

Managing the information that drives the enterprise

# STORAGE

**ESSENTIAL GUIDE TO**

# Bulletproof Disaster Recovery Planning

*Disaster recovery planning and testing tips that will have the most impact for enterprise data storage administrators are highlighted in this guide.*

## INSIDE

- 6 Eight DR planning tips
- 11 DR site options
- 14 DR testing basics
- 17 DR for virtual machines
- 22 Channel extra: SMB DR planning

# Are you prepared?



## Learn why over 28,000 businesses trust i365 EVault Disaster Recovery Solutions

Backup and Recovery

Bare Metal Restore

Remote Disaster Recovery

Data Replication


Disaster Recovery Assessments

Remote and Branch Offices

[DOWNLOAD OUR DR WHITEPAPER](#)



[www.i365.com](http://www.i365.com) | 1.877.901.DATA | [concierge@i365.com](mailto:concierge@i365.com)



# Some help for disaster recovery planning

*No business can afford to not have a DR plan, and new tools and techniques are making it possible for more companies to put effective business continuity plans in place.*

**D**EVELOPING a disaster recovery (DR) plan used to be a lot like taking out one of those hefty Lloyd's of London insurance policies—it might have given you some piece of mind, but it sure took a big chunk out of your budget. For years it was generally accepted that a DR plan was something that only well-heeled large enterprises could consider.

A decade of unexpected disasters—both manmade and acts of nature—have convinced most companies that having a set of procedures in place to help keep a business running in the face of adversity is less a luxury than good business sense. Fewer companies today rely on an up-to-date set of backups and crossed fingers to weather what could be business-crippling storms.

With greater awareness, companies of all sizes have realized how critical it is to have a DR plan in place, and many have given top priority to developing a plan. Still, disaster preparedness isn't easy or free, and if you're lucky you won't ever really know just how effective your plan is.

Most IT pros are familiar with the general steps involved in the DR planning process: Potential risks must be identified, the key corporate systems that are vital to your company's continuing operations must be identified, and the hardware and software infrastructure must be put in place to complete the safety net. But with each step along the way, there are a myriad of details that will arise along with easy to overlook interdependencies upon which success or failure may hinge.

But storage managers should be encouraged by recent technological developments that can make DR easier to plan and put in place. Replication is often a key element of a DR plan, ensuring that key data is safely tucked away at a distant site and ready to be recovered. In the past, implementing replication often meant duplicating primary site storage systems at recovery sites—an expense beyond the means of many companies. Today, there are many replication alternatives that

Disaster preparedness isn't easy or free, and if you're lucky you won't ever really know just how effective your plan is.

obviate the need for mirrored configurations, and they cost far less than previous alternatives. And virtualized servers can also reduce the reliance on duplicate hardware resources and make recovery an easier, more agile process.

Despite the new technologies that can take some of the sting out of DR configurations, some aspects of DR planning haven't changed all that much. Although there are some tools available that can assist, testing still remains a largely manual process—and is the only way to confirm that all the planning, systems implementation and recovery drills will actually work.

For this guide, we asked top DR experts to address some of the key issues related to DR planning, including incorporating virtual server technology, recovery site options and, of course, testing. Whether you're just embarking on DR planning or refining an existing plan, we think you'll find their advice useful. ☺

---

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



# Data Protection In Any Environment, On Any Platform, Across Any Distance.

**DR Planning and Best Practices:  
Everything you NEED to know -  
in 30 minutes**

Take 30 minutes out of your busy schedule to attend this live webinar and make sure you put the right DR and Business Continuity plan in place.


Connect with us:





## DISASTER RECOVERY STRATEGIES:

# Eight tips for better DR planning



*Successful DR doesn't happen by accident: Here's how to improve your chances.*

*By James Damoulakis*

**IN A CONVERSATION** I had once regarding disaster recovery (DR) planning, a CIO remarked that he'd like to achieve what he called "provable" disaster recovery. But achieving disaster recovery "provability," or at least greater predictability, remains a challenge. Fundamentally, disaster recovery has a number of moving parts. It's fairly easy to deal with one component of disaster recovery and for it to perform reasonably well. The hard part is coordinating and synchronizing the various elements so they function together. The following eight tips will help you establish more reliable disaster recovery:



**1. Clearly define organizational responsibilities.** Roles and responsibilities is a major area where organizations fall short with regard to disaster recovery. The DR process consists of much more than restoring or replicating data; it's also about ensuring that the applications and 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 groups and the business units impacted.



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



**3. Define and tier application recovery services.** When business executives hear IT people talking about disaster recovery strategy, they're thinking cost. With DR comes insurance, and because no one wants to spend too much on 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 business impact analysis requirements that provide 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 recovery time objective [RTO] and recovery point objective [RPO]), quality of service (including performance and consistency levels) and cost.

With multilevel recovery services, applications can be prioritized according to importance.



**4. Implement a comprehensive cost model.** While the business impact analysis determines the impact of downtime to 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 disaster recovery infrastructure.

The disaster recovery infrastructure must support the business impact analysis requirements and service-level targets. While disaster recovery 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.

The disaster recovery infrastructure must support the business impact analysis requirements and service-level targets.



### 6. Select the right target recovery site.

Disaster recovery 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 disaster recovery off the ground. Key concerns include the levels of protection needed, and whether to own or outsource disaster recovery (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 is a substantial difference in the fixed cost of each recovery site. 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.** Some people point out that one of the best ways to improve disaster recovery is to improve production. In other words, if normal day-to-day operations don't tend to function well, neither will your disaster recovery plan.

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 disaster recovery, by definition, occurs under seriously sub-optimal conditions, the need for well-documented standard operating procedures is clear. Organizations that have established and actively embraced standard frameworks, like the Information Technology Infrastructure Library (ITIL), are significantly improving their odds of recoverability in the chaotic atmosphere of a disaster situation.

The first sign of a potential operational deficiency is the lack of documentation for key processes.



**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 involves 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 consistently 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. Therefore, it's 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 disaster recovery decision-making wherever possible. This will help them avoid being placed in a situation where they must solve a problem they cannot control. ☉

James Damoulakis is CTO of GlassHouse Technologies, an independent storage services firm with offices across the United States and in the UK.



# Disaster Recovery Planning doesn't have to **Hurt**

There shouldn't be pain associated with your Business Continuity/Disaster Recovery plan.

Quest can help you consider the possibilities and plan for the unexpected. Through our time-tested processes and procedures, we can guide you in implementing an ongoing, sensible and cost-effective BCP/DR plan. We will show you how you can have your business back online in a matter of minutes, hours or days—based on your business needs, not a cookie-cutter approach.

## Get Started Now

Want a really pleasant surprise? Contact Quest and discover painless Business Continuity Planning and Disaster Recovery. Protect the future of your business. Call Quest.


## DISASTER RECOVERY SITE OPTIONS:

# Hot, warm and cold sites

*Depending on your needs, you can choose a hot, cold or warm DR site.*

*By Jacob Gsoedl*

**DISASTER RECOVERY (DR) TERMINOLOGY** can be confusing. Terms like hot site, warm site and cold site are common in DR parlance. Each option is a reliable disaster recovery site, but which one should you choose for your company? Here's a look at the differences between hot, warm and cold sites in disaster recovery and the pros and cons of each.



### **HOT SITES**

If the acceptable recovery time objective (RTO) for your company is a few hours instead of minutes, then a hot site is likely appropriate. The biggest difference between a hosted site and a hot site is the use of shared equipment for infrastructure components like servers and peripherals. Storage is dedicated and real-time data replication is used to get data from the production site to the disaster recovery site.

Because equipment in the DR site is shared by multiple customers, hot sites are significantly less expensive than hosted sites. "Hot sites and warm sites can be implemented less expensively through outsourcing than doing them in-house because of shared equipment,"

said George Ferguson, worldwide service segment manager for Hewlett-Packard (HP) Co.'s business continuity and recovery services. "DR service providers rely on the fact that not all customers have a disaster at the same time."

On the downside, the use of shared equipment makes hot sites less flexible because customers are limited by the equipment the disaster recovery service provider offers. While some service providers may have a limited selection of equipment, others are more flexible. "About 90% of the time we're able to use shared equipment, and the rest of the time we work with the customer to make it work," said Marc Langer, president at Recovery Point Systems Inc., a provider of backup, storage and disaster recovery services. Larger service providers may be less flexible, so the nature of the shared equipment is likely to be a determining factor when selecting a hot or warm site provider.

Another consequence of using a site with shared equipment is the time limit on how long customers can use the shared gear in the event of a disaster. The limit varies among service providers, but typically ranges between 30 days and 90 days. "Customers can use the shared equipment for 60 days before they need to get out or before they get migrated to a cold site," said Langer. Service providers with a larger number of data centers, like IBM Corp., can be more flexible. "We're pretty open-ended because we can shift workloads to other data centers," said John Sing, senior consultant, business continuity strategy and planning at IBM's Systems and Technology Group. To avoid unpleasant surprises, a clear understanding of the terms, conditions and limitations of managed disaster recovery services is required prior to committing to an agreement that may span several years.



### WARM SITES

In contrast to a hot site, a warm site relies on backups for recovery. As a result, it doesn't require dedicated storage, but instead can take advantage of less expensive shared storage. In other words, all components of a warm site, including storage, are shared among multiple customers. Therefore, most of the considerations of hot sites also apply for warm sites.

In the past, there was a huge difference between hot sites and warm sites because backups were limited to tape. As a result, warm site recoveries were typically measured in days. Warm sites that rely on tape-based backups for recovery are clearly at the lower end of the DR services spectrum.

Disk-based backups have narrowed the gap between warm sites and hot sites, and almost all disaster recovery service providers now offer an electronic vaulting option, which is disk-based backup of production data over the network. RTOs and recovery point objectives (RPOs) of warm sites with electronic vaulting are typically less than a day, which is very close to

the recovery times offered by hot sites, but at a fraction of the cost. “There has been about a 10 times price difference between a replicated DR infrastructure and a shared infrastructure with electronic vaulting,” explained HP’s Ferguson. “Electronic vaulting is closing the gap between tape-based recovery and a replicated DR infrastructure, and customers need to look at it because of its price and reliability benefits.”



### GOLD SITES

A cold site is rented space with power, cooling and connectivity that’s ready to accept equipment. With recovery times of a week or more, a cold site is only an option for business processes that can be down for an extended period. Cold sites are also used to complement hot sites and warm sites in case of long-lasting disasters. “Some of our customers sign up for a cold site as contingency to migrate equipment from the shared infrastructure to the cold site in case a disaster lasts more than six weeks,” said Recovery Point Systems’ Langer.

It’s the customer’s responsibility to provide equipment for the cold site during a disaster. A disaster recovery plan that relies on a cold site must clearly define the process of procuring and delivering equipment to the cold site when a disaster strikes. It’s a risky strategy to rely on purchasing the equipment on the open market when it’s needed because it may not be possible to get the equipment to the cold site in a timely fashion. A better option is to consider subscribing to a quick-ship service available from companies like Agility Recovery Solutions. “You can rent equipment for as little as \$50 a month with an option to buy it if needed,” said Recovery Point Systems’ Langer. ☉

---

**Jacob Gsoedl is a frequent contributor to *Storage* magazine.**



# Disaster recovery planning fundamentals:

# DR testing basics

*Before setting up a DR plan, you should determine your RTO and RPO.*

*By James Damoulakis*

**ONE OF THE FUNDAMENTAL** prerequisites of successful disaster recovery (DR) planning is to understand the requirements of the business. What does the business need, and is it capable of addressing this need with regard to both capabilities and cost? The key performance metrics to support this are recovery time objective (RTO) and recovery point objective (RPO). Briefly, RTO is the maximum acceptable time to resume operations—not just to data recovery—and RPO is a measure of acceptable data loss.

The failure to understand and agree upon these metrics for critical applications, and the subsequent inability to invest in and develop capabilities to support them, is the basis for the disaster recovery gap between business and IT. Bridging this gap requires IT to meet with business and application owners to understand recovery needs so that the financial impact of outages can be quantified and then weighed against the cost of providing the necessary service level. This may require some negotiation, but without this conversation, DR success is impossible.

Building this capability goes well beyond a technology exercise. It consists of planning, identifying dependencies, developing processes and, above all, testing.

## IF YOU FAIL TO PLAN, YOU PLAN TO FAIL

A disaster recovery plan represents an organization's detailed roadmap of where to go, what to do and when to do it in the event of a disaster. It should incorporate actions that need to be performed before, during and after a disaster is declared. The more basic elements include defining the criteria under which a disaster is declared, who can declare it and how individuals are notified. In the past, hurricane-related disasters reinforced the challenge and importance of communications. A good plan should include contingencies; you can't assume your email will work, or even that cell phone service will be available.

We know that processes and procedures need to be documented, but we also know that most people hate documentation. Even the most carefully crafted disaster recovery plans will become useless without proper attention. Disaster recovery needs to be baked into the standard change management process so that whenever systems are modified, software is patched or additional storage is assigned, then the impact on the DR plan is reviewed accordingly. Likewise, when reorganizations occur, the disaster recovery plan must be revisited.

**We know that processes and procedures need to be documented, but we also know that most people hate documentation.**

It's clear that double-digit data growth rates dramatically impact the ability to recover data within targeted time constraints, but application complexity and interdependence is often an overlooked factor that has a major impact on recoverability. Today, major applications are spread across multiple servers and architectures. It's not uncommon for a mainframe application to feed other applications or subcomponents that reside on Unix or Windows platforms. Based on the traditional server-centric recovery perspective, it's possible to successfully back up or snapshot each application component but be unable to fully recover the application due to inconsistencies among the various components.

You can avoid this by first understanding the interdependencies among applications and then applying the appropriate data protection approach. The method could be the use of split mirror/replication technology featuring consistency groups that encompass the interdependent elements, or it might be continuous data protection (CDP) technology that can ensure highly granular, synchronized time-based rollback.

## NO TESTING, NO DR

Planning disaster recovery is relatively easy compared to testing the plan. Testing the DR plan is often dreaded and, unfortunately, often avoided. Yet without proper testing, one might as well not bother with the planning

because the likelihood of successful execution is small if you have not tested your plan properly.

Some fundamental considerations for testing include:

- Test application recovery, not just data recovery (think application interdependency).
- Let nonprimary individuals perform the recovery to validate procedures and documentation.
- Construct multiple disaster scenarios and employ role-playing.
- Establish a positive disaster recovery testing mindset: uncovering (and fixing) problems is a good thing.
- Track metrics to measure and chart improvement.

The most common reason given for not doing more extensive testing is cost. This will inevitably be a point of contention because DR testing is viewed as an exception to what are commonly thought of as day-to-day operations. The only way to effectively address this issue and justify the cost is by closely linking the testing process to RTO/RPO service-level objectives. This means the disaster recovery business case, particularly the financial impact of RTO/RPO, must be accurate and complete. The message should be that comprehensive testing is an essential requirement to ensuring that those metrics can actually be met and is an integral part of the disaster recovery process. 🕒

---

**James Damoulakis is CTO of GlassHouse Technologies, an independent storage services firm with offices across the United States and in the UK.**

# Developing a disaster recovery plan for virtual machines: A tutorial

*Site Recovery Manager and geoclustering are among the best options for recovering virtual machines.*

*By Ray Lucchesi*


Vendor resources

Virtual DR

DR testing

DR site options

Eight DR tips



**V**IRTUAL MACHINE (VM) disaster recovery (DR) is a multifaceted activity that fails over a VM from a primary site to a remote location. There are a few approaches to facilitating disaster recovery in a virtual machine environment. One approach is VMware Inc.'s vCenter Site Recovery Manager (SRM) software that automates virtual machine failover. Alternatively, there are geographically disbursed clustering (geoclustering) services that support automatic failover, but can also recover more than just VMs. There are also standard data protection packages available that support varying levels of VM DR. While these packages are more manual than Site Recovery Manager or geoclustering, they cost substantially less.



## VMWARE SITE RECOVERY MANAGER

Facilitating recovery with VMware Site Recovery Manager automation depends heavily on array or storage area network (SAN) replication to copy datastore data between sites. SRM software executes on a SRM server or virtual machine at both the protected and DR sites, but also requires a vCenter to run at the remote site.

Once Site Recovery Manager is executed, an administrator should:

- Establish datastore replication
- Identify replicated datastores
- Select protected virtual machines
- Remap VM hardware
- Create a data recovery plan

Re-IP networking refers to the fact that the IP addresses at the remote site can't be the same as the primary site. Some of these are associated with the application and operating system running in the virtual machine and some are associated with VMware hypervisor interfaces like the server running vCenter Server, Site Recovery Manager, etc. As the VMs are brought up at the remote site, the IP addresses must be changed in order to run.

Moreover, multiple recovery plans can be defined and administrators may select which one to use for a specific failover. Alternative recovery plans such as these provide varying failover capabilities and supply recovery options for partial failures, e.g., a single datastore or ESX host failure at the protected site.

VMware Site Recovery Manager has several benefits. It supports DR testing at the local site and an administrator may modify an already existent recovery plan to support this testing. Also, SRM can have as many or as few recovery plans as you need. It's entirely conceivable that one would have a recovery plan for a total site failure and one or more for separate infrastructure failures.

VMware High Availability (HA) provides for ESX failover, but only to the local site. SRM is only involved when you want to failover to a remote site. Not every infrastructure failure would warrant a "disaster" being invoked, which would require SRM automated failover to a remote site.

VMware SRM currently has some limitations, including no support for:

- Raw Device Mode data
- Multi-LUN datastores
- Automated failback

VMs can access Fibre Channel (FC) storage in at least two ways. The first way is through normal VM hypervisor SCSI data access, which is virtualized to a VMware-defined VM cluster file system (VMFS datastore).



The second is through Raw Device Mode, where the VM actually owns the Fibre Channel port hardware and controls that link, and likely the storage attached at the other end of the link.

Non-support for Raw Device Mode data means that failover for virtual machines that have this data are more complex and less automated. SRM will not monitor replication of this data and will not automatically promote this data to active VM accessibility on failover. All of these steps have to be done manually or via data center scripting.

Raw Device Mode is normally used by performance-intensive virtual machines. These are typically high-profile applications, but are least likely to be virtualized. However, due to their criticality, they are very likely to warrant the highest form of disaster recovery.

Whether this is a concern for system admins/data centers depends on how much of their infrastructure and servers are virtualized. As more data centers move to 100% virtual machines, this will become more of a concern.

As a side note, VMware does supply support for Raw Device Mode in a beta version of SRM. Failback can still be accomplished, but an administrator would need to reconfigure SRM to perform the failback as an SRM failover.

While failover is typically unscheduled, failback is typically a scheduled activity once you have failed over. It takes time to bring the primary site back online, repair the infrastructure and power up the data center. These time-consuming activities can be scheduled to occur, so you would also be able to schedule the failback process.

It's possible that the recovery plan for failback could be in place beforehand, but Site Recovery Manager interrogates storage replication activity to validate that a protected datastore is being replicated. So a failback process identifying protected datastores and VMs and remapping inventory steps for SRM, might have to wait until the failover actually takes place before it starts, particularly when:

- Re-establishing datastore replication
- Re-identifying protected datastores
- Re-selecting protected VMs
- Remapping site inventories
- Creating a failback recovery plan

**Non-support for Raw Device Mode data means that failover for virtual machines that have this data are more complex and less automated.**

## **GEOCLUSTERING FOR VIRTUAL MACHINE DISASTER RECOVERY**

Many geocustering products are available that provide even more sophisticated cross-site recovery. In fact, geocustering can support

automated failback and failover, multi-destination DR sites and raw device mode data, and may not require a Virtual Center.

Symantec Corp. Veritas Cluster Services (VCS) allows for physical server to VM, VM to physical server and VM to VM failover.

For instance, VCS can failover a physical server at the protected site to a VM at the remote site, or vice a versa. Such capabilities go well beyond what VMware SRM was intended to support, but depending on data center needs, may be worthy of consideration. Also, VCS executes at the ESX service console level when supporting VM failover.

Windows HPC Server 2008 is another geoclustering product, but only supports server-to-server or VM-to-VM failover. As such, HPC Server must be executing in the Windows server at both the local and remote site, and only supports Windows-to-Windows failover.

### **SAN OR ARRAY REPLICATION FOR VM DISASTER RECOVERY**

Most failover automation depends heavily on SAN or array replication, and with this in place, automating failover can be accomplished with any number of approaches.

Once datastore replication is in place, administrators can build their own scripts using native VMware or other software to semi-automate virtual machine failover. However, this custom scripting must do all the work required to reconfigure the ESX servers to run the VMs, re-IP the VMs and promote replicated datastore copies.

### **DATA PROTECTION SOFTWARE FOR VIRTUAL MACHINE DISASTER RECOVERY**

Data protection packages such as EMC Corp. NetWorker, CommVault Simpana, IBM Corp. Tivoli Storage Manager (TSM) and Symantec Backup Exec and NetBackup all support DR at varying levels. The different levels of support may consist of bare-metal restore options and/or sophisticated independent backup data replication.

Tivoli Storage Manager supports a DR manager option that can be used to automatically replicate TSM protected data to a remote site. Once TSM is recovered at the remote site, data can be restored and VMs can be reconfigured with manual operator activity or hand scripted automation.

Alternatively, other backup packages support a bare-metal restore option. Such functionality can provide a one step, restorable version of all the data required by a server or VM. Once the VM data has been restored, one would need to reconfigure the VM to run at the remote site and re-IP its networking. After this is done, the VM can be powered on and recovered from its backup.

Furthermore, any backup package can be used to recover VM file data at a remote site. Without a bare-metal restore option, it may take more steps to recover all the VM data, but once it is restored, the rest of the

disaster recovery process will be similar.

VMware disaster recovery can be supported in multiple ways. But any failover automation will depend heavily on the data replication used and the software selected, specifically:

- VMware SRM can easily automate most VM failover, but has some current limitations.
- Geocustering software provides automatic failover functionality, except for VCS, which is limited to only a single operating system.
- SAN or array replication can also be used, but requires hand customized scripting to semi-automate failover.
- Most data protection packages support DR, but require customized scripting to semi-automate failover.

VM DR does not have to consist of only one approach alone. Due to replication expenses, automated failover may be limited to only a few critical virtual machines, with the rest relegated to less automated recovery. Such a multitier DR plan can easily be supported with combinations of the above products to support fully automated recovery for critical virtual machines and manual recovery for the rest. ☺

---

**Ray Lucchesi is president of Silverton Consulting, a storage, strategy and systems consulting services company, based in the USA offering products and services to the data storage community.**

# DISASTER RECOVERY TESTING: SMB VS. ENTERPRISE

*Larger firms are typically more interested in disaster recovery tests than SMBs.*

*By Sue Troy*

**R**ESearch on disaster recovery (DR) testing among end-user IT organizations shows that a large number of those organizations are—or at least say they are—testing their DR plans on a regular basis. For example, a March 2009 Snapshot Survey on DR testing conducted by *Storage* magazine showed that 59% of 139 survey respondents said they regularly perform DR tests and of those who do test, 65% said they perform DR tests at least twice a year.

Despite the attention to DR testing at IT organizations, whether or not your company has an easy shot at disaster recovery testing revenue will most likely depend on your target customers. Solution providers catering to small businesses oftentimes say that the vast majority of their customers aren't interested in testing, while those addressing the needs of enterprise customers say there's a vigorous business around DR testing.

The disparity between those two sectors makes sense, said DR expert Jon Toigo of Toigo Partners International. "The larger the company, the more complex it is, and the more they need to get religious

about continuity,” he said. Some small companies don’t even do DR planning, said Toigo, “even though their most critical data would fit on a USB key. It’s like [using] dental floss. They know they should do it, but they don’t do it.”

Mike Croy, author of *Are We Willing to Take That Risk?* and director of business continuity solutions for Forsythe Solutions Group Inc., whose customers are mostly enterprise-level IT organizations, agreed with Toigo. “Whether you’re large or small, testing is pretty important,” said Croy. At small companies, he said, testing is a challenge because of budgetary restrictions. “Larger firms have set aside more money and more staff for testing. They may often have contracts in place with large recovery site companies that help manage recovery exercises for them. Smaller firms obviously don’t have the same type of funding, but it’s as critically important to them to run a test.”

**"Whether you're large or small, testing is pretty important."**

—MIKE CROY, author, and director of business continuity solutions, Forsythe Solutions

Croy said that the majority of Forsythe’s customers have a DR plan in place, and “the larger customers have testing in place.”

Bob Gaines, technology marketing manager for All Covered Inc., a 250-employee solution provider based in Redwood City, Calif., said that just 2% of his customers are interested in DR testing, while about 15% have some form of a DR plan. All Covered’s target market is small- to medium-sized business (SMB) customers. “DR/BC is considered a luxury at those companies,” said Gaines. “They don’t worry about disasters. They’re just trying to dodge the bullet.”

Kyle Elworthy, owner and network engineer at Network Essentials, an MSP in Charlotte, N.C., echoed Gaines’ perception of customer interest in DR testing. Elworthy said that just one in 60 of his customers is interested in formal testing. “The customers with five to 25 users don’t want to go to the expense and trouble of doing DR testing,” said Elworthy.

### THE DR TESTING HAVES

For companies actually testing their DR procedures, Croy said he often finds that customers “plan their tests instead of test their plans.” Their tests become so scripted, he said, that they don’t end up with a valid assessment of how well the company would react to an actual disaster.

To get customers moving in a different direction, Croy said he tends to ask them questions along the lines of, “What are you trying to accomplish with that recovery? Have you determined the RPO [recovery point objective]/RTO [recovery time objective]?”

DR planning and testing shouldn’t just include plans for actual disasters and typical business interruptions, said Croy, who pointed to a possible



pandemic as an example of the kind of scenario that could cause a business interruption in the absence of an actual disaster. “We might have a great deal of absenteeism,” he said. “A pandemic could cause more employees to be working from home.” The questions he addresses to his customers for that scenario are: “Have you tested your remote access to see if it will support business functions?” and “Do you have enough bandwidth to make sure that they can get work?”

For VARs that are focused on the SMB market, that kind of DR plan sophistication may seem impossible to achieve with their customers, but at the enterprise level, according to Croy, Forsythe has customers that are thinking at that level of readiness. “They want to know if they’re ready for pandemic and absenteeism,” he said.

### THE DR TESTING HAVE-NOTS

According to Toigo, storage solution providers that don’t have a DR testing practice are simply leaving money on the table. Part of the problem, he said, is that people see DR testing as an onerous task. “There are a lot of ways to test tape backup without doing a full restore. A lot of tools will allow you to confirm that the tape is restorable.”

And rather than unplugging a system to see what happens when it fails, Toigo suggested a more simulated approach, where stakeholders in a DR project are told various systems are inoperable and need to determine how to address the purported outage. “I’ll go into a data center and put post-its on certain hardware and tell [the IT staff] that those systems are down” and that they need to react.

Croy also suggested that SMBs could suffice with a simplistic approach to DR testing. “This test may consist of simple restores at another location along with verification of connection capabilities,” he said. “By being much more selective in the parts of the plan that are tested and limiting it to mission-critical portions of their operations, a business can achieve some excellent results.”


When it comes to software to help ensure DR readiness, Toigo recommended that solution providers consider two classes of products for their customers: “aggregators” and “wrappers.”

The aggregators include products such as Continuity Software’s RecoverGuard, which monitors system health. “RecoverGuard gives

**“By being much more selective in the parts of the plan that are tested and limiting it to mission-critical portions of their operations, a business can achieve some excellent results.”**

—MIKE CROY, author and director of business continuity solutions, Forsythe Solutions

you a high degree of readiness should something happen. It's certainly a good tool," said Croy. But, he said, "there's no tool—period—that replaces testing. RecoverGuard will tell you the state [of the systems] but the actual recovery is something RecoverGuard won't give you."

On the other hand, wrapper applications, such as CA's XOsoft, Double-Take Software Inc.'s Double-Take, EMC Corp.'s RepliStor and Neverfail Group's Neverfail, monitor system health and coordinate data replication between platforms, according to Toigo. Of those applications, Toigo said XOsoft makes a lot of sense for VARs. "[CA's partner program] is the easy one to get into." 

---

Sue Troy is the Site Editor for [SearchStorageChannel.com](http://SearchStorageChannel.com).

## Check out the following resources from our sponsors:



Reduce your downtime and ensure your applications are always available, even in a disaster.

The secret to disaster recovery for any application -- revealed!

Reducing the costs and risks of branch office data protection



Five key questions for assessing backup and recovery solutions

Five cost-effective ways to enable fast recovery

The keys to disaster recovery planning: i365's EVault disaster recovery solutions help protect you from losing valuable data due to complete site outage



QuestFlex

DR

Cloud and managed services

## Regional Solution Providers

