IT Disaster Recovery Planning:

A Guide to “Getting It Done”

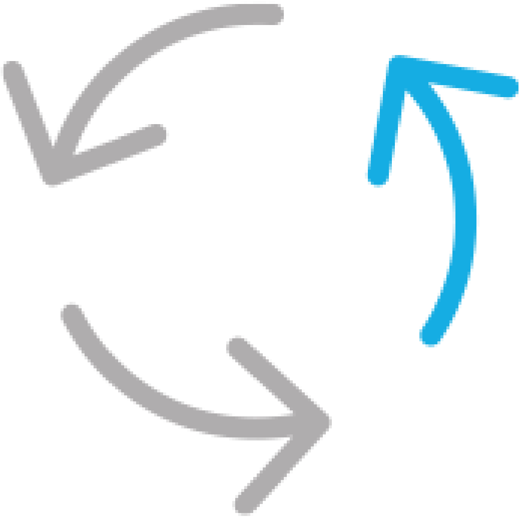

jelecos

Technology—Powered by People

With all the news about cyber-attacks taking down big company web sites, data breaches, theft of personal data from major organizations such as Target and Aetna, and the damage data centers endure due to natural disasters (flooding from Hurricane Sandy hobbled several data centers), you'd think every organization would have a disaster recovery (DR) plan in place. Surprisingly, many don't.

For some it's a matter of time and cost. The assumption is that DR planning doesn't come cheap, and can be time consuming. For others, it's a matter of thinking that disasters only happen to "someone else" or that there is little chance of some kind of serious business disruption happening. That kind of thinking can put organizations out of business if a disaster — natural or manmade — occurs and they aren't ready. DR planning doesn't have to be expensive or time consuming. It's a matter of knowing what to do and then making it happen. The information that follows can help you get your DR plan in place.

The Cost of Downtime

Before delving into the specifics of DR planning, let's consider a few basic realities that reinforce the need for a DR plan. Things happen. A power outage. A fire in the office suite next to your business. A cyber-attack that takes down your website. A flood that demolishes the data center where your servers are housed. Whatever the cause, the resulting downtime won't be cheap — and could possibly put your company out of business or, at the very least, cause you to lose revenue and customers.

A 2016 Ponemon Institute survey of Data Center Organizations reported that unplanned downtime cost $8,851 per minute, up from $7,908 per minute in 2013 [1]— and that the average cost of an outage was $740,357, up from $690,204 in 2013. Other surveys report similar findings. But the costs of downtime aren't just measured in dollar figures. Consider:

- **Brand damage:** With Twitter, LinkedIn and other social media channels at our fingertips (not to mention regular news outlets), it's virtually impossible to hide the fact that a major outage occurred. Your brand's reputation is at risk as soon as those affected share the news.
- **Loss of competitive edge:** We live in a "want-it-now, on-demand" world. If you can't deliver what your customers want when they want it, your competitors will. The unavailability of your services or products will send your customers to the next best provider.
- **Diminished customer loyalty:** All it takes is one bad experience for a customer to lose faith in your business. As in the previous bullet point, it's likely that when customers become unhappy with you, they'll go elsewhere and cease to be your customers.
- **Low employee morale:** Downtime doesn't just impact employee productivity, it can also result in stress and lower morale. Some employees may need to work longer hours — often without extra compensation — to mitigate the crisis and get systems back online. For others, it means finding ways to win back customers or make up for losses, increasing their workload or taking them away from their regular responsibilities. Still others will worry that lost revenue could put their jobs on the line.

These are just a few of the ways downtime can have a negative, even disastrous, effect on your business. Developing a DR plan for your IT assets can help you prepare for and mitigate downtime — and the potentially business-destroying effects.

[1] 2016 Cost of Data Center Outages. Ponemon Institute, January, 2016.

Assess Your Risks

Whether you take on DR planning yourself or outsource it, one of your first steps will be to conduct a risk assessment. A risk assessment examines the vulnerability of your IT assets to events that can cause downtime. The events can range from the mundane, such as hard drive failures, internet outages, application crashes, to the cataclysmic, such as malware or natural disasters. Each type of event has its own associated recovery measures and costs.

Geographical location and regional weather patterns also can play a role. For example, if you're located in area where tornadoes are common or where roads are frequently flooded during storms (which could prevent employees, as well as emergency services, from accessing your site), your IT operations could be at risk.

Other issues to consider are deficiencies in building construction or being located in a multi-tenant building in which a fire or broken water pipe in the suite next door could put your business as risk.

Inventory Your Assets

You know the potential risks. Now take stock of your IT assets that could be affected. Get input from others within your company as "shadow IT" could have introduced essential data and applications without the knowledge of your IT department.

Identify any application and system dependencies as well. Few applications are an island anymore. Most rely on several other applications and services within or beyond your IT infrastructure. For example, you may have a commerce application that incorporates an authentication server, a product database, or an inventory system from a partner or supplier. You can't successfully recover this application without ensuring that the applications on which it depends on are recovered as well. Most likely, you also have hardware or equipment dependencies such as an email server may that depend on a storage array, which in turn depends on specific network equipment.

Conduct a Business Impact Analysis (BIA)

The next step is to determine the operational, financial and reputational effects to your business if your IT assets were not available. Take into consideration the costs associated with downtime including, but not limited to, lost revenue, delayed sales, regulatory fines, customers who move on to your competitors, and brand damage. This will help you to calculate how much downtime your business can tolerate.

You'll also need to identify the resources required to maintain your business operations. Then establish an order of priority for restoring business functions and related data or applications. Lastly, don't forget about any compliance requirements.

Set Your Recovery Objectives

Once you understand the potential costs of downtime and the likelihood of particular threats, you can set recovery objectives for each of your critical applications.

Recovery point objective (RPO). Your RPO tells you how much time from the point of the outage you can afford to lose. RPO may vary based on the asset. If it's a server used for a microsite developed for an ad campaign, you may be able to tolerate a day or two of lost data. If it's your e-commerce system, any more than a few minutes of lost data may drastically reduce revenue. If you're in a healthcare organization, data loss could result in regulatory fines.

Your RPO will determine the frequency with which you'll need to replicate data from your production site to a DR site. If your RPO for an application is one hour or less prior to disaster, you'll need to replicate data at least hourly. If you can't afford to lose any data, ever, you'll need to implement synchronous replication for that application. In other words, you'll need to have your data written to a DR site at the same time it's being written at your production site.

Recovery time objective (RTO). Your RTO establishes how quickly you need to have data or applications back up and available after disaster. This answer is less about data loss, and more about having the application or data available. This will vary widely depending on the application and/or data and who uses it. Your internal data warehouse may need to come back online in several hours, whereas a customer-facing website may need to be back up immediately.

Recovery capacity objective (RCO). How much computing capacity do you need to recover and by when? This will depend on your collective RTOs, and as a result is usually phased over time. For example, suppose your IT infrastructure has 40 virtual machines (VMs) and 10 terabytes (TBs) of storage. A power outage takes the system down. At an RTO of one hour, you may need 20 VMs and 5 TBs to support your most critical applications but can get by without full capacity for several days.

Pick Your Strategy

Now comes the fun part — reviewing and evaluating different data backup, replication and recovery strategies.

Will tape backup work or is disk backup the better option? Should you keep your data backups on site or move them to an off-site facility? If you go with off-site backup, do you want a facility that is equipped with everything you need to get recovery started immediately – even if you are paying for everything even when it's not in use. Or, does it make more economic sense to bring in the equipment needed for data recovery at an off-site facility only when needed, even if that means the total time required for recovery could be quite long. And, what about cloud-based options? Each offers its own advantages and disadvantages.

Tape Backup. Among data backup options, tape is the least expensive strategy with a lower cost per gigabyte of storage than other methods. A tape drive or tape library is infinitely scalable. Just buy more tapes. It's also efficient, as it is only run when data is being read or written. Tape makes it easy for multi-site backups too.

However, the costs for tape backups can quickly rise when you factor in the expenses for securely transporting tapes and storing them in off-site facilities. Reading from and writing to tapes also takes longer than other backup methods. In addition, there's a greater chance of data corruption. Recovering from tape backup is slow as well. It can take days or even weeks if the hardware needs to be found before recovery can commence.

Disk Options. Disk backup is more efficient than tape, with faster recovery and more continuous protection. But if it still resides in your on-site data location, it may not be accessible if a disaster makes your site inaccessible or nonexistent.

To make disk backup safer, use a third-party provider for off-site backup or a replication of the on-site disk. Not going with an off-site provider can make it expensive to continue buying disk space. Plus, there's extra time required for doing the backups yourself — and time is money.

To mitigate high prices for disk space, you can back up your disk replicas to tape. However, any savings can be quickly diminished when you take into consideration the extra labor and infrastructure costs associated with maintaining the tape-based backup along with disks.

Virtual Tape. The use of virtual tape makes it possible to save data as if it were being stored on tape although it could actually be stored on another type of storage medium. A special storage device manages less frequently needed data so that it appears to be stored entirely on tape cartridges even though parts may actually be located in faster, hard disk storage.

Virtual tape can be used with a hierarchical storage management system in which data is moved as it falls through various usage thresholds to slower, less costly forms of storage media. It can also be used as part of a storage area network (SAN) where less frequently used or archived data can be managed by a single virtual tape server for multiple networked computers.

Hot, Cold or Warm Sites. There is also the matter of on-site versus off-site backup. On-site backup keeps your data close at hand. But if your site is not accessible, or is damaged or destroyed, chances are your backups will be too. Off-site backup, particularly to a location a good distance from your primary site, obviously offers greater protection of your data. The question is what type of facility will you need? The three options are hot, warm and cold facilities, with the differences between them a matter of recovery time and the cost.

- **Hot site:** A hot site means that customer production data is running and lives at two geographic separate Data centers. If you lose Data center a, Data Center B takes the full production load and there is no interruption to the user. There is also redundancy with the communication lines between the two data centers. The data centers are fully equipped with servers that can be online within hours. It's expensive but it's also a great way to minimize downtime and data loss.

- **Warm site:** A warm site provides basic infrastructure and replicated data. However, it also requires some lead time to prepare servers, so it could take up to a few hours or even a week to bring online. It costs less than a hot site but the lead time required may not be worth it.

- **Cold site:** This is a bare-bones approach to DR. A cold site has the basic infrastructure needed to run a data center, such as heating, ventilation and air conditioning (HVAC), power, network connectivity, and replicated data, but little else. Equipment must be brought in and configured, which can take weeks to be operational. It's the least expensive of the options if you can afford the down time.

There's also the matter of location. If possible, you'll want to make sure that the same regional disaster doesn't affect both your production site and your DR site. That means your DR site should be in a "geographically diverse" location. At the very least, it would be in a different flood zone, on a different power grid, and serviced by a different telcom network than the production site.

You'll also need staff on hand to help with your DR plan implementation. Will your DR site be in or near a city where you have staff who can help work on the recovery? Can you get them to the site if public transportation is disabled or otherwise unusable? Can the staff be housed on-site or nearby if they have to stay for a while?

Keep in mind that disasters or business disruptions don't negate your responsibilities for meeting your regulatory compliance requirements. Some actually specify that you ensure that any DR infrastructure you establish meets the same or similar regulatory requirements as your production infrastructure. Failure to do could result in penalties for non-compliance. Your DR site may also need to meet security, data privacy, or monitoring standards or requirements specific to your industry. While these standards may not carry costly penalties for violation, not meeting them could be costly to your company's brand or reputation, and may cost you customers.

Cloud-based DR. Yet another DR strategy is cloud-based DR. Cloud-enabled DR delivers a number of advantages over traditional DR, such as reduced capital expenses because it eliminates the need for investing in a remote DR facility. Ongoing operating expenses are lowered because you don’t have to pay to power and cool remote equipment. Capacity and performance can be allocated on demand, so you only have to pay for the resources consumed. Because the cloud is designed for remote management, it speeds up recovery. Compared to on- or off-site tape-based DR, such capabilities can make routine testing more practical, and help ensure the DR service works when needed.

In addition, the cloud also makes warm site DR a more cost-effective option. Backups of critical servers can be spun up in minutes on a shared or dedicated host platform. With SAN replication between sites, hot site DR also becomes a less expensive option. SAN replication provides rapid failover to the DR site with very short recovery times, offering the capability to return to the production site when the DR test or disaster event is over. That’s typically not feasible with traditional DR due to the cost and testing challenges. Cloud-based DR also offers the ability to finely tune the costs and performance for the DR platform. Applications and servers that are considered less critical in a disaster can be tuned down with less resources, while still assuring that the most critical applications get the resources needed to keep the business running through any business disruption or disaster.

There’s also data backup to consider. Cloud-based backup is essentially off-site backup to a third-party service provider or to your own cloud infrastructure using cloud enablement technologies or on-site appliances. Multi-site data redundancy is integral to cloud-based data backup as a local data copy can live in an on-site appliance, while the enablement technology replicates data to your service provider or your own data center. The appliances and enablement technologies continually run in the background of IT operations, eliminating some of the issues associated with manual IT processes. Administrators can provision VMs on the appliances, while using the appliance to back up the VMs off site. There are no tapes or disks to buy or refresh, and no need to spend hours each week physically managing backups or transporting tapes.

If you are concerned that cloud-based data backup could open your data streams to breaches from a third party or as a result of other customers residing in a given data center, there’s no need to worry. Many cloud service providers (CSPs) build high-level security features built into their clouds. Typically, CSPs that are audited to meet the requirements of the Healthcare Information Portability and Accountability Act (HIPAA), Payment Card Institute Data Security Standard (PCI DSS), and other regulations or industry standards, employ security best practices to help ensure data safety and integrity.

The Bandwidth Factor

Another consideration in setting up your DR plan and the tactics within it if the bandwidth needed between your production and DR sites. That will likely be determined by what you’re doing over the connection, such as:

- Seeding. This is the initial process of recreating your production infrastructure at your DR site, and involves moving a lot of data. The more bandwidth you have, the faster you’re protected by the DR site and the less catching up you’ll need to do during the seeding period.
- Transmitting and/or replicating changes. Once you’re set up, you’ll need enough bandwidth to transmit the volume of changes in time to meet your RPOs. If you have synchronous replication for any of your applications, you’ll need enough bandwidth to transmit that data. You’ll also need to throttle-up from that for periodic backup transmissions.
- Reseeding. If an application or database at your DR site is corrupted, you may be required to throttle-up to reseed the lost data.
- Failing over. Once you initiate failover to the DR site, you’ll need enough bandwidth to keep your business up and running from your DR site.
- Coming back. When you’re ready to take your IT operations back to your production site, you’ll need enough bandwidth to reseed your production environment from the DR site.

Test, Test Again, and Again

While there may be other components required to create a DR plan that meets your company's full needs, you should be off to a good start. However, the most comprehensive, thoroughly reviewed DR plan will do you no good if it doesn't work. That's why testing your DR plan is as essential as developing the plan. You'll need to determine how often you need to test your plan. If you've just created it, frequent tests — typically quarterly — are in order to work out any issues. Once you are confident with the testing process and the overall components of your DR plan, once or twice a year should be sufficient. Make sure to note the dates of your DR plan testing. Many organizations get caught up in their day-to-day operations and then find that they have neglected their DR plan tests for well over a year.

You'll also need to determine the particular scenarios to be tested and the type of test to be performed. Will you be testing for a tornado, a power failure, some kind of cyberattack or some other specific natural or manmade disaster that could disrupt your IT operations? What testing methodology will you use? The following are some of the most common testing methods.

- Walkthrough - Verbally go through every step of your DR plan with your team. This helps determine if everything is set up properly in the event you have to activate your DR plan. Review the plan to make sure all the information is correct, resources are available, copies of the plan are where they should be, and everyone is in compliance with the plan. It is a good first test for a newly created plan to identify any gaps. It's the quickest and least expensive method — but also the least thorough approach.
- Simulation – A simulation is more in-depth than a walkthrough, and typically doesn't affect day-to-day business operations. It is scenario-based, focusing on specific types of business interruptions. It may involve role-playing and actual physical testing of DR sites and equipment, as well as coordination with partners, vendors and others. You'll rehearse notification procedures and temporary operating procedures, as well as verify your backup and recovery operations. You'll test all hardware, software and personnel involved in addition to voice and data communications and the functionality of utilities.
- Parallel – Parallel testing can be done in conjunction with the walkthrough or simulation tests. Its purpose is to process historical transactions against the preceding day's backup files at your DR site. The reports generated at the DR site should align with those produced at your production site.
- Full interruption – With a full interruption test, actual production data and equipment are used to test your DR plan. This has the potential to disrupt your day-to-day operations. It can also be time-consuming and expensive. However, it is invaluable in identifying any gaps or problems in your DR plan. To minimize disruptions to your business or potential downtime, conduct this kind of test during your least busy operating hours. The idea is to simulate a variety of disaster scenarios such as hardware failure, DDoS and power outages.

Don't be surprised if during testing, failures occur. That can be good because it's better to identify the problems during testing rather than during an actual event. Determine what needs to be done to correct any gaps or problems uncovered by the test and who needs to perform the required actions. Make sure to include an estimated completion date and then follow up to make sure the corrective actions are completed. Remember that DR plan testing is not or ever should be a one-time activity. Periodically conduct tests to make sure any changes in your business needs, infrastructure or resources are accounted for and that the plan is effective.

Outsourcing

When it comes to DR planning, you don't have to do it alone. In fact, it may be a good idea to outsource if you lack the time or in-house resources and expertise. There are a number of service providers out there that specialize in DR solutions. The key is to find one that can best meet your needs. Here are some things to consider when looking at potential DR service providers:

- Knowledge and interest. The service provider should know about the industry you're in and the challenges and opportunities that could impact your DR plan. The provider should also want to learn about your business, your goals and business requirements, your back-up and recovery challenges, your data and application profiles, and growth expectations.
- Ability to deliver. Can the provider achieve your RPO and RTO for your most critical data, applications and business operations? Is the provider willing to put a guarantee in its service level agreements (SLAs)? Does it have references who can vouch for its follow through and performance?
- Compliance. Operating physical and cloud infrastructures according to compliance requirements requires a continuing commitment to staying current with the ever-increasing industry and government regulations; investment in physical plant and technology to assure compliance; and training the workforce to understand and carry out their duties accordingly. Look for providers with a portfolio credentials, such as SSAE-16, PCI DSS, and HIPAA/HITECH. Don't take a provider's word that it holds specific certifications or undergoes compliance audits. Ask to see relevant documentation.
- Security. It would seem that security and compliance would go hand-in-hand, but don't assume that if you have one you also have the other. Find out about the scale and scope of a provider's logical and physical security policies, programs, testing and auditing.
- Support. The best DR service providers will have 24/7/365 technical support. They'll also have skilled experts to fully manage server and storage recovery at all times.
- Geographically dispersed data centers. The provider should have data centers strategically placed so as not to be affected by the same single disaster, as well as to have the power, communications, networking, redundancy, staffing and security to sufficiently meet your requirements. If possible, tour at least one of the provider's data centers.
- DR testing. Ask any potential DR service provider to explain in detail its testing processes. That will include what is tested, how it is tested, and how often tests are conducted.

You'll likely have a number of other questions, depending on your organization's specific needs and requirements. What data replication service options does the provider offer? How does it handle encryption? Can the provider handle hybrid cloud strategies? Can the provider customize your DR solution or does it only offer “off-the-shelf” services? Assemble all your questions before meeting with prospective DR service providers so you'll be better prepared to assess their capabilities.

Let Jelecos Help

If you need a place to start in terms of looking at DR service providers, consider Jelecos.

Jelecos can work with you to devise a customized DR solution to help minimize data loss and mitigate other potentially disastrous effects of any kind of business disruptions. Learn more about Jelecos' Disaster Recovery as a Service (DRaaS).

Jelecos is also well versed in helping customers cost effectively meet their regulatory and business needs. Jelecos invests heavily in compliance to provide clients with turn-key solutions including PCI-DSS, HIPAA, SOX, GLBA and more. Jelecos is one of the only service providers in the region with third-party PCI certifications in both infrastructure and application development.

To learn more, visit us at: http://jelecos.com, or Call 402.955.0489