CloudEndure®

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 plan.

Whether you already have a disaster recovery plan or you're just beginning the process of creating one for your organization, the Definitive IT Disaster Recovery Plan Checklist below 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 IT disaster recovery is to keep your business operating as usual, all the time, so you need to determine which IT services are most critical to the running of your organization and what Recovery Time Objective (RTO) and Recovery Point Objective (RPO) are required for these services/machines.

RTO is the amount of time required to recover from a disaster after notification of business disruption. A reliable disaster recovery plan must contain a clearly-stated allowed RTO. If your business cannot withstand an hour of downtime without losing customers to your competitors or paying penalty fees due to service-level agreements (SLAs), it's crucial that it be back online before an hour expires. 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 survive a data loss of four hours, yesterday's full backup in a situation where disaster strikes in the afternoon, after a full day of business, can mean a catastrophic loss of important data. In this case, your RPO would be four hours.

A company's RTO and RPO will affect its disaster recovery strategy as well as associated expenses. More aggressive recovery objectives can require higher costs. In order to reduce the total cost of your disaster recovery strategy, we recommend dividing applications into tiers. The highest tier, reserved for mission-critical applications, should have near-zero RPO and RTO. Therefore, you will likely need a disaster recovery technology based on real-time continuous data replication. The mid-level tier may need less aggressive RPO and RTO so snapshot-based replication might be good enough. The lowest tier, which often includes applications used for data retention, may get by with a simple file-level backup system.

#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 disaster recovery (e.g. engineers, support, executives), you need to pinpoint others such as members of the PR and marketing team, vendors, third-party suppliers, and even key customers. Many companies keep a register of stakeholders in their project office documentation, a good starting point for identifying all stakeholders you'll want to notify in the case of a disaster.

#3: Establish Communication Channels

Organizations should keep a list of all teams responsible for disaster recovery, 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, storage). Assign a designated contact person from support, as well. You should also set up dedicated communication channels and hubs, whether it's an on-site "war room" in which everyone gathers or an online information sharing tool to use for instant messaging.

#4: Collect All Infrastructure Documentation

Even though you'll dispatch your engineering teams to activate the disaster recovery procedures, and they presumably possess the required skills and knowledge for shifting operations to your disaster recovery site, infrastructure documentation is a must, especially because of the pressure that everyone will be under during a disaster. Even the most highly trained engineers prefer to follow infrastructure documentation line by line, command by command while performing disaster recovery.

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 practically everything IT-related on which your business relies. Of course, always keep hard copies of the documentation as outages may knock your internal wiki offline.

#5: Choose the Right Technology

Do-it-yourself on-premises disaster recovery and outsourced Disaster Recovery as a Service (DRaaS) are no longer the only viable solutions for business continuity. Another option is to utilize cloud-based disaster recovery, where you can spin up your disaster recovery site on a public cloud such as AWS, Microsoft Azure, and Google Cloud in minutes using an automated disaster recovery solution.

Before selecting your solution, you should consider total cost of ownership (which is much higher with an on-premises disaster recovery strategy because of duplicate hardware and software licensing costs), scalability, recovery to previous points in time, maintenance requirements, recovery speed, continuous data replication, and ease of testing. You should also take into consideration your current production setup (the hardware and software that you run in a production environment every day). As there are plenty of choices, be sure to research and carefully consider all your options before selecting one disaster recovery solution.

#6: Define Incident Response Procedure

An incident response procedure is a must in every disaster recovery 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 really happening, 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, if customers cannot 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 RTO.

#7: Define Action Response Procedure

After declaring a disaster, a disaster recovery environment should be activated as soon as possible. An action response procedure will outline how to failover to the disaster recovery site, with all the necessary steps. Even if your recovery process is utilizing a disaster recovery tool or DRaaS that will launch your disaster recovery site automatically, you should still prepare the action response procedure in writing to be completely sure how the necessary services will be started, verified, and controlled. Also, simply spinning up production services in another location is not enough. 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 disaster recovery site is not the one 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 location once the disaster is over. Plan for possible downtime (if required) or a potential partial disruption of your business during the revert process. Luckily, there are disaster recovery solutions that provide seamless failback to the primary location, triggered either automatically or manually once you complete the verification of the primary IT location.

#9: Perform Extensive Tests

Testing your disaster recovery plan in action is essential, but often neglected. Usually, failover procedures are too complex and there are legitimate concerns from technology leads that failover tests will lead to a disruption of production services, or even data loss, so many companies don't test their disaster recovery plan on a regular basis.

Despite these concerns, you should schedule regular (minimally, once a quarter) failover tests to the disaster recovery site to see if and how well your disaster recovery plan works. If you never test your disaster recovery plan, you are putting at risk your entire business once disaster strikes, and you might end up not able to recover in time, or recover at all. Not only will disaster recovery drills demonstrate if your disaster recovery solution is adequate, but it will also train 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. Some companies utilize their disaster recovery site every day, running non-essential applications on it, which is another way to verify that your environment is operational.

#10: Stay Up-to-Date

As important as regularly scheduled testing of your disaster recovery infrastructure is, so is keeping all of your disaster recovery documents updated. 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 post-mortems of previous disasters and lessons learned.

Disaster Recovery Plan Example

As an example, let's review a summary of a disaster recovery plan for a modern IT enterprise company, running a total of 200 servers (physical and virtual) in an on-premises data center. (Note: The below chart is just an overview. Your organization's full disaster recovery plan would run anywhere from 10 to more than 100 pages.) The company relies on its Internet production environment, available 24/7 to customers, which is why their disaster recovery strategy needs to function perfectly with minimal downtime. The company also recently switched their disaster recovery site to a cloud-based disaster recovery infrastructure, in order to cut costs and improve their RTO and RPO.

Recovery Objectives	**RTO:** < 5 minutes RTO states that the entire production is shifted from local data center to AWS cloud, with production services dynamically switched to secondary DNS records based on AWS Route53 health checks. Disaster recovery site is a hot standby, with auto-scaling of application services based on load, while the data is synchronously replicated from primary DB to one located in the cloud. **RPO:** 0 minutes RPO should be near-zero because the business cannot tolerate any data loss. That's why data is synchronously replicated from the on-premises database cluster to the managed database services in the cloud. **Documents required here:** - Stakeholder register - Risk register - Communication plan

Incident Reporting	**Sources of incident reporting:** Automatic monitoring service External (customers) or internal incident reporting (support, engineering) **When incident is reported:** Gather responsible teams and implement chain of command Do required production checks to establish that it's a real threat Determine if the production can be repaired in RTO, or if disaster recovery plan should be triggered **Documents required here:** - Incident handling documentation
Action Response	**Ops/SysAdmin teams should:** 1. Verify database replication and diagnose potential loss of data 2. Start minimum number of required servers and check auto-scaling operation on AWS 3. Route traffic to disaster recovery site 4. Verify secondary production before starting **Documents required here:** - Infrastructure documentation of physical environment - Failover procedures - AWS infrastructure documentation and logging procedure
Operation Restore	**Fallback procedure:** 1. Verification of primary site when disaster is finished 2. Data verification of the primary database 3. Verification of other components (application/web servers, load balancers, network connections) 4. Final run tests before going live on the primary site

Summary

This ten-point checklist provides you with a solid starting point for developing a strong disaster recovery plan. That said, as every business has its own individual processes and procedures, you will need to tailor these guidelines to fit your organization's needs. For years, disaster recovery sites were managed in-house. However, the cloud is becoming a much more reliable solution for secondary disaster recovery sites, just as it is for the one in production. (If you're thinking about transitioning your disaster recovery site to the cloud, this white paper can be a good starting point.)

While implementing disaster recovery in the cloud, a good approach is to invest in third-party SaaS solutions that run on top of public cloud services and can provide you with the disaster recovery functionality you need. CloudEndure, which runs on top of services provided by all major cloud vendors, gives you automated disaster recovery with near-zero data loss (due to continuous data replication) and reduces total cost of ownership since you only pay for cloud services when you use them. Additional advantages to this solution include easy implementation of non-disruptive disaster recovery tests as well as automatic failover and failback procedures, ensuring protection of your business services. With the help of CloudEndure's disaster recovery solution, you can be certain that your business is protected from any kind of disaster and that your production services will be up and running in the cloud seconds after disaster strikes.

About CloudEndure

CloudEndure provides **Disaster Recovery** and **Live Migration** for all applications, allowing enterprises to mobilize entire workloads to and across clouds with **near-zero downtime** and **no data loss**. Our **Live Workload Mobility** technology provides continuous, block-level replication and application stack orchestration — at the touch of a button, within minutes, and with the latest data.

CloudEndure supports physical, virtual, and cloud-based infrastructure as sources and Amazon Web Services (AWS), Google Cloud Platform (GCP), Microsoft Azure, OpenStack, and Oracle Cloud as target cloud locations. Whether you are seeking a Disaster Recovery solution or Migration tool, or both, CloudEndure ensures all your systems are always available. **CloudEndure – All Systems Go.™**

For more information, visit www.cloudendure.com.