay is based on the rate of change and the size of the volume itself.

In Figure 2, a Replay represented by row T_0 is the oldest point in time. The columns in the figure show the pages in this Replay. At T_0, data is written to Page 0, Page 2, and Page 3 shown with an X in each of these pages. Due to the design of Dynamic Block Architecture, disk space, in the form of pages, is consumed only as data is written to the disk drives.

Now, assume that at the next time interval after the initial Replay was created, a second Replay is taken at T_1. Some of the data in this Replay will update existing pages, Page 0 and Page 1. At time T_1, since the data in Page 2 and Page 3 did not change, no additional writes are required, nor are any additional pages needed.

The far right column in Figure 2 shows the number of 2 MB pages required by the Replays. The only pages used are those that contain new data written to the volume. In typical computing environments, most of the data on a volume does not change, so this approach allows Replays to consume very little space relative to the size of the original volume.

At time T_2, a third Replay is taken. At this time interval, Pages 0, 1, 2, and 3 did not change. A new page is required to represent an increase in the size of the original volume, shown by a write at Page 4.

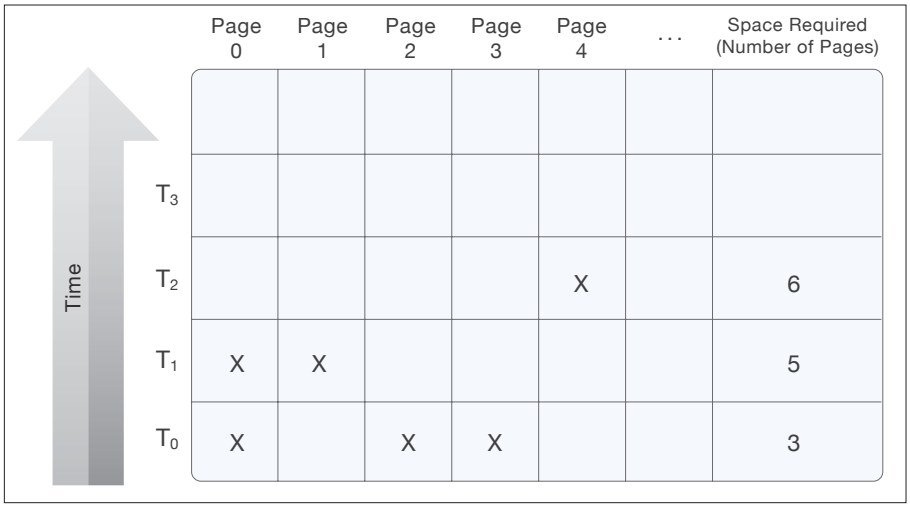

Figure 2: *Replays consume disk space only to record writes to the volume.*

At any given time interval, a complete Replay can be recovered through the combination of the most current writes from all the pages in the Replay. To illustrate this concept, suppose that it is necessary to restore the volume to its state at time T_3. Storage Center would simply provide the pages shown in the table below:

Page	Time (T)
0	T_1
1	T_1
2	T_0
3	T_0
4	T_2

Restoring a volume based on a collection of Replays is extremely fast and efficient with Dynamic Block Architecture. The volume is restored in seconds by assembling the pointers to the pages. Most traditional systems restore a snapshot by creating a complete copy of the volume. This lengthy process often takes hours. In most environments, data recovery is a time-critical operation and users do not have time to wait while an entire clone of the volumes is re-created.

Automatic Coalescence of Replays Conserves Storage

Typically, Replays have a predetermined lifetime. For example, if Data Instant Replay generates Replays every hour, the administrator may decide to retain them for only one day on the assumption that, if a problem develops, users will report it during that time. Once a Replay reaches the predetermined lifetime, the data from the Replay is coalesced with existing Replays and the space consumed by the Replay is returned to the available pool of storage.

Storage Center automatically uses coalescence to make the most economical use of storage capacity as Replays expire. Figure 3 shows how this process operates. Pages that have been replaced with newer data are deleted, while pages that contain the most current version of data are designated as the original point in time.

The process is illustrated in Figure 3. The Replay at T_0 has expired and is no longer needed. The T_0 Replay is comprised of Pages 0, 2, and 3. Since Page 0 has been updated with newer data at T_1, the page representing T_0 can be deleted. This frees disk space which is returned to the free storage pool. The data in Pages 2 and 3 has not changed. The pointers to these pages are changed to time T_1. Once the pointers have changed, all of the information about the T_0 Replay can be deleted.

All of the Replay management occurs without any performance impact on the operating volume since data is not moved on the disk drives, only pointers to the data. Through Dynamic Block Architecture, Replay management is extremely fast and efficient, saving valuable storage resources.

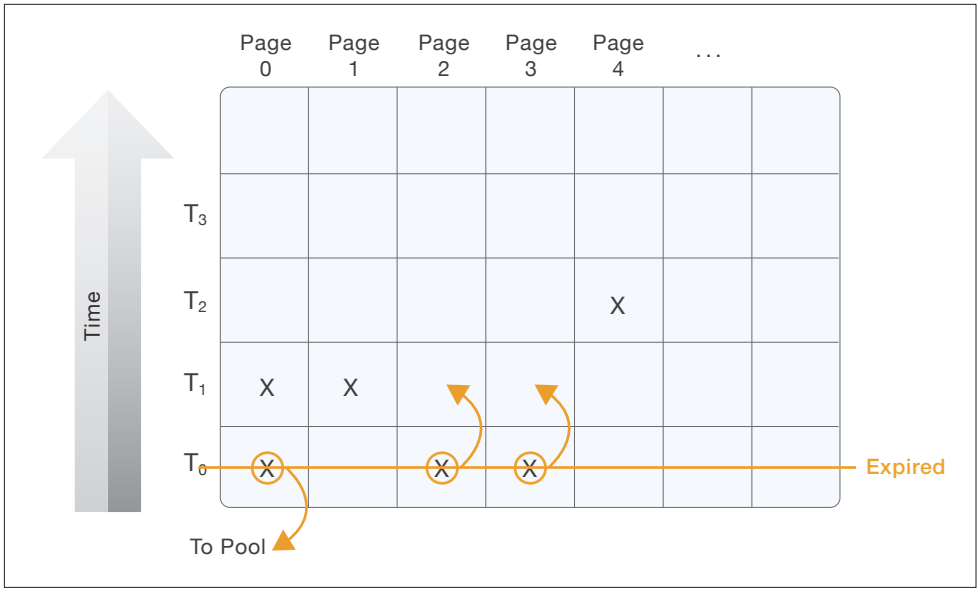

Figure 3: *Pages are coalesced as Replays expire.*

New Volumes Are Easily Deployed

Storage Center provides the flexibility to deploy Replays back into online systems. Figure 4 shows one of the Replays from Figure 3 being deployed as a new volume at time T_1. As time progresses, a Replay is created at time T_2 for both the original and the new volume.

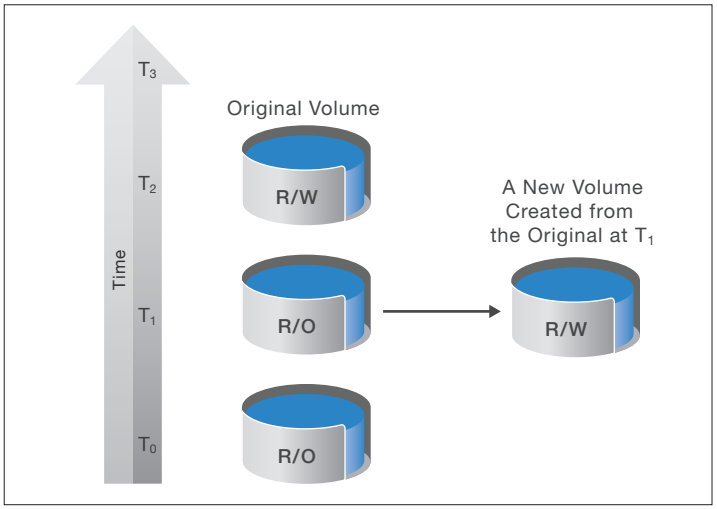

Figure 4: *A Replay taken at time T1 is deployed as a new volume.*

Mapping Replays To Multiple Servers Increases Efficiency

Multiple new volumes can be created from a single Replay. In fact, multiple Replays can be created from Replays that are mapped to servers—all while consuming minimal additional storage space.

Figure 5 shows a Replay taken at time T_1 mounted to two servers. This can prove useful in a number of common circumstances. For example, it could be used to create the boot volume within a Boot from SAN server environment, where each server boots from the same image. Or it might be done to allow two programmers to independently test an updated application on the same live data to ensure it works correctly before deploying it into production. In this scenario, one Replay is used to create two versions of the original volume. In Compellent terminology, this is called a View Volume. These two new View Volumes are mapped to different servers, one to server 1 and the other to server 2. Now, each server has a fully functioning version of the original volume with current data for testing.

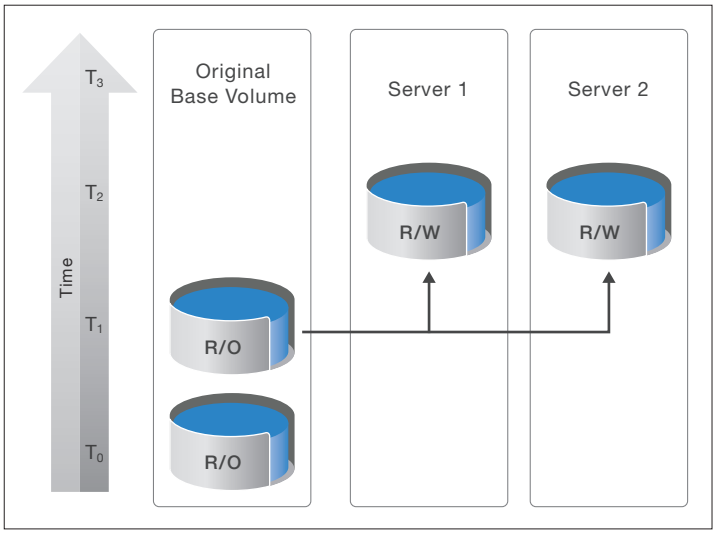

Figure 5: *A Replay can be mapped to multiple servers.*

The amount of overhead space required for each mapping is insignificant. In a Boot from SAN server environment, administrators can install the operating system on one server, configure it the way they want, and then take a snapshot of that volume to be used by other servers. The multiple Replays of that volume consume minimal storage space. Simply mapping those Replays to other servers enables the servers to boot off the Replay. This is a very efficient means of maintaining centralized control over server boot images while also making efficient use of storage space as shown in Figure 6.

In this scenario, each server accesses the files from the original volume, called the Gold Image, and changes made to each server's volumes are recorded independently. But the original volume remains unchanged as read-only until it expires. There is no limitation on how many servers can access that base volume. Very little space is required to map the volume to additional servers—unlike other systems that require a complete copy of the volume, which wastes space and takes time to create. And the mounting process is quick and painless; with Storage Center, a multi-terabyte Replay can be mounted in seconds without any performance degradation to the original volume.

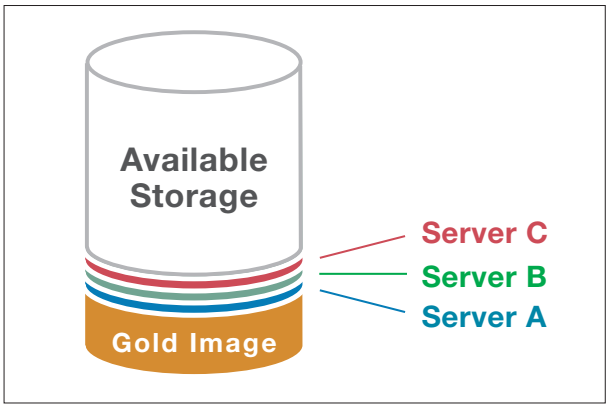

Figure 6: *Booting servers from a Gold Image on the SAN reduces storage space requirements.*

Being able to create multiple volumes from a Replay can also shorten the time to find a problem plaguing a server. One of the first steps in recovering from a virus attack, for example, is to determine the origin of the virus and identify when it infected the server. If several people could work on that at once, it might reduce downtime. Using Storage Center to mount the Replays on any number of servers enables as many staff as necessary to work on the volume in question.

Another scenario in which multiple Replays are invaluable is the creation of test environments, where data may be destroyed or corrupted by the testing system. With Data Instant Replay, it is quick and easy to revert back to the original volume. The original data stays unchanged as the test servers write information to the new mapped volumes.

Replays Can Reduce Backup Windows

The use of Replays virtually eliminates the need for backup windows. Administrators can simply make a Replay of the production system and then backup the Replay to tape while the production system continues to operate uninterrupted. The time and effort associated with backup windows becomes a distant memory.

Some customers using Data Instant Replay have created powerful systems in which Replays are automatically taken on a nightly basis. By using Storage Center's scripting utilities, these Replays are automatically archived to tape. This implementation, using a Replay of the actual data, dramatically reduces the backup window. In addition, the entire process is automated to run during hours of reduced access, eliminating impact to end users.

Conclusion

Storage Center Data Instant Replay utilizes Dynamic Block Architecture to provide an improved method for using snapshots. Through this innovative approach, Compellent provides a faster and more efficient method for creating, managing, and recovering snapshots. These innovations provide significant cost savings and improved data protection to businesses of all sizes.

For more information about Data Instant Replay and the Compellent's Storage Center, please contact 952-294-3300 or visit www.compellent.com.

COMPELLENT

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

www.compellent.com