DISASTER RECOVERY Planning Workbook

Preparing for a disaster can be just as stressful as dealing with an actual disaster and its impact on your business. This template is designed to provide a simple framework under which you can develop a successful disaster recovery plan that is easy to execute for your business. Each section provides information about the purpose of that particular component, and a table for your business to complete and satisfy the objective of that component. Examples are provided to guide you in understanding the types of data that should be input.

RISK LEVEL ASSESSMENT

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

Threat	Probability	Impact	Risk Rating	Recommended Action
Hurricane	0.3	0.9	0.27	Plan
Misconfiguration	0.7	0.8	0.56	Mitigate

Threat: These are potential events (including natural disasters, security risks, and human errors) that your organization has identified can impact business operations.

Probability: This 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 occur; whereas values closer to [1] indicate the event is highly likely to occur.

Impact: This is a rating that 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 there would be little to no impact to business operations; whereas a value closer to [1] indicates that the threat would be extremely detrimental.

Risk Rating: The risk rating is a value derived by multiplying the Probability by the Impact of the particular threat. The higher the risk rating, the larger the need for the business to take action.

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

SYMMETRY

BUSINESS IMPACT ANALYSIS

Once you have determined which risks your business intends to plan for, you can begin examining the systems used for business processes. This helps to define the scope of the plan 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 the functions, owners, users and components of each of the systems.

System Name	Business Process	System/ Process Owner	System/ Process Users	System Components
SAP	Supply Chain Management	VP of Business Systems	Line Technicians, Finance, Human Resources	ECC01.prod, Solman.prod, BW.prod, bobj.prod

After detailing the system inventory, your company then needs to understand how the unavailability of that system or business process impacts your company in an effort 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 important metrics to measure the success of the plan. The table below seeks to understand what the impact is monetarily for each system. This example is measured per hour and based on things such as lost or delayed sales, fines and contract breaches, lost productivity and restore costs.

System	Lost Sales	Delayed Sales	Fines	Contract Breach	Lost Productivity	Restore Expenses	Other	Total Cost/Hour
CRM	$500	$250	-	-	$1000	$100	-	$1850

By determining the impact of downtime for each system and/or business process you can now classify each system or process 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 for 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 for 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 set for recovery during the disaster recovery plan. These objectives are illustrated and explained below:

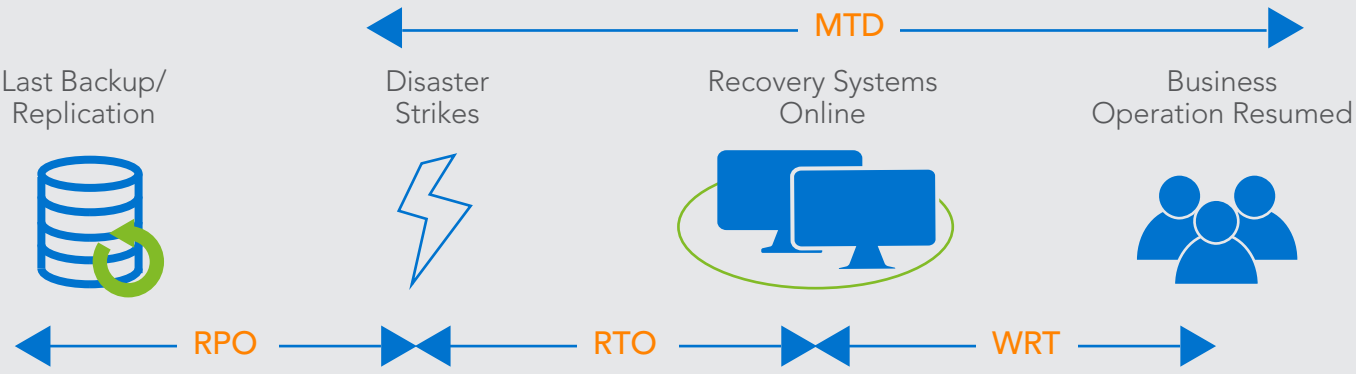

Recovery Point Objective (RPO): The amount of data (measured in minutes or hours) that could be lost from the system in the event of a disaster. Typically identifies the frequency at which data must be backed up/replicated to a recovery system.

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

Working Recovery Time (WRT): The amount of time from when the system comes back online until users are able to perform their regular job functions again.

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

System	Tier	RPO	RTO	WRT	MTD
SAP	1	15 minutes	2 hours	2 hours	4 hours

SYMMETRY

TECHNOLOGY SELECTION

Once the organization has set objectives that are to be met by the disaster recovery plan, you will likely need to select technology or tools that provide a means by which to accomplish the objectives. There are many options available to companies that meet various different 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 reconstruction of data 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.

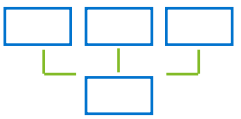

Hypervisor Replication: Transmission of protected virtual machine images leveraging technology built into, or layered on-top of, hypervisor platforms. Virtual machine images at a secondary location can be started following a failure of the primary location either manually or automatically.

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

It is important that your disaster recovery plan identify how each system is protected, which technologies are leveraged, the frequency at which the protection occurs and where protected data is transmitted to.

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

SYMMETRY

DISASTER RECOVERY TEAM

Another critical part of a disaster recovery plan is identifying the people responsible for executing the plan. Each person must understand his or her role to properly prepare and execute. Keeping a contact list with all of the internal parties as well as the external parties will ensure that when the plan is invoked it is easy to get a hold of those involved.

Role	Name	Email	Primary Phone	Secondary Phone
Communication Lead				
Comm. Backup				
Technical Lead				
Technical Backup				
Customer Lead				
Customer Backup				
Human Resource Lead				
HR Backup				
Marketing Lead				
Marketing Backup				
Data Center Provider(s)				
Network Provider(s)				
Software Provider(s)				
Police				
Fire Department				
Remote Work Location				

EMERGENCY COMMUNICATION

Not only is it imperative that 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 actual event occurs. Below is an example of a communication plan that shows how you might task various parties with cascading communication to other members across your organization.

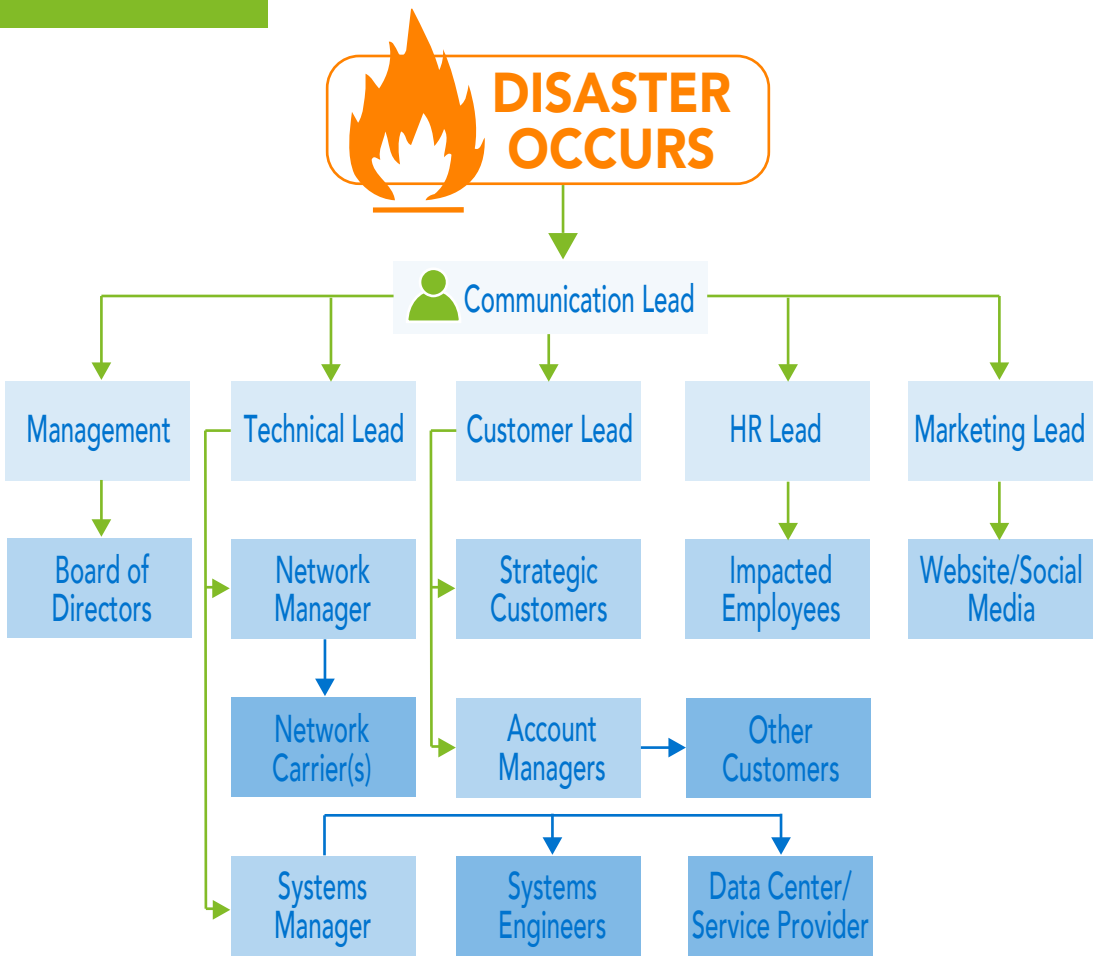

SYSTEM RECOVERY PLAN

A detailed disaster recovery plan should document step-by-step instructions indicating the actions that need to be executed in order to restore systems. It is possible that 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 not only will perform the step, but also those that have ownership over its success. Anyone who should be consulted and others who should be informed the step has been executed need to be documented.

- Detailed instructions regarding "how to execute the step" are essential. The more detail the better - it is possible that someone less familiar with the plan may need to help with execution and may need more explicit directives.

- Particulars around how the step can be validated as successful are very useful. In the event that the activity is not successful contingency plans where possible are preferred especially with dependencies between steps. A failure in a step early in the plan can cascade throughout the entire plan.

• Expected execution time of each step (when totaled) will help you determine if your plan will meet the RTO goal that you specified in the overall plan.

Step #1	Establish Connectivity to Secondary Location		
Responsible: Network Engineer	Accountable: Network Manager	Consulted: Network Carrier	Informed: Systems Engineer, Storage Engineer
Details	Reconfigure IPSEC tunnel from primary data center peer IP address to IP address of secondary data center (x.x.x.x)		
Validation	Ping the address of the router at the secondary data center (x.x.x.x)		
Contingency	If IPSEC tunnel cannot be configured leverage software VPN client.		
Expected Execution Time: 15 minutes			

Step #	Action		
Responsible:	Accountable:	Consulted:	Informed:
Details			
Validation			
Contingency			
Expected Execution Time:			

Step #	Action		
Responsible:	Accountable:	Consulted:	Informed:
Details			
Validation			
Contingency			
Expected Execution Time:			

Step #	Action		
Responsible:	Accountable:	Consulted:	Informed:
Details			
Validation			
Contingency			
Expected Execution Time:			

SYMMETRY

TESTING THE DR PLAN

The most central, yet often overlooked, portion of a disaster recovery plan is the testing. Many organizations prescribe to the "set it and forget it," mentality of disaster recovery. A good disaster recovery plan includes how and when it will be tested. Here are some of the options and considerations for testing.

Testing Frequency: Most companies decide to test their plan annually at a minimum. The frequency that 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 true disaster where to occur.

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

Active or Passive: Most companies choose to 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 fashion then tested against. The problem with a passive test is that it also doesn't emulate an actual disaster scenario. Because the production system remains online, potential exists for unknown dependencies to be satisfied. An active test where the production systems are first shut down and then the recovery systems brought online provide a more accurate test result.

Type of Test	Schedule of Test	Frequency
Active	Planned	Annual
Passive	Ad-Hoc	Bi-Annual

When actual tests are being performed it's important to document the results of the test to 1) ensure that the plan is meeting the original goals established for it and 2) identify ways that the plan could be improved. This table articulates a way that you can record the execution of your disaster recovery plan in order to meet these ends.

	Execution Time			
Step	Expected	Actual	Errors	Remediation
1	15 minutes	20 minutes	Couldn't ping server	Yes

SYMMETRY

Disaster Recovery Made Easy

If your organization doesn't have a great business continuity plan the repercussions can range from guaranteed revenue loss to potential failure. Taking the steps outlined in this guide to prepare your executives and IT professionals with careful evaluation and benchmarking of your current business continuity plans will jumpstart the process.

However if the thought of building and executing a comprehensive disaster recovery solution is overwhelming, you can work directly with a service provider, like Symmetry, to build the disaster recovery plan you need. At Symmetry, we work hand-in-hand with you to create a tailor-made disaster recovery plan that is tested and ready to protect your mission-critical applications.

Your disaster recovery plan is the ultimate insurance policy for your data and business should something out of your control happen. With a solid disaster recovery plan in place you'll save money, save your customers, and ultimately save your business.

About Symmetry:

Symmetry Corporation is a leading applications management and hybrid cloud hosting solution provider. An SAP certified partner since 2005, Symmetry is certified in SAP Hosting, Cloud and SAP HANA® Operations. As a true extension of your team, Symmetry places a laser focus on our customer's experience and is one of the only managed IT providers with flexible solutions built to meet each customer's unique business needs. Headquartered in Milwaukee, WI., Symmetry supports global customers through our 24/7 operations support model and its extensive worldwide datacenter network. With a proven methodology for delivering technical managed services and complete hosting solutions, Symmetry delivers flexible, high-quality solutions that help reduce the total cost of ownership and enable high-performing and secure environments of customers' most mission critical systems.

For more information please contact our business development department at:

www.SymmetryCorp.com | salesinfo@symmetrycorp.com | 888-796-2677