



RELDATA WAN Replication

Remote Replication with Integrity

A Technical Whitepaper

December 2006

Abstract

The RELDATA 9240 IP unified storage system provides volume and file-based replication solutions to address many of the data availability, disaster recovery and data protection challenges faced by today's enterprises. Since replication technology can significantly shorten recovery time and recovery point windows, replication has become a core component in protecting vital business data from loss.

RELDATA WAN replication technology is an integral part of the company's SAN-level embedded virtualization software platform. As a result, users have the freedom to deploy replication, snapshot, and storage volume virtualization to achieve superior data management and to reduce overhead and administration normally associated with deploying these technologies.

Combined with its price advantage over proprietary Fibre Channel SAN or branded NAS, RELDATA-based storage systems offer businesses a storage-based solution to many challenges faced by contemporary IT organizations.

Content

Why Data Replication?	3
Why iSCSI is Ideal for Replication.....	3
Application Aware vs. RELDATA SAN-level Virtualization.....	4
RELDATA Consistency Groups	5
Volume Mirroring	5
Asynchronous Replication using RELDATA SAN-level snapshots.....	5
Snapshots Integrate with RELDATA WAN replication Asynch Mode.....	6
RELDATA Data Coalescing and Snapshots	7
Data encryption and compression.....	7
Summary	8

December 2006

Copyright ©2006 RELDATA Inc., all rights reserved. The copyright and all intellectual property rights in this document belong to RELDATA. It is strictly forbidden to copy, duplicate or otherwise use this document or any part thereof in any way shape or form without the prior written consent of RELDATA. RELDATA reserves the right to make changes to this document at any time, without notice, and assumes no responsibility for its use. This informational document describes features that may not be currently available. RELDATA and the RELDATA Logo are trademarks of RELDATA Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

U.S.A.:

RELDATA Inc.
www.reldata.com
Phone +1 (973) 644 2770
Fax +1 (973) 644 3385
Email sales@reldata.com

Europe:

RELDATA Europe GmbH
Waldburgstrasse 15
D-70563 Stuttgart, Germany
Phone +49 (711) 2272020
Fax +49 (711) 22720210
Email sales.emea@reldata.com

Why Data Replication?

In today's modern business environment, the availability of data is paramount as it is the engine behind critical applications or databases that drive successful IT and business operations. Why? The advent of cheap, high performance disk storage over the last decade has fueled a transformation of many businesses into paperless offices and entirely electronic communications. The majority of relevant and important business information is now kept in the form of digital files.

This online, 24x7x365 digital world demands access to data under any and all circumstances. Traditional backup and recovery methods are effective for many kinds of data loss scenarios. But the major drawback of backup is the relatively long time it takes to restore a large amount of lost data. Moreover, backup is not designed to improve availability of data or provide fast, redundant disaster recovery. The high cost of business interruption plus a new more regulated commercial world is providing the impetus for IT organizations to deploy a layer data protection strategy including redundant disaster recovery.

Contemporary layered data protection approach uses different technologies to provide the appropriate recovery point and recovery time objective on disk and tape depending on the value of the data to the business unit. Backup and recovery remains the mainstay of data protection operations; yet, disk to disk recovery like snapshots and replication are used to improve recovery and availability of data beyond what backup alone can achieve.

Replication comes in many sizes, styles, and types just like blue jeans. Many considerations are necessary when selecting the optimal replication technology and strategy. Considerations as diverse as geography, number of replicas, type of data or application, recovery scenarios, even the type of disaster must be carefully taken into account.

RELDATA WAN replication technology offers the local and remote replication capabilities designed to address disaster recovery and data availability scenarios. Because WAN replication is based on the foundation of our SAN-level unified storage virtualization platform, IT can implement data protection as part of storage provisioning for file systems, relational databases, even email applications.

Why iSCSI is Ideal for Replication

In a long-distance replication scenario, the data being replicated must be ultimately packaged into some sort of a WAN communication protocol and sent from origin system to replica system(s).

RELDATA WAN Replication – Remote Replication with Integrity

Most vendors tend to use standard communication links, such as Internet and TCP/IP for such transfers, although in many cases they choose a proprietary communication protocol for transferring data and control commands. Such approach tends to limit the customers to a single-vendor lock-in solution because the replication can only take place between the vendor's hardware.

Since the iSCSI is a protocol developed specifically for wrapping storage traffic into TCP/IP, it is obviously the best choice for a standards-based replication solution. The replica volumes are exposed as iSCSI targets that are connected via the IP link to the origin system running iSCSI initiator. That way, all replica volumes are available directly at the origin system, and the process of replication is reduced to simply writing the blocks to the locally available replica volumes that appear no different from regular locally connected disk drives.

Application Aware vs. RELDATA SAN-level Virtualization

Many available 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. This checkpoint or guaranteed consistent state is used to create an image of the state on the replica. Without this type of operation, the resulting files at the replica won't be consistent and therefore won't be useful.

Creating a consistent state of the file system typically requires some assistance from the application, which is notified about the need to establish a *checkpoint*. The application then suspends its operations and writes out all transient data creating the required consistent state that could be used as a replication checkpoint. Once the replication checkpoint is registered, the application is notified accordingly and then resumes the 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, Windows includes a component called Volume Shadow 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.

Performing replication on a block level, like RELDATA SAN-level Virtualization technology, as opposed to file level may remove the requirement for software agents and VSS support. If the replication occurs on a block level, any point in time can be used as a checkpoint because such a "checkpoint" would be identical to pulling the disk drive, rebooting the system, reconnecting the drive and restarting the application. 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 transaction journal built into the core of the operating system or application.

RELDATA WAN Replication – Remote Replication with Integrity

RELDATA Consistency Groups

If the application is simultaneously accessing several RELDATA SAN-level Virtualization storage volumes, the replication checkpoint will be synchronized across all volumes involved. This is 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 been registered.

Volume Mirroring

Mirroring of volumes is the only way to protect all transactions committed to the origin volume and maintain a zero RPO lag at all replica volumes at all times. RELDATA WAN replication technology offers volume mirroring functionality designed to offer a “no loss” Recovery Point capability for business critical application or other data sets. This feature is equivalent to expensive “continuance” volumes from competitive and very expensive Fibre Channel array vendors.

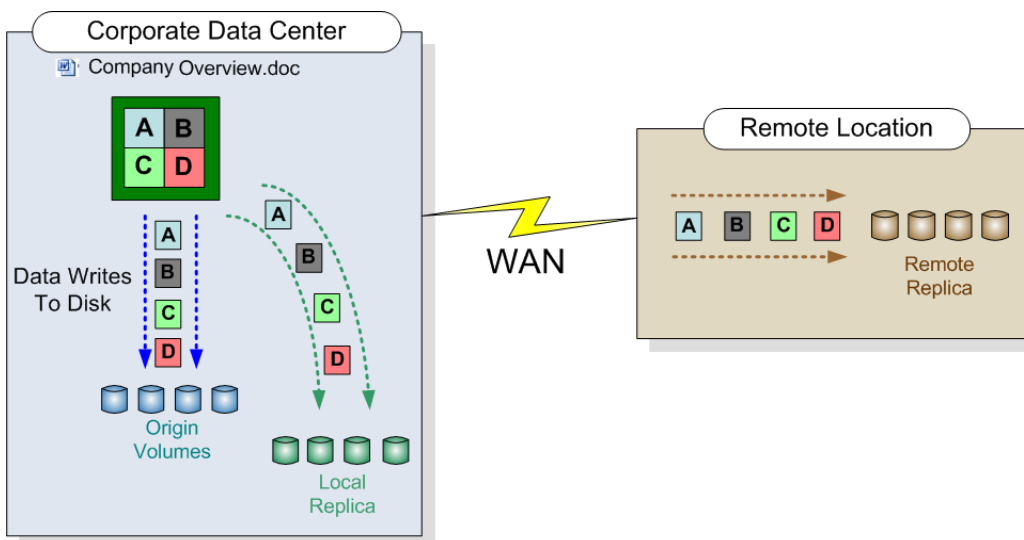


Figure 1: RELDATA Volume Mirroring to Local and Remote Locations

Asynchronous Replication using RELDATA SAN-level snapshots

Snapshots offer simple means for obtaining instant, consistent and stable copies of the origin volume that can serve as asynchronous replication checkpoints. RELDATA SAN-level snapshot technology is an integral component of the SAN-level virtualization platform.

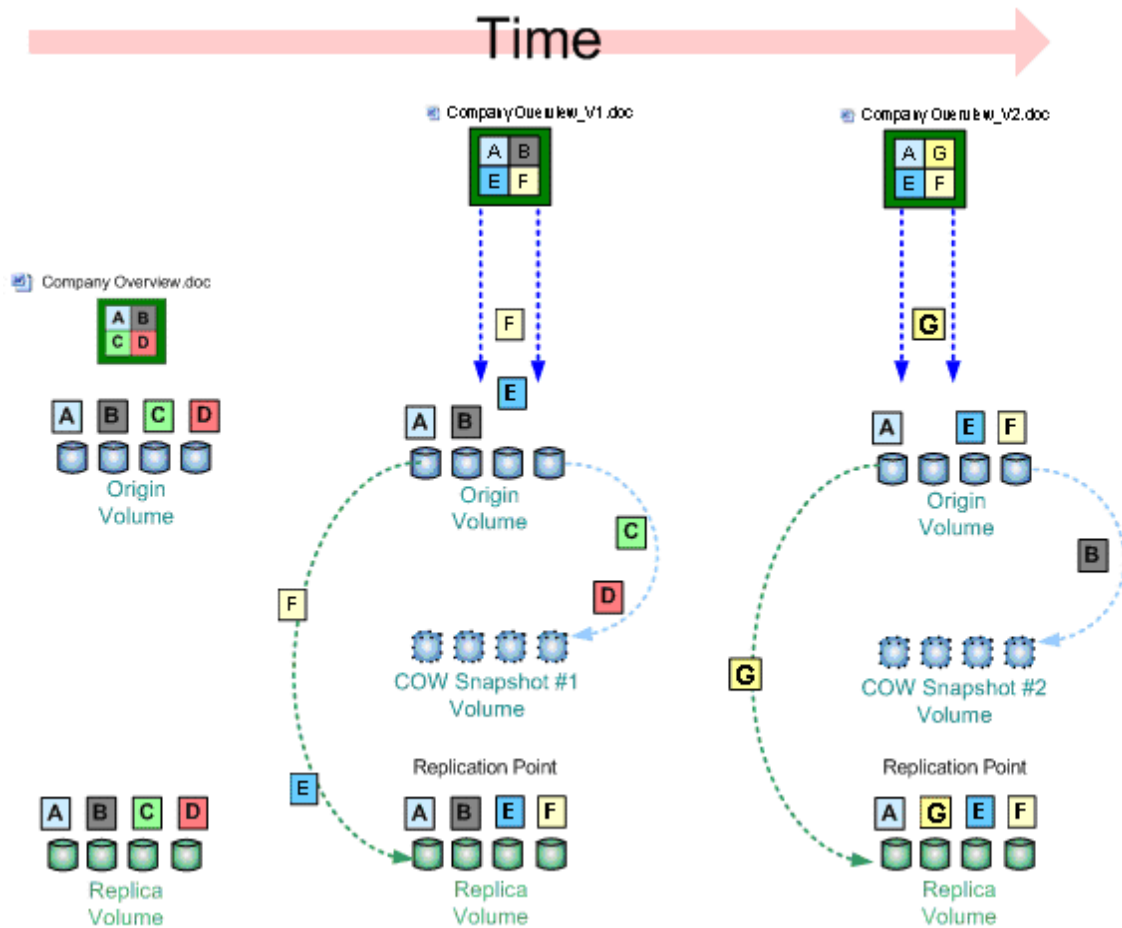


Figure 2: RELDATA Snapshots Aid in Replication

Snapshots Integrate with RELDATA WAN replication Asynch Mode

Two consecutive RELDATA snapshots can also serve as the means for deriving the incremental difference between two points in time or 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 no longer change (i.e. are “frozen”), the comparison between the two snapshots is constant even as the application continues to access and update the origin volume(s). This allows the incremental changes to be written out to the replica volumes at virtually any speed – addressing bandwidth constraints and other considerations.

RELDATA Data Coalescing and Snapshots

Among several benefits of the RELDATA snapshot-based incremental replication is the bandwidth utilization efficiency resulting from coalescing multiple updates of the same block of the origin volume. This RELDATA data reduction feature becomes very important when the 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. Reducing the amount of data to be transmitted over the WAN to a single update to the replication locations consisting of only the latest contents of the block can save IT large expenses for unnecessary bandwidth for disaster recovery purposes.

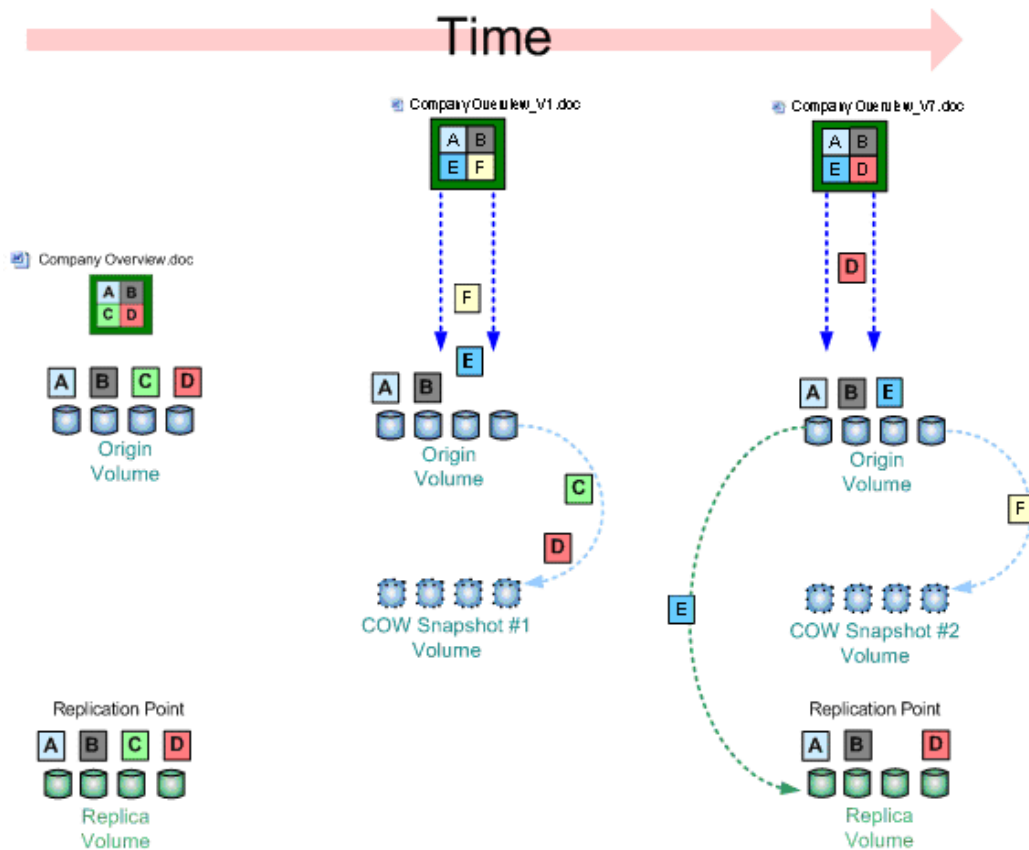


Figure 3: Coalescing Reduces Amount of Data Transferred

Data encryption and compression

One of the aspects of the replication is a consideration for potential security breaches, both for the data at-rest and the data in-flight. In a modern enterprise environment, the security of data across multiple locations is usually provided by

RELDATA WAN Replication – Remote Replication with Integrity

using a Virtual Private Network (VPN) for remote communications. VPNs typically include data encryption technologies.

In addition, the iSCSI protocol offers securing connections over IP by leveraging its mandatory support for the IPsec. IPsec specification includes encryption and optional compression. The compression can be used to further alleviate high bandwidth requirements for WAN replication.

Summary

The use of a standards-based technology adopted by RELDATA unified IP storage greatly expands interoperability of storage. This approach offers many advantages over proprietary vendor approaches, such as:

- greater choice of storage type and performance,
- ability to optimally match storage performance with application requirements,
- built-in data security,
- seamless support for compression and encryption,
- avoids vendor lock-in,
- industry-standards are well vetted, tested and reliable,
- vendors can leverage standards for performance optimization or data reduction.

In addition, RELDATA delivers its IP 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 proprietary competitor's bundled arrays. That's why RELDATA allows IT to align storage with business objectives, whether it's performance, access, availability, disaster protection at a TCO that makes business sense.