

Cool Vendors in Data Protection, 2008

Dave Russell, John P Morency

This research details five disk-based data protection solutions from emerging data protection vendors.

IT directors and storage administrators responsible for backup/recovery and/or disaster recovery will want to consider these vendors for applicability in their environments.

Key Findings

- Disk-based backup and recovery solutions offer many benefits, and as the number of vendors expands in this growing market, approaches and implementations can vary widely, but the net effect is greater choice for the buyer.
- Virtual tape libraries (VTLs), data deduplication and replication techniques are some of the methods that vendors use to offer solutions to modernize the recovery and continuity infrastructure.

Recommendations

- Large and midsize companies looking to leverage current storage assets to create a scalable VTL solution with data deduplication, or those that look to Hitachi Data Systems, Sun Microsystems or Overland Storage for storage solutions, should consider Diligent's offering.
- Midsize enterprises should consider ExaGrid Systems deduplication as a cost-effective means of significantly reducing the disk requirements and, therefore, the expense of a disk-based backup solution that includes built-in replication for tape elimination.
- Midsize enterprises and remote-office installations should consider FilesX for a complete recovery solution for Microsoft Windows environments that provides replication, continuous data protection (CDP) and bare metal restore (BMR) capabilities for files, applications and the operating system.
- Small and midsize businesses that require a heterogeneous and simplified set of data recovery capabilities for a broad set of mission-critical application and data recovery requirements should evaluate InMage's DR-Scout.
- Large and midsize enterprises should consider Sepaton VTLs as a means of easily incorporating disk into their recovery infrastructures through its fast, scalable, well-integrated and easy-to-install-and-configure solutions.

ANALYSIS

This research does not constitute an exhaustive list of vendors in any given technology area, but rather is designed to highlight interesting, new and innovative vendors, products and services. Gartner disclaims all warranties, express or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

What You Need to Know

In early 2006, Gartner stated that "Disk-to-disk backup and, therefore, disk-based recoveries will be the most significant technology impact on recovery in the next five years." In 2008, Gartner reiterates that this trend will continue. In fact, we have seen a significant uptake in the adoption of disk-based recovery solutions in 2007 and early 2008. End-user inquiry and conference polling results show that disk-based backup, heterogeneous replication, deduplication and VTL offerings are being considered and deployed at an accelerating rate. Deduplication and VTL discussions account for a significant amount of the Gartner inquiry load in the data recovery area.

New and less-expensive disk options make the use of disk for faster backup and recovery a more viable option than traditional backup-to-tape methods. This is not to suggest that tape is no longer relevant for data protection. Tape can still play a vital role for longer-term retention and for a tertiary copy of data; however, directionally, the market is moving quickly to solutions that enable faster recovery solutions that, in some cases, allow for more-frequent recovery points.

In recent years, some of the most interesting innovations in backup/recovery have come from smaller vendors at the periphery, and not from larger vendors at the core. The giant vendors may consume the new vendors, but it is also possible that, during the next three to five years, some emerging vendors may survive and thrive in their own right, and make a minor, but meaningful, impact on the storage landscape. Risk-averse organizations may want to selectively consider a trial with one of the emerging vendors to better access technical innovation and vendor support, while keeping much of the data on established vendor solutions. As the technology and vendor prove themselves over time, future purchases could be directed toward the emerging vendor, as warranted.

The following are five Cool Vendors that provide disk-based recovery solutions.

Diligent, Framingham, Massachusetts (www.diligent.com)

Analysis by Dave Russell

Why Cool: Diligent began shipping a mainframe, software-based VTL solution in late 2002. Later, the company introduced an open-systems, software-based VTL. It was the introduction of the ProtecTIER product for open-systems deduplication at the end of 2005 that really established the company. Diligent claims that its HyperFactor proprietary, in-line deduplication engine is not only the fastest such offering (up to 450MB-per-second, sustained in-line throughput), but that backup workloads can be reduced by a ratio of up to 25:1 or higher. Gartner clients have validated the performance and deduplication claims, as well as attested to the solution's scalability up to 1PB of storage.

As with all deduplication products, the largest data reduction ratios may take up to four months to achieve. However, ProtecTIER enables significant reduction immediately on the first backup. ProtecTIER is a "target-side" deduplication solution, meaning that the data reduction occurs as the backup is written to disk after the backup software, so any application, any operating system and any backup software is supported, and data is available immediately for replicating off-site.

Challenges: ProtecTIER was once perceived as an expensive solution that required large amounts of Fibre Channel (FC) disk spindles. Recently Hitachi Data Systems, Overland Storage and Sun Microsystems have announced midmarket bundles based on Diligent ProtecTIER that offer broader distribution, Serial Advanced Technology Attachment (SATA) disk drives and a lower price point, which have the potential to address these challenges. The lack of built-in replication capabilities and only offering a VTL interface may be issues for some organizations; however, expect Diligent to offer its own replication in the near future that complements the use of embedded disk-based replication.

Who Should Care: Large and midsize enterprises seeking a scalable VTL solution with fast in-line deduplication that supports any backup application.

ExaGrid Systems, Westborough, Massachusetts (www.exagrid.com)

Analysis by Dave Russell

Why Cool: ExaGrid offers its target-side deduplication systems of the same name as the company. Starting at around \$20,000, ExaGrid solutions frequently are much lower in cost than other disk-based solutions. Interesting features include target-side deduplication, "post-processing" deduplication after the backup data has been written to disk and a grid architecture. Post-processing means that the data lands as-is on disk first, for faster backup performance, and then a process runs to compare the latest data to bytes of previous backups. The grid architecture enables customers to start with solutions as small as 1TB and then add up additional nodes that range from 1TB to 5TB to achieve a maximum of 60TB. The grid architecture also automatically load balances workload across the ExaGrid systems. The highest deduplication ratios may require three to four months to achieve, as retention grows; however, with standard compression and using the common "full + incremental" backup methodology, ratios in the 5:1 to 10:1 range are achievable in a few weeks.

Off-site tape can still be created or, through the included built-in replication feature, off-site copies can be reduced or eliminated, as ExaGrid replicates deduplicated data across the wire to a remote location. ExaGrid is content-aware, meaning that it must understand the backup application. Currently supported backup programs include ARCserve, Backup Exec, CommVault, NetBackup and NetWorker. ExaGrid claims that installation typically takes less than an hour.

Challenges: As a small, independent company with only recent international channels, ExaGrid lacks the visibility of other vendors; however, recent new funding (from Lehman Brothers) is aimed at improving marketing and channels. Currently, only a network-attached storage (NAS) interface is supported — no VTL, FC or Internet Small Computer System Interface (iSCSI) storage area networks are supported. Given the target market, this may be appropriate; however, the lack of iSCSI may be an issue for some organizations. ExaGrid is rumored to be addressing the iSCSI issue in the near future. Lack of support for IBM's Tivoli Storage Manager and some other backup applications may be a gate, as well.

Who Should Care: Small-to-midsize enterprises that are looking to move into or further expand their disk-based backup/recovery capabilities with a NAS Common Internet File System (CIFS) or Network File System (NFS) interface. Also, small-to-midsize organizations seeking to eliminate tape will want to investigate ExaGrid for applicability in their environments.

FilesX, Newton, Massachusetts (www.filesx.com)

Analysis by Dave Russell

Why Cool: FilesX began shipping its Xpress Restore solution in 2004. The Xpress Restore portfolio focuses exclusively on Microsoft Windows files and application data protection. The offering provides many capabilities: data replication, true CDP, remote vaulting and, as of the 3.5

release in early 2008, a BMR functionality that enables dissimilar hardware support. All features share a common agent, server, console and policy set, making for a versatile solution that provides various data recovery techniques that protect files, applications and the base operating system. Windows files and applications, such as Exchange, Oracle, SAP, SharePoint and SQL Server, are supported.

Although CDP is a term that has experienced some backlash in the industry, Gartner inquiry and conference polling show that there is rising interest in this capability for Windows files and Exchange environments — exactly where Xpress Restore is focused. Another compelling feature is the use of block-level, incremental data capture technology, whereby after the initial full backup, only changed blocks are transmitted, making the solution particularly effective in remote-office installations. Recoveries are fast for all sizes of objects, from files to entire servers.

Challenges: Although FilesX has relationships with larger vendors, such as HP and IBM, the company needs to establish stronger routes to market to further penetrate the recovery market. The lack of non-Windows support may be an issue for some organizations; however, Windows is a fast-growing platform and the most common remote-office, midmarket and departmental operating system. The larger concern for some organizations may be the lack of robust reporting, as well as the fact that the solution is only available in English.

Who Should Care: Small-to-midsize enterprises or organizations with branch offices needing a robust, Windows-specific solution that offers traditional backup, replication, CDP and BMR capabilities in one offering. Also, managed-services providers looking for a data protection platform on which to base their service should consider FilesX.

InMage, Santa Clara, California (www.inmage.net)

Analysis by John Morency

Why Cool: InMage started in late 2001. Several years of R&D have resulted in the creation of its DR-Scout product family. DR-Scout was originally little more than a CDP offering, but strong technology "under the hood," combined with the development of complementary offerings, has led to a diverse product set that typically would be associated with larger vendors. The primary technology differentiators are a highly granular volume and file recovery logic that enable heterogeneous replication, CDP and snapshots via lightweight, multiplatform agents (DR-Scout FX for file recovery and DR-Scout-VX for volume recovery) and a dedicated appliance (DR-Scout CX). The result is product functionality that goes beyond traditional replication and mirroring offerings in the level of supported recovery granularity, as well as in additional predictive analytics and modeling that enable disaster recovery and storage managers to assess overall recovery health and to perform more-effective capacity planning for storage and the wide-area network circuits that connect the primary production and secondary recovery facilities.

Recovery of physical-to-virtual and virtual-to-physical servers can be equally supported by the same product. The result is an opportunity for the support staff to more-effectively plan data center failover without having to be as concerned about which applications and data can and cannot fail over effectively. Other potential benefits of more-flexible virtual-machine failovers include improved recovery service levels, as well as reduced capital and operations expenses. Support is also included for application-specific recovery if point failover for key production software platforms, such as Microsoft Exchange, SQL Server 2005 and Oracle Financials, is also a high priority.

Challenges: InMage's functionally rich product set typically would be associated with much larger vendors. This imposes the burden of explaining, positioning, marketing and supporting a fairly complex feature set on a limited marketing, sales and support staff. This challenge can be overcome with time, but an open question remains regarding whether the company can scale

these critical competencies before larger players, such as EMC, Hitachi, IBM, Symantec and others, turn these advanced features into more-broadly available product commodities that fit comfortably into the remediation and consulting "sweet spots" of much larger and more-mature vendor support organizations.

Who Should Care: IT operations that have limited funds for recovery and continuity will find the alternative implementation options supported by DR-Scout especially intriguing. These options range from a lower-cost, application-specific recovery scenario (for example, for Microsoft Exchange) up to and including the partial or complete recovery of a virtualized data center. In addition, IT managers concerned about the effect of increased recovery capability on wide-area circuit bandwidth costs will find the storage throughput modeling and capacity planning features supported by the DR-Scout Profiler worth a second look.

Sepaton, Marlborough, Massachusetts (www.sepaton.com)

Analysis by Dave Russell

Why Cool: Sepaton's company name is "no tapes" spelled backward. In late 2003, its first VTL product shipped. Sepaton is on its seventh generation of VTL code. The company is predominately focused on the Fortune 1000 or higher. Sepaton's largest VTL, the S2100-ES2, can be configured with up to 1.2PB of physical disk that with optional hardware-assisted compression can bring the single-unit capacity to approximately 2.5PB. Large scale and fast performance are what Sepaton's VTLs are known for; however, in 2007, a smaller appliance, the S2100-DS2, was introduced for remote-office and departmental installations.

Sepaton was early to offer features such as hardware compression, built-in replication across VTLs for off-site vaulting of data and an oversubscription feature that is essentially thin provisioning for the VTL. Sepaton's VTL features a grid architecture that automatically balances workload across up to 16 computing nodes. The larger and smaller units come with a single integrated console for managing the VTL engine and the disk, and both solutions are preassembled and largely preconfigured, enabling rapid deployment. HP uses the Sepaton VTL engine for use in its VLS disk library systems.

Challenges: Sepaton was also early in announcing a deduplication feature called DeltaStor in early 2007. More than a year later, deployments of DeltaStor deduplication are modest, and only NetBackup as a backup application is supported. Expect Sepaton to update its deduplication capabilities in mid-2008. Some organizations may be concerned by the lack of a high-availability feature for the larger VTL, given the amount of data that can be stored on it.

Who Should Care: Large and midsize enterprises that are seeking to move into or further expand their disk-based backup/recovery capabilities via a well-integrated, fast and scalable VTL will want to investigate Sepaton.

RECOMMENDED READING

"Data Deduplication Is Poised to Transform Backup and Recovery"

"Virtual Tape Libraries for Open Systems Overview"

"The Half-Life of a Virtual Tape Library"

REGIONAL HEADQUARTERS

Corporate Headquarters

56 Top Gallant Road
Stamford, CT 06902-7700
U.S.A.
+1 203 964 0096

European Headquarters

Tamesis
The Glanty
Egham
Surrey, TW20 9AW
UNITED KINGDOM
+44 1784 431611

Asia/Pacific Headquarters

Gartner Australasia Pty. Ltd.
Level 9, 141 Walker Street
North Sydney
New South Wales 2060
AUSTRALIA
+61 2 9459 4600

Japan Headquarters

Gartner Japan Ltd.
Aobadai Hills, 6F
7-7, Aobadai, 4-chome
Meguro-ku, Tokyo 153-0042
JAPAN
+81 3 3481 3670

Latin America Headquarters

Gartner do Brazil
Av. das Nações Unidas, 12551
9º andar—World Trade Center
04578-903—São Paulo SP
BRAZIL
+55 11 3443 1509