COMMONd

CYBER SECURITY
NETWORK ARCHITECTURE
REGULATORY COMPLIANCE

Disaster Recovery Planning Template
March 2020

Critical Asset Identification in IT Disaster Recovery

Disaster Recovery (DR) depends upon the ability to recover damaged assets because assets are needed to recover function, and restoring functionality is the goal of DR.

Business Continuity and Disaster Recovery (BC/DR) planning is the process of developing the plans, processes and procedures to respond to the range of incidents. We start with understanding the essential functions of an organization, called Business Impact Analysis (BIA). In life, we set the same priorities: protection of family and friends, shelter, food and water and other life-giving essentials.

Having performed the BIA and identified the critical assets to support our business or our lifestyles and priorities and budget, it's time to get to work protecting those things. We also need to understand that we'll need plans to respond to incidents, and that we'll need to execute them at some point.

IT assets include: (1) Hardware, (2) Software, (3) Data and (4) People. Each of these assets has a financial worth, but together they're worth much more than that. Combinations of these things gives you the functionality you need. For example, networking and communications requires a combination of all four.

Risk Assessment

The first step of a risk assessment is to assess and identify all possible threats as well as their likelihood of impacting your business.

Once you've analyzed the potential risks, it's time to create a business impact analysis (BIA). This helps you predict the consequences of disruption and gathers data needed to develop various recovery strategies.

Know Your Critical Assets

When building a DR strategy, the first step is to identify and understand what you need to protect. In most organizations, this means identifying critical assets—assets that impact confidentiality, integrity, and/or availability and support the business mission and functions.

Critical assets can include patents/copyrights, corporate financial data, customer sales information, human resource information, proprietary software, scientific research, schematics, and internal manufacturing processes.

For example, an ecommerce business might identify its website, inventory system, sales and accounts receivable system, any proprietary products it produces, and interfaces with delivery systems, either electronic or physical.

You can identify critical assets using different methods, including risk assessments, asset tracking through a service or hardware inventory, and network traffic monitoring that reveals the most frequently used network(s) and system components.

Once you identify your critical assets, you must determine which ones are at the most risk of being attacked by authorized insiders and how these assets should be protected and monitored. From an insider threat perspective, for each critical asset, risks should be identified such as privileged users, employees, contractors, trusted business partners, and others. The insider threat team works in collaboration with other parts of an enterprise (e.g., human resources, risk management, information technology, legal, etc.) to identify high-risk users who most often interact with these assets.

To protect critical assets, mitigation strategies are prioritized and implemented to ensure the highest value assets have the most comprehensive security. Actions include putting appropriate configurations, controls, training, and defenses in place. Often protections for critical assets also provide protections for other assets within the enterprise.

Although identifying critical assets is directly tied to an insider threat program, the asset inventory and tracking are not usually done by the insider threat team. Critical asset identification is usually done by a risk management group or similar team. Working with the critical asset owners, the risk or inventory team ensures it has the most up-to-date information about the assets. This information then needs to be passed to the insider threat team in a timely manner.

Identifying your assets is not easy. It takes knowledge, funding, and resources to collect information, conduct the inventory, and keep it current. Failing to follow this practice can result in the inadequate protection of key resources, delayed response to critical breaches or data exfiltration, and impediments to mission success.

RTO and RPO

RTO and RPO (Recovery Time Objective and Recovery Point Objective) are two key metrics that organizations must consider in order to develop an appropriate disaster recovery plan that can maintain business continuity after an unexpected event.

Although only one letter separates RTO from RPO, it’s important not to confuse or conflate these two metrics. Both help to determine maximum tolerable hours for data recovery, how often data backups should occur and what your recovery process should be. Both need to be considered when creating a disaster recovery plan.

RTO vs. RPO:

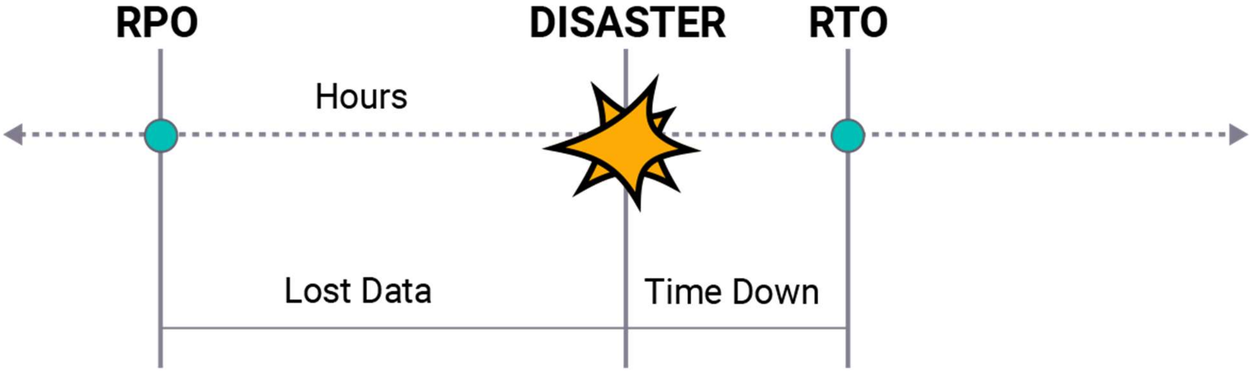

RTO

RTO stands for Recovery Time Objective. It’s a metric that helps to calculate how quickly you need to recover your IT infrastructure and services following a disaster in order to maintain business continuity.

RTO is measured in terms of how long your business can survive following a disaster before operations are restored to normal. If your RTO is twenty-four hours, it means you’ve determined that the business can maintain operations for that amount of time without having its normal data and infrastructure available. If data and infrastructure are not recovered within twenty-four hours, the business could suffer irreparable harm.

RPO

RPO, or Recovery Point Objective, is a measurement of the maximum tolerable amount of data to lose. It also helps to measure how much time can occur between your last data backup and a disaster without causing serious damage to your business. RPO is useful for determining how often to perform data backups.

RPO is significant because in most cases, you will likely lose some data when a disaster occurs. Even data that is backed up in real-time has a risk of some data loss. Most businesses back up data at fixed intervals of time -- once every hour, once every day or perhaps as infrequently as once every week. The RPO measures how much data you can afford to lose as the result of a disaster.

For example, imagine that you back up your data once every day at midnight and a disaster occurs at eight in the morning. In that case, you would lose eight hours’ worth of data. If your RPO is twenty-four hours or longer, you’re in good shape. But if your RPO is, say, four hours, you're not.

Differences Between Recovery Objectives

RTO and RPO are both business metrics that can help you calculate how often to perform data backups. However, there are some key differences:

Assessment basis. RTO reflects your overall business needs. It’s a measure of how long your business can survive with IT infrastructure and services disrupted. In contrast, RPO is just about data. It determines how often to back up data and does not reflect other IT needs.

Cost relevance. The costs associated with maintaining a demanding RTO may be greater than those of a granular RPO. That’s because RTO involves your entire business infrastructure, not just data.

Automation. Meeting your RPO goals simply requires you to perform data backups at the right interval. Data backups can easily be automated, and an automatic RPO strategy is therefore easy to implement. RTO, on the other hand, is more complicated because it involves restoring all IT operations. It is virtually impossible to achieve RTO goals in a completely automated way (although you should automate as much of your recovery process as possible).

Ease of calculation. In some ways, RPO is easier to implement because data usage is relatively consistent and there are fewer variables. Because restore times involve your entire operation, not just data, it is more complicated. Restore times can change based on factors such as the time of day or the day of the week at which a disaster occurs. The RTO must be aligned with what is possible by the IT organization. If the minimum restore time possible is 2 hours, then an RTO of 1 hour will never be met. IT administrators must have a good understanding of the speeds with which different types of restores can take place. Only then can an RTO be properly negotiated and met based on the needs of your organization.

Critical Assets

1.	
2.	
3.	
4.	
5.	
6.	
7.	
8.	
9.	
10.	
11.	
12.	
13.	
14.	
15.	
16.	
17.	
18.	
19.	
20.	
21.	
22.	

Potential disasters have been assessed as follows:

Potential Disaster	Probability (1-5) 1=Very High 5=Very Low	Impact Rating (1-5) 1=Total Destruction 5=Minor Annoyance	Brief Description Of Potential Consequences & Remedial Actions
Ransomware			
Hacking			
Cyber Crimes			
Data Breach			
Theft			
Act of Terrorism			
Act of Sabotage			

Tornado			
Electrical Storms			
Power Failure			
Loss of Communication Systems			

Probability:
1=Very High   5=Very Low

Impact:
1=Total destruction, 5=Minor annoyance

Plan Triggering Events

Key trigger issues at headquarters that would lead to activation of the DRP (Disaster Recovery Plan) are:

Plan Triggering Events	
Ransomware	**1.** **2.** **3.** **4.**
Hacking	**1.** **2.** **3.** **4.**
Cyber Crimes	**1.** **2.** **3.** **4.**
Data Breach	**1.** **2.** **3.** **4.**
Theft	**1.** **2.** **3.** **4.**

Act of Terrorism	**1.** **2.** **3.** **4.**
Act of Sabotage	**1.** **2.** **3.** **4.**
Tornado	**1.** **2.** **3.** **4.**
Electrical Storms	**1.** **2.** **3.** **4.**
Power Failure	**1.** **2.** **3.** **4.**
Loss of Communication Systems	**1.** **2.** **3.** **4.**

Activation of Emergency Response Team

When an incident occurs, the Emergency Response Team (ERT) must be activated. The ERT will then decide the extent to which the DRP must be invoked.

1. Respond immediately to a potential disaster and call emergency services
2. Assess the extent of the disaster and its impact on the business and data center
3. Decide which elements of the DRP should be activated;
4. Establish and manage the disaster recovery team to maintain vital services and return to normal operation
5. Ensure employees are notified and allocate responsibilities and activities as required

Disaster Recovery Team

The team will be contacted and assembled by the ERT. The team's responsibilities include:

1. Establish facilities for an emergency level of service within ___ business hours
2. Restore key services within ___ business hours of the incident
3. Recover to business as usual within ___ to ___ hours after the incident
4. Coordinate activities with disaster recovery team, first responders, etc.
5. Report to the emergency response team

Disaster Procedures for Management

Members of the management team will keep a hard copy of the names and contact numbers of each employee in their departments. In addition, management team members will have a hard copy of the company's disaster recovery and business continuity plans on file in their homes in the event that the headquarters building is inaccessible, unusable, or destroyed.

Contact with Employees

Managers will serve as the focal points for their departments, while designated employees will call other employees to discuss the crisis/disaster and the company's immediate plans. Employees who cannot reach staff on their call list are advised to call the staff member's emergency contact to relay information regarding the disaster.

Backup Staff

If a manager or staff member designated to contact other staff members is unavailable or incapacitated, the designated backup staff member will perform notification duties.

Alternate Recovery Facilities / Hot Site

If necessary, the hot site at _______________ will be activated and notification will be given via recorded messages or through communications with managers. Hot site staffing will consist of members of the disaster recovery team only for the first 24 hours, with other staff members joining at the hot site as necessary.

Personnel and Family Notification

If the incident has resulted in a situation which would cause concern to an employee's immediate family such as hospitalization of injured persons, it will be necessary to notify their immediate family members quickly.

Insurance

As part of the company's disaster recovery and business continuity strategies a number of insurance policies have been put in place. These include errors and omissions, directors & officers liability, general liability, and business interruption insurance.

If insurance-related assistance is required following an emergency out of normal business hours, please contact: ________________________

Policy Name	Coverag e Typ	Coverage Period	Amount Of Coverage	Person Responsible for Coverage and Contact Info.	Next Renewal Date

Legal Actions

The company's legal department and ERT will jointly review the aftermath of the incident and decide whether there may be legal actions resulting from the event: in particular, the possibility of claims by or against the company for regulatory violations, etc.

DRP Testing

Disasters don't occur very often, but when they do, the effects can be devastating. Your DRP should be in place to guarantee continued viability in case of a catastrophic event. Your business will continue, processes will continue, and your DRP will work.

Your business will operate, and your support and critical assets will go back online with minimal downtime. Without regular testing, your DRP is just a book of paper. Ongoing testing is a must.

At a minimum, the following tests will need to be considered and executed:

1. Test your RPO and RTO (Recovery Time Objectives and Recovery Point Objectives) bench mark. These tests need to make sure you're reaching your objectives while also detailing the processes that account for success.

2. Simulation: Your team can go through a simulated disaster to identify whether emergency response plans are adequate. You can test for:
    a. Ransomware
    b. Cyber attack
    c. Thief
    d. Data breach

3. System/Hardware Failover Test: Your system should have an automatic failover built in. I would go to the data center and ask my IT Administrator: What would happen if I physically remove one of the hard drives from one of the mission critical servers? The answers you will receive from the IT Administrator and their faces are priceless! Your system should be designed to handle this issue and still carry the full production workload.

4. Data Integrity Test: This test will make sure that you have 100% data integrity.

5. Complete Cutover Test: This is to make sure that your secondary infrastructure will perform in the event of complete shutdown of your primary infrastructure.

There’s no standard suggestion for how often you should conduct a full DR readiness test, but twice yearly is a good place to start. Additionally, it is important to conduct testing following changes to your environment. The scope of these tests will depend on the changes made and may not require testing every aspect of your disaster recovery plan.

“Set it and forget it” is another approach to disaster recovery. Organizations assume that their plans will work. This is an interesting but dangerous assumption. Unfortunately, many businesses fall within this category.

***** *This document is intended to be a starting point for a potential disaster recovery plan and is a not, in any way, a complete disaster recovery plan.*

COMMONd does not take responsibility for any disaster recovery plans that are based off of this document, nor any failure due to implementation or execution of this template.