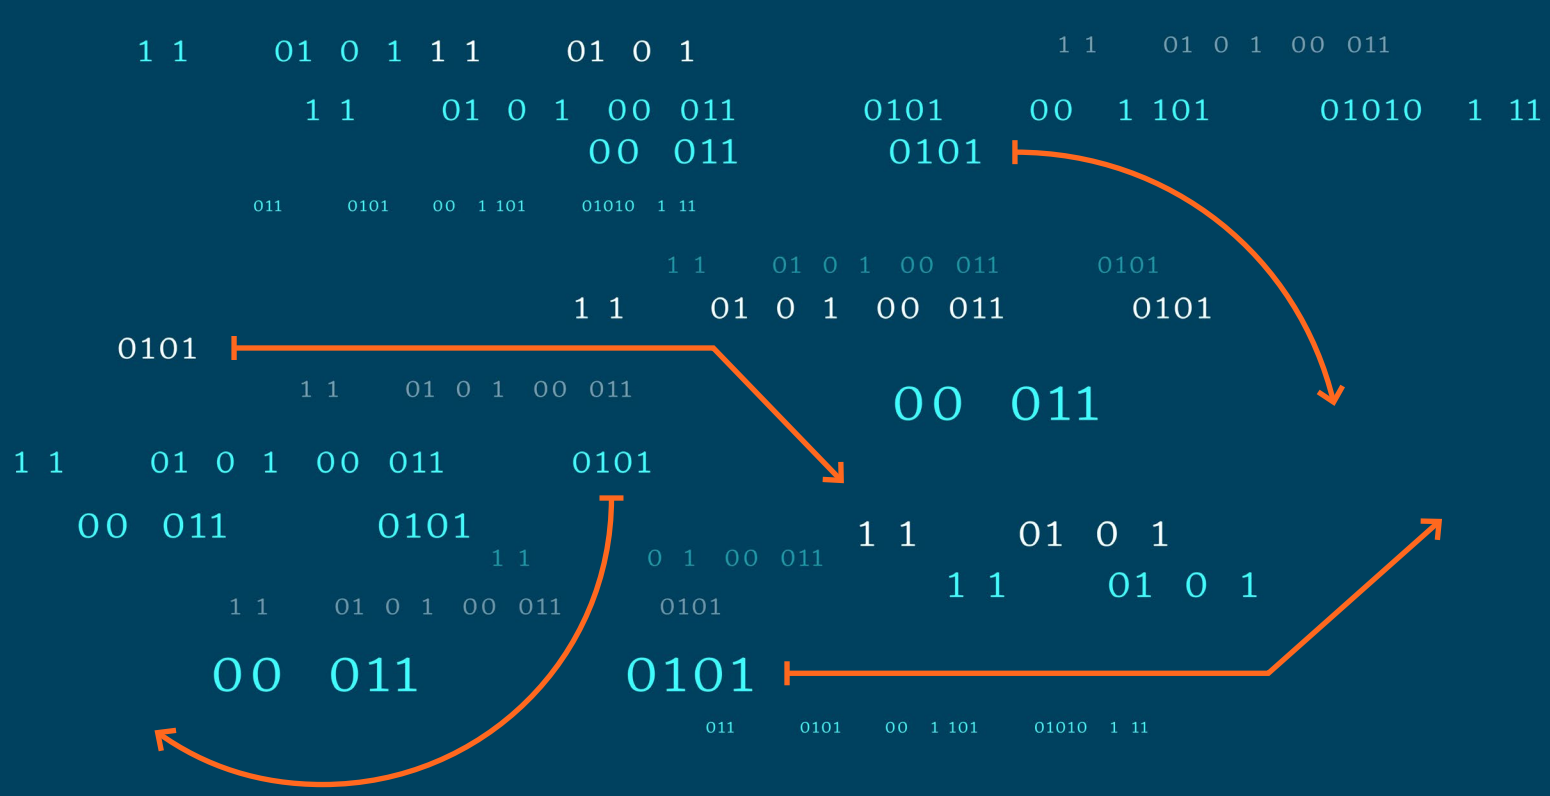

DISASTER RECOVERY

The Essential Recovery Playbook

WHAT IS YOUR TOLERANCE FOR DOWNTIME?

4 → 4 → 4 → Forever?

Minutes Hours Days

DATABANK
Data Center Evolved™

ROADMAP

1 Your Guide to Resilience

Preparing for a disaster can be as stressful as facing the disaster and managing its aftermath. As IT professionals, crafting a disaster recovery plan often falls on your shoulders—but where do you begin?

It's not just natural disasters you need to worry about. Your organization faces many risks, from cyberattacks and ransomware to software misconfigurations that could render your systems inaccessible. These threats could lead to significant downtime, data loss, and operational disruption—putting your organization at risk of financial losses, reputational damage, or even total business failure.

Business stakeholders believe they're covered. You've got this—right?

You may have onsite backups for quick recovery, offsite options for easy access, or cloud storage available anywhere, anytime. But what if those backups are sitting in the trunk of your car? More importantly, what will you recover those files to if you can't physically access your server room or connect over the network? Are offsite servers ready to go? Is there a live network at a secondary site to keep your users connected to critical applications? And do those users have secure access (VPN/MFA)?

If you've been in IT long enough, you know **it's not a matter of if, but when.**

That's why you need more than just a disaster recovery plan—an **executable plan**. A binder on a shelf might look reassuring, but printed plans are outdated the moment they're published and often provide little more than general guidance when disaster strikes. Your business truly needs a practical guide that ensures you're prepared—so you can attend your kid's baseball game or take that vacation, without being glued to your computer for fear of the unknown.

This template offers a straightforward framework to build a disaster recovery plan that's easy to execute. Each section details the purpose of a specific component, complete with tables for you to fill in, ensuring your recovery strategy tackles vital threats. With clear examples and practical guidance, this resource helps you safeguard critical data, maintain customer trust, and protect your business's reputation.

Being ready for a disaster makes all the difference—and a strong recovery plan is your key to bouncing back swiftly and smoothly.

2 Practical Approach to Hybrid Workloads

Most organizations operate with a mix of infrastructure setups. Some maintain their own data centers, while others have adopted the flexibility of public cloud services, and many rely on third-party colocation facilities for hosting critical applications and data. This hybrid approach allows companies to optimize costs, performance, and scalability, but it also introduces complexity regarding disaster recovery (DR). With workloads and data dispersed across different environments, a one-size-fits-all recovery plan is no longer sufficient. A tailored, multi-faceted strategy is essential to ensure resilience in the face of disruptions.

Understanding the Unique Challenges of a Hybrid Infrastructure

When a company's data and applications are spread across multiple platforms—onsite data centers, public clouds, and colocation facilities—each environment presents its challenges in terms of data protection, recovery speed, and business continuity:

Onsite Data Centers: Organizations that own their data centers often have direct control over their infrastructure, including physical security, hardware, and network configurations. However, these data centers can be vulnerable to local disasters such as power outages, natural disasters, or hardware failures. A robust DR plan for onsite systems should include offsite backups, redundancy, and a clear protocol for transferring operations to a secondary location if the primary site becomes unavailable.

Public Cloud: Cloud services offer unparalleled scalability and flexibility, allowing businesses to quickly adjust resources and storage needs. However, relying on a third-party provider also means depending on their availability and security measures. DR strategies for public cloud workloads should focus on replicating data across different regions or availability zones, establishing automated failover processes, and ensuring that snapshots and backups are regularly updated and easily accessible.

Colocation Facilities: Colocation provides a balanced solution where businesses control their hardware while using a third-party facility for essential services like power, cooling, and physical security. This approach offers a mix of control and outsourcing but comes with specific challenges, such as ensuring geographic diversity and having the infrastructure needed to recover applications and data. Many colocation providers offer solutions to address these challenges, including secondary site options - and the servers, networks, and services needed for recovery in case of failure. Additionally, some providers offer testing services, allowing businesses to test their failover and recovery plans efficiently.

Other Considerations, Tips and Traps Before You Get Started

Before designing a technical solution to what is ultimately a business challenge, it's crucial to address a few key considerations. A successful disaster recovery plan goes beyond technology and must account for the following:

1. **People:**
    - Access: Understand how employees access systems today and where they are located. Plan how they will access these systems during a disaster, especially if physical access is restricted.
    - Communication: Develop a communication strategy for when systems are down. Do you have a chat system or other tools that allow employees to receive status updates on their cell phones? A people plan is just as essential as your IT plan.
2. **Engage Business Stakeholders Early:**
    - Involve key C-level stakeholders, particularly the CFO, before designing a recovery plan. They manage risk within the organization and can help align the budget with the business's risk tolerance.
    - Gather input from department heads and other key stakeholders to understand what they believe is necessary for resilience. Be prepared to present the potential risks to the business and how your proposed plan can mitigate those risks effectively.
3. **Work Requirements Assessment:**
    - Identify what employees across different departments need to continue their work during a disaster. Some teams may need to be physically present in the office, while others can work from home.
    - For those typically in the office, ensure they have the necessary equipment to work remotely. This might require purchasing additional laptops, monitors, or other resources.
4. **Plan for Network Access:**
    - Recovering applications and data is only part of the solution—employees also need secure network access to reach those systems. This requires planning well in advance.
    - Remember that telecommunications providers may take 30-60 days to set up new services, especially during widespread disruptions when demand is high. Being proactive ensures you're not left waiting when time is critical.

By addressing these considerations before implementing a technical solution, organizations can create a disaster recovery plan that is both effective and aligned with business needs, ensuring a smoother recovery when disaster strikes.

3 The Lights Out Scenario

Consider this: On this sunny Monday morning, everything seems perfect. Sales are soaring, clients are happy, and business is booming.

But then, **everything goes dark.**

A sudden power outage hits your city, taking down your servers. Simultaneously, a cyberattack slips through your defenses, locking your employees out of their systems and encrypting all your data. The attackers demand a hefty ransom, threatening to expose sensitive customer information if you don't comply.

Panic sets in. Orders grind to a halt, your website is offline, and frustrated clients flood your customer service lines and social media with complaints. Your team scrambles for a solution, only to discover that your backup systems aren't working as expected.

Hours drag into days. Without access to systems, your employees are at a standstill. Sales plummet, and once-loyal customers lose trust. News of the breach spreads quickly, raising concerns from your CEO about lasting brand damage.

Financial losses mount, and you face the possibility of legal action or penalties for failing to meet customer SLAs.

Now, the disaster recovery plan that looked great on paper proves worthless—and your replication and failover plan was never fully implemented.

No matter how hard you try, you can't get your systems back online.

Will a lack of a reliable, executable disaster recovery plan force you to shut down operations for good?

> "93% of organizations have suffered a data-related business disruption in the last 12 months yet, only 29% of the surveyed IT and business leaders expressed complete confidence in their current disaster recovery technology to recover data."
>
> – IDC Research, "The State of Ransomware and Disaster Preparedness: 2022"

Sections 1 → 2 → 3 → Notes

4 Common Threats

System Misconfiguration: Unpatched systems, inadequate access controls, default account settings, and poor coding practices are just some common exposures.

Cyberattacks: Ransomware, malware, DDoS attacks, phishing, social engineering, and other cyber threats are increasingly sophisticated and can paralyze your IT systems, resulting in data breaches, financial loss, and reputational damage.

Human Error: Mistakes made by employees, such as accidental data deletion or misconfigurations, can cause significant disruptions, especially if backup retention is not appropriately set – or if those systems are not actively replicated to a secondary site.

Power Outages: Unexpected power failures can bring business operations to a halt, especially if backup power solutions are not in place. UPS is also important in preventing database servers from crashing and causing data corruption.

Hardware Failures: Hardware can fail without warning. Virtualization has solved some of these issues, so be sure you have high availability (HA) within your production environment to the extent possible in every tier of your infrastructure for your business-critical applications.

Cooling Issues: The equation is Power=Heat. Heat=Downtime, Barn fans are not the solution.

Not Having a Secondary Recovery Site: Not having a secure, reliable secondary data center location with servers, storage, and network elements to recover to in the event of a disaster is one of the most common risks. Consider your geodiversity risk.

Natural Disasters: We list this last because this is what most of us think of when the topic of disaster recovery comes up, but this is not the most significant risk to most businesses.

THE RISKS ARE REAL

Your risk is downtime due to ANY threat, typically the threats mentioned in this document other than natural disasters. These threats have been compounded by the on-demand and real-time nature of everything we do. Risks have increased exponentially over the past 10-20 years, so it's no longer and insurance policy discussion about buying a solution to protect against something that may never happen.

5 → Priority 1: Risk Level Assessment

A risk assessment is the most crucial aspect of developing a disaster recovery plan. The purpose is to identify possible events that could negatively impact your business operations. You can then determine the likelihood of occurrence and the impact (should the event occur) to assign the appropriate risk factor. After determining the risk level, your organization can decide what actions to take. Below are some essential terms to understand from the following template.

Threat: Consider the threats mentioned above and discuss other threats that may be unique to our business.

Probability: What is the likelihood that the threat/event would occur (generally expressed as a value between 0 and 1. Values closer to 0 indicate that the event is not likely to happen, whereas values closer to 1 indicate the event is highly likely to occur.

Impact: This rating suggests how significantly the threat/event would impact the operation of your business—generally expressed as a value between 0 and 1. A value closer to 0 indicates little to no impact on business operations, whereas a value closer to 1 indicates that the threat would be highly detrimental. Discuss these impacts with various departments to ensure that all operational departments are considered.

Risk Rating: The risk rating is derived by multiplying the Probability by the impact of that threat. The higher the risk rating, the larger the need for the business to act.

Recommended Action: Once you outline threats and establish risk ratings, the organization must decide on the next steps. Options could include Mitigation (taking actions to prevent the event from occuring), Transfer (moving the risk to a third party like an insurance company), Plan (identifying how to restore operations if the event occurred), or Accept. Accepting some risk is a normal part of business when the cost of mitigating, transferring, or planning exceeds the impact cost.

Threat	Probability	Risk Rating	Impact	Recommended Action *
Cyberattack	.7	.9	.63	Mitigate
Misconfiguration	.3	.7	.21	Plan
Natural Disaster	.2	.6	. 12	Plan
Power Outages				Plan
No Secondary Site				Transfer

*Chart examples shown in light blue.

6 Business Impact Analysis

Once you have determined which risks your business intends to plan for, you can examine the systems used for business processes. This helps define the plan's scope while setting goals and objectives. Evaluating the systems and business processes is done through a Business Impact Analysis. The first step in the business impact analysis is to take a system inventory and detail each system's functions, owners, users and components.

System Name	Business Process	System/Process Owner	System Users	System Components
Oracle	Supply Chain	VP of Business Systems	Operations, Purchasing	supply.prod, vendor.prod

*

After detailing the system inventory, your company needs to understand how the unavailability of that system or business process impacts your organization to measure the cost of downtime. The cost of downtime will later help your organization determine an appropriate budget for the disaster recovery plan and guide some of the essential metrics to measure the plan's success. The table below seeks to understand the monetary impact of each system. This example is calculated per hour based on lost or delayed sales, fines, contract breaches, lost productivity, and restoration costs.

System	Lost Sales	Delayed Sales	Fines	Contact Breach	Lost Productivity	Respore Expense	Other	Total Cost/Hour
CRM	$500	$250	--	--	$1000	$100	–	$1,850

*

*Chart examples shown in light blue.

By determining the impact of downtime for each system and/or business process, you can classify each system or method into a tier. Tiers are usually defined as follows:

Tier 0 (Life Critical): Any system where the unavailability can result in the loss of human or animal life (i.e., medical systems, safety systems, 911, etc.).

Tier 1 (Mission Critical): Any system where unavailability can result in significant financial losses and halts the operation of the business and service to customers (i.e., ERP, e-commerce websites, etc.).

Tier 2 (Business Critical): Any system where the unavailability can result in some financial losses and degrades the ability of the business to operate and serve its customers. (i.e., email, CRM, etc.).

Tier 3 (Non-Critical): Any system where unavailability does not impact the ability of the business to operate and serve its customers, but data should be protected from loss (i.e., file servers, data warehouses, etc.).

Based on the tiers, different systems should have different objectives for recovery during the disaster recovery plan. These objectives are illustrated and explained on the next page.

Based on the tiers, different systems should have different objectives for recovery during the disaster recovery plan. These objectives are illustrated and explained below.

Recovery Point Objective (RPO): The amount of data (measured in minutes or hours) that could be lost from the system during a disaster. Typically identifies the frequency at which data must be backed up/replicated to a recovery system.

Recovery Time Objective (RTO): The time from when the disaster occurs until the system needs to be back online and available to users.

Working Recovery Time (WRT): The time from when the system returns online until users can perform their regular job functions again.

Maximum Tolerable Downtime (MTD): Total amount of time from when disaster strikes until business is operational again.

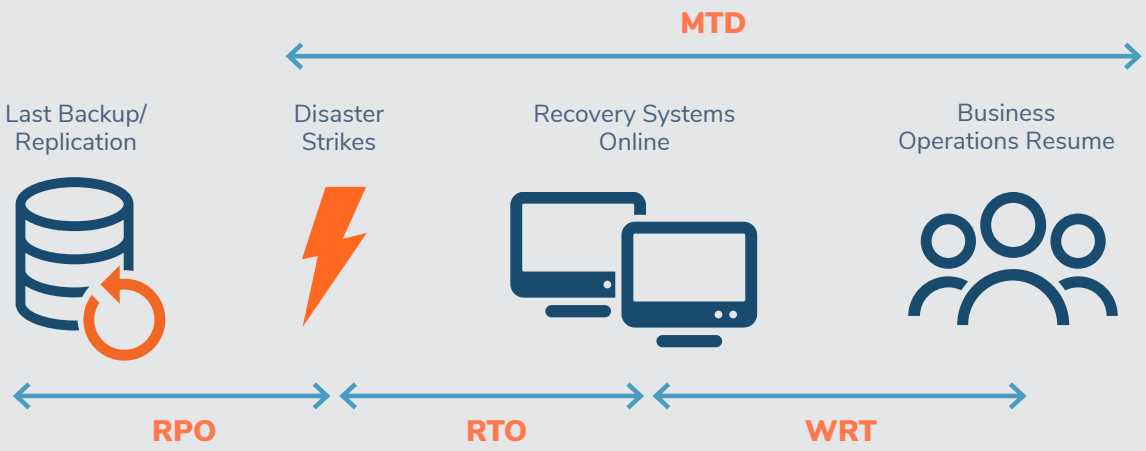

System Name	Tier	RPO	RTO	WRT	MTD
SAP	1	15 Min.	2 Hrs.	2 Hrs.	4 Hrs.

*

*Chart examples shown in light blue.

Sections 4 → 5 → 6 → Notes

7 Technology Selection

Once your organization has set objectives to be met by the disaster recovery plan, you will likely need to select technology or tools that provide a means to accomplish the goals. Many options are available to companies that meet various recovery time and point objectives.

Some of these technologies include:

Backups: Copies of data that are stored in a separate location from where the master data resides to provide for restoration of data following data loss or corruption.

Snapshots: Point-in-time record of a data volume that tracks changed bytes to provide for rollback in the event of data loss or corruption.

Log Shipping: Transmission of journals/logs from a primary system to a recovery system that allows for data reconstruction in the event of a failure of the primary system.

Storage Replication: Transmission of data volumes from a primary storage frame to a secondary frame that can be connected to standby servers or cloud platforms. This is used to activate systems in the secondary location in the event of a failure at the primary location.

Hypervisor Replication: Transmission of protected virtual machine images leveraging technology built into, or layered on top of, hypervisor platforms. Virtual machine images at a secondary location can be started following a failure of the primary location either manually or automatically.

Geo-clustering: Groups of servers load balanced across multiple active locations, allowing for the least disruption in the event of a failure at one of the active sites.

Your disaster recovery plan must identify how each system is protected, which technologies are leveraged, the frequency at which the protection occurs, and where protected data is transmitted.

System Name	Technology Used	Frequency/Schedule	Target
Oracle	Hypervisor Replication	5 minutes	Secondary Data Center
SAP	Backup	Daily Full	Tape/Vault

*

*Chart examples shown in light blue.

8 Your Disaster Recovery Team

Another critical part of a disaster recovery plan is identifying the people responsible for executing the plan. Each person must understand their role to prepare and perform properly. Keeping a contact list with all the internal and external parties will ensure that when the plan is invoked, it is easy to contact those involved.

Role	Name	Email	Primary Phone	Secondary Phone
Comms Lead				
Comms Backup				
Technical Lead				
Technical Backup				
C-Level/Mgt. Lead				
Customer Service Lead				
HR Lead				
Ops Lead				
Customer				
Finance Lead				
General Counsel				
Data Center Provider(s)				
Network Providers				
Software Providers				
Police				
Fire Department				
Remote Work Location				
Marketing Lead				
Customer Lead				

*

*Chart examples shown in light blue.

9 Emergency Communications

Not only must you maintain a list of disaster recovery team members, but you also need to ensure you have designated how the proper communication should flow if an event occurs. Below is an example of a communication plan that shows how you might task various parties by cascading communication with other members across your organization.

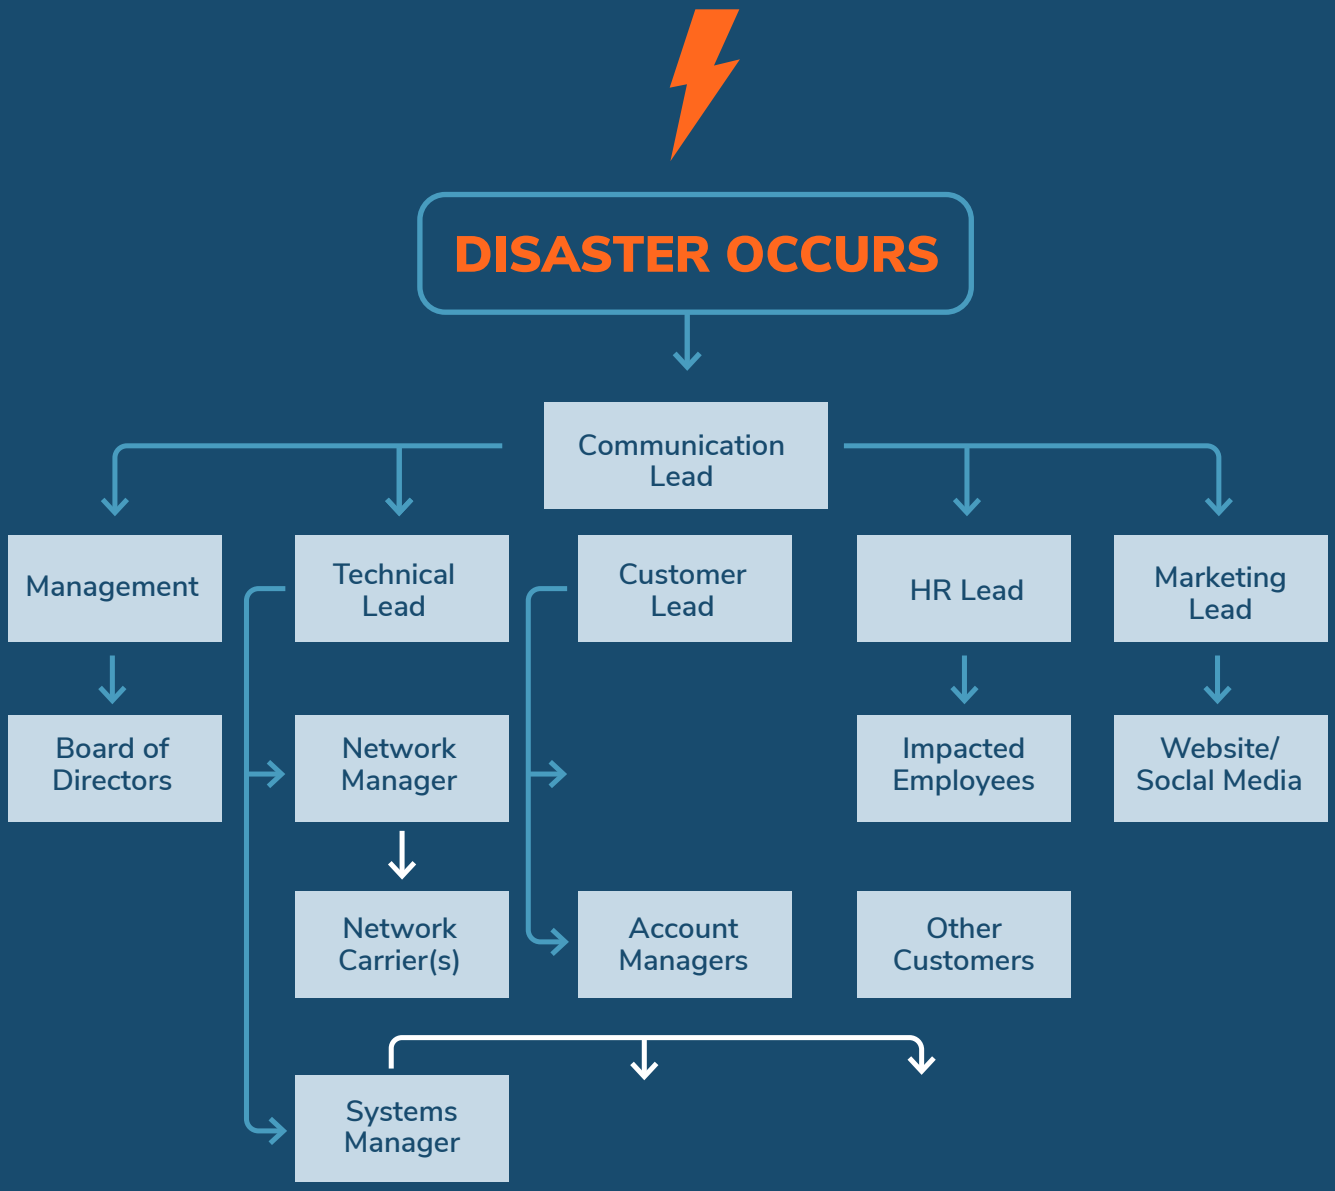

Sections 7 → 8 → 9 → Notes

10 System Recovery Plan

A detailed disaster recovery plan should document step-by-step instructions indicating the actions needed to restore systems. You may have a separate recovery plan for each major system. If that is the case, the dependencies between the systems need to be defined so that each system plan is executed in an appropriate order. The following is a framework for information that should be included for each step of your system recovery plan.

- Identify the stakeholders who will perform the step and those who have ownership over its success. Anyone who should be consulted and others who should be informed that the step has been executed needs to be documented.
- Detailed instructions regarding "how to execute the steps" are essential. The more detail, the better. It is possible that someone less familiar with the plan may need help with execution and may need more explicit directives.
- Describing how each step can be validated as successful is very useful. If the activity is unsuccessful, contingency plans are preferred, especially with dependencies between steps. A failure in a step early in the plan can cascade throughout the entire plan.
- RUNBOOKS: Mapping out your recovery plan on paper is an important part of any plan. However, these paper plans can become outdated quickly as system interdependencies change. It would be best to have an automated and tested failover plan to achieve an EXECUTABLE plan during a disaster. Disaster recovery runbooks are essential tools that streamline the management of your disaster recovery (DR) environment. They consist of Recovery Groups combined into a single framework, making executing failover tests or performing live failovers during a disaster more manageable. These runbooks are pre-configured, automated, online software tools that include systems needing step-by-step or simultaneous failover, even if they belong to different recovery groups, so you can quickly choose the appropriate systems for setup and testing. By utilizing disaster recovery runbooks, you can significantly reduce the time and effort required for failover management, providing an essential layer of automation to your disaster recovery strategy. Once you have these runbooks in place, you can rest comfortably so that your organization can recover when the time comes.

Note: The expected execution time of each step (when totaled) will help you determine if your plan will meet the RTO goal specified in the overall plan.

Step #1	**Establish Connectivity to the Secondary Data Center**		
Responsible: Network Engineer	Accountable: Network Manager	Consulted: Network Carrier	Informed: Systems Engineer Storage EngineerIS Manager
Detals	Reconfigure IPSEC tunnel from primary data center peer IP address to IP address of secondary data center		
Valadation	Ping the address of the router at the secondary data center.		
Contingency	Leverage software VPN client		

*

Step #	**Task**		
Responsible:	Accountable:	Consulted:	Informed: Systems Engineer
Details			
Valadation			
Contingency			
Expected Execution Time:			

*

Step #	**Task**		
Responsible:	Accountable:	Consulted:	Informed: Systems Engineer
Details			
Valadation			
Contingency			
Expected Execution Time:			

*

*Chart examples shown in light blue.

11 → Testing the Plan

Testing is the most critical part of having an executable plan. Without testing, you are relying on hope. As you've often heard, hope is not a strategy. The testing frequency depends on how frequently you change your systems, employee turnover in key areas, and the timing of your specific compliance requirements. It would help if you targeted at least once per year. If you genuinely want an executable plan that will be valuable when you need it, a "set it and forget it" mentality won't suffice. If you are like many organizations, systems, and processes change so fast that you barely have time for production documentation – but keeping production and recovery procedures up to date is paramount to achieving success.

Options and Considerations for Testing:

Testing Frequency: Most companies decide to test their plan annually at a minimum. The frequency at which you should perform a test depends on how dynamic your IT environment is. The more your environment changes, the more frequently you should initiate a test of the plan to ensure it is executed as expected if a significant outage occurs.

Planned or Ad-Hoc: Most companies choose to perform planned and controlled tests of their disaster recovery. This is scheduled in advance, and all involved parties know the schedule. The challenge with planned tests is that they do not simulate a disaster. Ad-hoc tests, on the other hand, are tests in which the schedule is determined only by the leader of the disaster recovery team ideal. Since most employees executing a disaster recovery plan will not expect the activity in advance, these results more accurately reflect what the business should expect.

Active or Passive: Most companies perform passive tests of their disaster recovery plan where technology allows. In a passive test, the production systems remain online, and the recovery systems are brought online in an isolated environment and then tested against production. The problem with a passive test is that it does not emulate an actual disaster scenario. Because the production system remains online, the potential exists for unknown dependencies to be satisfied. An active test where the production systems are shut down and then the recovery systems brought online provides a more accurate test result.

Step #1 Type of Test	Schedule	Frequenchy
Active	Planned	Annual
Passive	Ad-Hoc	B-Annual

*

When actual tests, it is essential to document the results for two key reasons: first, to verify that the plan is achieving its original objectives, and second, to identify areas for improvement. This table provides a structured way to record the execution of your disaster recovery plan, helping you accomplish both goals. significant outage occurs.

Step #1 MM/DD/YYYY Expected Time	Actual Time	Errors	Remediation
15 minutes			

*

Disaster Recovery Made Easy

Without a robust business continuity plan, your organization faces significant risks, including guaranteed revenue loss and potential operational failure. Taking proactive steps to evaluate and benchmark your current plans is crucial, and ensuring that your executives and IT professionals are prepared. This guide catalyzes and jumpstarts that essential process.

If building and executing a comprehensive disaster recovery solution feels overwhelming, don't hesitate to contact our qualified engineers and professionals. Partnering with a service provider like DataBank can make all the difference. We collaborate closely with you to design a customized disaster recovery plan that is thoroughly tested and ready to safeguard your mission-critical applications.

Your disaster recovery plan is your ultimate insurance policy against unforeseen events. With a solid plan in place, you'll protect your data and business, save money, and preserve customer trust.

*Chart examples shown in light blue.

12 Next Steps

Ready to Secure Your Business?

Take the next step in safeguarding your organization with a complimentary consultation from one of our expert data protection solutions engineers. Let us help you assess your current disaster recovery strategy and tailor a plan that meets your unique needs.

Reach out today and ensure your business is prepared for whatever comes next!

Secure Your Future. **Schedule a Consultation Today**

About DataBank's Managed Services

Free Up Your Staff and Extend Your Capabilities

DataBank has strategically assembled a suite of managed services and expertise to target the most challenging and time-consuming tasks that typically burden your IT resources. Leverage our Disaster-Recovery-as-a-Service capabilities to diversify your workloads and operations across multiple data centers and metros. Tap into our Storage-as-a-Service platform to scale your most data-intensive applications.

DRaaS: Continuous Availability for Mission-Critical Workloads

Ensuring the continuous availability of data and applications doesn't have to be an expensive, time-consuming process. There's no need for your technical resources to set up and manage a duplicate production system in a secondary data center or rely on backups at an offsite location. DataBank DRaaS is an effective, easily executable DR service that protects your IT environment, enabling your business to resume operations quickly in the event of a disruption. DataBank DRaaS enables setting your own RTOs and RPOs. We help businesses achieve near active/active outcomes for a fraction of the cost, backed by an expert engineering team, proven methodologies, compliance, security, 24x7x365 support, and an SLA that meets and exceeds your uptime requirements.

Sections 10 → 11 → 12 → Notes