

# The Definitive Disaster Recovery Plan Checklist

Every day there are new horror stories of tech outages, downtime, and data loss — even at the best of companies. When disaster strikes, engineering teams are dispatched to repair the damage, while PR teams work overtime to restore customer confidence. It's 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 or you are just beginning the process of creating one for your organization, this Definitive Disaster Recovery Plan Checklist will help you ensure you've included all the crucial components in your plan.

## #1: Determine Recovery Objectives (RTO and 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) is 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 must contain a clearly-stated allowable RTO. 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's mission-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 that data loss is tolerable. If your business can only survive four hours of data loss and you perform nightly backups, you would have a catastrophic loss of important data if disaster strikes the afternoon after a full backup. 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 technology based on real-time continuous data replication, to enable you to achieve near-zero RPO and RTO.

## #2: Identify Stakeholders

The next step is to identify all those who need to be updated once disaster strikes. In addition to those involved with performing the actual recovery from a disaster (e.g., engineers, support, executives), you should also pinpoint members of your PR and marketing team, vendors, third-party suppliers, and even key customers. Many companies keep a register of stakeholders, a good starting point for identifying all of the stakeholders you'll want to notify if there is a disaster.

## #3: Establish Communication Channels

Organizations should keep 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 (e.g., 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 an online 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 a must, especially with the pressure that comes with a disaster. Even the most highly trained engineers 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 (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. Of course, always keep hard copies of the documentation, as outages may knock your internal systems offline.

## #5: Choose the Right Technology

There are many effective solutions for business continuity beyond traditional in-house, on-premises or outsourced Disaster Recovery as a Service (DRaaS) solutions. Another option is to utilize cloud-based DR, where you can spin up your DR site on a public cloud such as AWS in minutes using an automated DR solution. Before selecting a DR solution, you should consider total cost of ownership (which is much higher with an on-premises DR strategy because of the duplicate hardware and software licensing costs), scalability, ability to recover to previous points in time, maintenance requirements, recovery objectives, and ease of testing. You should also take your current production setup into consideration (the hardware and software that you run in a production environment every day).

## #6: Define Incident Response Procedure

An incident response procedure is a must in every DR plan. This is where companies define in detail what is considered 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 plan should indicate how you will verify 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. If something is odd or not working, such as customers being unable to reach your online shop or access their data, then you definitely have a disaster on your hands. Being able to quickly detect the failure and verify that it's not a false alarm will impact your ability to meet your RTO.

## #7: Define Action Response Procedure

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 or DRaaS provider that launches your DR site automatically, you should still prepare the action response procedure in writing to be completely certain 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. A verification process in which you make sure that all the required data is in place, and all the required business applications are functioning properly, is critical.

## #8: Prepare for Failback to Primary Infrastructure

For most companies, the DR site is not designed to run daily operations, and a lot of effort may be required to implement the moving of 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. Fortunately, there are DR solutions that provide seamless failback to your primary location, triggered either automatically or manually after you have completed verification of your primary environment.

## #9: Perform Extensive Tests

Testing your DR plan in action is essential, but is often neglected. Many organizations don't test on a regular basis because their

failover procedures are too complex and there are legitimate 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 (minimally, once a quarter) failover tests to your DR site. If you never test your DR plan, you are putting your entire business at risk, since you might not be able to recover in time (or at all) if disaster strikes and your recovery plan doesn't work. Not only will DR drills demonstrate whether your DR solution is adequate, but 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 or not your secondary location is sufficient to withstand the business load.

## #10: Stay Up-to-Date

Keeping all of your DR documentation updated is as important as regularly scheduled testing of your target infrastructure. After every test (or worse, every incident), review what happened, how your teams handled the test or event, and document your findings. Many companies keep a risk register that, in addition to listing potential risks to business continuity, include analyses of previous disasters and lessons learned.

## Disaster Recovery Plan Example

Here's an example summary of a DR plan for a modern company, running 200 physical and virtual servers in an on-premises data center. (Note: The plan below is an overview. An organization's full DR plan would run anywhere from 10 to more than 100 pages.) The company relies on its production environment being available 24/7 to customers, which is why their DR strategy needs to function perfectly with minimal downtime. This company uses AWS as their target DR infrastructure in order to cut costs and improve their RTO and RPO.

<p><b>Recovery Objectives</b></p>	<p><b>RTO:</b> 5 minutes According to their RTO, the production environment must be shifted from the on-premises data center to AWS within 5 minutes.</p> <p><b>RPO:</b> 0 minutes RPO is near-zero because the business cannot tolerate any data loss. This is why data is continuously replicated from the on-premises environment to the cloud.</p> <p><b>Required Documents:</b></p> <ul style="list-style-type: none"> <li>- Stakeholder Register</li> <li>- Risk Register</li> <li>- Communication Plan</li> </ul>
-----------------------------------	--

<p><b>Incident Reporting</b></p>	<p><b>Sources of Incident Reporting:</b></p> <ul style="list-style-type: none"> <li>- Automatic monitoring service</li> <li>- External (customers) or internal incident reporting (support, engineering)</li> </ul> <p><b>When Incident Is Reported:</b></p> <ul style="list-style-type: none"> <li>- Gather responsible teams and implement chain of command</li> <li>- Perform required production checks to establish whether it is a real threat</li> <li>- Determine if the production environment can be repaired within the defined RTO, or if the DR plan should be triggered</li> </ul> <p><b>Required Documents:</b></p> <ul style="list-style-type: none"> <li>- Incident handling documentation</li> </ul>
<p><b>Action Response</b></p>	<p><b>Ops/Sys Admin Teams Should:</b></p> <ol style="list-style-type: none"> <li>1. Verify data replication and diagnose potential loss of data</li> <li>2. Check network connectivity</li> <li>3. Route traffic to disaster recovery site</li> <li>4. Verify secondary production before starting</li> </ol> <p><b>Required Documents:</b></p> <ul style="list-style-type: none"> <li>- Infrastructure documentation of physical environment</li> <li>- Failover procedures</li> <li>- AWS infrastructure documentation and logging procedure</li> </ul>
<p><b>Operation Restore</b></p>	<p><b>Fallback Procedure:</b></p> <ol style="list-style-type: none"> <li>1. Perform verification of primary site when disaster has finished</li> <li>2. Perform verification of other components, such application/web servers, load balancers, network connections</li> <li>3. Prepare for fallback by reversing data replication from the target to the source environment</li> <li>4. Perform final run tests before going live on the primary site</li> </ol> <p><b>Required Documents</b></p> <ul style="list-style-type: none"> <li>- Fallback procedures</li> <li>- Findings and lessons learned documentation</li> </ul>



## Summary

This ten-point checklist provides you with a starting point for developing a solid DR plan. That said, as every business has its own processes and procedures, you will need to tailor these guidelines to fit your organization's needs. Although DR used to be something that organizations tended to manage in house, security advances have made the cloud a trusted target site for DR, just as more organizations are choosing to run their primary workloads in the cloud. If you're considering transitioning your DR site to the cloud, the [Affordable Enterprise-Grade Disaster Recovery Using AWS](#) white paper is a helpful resource to evaluate DR strategies. CloudEndure Disaster Recovery is an automated DR solution that can spin up thousands of machines in your target AWS Region from any source infrastructure within minutes, and with minimal data loss (due to block-level, continuous data replication). Additional benefits include unlimited, non-disruptive DR tests that are easy to implement, as well as reduced total cost of ownership (since you only pay for the more expensive cloud resources when you use them in a disaster or drill). With the help of CloudEndure Disaster Recovery, you can recover your entire environment in its most up-to-date state or from a previous point in time, ensuring that your business will run as usual in the event of a disaster.

---

### About CloudEndure

CloudEndure accelerates the journey to the AWS cloud with solutions that provide business continuity during the migration process and additional protection once there. CloudEndure Migration simplifies, expedites, and automates large-scale migrations from physical, virtual, and cloud-based infrastructure to AWS. CloudEndure Disaster Recovery protects against downtime and data loss from any threat, including ransomware and server corruption. With CloudEndure it's business as usual, always.

---

