

RELDATA



Mirroring and Replication

*Local Mirroring and
Remote Replication
With Integrity*

1719 Route 10, Suite 209
Parsippany, NJ 07054
1-888-RD-IP-SAN
(1-888-734-7726)
www.reldata.com

ABSTRACT

The RELDATA 9240i storage system provides block- and file-based replication solutions to address data availability, disaster recovery, and data protection challenges faced by today's enterprises. Since replication technology significantly increases availability and shortens recovery time and recovery point windows, replication has become a core component in protecting vital business data against loss.

RELDATA mirroring and replication technology is an integral part of the RELDATA 9240i platform. With the RELDATA 9240i, users can deploy mirroring, replication, snapshot, and storage volume virtualization to achieve superior data management while reducing overhead and administration costs normally associated with deploying these technologies.

Combined with its price advantage over proprietary Fibre Channel SAN or NAS products, a RELDATA 9240i offers enterprises a wide range of solutions to meet the needs of their IT organizations.

THE NEED FOR REPLICATION

In today's modern business environment, data is the engine behind critical applications or databases that drive successful IT and business operations, so the availability of all business data is paramount. Access to data, under all circumstances, is a requirement for all businesses.

Although traditional backup and recovery methods are effective for many kinds of data loss scenarios, the major drawback of conventional backup procedures is the length of time required to restore large amounts of lost data and the unavoidable impact on business processes during both backup and recovery. In other words, conventional backup does not enhance availability of data or provide fast, disaster recovery. In addition to the high cost of business interruption, increased regulations in the commercial world are driving IT organizations to implement more effective data protection strategies, with particular emphasis on disaster recovery initiatives.

Depending on the value of the data to the business unit, traditional data protection schemes use different technologies to provide the appropriate recovery point and recovery time objectives for disk and tape backups. Whereas backup and recovery remain the mainstay of data protection operations, advanced technologies such as snapshots and replication can be used to improve recovery and availability of data beyond what backup alone can achieve.

A careful analysis of business requirements allows an optimal replication technology and strategy to be selected, with consideration of such issues as type of data and applications, number of primary and secondary storage systems, location and geography details, recovery scenarios, and the types of potential disasters that may be encountered.

RELDATA mirroring and replication technology offers local and remote replication capabilities designed to address data availability and disaster recovery scenarios. Because mirroring and replication services are an integral part of the RELDATA 9240i, enterprises are able to implement data protection as part of the storage provisioning for all business- and mission-critical applications across the enterprise.

Why iSCSI is Ideal for Mirroring and Replication

For long distance replication scenarios, the data being replicated must be packaged into some sort of a WAN communications protocol and transferred from primary systems to secondary systems. For such transfers, most vendors tend to use standard communications links, along with a TCP/IP protocol; in many cases, vendors may even have designed proprietary communications protocols for data transfer and control commands. Such an approach tends to lock customers into a single-vendor solution because the replication can only be used with the vendor's hardware.

Local mirroring solutions have been traditionally dominated by Fibre Channel based topologies. With the increasing number of IP-based storage systems being deployed, customers are growing reluctant to build Fibre Channel islands and seek out alternatives.

A better solution is a standards-based mirroring and replication solution, using iSCSI, a protocol specifically designed for storage traffic using TCP/IP. Primary systems with volumes that require replication are identified at the primary site. Secondary replica volumes of these primary volumes are created as iSCSI target volumes and linked to the primary site systems by the iSCSI initiator. Using this methodology, all primary source volumes are available directly at the primary systems and the process of replication is reduced to simply writing the blocks to the locally-available mirror or replica volumes that appear no different from locally-connected disk drives.

Application Replication Considerations

Many replication solutions are based on replication of the file system state, as opposed to the blocks of the storage device that contains the file system. The main drawback of the file-based replication is that it requires a special effort to create a checkpoint—or guaranteed consistent state—of the application's files. The checkpoint is then used to create an image of the file state on the secondary system. Without this type of operation, the resulting files at the secondary system won't be consistent with those on the primary system and, therefore, will be impossible to recover from a disaster.

Creating a consistent state of the file system usually requires assistance from the application, which is notified about the need to establish a checkpoint. Using this checkpoint notification, the application must suspend its operations and write out all transient data creating the required consistent state. Once the replication checkpoint completes, the application is notified accordingly and resumes normal operation.

The application integration is typically attained through application-specific software agents that have access to both the application and to the replication system. Some operating systems include generic support for such agents. For example, Microsoft Windows includes a component called Volume Shadow Copy Service (VSS) that facilitates suspend and resume of the application during snapshot creation. However, it still requires lightweight agents to talk to the applications and storage hardware.

These restrictions that are imposed by file-level replication are eliminated by replication on the block level using the RELDATA 9240i. In addition, it will reduce the requirement for software agents. Block level replication allows the possibility of check pointing at any time because the process persists the data in a consistent manner that allows recovery. In such cases, the application simply

recovers from what appears to be a sudden shutdown, and proceeds normally. The vast majority of modern file systems and databases support such a recovery by using built-in transaction journals or logs provided by the operating system or application.

RELDATA 9240i Mirroring and Replication Modes

The RELDATA 9240i provides both synchronous mirroring and asynchronous replication modes. Synchronous mirroring is usually the preferred mode for short distance, high volume replication. Asynchronous replication is available for long distance scenarios.

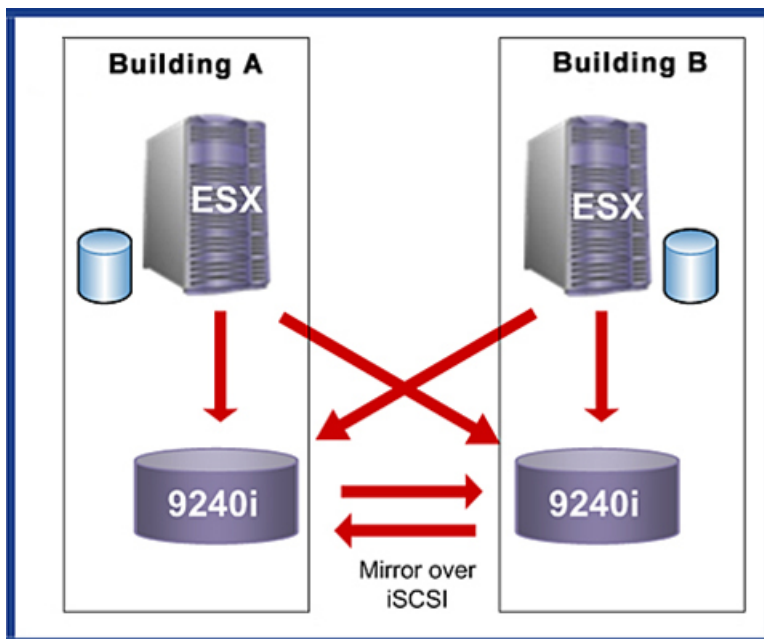


Figure 1: Campus Mirroring for HA
Maintain Multiple Synchronous LUN copies across LAN/WAN
with load balancing & zero TTR if one side of mirror fails.

Synchronous mirroring provides the ability to synchronize writes between the primary site data volumes and their designated secondary site mirror volumes. With each write to the primary data volume, the RELDATA 9240i writes the same data to the secondary site mirror volumes. All writes must complete before a completion indicator is returned to the host application. Since the distance over which data is transmitted can introduce latency effects, synchronous mirroring is limited to short distances between the primary and secondary sites.

Asynchronous replication is available for those scenarios where the secondary sites are geographically distant and where the communications link has significantly lower bandwidth capability. The latencies introduced by long distances prevent synchronization of the writes to the primary and secondary volumes, thus requiring a separate queuing function for the remote replica

writes. This separate queuing function allows the writes to the secondary volumes to be de-coupled from the writes to the primary volumes; hence, the term asynchronous. The RELDATA 9240i creates a snapshot of the replica writes and sends this to the secondary systems. Using snapshot technology, the RELDATA 9240i is able to manage and pace the writes without regard to the primary volume activity, in addition to optimizing the use of lower bandwidth communications links to the secondary sites.

Volume Mirroring vs. Replication

The choice between mirroring and replication of volumes provides the ability to set Recovery Point Objectives (RPO) and Recovery Time Objectives (RTO) for each and every replication scenario. For synchronous mirroring, the primary and secondary mirror volumes are lock-stepped for absolute synchronization. For asynchronous replication, the granularity of the snapshots is configured for the desired synchronization necessary to meet RPO and RTO objectives.

RELDATA Consistency Groups

If the application is simultaneously accessing several RELDATA 9240i volumes, the replication checkpoint is synchronized across all volumes involved. Database applications require that all volumes be synchronized in replication scenarios so that all I/Os maintain their relative sequence both within a volume and between volumes. This ensures the secondary system sees the I/Os in the same order that occurred in the original primary system. The replication checkpoint is also necessary to ensure that no write operations have been reported to the application as completed between the checkpoints on different volumes. For this purpose, the volumes are gathered into a consistency group. The checkpoint is declared at exact same time for all volumes in the group, putting the completion of write operations on hold across the entire group and resuming them after the checkpoint has completed.

Asynchronous Replication using Snapshots

Snapshots offer a simple means for obtaining instant, consistent, and stable copies of the primary data volumes that serve as asynchronous replication checkpoints. RELDATA snapshot technology is an integral component of the RELDATA 9240i.

Asynchronous Mode Snapshots

Two consecutive RELDATA snapshots can also serve as the means for deriving the incremental difference between two points in time of two RELDATA replication checkpoints. The incremental update is a list of block numbers modified after the first snapshot was taken but before the second snapshot is taken. Since both snapshots are "frozen (that is, they no longer change), the

comparison between the two snapshots is constant, even as the application continues to access and update the original primary volumes. This allows the incremental changes to be written out to the secondary replica volumes at virtually any speed, while addressing and mitigating bandwidth constraints and other communications considerations.

RELDATA Snapshot Efficiency

One of the many benefits of the RELDATA snapshot-based incremental replication is its bandwidth utilization efficiency. This efficiency results from the elimination of multiple updates of the same block of the primary data volume, a feature that becomes very important when an application tends to update the same blocks over and over. Such behavior is often observed with transaction-oriented applications such as database management systems or email servers. By eliminating all of the duplicate updates to a single block and repeating this reduction process for all blocks, the snapshot makes the most efficient use of limited bandwidth capacities. Reducing the amount of data to be transmitted over the WAN to a single update consisting of only the latest contents of the block can save IT the expense of unnecessary bandwidth for disaster recovery purposes.

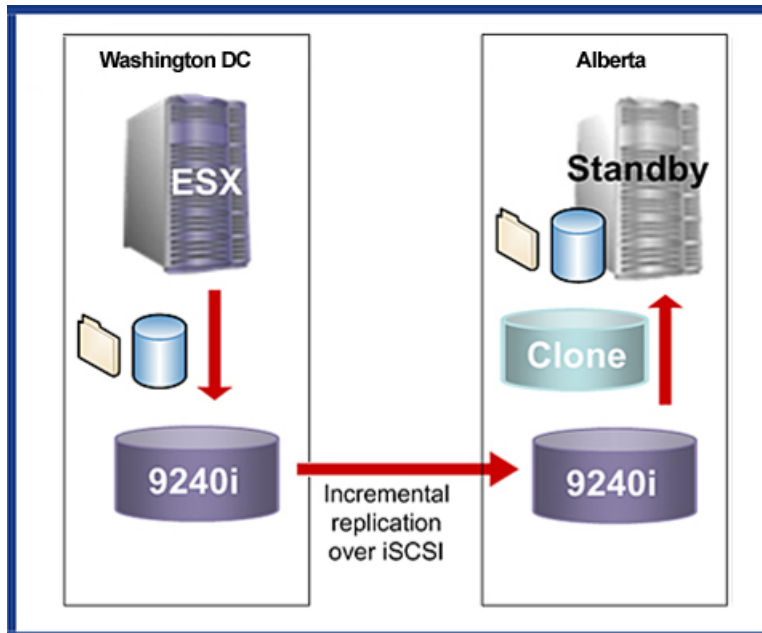


Figure 2: Replication for Disaster Recovery
Incrementally replicate LUNs to a remote location via WAN.
Leverage remote clone for testing & failover to the DR site.

SUMMARY

The use of the iSCSI protocol by the RELDATA 9240i expands interoperability of storage through an industry-standard communications infrastructure. This approach offers many advantages over proprietary vendor approaches, such as:

- The use of low-cost communications infrastructure
- Highly reliable and available replication
- High throughput for local mirroring and remote replication

In addition, RELDATA delivers its 9240i storage solution in the form of an extremely easy-to-use, cost-effective appliance. This reduces initial deployment, on-going maintenance, and administration costs, all at a price that's only a fraction of a competitor's proprietary bundled arrays. That's why RELDATA allows enterprises to align storage with business objectives, whether it's performance, access, availability, or disaster protection at a TCO that makes business sense.