Managing the information that drives the enterprise

STORAGE

ESSENTIAL GUIDE TO

Developing and Refining a DR Plan

Preparing a disaster recovery (DR) plan can be a complex and resource-consuming project. Use these tips to simplify the process and put together a predictable and comprehensive DR strategy.

INSIDE

TechTarget
The IT Media ROI Experts

Be a DR doctor

Give your DR plan a periodic checkup to see if it's still in good shape.

YOU NEED TO TREAT your disaster recovery (DR) plan as if it was a living thing—your life might depend on it. Jokes aside, just having a plan isn't enough. You need to have a strategy for keeping your DR docs, expertise and staffing up to date. It's no small effort, but some hard work and cold cash now can help you avoid some pretty serious stuff later.

Testing is the foundation of any DR plan health regimen. The more frequently and completely you test, the more likely you'll be able to reveal those hidden gotchas that could cause a recovery to stumble. But re-examining some of the basic premises of your plan is good practice, too. For instance, the value of certain applications and data may change over time, making the application that your initial business impact analysis classified as "critical" not quite as important now. You should revisit your analysis results on a regular basis and see if they still stand up in the context of current business operations.

Testing is the foundation of any DR plan health regimen.

Some other crucial parts of your plan should also get a periodic checkup. A big chunk of your DR budget probably goes to maintaining a recovery site, whether it's company owned or a facility provide by a service. What was an adequate setup a year ago might strain under today's recovery load. Conversely, you may have overestimated your needs back then and current testing indicates the need for a less elaborate (and cheaper) facility. When you engage a site provider, keep your contracts short or make sure there's enough wiggle room in the agreement to make adjustments—up or down—later.

You also need to be vigilant for changes in the "patient's" constitution. Adding new technologies like data deduplication might keep your production environment slim and trim, but they can also make recoveries more complex.

We've touched on a few key points here, but a DR plan needs a lot more care and feeding to stay in shape. Read on for some great expert advice and tips to help you keep your DR plan healthy. ⊙

Rich Castagna (rcastagna@storagemagazine.com) is Editorial Director of the Storage Media Group.

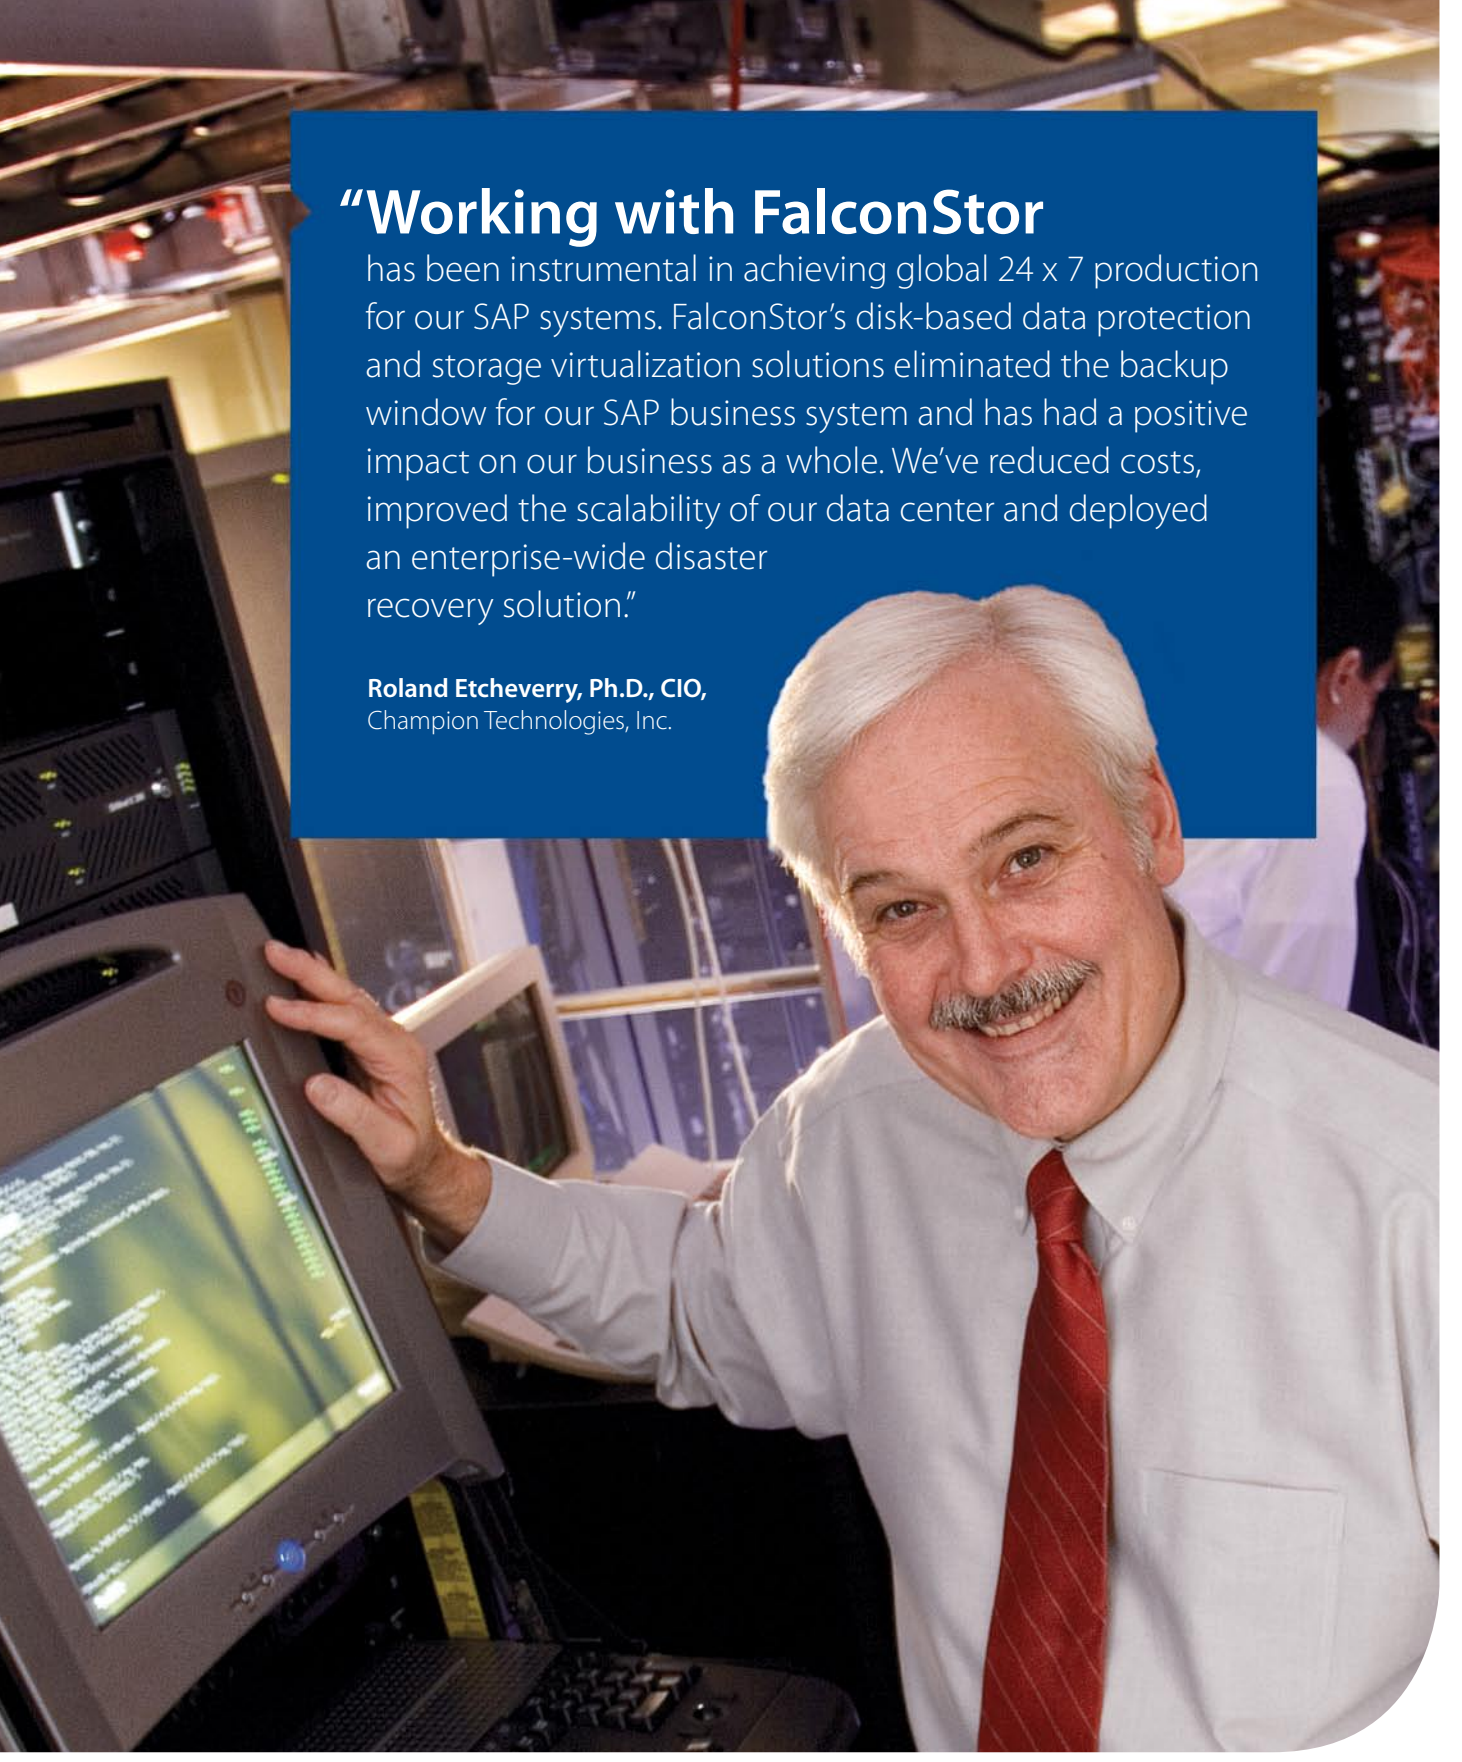

"Working with FalconStor

has been instrumental in achieving global 24 x 7 production for our SAP systems. FalconStor's disk-based data protection and storage virtualization solutions eliminated the backup window for our SAP business system and has had a positive impact on our business as a whole. We've reduced costs, improved the scalability of our data center and deployed an enterprise-wide disaster recovery solution."

Roland Etcheverry, Ph.D., CIO,
Champion Technologies, Inc.

To learn how to eliminate your backup window, reduce costs and deploy an enterprise-wide disaster recovery solution:

Contact FalconStor at 866-NOW-FALC (866-669-3252)
or visit www.falconstor.com

FalconStor Software

Foolproof DR is still a moving target

Keep these eight key steps in mind when designing and testing your disaster recovery strategy.

STEPS 8

IN A RECENT CONVERSATION I had regarding disaster recovery (DR), a CIO remarked that he'd like to achieve what he called "provable" DR. I heard this as evidence of a positive development I've noticed in recent years: Many companies have finally become serious about DR. There are several reasons for this, including the heightened awareness of data protection issues among the general public.

Technology developments have also played a critical role, with more data protection options available than ever before. The changing nature of business applications means that expectations regarding performance and availability have also been raised.

But achieving DR "provability"—or at least greater predictability—remains a challenge. Fundamentally, DR is a holistic endeavor with a number of moving parts. It's fairly easy to deal with one component of DR and for it to perform reasonably well. The hard part is ensuring the coordination and synchronization of the various elements so they function together. To establish more predictable DR, I've outlined the following eight necessary elements.

1 **Clearly defined organizational responsibilities.** Roles and responsibilities is a major area where organizations fall short with regard to DR. The DR process is much more than restoring or replicating data; it's about ensuring that applications and the systems they support can be returned to functional business usage. Accomplishing this requires participation from groups outside of IT, including corporate governance and oversight groups, finance and the business units impacted.

While IT may drive the planning and execution of DR, it's imperative that it's coordinated with a broader business-continuity planning effort. A solid DR strategy needs the strong endorsement of the highest level executives.

2 **Validate the business impact analysis (BIA) process.** Technically, the BIA isn't part of the DR process—it's a prerequisite that forms the foundation of DR planning. In a perfect world, the output of a BIA would define the kinds of recovery capabilities IT must design

Photograph by Kathleen Dooher

and deliver in support of the business. The real world, unfortunately, isn't so simple. Information is often incomplete, and we need to make assumptions to fill in the gaps.

However, there needs to be a level of confidence in the validity of BIA requirements. To design an effective DR strategy, IT needs two pieces of data: recovery requirements (e.g., RTO and RPO), and a reliable estimate of the real cost of downtime to the business. Recovery requirements supplied without valid financial impact data can result in a great deal of effort for a project that isn't funded. Passing requirements through the financial prism adds realism to the process.

3) **Define and tier application recovery services.** When business executives hear IT people talking DR strategy, they're thinking cost. DR represents insurance and because no one wants to buy too much insurance, efficiency is vital. While there are significant fixed costs inherent to DR—a recovery site, for example—there are also a substantial number of variable costs that can be controlled. The key is to realize that not every application requires a two-hour recovery time. Establishing a catalog of services, based on BIA requirements, that provides several levels of recovery and then aligning applications appropriately is one way to contain costs. With multilevel recovery services, applications can be prioritized according to importance. Among the business attributes that should be defined within the service catalog are risk (usually expressed in terms of RTO and RPO), quality of service (including performance and consistency levels) and cost.

4) **Implement a comprehensive cost model.** While the BIA determines the impact of downtime on a line of business, and tiered recovery services provide a catalog of services that align with business requirements, there also needs to be a method to determine and allocate the cost of those services. Corporate governance may help set thresholds for recovery and imply minimum levels of protection, but the service level is greatly influenced by cost. The cost model should calculate the per-unit total cost of ownership that would be charged to the business for any given service offering. Among the items included in such a cost model are personnel, facilities, hardware and software, maintenance and support. Having this data available helps significantly in aligning "want" with "need," and is a critical success factor in delivering these services efficiently.

5) **Design an effective DR infrastructure.** The DR infrastructure must support the BIA requirements and service-level targets. While DR is an extension of operational recovery capability, factors such as distance and bandwidth also come into play. The good news is that the number

of remote recovery options available to architects and designers has increased dramatically over the past few years. Traditional storage mirroring and replication are more broadly available on a wide range of systems, and compression and deduplication technologies can reduce bandwidth requirements. In addition, technologies like server virtualization can dramatically improve remote recoverability.

6) **Select the right target recovery site.** DR site selection often presents a challenge. Organizations with multiple data centers can develop cross-site recovery capabilities; if you don't have that option, selecting a DR site can easily become the biggest challenge in getting DR off the ground.

Key concerns include the levels of protection needed, and whether to own or outsource (and to what degree). The two chief (and often competing) factors to consider are risk and convenience. Planning for protection against a regional disaster means that many DR sites get pushed far away from headquarters, where most of the IT staff is housed. Service recovery levels will determine whether the site is a hot, warm or cold site. This is a critical designation because there's a substantial difference in the fixed cost of each. Generally, RTOs of less than a day require a hot site. The question of outsourcing depends on the desired degree of control, guarantees of infrastructure availability at a given location and, of course, cost.

7) **Establish mature operational disciplines.** One of my colleagues is fond of pointing out that one of the best ways to improve DR is to improve production. Put another way, if normal day-to-day operations don't tend to function well, your DR isn't likely to either. Therefore, operational discipline is an essential element of predictable DR. The first sign of a potential operational deficiency is the lack of documentation for key processes. Given that DR, by definition, occurs under seriously sub-optimal conditions, the need for well-documented standard operating procedures is clear. Organizations that have established and actively embrace standard frameworks, like the Information Technology Infrastructure Library (ITIL), are significantly improving their odds of recoverability in the chaotic atmosphere of a disaster situation.

8) **Develop a realistic testing methodology.** Given the operational disruption, practical difficulties and costs involved, we tend to focus our testing on those components that are easy to test. But realistic testing is just that—testing real business function recovery. While it's necessary to perform component testing on a regular basis, it's equally important to test the recoverability of large-scale functions to ensure that interoperability and interdependency issues are

addressed. The closer to a real production environment a test can get, the more "provable" the DR capability.

The elements outlined here transcend the boundaries of the IT infrastructure. It's therefore critical for IT administrators to have a strong understanding of the problems at hand and to learn how to address them so they can influence strategic decision-making wherever possible. This will help them avoid being placed in the Catch-22 situation of solving a problem over which they have no control.

Jim Damoulakis is CTO at GlassHouse Technologies, a leading independent provider of storage and infrastructure services. He can be reached at jimd@glasshouse.com.

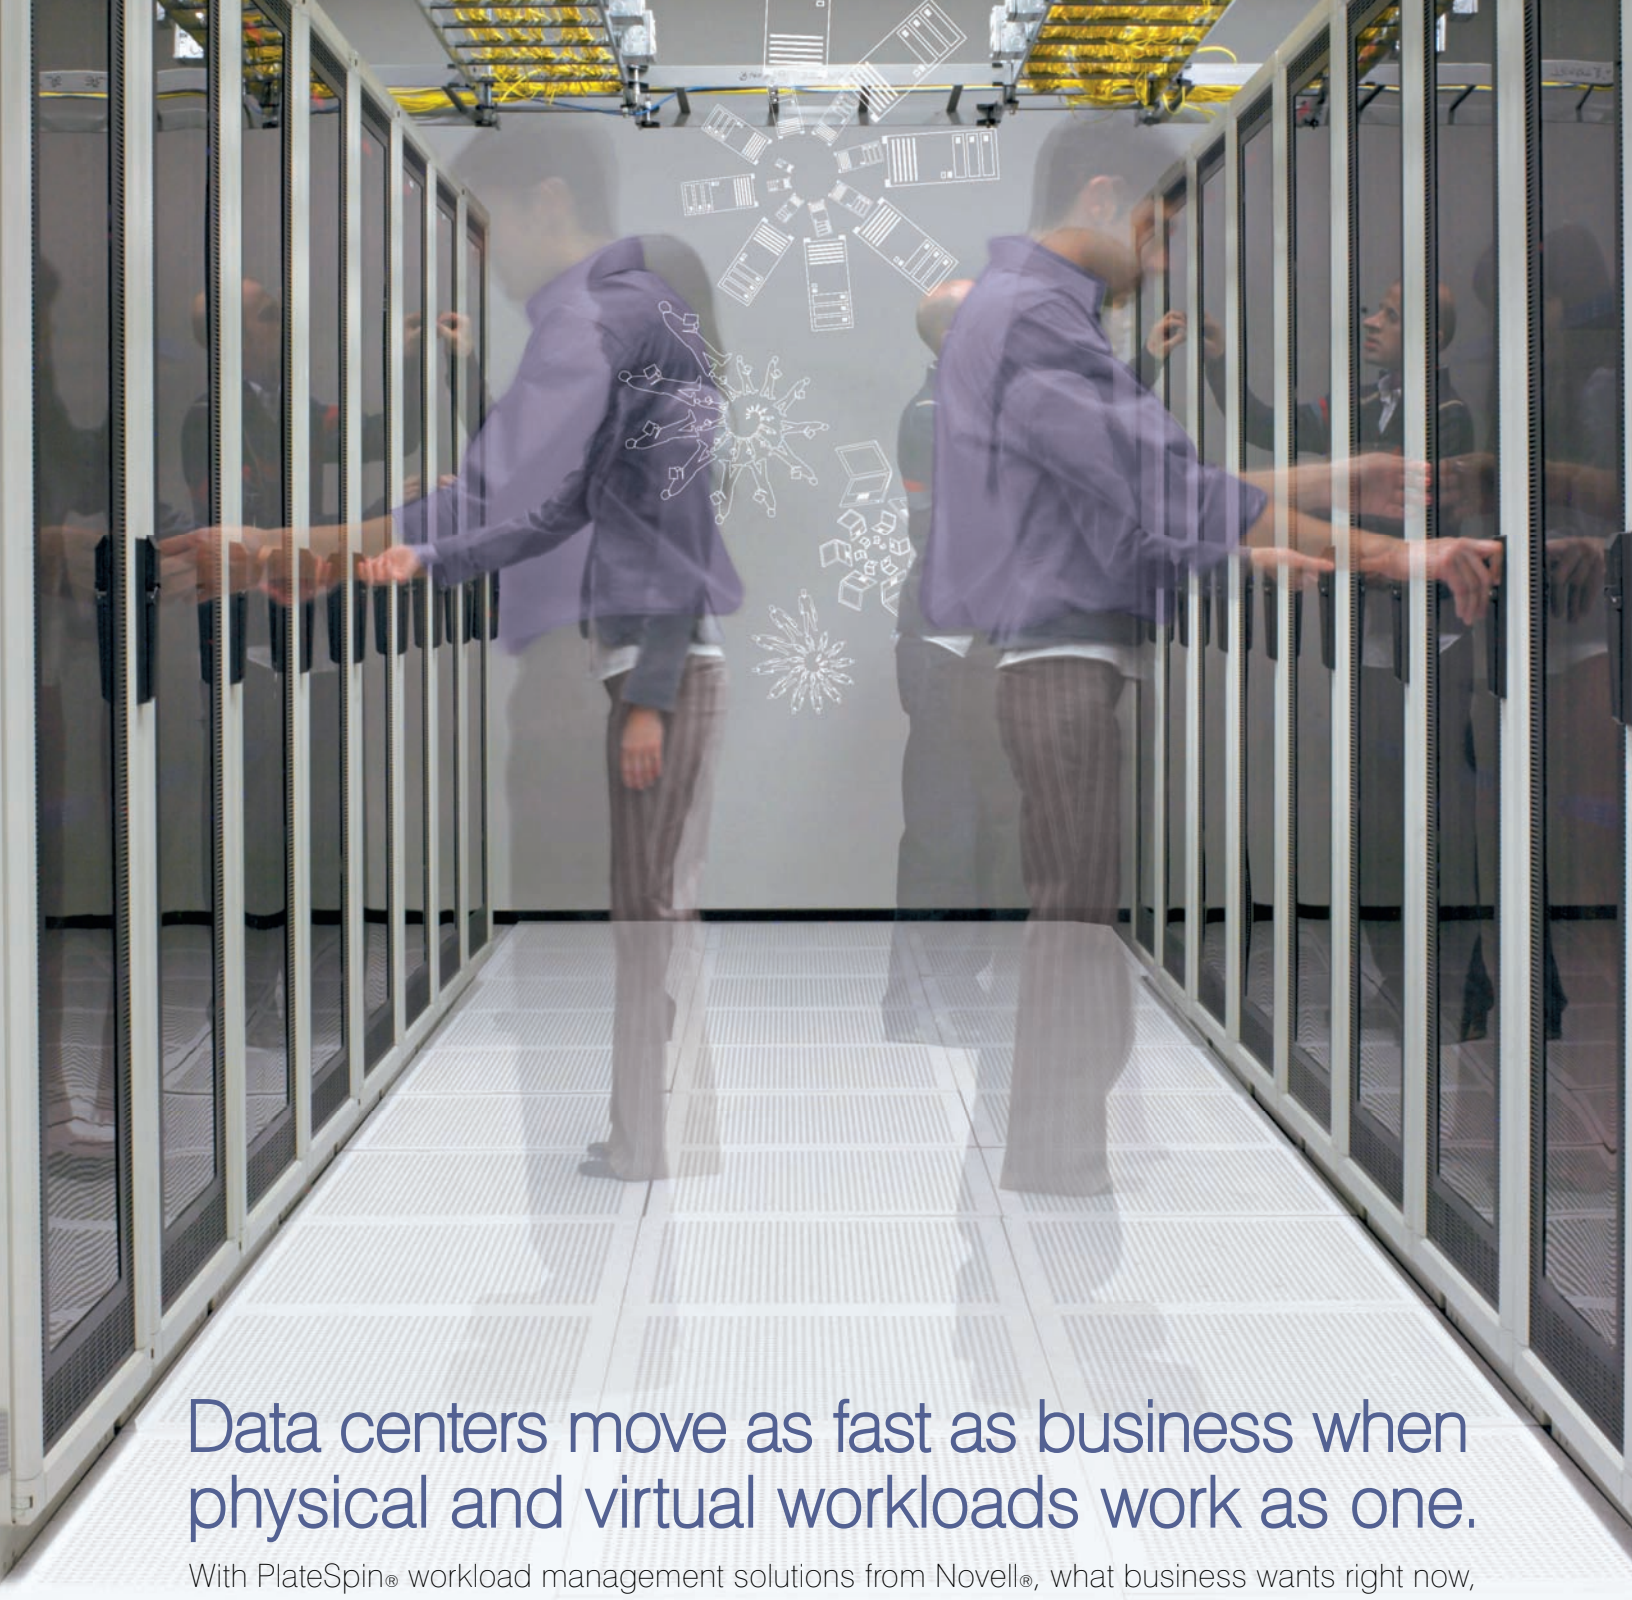

Data centers move as fast as business when physical and virtual workloads work as one.

With PlateSpin® workload management solutions from Novell®, what business wants right now, your data center can deliver right now. A single suite of products centrally monitors, manages and optimizes physical and virtual servers for you, automatically shifting workloads to the right server at the right time. Improve server utilization, reduce costs and make your data center more agile so you can respond to business demands in real time. Let us make IT work as one for you.

For more information on PlateSpin Workload Management from Novell, please visit us at www.novell.com/promo/hp/platespin.html

Novell®
Making IT Work As One™

Recovery site options

Recovery sites are typically classified using a temperature methodology of hot to cold to designate how quickly they can be up and running when the primary IT site goes down.

By Ed Tittel

WHEN PLANNING, SETTING UP and operating a disaster recovery (DR) plan, having alternate sites for IT operations is critical. DR site options include hot sites, warm sites, cold sites and mobile sites. Each has its own features, functions and costs.

hot sites

A hot site is a full or partial duplicate of a primary IT operation, including complete computer systems and near-real-time back-ups for systems, applications and data. In its most expensive form, mirroring software is used to keep a hot backup site and a primary site synchronized.

A hot site is used when an organization can tolerate little or no downtime. Switchover typically takes no more than a few hours, and may occur more quickly than that. Because staff must travel from a primary site to a hot site after a disaster, data and services

may be available to users sooner than staff. Hot sites make sense for businesses where ongoing operations are critical, or where the cost of downtime (calculated using standard risk assessment) meets or exceeds the cost of hot site acquisition, outfitting and maintenance. This includes government agencies (especially FEMA and the DoD), financial institutions, and large e-commerce and trading operations (stock or commodities markets).

warm sites

A warm site is best understood as a hot site minus data replication. A warm site offers access to space, utilities and equipment, but it requires current backups to be installed, and systems and services to be brought online to become operational. A warm site may be a complete duplicate of an original site, but it will typically provide only a subset of mission-critical equipment, services and data.

A warm site is best understood as a hot site minus data replication.

A warm site works for businesses or organizations that can tolerate one day or two days of downtime, which is the typical delay between when a primary site goes down and a recovery site comes up. Many major and medium-sized businesses opt for warm sites because the costs are significantly lower than for hot sites. Although equipment and siting costs are similar, data synchronization, and ongoing maintenance and monitoring costs don't apply to warm sites.

cold sites

A cold site is best understood as a DR plan that will be enacted when necessary. A business or organization typically makes arrangements for access to a recovery site with adequate utilities and services, and purchases necessary equipment to mount and restore essential IT operations. Aside from the cost of planning and arranging for DR, and access to sufficient funds or lines of credit necessary to cover equipment acquisition and siting costs, cold sites don't require substantial upfront outlays.

Bringing a cold site up usually takes one week to two weeks, assuming a site, equipment and backups can be acquired and activated that quickly. This can take longer, particularly after a natural disaster (think about businesses in New Orleans following Hurricane Katrina) when many organizations must compete to meet similar needs. That's why cold sites make sense primarily for small- and medium-sized businesses (SMBs) that can continue to function for some time without a fully operational IT infrastructure.

mobile sites

A mobile site sits somewhere between a warm site and a cold site, and is where a site operator makes portable structures equipped with computing equipment available to customers. The degree to which the computing infrastructure is decided and funded in advance decides the relative "temperature" of a mobile site (more advance funding and preparation is warmer; less is colder).

Bringing a mobile site online depends on how quickly it can be delivered to its desired location, and how quickly backups are restored and made operational. Because mobile sites are limited in terms of scope and scale, they make the most sense for SMBs, albeit those with less tolerance for delay in resuming IT operations.

COST FACTORS

Costs may be understood in terms of purchase and monthly expenditures for a basic IT server at a recovery site, plus square footage for housing; cooling; power; utilities; and high-speed Internet access (100 Mbps or better). Cost and capabilities scale as the number of units involved increases. Upfront acquisition and ongoing maintenance and operation costs for hot, warm and cold DR sites are the biggest differentiators (hot and warm sites must upgrade and replace recovery gear in synch with primary gear). Dollar entries compare relative costs (for every dollar spent on cold sites, plan to spend $2.50 on warm sites, and $10 on hot site equivalents). ⊙

Ed Tittel writes regularly for numerous TechTarget websites on networking, IT security and developer topics. His most recent books are *Windows 2008 Server for Dummies* (Wiley, 2008) and the *CISSP Study Guide, 4e* (Sybex, 2008).

What Happens When Email Stops?

You rely on email to communicate, to action immediate service and to collaborate across the organization.

Without email everything stops. Productivity dies, employees are isolated and information flow ends.

Keep Lines of Communication Open

In an emergency the ability to communicate is vital. There is no acceptable downtime window for email, it must be available 24x7.

Planned maintenance, storage failures, power outages and user errors are all reasons email stops. Factor these into service continuity plans. Service continuity plans should have protection of email as a high priority. Lives may depend on it.

Keeping Email Flowing

Neverfail is an award winning solution to keep users connected to email. Disaster recovery, high availability and data protection comes as standard. Out-of-the box your email is protected. Predictive monitoring ensures best practice. Replication ensures data is always protected. Automated failover keeps email flowing when things go wrong.

Can you afford to be without email for a day?

Visit www.neverfailgroup.com/resources/whitepapers.aspx for your free copy of the Osterman Research White Paper: 'Planning for Improved Email Availability'

Or, better still, email us at info@neverfailgroup.com today or call 512.327.5777 to join organizations across the World who've chosen Neverfail for the most effective disaster recovery, data protection and high availability solutions in the industry.

EXCHANGE • SQL SERVER • FILE SERVER • IIS • SHAREPOINT • BLACKBERRY • LOTUS DOMINO • RIGHTFAX

4 Dedupe strategies

Here are four strategies to consider when developing a disaster recovery plan in an environment using data deduplication.

By Pierre Dorion

DATA DEDUPLICATION HAS MADE its way into data centers around the world and is in the process of replacing tape as the media of choice for back-up data storage. This evolution must be taken into consideration when developing IT disaster recovery (DR) strategies. Here are some of the strategies worth considering.

STRATEGY 1: DISK-BASED BACKUPS

While data deduplication usually leverages disk for storage, it shouldn't be confused with data mirroring or snapshot technologies. In most cases, data is written to disk using backup software and must be written back (restored) to a host in its native format before it can be accessed again. Although data deduplication vendors remind us that disk is faster than tape, backing up to disk isn't data mirroring. In other words, if an application can tolerate little to no downtime, data deduplication isn't the best choice as a primary data protection target.

STRATEGY 2: REPLICATION IS A MUST

Unless deduplicated data is also replicated offsite, it offers only limited DR capability. Some organizations choose to implement deduplication onsite for backup data, but still use tape for offsite storage and disaster recovery. In many cases, data is no longer deduplicated once it's copied to tape. This will eventually be addressed when all backup applications are dedupe-aware or -capable. In the meantime, using tapes for offsite storage will undo the benefits of data reduction and disk-based backups, which brings recoverability back to the same level as traditional tape backups.

STRATEGY 3: NETWORK BANDWIDTH

One of the advantages of data deduplication is the ability to replicate a reduced data set to a remote location without the same network bandwidth requirements as conventional replication. However, even with this reduced bandwidth requirement, the initial replication is still likely to take a significant amount of time or bandwidth since data-reduction gains are usually not immediate and typically improve over time following multiple backups. In some cases, the first replication pass is done with the replication target installed locally to work around possible network bandwidth limitations; subsequently, the secondary data dedupe appliance is sent offsite to resume replication of deduplicated data.

Any potential bandwidth limitation must be taken into consideration when planning for large restore operations typically associated with disaster recovery. It's also important to choose a suitable DR location for the remote replication target to avoid having to relocate the storage to accommodate large restores due to a lack of bandwidth or space.

STRATEGY 4: PERFORMANCE

There are some differences worth noting in the way data deduplication products process data. These differences can have a significant impact on recovery capabilities and must be taken into consideration. Some deduplication technologies are referred to as out of band or offline, which means data is first written to disk and then processed for data deduplication before the final write. While this offers a certain performance advantage during the backup process, it creates a delay in the replication process that can affect the recovery point objective (RPO) for some data. In the event a catastrophic failure affecting the primary storage target took place before the data was replicated offsite, this situation would result in data loss, forcing a restore from the last known good copy stored offsite.

Data deduplication vastly improves backup and archive data storage. By considering external factors and selecting a solution that will meet the organization's recovery requirements, deduplication definitely has its place in a DR strategy. ⊙

Pierre Dorion is a certified business continuity professional at Mainland Information Systems Inc.

5 outsourcing questions

Here are the five questions you should ask any potential disaster recovery outsourcing provider.

By Ray Lucchesi

FOR MANY COMPANIES, disaster recovery (DR) and business continuity (BC) planning are rushed reactions to a disaster. Proactive DR and BC support can be established with relatively little time investment by employing DR outsourcing services from companies such as EDS (a Hewlett-Packard company), IBM Corp., SunGard Data Systems Inc. and others. But before you choose a DR service provider, consider the following questions.

Q What recovery time objective/recovery point objective (RTO/RPO) timeframes are offered by your service?

Most large DR outsourcers offer anything from a mirrored site with an RTO of less than one hour to time-share sites with an RTO/RPO of more than 48 hours. For mirrored sites, there are severe distance constraints and DR site storage equipment is typically dedicated to one organization. For time-share sites, tapes or other media are shipped to the DR site and outsourcers supply servers, storage and networking to use. Some outsourcers also offer shell sites with only a raised floor, and power and

cooling that can have an RTO/RPO of more than one week. For shell sites, hardware must be shipped and installed prior to restoring applications and data.

What happens when a regional disaster hits?

Time-share sites are first come, first served. Many outsourcers provide alternate DR sites to be used when a designated DR site is full. However, any telecommunications and/or networking bandwidth services previously dedicated to support designated DR site business requirements must be re-routed to the alternate site(s) at the outsourcer's expense.

DR outsourcers should actively manage regional disaster risk, especially with respect to a 25 mile-wide disaster centered around the nearest large city. Most outsourcers limit this risk by not offering more services than can be simultaneously supported, or they oversubscribe and offer enough alternate sites to cover regional disasters. Oversubscription rates also often impact DR test scheduling availability.

Are the designated and alternate DR sites high-availability data centers?

The Uptime Institute Inc., the industry leading data center certification authority, defines the four tiers of data center sites as follows:

- Tier I has basic infrastructure
- Tier II has redundant capacity components with a single, non-redundant distribution infrastructure
- Tier III has redundant capacity components and multiple independent distribution paths providing a concurrently maintainable infrastructure
- Tier IV has a fault-tolerant infrastructure

Outsourcers should supply at least Tier III DR sites as defined by the Uptime Institute, and should match or exceed the primary site availability level.

How much DR test time is supplied with your DR service?

Periodic tests or exercises are critical to any successful DR plan. Most contracts offer two days of test time per year, but more time can often be purchased. Testing is critical to ensure the viability of the overall DR plan, both to the business and the outsourcer. Some service firms also offer detailed audits on DR exercises.

What hardware and software is supplied at the DR sites?

An outsourcer must supply hardware at DR sites that's compatible to hardware available at the primary data center. For instance, where data security is used at the primary site, compatible data encryption and key management services must be available at the DR site. Additionally, the outsourcer should supply software for all common operating systems and database servers. Special applications software is typically licensed to the primary data center owner and needs to be transferred to the DR site while running or testing there.

These five questions constitute a starting point for an informed discussion with prospective outsourcers. More questions should be asked regarding costs, DR and BC planning support, DR site personnel and so on. Disasters do happen; however, a DR plan using outsourced service providers may be a viable alternative.

Ray Lucchesi is president of Silverton Consulting, a Storage, Strategy & Systems consulting services company. Based in the U.S., the firm offers products and services to the data storage community.

Check out the following resources from our sponsors:

FalconStor Software

White Paper: TOTALLY Open™ Disaster Recovery

Replicate with Integrity: Protecting VMware Data to Ensure Fast Recovery, Business Continuity

WAN-Optimized Replication: Built into FalconStor Solutions

neverfail™
WWW.NEVERFAILGROUP.COM
PREDICT · PROTECT · PERFORM

Business Continuity: Choosing the Right Technology Solution Whitepaper

Planning for Improved Email Availability Whitepaper

Enhancing Exchange Server 2007: High Availability with Neverfail Whitepaper

Novell®

A Pragmatic Approach to Server and Data Center Consolidation

Consolidated Disaster Recovery Using Virtualization—Affordable Workload Protection and Recovery