

Whitepaper

HERO™ BCA: Better Than Tape

Business-critical data is constantly growing and most Business owners are responsible for protecting it. But when they consider a replication and backup product, their assumption is that they must choose between high availability and disaster recovery. While that can be true, the routine and automated protection of data will not always include a tape aspect, as times have changed for the better. In recent years, the industry has learned that tape backup and replication software are not mutually exclusive. In fact, replication technologies with existing software and off the shelf hardware can and will provide a better backup.

Introduction

Today most companies understand that the only way to ensure data protection and business continuity in the face of the worst sorts of disasters - floods, tornados, earthquakes, terror attacks, massive power outages - is to establish a remote recovery site at a significant distance from their main and branch offices. As a result, every night many companies are already backing up their main and branch office systems to tape, and transporting them to a site anywhere from 50 to over 1,000 miles away. What they don't understand is just how vulnerable their data, and therefore their business, remains to these threats, even after such a huge outlay of administrative effort and cost. This paper explores the high cost, complexity and potentially dangerous shortcomings of a recovery strategy based only on traditional tape backup and demonstrates how an alternative solution - continuous data replication to a remote recovery site over existing WAN connections - provides exponentially better remote disaster protection without adding significant cost or complexity. Finally, it introduces new data acceleration technologies that can optimize the performance of remote recovery solutions.

What's Wrong With Tape Backup?

The problems with tape backup are well known; companies have been dealing with hardware and backup software are expensive, as is the labor required to set up and maintain them. Tape cartridges are a continuing cost and completing daily tape backups requires heavy administrative intervention. Even once the equipment is in place, making backups is inconvenient – bordering on impractical. Tape backup can involve downtime, known as backup windows, since the system being backed up cannot be used during the process. Given the ever-increasing demand for around the clock data access, it gets harder and harder for companies to complete nightly backups within the time window provided. In many cases, it is so hard that the once-nightly backup goal often slips to every other night for many machines. Foregone backups are even a common problem in remote branch offices where backing up is left to non-IT staff. Most companies don't understand how vulnerable their data and business remain to disaster – even after they've made a huge up-front and ongoing investment in tape-based disaster recovery. An article in Search Security reports that in a survey of 500 IT departments, as many as 20% of routine nightly backups fail to capture all data. Among participants of another survey cited in this article, 40% of IT managers were unable to recover data from a tape when they needed it. This is a significant concern for corporations that are regulated by industry or government requirements as they can face the risk of being out of compliance if they cannot produce required data when they need it. Tape backup also places limits on your recovery point objective (RPO), the point in time to which you can recover your systems should disaster strike. Periodic tape backup guarantees hours of lost data in the event of a disaster. Suppose, for example, that a critical system fails anytime today; the best you can do is recover to yesterday's data, which will be at least twelve hours old. The later in the day disaster strikes, the older the data from which you'll recover. In addition, recovering from a disaster, any data not backed up is lost for good – unless you recreate it. The cost of permanently lost data is high and includes the cost of the revenue that the data represents, the business value you can extract from it, and the cost to recreate it.

Consider:

- How much money would your company lose if you lost all your transaction data for the last twelve hours, or even the last ten minutes?
- What is the value of the knowledge contained in your company's last twelve hours worth of e-mails and e-mail attachments? What would it cost to have your engineers recreate the last twelve hours worth of original or edited CAD/CAM drawings?
- What's your exposure if you can't produce this data in compliance with Sarbanes-Oxley, HIPAA, SEC and other regulations?

In *The Cost of Lost Data*, a Pepperdine University report updated in 2003 – before the advent of Sarbanes-Oxley – Dr. David Smith estimates the average cost of irrecoverably lost data at more than \$10,000 per megabyte lost. But if the data lost is business transaction data or data that's especially expensive to reproduce and key to your company's regulatory compliance, your costs could be much, much higher.

Cost of downtime

When a large-scale disaster strikes, with tape backup you're out of business until you can restore your systems and your data from your tapes. This kind of restoration takes a minimum of several hours, and can easily take days or even weeks. Gartner Group estimates that the average cost of network downtime for larger corporations is \$42,000 per hour; Contingency Planning Research pegs the average hourly downtime costs for small businesses at roughly \$18,000. But the cost of downtime can be significantly higher depending on the business. In fact, it can be in the hundreds of thousands per hour for health care, consumer products and banking businesses, and in the millions per hour for brokerage, energy, and manufacturing and telecommunications companies. The key to a successful disaster recovery plan is to focus not just on the data (RPO) but also on the applications that end users run to gain access to that data. Recovery Time Objective (RTO) is generally defined as the amount of time it takes to regain access to business-critical data. Solutions like tape backup, which have an RTO of hours or days, don't provide the level of recoverability that most companies require.

Better Backups with Replication

Data replication has long been considered an impractical solution to the data protection problem. Historically, it required expensive hardware and large investments in bandwidth to protect data in real-time. The evolution of software-based, asynchronous replication has dispelled this long-held belief that continuous data replication isn't feasible - especially for small or medium-sized business with limited resources. And this new breed of data replication offers benefits that more traditional solutions such as tape-based periodic backup cannot:

- Data replication provides a continuously updated copy of critical data at a remote site which minimizes data loss should a recovery be necessary.
- Disk-based recovery is more reliable, less complex and takes less time, improving the RTO of the disaster recovery solution.

Even within the realm of software-based data replication, there are opposing approaches: synchronous and asynchronous replication. It's important to understand the benefits and drawbacks of each. In synchronous replication, the replication software intercepts data being written to disk and sends it to both the primary and secondary disk arrays at the same time. Only when both arrays confirm receipt of the data does the software accept another write. Asynchronous replication can deliver recovery point objectives (RPO's) measured in minutes, and recovery times measured in seconds. With synchronous replication, data loss approaches zero because both the primary and secondary disk arrays must contain the same data. But the confirmations required for each data write can cause performance problems, especially in applications that process lots of transactions. Acceptable performance often requires connecting the arrays with high-bandwidth fibre channel, which is very expensive and which has an effective range of about ten miles. As a result, synchronous replication is not ideal for remote disaster recovery, and is most often used to create a local backup of data in situations where having an exact copy of the data is essential. In asynchronous replication, the replication software grabs data once it is written to disk, and rewrites it to a second array. In

asynchronous replication, the application doesn't have to wait for any confirmations and can continue to operate. As a result, it has little or no impact on application performance, and can work effectively and economically over low bandwidth connections and long distances. While it can't deliver the zero data loss available through synchronous replication, it can be configured to deliver RPO's measured in minutes, and recovery times measured in seconds, both of which are more than acceptable for most businesses. This combination of excellent data protection, minimal performance impact, long-distance effectiveness and low-cost deployment makes asynchronous replication an ideal solution for backing up data to a remote recovery site.

What to Look For in an Asynchronous Replication Solution

The asynchronous replication solution that makes the most sense for remote recovery implementation is the one that lets you implement the highest degree of data protection while making the most cost-effective use of your existing infrastructure. Specifically, you want a solution that works as-is with your existing applications and infrastructure, that poses no distance limitations, and that makes the most economical use of your existing bandwidth, enabling you to maximize data protection while minimizing the performance hit on your network overall. One solution that clearly meets these requirements is HERO from HEROware (powered by Double-Take Software). HERO combines Server hardware with patented asynchronous replication and failover technologies; it captures and replicates changes, as they happen, to a secondary Server at any location, and then lets you recover from that location in seconds in the event of disaster.

Several HERO features combine to enable the highest level of data protection while maximizing your existing application and infrastructure investments:

- **Incremental, byte-level replication.** HERO monitors all files and replicates only the bytes that change, as they change, which reduces replication traffic on your network to an absolute minimum.
- **Unlimited distance replication over standard IP networks.** With HEROguard™ you can replicate to the HEROware Data Center to minimize your vulnerability to natural or man-made disasters. And it replicates over any existing IP LAN,WAN, VPN or NAT.

The numbers speak for themselves


Not only is relying on traditional tape backup methods costly and complex, it can negatively impact your ability to continue doing business after a disaster and can cause a company to incur additional expenses related to recreating critical lost data and employee productivity. While tape backup is the most common and cost effective method for protecting and recovering large amounts of data, it may be woefully unable to meet your established recovery goals. Other solutions such as high availability, disk-based snapshots and data replication must become part of a company's overall data protection solution to be successful in meeting its RTO and RPO goals.

What will have to change for me to implement?

Implementing HERO does not require changes to the existing user environment. Permissions to production files will also be applied to the replicated copy on the HERO (BCA) machine. HERO does use the existing infrastructure, requiring only native IP connectivity between sources and targets. Specifically, two defined TCP/UDP ports are used for all HERO traffic; thus allowing network management and monitoring, as well as "quality of service" or packet-prioritization to be optionally used.

Beyond Protecting the Data – Recovering the Server

The complexity of traditional recovery solutions compounds an already difficult situation, and heightens the opportunity for human error. Speed and quality of recovery are extremely important when customers and employees are relying on access to critical data, but the average restoration takes hours at best. And with solutions like tape backup, even a successful recovery can result in the loss of any data that is new or has changed since the backup was made. The HERO (BCA) Server is a whole-server data protection solution that is a full mirror image of your entire production server - its operating system, applications and data will be protected and easily and quickly recovered to your HERO system.



Whole dataset recovery

For the scenario where a data volume or disk set have been damaged and need to be restored, the HERO mirroring and replication processes can be put "in reverse" - pushing the data from the target to the source. Simply repair and then replace the storage on the production source server. The HERO database is aware of where the various target data files came from. Then the Restoration Manager can be used to select a set of files and then use the HERO engine to put the files back where they came from. The difference between using the replicated files for restoration and last night's tape is the currency of the restore. A copy of the files will be seconds away from what the production source had at the moment of failure. Last night's tape would have lost all the files that had been changed during the entire business day.

Individual File Recovery

For the scenario where the source server has simply lost a few files, there are two options.

1. HERO can be configured to "burst" the changes (instead of real-time replication). The result is a copy of the files on the target, which can be minutes to hours behind the source server. This allows a redundant copy from which to quickly restore.
2. Disk snapshots can be configured to protect the files on the target server even while the production source files are in use. With this approach, you can go to a snapshot from this morning - and recover the file all without impacting the production users. Restoring the errant file directly to the source server via snapshot UI or backup console will provide the recovered file to the users. It will also be immediately replicated back to the target, to provide consistency for all copies.

About HEROware® Inc.

HEROware® Inc. provides the world's most relied upon solution for accessible and affordable data protection for Microsoft® Small Business Server (SBS). The HERO product is the standard in data replication, enabling customers to protect business-critical data that resides throughout their Microsoft Small Business server environment. With its partner programs and service offerings, HERO delivers unparalleled data protection, centralized back-up, high availability, and recoverability. It's the solution of choice for finance, legal services, retail, manufacturing, government, education and healthcare markets. HERO is an integral part of their disaster recovery, business continuity and overall storage strategies. For more information, please visit www.heroware.com or call 866-810-HERO(4376).

© HEROware, Inc. All rights reserved.

HEROware, HERO, HERO BCA AND HEROguard are Trademarks of HEROware, Inc. Double-Take, is a registered trademark of Double-Take Software, Inc. Microsoft, Microsoft Windows Small Business Server, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. All other trademarks are the property of their respective companies.

