Kick-Start Your Disaster Recovery Plan

Disaster recovery planning can mean a lot of different things and that's a key point to remember; your organization's IT disaster recovery plan should be tailored to its specific needs.

Getting Started With Disaster Recovery Planning

Year after year, we read survey results that indicate that organizations do not feel confident in their ability to resume operations within a reasonable time frame following a disaster. Many indicate that they have a hard time getting management buy-in for DR planning because it can be a time-consuming, expensive endeavor that has no immediate return.

However, in the past few years, new technologies have enabled some organizations to develop DR strategies that they would not have been able to afford a few years ago. These technologies are not a panacea, of course. They are merely tools that can help you achieve disaster preparedness.

Which brings us to a basic but important point—DR is not just about technology. At its core, it is about identifying risks and addressing them by any means necessary.

Your DR plan must address everything from data protection to personnel issues. It can be like peeling back the layers of an onion—every organization has its unique set of interdependencies and needs. These must be considered when getting started with disaster recovery planning, and that can be pretty overwhelming.

This Drill Down is a good place to start. You'll find an article detailing 10 common mistakes to avoid when you are creating your DR plan. There is also a piece on DR as a service, which is becoming increasingly popular but comes with its own set of concerns. Finally, you'll learn how to integrate disaster recovery planning into your organization's culture.

OK, that's it from me. Read up and get going. Disaster preparedness and business continuity is important stuff. Don't put it off any longer. ■

Andrew Burton
Senior Editor
SearchDisasterRecovery

10 Mistakes to Avoid in Your Disaster Recovery Planning Process

AT THE START of the new year, many IT folks (and perhaps a few business managers) resolve to take steps to prevent avoidable interruption events and to cope with interruptions that simply can't be avoided. In short, they decide to get serious about data protection and disaster recovery planning for business IT operations.

WHY THE DISASTER RECOVERY PLANNING PROCESS CAN BE SO TOUGH

Disaster recovery (DR) planning is a complex and time-consuming task when done properly. This helps to explain why recent surveys have shown a decline in the number of companies with continuity plans. In one annual PricewaterhouseCoopers study, companies with DR plans are down from roughly 50% of those previously surveyed to approximately 39% last year. Of these companies, the ones that actually test their plans are usually a fraction of those that claim to have a plan, raising further concerns about the actual preparedness of those firms with documented, but untested, plans.

Planning activity has also dropped off because of misperceptions about its necessity and value. It may seem obvious that "doing more with less" means "doing more with computers" and that downsizing staff actually increases dependency on automated resources and reduces tolerance to interruptions, even short-term ones. However, organizations aren't making the connection between these insights and the need to ensure that automation is resilient and continuous.

Money is also a hurdle, as it always is. Managers can always think of ways to invest money so that it makes more money for the organization—an option that's generally preferred to spending money on a continuity capability that may never be needed. With some economic

uncertainty in today's marketplace, this normal preference to focus spending on initiatives with revenue-producing potential is even more distorted, often at the expense of initiatives focused solely on risk prevention.

DR IS AN INVESTMENT

Common sense regarding the need to allocate budget, resources and time to the DR planning process may also be diminished by the marketecture and hype around technologies such as server virtualization, data deduplication and clouds.

Over the past few years, vendors have spent considerable effort trying to convince users that a side benefit of those technologies is increased protection for data and operations. "High availability trumps disaster recovery," according to one server virtualization hypervisor vendor's brochure. "Tape Sucks. Move On" was emblazoned on bumper stickers distributed at trade shows by a dedupe appliance vendor. "Clouds deliver Tier 1 data protection," claimed a service provider's PowerPoint.

These statements suggest that disaster recovery planning is old school, replaced by resiliency and availability capabilities built into new products or services. Most of these claims are downright false or, at least, only true with lots of caveats.

1. Don't think high availability equals DR. Perhaps the first and most important mistake to avoid when undertaking to build a disaster avoidance and recovery capability is to believe vendor hype about the irrelevancy of DR planning. While improvements might be made in high-availability (HA) technology, this changes nothing about the need for continuity planning. Although HA is part of the spectrum of alternatives for recovering from a disaster event, the use of HA strategies is constrained by budget. HA (failover between clustered

These statements suggest that DR planning is old school, replaced by resiliency and availability capabilities built into new products.

components) tends to be much more expensive than alternatives and is inappropriate for workloads and data that don't need to be available continuously. For most companies, only about 10% of workloads actually fall into the "always on" category.

2. Don't try to make all applications fit one DR approach. A second common mistake in planning, and one closely related to the first mistake, is to try to apply a one-size-fits-all data protection strategy. For the same reason that failover clustering isn't appropriate for all workloads, all data doesn't require disk-to-disk replication over distance, disk-to-disk mirroring, and continuous data replication via snapshots or some other method. The truth is that most data can be effectively backed up and restored from tape. Using disk for everything, including backup data, may seem less complex, but it tends to be far more costly and far less resilient. Given the numerous threats to disk storage, the problems with vendor hardware lock-ins for inter-array mirroring and replication, the costs of WANs and their susceptibility to latency and jitter, and many other factors, disk-to-disk data protection may not be sufficient to protect your irreplaceable information assets. At a minimum, tape will provide resiliency and portability that disk lacks. Think "defense in depth."

3. Don't try to back up everything. Expecting all your data protection needs to be included in a single backup process is another common mistake. The truth is that much of your data, perhaps as much as 40% to 70%, is a mix of archival-quality bits—important, but static, data that should be moved into an archive platform and dreck (duplicate and contraband data that should be eliminated from your repository altogether). Only approximately 30% of the storage you have today requires frequent backup or replication to capture day-to-day changes; the other 70% requires very infrequent backing up, if at all. You can take most of the cost out of data protection and shave precious hours off recovery times if you segregate the archive data from the production data. Doing so will also reclaim space on your expensive production storage environment, bending the cost curve on annual storage capacity

expansion and possibly saving enough money to pay for the entire data protection capability that you field.

4. Don't overlook data that's not stored centrally. This mistake is forgetting about outlying data repositories. Not all important data is centralized in an enterprise SAN or some complex of scale-out network-attached storage boxes. Mission-critical data may exist in branch offices, desktop PCs, laptops, tablets and, increasingly, smartphones. Recent surveys by TechTarget's Storage Media Group reveal that even before the rise of the bring-your-own-device (BYOD) era, companies weren't doing a very good job of including branch offices or PC networks in their data protection processes. In another study published this year, 46% of 211 European companies admitted they had never backed up user client devices successfully and that BYOD looms on the horizon as a huge exposure to data loss. You need to rectify this gap and may find it possible to do so with a cloud backup service, provided you do your homework and select the right backup cloud.

5. Don't mismanage data and infrastructure. Another mistake DR planning newcomers often make is ignoring root causes of disaster, such as lack of data and infrastructure management. Lack of data management, the failure to classify data according to priority of restore, is a huge cost accelerator in the disaster recovery planning process. Without knowing which data is important, all data needs to be protected with expensive techniques. As for infrastructure, you can't protect what you can't see. The failure to field any sort of infrastructure monitoring and reporting capability means that you can't respond proactively to burgeoning failure conditions in equipment, inviting disaster. These gaps can be addressed by deploying data classification tools (and archiving) to manage

The truth is that much of your data, perhaps as much as 40% to 70%, is a mix of archival-quality bits—important, but static, data.

data better, and resource management tools to manage infrastructure better. And, with respect to infrastructure management, tell your equipment vendors that you will no longer purchase their gear if you can't manage it using the infrastructure management software you've selected. That will also drive some cost out of your normal IT operations.

6. Don't try to duplicate equipment configurations at the recovery site. Given that only a subset of applications and data typically need to be re-instantiated following a disruptive event, you don't need to design a recovery environment that matches your normal production environment on a one-for-one basis. Minimum equipment configurations (MECs) help reduce DR environment cost and simplify testing. Often, you can use server virtualization technology to host applications in the recovery environment that you may not trust to a virtual server under normal circumstances. Testing is key to making the transition, whether from a physical host to a MEC host, or physical to virtual.

7. Don't forget to fortify your WAN connections. Vesting too much confidence in WANs and underestimating the negative effect they can have on recovery timeframes is another mistake. WANs must be properly sized and configured, and perform at peak efficiency to facilitate data restoration or to support remote access to applications. Regardless of the service-level agreement promised by your cloud host or cloud backup service provider, your actual experience depends on the WAN. Don't forget about providing redundancy (a supplemental WAN service supplied via an alternative point of presence) in case your primary WAN is taken out by the same disaster that claims your production environment. And also keep in mind that your WAN-connected remote recovery facility or backup data store should be at least 80 kilometers from your production

You don't need to design a recovery environment that matches your normal production environment on a one-for-one basis.

site and data as a hedge against both sites being disabled by a disaster with a broad geographical footprint. Most metropolitan-area networks that provide lower cost, high-bandwidth multiprotocol label switching connections do *not* provide sufficient separation to survive hurricanes, dirty bombs or other big footprint disasters.

8. Don't put too much trust in a cloud provider. While not yet as prominent as some of the aforementioned potential pitfalls, placing too much trust in a cloud service provider to deliver disaster application hosting or post-disaster data restoration can backfire. If you're using an online backup provider, for example, you've probably moved data to the backup cloud over time. You might be surprised how much data has amassed at the service provider, and at the length of time and the amount of resources required to transfer it back to a recovery environment. Remember: Moving 10 TB over a T1 network takes at least 400 days. Alternatively, if your plan is to operate applications at a cloud infrastructure provider, using the latter as a "hot site" for example, then be sure to visit the cloud provider's facility in person. In the 1970s, when hot site facilities were first introduced, there was a guy selling subscriptions to a non-existent hot site who, once his scam was discovered, retired to a non-extradition country before he could be arrested. At a minimum, if you plan to use a cloud to host your recovery environment, make sure that it actually has all the bells and whistles listed in the brochure, including that Tier-1 data center.

9. Don't let app designs foil DR. This mistake is procedural. Planners need to stop accepting the notion that DR planning is a passive activity—that you're dealt some cards and are required to play the hand as it was dealt. For business continuity capabilities to be fully realized, resiliency and recoverability should be built into applications and infrastructure from the outset. However, few DR-savvy folks have been given seats at the tables where applications are designed and infrastructures are specified. This must change going forward. Put bluntly, bad design choices are being made right now that will obfuscate some company's recovery

efforts in the future, including the platforming of applications and data in proprietary server hypervisors or storage platforms; coding applications using insecure functions; and employing so much caching that significant amounts of critical data will be lost if an interruption occurs. If DR planners can get involved early on, better design choices can be made and IT can be much more recoverable at a much lower cost.

10. Don't forget to follow the money. Management holds the purse strings, so it could be a big mistake if you don't make the case for your DR plan based on business value rather than technical terms. You need to show management that you're doing everything possible to drive cost out of the continuity capability without sacrificing plan efficacy. You also need to emphasize the investment risk reduction and improved productivity the plan offers, thereby providing a full business value case. Only then will you have a chance of overcoming the natural reluctance of management to spend money on a capability that in the best of circumstances will never be used.

For the record, the greatest expense in DR planning isn't the cost for data protection, application re-instantiation or network re-routing; it's the long-tail cost of testing. So, try to build a capability that can be tested as part of day-to-day operations, alleviating the burden on formal test schedules, which should serve as logistical rehearsals (not tests) of whether or not data can be restored.

—Jon William Toigo

DR as a Service Scales 'Big Company' DR Down to Size

HISTORICALLY, DISASTER RECOVERY (DR) offerings have run the gamut from ways to simply get data off-site, to replicating data sets to a second physical location or completely mirroring an infrastructure that runs in a failover mode. The thing that differentiated these products was recovery time, or how quickly an application could be back online. These services varied widely in what they provided.

The faster a company needed its applications back up and running, the more complexity and cost they had to assume. And lower-end products were essentially limited to restoring data, not actually recovering applications. For many companies, the cost of a more sophisticated DR as a service offering was just too high.

Today, the cloud has greatly enhanced those options. If recovery time is the metric that determines how effective the DR as a service provider is, then good DR is now available for a lot more companies.

WHAT IS DR AS A SERVICE?

The cloud has made "as a service" part of the IT lexicon, joining the list of options that IT managers can deploy. These offerings are essentially various combinations of compute and storage infrastructure and applications running in the cloud that cloud companies sell access to.

Server virtualization allows providers to package the entire IT environment into a cloud service, like Amazon has done with EC2, allowing companies to exist without ever owning a server. This capability of running virtual machines (VMs) in the cloud has enabled DR to join the list of "as a service" products, dramatically bringing down the cost of high-quality disaster recovery protection.

But rather than just providing a safe place to transfer a company's most critical data and (hopefully) get it back soon enough, DR as a service also provides a platform on which to actually run those applications in the cloud.

This addresses the primary challenge of DR: recovery time. Here's how it works:

The company's critical VMs are replicated to a host in the cloud and kept current with regular updates. Then, when a failure occurs on the local host, a failover mechanism kicks in and points users to the VMs running in the DR as a service provider's cloud. Following the outage when local infrastructure has been restored, those VMs can be migrated back to the company's data center or resynchronized with the local host. This can be done over the Internet if the amount of data is relatively small, or can involve shipping a storage device.

Obviously, running these applications directly from the cloud subjects those users to more latency than when running locally, but the impact of this would depend on the application itself and how much actual data transfer it had to do.

Keeping a VM updated in the cloud can require a lot of bandwidth and restoring it locally after a server goes down can take hours (or more). And more importantly, in a pure cloud service, all this data movement must be handled by the application server. For these reasons a hybrid cloud model is especially effective for DR as a service.

Hybrid cloud DR as a service includes putting a server on-site that functions as the local replication target and local failover appliance. When an application fails, its corresponding VM can be started on the appliance by an administrator, giving users a recovery measured in minutes, depending on backup frequency, instead of hours. For events that don't take out the company's building, having the local failover host eliminates the latency issues that the pure cloud DR as a service product has.

When a failure occurs, a failover mechanism kicks in and points users to the VMs running in the DR as a service provider's cloud.

On the back end, the hybrid appliance replicates VMs regularly to a cloud host, handling all the communication and data transfer and removing this load from the primary application servers. Also, many hybrid DR appliances

also offer data reduction functionality to maximize bandwidth utilization.

Some hybrid DR-as-a-service products can also provide protection for non-virtualized servers as well. They do this by running a physical-to-virtual conversion of the local bare-metal server, creating a virtual recovery node that is kept updated like all the other VMs on the target appliance.

Of course, hybrid cloud DR as a service isn't perfect. Running an application from the cloud does involve latency that's simply not there on the primary production servers. And, depending on the product chosen, failing back from the cloud can be a complex process and involve some downtime.

But the alternatives can involve more downtime, more complexity and certainly more cost. DR as a service can provide a highly functional disaster recovery product that is also very cost-effective. The right hybrid recovery appliance can give a smaller organization "big company" DR, with near-real-time failover and a reasonable recovery window from the cloud. Hybrid cloud can also provide a high-availability infrastructure for multiple applications (on physical or virtual servers), offering protection against the much more common "disasters" in which a server fails, without the site going down.

—Eric Slack

Business Continuity and Disaster Recovery Plans Must Affect IT Culture

THE WAY A company handles its information technology—systems, support and strategy—is part of its culture. One could even argue that anything ingrained in an organization's culture gets done and the rest doesn't.

Backup tasks often don't affect corporate culture. They probably should, but they typically don't. Instead, backup is often seen simply as a bunch of after-production tasks to make one or more (often a lot more) copies of what's in production. That explains why so many business continuity and disaster recovery (BC/DR) plans suffer atrophy—they're often developed within the vacuum that is IT, and therefore don't affect the ongoing organic culture of the organization.

BC/DR preparedness has to affect corporate culture. Why? Because if you have developed a BC/DR plan that hasn't affected corporate culture, that plan became out of date the day after you published it. You need to recognize that production environments continually change: new servers are added, machines get moved and the critical nature of services changes. If you haven't made preparedness part of production, then when the changes happen in production, they won't be organically reflected in your preparedness plan. And when it comes time to actually fail them over, you won't know about them because your documentation effort stopped the day your plan was published.

BC/DR planning has to affect corporate culture so that as production evolves, your BC/DR plan evolves as well. For example, whenever IT decides to stand up a server or a new service, the first questions its operations person should ask are: "What do we need to do to update our BC/DR plan accordingly? Should we start replicating that virtual machine? How often should that server's data be protected? How long should the data on the server be retained?"

Answering those questions takes more than the backup admin's opinion, which is just one of the many reasons why BC/DR planning requires a wider effort. In the broader sense, the initial BC/DR initiative, the first BC/DR plan, the ongoing "preparedness as part of production" culture shift, and the recurring BC/DR plan testing and maintenance all take a multi-member, cross-functional team:

- Executive sponsorship is needed to ensure that the plan does affect culture. The backup administrator isn't going to change the culture of the IT team, much less the culture of the whole business.

- In many cases, you, as the backup manager, don't have enough information to understand (as production changes) what related changes the BC/DR plan needs to receive. Often, that's where tech tools can help; they can assess what's on the wire through discovery and potentially tell you what the interdependencies are, which is hard to discover otherwise. You have to know how your production environment is evolving, so your BC/DR team can sustain and evolve the BC/DR plan accordingly.

Preparedness has to be part of production. It's a cultural change that's dependent on a technology-level understanding of what is in, and what is evolving with, the IT infrastructure.

This emphasis on needing a broader team than just the backup staff won't diminish the value of their role. After all, no amount of process or procedure will help if the data doesn't survive the calamity. The good news for backup administrators considering their participation within a BC/DR framework is that the conversations BC/DR planning drives may actually help a backup administrator get to a managers' desk or a corner office.

When you think about ways to affect culture, to convert technical challenges into business challenges and solutions, that's when you go from being a manager of backup tactics to a leader of BC/DR strategy. Not only will your company benefit from the better preparedness as part of production, but the view from your desk might improve as well. —*Jason Buffington*

JON WILLIAM TOIGO *is a 30-year IT veteran, CEO and managing principal of Toigo Partners International, and chairman of the Data Management Institute.*

ERIC SLACK *is a senior analyst at Storage Switzerland.*

JASON BUFFINGTON *is a senior analyst at Enterprise Strategy Group. He focuses primarily on data protection, as well as Windows Server infrastructure, management and virtualization. He blogs at CentralizedBackup.com and tweets as @Jbuff.*

TechTarget

Kick-Start Your Disaster Recovery Plan
is a SearchDisasterRecovery.com e-publication.

Rich Castagna | *VP of Editorial/Storage Media Group*

Andrew Burton | *Senior Site Editor*

Ed Hannan | *Managing Editor*

Dave Raffo | *Senior News Director*

Linda Koury | *Director of Online Design*

Neva Maniscalco | *Graphic Designer*

Jillian Coffin | *Publisher*
jcoffin@techtarget.com

TechTarget
275 Grove Street, Newton, MA 02466
www.techtarget.com