Managing the information that drives the enterprise

STORAGE

ESSENTIAL GUIDE TO

Replication for DR

Replication for disaster recovery (DR) is no longer a "nice to have" technology, but a necessary part of every business continuity strategy.

INSIDE

TechTarget
The IT Media ROI Experts

Replication is essential to DR

Once considered a DR solution only for well-heeled enterprises, replication is now a flexible and affordable alternative for all companies.

NOT TOO LONG AGO, disaster recovery (DR) was an exercise in logistics involving hundreds of tapes, elaborate transportation schemes and numerous personnel assignments. But no matter how well a company prepared, it was almost inevitable that the sheer complexity of the plan meant something would fall through the cracks. A single missing tape could cripple a recovery.

Tape has its place in data protection, including disaster recovery, but not for the mission-critical apps that a company relies on to maintain business as usual—or as close to usual as possible. Today's business environment is almost a cliché: 24/7 operations, global reach, interwoven networks of suppliers and customers, and a sense that a day in the docks could spell doom for a company. With that kind of urgency, and recovery time objectives (RTOs) that are barely a blink of the eye, tape just doesn't cut it.

Dozens of newer replication products sidestep the constraints of earlier generations of replicators, principally by allowing data to be replicated between different storage configurations.

Data replication is the obvious alternative to tape recoveries. Typical replication setups called for array-to-array replication, which required duplicating the production site configuration at the recovery site. To say the least, duplicating primary storage for recovery purposes often meant an extraordinary expense to ensure business continuity. Other approaches, like host-based replication, are cheaper to implement, but they come at a cost, as well with greater maintenance and administration.

In the last few years, storage managers have received some relief on the budgetary and administrative sides of the replication equation. Dozens of newer replication products sidestep the constraints of earlier generations of replicators, principally by allowing data to be replicated between different storage configurations.

On the server side of the shop, virtualization has radically redefined disaster recovery by breaking the one-to-one, physical server-to-physical

server setup that had been required for a recovery site. Now, a single server hosting many virtual servers may serve as the DR mechanism to support multiple primary servers. And with virtual server setups at both primary and recovery sites, the flexibility is almost endless. Fortunately, newer replication products are up to the task and can match the resiliency and flexibility of virtualized server environments.

That's not to suggest that data replication for disaster recovery is necessarily a lot easier to implement. There are still a lot of decisions to make, like when to use synchronous vs. asynchronous replication, and the appropriate types of recovery site storage that won't bust budgets but will still be up to the task of maintaining business processes at an acceptable level.

On the following pages, you'll get a good look at the various replication alternatives, where they might be most effectively deployed, and how server and storage virtualization can be best exploited. Replication may never entirely eliminate tape from DR scenarios, but it has certainly provided a faster and more reliable alternative. ⊙

Rich Castagna (rcastagna@storagemagazine.com) is editorial director of the Storage Media Group.

Replication ALTERNATIVES

Data replication is great for protecting critical data and ensuring quick recoveries. Find out where you should deploy replication: in your array, network or servers.

By Jacob Gsoedl

DATA REPLICATION as a means of data protection has seen continuous and increasing adoption since it first emerged in storage systems after the first World Trade Center bombing in 1993. Over time, it has evolved into an indispensable component of disaster recovery (DR), as well as for operational backup for applications that require shorter recovery point objectives (RPOs) and recovery time objectives (RTOs) than what traditional tape backups can offer. Firms are also adopting data replication for remote- and branch-office data protection; in a hub-and-spoke architecture, branch-office data can be replicated back to central data centers, thus eliminating unwieldy tape-based backup procedures at the branch sites.

The growing adoption of replication services has been driven by a wide array of data replication products, more lower cost replication offerings, faster and less-expensive networks, and an overall maturing of the technology itself. “Replication-based data protection is among the top three priorities of 60% of our clients, which is very different from only a few years ago,” said Tim Bowers, global product manager, storage services at EDS, a Hewlett-Packard (HP) company.

NOT ALL REPLICATION IS EQUAL

At a macro level, data replication copies data from one storage location to one or more other local or remote storage systems. But venture beyond that basic task and you'll find that data replication products vary in several key aspects:

LOCATION: One of the main differentiators among products is where replication occurs. The replication service or software can reside on the storage array, in the network or on the host (server). Array-based replication has been dominating the replication market up to now.

"We did a recent study that shows that in 2007, 83.7% of worldwide revenue for storage-based replication was done using array-to-array replication, followed by host-based replication with 11.5% and network-based replication with 4.8%," said James Baker, research manager, storage software at Framingham, Mass.-based IDC. But according to the same study, both host- and network-based replication are catching up. Host-based replication is expected to grow at a compound annual growth rate (CAGR) of 18.2% until 2012, while a CAGR of 15.4% is anticipated for network-based replication. Both are expected to expand significantly faster than the 10% forecasted annual growth for array-based replication.

MODE: Replication can occur synchronously, where data is written to the primary and secondary storage systems simultaneously; or it can be performed asynchronously, where data is replicated to replication targets with a delay. In synchronous replication, the primary storage system only commits I/O writes after the replication target acknowledges that data has been written successfully. Synchronous replication depends on sufficient bandwidth and low latency, and supported replication distances range from 50 km to 300 km. It's typically used in apps where zero RPOs and RTOs are required, such as high-availability clusters and mission-critical applications that demand 100% synchronicity between the primary and target systems. Conversely, asynchronous replication writes data to the primary array first and, depending on the implementation approach, commits data to be replicated to memory or a disk-based journal. It then copies the data in real-time or at scheduled intervals to replication targets. Unlike synchronous replication, it's designed to work over long distances and greatly reduces bandwidth requirements. While the majority of array- and network-based replication products support both synchronous and asynchronous replication, host-based replication offerings usually only come with asynchronous replication.

TYPE: Replication products can replicate blocks of data on volumes or logical unit numbers (LUNs), or replication can be performed at the file level. With the exception of network-attached storage (NAS), which can support both block- and file-based replication, array-based replication products usually operate at the block level. The same is true for network-based replication products. In contrast, most host-based replication offer-

ings operate at the file-system level. Block-based replication is platform-agnostic and will work seamlessly across various OSes. File-based replication products are very much OS-specific and the majority of available host-based replication products are written for Windows. Unlike file-based replication, block-based replication products have no knowledge of the attached platform, file system or applications, and depend on auxiliary services like snapshots for any type of app integration. As a result, most storage arrays with replication support also provide snapshot capabilities that are more or less integrated with the file system and key apps like Exchange and SQL Server databases.

ARRAY-BASED REPLICATION

In array-based replication, the replication software runs on one or more storage controllers. It's most prevalent in medium- and large-sized firms, mostly because larger firms have deployed higher end storage arrays that come with data replication features.

With more than 15 years of history, array-based replication is the most mature and proven replication approach, and its scalability is only constrained by the processing power of the array's storage controllers. "Customers scale replication performance in both our Clariion and Symmetrix arrays by distributing data replication across a larger number of storage processors," said Rick Walsworth, director of product marketing replication solutions at EMC Corp.

With the replication software located on the array, it's well suited for environments with a large number of servers for several reasons: it's operating system-agnostic; capable of supporting Windows and

DATA REPLICATION trends

These data replication-related trends are gradually changing data protection and disaster recovery.

- The use of replication-based data protection is increasing.
- Replication-based data protection is merging with traditional data protection, and traditional backup products are increasingly able to manage replicas and snapshots along with backups.
- The integration of replication and applications to provide application-consistent recovery is on the rise.
- Replication-based data protection is becoming an important option for protecting virtualized server environments.
- As storage-as-a-service (SaaS) and cloud-based computing become more prevalent, hosted replication offerings will become more common.
- Because replication relies on available bandwidth, wide-area network (WAN) optimization offerings from the likes of Blue Coat Systems Inc., Cisco Systems Inc., Citrix Systems Inc., F5 Networks Inc., Juniper Networks Inc., Packeteer Inc. (now a Blue Coat company), Riverbed Technology Inc. and Silver Peak Systems Inc. are used to complement replication products to preserve valuable WAN bandwidth.

Unix-based open systems, as well as mainframes (high-end arrays); licensing fees are typically based on the amount of storage rather than the number of servers attached; and it doesn't require any administrative work on attached servers. Because replication is offloaded to storage controllers, processing overhead on servers is eliminated, making array-based replication very favorable for mission-critical and high-end transactional applications.

The biggest disadvantage of array-based replication is its lack of support of heterogeneous storage systems. And unless the array provides a storage virtualization option—as Hitachi Data Systems does for its Universal Storage Platform (USP)—array-based replication usually only works between similar array models. Besides a high degree of vendor lock-in, entry cost for array-based replication is relatively high, and it could be particularly expensive for companies that have to support a large number of locations. In general, array-based replication works best for companies that have standardized on a single storage array vendor.

COMPARING DATA REPLICATION METHODS

	TYPE OF REPLICATION		
	ARRAY BASED	HOST BASED	NETWORK BASED
Support of heterogeneous environments	Low; only works between similar arrays	High; storage-agnostic and works with networked- and direct-attached storage	High; storage array- and platform-agnostic
Performance and scalability	Depends on the storage array; very good in high-end arrays	Good; workload is spread across servers; limited scalability because of management challenges	Very good
Cost	Requires similar arrays; high entry cost; expensive for a large number of locations	No hardware required; low entry cost; cost rises proportionally to the number of servers	Requires intelligent switches or inline appliances; high entry cost; expensive for a large number of locations
Complexity	Medium to high	Low	Medium to high
Replication modes	Synchronous and asynchronous	Asynchronous	Synchronous and asynchronous
Predominant replication type	Logical unit number (LUN) or volume block-level based	File-system based	LUN or volume block-level based

Almost all vendors of midsized to high-end arrays provide a replication feature. The replication products of these leading array vendors have made significant inroads and gained market share:

- EMC Symmetrix Remote Data Facility (SRDF) for both synchronous and asynchronous replication, and EMC MirrorView for synchronous and asynchronous replication of Clariion systems.
- Hitachi Data Systems TrueCopy for synchronous replication and Hitachi Data Systems Universal Replicator software for asynchronous replication.
- HP StorageWorks XP Continuous Access and HP StorageWorks Continuous Access Enterprise Virtual Array (EVA) for both synchronous and asynchronous replication for HP XP and EVA arrays.

CHOOSING A DATA REPLICATION SOLUTION

1. Selection of a data replication method should start with a business impact analysis to determine required recovery time objectives (RTOs) and recovery point objectives (RPOs).
2. For applications that can't accept data loss (RTO equals zero), synchronous replication is required. Heed latency in synchronous replication because it will drag down application I/O performance. If there's any risk of latency or unreliable bandwidth, or for replication beyond certain distances (50 km to 300 km), asynchronous replication is the way to go.
3. Besides the replication mode, application performance can be impacted by the replication platform. Host-based replication competes with applications for valuable processor, memory and I/O resources.
4. Have a clear understanding of the bandwidth requirements, impact on bandwidth cost, and how data replication will impact other applications and users. Clearly understand and take advantage of replication features related to bandwidth such as compression, bandwidth throttling and configurable bandwidth usage depending on the time of day. Consider wide-area network (WAN) optimization devices to preserve bandwidth.
5. Replication products that support heterogeneous environments can substantially reduce cost by supporting less-expensive or legacy arrays. They also limit vendor lock-in.
6. The disadvantage of vendor lock-in of array-based replication is offset by the advantage of close integration between replication and the storage platform and easier support, eliminating the risk of finger-pointing in multivendor configurations.

- IBM SAN Volume Controller (SVC) Metro Mirror for synchronous replication and SVC Global Mirror for asynchronous replication.
- NetApp SnapMirror for synchronous and asynchronous block-based replication, and NetApp SnapVault for file-based replication.

Even though these replication products are similar in many aspects, a close technical analysis reveals subtle differences. For instance, the efficiency of the handshake between primary and target storage systems used during synchronous replication greatly impacts the distance a replication product can support. “Metro Mirror is able to write data to the target system with a single handshake, enabling it to support distances of up to 300 km,” said Vic Pelz, consulting IT architect at IBM. That distance goes well beyond the 50 km to 200 km cited by other storage vendors.

Differences can also be found among asynchronous replication implementations. While EMC buffers data to be replicated in memory, IBM Metro Mirror tracks changes with so-called bitmaps, continuously transmitting changes and periodically re-synchronizing the source and target to ensure they stay in sync. On the other hand, Hitachi Data Systems uses change journals stored on disk in its Universal Replicator software.

> "Deploying host-based replication only requires installing the software on source and target servers and you're ready to go."
>
> —Bob Roudebush, director of solutions engineering, Double-Take Software Inc.

“The combination of disk-based change journals that are pulled by the replication targets instead of pushed by the source makes it extremely resilient, capable of automatically recovering from elongated disruptions,” said Christophe Bertrand, senior director, solutions and product marketing business continuity at Hitachi Data Systems. “Because changes are pulled by replication target arrays, valuable processing cycles are offloaded from primary arrays to secondary target arrays.”

HOST-BASED REPLICATION

In host-based replication products, the replication software runs on servers so, unlike array- and network-based replication, it doesn’t depend on additional hardware components. That makes host-based replication the least-expensive and easiest replication method to deploy.

“Deploying host-based replication only requires installing the replication software on source and target servers and you’re ready to go,” noted Bob Roudebush, director of solutions engineering at Double-Take Software Inc. It’s well suited to work in heterogeneous environments, supporting

the widest range of storage options that include both network- and direct-attached storage. While most products support Windows, Linux and Unix support is more tenuous and, therefore, platform support is clearly one of the critical evaluation criteria when selecting a host-based replication product.

On the downside, host-based replication adds processing overhead to servers and the installed replication software carries the risk of introducing unknown behavior. "For critical and high-end application servers, IT managers tend to favor array-based replication over host-based replication because it keeps server resources dedicated to the app and doesn't expose it to potential bugs or flaws in the replication software," said Lauren Whitehouse, an analyst at Milford, Mass.-based Enterprise Strategy Group. Furthermore, licensing costs and system administration duties increase proportionally with the number of servers, giving both array- and network-based replication an advantage in environments with a large number of servers. In addition, visibility in host-based replication is typically limited to source and target servers. This is very different from the centralized architectures of array- and network-based replication offerings that enable a more holistic view into the replication infrastructure.

On the downside, host-based replication adds processing overhead to servers and the installed replication software carries the risk of introducing unknown behavior.

The target markets for host-based replication products are typically small- to medium-sized businesses (SMBs) that can't afford more expensive replication alternatives, enabling them to deploy data protection and disaster recovery architectures that, until a few years ago, were only seen in larger firms. CA, Double-Take, InMage Systems Inc., Neverfail Ltd. and SteelEye Technology Inc. are some of the vendors that have enabled smaller companies to deploy replication-based DR and data protection at a fraction of the cost of array- and network-based replication. Although each of these products replicates data from one location to another, they differ in features such as efficiency, bandwidth throttling, management, high-availability failover capabilities, platform support and application integration. Only a thorough product evaluation will reveal which product offers the best fit for a given environment.

In addition to these standalone offerings, backup software vendors are integrating host-based replication into their backup suites with the hope of expanding their reach into the lucrative remote- and branch-office data protection business.

"We see a convergence of DR and data protection, and consider replication to be a feature and not a standalone product," said Marty Ward, senior director, product marketing for the Data Protection Group at Symantec Corp. Most backup software vendors already offer host-based data replication options for their backup suites; some examples include BakBone Software Inc.'s NetVault:FASTRecover (formerly NetVault:Real-Time Data Protector); CommVault Systems Inc.'s Continuous Data Replicator (CDR); EMC RepliStor to complement EMC NetWorker; and Symantec Backup Exec Continuous Protection Server (CPS) and Symantec NetBackup PureDisk with a deduplication option, as both a standalone product and a NetBackup option.

The main advantage of combining traditional backups and replication is the ability to manage replicas and backups within a single tool. Aside from their host-based replication options, backup software vendors have been working on integrating their backup suites with leading storage arrays and network-based replication products to enable customers to manage all replicas and backups with the same tool.

The main advantage of combining traditional backups and replication is the ability to manage replicas and backups within a single tool.

"Just like with Continuous Data Replicator, array-based replicas of supported arrays are integrated into the backup application index and catalog, allowing users to restore an array-based snapshot by simply right-clicking it within our application," said Brian Brockway, vice president of product management at CommVault. Similarly, Symantec's Veritas NetBackup is integrated with more than 40 arrays and virtual tape libraries (VTLs), and EMC NetWorker offers tight integration for EMC's RecoverPoint network-based replication product.

NETWORK-BASED REPLICATION

In network-based replication, the replication occurs in the network between storage arrays and servers. I/Os are split in an inline appliance or in a Fibre Channel (FC) fabric; the I/O splitter looks at the destination address of an incoming write I/O and, if it's part of a replication volume, forwards a copy of the I/O to the replication target. Network-based replication combines the benefits of array-based and host-based replication. By offloading replication from servers and arrays, it can work across a large number of server platforms and storage arrays, making it ideal for highly heterogeneous environments. Most network-based replication products also offer storage virtualization as an option or as part of the core product.

Contemporary network-based replication offerings are either inline appliances or fabric based. With inline appliances, all I/Os need to pass through the replication device. Technically, the appliances terminate all incoming I/Os and initiate new I/Os that are forwarded to the primary storage targets and, in case of write I/Os, to replicated storage targets. The inline approach has been plagued by performance and scalability issues. The poster child for inline appliances is IBM's SVC.

A scalable architecture and plenty of cache have not only enabled it to overcome performance and scalability limitations but, aided by the simplicity of the inline appliance approach compared to the more complex fabric-based implementations, it has become one of the successes in the network-based replication and virtualization market.

In fabric-based replication products, the splitting and forwarding of I/Os is performed within an FC fabric. By taking advantage of FC switching and the separating data and control path, it's the best performing and most scalable approach. The majority of fabric-based replication products run on intelligent switches from Brocade Communications Systems Inc. and Cisco Systems Inc. Even though both Brocade and Cisco offer Data Mobility Manager (DMM) for local data center replication, third-party vendors like EMC and FalconStor Software Inc. offer more advanced fabric-based replication products that run on Brocade and Cisco intelligent switches. A case in point is EMC RecoverPoint, which provides fabric-based, asynchronous continuous data protection (CDP) with app integration that's on par with commensurate host-based CDP products. Despite obvious benefits, fabric-based replication has seen lackluster adoption.

"Switch-based replication and virtualization have been over-hyped, but there are people who are working on it and over time it will become more common."

—Greg Schulz, founder and senior analyst, StorageIO Group

"Switch-based replication and virtualization have been over-hyped, but there are people who are working on it and over time it will become more common," said Greg Schulz, founder and senior analyst at Stillwater, Minn.-based StorageIO Group.

LSI Corp.'s StoreAge Storage Virtualization Manager (SVM) straddles the line between inline appliances and fabric-based products that depend on expensive intelligent switches. The combination of SVM and LSI's Data Path Module, which plugs into existing Fibre Channel switches to perform switch-based forwarding and eliminates the need for intelligent switches, combines the simplicity of IBM SVC with the performance and scalability benefits of a split-path architecture. HP seems to concur, and is offering

the LSI product as HP StorageWorks SAN Virtualization Services Platform (SVSP) to complement its host- and array-based replication offerings with a network-based replication and virtualization product.

Even though the market share for array-, host- and network-based replication will shift over time, there will be appropriate places for all three approaches. While each has its own set of advantages and shortcomings, specific environments and situations will best determine where replication should occur.

Jacob Gsoedl is a freelance writer and a corporate director for business systems. He can be reached at jgsoedl@yahoo.com.

FAQ: DR replication

W. Curtis Preston answers some of the most important questions about replication in the disaster recovery space, including how replication works, as well as the pros and cons of synchronous and asynchronous replication.

By W. Curtis Preston

WITH THE INCREASING availability of data replication products, many are turning to this technology as part of a disaster recovery (DR) strategy to address tighter recovery time objectives (RTOs). W. Curtis Preston, an executive editor in TechTarget's Storage Media Group, as well as an independent backup expert, discusses replication in the DR space in this FAQ.

What is replication and how does it work?

Replication is simply the concept of copying data as it's changing over to another site. So as you change a file that

changes a few blocks, then the idea is that those blocks get sent incrementally to another system that then overwrites the same blocks on the target system. So you continually have an updated copy of what you have on one site living on another site.

How are organizations leveraging replication for DR?

For those people who have tried to use tape and traditional methods of backup and have found that they're simply not able to meet their RTOs, those companies have gone to replication for DR, especially for their critical applications. Basically, they're using synchronous or asynchronous replication so that they constantly have a copy of their data offsite. This way, they don't have to do a restore when they're in a DR situation. They simply have to switch whichever copy of the data they're using.

Can you discuss the pros and cons of synchronous replication vs. asynchronous replication?

It's not so much pros and cons as it is that each one is appropriate for a different application. Synchronous replication is meant to have the updated copy continuously, completely up to date. It doesn't give the acknowledgement of the write to the application until that write has been copied over to the target site. So that's going to ensure that the target system is always up to date with the source system.

The disadvantage is that it can create latency in the app because if the round trip from the source system to the target system takes a long time, you're going to be slowing down your primary application. That's what asynchronous replication is for; it goes ahead and tells the application that the bite has been stored and then it asynchronously is copying that over to another site.

The challenge with that is [depending on factors such as how busy your application is vs. how well your connection is] that you can be anywhere from seconds to hours behind the source machine. That will result in a loss of data. It's a tradeoff of how much data you want to lose vs. how much you want to slow down your primary application.

Generally, people use synchronous only when they're going local or almost local, where something is only a couple of miles away. As soon as we start talking about a double-digit number of miles, people generally switch to asynchronous replication because it's simply too much of a hit on their primary application.

Is asynchronous replication a legit alternative for DR? Can you outline some strategies for overcoming its limitations?

For true DR, you have to be outside of the blast radius. So if this is Florida and you're liable to get hit anywhere in Florida with a hurricane, then your DR sites need to be outside of Florida. If this is California and you want to be outside of the wildfire range, then you have to be outside of California.

As soon as you start doing that, you have to go to asynchronous replication. The challenge again is "How far behind do you want it to go?" There are a lot of technologies to help make long-term trips faster to minimize things like the number of round trips. You don't want your asynchronous replication to work exactly like it was local because when you believe the storage is right next to you, you do things like ask eight questions when you could ask all eight questions as one transmission.

I would say that it's absolutely an alternative. What's the other alternative: Go to tapes and FedEx them? That certainly isn't going to be as good as asynchronous replication. I would add to that CDP as well, which is basically replication with a back button, so it has the ability to go back in time.

Can you discuss any management best practices for ensuring effective replication?

At a minimum, this stuff needs to be monitored. Probably the biggest problem most people have is that they set up that initial replication and then they're not doing regular monitoring. I would say that it starts with having consistent replication storage.

There are lots of different ways to replicate: whether it's storage array replication, hardware in the SAN replication or software on the host replication. If at all possible, pick one or at least stay with a common brand. One of the things to do when you select

that brand is to check out what they have in terms of monitoring and reporting and letting you know when you're behind, a link has been broken or when the target system that you're replicating to isn't there anymore because somebody turned it off.

Also, to make replication affordable, a lot of people are making the target system an older system that has been moved out of the data center or a less-expensive storage system using Fibre Channel in the primary site and SATA in the secondary site. You must understand that there will be a significant difference in performance when changing the storage, because it's where a lot of the performance comes from. ⊙

W. Curtis Preston is an executive editor in TechTarget's Storage Media Group, as well as an independent backup expert.

Data replication tools

Wading through replication products is tough enough, but determining which products best fit your system is even harder. We provide readers with a list of available replication products and how they differ from each other. *By Ed Tittel*

ON THE SURFACE, data replication tools represent an efficient strategy for disaster recovery (DR). A data replication approach relies on ferrying updates, transactions and changes across a network—from a primary production server to a secondary backup server. If the primary server goes down, the secondary server can take over when a failover operation kicks in.

For firms whose recovery time objective (RTO) and recovery point objective (RPO) for regaining business-critical functionality is a matter of minutes, data replication is a truly workable strategy. However, in researching vendors for replication tools, it soon becomes apparent that there are a few catches with the best and most workable data replication tools:

- The core of any data replication system is the data replication software, along with the links and hardware necessary for primary

and secondary systems to maintain constant contact and communication. Some of the major players in this arena include Dell Inc. EqualLogic (iSCSI SAN arrays remote replication facility), Double-Take Software Inc., EMC Corp. (SRDF/A and MirrorView/A) and NetApp (SnapVault and SnapMirror).

- Where replication data volumes are significant or part of heavy wide-area network (WAN) traffic, a vital component in any Internet-based solution is some kind of WAN accelerator. Over time, the cost of bandwidth for replication becomes a major factor. Consequently, all of the major WAN accelerator players, including Blue Coat Systems Inc., Cisco Systems Inc., Juniper Networks Inc., Riverbed Technology Inc. and Silver Peak Systems Inc., have extensive partnerships with the leading replication players.
- Likewise, replication touches deeply into any enterprise's storage infrastructure, particularly where server farms, data centers or cluster/grid computing environments are involved. That's why you see storage vendors doing replication business (EMC and Dell EqualLogic), and why so many other data storage vendors (Brocade Communications Systems Inc., Emulex Corp. and Symantec Corp., as well as major system vendors like Dell, Fujitsu, Hewlett-Packard [HP] Co., IBM and Sun Microsystems Inc.) often come into this picture as well. It's possible to get into this market for as little as $20,000 to $30,000 (two WAN endpoints with accelerators, plus replication software). Organizations with more WAN endpoints and lots of servers will probably end up spending $8,000 to $10,000 per endpoint and approximately $1,000 per server for software (but can drop the endpoint costs for solutions that integrate with existing WAN accelerators).

Over time, the cost of bandwidth for replication becomes a major factor.

DATA REPLICATION AND WAN ACCELERATORS

Given that the cost of bandwidth can easily consume 30% or more for annual data replication budgets, making the most of Internet bandwidth between primary and secondary systems looms very large in the minds of DR and business continuity (BC) planners. This also explains why so many enterprises build

data replication solutions around combinations that work with their existing WAN and server infrastructures, which together dictate what kinds of replication solutions get purchased and deployed.

Sometimes, when always-on or hot recovery sites are contracted from third-party suppliers, those site operators will specify replication solutions. Because duplicating IT assets essentially means multiplying IT costs by the number of recovery sites set up, this gets expensive fast. This also explains why it's crucial to limit recurring costs for any replication solution, where bandwidth and personnel costs quickly swamp equipment and software outlays. As you examine your options for data replication, the most important considerations that emerge will most often come from a systems integration perspective. After that, you'll also want to compare features and functions, talk to reference accounts and make sure your candidates' products will work well in your production environments, often through a pilot test of some kind.

As you examine your options for data replication, the most important considerations that emerge will most often come from a systems integration perspective.

Very few companies have the luxury of crafting such products from scratch, and must work with or around existing components and technologies. That's what makes replication so interesting for those businesses and organizations that decide they need it for disaster recovery or BC reasons. It's not a difficult option to select, as much as it can be a challenging option to implement. That's why working from the WAN accelerator out to the servers and storage elements appears to be the most commonly practiced and productive approach.

Take Silver Peak Systems as an example of this approach. Its data replication products embrace a range of offerings, including Dell EqualLogic, Double-Take, EMC, NetApp and Symantec. Thus, companies with existing investments in those tools and tech-

nologies will find it easy to add (or build upon) Silver Peak WAN accelerators as part of a coherent data replication strategy. Similar investigations of offerings from Blue Coat, Cisco, Juniper Networks and Riverbed will turn up similar alliances and concomitant possibilities for implementation. That said, it's probably a good idea to carefully manage your costs for software and integration to prevent them from ballooning out of control.

Ultimately, the best data replication solution is one that integrates most easily (and inexpensively) with existing infrastructure investments, and that shows itself to support BC/DR planning recovery objectives. Because the WAN accelerator components are so important to controlling recurring costs for data replication, they'll often drive other elements that come into this picture. That's why we recommend that you start with your WAN accelerator vendor when it comes to scoping out a suitable product, or that you base your choice of WAN accelerator at least in part on the data replication options it supports. ⊙

Ed Tittel is a long-time freelance writer and trainer who specializes in topics related to networking, information security and markup languages. He writes for numerous TechTarget websites, and recently finished the 4th edition of *The CISSP Study Guide* for Sybex/Wiley.

Disk array-based data replication: PROS AND CONS

There are three main types of data replication: application based, host based and array based. Learn the differences between the three technologies, how to get started with array-based replication and find out about available products.

By Pierre Dorion

AT THE HIGHEST LEVEL, there are three main types of data replication commonly used: application-based, host-based and storage array-based data replication. We could get even more granular and further subdivide array-based replication depending on whether it takes place at the controller level, the storage-area network (SAN) level or is controlled by an appliance.

Purists could argue that SAN and appliance-based replication aren't "true" array-based replication because they're independent of the disk array, but for the purpose of this article, we can agree that replication takes place at the storage level rather than being host or application based. What distinguishes storage-level replication is that it relieves the application and server resources from the processing overhead associated with replication.

THE DOWNSIDES OF ARRAY-BASED REPLICATION

A few years ago, it was a lot easier to outline the downsides of array-based replication; it was a very low-level technology that replicated blocks of data without much ceremony. Many times you had to take applications down to preserve data integrity because the application wasn't aware of the replication process. The technology offered very little support for heterogeneous storage hardware, which made it pricey. It typically worked at the array controller level, and remote replication was costly due to network bandwidth requirements and often required some proprietary form of protocol conversion. Mirroring or RAID 1 is probably one of the earliest and best known forms of controller-based (array-based) replication. It's also known for its cost and limitations.

Mirroring or RAID 1 is probably one of the earliest and best known forms of controller-based (array-based) replication.

Not that these days are long gone, but data replication options are now much more flexible and affordable. Once reserved for enterprise-class organizations, there are now many offerings available to small- and medium-sized businesses (SMBs) without breaking the bank. This has made array-based replication a very popular option because it's centralized and operating system-agnostic. However, interoperability between most vendor solutions still remains a challenge.

That being said, if a very low-cost replication solution for a limited amount of data is what's needed, array-based replication is often still too expensive for small environments or small remote offices.

GETTING STARTED WITH DISK ARRAY-BASED DATA REPLICATION

The best way to get started with array-based replication is by first answering a number of high-level questions:

- **What are you trying to achieve?** The need to implement array-based replication may be to address slow backups, increase the frequency of backups, capture frequently changing data, reduce or eliminate tape management, etc. If host-based or application-based replication is too resource intensive, array-based replication is likely the right approach. If large volumes of data are replicated, host-based or application-based replication can also interfere with network traffic.
- **What do you need to replicate?** The answer to that question can influence the cost of the solution so it's always good to address this realistically. Because technology is deployed to support the business (in most cases), there should be a business requirement for replication benefits rather than a "nice-to-have" capability.

This takes us back to a very familiar discussion about data availability, recovery time objectives (RTOs) and recovery point objectives (RPOs), which ultimately drive the need for a particular technology. There might be a need for a certain application to have access to the latest possible copy of data. In such cases, a technology like EMC Corp.'s RecoverPoint is well suited. The I/O is split at the SAN level and written simultaneously on two different supported storage arrays, local or remote. This is done transparently and provides continuous data protection (CDP).

However, the ability to access point-in-time copies might be desirable. Data corruption or unintentional data deletion are often cited as situations where this functionality is required. Snapshot technology is a good fit to create multiple point-in-time copies without requiring an amount of disk space that's a multiple of the space the original data occupies by only replicating changed blocks. NetApp Inc.'s Snap software suite offers a comprehensive set of options, including local/remote snapshots, mirroring, vaulting and application-specific snapshot capabilities.

APPLICATION-AWARE REPLICATION

Application-aware replication is probably the most significant development in the field of array-based replication. The integration of replication with applications has made it possible to create copies of the data while the application is up and user access is maintained. Whether the product leverages snapshots, mirroring or volume-based replication, the ability to create a replica of a particular data set without affecting user access is very much aligned with today's availability requirements. This capability is also far superior to the traditional daily backup because it can support much tighter recovery point objectives by allowing multiple daily copies or even continuous protection.

LOCAL VS. REMOTE REPLICATION

One more significant feature of array-based replication is the ability to provide local or remote copies of the data. While this capability certainly isn't unique to array-based replication, it's not as intrusive or resource intensive as host- or application-based replication, which provides an opportunity to replicate more frequently.

VENDOR CHOICES IN ARRAY-BASED REPLICATION

There are many vendors with array-based replication offerings and their costs vary. Some of the best-known vendors include:

- EMC's TimeFinder, SRDF, MirrorView, SAN Copy and SnapView
- IBM's Tivoli Storage Manager for Advanced Copy Services
- Hitachi Data Systems Universal Replicator, TrueCopy and ShadowImage
- Hewlett-Packard (HP) Co.'s StorageWorks suite of software

Some vendors with specialized offerings have smaller market shares but offer innovative data-protection technologies with an array-based replication component. For example, Data Domain Inc. (an EMC company), leverages data reduction through data deduplication as a backup data storage target enhanced with array-based replication for disaster recovery (DR) purposes.

That said, with the exception of EMC's SAN Copy, which replicates at the logical unit number (LUN) level, and Hitachi Data Systems' Universal Replicator, which leverages TagmaStore, heterogeneous support remains the main challenge for array-based replication.

The line between host-based or application-based replication and array-based replication isn't as clear as it once was. The emergence of virtualization appliances is becoming the answer to interoperability between heterogeneous platforms while remaining off the host.

Pierre Dorion is the data center practice director and a senior consultant with Long View Systems Inc. in Phoenix, specializing in the areas of business continuity and disaster recovery planning services, and corporate data protection.

Channel Perspective: Benefits of thin replication

Replication coupled with thin provisioning is making its way into mainstream storage products. Learn how thin provisioning can enhance your replication product.

By George Crump

USUALLY, when a customer asks a reseller to provide a product for data replication to a disaster recovery (DR) site, the reseller thinks, "Well, that's good news and bad news." The good news is that the customer will need extra storage. The bad news is that the extra storage may break the budget, and the customer may have to put the purchasing process on hold to find additional funds.

Beyond potential budget problems, there's the technical difficulty. Setting up an efficient and easy-to-use replication topology is tedious and time-consuming. Managing and implementing replication via products such as EMC Corp.'s SRDF and Hitachi Data Systems' TrueCopy and Universal Replicator is a very complex undertaking; it's also costly, and difficult to adopt without some sort of professional services engagement.

Despite both of these problems, a replicated copy of data for DR is no longer a "nice to have" for most data centers, it's a requirement. So how can a solution provider address the need for replication, have a satisfied customer and live to tell the tale?

Until recently, there was no easy answer. But next-generation replication techniques that will make life easier for resellers and their customers are on the way. Companies like Compellent Technologies Inc. and 3PAR Inc. are leveraging thin provisioning capabilities into their replication offerings. Of these two companies, Compellent has gone the furthest by providing advanced quality-of-service (QoS) and bandwidth shaping. Armed with these "thin replication" products, you'll no longer have to spend day upon day at your customer's site configuring the replication topologies with complex Fiber Channel over IP (FCIP) interconnection schemes, full disk-based copies, little bandwidth control and cumbersome management interfaces.

WHAT MAKES THIN REPLICATION THIN?

Traditional storage-area network (SAN) replication involves a significant planning cycle to determine which protocol to use, acquire the needed conversion hardware and decide how much bandwidth is required. And the actual implementation is complicated because those scripts have to be written, tested and fine-tuned. Finally, ongoing management is a challenge, as no data center is static; changes to the application infrastructure will mean changes to the replication environment and a call to the storage reseller for support.

One of the benefits of a thin replication product is that it provides the ability during the initial replication phase to replicate only the data that has been written to a given area rather than the entire allocated amount of capacity. Most other techniques require a block-for-block copy. For continuous replication, only the changed data sets are replicated to the alternate locations. This gives your customer the ability to conserve costly bandwidth segments that would normally be overtaken by replicating needless amounts of allocated but unused capacity (white space). The target systems can be deployed with a diminished capacity, saving customers money, and allowing you to sell only the bandwidth and storage your customers need at the time they request it.

Thin replication reduces the replication process to just a few mouse clicks. It's no longer necessary to undertake large scripting efforts to automate replication; instead, the process is driven through simple and easy-to-create replication templates. You can install the product without needing office space while you make it work.

QUALITY OF SERVICE

Most replication techniques don't let you scale bandwidth prior to installation, though it can be accomplished via integration or added soft-

ware functions. Another benefit of thin replication is the variety of unique tools that allow you to model the customer's bandwidth and provide a QoS level—capabilities missing from other replication techniques. Being able to look at your customer's source data and model the required replication bandwidth up front, before they procure the bandwidth and/or target systems, means you can sell only what your customer needs for their DR replication.

In addition to enabling users to estimate the amount of bandwidth a data center will need in the future, thin replication provides the ability to throttle bandwidth usage throughout the day and monitor that bandwidth usage over time. This gives your customer the flexibility to replicate data without burdening business operations during peak hours.

Even if your customers don't start with replication on the initial installation of these products, these tools let them continuously model their data sets over time. When they're ready to implement replication, they can provide you with all of the necessary data to do so. In a typical scenario, you would need to use host-based tools along with array-monitoring processes to determine the upfront and ongoing change rates. The QoS offered in thin replication lets storage administrators change a number of very useful parameters, such as:

- Scaling replication bandwidth up and down during peak and off-peak hours to ensure adequate usage of the links based on customer service-level agreements (SLAs).
- Giving priority to one data set over another based on the business-critical nature of the data contained within a given volume. This allows an administrator to provide, for instance, a higher spec to an SAP volume than a file server one.

On the topic of QoS, legacy replication offerings have been slow to adapt to handling any real QoS concerns solely by the functionality contained within the replication and array areas. The task of collecting the metrics and then performing automated functions based on those metrics would be very slow and painful, based on the number of areas you would need to go to gather information (such as servers, storage, SAN, IP network, wide-area network and so on). This makes traditional replication implementations that have to meet a certain performance or utilization requirement very risky for resellers to guarantee.

THE BENEFITS OF INTERCONNECT PATHWAYS

Using simple networking technologies to provide the interconnect pathways between the target and remote systems allows for even more flexibility in the replication solution. With thin replication, there's no need for complex FCIP engines on both sections of the replication segments. The deployment is as simple as plugging into the iSCSI connection on the storage system and then into your customer's network. It's then a matter

of using already available pathways to gain access to the target systems. While FCIP is supported even by thin replication products, most customers generally default to iSCSI as the transport of choice because it's cheap and easy to use.

In more legacy-like deployments, you would need to purchase the FCIP engine and integrate it into your Ethernet and Fibre Channel networks. For a robust and redundant deployment, these devices are typically doubled to ensure physical-layer high availability in the replication link. The FCIP devices are fairly complex to configure and maintain over time.

THIN REPLICATION IMPROVES VALIDATION

Thin replication also enables storage administrators to test and verify that the DR and business continuity sites are in solid working order without affecting the production volumes at the source location. This allows a customer to perform a simple procedure to determine that the recovery time objective (RTO) and recovery point objective (RPO) can be met with no impact to the primary facility. These verification processes could be added to the daily, weekly or monthly checks to determine data integrity at the remote location, delivering a simulated DR process.

In a traditional replication environment, this functionality is almost nonexistent. Testing DR plans in real-time requires a significant time investment from a number of IT groups within your customer's organization or the purchase of specific aftermarket products that provide that functionality.

With thin replication, no matter what topology the customer selects, the replication process can be configured, managed and monitored from one management console, minimizing the number of full-time employees needed to maintain and manage the system.

Disaster recovery is no longer a "someday" project, as most medium-sized businesses require some form of near-real-time data replication. The good news is that the benefits of technologies like thin replication have significantly lessened both the cost of entry and the technical learning curve. The advanced features and functionality built into thin replication provide you and your customers with a flexible, simple and cost-effective replication product to meet many enterprise needs.

George Crump is president and founder of Storage Switzerland, an IT analyst firm focused on the storage and virtualization segments. Prior to founding Storage Switzerland, Crump was chief technology officer at one of the nation's largest storage integrators, where he was in charge of technology testing, integration and product selection.

Check out the following resources from our sponsor:

A Practical Guide to Disaster Recovery Planning: The Basics to Getting Started

Real-Time Data Protection for Microsoft Hyper-V

Next Generation Backup for Remote Offices