[COMPANY NAME]
[Add your company logo / letterhead and registered address here]

Disaster Recovery (DR) Policy

HOW TO USE THIS TEMPLATE
This template is mostly complete and pre-filled with standard Indian practice. You should not have to fill many blanks.
Text in blue is a default that companies commonly change. Skim the blue text and edit only what differs for you.
Add your company name and letterhead once, in the header above. The body refers to "the Company".
Set realistic RTO/RPO targets per system tier (not one number for everything), keep runbooks current, and prove the plan works by testing it on the schedule below rather than only after a real outage.
Have it reviewed by a qualified HR or legal professional before you adopt it, and delete this box.
Provided by CFOmatrix (cfomatrix.in). General template, not legal advice.

Policy owner	**[Human Resources / IT / Compliance]**
Effective date	**[DD MMM YYYY]**
Version	1.0
Approved by	**[Name, Title]**

1. Purpose

This Policy establishes the framework by which the Company prepares for, responds to, and recovers from events that disrupt its information technology (IT) systems, applications, and data. A disaster, for the purposes of this Policy, is any unplanned event that renders one or more critical IT services unavailable beyond their tolerable downtime, including hardware or infrastructure failure, cloud region or provider outage, cyber incidents (ransomware, data corruption, denial of service), data centre loss, natural disasters (flood, fire, earthquake), prolonged power or network failure, or supply-chain and third-party service failure.

The objectives of this Policy are to:

- Minimise downtime and data loss for critical systems following a disruptive event.
- Define clear, pre-approved recovery objectives (RTO and RPO) so that recovery decisions are made against agreed targets, not improvised during a crisis.
- Assign unambiguous roles and responsibilities for declaring, managing, and standing down a disaster.
- Maintain tested, current runbooks that allow recovery to proceed even if key individuals are unavailable.
- Support the Company's wider Business Continuity Plan (BCP), regulatory obligations, and customer commitments, including the security control expectations of frameworks such as SOC 2 and ISO/IEC 27001 (control domains for availability, backup, and resilience).

This Policy is technical and IT-focused. It sits underneath, and is invoked by, the Company's Business Continuity Plan, which addresses the broader business response (people, premises, communications, finance). See Section 12.

2. Scope

This Policy applies to:

- All production IT systems, applications, databases, and infrastructure owned, operated, or contracted by the Company, whether hosted in the cloud, in a co-location facility, or on premises.
- All data classified as critical or important to business operations, including customer data, financial records, source code, and configuration.
- All employees, contractors, and managed service providers who design, operate, support, or recover these systems.
- All third-party and cloud service providers whose services the Company depends on, to the extent of their contractual recovery commitments.

This Policy covers recovery of IT services. It does not by itself cover non-IT continuity matters (alternate office space, manual workarounds, crisis communications to the public), which are governed by the BCP.

3. System Tiers and Criticality

Recovery effort and targets are driven by how critical a system is. Every production system must be assigned a tier in the system inventory maintained by **IT / Infrastructure Team**. The default tiering is:

Tier	Description	Examples (edit per Company)	Recovery priority
Tier 1 (Critical)	Revenue-generating or customer-facing; an outage causes immediate business, financial, or contractual impact	**Core product / application, customer database, payment processing, primary website**	Recover first
Tier 2 (Important)	Internal operations depend on it; tolerable for a short period	**ERP, internal CRM, email, ticketing, analytics**	Recover after Tier 1
Tier 3 (Standard)	Useful but non-urgent; business can operate without it for days	**Internal wikis, reporting dashboards, dev/test environments**	Recover last

The system inventory must record, for each system: tier, business owner, technical owner, hosting location, dependencies, backup method, and assigned RTO and RPO.

4. Recovery Objectives (RTO and RPO)

Two objectives govern every recovery:

- RTO (Recovery Time Objective): the maximum acceptable time to restore a system to working order after a disruption is declared.
- RPO (Recovery Point Objective): the maximum acceptable amount of data loss, measured as the time between the last usable backup or replica and the point of failure.

The Company's default targets by tier are below. Individual Tier 1 systems may have tighter, system-specific targets recorded in the inventory; those override the table.

Tier	RTO target	RPO target	Implication
Tier 1 (Critical)	**4 hours**	**15 minutes**	Warm standby or continuous/near-continuous replication required
Tier 2 (Important)	**24 hours**	**4 hours**	Frequent backups; restore from backup acceptable
Tier 3 (Standard)	**72 hours**	**24 hours**	Daily backup and rebuild acceptable

RTO and RPO targets must be reviewed at least annually (Section 13) and whenever a system's tier, architecture, or business importance changes. Targets are commitments the recovery design must be capable of meeting; testing (Section 9) verifies whether they are actually met.

5. DR Strategy

The Company's recovery strategy is risk-based and matched to each tier's RTO and RPO. The default approach is cloud-first, with a defined secondary location for resilience.

5.1 Cloud and multi-region

- Primary production runs in **the Company's primary cloud provider and region**.
- Tier 1 systems must have a defined recovery target in a separate availability zone or, for region-level resilience, a **secondary region**, using one of: active-active, warm standby, or pilot-light architecture as recorded per system.
- Infrastructure must be reproducible through Infrastructure-as-Code (**Terraform / CloudFormation / equivalent**) so that environments can be rebuilt quickly and consistently in the recovery location.

5.2 Secondary site / failover

- Where systems run in a co-location or on-premises facility, a **secondary site** or cloud failover target must be identified and kept in a state consistent with the system's RTO.
- DNS, load balancer, and routing changes required to fail over must be documented in the runbook and access to make them must be pre-provisioned to the recovery team.

5.3 Replication and standby

- Tier 1 databases use **continuous or near-continuous replication** to the recovery location to meet the RPO.
- Standby capacity (compute, storage, licences) sufficient to run Tier 1 workloads must be available or rapidly provisionable in the recovery location.

5.4 Dependencies

- Recovery design must account for upstream dependencies (identity provider, DNS, secrets manager, third-party APIs, payment gateways). A system cannot be considered recoverable if a critical dependency is not.

6. Data Backup (Link)

Backups are the foundation of recovery. The detailed controls (frequency, encryption, retention, off-site copies, integrity checks) are defined in the Company's **Data Backup Policy**. This Policy relies on those backups and adds the recovery-specific requirements below.

- Backups for Tier 1 and Tier 2 systems must be replicated to a location separate from the primary (different region or off-site) so a single-site disaster does not destroy both production and backups.
- The Company maintains a **3-2-1** posture: at least three copies of critical data, on at least two media or platforms, with at least one copy off-site or in a separate cloud account, and at least one immutable or logically air-gapped copy to resist ransomware.
- Backup restore must be tested as part of DR testing (Section 9). A backup that has never been restored is not a proven backup.
- Backup encryption keys and credentials must be recoverable independently of the systems they protect.

7. Roles and Responsibilities

Role	Responsibility
DR Coordinator / **Head of Infrastructure**	Owns this Policy; maintains runbooks and the system inventory; coordinates testing; leads technical recovery
Disaster Declaration Authority / **CTO**	Has authority to formally declare a disaster and to declare it resolved; approves invocation of the recovery plan
Incident Commander	Single point of control during an active disaster; directs the recovery team, manages the timeline, and decides on failover and failback
Recovery Team (system / database / network owners)	Execute runbook steps for systems they own; report status to the Incident Commander
Communications Lead / **HR / Internal Comms**	Manages internal and customer communications during the event, in line with the BCP
Security / IT Security	Assesses whether the disaster is a security incident (e.g. ransomware), ensures recovery does

	not reintroduce the threat, and handles any breach notification obligations
Executive Sponsor / **CEO or Management**	Final escalation point; authorises major spend and external communications
Third-party / cloud providers	Deliver recovery within their contracted SLAs; the DR Coordinator manages these relationships

A named primary and at least one named alternate (deputy) must exist for every critical role so that recovery is not dependent on one person being reachable. Current names and 24x7 contact details are kept in the call tree (Section 8).

8. Disaster Declaration and Escalation

8.1 Detection and reporting

- Monitoring and alerting (**your monitoring tool**) must detect outages of Tier 1 and Tier 2 systems and alert the on-call engineer automatically.
- Any employee who suspects a disaster-level disruption must report it immediately to **the on-call / IT helpdesk: it-oncall@company.com**.

8.2 Declaration

- The on-call engineer triages and escalates to the Incident Commander. If the disruption exceeds, or is forecast to exceed, the affected system's RTO, the Disaster Declaration Authority formally declares a disaster and invokes this Policy.
- Declaration triggers activation of the recovery team, the call tree, and the relevant runbooks.

8.3 Call tree and communication

- A current call tree (primary and alternate contacts for every recovery role, plus key vendor contacts) is maintained by the DR Coordinator and stored where it is reachable even if Company systems are down (e.g. **a secured offline / out-of-band location**).
- A designated out-of-band communication channel (**secondary messaging tool / phone bridge**) is used if primary email and chat are unavailable.
- Status updates are issued to stakeholders at a defined cadence (**every 60 minutes** for Tier 1 incidents) by the Communications Lead.

9. Recovery Procedures and Runbooks

9.1 Runbooks

Every Tier 1 and Tier 2 system must have a written recovery runbook, maintained by its technical owner and stored where it remains accessible during an outage. Each runbook must contain, at minimum:

- System name, tier, owners (primary and alternate), and its RTO and RPO.
- Dependencies that must be available before recovery can start.
- Step-by-step recovery instructions (failover, restore from backup, rebuild from Infrastructure-as-Code), written so a competent engineer who does not own the system can follow them.
- Exact commands, console paths, and the location of credentials/secrets (referenced, not embedded).
- Validation steps to confirm the system is genuinely working (health checks, data integrity checks, smoke tests) before it is declared recovered.
- Failback steps to return to the primary environment once it is safe.

9.2 General recovery sequence

1. Assess and contain: confirm scope and root cause; if a security incident, isolate affected systems before recovering so the threat is not carried into the recovery environment.

2. Prioritise: recover in tier order (Tier 1 first), respecting dependencies.

3. Recover: execute the relevant runbooks (failover to standby/secondary site, or restore from backup, or rebuild from code).

4. Validate: run validation steps; confirm data is consistent and within the RPO.

5. Resume service: redirect traffic/DNS; confirm with the business owner that the service is usable.

6. Stabilise and fail back: once primary is restored and verified, fail back in a controlled window and confirm replication/backups have resumed.

7. Stand down: the Disaster Declaration Authority formally declares the disaster resolved.

9.3 Documentation during the event

- A timeline log (actions, decisions, timestamps, who did what) must be kept throughout the event to support the post-incident review and any regulatory or customer reporting.

10. Security and Regulatory Considerations During Recovery

- If the disaster involves a personal data breach, the Company must meet its obligations under the Digital Personal Data Protection Act, 2023, including notifying the Data Protection Board of India and affected data principals as required; the response is coordinated with the **Data Protection Officer / privacy contact** and the Company's data breach procedure.
- Where the event is a reportable cyber incident, it must be reported to CERT-In within 6 hours as required by the CERT-In Directions, 2022.

- Recovery must not restore compromised data or reintroduce malware; backups used for restore must be verified clean. Security sign-off is required before a Tier 1 system compromised in a security incident is returned to production.
- Access provisioned for emergency recovery (break-glass accounts) must be logged and revoked once the disaster is resolved.

11. Testing and Exercises

A plan that is not tested cannot be relied upon. The Company tests this Policy and its runbooks at least **twice a year** (semi-annually), and additionally after any major architecture change.

Test type	Frequency	What it proves
Backup restore test	**Quarterly**	Backups are usable and meet the RPO
Tabletop exercise	**Twice a year**	Roles, call tree, and decision-making work; gaps in runbooks surface
Failover / functional DR test	**Twice a year**	A Tier 1 system can actually be recovered within its RTO in the recovery location
Full DR simulation	**Annually**	End-to-end recovery of critical services under realistic conditions

- Each test must record: scope, date, participants, the RTO/RPO achieved versus target, issues found, and corrective actions with owners and due dates.
- Corrective actions are tracked to closure by the DR Coordinator and reviewed at the next test.
- Where possible, tests should be conducted without notice to better reflect real conditions, subject to avoiding impact on production.

12. Link to Business Continuity Plan (BCP)

This DR Policy is the IT-recovery component of the Company's broader **Business Continuity Plan (BCP)**. The relationship is:

- The BCP addresses the whole business response to a disruption: people safety, alternate premises, manual workarounds, finance, supplier and customer continuity, and external communications.
- This DR Policy addresses recovery of IT systems and data and is invoked by the BCP whenever a disruption affects IT services.
- RTO and RPO targets in this Policy must be consistent with the recovery priorities set out in the BCP's Business Impact Analysis. Where they conflict, the BCP's business impact assessment prevails and this Policy is updated accordingly.

- A disaster declaration under this Policy and a BCP activation may occur together; the Incident Commander (technical) coordinates with the BCP crisis lead (business).

13. Post-Incident Review

After every declared disaster, and after every full DR test, the DR Coordinator conducts a blameless post-incident review within **10 business days**. The review must:

- Reconstruct the timeline from the event log.
- Compare actual recovery against RTO and RPO targets and explain any gap.
- Identify root cause and contributing factors (technical and process).
- Produce specific corrective and preventive actions, each with an owner and a due date.
- Feed lessons back into runbooks, recovery design, and, where relevant, the BCP and the system inventory.

Findings and the status of corrective actions are reported to **the CTO / management** and retained as evidence for SOC 2 / ISO 27001 audits.

14. Roles in Maintenance, Compliance and Enforcement

- The DR Coordinator is accountable for keeping this Policy, the system inventory, runbooks, and the call tree current.
- System and database owners are responsible for maintaining their own runbooks and for participating in tests.
- Failure to maintain runbooks, meet backup requirements, or participate in scheduled tests is a compliance gap and may be treated as a performance matter; repeated or wilful non-compliance may attract disciplinary action under the Company's HR policies.
- Compliance with this Policy is subject to internal and external audit.

15. Review and Governance

- This Policy is owned by **the CTO / Head of Infrastructure** and approved by **the Management / Board**.
- It is reviewed at least annually, and additionally after any major incident, significant change in IT architecture, change in hosting strategy, or change in applicable law (for example, finalisation of the DPDP Rules, which remain evolving).
- All changes are version-controlled with a documented change history (version, date, author, summary of change).
- The current approved version supersedes all previous versions.

Field	Value
Policy owner	**CTO / Head of Infrastructure**
Approved by	**Management / Board**
Version	**1.0**
Effective date	**DD-MMM-YYYY**

[COMPANY NAME]

[Add your company logo / letterhead and registered address here]

Next review due	**DD-MMM-YYYY**