

Affordable Replication for both File and Block Storage is Attainable for Mid-sized Organizations

By DCIG Lead Analyst Jerome M Wendt



Continuing declines in the cost per GB for disk coupled with the availability of higher bandwidth, lower cost WAN connections and new products like the RELDATA 9240i is making replication an attainable objective for almost any mid-sized organization.

RELDATA

Company

RELDATA Inc.
1719 Route 10, Suite 209
Parsippany, NJ 07054 USA
Founded 2005
www.reldata.com

Industry

Data Storage

Challenges

- Identifying an affordable storage system for mid-sized organizations that replicates both file-based and block-based data
- Providing an enterprise quality data protection plan for high availability as well as local and remote disaster recovery
- Recovering mission-critical applications in a virtual environment without negatively impacting the physical host

Solution

- RELDATA 9240i Unified Storage

Benefits

- Replicates file-based and block-based data
- Leverages existing Ethernet networks to deliver replication over TCP/IP
- Supports synchronous or asynchronous replication
- Provides high availability for applications
- Enables a zero RPO/RTO when either side of the mirror is lost
- Provides automated, non-disruptive incremental application failback

Data is the lifeblood of any organization and, in the last decade, increasing emphasis has been placed on protecting that data so organizations can recover the information that they need in the time frame they need it. Replication is now rapidly emerging as a viable form of data protection even for mid-sized organizations because of its lower costs, faster recovery times and the level of protection that it provides. However the process of selecting the right replication solution for the different data types that mid-sized organizations need to protect is anything but an intuitive exercise.

Replication in its simplest form is the process of copying data to another location, be it disk or some other media. Then if the production data becomes corrupted, destroyed or inaccessible, a copy of that data can be quickly accessed to minimize or even eliminate a disruption in service to the application or end-user. In addition to being used for data protection or disaster recovery, replication can also be used to provide a secondary copy of data for testing and/or analysis.

Due to continuing declines in the cost per GB for disk, coupled with the availability of higher bandwidth and lower cost WAN connections, **replication is becoming a viable option for almost any mid-sized organization.** Unfortunately these organizations are often put in a position where they must make trade-offs as to what types of data they can afford to replicate.

Making The Right Replication Decision

Selecting the right replication process presents problems as several different methods of replication exist. One common method of replication is **file-based replication**, of which there are at least two forms. If all an organization wants to protect is a few folders and/or files on an NFS file share, then a simple Linux utility such as rsync will do the job quite nicely as it copies files from one file server to another. On the other hand, if an organization wants to accomplish the same thing on a Microsoft Windows file server, then the distributed file system (DFS), a built-in feature of Microsoft Windows, will suffice.

However storage systems can also perform what is known as **block-based replication**. Block-based replication can occur between two physically separate storage systems as well as in the form of snapshots or mirrors which are copies of production data made internally on disk on the storage system.

The advantage of block-based replication is that it does not care what protocol or file system (NFS, CIFS, ext3, xfs, NTFS, etc.) an applications uses. It operates below the file system stack and replicates raw blocks of data irrespective of the file system that sits above it. This method is most often used for replicating application databases where the file system replication is not the best choice.

The RELDATA 9240i gives mid-sized organizations access to replication services for both file and block data types at a fraction of the cost of other enterprise solutions."

— Jerome Wendt, DCIG Lead Analyst

Addressing The Special Needs Of Midsize Enterprise Replication

Here is where the problem comes in for midsize organizations. They may want to introduce replication into their environment to take advantage of its benefits as a data protection solution or as a tool for testing and/or analysis. However *it is also possible that they have a need for both forms of replication*: file-based (protecting data on NFS and CIFS shares) as well as block-based (creating a less-expensive second shelf behind primary storage), for local and remote disaster recovery and/or for higher availability. But to purchase two or three separate storage systems that support each of these different types of replication is beyond what midsize organizations can pay.

This is where a single storage system such as the **RELDATA 9240i Unified Storage solution** comes into play as the 9240i *supports both file-based and block-based replication*. This enables an organization to choose a single storage solution to meet these disparate replication requirements. In addition, the RELDATA solution uses Ethernet for replication, which is a less expensive alternative to costly fibre channel switches and cables, and can leverage existing network resources.

The 9240i uses iSCSI to support block-based replication over a regular TCP/IP connection and can do either synchronous or asynchronous replication. Configured for synchronous replication, organizations can use the 9240i to mirror critical data over a short distance.

In this fashion, primary and secondary copies of data can be kept in lock-step with each other and provide real-time copies of data in two locations for high availability of applications utilizing file or

block based data. The failover is instantaneous, *enabling a zero RPO/RTO* when either side of the mirror is lost, along with *fully-automatic, non-disruptive incremental fallback*.

Enterprise-quality Solution With Mid-sized Affordability: The Reldata 9240i

If an organization is looking to provide a disaster recovery solution at a remote location, the 9240i also supports asynchronous replication. Used this way, mid-sized organizations can replicate either their file- or block-based data to one or more remote locations without needing multiple storage solutions.

Incremental replication dramatically reduces RPO compared with traditional backup solutions. Since remote copies are ready for production deployment immediately, the organization can enjoy a near-zero RTO in the event of contingency. Asynchronous replication can also be combined with synchronous mirroring to deliver a comprehensive DR solution.

Replication is a powerful new tool that was once the exclusive domain of enterprise organizations that had big budgets, big data centers and lots of IT staff to manage it.

Not anymore.

The RELDATA 9240i gives mid-sized organizations access to replication services for both file and block data types at a fraction of the cost of other enterprise solutions.

About DCIG

DCIG analyzes software, hardware and services companies within the storage and ESI industries. DCIG distributes industry, company and product analysis by way of viral marketing and community building using the burgeoning BLOG infrastructures created worldwide.

About RELDATA

RELDATA is the innovation leader in unified block and file storage solutions over iSCSI which cost-effectively enable mid-size and distributed enterprises to nondisruptively consolidate, virtualize and protect business data across multi-vendor platforms, providing highly scalable, cost-effective integrated storage, disaster recovery and high availability solutions. The RELDATA Unified Storage Engine (USE™) platform allows companies to integrate their existing storage infrastructure to minimize the functional compromise of siloed block and file solutions while reducing administration costs and data protection compliance risk, and preserve current storage investments at lower total costs. Founded in 2005 and headquartered in Parsippany, New Jersey, RELDATA solutions are available through a global network of value added resellers. For more information, visit www.reldata.com.

RELDATA



DCIG, LLC | 7511 Madison Street | Omaha NE 68127 | 402.884.9594
dciginc.com