



compellent

Compellent Source Book

Remote Instant Replay

Thin Replication

Compellent Corporate Office

Compellent Technologies
12982 Valley View Road
Eden Prairie, Minnesota 55344

www.compellent.com



Contents

Contents	2
Document Revision.....	Error! Bookmark not defined.
Product Description	3
Overview.....	3
Detailed Description.....	4
Architecture	12
Platform / Operating System.....	13




Product Description

Overview

Remote Instant Replay replicates continuous snapshots, called “Replays: between local and remote sites, offering easy to implement disaster recovery at a very affordable price. With Remote Instant Replay, companies can establish business continuity with the granularity required to recover to any point in time. Remote Instant Replay supports traditional Fibre Channel replication or cost-effective IP-based replication, bringing disaster recovery within reach for every business.

Remote Instant Replay fully integrates with all Compellent products including Data Progression, Dynamic Capacity and Data Instant Replay. These combinations enable large and small enterprises to quickly implement business continuance and disaster recovery strategies that were previously too expensive or complex.

With Remote Instant Replay, remote Replays are created using only the actual written data during initialization and only blocks that have changed after that, optimizing bandwidth and cost. Any number of Replays at any interval can be created between multiple sites, allowing multiple sites as recovery points. Sites can be configured as bi-directional, point-to-point, multi-point, or peer-to-peer, with all sites active and available.



<i>Feature</i>	<i>Benefit</i>
Thin Provisioning and Snapshot Replication	<ul style="list-style-type: none"> Intelligently optimize replication at the block level and transfer only the written or changed data Unlimited Replays (snapshots) enable granular recovery intervals for continuous data protection
Bandwidth Simulation	<ul style="list-style-type: none"> Eliminate traditional process of guessing or spending hours trying to determine bandwidth requirements Provide accurate assessment of bandwidth requirements based on actual data requirements, optimizing initial bandwidth purchase
Bandwidth Shaping	<ul style="list-style-type: none"> Allow utilization of a single bandwidth connection for both regular internet traffic and replication traffic by managing replication priority
IP-based Replication	<ul style="list-style-type: none"> Reduce cost by replicating over long distances using existing Ethernet networks
Multi-site Replication	<ul style="list-style-type: none"> Establish multiple locations as recovery points in bi-directional, point-to-point, multi-point or peer-to-peer configurations Replicate using synchronous (short distance) or asynchronous (long distance) transfer modes Local and remote sites can be different configurations
Centralized Replication Management	<ul style="list-style-type: none"> Reduce administration time by setting up replications in as few as 6 clicks Recover more quickly from any form of disaster Validate Disaster Recovery plan
Flexible Replay Management	<ul style="list-style-type: none"> Active Replay / One behind replication provide additional control over bandwidth utilization
Replay De-duplication	<ul style="list-style-type: none"> Reduce bandwidth requirements by eliminating transmission of duplicate data blocks

Detailed Description

Thin Provisioning and Snapshot Replication

Remote Instant Replay utilizes Storage Center's Dynamic Block Architecture to provide a more efficient replication solution called Thin Replication. Storage Center intelligently optimizes replication at the block-level and Thin Replication transfers only the written or changed data, instead of transferring allocated but unused space.

Any number of Replays can be created and easily replicated to one or more locations, with independent Replay schedules at each location. The unlimited granularity enables small recovery intervals, allowing more recovery points to instantly recover from any point in time. Replication does not require a pre-allocation of space and only sends written blocks of data.

Replays can be tested and validated daily or even hourly without downtime, which allows System Administrators to easily verify that volumes are being replicated successfully.

One of the greatest challenges to implement an initial remote replication installation is the initial determination of the actual bandwidth requirements. There are many variables that determine the amount of bandwidth required including; recovery time objective, recovery point objective, budgets, available bandwidth, and amount of data to be replicated. These variables require many hours / days of analysis with a spreadsheet to determine actual requirements. Even with careful analysis, the effort is slightly better than an educated guess. However, since the largest cost associated with replication is the recurring bandwidth costs, this is an important exercise.

The Bandwidth Simulation feature introduced in Storage Center 3.5.1 eliminates the traditional guesswork by using actual system parameters and data to graphically show the required bandwidth that meet currently defined bandwidth shaping parameters.

Creating Simulated Replications

The Storage Center GUI is used to create a simulated replication on a volume. When the replication options are set, Storage Center creates the simulated Copy-Mirror-Migrate (CMM) of the data and begins the replication process. Figure 1 displays the Replication to Simulation options.

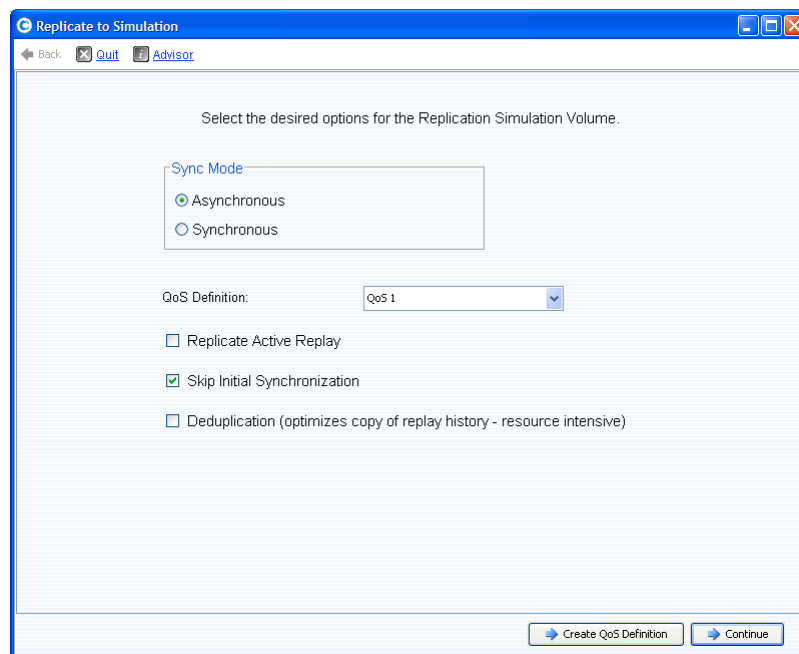


Figure 1 Replicate to Simulation

Viewing Simulated Replications

Enterprise Manager 1.3 (or later) is used to monitor the simulated replications and QoS Node link utilization over time. Enterprise Manager provides Bandwidth Usage and Replication I/O Per Second reports that provide a strong indication of the link utilization and bandwidth limits over time as shown in Figure 2.

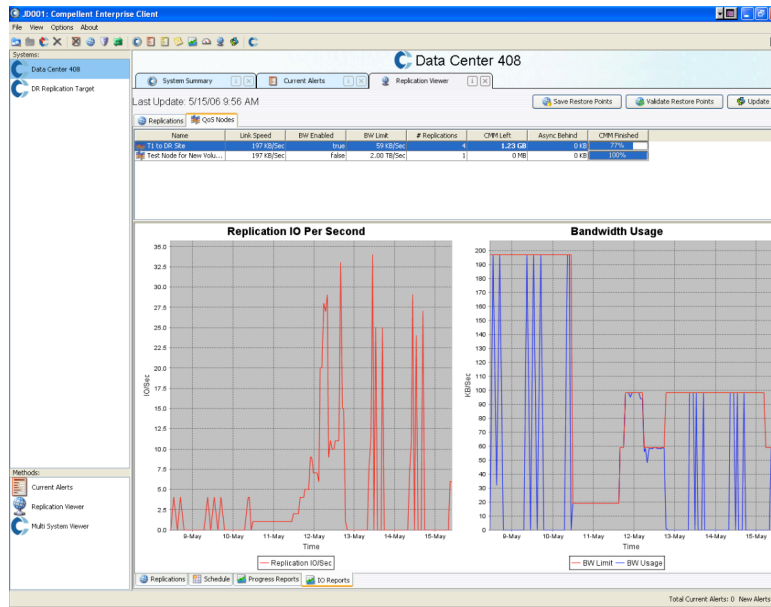


Figure 2 Enterprise Manager Bandwidth Simulation Reports

Bandwidth Shaping

Bandwidth shaping enables an available better utilization of a high speed connection during a 24 hour period. Since most businesses have a single connection used for both general internet access and replication, it is preferred to maximize internet traffic during the day while maximizing replication traffic during off hours. Bandwidth shaping determines on an hourly basis how a single internet connection can be optimally used. This feature allows the admin to set bandwidth usage parameters on an hourly basis to best use limited resources. Figure 3 illustrates a sample Bandwidth Shaping schedule.

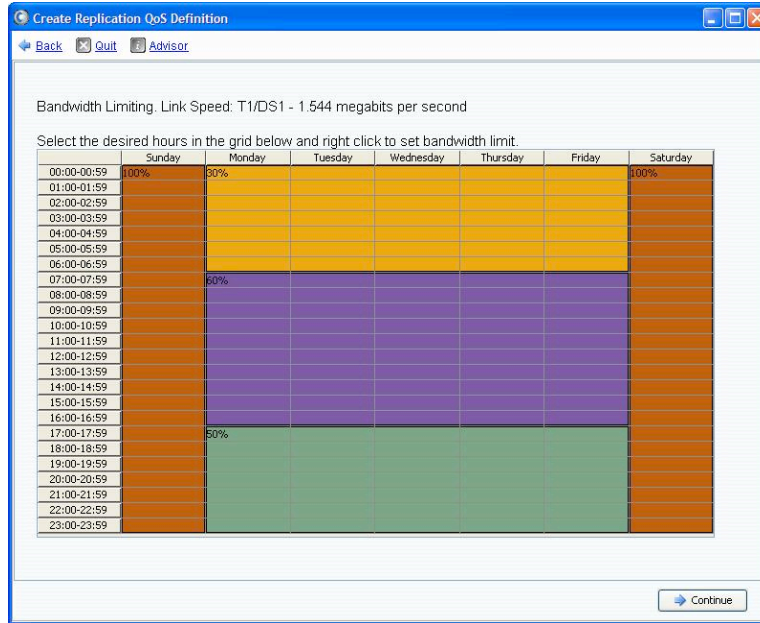


Figure 3 Bandwidth Shaping Schedule

IP-based Replications

With a native iSCSI interface, Remote Instant Replay allows replications over long distance using existing Ethernet networks rather than Fibre Channel HBAs and Fibre Channel switches. iSCSI takes SCSI protocol data and encapsulates it in IP packets so it can travel over any IP-based network.

IP-based replication is simply enabled by connecting an iSCSI card directly into the controller for each system.

Remote Instant Replay also supports traditional Fibre Channel replication.

Multi-site Replication

With Remote Instant Replay, unlimited Replays deliver remote replication capabilities to both primary and remote sites, allowing multiple locations as recover points. Any number of locations can actively and simultaneously replicate to a single primary location or to each other, without the need and expense of a dedicated site and duplicate configurations. Replications can be enabled in bi-directional point-to-point or multi-point configurations (point-to-point, peer-to-peer).

Synchronous and Asynchronous Replication

Remote Instant Replay offer two forms of replication: **Synchronous** and **Asynchronous**

Synchronous Replication

In a Synchronous Replication scheme, all data is sent to the local volume and the replication. Storage Center does not acknowledge completion of the written data until the data has been written to the local volume and the replication. This means the local volume and replication are fully synchronized. There will be no loss of data in the event of a failure on the local system.

Synchronous replication does not maintain recovery points between the local volume and the replication. If connectivity is lost between the replicating system and the target system, the entire data volume must be re-copied to ensure data is consistent in both locations. Historical Replay information will not be available from the replication volume.

Asynchronous Replication

In the Asynchronous Replication scheme, data is acknowledged as soon as it is written to the local volume. The data is set for delivery to the replication. This allows for more efficient link utilization and data transfer optimization. The downside of this is the possibility of data loss in case of a local failure.

Asynchronous Replication utilizes Replays to create recovery points between the local volume and the replication. A Replay created on the local system is sent intact to the remote system.

Use of scheduled Replays is strongly recommended for Replay creation on the local system. Specifying Replay schedules on the remote system is not recommended as the schedules are provided by the local system.

Local and Remote Site Configuration

With Remote Instant Replay, each replication site can have its own configuration. The Enterprise Manager Replication wizard provides the ability to select Data Progression attributes for each replicated volume as displayed in Figure 4. Data Progression automatically moves data between tiers of storage. For example, the primary site could consist of high-performance Fibre Channel disk while a remote configuration is configured with lower-cost SATA.

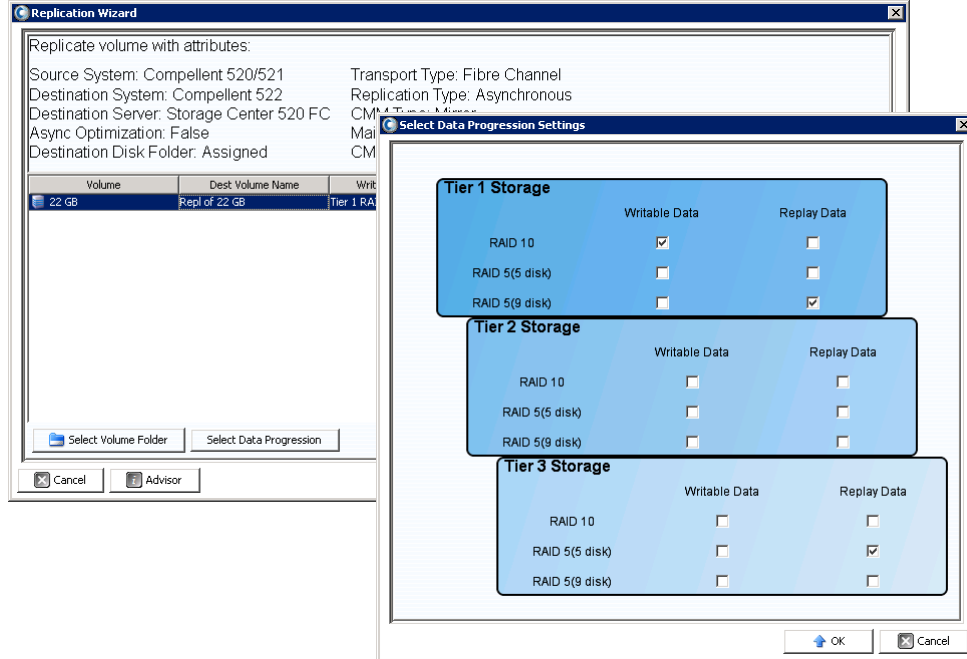


Figure 4 Data Progression Configuration

Centralized Replication Management

Creating Replications

Enterprise Manager is capable of creating multiple replications between Storage Center systems running code versions of 3.3.7 or later. Enterprise Manager provides a wizard to create a replication by simply selecting a destination system, the volumes to replicate and the replication attributes. Figure 5 displays the Replication Wizard's replication attributes.

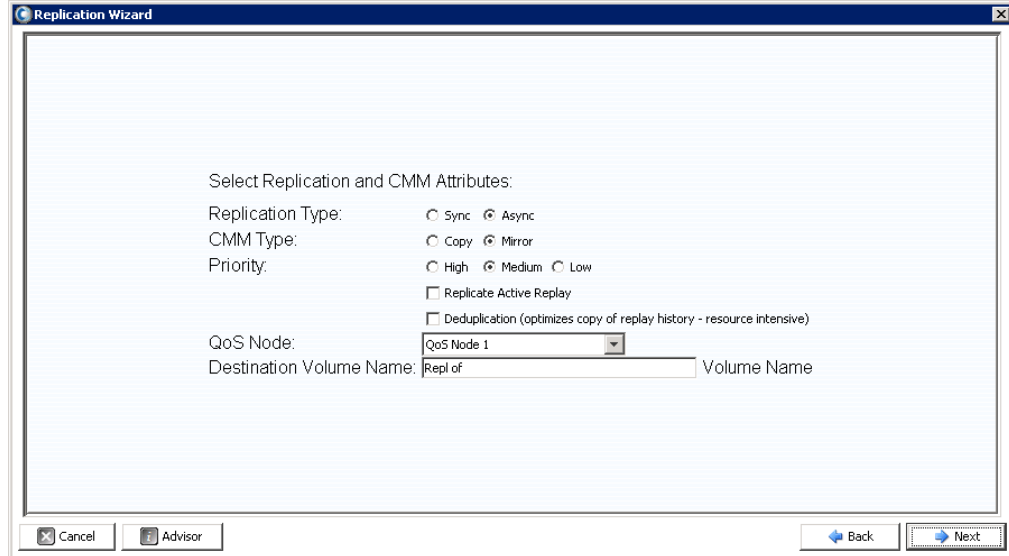


Figure 5 Enterprise Manager Replication Attributes

Recovering Data

When recovering data is required, Enterprise Manager provides a point-and-click interface to:

1. Save Restore Point

Enterprise Manager provides the ability to save the current replications that are occurring on all systems. This data is used for recovery of data that has been replicated.

2. Validate Restore Point

All saved restore points for replications can be validated. A wizard is presented with options for replications that are no longer valid: either restore the replication to the original system or activate the destination volumes on the disaster recovery site.

3. Restoring to Original Source (Non-active DR site)

When the Source volume no longer exists, Enterprise Manager provides the option to recreate the Source volume from the destination replication copy.

4. Restoring to Different Source (Non-active DR site)

Restoring to a different source occurs when the original source system no longer exists and it is desired to replace it with a new system. The same process is invoked to restore to original source where the source volume no longer exists, but to a new target.

5. Activating the Destination Site (DR site)

When the source volume is no longer available (or in most cases when the source system is no longer available), Enterprise Manager provides an option to activate the destination volume on the disaster recovery site.

6. Restoring to Original Source (Active DR site)

When a site has been activated on the disaster recovery site, it can be restored back to the original volume.

7. Restoring to Different Source (Active DR site)

When a site has been activated on the DR side, it can be restored back to another system's volume as well. This occurs when a new system is needed to replace the original source.

Flexible Replay Management

Replays can be sent via two different methods: **Active Replay** and **One Behind**

One Behind is a new option for Storage Center 3.5. Active Replay sends blocks of data across the network to the remote site as soon as the data is written. In some cases, excessive server traffic associated with the active Replay can monopolize the replication link. To avoid this behavior, One Behind waits for a Replay to be completed on the primary site before transmitting to the remote site. Once the Replay is complete, the entire Replay is sent.


Replay De-duplication

Reduction of recurring bandwidth costs is a key priority for system administrators. To further reduce bandwidth, this feature analyzes the blocks within sequential replays to reduce the amount of data sent across the pipe.

The current replay is compared to the previous replay. Any blocks that were sent with the previous replay do not need to be re-sent with the current replay.

Architecture


IP-based Replication



Notes:

- Requires iSCSI card in each controller
- Flow control handled by TCP/IP
- Compression / encryption handled by router
- Remote Instant Replay license required at each site for bi-directional replication
-

Fibre Channel-based Replication




Notes:

- Flow control handled by buffer credits on identical FC switch
- FC connection required on each side
- Remote Instant Replay license required at each site for bi-directional replication
- Primary and remote system can have different configurations

iFCP/FCIP-based Replication

Primary System



Notes:

- Flow control handled by identical iFCP devices
- FC connection required on each side
- Remote Instant Replay license required at each site for bi-directional replication
- Primary and remote system can have different configurations

Platform / Operating System

Storage Center seamlessly operates with the following operating systems:

- Microsoft Windows, Sun Solaris, HP-UX, Linux, IBM AIX, Novell Netware, Apple, Tru64, VMware