The Disaster Recovery Playbook

In the immortal words of Arthur Conan Doyle, the creator of Sherlock Holmes, "It is easy to be wise after the event." Every IT team knows that feeling: the flood or breach or hardware failure has hit and left in its wake a map of actions that should have been taken much earlier.

After a disaster has taken down your datacenter and crippled your infrastructure, it's easy to wish you had invested in stronger preparation. Maybe you'll wish that you chose a more advanced backup and disaster recovery solution, with tested backups safely offsite and all critical data, apps and servers protected. Or that you spent more time training your staff so they knew exactly how to respond and who to contact. But most likely you'll wish your team had documented and tested a solid a disaster recovery (DR) plan that was available in an externally-hosted location.

Maybe you know that outages cost enterprises $700 billion a year, with each downtime event costing between $1 million a year for a midsize company to more than $60 million for a large enterprise. That price tag is impacted by many things: the ability to recover quickly, whether you're using a unified or multi-vendor backup and disaster recovery solution, the quality and security of your backups and other factors. But in the heat of the disaster moment, one factor that especially impacts your survival is the viability of your disaster recovery plan.

A good plan ensures your team acts swiftly and efficiently in a crisis, rather than getting lost in chaos. But that's only if your plan is good. Many teams find themselves under-protected if their plan hasn't been tested or kept current or if it only accounts for some risks and not the actual cause of the incident.

The High Stakes of Your Disaster Recovery Plan

Every disaster comes with a dollar sign attached. Each minute of chaos, downtime and cleanup drives the amount higher. Your disaster recovery plan must go beyond downtime and account for the following factors:

- **The IT team's ability to operate.** If your servers, networks, and apps are down, is the team still able to fix the issue – or access the response plan?
- **Your employees.** How could each disaster take your staff offline? Would they be able to work remotely?
- **Business as usual.** Can customers keep buying – can sales people keep selling? How hard hit will your brand image be if your site stays dark?
- **The cleanup**. How extensive will the recovery process be? What kind of new equipment will you need in case of flood, fire or other kind of loss?

Making A Speedy Recovery: Best Practices

Create a unique plan. There are some great guidelines for disaster recovery plans out there (some are included at the end of this guide.) But generic DR plans should serve as springboards only. Your organization's disaster recovery plan should address your unique needs, from your system configurations to your buyers to your team skill set. Even competitors in your industry will have different nuances that take their plans in divergent directions.

Translate recovery into financial terms. Assessing the Total Cost of Ownership (TCO) of your DR solution is a must when it comes to getting executive buy-in for new technology, additional staff or an overhaul of your existing system. You'll want to factor in your capital expenditures (the cost of your hardware, software and implementation) and your operational expenses (such as datacenter maintenance and labor time.) Once you know how much you're spending to recover, stack up that amount next to the hypothetical cost of a disaster to see if your solution is truly cost-effective. The financial damage will vary depending on multiple factors, but you can get a workable idea by using this formula: *Lost Revenue + Lost Productivity + Recovery Costs + Cost of Brand Damage = Cost of Downtime.*

Test like your survival depends on it. Because it does. Defining your RTO and RPO metrics is important but you need to test your plan to know if you can achieve either when it counts. Remember that a "failed" plan is a good thing, as the failures will show you exactly what you need to fix before a drill becomes a reality. With continual changes in staffing and systems, your plan should be tested on a regular basis even after the results are successful.

Train and test your employees. The DR focus on making technology work often means that IT leaders forget to test their employees. Make sure your team and every relevant staff member and partner understands their role in disaster recovery and can jump into action at a moment's notice.

Check your service-level agreements (SLAs) for disaster service. Most organizations have some type of vendor partnership when it comes to backup and disaster recovery and other IT needs. Whether you have a DRaaS vendor or simply rent datacenter space, you'll want to ensure your contracts cover service during a disaster – and that it's adequate.

Practice dispersed redundancy. While most IT leaders worry more about human threats than natural disasters, floods, earthquakes, hurricanes and other calamities can take out a nearby secondary datacenter as easily as your primary datacenter. Make sure you have mission-critical data stored a good distance away, ideally on a different power grid. Cloud backups offer another layer of protection.

Store your plan in multiple locations, including the cloud or an externally hosted site. If your disaster takes out your central site and systems, your team may not be able to access the disaster recovery plan – rendering it useless. Document the plan and store it in an offsite location that will be accessible even if your own infrastructure becomes impenetrable. Be sure to include any vital documents like maintenance contracts, the call chain, instructions for the BDR solution and configuration diagrams.

Maintain a full bench of staff who understand both the BDR solution and plan. There is no security in having a single person – or even trio – responsible for the entirety of disaster recovery. Employees leave for other companies, get sick, go on leave or simply become unavailable during a disaster. Train multiple people on the recovery plan and process – including employees in other geographic regions if possible.

Remember your outside and hosted applications. If your payroll, HR or other applications are external, you'll need to make a plan for them too. If your employees are working from home or routed through your failover facility, those unfamiliar IP addresses may not be validated for critical applications. To avoid this, make sure those hosted applications have an alternative set of IP ranges and secondary Active Directory copy. The same should apply to any data flowing to and from failed servers.

Keep your plan agile and adaptable. Nothing in IT or business stays unchanged for long. Vendors, leadership, tools and key personnel will shift, which can render your plan obsolete in just a few months' time. As new applications rise in importance and your mission-critical systems and data change, your recovery plan will need to shift to match your production environment.

Have a Plan B. We all know the quote about plans not surviving contact with the enemy. While testing should help prevent your plan from failing during a crisis, it's wise to create a backup plan. Maybe the right team members won't be available. Maybe a snafu takes down a backup datacenter. Make sure you have a Plan B that can be used in an emergency or identify a partner or vendor who can assist.

7 Components of a Disaster Recovery Plan

Updating your backup and recovery strategy doesn't have to mean re-creating the wheel. In many cases your team can leverage your existing infrastructure by partnering it with more advanced new tools and backup appliances that speed recovery and make your life easier.

Creating and maintaining your DR plan, however, is an ongoing endeavor. Whether you're starting from scratch or you have an existing plan that you'd like to refresh, you'll need to go through 7 steps.

Take an Inventory.

Before you can decide *how* you'll recover, you need to know *what* you're recovering. What software and hardware are you dependent on? What kinds of data are your systems receiving, storing and transmitting – and how much of it is mission critical? What information do your customers, vendors and partners rely on from you? Your DR plan must account for a complete inventory of your data, applications and servers.

Once you have an accurate and comprehension overview, you'll need to take a critical look at your assets in terms of their protections and vulnerabilities. Examine who has access to them and which authentication and transmission controls are in place. You'll also want to consider any technical weaknesses or configuration flaws that could open the door to a breach. Finally, you should organize support and contract information for any vendors.

Conduct a Risk Assessment.

Once you know what you'll be recovering, you'll want to identify the risks you're facing. Most threads can be grouped into one of four categories:

- Natural threats such as floods or earthquakes

- Intentional human attacks or accidental errors
- Environmental threats such as power outages or leakages
- Server and other hardware failures

Once you've outlined your risks, you'll want to assess potential disasters in terms of:

Impact: How long would your downtime last? How many sales would be lost? How many customers, patients, employees or partners would be affected? Would you face regulatory fines or need to buy new equipment?

Likelihood: How likely is an earthquake compared to a ransomware attack? How about a critical server failure as opposed to an electrical fire or an administrative error?

One threat might carry a higher likelihood of occurrence but with a lower impact, while another threat could be the opposite. By averaging the combined factors, you'll be able to accurately assign risk levels to each area.

Define Your RTO and RPO.

Your Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) will steer your recovery strategy. Your RPO is the point in time from which you need to recover. It could be hours, it could be days, but it will represent the moment of failure. You'll aim to retrieve backups dating back to that point in time. Your RTO is the downtime your organization can tolerate without severe impact. This too is measured in minutes, hours or more, but goes forward from the moment of disruption, rather than backward.

Your RTO and RPO may vary by application, region, time of day and other factors. As you calculate objectives based on primary versus secondary or tertiary systems for seamless continuity, you'll prioritize your recovery based on tiers:

Tier 1: These are the servers and applications required for business operations. Transactional systems and email are typically considered mission critical apps.

Tier 2: These systems and applications can usually stay dark for half a day or even a full day. For instance, a content library or an HR requisition system might be able to be down for a day without causing too much disruption.

Tier 3: These are the back-burner apps and data that you can go back and retrieve a few days after the lion's share of recovery work has been done.

Devise Your Recovery Strategy.

Now that you know exactly what you'll need to recover and when in every situation, you'll want to craft multiple recovery strategies. Depending on the type of disaster that strikes, you'll need an assortment of procedures and tactics, and need to call on different staff members.

This is the area of your plan that you'll test scrupulously later. For now, go through each asset for each risk and map out:

Current preventive controls: how are you preventing these threats?

Monitoring and detection systems: how will you be alerted when a disaster is brewing or a failure has already occurred?

First response and recovery strategies: what steps will you take to control the disaster and mitigate the damage?

Failover procedures: what is the process for failing over and who understands how to do it?

The system restart and failback process: how can you resume normal business operations? Who is trained in the process?

Define Staff Roles and Responsibilities.

Once you have an idea of which tasks need to be completed, it's time to identify which staff members will handle which process. Include everyone from your CIO to your project managers and any third-party partners and vendors. Each player will need to understand not only their role but the others' responsibilities to step in if the chain breaks down or the recovery process goes off track.

Once you've identified every key player in the process, create a backup for each person. Even critical leaders and engineers can go on vacation or leave for another job the day before disaster hits.

Create a Communication Plan.

Because situations are liable to change quickly during a crisis, a detailed communication plan is critical. You'll want to create a call chain that outlines which leaders, vendors and team members need to be notified and in what order; you'll also explain the steps for alerting everyone from the media to employees to clients.

Develop contingency plans for situations impacting your phone, sites, social media accounts or email; you'll need to make sure not only that you have an alternate communication channel but that your customers and employees know to check it.

While the types of communications will vary according to each organization's needs, the following should all be issued during any significant disaster. Whenever possible, prepare templates in advance that can be updated for each situation.

- A formal statement of disaster
- Initial communications with employees and customers at the beginning
- Ongoing status updates
- A brief statement to publish on social media and your site
- Downtime time estimates for customers
- A statement to the media
- Contact information to have questions answered

Test and Test Again.

Putting your DR plan through its paces is the best way to make tougher and more accurate. Even a meticulously crafted plan could turn out to be built on a shaky foundation; maybe you'll find out you're not snapshotting frequently enough for your important data or you realize some of the key players have changed their phone numbers. With so many moving parts in any BDR ecosystem, it's highly likely some changes have been missed.

The chief objective of testing is to confirm you can meet your RTO and RPO goals. From tabletop testing to drills, you'll want to regularly validate your ability in those two areas. Your methods may vary depending on the type of theoretical scenario you're testing against, but the below process can help you carry out an effective and low-impact test.

- Schedule the test on a date that won't disrupt an important business event or launch. Make sure all key players can attend.
- Calculate how much the test will strain your systems or impact vital workloads. If you can't find a way to test your entire plan safely, test one area at a time.
- Notify your IT staff several weeks in advance.
- Turn a conference room or other area into your war room.
- Document the step-by-step procedures of the test and hand them out to all players.
- Launch the test and instruct all players to perform their tasks.
- Time each component of the test.
- Identify the components that worked and failed, how long each phase took and which unexpected obstacles took the process off track.
- Update your plan based on those findings.

Creating a Safer Future

A smarter disaster recovery plan allows you to substitute the much more valuable foresight for the wise – but less actionable - hindsight. The ultimate casualty of any disaster is the loss of the profit and innovation that would have happened in its stead. Planning ahead does more than save money and reputations – it can staunch the damage and preserve the transactions, productivity and ideas that would have come to fruition.

Just as secure and tested backups can help you sleep better at night, a documented and tested plan can help your leadership and team feel more confident about the future. Disasters small and large may one day find their way to your door – but you'll be ready when they do.

Disaster Recovery Planning Resources

The ISO/IEC 27031:2011

A methodology framework for information and communication technology (ICT) for business continuity, including design, evaluation and implementation.

NIST Special Publication 800-34

Guidelines for government IT contingency planning.

Rule 4370 of the Financial Industry Regulatory Authority

Requirement for business continuity plans (BCPs) for financial services institutions.

The Business Continuity Institute Good Practice Guidelines (GPG)

Information on sustaining continuity.

Quorum Resources

Quorum onQ Ransomware Edition

Backup and Disaster Recovery Justification Spreadsheet

Waking Up from the Ransomware Nightmare

The Backup and Disaster Recovery Security Toolkit

The 5 Downtime Strategies You Need to Change