

IT Disaster Recovery Plan

Owner	[Engineering owner role]
Approved by	[Policy approver role]
Effective date	[Effective date]
Next review	[Next review date]
Version	[Version]

1. Purpose and the criteria it serves

This plan sets how [Company name] (the "Company") restores [System name] and the systems it depends on after a disaster: the loss of a cloud region, a database, an account, a critical vendor or the data itself. It is the technical runbook. The Business Continuity Plan sets how the Company keeps serving customers while this plan runs, and the Incident Response Plan sets how an incident is declared, handled and reported.

Criterion	What it expects	Where this plan answers it
CC7.5	Activities to recover from identified security incidents	Activation, restore order, recovery procedure, verification and return to normal
CC9.1	Risk mitigation activities for potential business disruptions	System inventory, backup inventory and testing
A1.2	Data backup processes and recovery infrastructure, when Availability is in scope	Backup inventory and recovery procedure
A1.3	Testing of the recovery plan procedures, when Availability is in scope	Testing and the record it leaves

A SOC 2 report that covers only the Security category is tested against CC7.5 and CC9.1. A1.2 and A1.3 apply only when the Availability category is in scope. Delete their rows if it is not, and keep the testing section either way: a backup nobody has restored is not evidence of recovery.

2. Scope and system inventory

This plan covers [In-scope systems]. The Company targets a recovery time objective of [Recovery time objective] and a recovery point objective of [Recovery point objective] for the production service. Each system below has its own targets, and no row promises a shorter time than its last restore test achieved.

System	Tier	Recovery time	Data loss allowed	Owner
[Production application]	Tier 1	[Recovery time objective]	None held	[Engineering owner role]
[Production database]	Tier 1	[Recovery time objective]	[Recovery point objective]	[Engineering owner role]

Tier 1: customers cannot use the service without it. Tier 2: the service runs degraded. Tier 3: staff work slows, customers are not affected.

2.1 Worked example

A twelve person software company whose production runs in one AWS region. The full rows, with backup source, dependencies and test dates, are in the Systems sheet of the recovery runbook register.

System	Tier	Recovery time	Data loss allowed	Owner
Staff identity provider and break glass cloud access	Tier 1	1 hour	Not applicable	Chief technology officer
DNS zone and TLS certificates	Tier 1	1 hour	Not applicable	Engineering lead
Secrets and encryption keys	Tier 1	2 hours	Not applicable	Chief technology officer
Production database (managed PostgreSQL)	Tier 1	4 hours	15 minutes	Engineering lead
Customer file storage (object storage)	Tier 1	8 hours	24 hours	Engineering lead
Application containers and background workers	Tier 1	4 hours	None held	Engineering lead
Internal wiki and ticketing	Tier 3	3 business days	24 hours	Operations lead

3. Roles and contacts

Role	Held by	Responsibility
Recovery lead	[Engineering owner role]	Declares a disaster, chooses between waiting, failing over and restoring, runs the recovery and signs off before reopening.
Backup recovery lead	[Named second person]	Takes every recovery lead duty when the lead cannot be reached within 30 minutes.
Security owner	[Security owner role]	Runs the security incident under the Incident Response Plan, checks that security controls are back in place, and co-signs the reopening.
Communications lead	[Customer-facing role]	Writes and sends status updates and customer notices, approved by the recovery lead.
Executive approver	[Policy approver role]	Approves spending, contractual notices and any decision to stay down longer than the recovery time objective.

Names, phone numbers and the out of band channel live in the Contacts sheet of the recovery runbook register. The recovery lead and the backup each keep a copy outside company systems, because the wiki and email may be down too.

At least two people hold break glass access to the cloud account and the domain registrar. The credentials are stored in [Secrets manager] and in a sealed offline copy, every use is logged and reviewed, and they are changed after each use.

4. Activation: when, and who declares

The recovery lead declares a disaster when any of these is true:

- A Tier 1 system is down or unusable and the provider shows no fix inside its recovery time objective.
- Production data is lost, corrupted or encrypted, including by ransomware.
- The cloud account or the production region is unavailable or under an attacker's control.
- A vendor the service cannot run without is unavailable and its fallback has to be used.

The declaration is recorded in [Ticketing system] with the time, the trigger, and the decision taken: wait for the provider, fail over, or restore from backup, with who decided. A security incident is opened under the Incident Response Plan at the same time, because it sets the notification duties.

5. Backup inventory

Backups are taken [Backup frequency], kept for [Backup retention period] and stored in [Backup location], separate from the primary data. At least one copy of each Tier 1 data set sits in a different account or behind different credentials, so that one compromised account cannot delete both the data and its backups. Backups are encrypted with [Encryption standard at rest]. A failed backup job raises an alert in [Monitoring tool].

Data	Method	How often	Kept for	Where
[Production database]	[Snapshots and point in time recovery]	[Backup frequency]	[Backup retention period]	[Backup location]
[Customer files]	[Versioning and replication]	[On every write]	[Backup retention period]	[Backup location]
Infrastructure definition	Infrastructure as code	On every merge	Full history	[Code repository]
Secrets and keys	A list of secret names and how to reissue each. Values are reissued, not restored	On every change	Current version	[Secrets manager]
DNS zone	Zone file export	On every change	Every version	[Code repository]

5.1 Worked example

The managed PostgreSQL database keeps point in time recovery for 35 days and takes a daily snapshot that is copied to a second region in a separate backup account. Customer files are versioned, prior versions are kept for 30 days, and objects are replicated to the same second region. A restore can therefore lose at most 15 minutes of database writes and one day of file changes, which is what the inventory states.

6. Restore order and dependencies

Restore in this order. Each step needs the one before it, so starting further down the list wastes the recovery time objective.

1. Access: the identity provider and the break glass cloud credentials. Nothing else can be restored without them.
2. DNS and certificates: the domain must resolve to wherever the service will run, and certificates must be valid there.
3. Secrets and keys: confirm the encryption keys exist before restoring encrypted data. A backup encrypted with a deleted key cannot be read.

4. Network: virtual networks, subnets, firewall rules and load balancers, from the infrastructure code.
5. Database: restored to the recovery point, as a new instance.
6. Object storage: buckets restored or switched to the replica.
7. Compute: the application and workers, deployed from the last released build through [Deployment pipeline].
8. Third parties: email, payments and other vendors reach the restored environment, with webhooks, allowed addresses and API keys updated.
9. Monitoring and logging: [Logging system] and [Monitoring tool] receive data from the restored environment before customers are let back in.

7. Recovery procedure by layer

7.1 Database

1. Choose the recovery point: the last moment before the loss or corruption, from the incident timeline.
2. Restore to a new instance, never over the damaged one, so the damaged copy stays available for investigation.
3. Run the verification checks in the next section against the new instance.
4. Point the application at the new instance by changing its connection secret, and record the time.
5. Keep the damaged instance, with access restricted, until the incident is closed.

7.2 Object storage

1. For deleted or overwritten objects, restore prior versions from the version history.
2. For a lost bucket or region, switch the application to the replica and record which objects changed after the last replication.

7.3 Application and compute

1. Redeploy the last released build through the deployment pipeline. Do not rebuild by hand from a laptop.
2. Scale to normal capacity and confirm background jobs are processing.

7.4 Loss of the production region

1. Recreate the network and compute in the recovery region from the infrastructure code.
2. Restore the database from the cross region copy and point object storage at the replica.
3. Switch DNS to the recovery region once the verification checks pass.

7.5 Compromised cloud account

1. Handle it as a security incident first, under the Incident Response Plan.
2. Rebuild in a clean account from the copies held outside the compromised one, and reissue every credential and secret.

3. Do not restore a backup taken after the earliest suspected compromise until it has been checked.

7.6 Vendor outage

1. Follow the fallback recorded for that vendor in the runbook register, or wait for the vendor if the outage will end inside the recovery time objective.
2. Record the decision and the time in the incident record.

8. Verification before reopening

The service reopens to customers only when every check below passes, and the recovery lead and the security owner have both signed off in the incident record.

- Row counts on the key tables match the last known counts, allowing for the recovery point.
- The newest record is from within the recovery point objective of the failure time.
- A test customer account signs in, opens a record and downloads a file.
- Background jobs and scheduled tasks run, and outbound email sends.
- Security controls are back: single sign on and multi-factor authentication, encryption at rest, logging, alerting and network rules match production as documented.
- The actual recovery time and data loss are recorded against the targets in the system inventory.

9. Communication

Staff coordinate in the out of band channel named in the Contacts sheet. Customers hear through [Status page] within one hour of the declaration and at least every hour while a Tier 1 system is down. Customers with a contractual notice period are notified directly within [Notification deadline] hours, and the executive approver approves that notice.

Every update says what is affected, what is being done, whether any data may be lost, and when the next update will come. It never guesses a cause that has not been confirmed.

9.1 Worked example

"We are restoring our customer database after a storage failure in our primary hosting region. Sign in is unavailable. Data saved before 09:40 UTC is safe, and we are checking whether anything saved after that time was lost. Next update by 11:00 UTC."

10. Return to normal

1. Decide whether to stay in the recovery environment or move back, and plan any move back as a normal reviewed change.
2. Turn backups and replication on for every restored system, and confirm the first new backup completed.
3. Remove temporary access, reissue any break glass credential that was used, and review its use.

4. Close the incident and hold a review within 5 business days. Record what worked, what did not, and each action with an owner and a date.
5. Update this plan, the system inventory and the runbook register with what the recovery taught.

11. Testing and the record it leaves

Test	How often	What it proves	Record it leaves
Restore test of the production database	Every quarter	The backup restores and the data is complete	Test log row: backup used, time taken, data checked, result
Recovery exercise: a region rebuild in a separate account, or a tabletop walk through this plan	Once a year	The plan works end to end and people know their part	Exercise report with findings and actions
Contact list check	Every quarter	The people and numbers in the Contacts sheet are current	Date checked, in the Contacts sheet
Backup job monitoring	Daily, by alert	Backups are being taken	Alert history and the fix for each failure

Every test is recorded in the Test log sheet of the recovery runbook register. A missed target or a failed test is fixed and retested within 30 days, and both rows stay in the log. A test dated before or after an audit period is not evidence for that period.

11.1 Worked example

The annual exercise rebuilt the database, file storage and application in the recovery region in 2 hours 50 minutes against a 4 hour objective, with 6 minutes of data loss against a 15 minute objective. The first deploy failed because the recovery region had no copy of the application container image. Image replication was turned on and the deploy was retested the same day. Both the failure and the retest are in the log.

12. Records an auditor asks for

Record	Where it is kept
This plan, approved, with its version history	[Documentation system]
The recovery runbook register: systems, contacts and test log	[Evidence storage location]
Backup configuration and retention settings, as screenshots or exports	[Evidence storage location]
Restore test and exercise records from inside the audit period	[Evidence storage location]
Incident records for any real recovery, with the review	[Ticketing system]

13. Plan maintenance and approval

The [Engineering owner role] reviews this plan at least once a year, and after any new Tier 1 system, vendor or region, any test finding, and the departure of anyone named in it. The [Policy approver role] approves each version.

Version	Date	Summary of change	Approved by
[Version]	[Effective date]	First version.	[Policy approver role]

Plan approver

Name:

Signature:

Date:

14. Fields to complete

Replace every bracketed field in this document with your own detail, then delete this section.

Field	What to enter
[Company name]	Your legal company name.
[Effective date]	The date the approver signs this version.
[Version]	Start at 1.0 and increase it each time the document changes.
[Security owner role]	The role that runs the security program day to day, for example Chief Technology Officer.
[Policy approver role]	The officer who approves policies and accepted risks, for example Chief Executive Officer.
[Evidence storage location]	The shared folder or system where evidence is kept.
[Documentation system]	Where policies and procedures are published, for example the company wiki.
[In-scope systems]	The product, cloud accounts and tools inside the audit boundary.
[Ticketing system]	Where changes and incidents are tracked, for example Jira or Linear.
[Code repository]	Where source code is hosted, for example GitHub or GitLab.

Field	What to enter
[Deployment pipeline]	The system that builds and deploys code, for example GitHub Actions.
[Secrets manager]	Where keys and secrets are stored, for example a cloud key management service.
[Logging system]	Where logs are collected and searched.
[Monitoring tool]	The system that raises alerts.
[Encryption standard at rest]	For example AES-256, the Advanced Encryption Standard with 256-bit keys, through the cloud provider key service.
[Recovery time objective]	The longest the service may be down after a disaster, for example 4 hours.
[Recovery point objective]	The most data you can afford to lose, for example 24 hours.
[Backup frequency]	For example daily.
[Backup location]	Where backups are kept, separate from the primary data.
[Backup retention period]	For example 30 days.
[System name]	The product or service the report covers.
[Notification deadline]	The shortest customer or legal notice period you are bound by, in hours.
[Status page]	Where customers see service status, if you publish one.
[Engineering owner role]	The role that owns production systems, for example engineering lead.
[Next review date]	The date this version must be reviewed by, at most 12 months after the effective date.

Source

Source: free template from www.polaralabs.com/templates/disaster-recovery-plan (CC BY 4.0).