Checklist for your IT disaster recovery plan

The key to managing IT disruptions effectively is preparation

Disaster recovery plan checklist

1. Determine recovery objectives (RTO and RPO)
2. Identify stakeholders
3. Establish communication channels
4. Collect all infrastructure documentation
5. Choose the right technology
6. Define incident response procedure
7. Define action response procedure and verification process
8. Perform regular disaster recovery drills
9. Stay up to date
10. Prepare for failback to primary environment

1. Determine recovery objectives (RTO & RPO)

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

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 keep your business running and achieve minimal RPOs and RTOs.

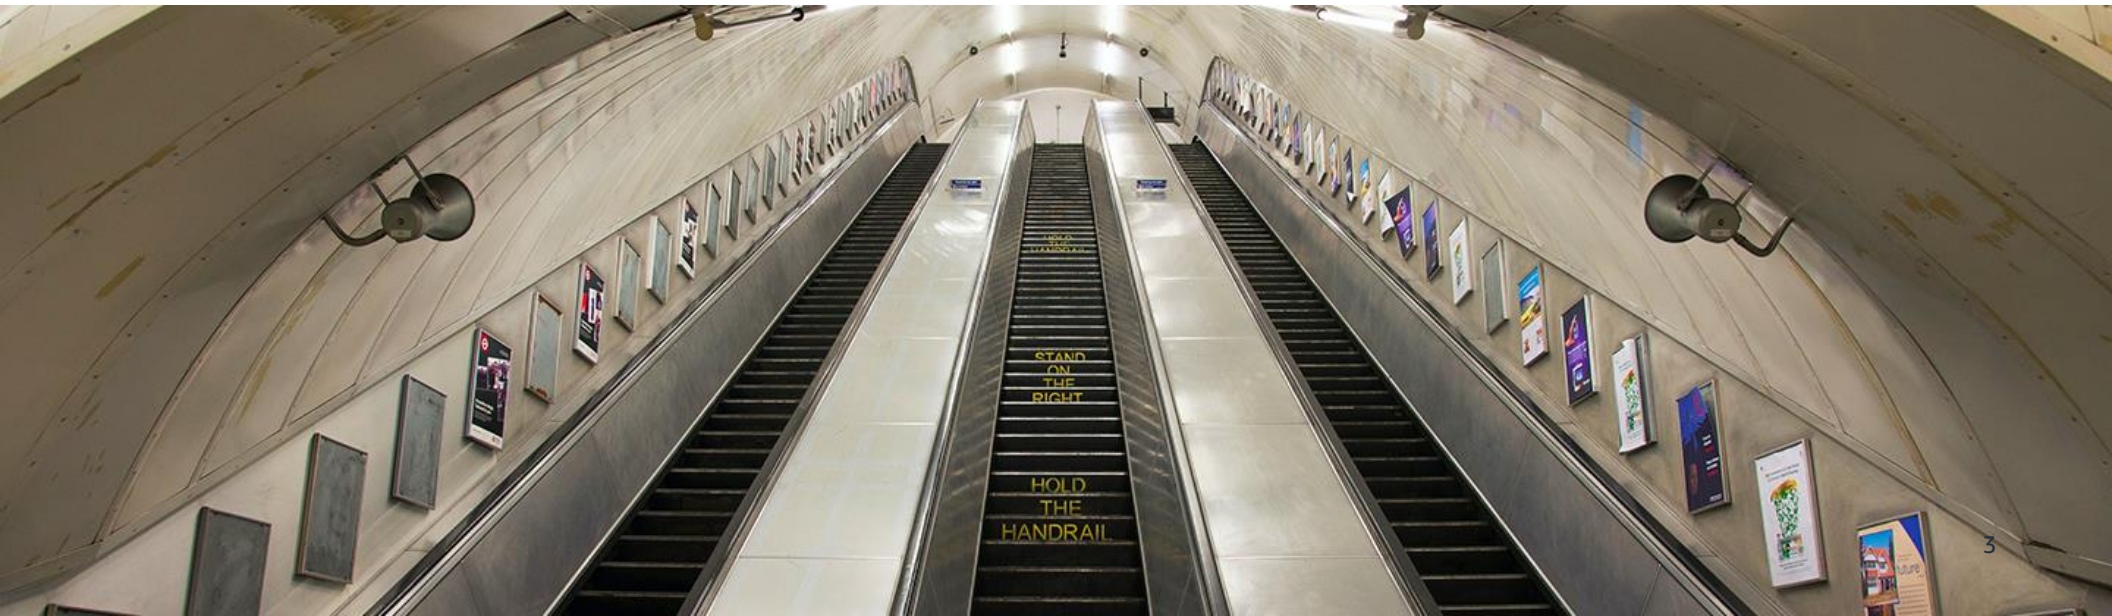

2. Identify stakeholders

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 **public relations 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.

3. 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.

4. 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
- Storage and databases
- Cloud templates
- Setup of systems and usage

PRO TIP:

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

5. Choose the right technology

There are many effective solutions for business continuity beyond traditional, on-premises DR. Using the cloud as a DR site can provide additional flexibility and security, and can reduce the cost of recovery site hardware and maintenance.

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
- Includes point-in-time recovery
- Meets defined recovery objectives
- Simple to monitor and maintain
- Can be tested easily
- Is scalable

6. 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 disasters, 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.

Identifying 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?

7. 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 all of the necessary steps for how you will fail over to your DR site. 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.

8. Perform regular DR drills

Testing your DR plan in action is essential, but is often neglected. Many organizations do not perform DR drills 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 frequent DR drills.

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.

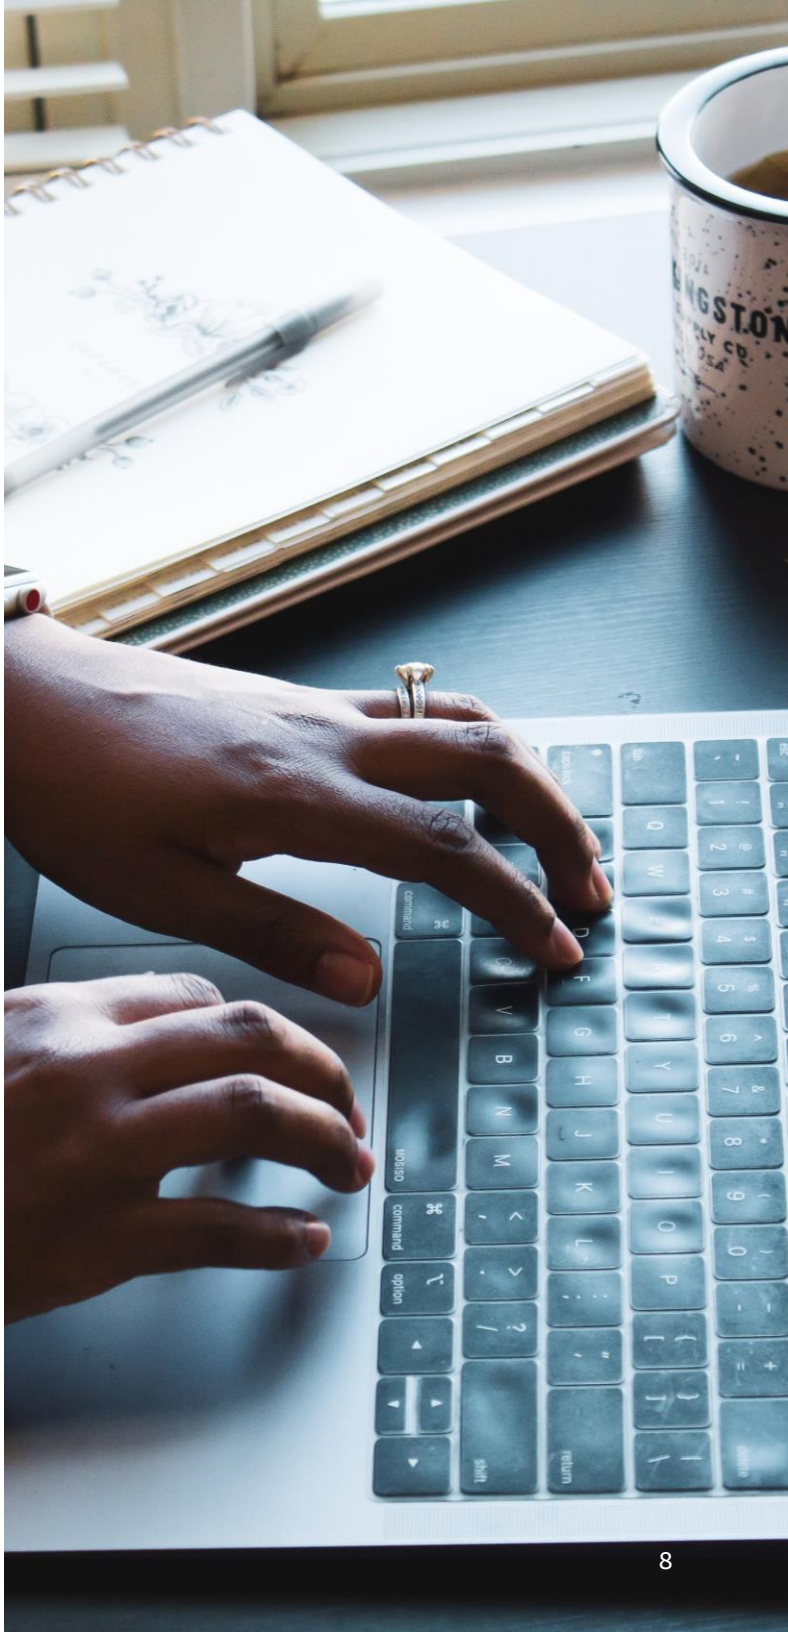

9. Stay up to date

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 drills or disaster events and document your findings. In addition, continue to **update your DR strategy** to reflect the changes you make to your primary environment.

10. Prepare for failback to primary environment

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 environment after the disaster, once you have verified that it is operational.

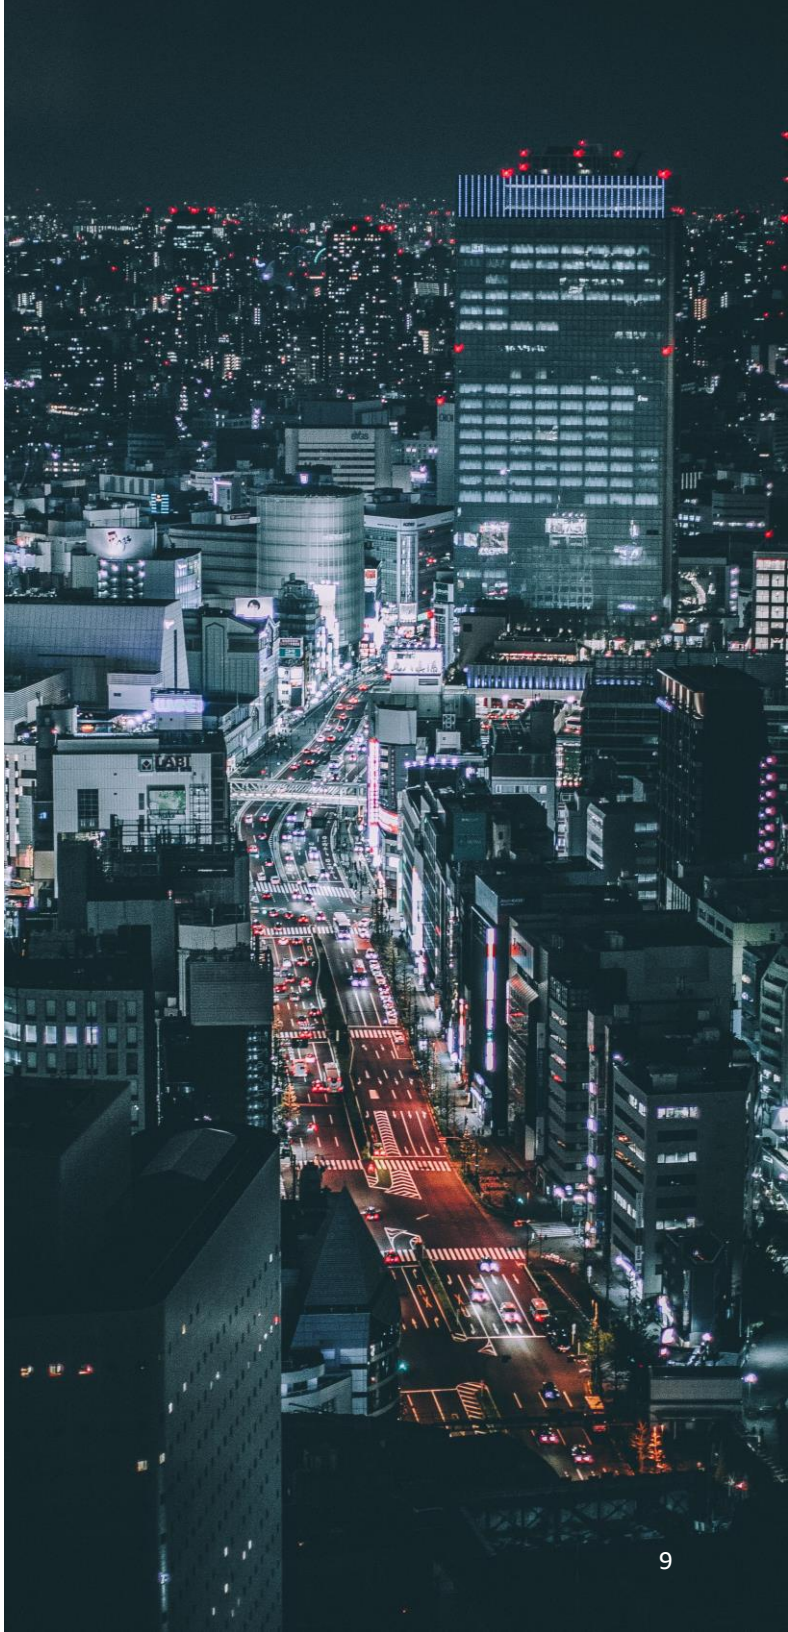

About AWS Elastic Disaster Recovery

AWS Elastic Disaster Recovery minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications. It uses cost-effective AWS resources to maintain an up-to-date copy of your source servers on AWS, thereby removing idle disaster recovery site resources.

During normal operation, use Elastic Disaster Recovery to maintain disaster readiness by performing non-disruptive recovery and failback drills. In the event of an IT disruption, recover your applications on AWS within minutes, at their most up-to-date state or from a previous point in time. Point-in-time recovery is useful for recovery from data corruption events such as ransomware. After the issue is resolved in your primary environment, you can use Elastic Disaster Recovery to fail back whenever you are ready.