carahsoft.

Checklist for Your IT Disaster Recovery Plan

The key to managing IT disruptions effectively is preparation

Thank you for downloading this checklist. Carahsoft represents proven DevSecOps solutions, delivering agencies the innovative solutions needed for every phase of the DevOps and DevSecOps lifecycles and with security built-in every step of the way. These solutions provide support for collaborative planning, rapid code builds, iterative testing, rapid release, optimized deployment and ongoing monitoring that continuously feeds into the next wave of planning.

Carahsoft combines extensive knowledge of the technologies we provide with a thorough understanding of the government procurement process to analyze needs, provide configuration support, simplify the ordering process, and offer special government pricing. Speak to a Carahsoft representative today about achieving your DevSecOps objectives.

For more information, contact Carahsoft or our reseller partners:
carah.io/solve/devsecops |devsecops@carahsoft.com | (703) 871-8570

Checklist for Your IT Disaster Recovery Plan

The key to managing IT disruptions effectively is preparation

Even the most prepared companies can experience IT disruptions — caused by application failures, cyberattacks, or human error — that lead to downtime and data loss.

When this happens, engineering teams are dispatched to repair the damage, while public relations (PR) teams work overtime to restore customer confidence. It is a time-consuming and often expensive effort. No matter what the cause of the disaster, the organizations that manage them most effectively, and with the least amount of collateral damage, are those with a comprehensive, easy-to-follow, and regularly tested disaster recovery (DR) plan. Whether you already have a DR plan in place or you are just beginning the process of creating one for your organization, you can use this checklist to identify key components of a successful plan.

Disaster Recovery Plan Checklist

Actions	✓		✓
Determine Recovery Objectives (RTO and RPO)	○	Define Incident Response Procedure	○
Identify Stakeholders	○	Define Action Response Procedure and Verification Process	○
Establish Communication Channels	○	Perform Regular Testing	○
Collect All Infrastructure Documentation	○	Stay Up to Date	○
Choose the Right Technology	○	Prepare for Failback to Primary Infrastructure	○

Determine Recovery Objectives (RTO & RPO)

The main goal of DR is to keep your business operating as usual, all the time. This means you need to determine which workloads are the most mission-critical to your organization, and what recovery time objective (RTO) and recovery point objective (RPO) are required for these workloads.

RTO is the amount of time required to recover from a disaster after notification of business disruption.
A reliable DR plan contains a clearly stated allowable RTO for each application group. If your business cannot withstand an hour of downtime without losing customers to competitors or paying penalty fees due to service-level agreements (SLAs), it is critical to your business to be operational before an hour has expired. In this case, your RTO would be one hour.

RPO is the window of time in which data loss is tolerable.
If your business can only withstand four hours of data loss and you currently perform only nightly backups, you would have a catastrophic loss of important data if disaster strikes in the afternoon. In this case, your RPO would be four hours.

A company's RTO and RPO will affect its DR strategy as well as associated expenses. While a simple file-level backup system might be sufficient for some applications, your mission-critical applications will likely need a DR solution with continuous data replication and rapid recovery to enable you to keep your business running and achieve minimal RPOs and RTOs.

Identify Stakeholders

Identify all those who need to be updated once disaster strikes. In addition to stakeholders involved in performing the actual recovery from a disaster (such as **engineers, technical support, and executives**), you should also pinpoint members of your **PR and marketing teams, vendors, third-party suppliers**, and even **key customers**.

Many companies keep a register of stakeholders, which is a good starting point for identifying everyone you will want to notify if there is a disaster.

Establish Communication Channels

Create **a list of all teams** responsible for DR, along with their roles and contact information. Establish **a complete chain of command**, including relevant executive leadership and accountable individuals from each of the engineering teams (such as network, systems, database, and storage).

Assign a designated contact person from the support team as well. You should also set up **dedicated communication channels and hubs**, such as an on-site room where everyone will gather, or a **remote information-sharing tool** to use for instant messaging.

Collect All Infrastructure Documentation

Although your engineering teams that are dispatched to activate DR procedures possess the required skills and knowledge for shifting operations to your target DR site, infrastructure documentation is still recommended, especially given the pressure that comes with a disaster.

Even highly trained engineers often prefer to follow infrastructure documentation line by line and command by command during a disaster. The documentation should list all of your mapped network connections (with functioning devices and their configurations), the entire setup of systems and their usage (operating system (OS) and configuration, applications running, installation and recovery procedures), storage and databases (how and where the data is saved, how backups are restored, how the data is verified for accuracy), and cloud templates. It should contain everything IT-related that your business relies upon. Keep hard copies of the documentation, as outages may knock your internal systems offline.

What you will need

- Mapped network connections
- Setup of systems and usage
- Storage and databases
- Cloud templates

PRO TIP:
Store hard copies of the documentation in a safe place that is accessible in the event of an emergency.

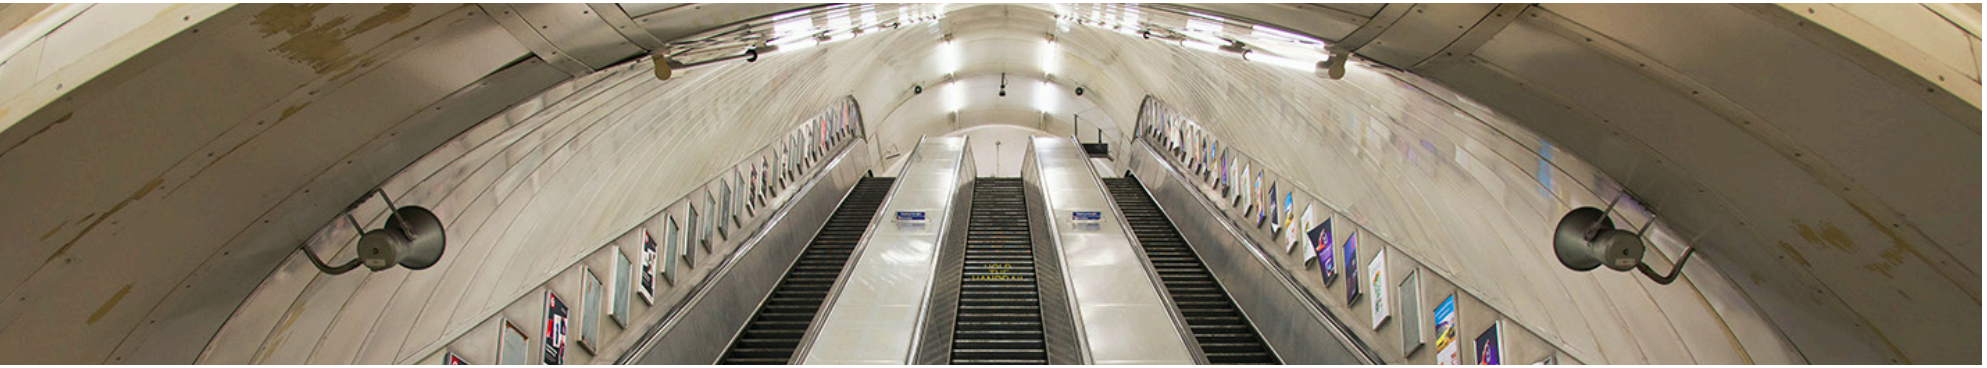

Choose the Right Technology

There are many effective solutions for business continuity beyond traditional, on-premises DR. Cloud-based DR solutions can enable you to securely spin up your DR site on a public cloud, such as Amazon Web Services (AWS), in minutes.

Before selecting a DR solution, consider total cost of ownership (TCO), which can be higher for on-premises DR than cloud-based strategies because of duplicate hardware and software licensing costs. In addition, take into consideration the ability to recover to previous points in time, meet recovery objectives and maintenance requirements, scalability, and ease of testing. You should also consider how the solution will work with the hardware and software you currently run in your production environment.

What to look for

- Reduces total cost of ownership
- Meets defined recovery objectives
- Includes point-in-time recovery
- Enables monitoring and maintenance
- Can be tested easily
- Is scalable

The makings of a disaster

- How long do your systems have to be affected?
- What are considered trigger events or causes?
- Who can report an incident? How is it reported?
- How do you verify it is not a false alarm?

Define Incident Response Procedure

An incident response procedure defines in detail what your company considers to be a disaster. For example, if your system is down for five minutes, should you declare a disaster? Does it matter what the cause is?

In addition to listing the events that will be declared a disaster, the procedure indicates how you will **verify that the disaster is really happening and how the disaster will be reported** — by an automatic monitoring system, raised by calls from site reliability engineering (SRE) teams, or reported by customers?

To verify that a disaster is taking place, check the status of critical network devices, application logs, server hardware, or any other critical components in your production system that you monitor proactively. Being able to quickly detect the failure and verify that it is not a false alarm will impact your ability to meet your RTO.

Define Action Response Procedure & Verification Process

After declaring a disaster, the recovery environment should be activated as soon as possible.

An action response procedure outlines how to perform failover to the DR target site with all necessary steps. Even if your recovery process uses a DR tool with automated components, prepare the action response procedure in writing to define how the necessary services will be started, verified, and controlled.

In addition, it is not enough to simply spin up production services in another location. It is critical to have a verification process that tests that all of the required data is in place, network traffic has been redirected, and all of the required business applications are functioning properly.

Perform Regular Testing

Testing your DR plan in action is essential, but is often neglected. Many organizations do not test on a regular basis because their failover procedures are too complex and there are concerns that failover tests will lead to a disruption of their production environment or even data loss.

Despite these concerns, it is important to schedule regular failover tests to your DR site.

Not only will DR drills demonstrate whether your DR solution is adequate, it will also **prepare your engineers and supporting teams** to respond quickly and accurately to a disaster. Performance tests are also important to **assess whether your secondary location is sufficient** to withstand the business load.

Stay Up to Date

Many companies **keep a risk register** that lists potential risks to business continuity and contains analyses of previous disasters and lessons learned. Review how your teams handled past tests or disaster events, and document your findings. In addition, continue to **update your DR strategy** to reflect changes you make to your primary production environment.

Prepare for Failback to Primary Infrastructure

For most organizations, the DR site is not designed to run daily operations, and a lot of effort may be required to move data and business services back to the primary environment once the disaster is over.

You may need to plan for downtime or a partial disruption of your business during the failback process to your primary site.

Fortunately, there are DR solutions that simplify failback to your primary infrastructure after the disaster, once you have verified that your primary environment is operational.

About CloudEndure Disaster Recovery

CloudEndure Disaster Recovery minimizes downtime and data loss by enabling fast, reliable recovery of physical, virtual, and cloud-based servers into AWS in the event of IT disruption. Meet stringent recovery objectives and reduce your disaster recovery TCO with a single tool for your entire environment.

CloudEndure Disaster Recovery continuously replicates your workloads into a low-cost staging area in AWS, which reduces compute costs by 95% and eliminates the need to pay for duplicate OS and third-party application licenses. You only pay for fully provisioned workloads during a disaster or test. With CloudEndure Disaster Recovery, you can recover your environment in its most up-to-date state or a previous point in time for cases of data corruption, accidental system changes, or cyberattacks.

aws.amazon.com/cloudendure-disaster-recovery | cloudendure-info@amazon.com