CONTROLLED DOCUMENT

Disaster Recovery Plan

Sample Company LLC

Prepared as a structured operating document for internal use, review, and controlled distribution.

CIS Critical Security Controls v8.1

Review this document after major operational, regulatory, or technology changes.

Version 1.0

DOCUMENT CONTROL

RECOMMENDED

PREPARED FOR Sample Company LLC	**DOCUMENT OWNER** IT director role
CLASSIFICATION Confidential	**DISTRIBUTION** Executive Leadership, Disaster Recovery Team, IT Operations
GENERATED June 12, 2026	**NEXT REVIEW DATE** 2027-06-12
STATUS Internal Review Draft	**DOCUMENT VERSION** 1.0
TEMPLATE VERSION 1.1.0	**TARGET CONTROL LEVEL** Recommended
FRAMEWORKS CIS Critical Security Controls v8.1	

Table of Contents

Preparation
- Executive Summary
  - CIS Controls v8.1: Disaster Recovery Requirements
- Critical Systems Inventory

Recovery Discipline

Response & Failover
- Backup Strategy
  - CIS Controls v8.1: Backup and Recovery Safeguards
- Recovery Procedures
- Communication Plan
  - CIS Controls v8.1: Incident Response Communication
- Cloud Disaster Recovery
- On-Premises Disaster Recovery

Recovery & Restoration
- Identity and Directory Recovery
- Data Recovery Procedures

Testing & Maintenance
- Testing and Maintenance

Appendix: Revision History

Appendix: Glossary of Terms

Appendix: Identified Gaps and Improvement Opportunities

Appendix: CIS Controls v8.1 Regulatory Control Mapping

Appendix: Internal Review and Local Completion Checklist

INFO: IMPORTANT NOTICE — READ BEFORE USE

This document was generated as a structured starting template from the information provided. It does not constitute legal, regulatory, or professional compliance advice, and it does not by itself establish or prove compliance with any framework or regulation. Specific obligations vary by jurisdiction, industry, and circumstance. Before adopting or relying on it, have this document reviewed by qualified legal counsel and the appropriate internal owners, and adapt it to your organization's actual environment.

Preparation

Executive Summary

This Disaster Recovery Plan establishes Sample Company LLC's framework for recovering critical IT systems and infrastructure following a disruptive event. Sample Company LLC is an organization operating in the Financial Services sector with approximately 450 employees across 3 locations. Infrastructure: Hybrid (cloud + on-premises); the organization additionally depends on vendor-hosted platforms (1 SaaS category in scope).

Plan posture: 5 critical systems are in scope, the most time-sensitive of which has an RTO of < 1 hour, offsite backup copies are maintained, a DR site is designated, and the plan is tested on a documented cadence.

DR Coordinator: Disaster recovery coordinator role.

Plan owner: IT director role.

This plan applies to all IT systems, applications, and data that support critical business operations. It does not address the continuity of non-IT business processes (see the Business Continuity Plan), physical facility recovery, or employee safety procedures during a natural disaster.

This plan integrates framework-specific control content for the following selection(s) attached to this document:

- CIS Critical Security Controls v8.1

If the organization is subject to additional compliance frameworks beyond this integration, complete the Organization Information regulatory-framework step so declared scope matches integrated coverage.

CIS Controls v8.1: Disaster Recovery Requirements `cis-v8`

This Disaster Recovery Plan aligns with CIS Controls v8.1, specifically Control 11 (Data Recovery) and Control 17 (Incident Response Management). Together these controls establish the safeguards necessary to recover IT systems and data following a disruptive event.

Key CIS safeguards addressed by this plan:

- CIS Control 11.1–11.5: Data recovery process, automated backups, protection and isolation of recovery data, and recovery testing.
- CIS Control 17.1: Designated incident handling personnel with defined roles and escalation authority.
- CIS Control 17.8: Post-incident reviews to capture lessons learned and improve future recovery.
- CIS Control 17.9: Security incident thresholds that define when the disaster recovery plan is activated.

Critical Systems Inventory

The organization has identified 5 critical systems. Each carries a Recovery Time Objective and Recovery Point Objective that drive recovery prioritization, backup frequency, and infrastructure investment.

Critical Systems Inventory

SYSTEM	TYPE	HOSTING	RTO	RPO
Core transaction platform	Application	On-premises	< 1 hour	< 15 minutes
Primary database cluster	Database	On-premises	< 1 hour	< 15 minutes
Customer Portal	Web Application	Cloud platform	1-4 hours	1 hour
Messaging suite	SaaS	Hosted collaboration suite	4 hours	1 hour
Active Directory	Identity	On-premises	< 1 hour	< 1 hour

INFO: SYSTEMS WITH AGGRESSIVE RTO TARGETS

4 systems have an RTO of 4 hours or less: Core transaction platform, Primary database cluster, Customer Portal, Active Directory. These systems require pre-provisioned recovery infrastructure (hot standby, warm site, or cloud-based failover with pre-configured templates) to achieve recovery within the target window. Cold recovery approaches are unlikely to meet these targets.

System Ownership

Each critical system has a designated owner responsible for validating recovery procedures, participating in DR exercises, confirming successful recovery during actual events, and approving changes to the system's RTO and RPO targets. System owners are expected to maintain current knowledge of their system's architecture, dependencies, and recovery procedures.

System Ownership

SYSTEM	OWNER
Core transaction platform	VP of Engineering
Primary database cluster	Database Administrator
Customer Portal	Digital services lead
Messaging suite	IT Director
Active Directory	Systems Administrator

System Dependencies

System dependencies determine the order in which systems must be recovered. A system cannot be restored until all of its dependencies are available. Attempting to recover systems out of dependency order wastes time and can cause confusing failures that slow the overall recovery effort.

System Dependencies

SYSTEM	DEPENDENCIES
Core transaction platform	Directory services, primary database cluster, shared storage
Primary database cluster	Shared storage, directory services
Customer Portal	Managed database service, load balancer, identity provider
Messaging suite	Cloud directory, DNS

SYSTEM	DEPENDENCIES
Active Directory	DNS, DHCP

Infrastructure Distribution

Critical systems are distributed across multiple hosting environments. This distribution affects recovery strategy — each hosting environment may require different recovery procedures, different credentials, and different vendor contacts. The DR team must be prepared to execute recovery procedures for each environment type simultaneously.

- On-premises: 3 systems
- Cloud platform: 1 system
- Hosted collaboration suite: 1 system

Recovery Discipline

Preconditions Discipline

Every recovery activity in this plan opens with preconditions the responsible coordinator clears before work begins. A missing precondition pauses the activity; it does not lower the bar. Entering execution with unresolved preconditions is the single most common source of compounding errors during recovery, so each operational section lists its preconditions as a checklist and this discipline applies to all of them uniformly.

Failed Checks and the Three Outcomes

When an acceptance or readiness check fails, the affected system is held and the response is one of three outcomes — never a fourth:

- Close the gap and re-run the failed check
- Escalate to an executive-authorized risk-acceptance decision, bounded in time and assigned a corrective-action owner
- Remove the affected system from the plan's recoverable scope (or re-restore from an earlier known-good point where one exists)

"Proceed cautiously" is not an option. This rule appears wherever acceptance gates appear in the operational sections below.

The Restricted State

A recovered environment that has not passed acceptance is held in a restricted state — reachable by responders for investigation and validation traffic, not by users or upstream systems. Partial exposure of an untrusted environment is the failure mode that produces inconsistent behavior visible to customers, so no system leaves the restricted state until its acceptance checks pass.

Online Versus Trusted

This plan distinguishes between a service being online and being trusted for recovery. Identity services, backups, and restored datasets each cross an explicit trust gate — a verification step that separates "it responds" from "downstream recovery may depend on it." The operational sections name each trust gate where it applies.

Response & Failover

Backup Strategy

The backup strategy below sets the floor for recoverable data loss. Each parameter is paired with the operational verification needed to move from "backup exists" to "backup is trusted for recovery."

Backup Configuration Summary

Backup Configuration Summary

PARAMETER	CURRENT CONFIGURATION	OPERATOR NOTES
Primary backup solution	Centralized backup platform with immutable object storage	Documented at a generic category or placeholder level in this plan; operator-specific tooling and versions are maintained in the organization's backup runbook
Backup frequency	Every 4 hours for Tier 1 systems, daily for Tier 2	Sets the floor of the effective RPO for the systems covered; the Critical Systems Inventory records the per-system RPO this frequency is expected to satisfy
Offsite / off-network copies	Yes: encrypted copies replicated to immutable offsite object storage	Guards against site-level events and against reachable-storage destruction; independence from the primary copy is verified rather than assumed
Retention period	30 days for daily backups, 12 months for monthly archives	Must cover delayed-discovery scenarios — data corruption or intrusion that is not detected until days or weeks after it occurs
Encryption	Yes: AES-256 encryption at rest and TLS 1.3 in transit	Applies to both at-rest and in-transit paths where the plan requires it; encryption keys are held in a survivability path separate from the backups themselves
Restore testing cadence	Monthly: automated restore verification with integrity checks	Execution detail is owned by the Testing & Maintenance section; this row records the cadence the backup program is measured against

Backup Methods

The organization uses the following backup methods. Each method carries different restore characteristics — time to first usable data, granularity of point-in-time recovery, and dependence on other methods in the chain — and the Data Recovery Procedures section documents the per-method restore path:

- Full backup (weekly)
- Incremental backup (every 4 hours)
- Transaction log backup (every 15 minutes for databases)

Backup Readiness Dimensions

The readiness of the backup program is verified as an ongoing invariant, not only during restore exercises. Each row below names a dimension the plan guards against, the question that must have a documented answer, the most common silent failure for that dimension, and how verification is performed:

Backup Readiness Dimensions

DIMENSION	READINESS QUESTION	COMMON SILENT FAILURE	VERIFIED BY
Coverage	Are every system and dataset in the Critical Systems Inventory actually included in the backup program, rather than assumed to be?	A system was added since the last inventory review and is not yet covered by any backup job; the gap is invisible until a restore is attempted	Backup-coverage reconciliation against the Critical Systems Inventory on a documented cadence; additions and removals trigger coverage review
Retrievability	Can responders reach the backup data during an event — including when the primary environment, primary network, or primary identity plane is unavailable?	Backup-console credentials are stored only in the system that is offline; backup storage is reachable from the primary network but not from the responders' survival-mode location	Backup-retrieval path is exercised from a responder location that does not depend on the primary being up; credentials are stored in a survival-mode path
Restorability	Has the most recent backup actually been restored successfully, or is "restorable" inferred from the backup job exit code?	Backup jobs report success for months; the first attempted restore fails on a condition the job never checked (corrupted archive, missing dependency, incompatible software version)	A restore test has been completed within the documented testing cadence owned by the Testing & Maintenance section, with results captured in the recovery record
Integrity	Is each backup confirmed as internally consistent and not silently corrupted between creation and the next restore attempt?	Bit rot on cold storage or a transient backup-software defect produces backups that pass signature checks but fail application-level consistency when restored	Integrity checks at the backup layer (checksums, manifest validation) plus application-level checks at restore time; the Data Recovery Procedures section owns the restore-time acceptance rules
Encryption posture	If backup media is lost, stolen, or copied without authorization, is the data readable to whoever has it?	Encryption was enabled once and silently regressed when a backup job was reconfigured; keys are stored on the same system as the encrypted backups	Encryption state is a first-class monitored attribute of every backup job; keys are stored in a survivability path separate from the backups they protect
Offsite / off-network copy	Does at least one backup copy survive a site-level event, a network-wide compromise, or a ransomware-initiated deletion of reachable storage?	Offsite copy is on the same network segment as primary backups; ransomware reaches it through the same credentials or API as the primary copy	Offsite copy is reachable only through a separate credential set, separate network path, or physical-media flow; its independence is exercised, not assumed

DIMENSION	READINESS QUESTION	COMMON SILENT FAILURE	VERIFIED BY
Immutability and access control	Can an attacker with backup-administrator access delete or encrypt every copy of the backups, or is at least one copy protected by a write-once or time-locked policy?	Backup-administrator credentials are the same identity used for day-to-day administration; an identity compromise reaches both production and every backup	Backup-administrator identities are separate from general administrative identities; at least one backup copy is subject to an immutability or retention-lock policy that the administrator identity cannot override

Offsite Copy Status

The organization maintains offsite copies as documented in the configuration summary above. Offsite copies protect against site-level events; their effectiveness depends on independence from the primary copy — a separate credential set, a separate network path, or a physical-media flow. Independence is verified, not assumed.

Encryption Posture

Backup data is encrypted as documented in the configuration summary. Encryption prevents unauthorized access to backup media in the event of loss, theft, or improper disposal. Encryption keys are held in a survivability path separate from the backups they protect — a key lost alongside its backups renders the backups unrecoverable regardless of the encryption strength used.

Restore Testing Status

Backup restore testing is performed on a monthly: automated restore verification with integrity checks cadence. The Testing & Maintenance section of this plan owns the exercise discipline — what is tested, who attends, how results are captured — and the Data Recovery Procedures section owns the per-source restore acceptance rules applied during each test. This section records only that the cadence is in place.

Acceptance — From "Backup Exists" to "Backup Is Trusted for Recovery"

A backup "exists" when the backup job completes with a success signal. It is "trusted for recovery" only after the checks below are satisfied. A backup that has not cleared these checks is treated as present but not trusted — it is not counted toward recovery planning, and the gap is tracked as an open finding against this plan until closed:

- Coverage reconciled — the system or dataset is actually in scope of a backup job, verified against the Critical Systems Inventory, not assumed from the existence of adjacent coverage
- Retrievability path usable — the backup can be reached from a responder location that does not depend on the primary environment being up; credentials are reachable from a survival-mode path
- Restore demonstrated — at least one recent restore has completed successfully against the current backup configuration, within the cadence owned by the Testing & Maintenance section; backup-job success codes are not accepted as a substitute
- Integrity verified — backup-layer integrity checks (checksums, manifest validation) are in place and current; application-level consistency is verified at restore time per the Data Recovery Procedures section
- Encryption posture matches the plan requirement — state is monitored, not assumed, and key custody is in a survivability path separate from the backups
- Offsite copy present and independent — an offsite or off-network copy exists and is reachable through a credential set and network path distinct from the primary copy
- Owner acceptance recorded — the accountable system owner acknowledges the backup posture for each in-scope system in a documented review, rather than inferred from the absence of complaints

Failed backup checks follow the three outcomes in Recovery Discipline.

Retention and Disposal

Backups are retained for 30 days for daily backups, 12 months for monthly archives. The retention period accommodates delayed-discovery scenarios — data corruption, intrusion, or policy events that are not detected until days or weeks after they occur — so the retention window has to exceed the longest plausible detection lag, not the average one. Disposal of backups at end of retention follows the organization's data-handling policy for the media class involved; this plan records the retention window and defers the disposal method to that policy so a change in media or tooling does not require a plan revision to stay accurate.

CIS Controls v8.1: Backup and Recovery Safeguards cis-v8

CIS Control 11 defines five safeguards that form the foundation of a resilient backup and recovery strategy. Each safeguard must be addressed in the disaster recovery plan to ensure data can be restored reliably.

- CIS Control 11.1: Document the data recovery process including scope of covered systems, prioritization criteria, and personnel responsible for recovery execution.
- CIS Control 11.2: Automate backups for all in-scope enterprise assets. Manual backup processes introduce human error and are insufficient for CIS compliance.
- CIS Control 11.3: Protect recovery data with encryption at rest and in transit. Apply access controls equivalent to those on production data.
- CIS Control 11.4: Maintain at least one isolated (air-gapped or offline) copy of recovery data. This safeguard is critical for ransomware resilience.
- CIS Control 11.5: Test data recovery at least quarterly. Recovery tests must validate that data can be restored within the defined RTO and RPO.

Recovery Procedures

Recovery procedures must be specific enough to execute under pressure by personnel who may be stressed, sleep-deprived, and working without access to their usual tools. "Restore from backup" is not a procedure; it is a placeholder. A usable procedure names the backup location, the restore command or console path, the expected completion time, and the validation step that confirms the restore was successful. Systems are restored in tiers, each tier has explicit prerequisites, and no tier advances until the tier beneath it is healthy under the same monitoring coverage the organization had before the event.

Preconditions — Before Beginning Recovery

The DR Coordinator clears the following preconditions before the recovery execution begins.

A missing precondition pauses the execution (see Recovery Discipline).

- Plan formally activated by the authorized role; the activation is recorded in the incident record with timestamp and declaring party
- Recovery team assembled on a known communication channel that is reachable even if the primary environment is unavailable — see the Communication Plan for the designated channels
- DR-site access confirmed — credentials, break-glass accounts, console URLs, and retained-provider contacts are reachable from a location that survives the disaster
- Most recent backup confirmed present, readable, and within the documented RPO window; a backup that has never been restored is assumed broken until a restore proves otherwise
- Primary environment confirmed unavailable rather than degraded — a failover executed against a partially working primary causes split-brain recovery and data divergence
- Executive authorization captured where the plan requires it for customer-facing disruption, spend outside existing retainers, or public statements
- Recovery record opened in the incident or DR tracker with timestamp, responders, and the decision to enter execution

Recovery Site

The organization uses cloud-based failover as its primary disaster recovery site. Recovery time depends on the speed of provisioning cloud resources, restoring data from backups or replicas, and re-establishing network connectivity. Pre-staged infrastructure-as-code templates reduce provisioning time by replacing ad-hoc console work with a reviewed, version-controlled deployment path; the templates themselves are treated as part of the recovery artifact set and are tested on the same cadence as the rest of this plan.

Recovery Sequencing

Systems are restored in dependency order — infrastructure first, then identity, then data, then applications, then secondary services. Each tier has explicit prerequisites and an advancement check; no tier advances until the tier beneath it is healthy under the same monitoring coverage the organization had before the event.

Recovery Sequencing

TIER	REPRESENTATIVE SYSTEMS	PREREQUISITES	ADVANCEMENT CHECK
1. Network and edge	DNS, DHCP, firewalls, routing, VPN, site-to-site links, out-of-band management	Break-glass console access to core network devices is reachable; pre-staged low-TTL DNS entries are present	Responders can reach the DR environment from a trusted client and from each other over the recovery channel
2. Identity and authentication	Directory services, SSO, MFA, privileged-access tooling, secrets store, certificate authority	Tier 1 complete; recovery credentials for identity systems are available from a location that survives the disaster	Responders and downstream systems can authenticate using the recovered identity plane without using emergency shared credentials
3. Data tier	Databases, object storage, shared file systems, message queues, replication endpoints	Tier 2 complete; most recent backup confirmed present and readable; application-level integrity checks are documented for the datasets being restored	Restore completes, integrity checks pass, and expected row counts or size reconcile against a known-good baseline
4. Core application tier	Revenue-generating and customer-facing applications, line-of-business systems, email and collaboration	Tier 3 complete; application configuration points at the recovered identity and data endpoints; secrets are present in the recovered secrets store	Each application passes the pre-defined recovery smoke test against the recovered dependencies, and its owner acknowledges the result
5. Secondary and supporting services	Reporting, analytics, monitoring dashboards, support tooling, non-critical integrations	Tier 4 complete; monitoring and alerting re-pointed at the recovered environment before these services take load	No outstanding tier-specific findings blocking return to steady state; all earlier tiers remain green under production load

The failed tier stays restricted until acceptance passes (Recovery Discipline).

Dependent tiers do not begin restoration until the failed tier passes. The DR Coordinator records the hold, the gap, and the mitigation path in the recovery record rather than deciding to proceed verbally; a tier that is silently skipped is the failure mode that surfaces hours later as data integrity or identity inconsistency.

Failover Procedure

The following procedure describes the high-level failover process for the organization's most critical system. Responders record the actual time each step takes during execution so that variance against the documented RTO can

be measured in the recovery record, rather than reconstructed from memory afterward:

1. DR Coordinator declares plan activation and assembles recovery team
2. Verify DR region connectivity and IAM permissions
3. Execute Terraform apply to provision recovery infrastructure
4. Restore databases from most recent S3 backup
5. Update DNS records to point to DR environment
6. Validate application functionality with smoke tests
7. Notify stakeholders that services are restored on DR infrastructure

Each step should name the system or endpoint it targets, the command or console path used, and the success signal that confirms the step worked. A step whose success signal is "it seems fine" is a step the next responder cannot repeat; rewrite it before the next revision of this plan.

Network Recovery

Network connectivity is restored first because all other recovery activity depends on it. The procedure below must be executable even if the primary network-management tooling is unavailable — console and out-of-band management paths for core network devices are pre-staged, with credentials reachable from a location that survives the disaster.

1. Verify VPN tunnel to DR site is operational
2. Update DNS records with pre-staged low-TTL entries
3. Confirm firewall rules and security groups in DR region
4. Test connectivity from client networks to DR environment
5. Verify load balancer health checks are passing

Post-Recovery Data Validation

Data validation is performed per tier, not only at the end of the recovery. The following validation methods are used for restored data:

- Verify database consistency checks pass (pg_isready, amcheck verification on critical indexes)
- Compare row counts for critical tables against last known baseline
- Run automated smoke test suite against restored applications
- Confirm API integrations with external partners are functional
- Validate file integrity using SHA-256 checksums against backup manifest

A system is considered successfully recovered only when all validation checks applicable to its tier pass and the accountable system owner acknowledges the result against a pre-defined recovery smoke test. The DR Coordinator captures the sign-off in the recovery record before advancing to the next tier.

The unvalidated system stays restricted until acceptance passes (Recovery Discipline).

Failed validation checks follow the three outcomes in Recovery Discipline.

For data validation, the risk-acceptance outcome takes the form of a documented data-loss decision with explicit executive authorization.

Recovery Team Cadence

During active recovery, the DR Coordinator runs a brief responder sync at a cadence proportional to how fast the situation is changing — more frequent while tier transitions are in flight, tapering once the recovered environment is under load. Each sync produces a short written entry in the recovery record naming which tiers are green, which are in progress, the estimated time to next milestone, and any blockers requiring escalation.

External communication — customer notification, regulator notification, press statements, and partner notification — is owned by the Communication Plan section of this plan and is not improvised by responders during execution. The DR Coordinator routes status upward to the executive sponsor at the cadence set in the Communication Plan, and the Communication Plan is the single source for external-facing messaging decisions during the event.

Communication Plan

A single designated DR Coordinator owns all communications during a recovery event. This section names that role, the notification channels in use, and the cadence at which stakeholders receive updates.

DR Coordinator

The DR Coordinator is the single point of authority during disaster recovery operations. This individual is responsible for declaring plan activation, assembling the recovery team, directing recovery priorities, authorizing changes to the recovery sequence, and communicating status to executive leadership. All recovery decisions are routed through the DR Coordinator to prevent conflicting actions.

DR Communication Contacts

ROLE	TITLE / NAME	PHONE	EMAIL
DR Coordinator	Disaster recovery coordinator role	[Fill locally: Disaster recovery coordinator role emergency phone]	[Fill locally: Disaster recovery coordinator role shared mailbox]

If the DR Coordinator is unavailable, the designated alternate should assume this role within 30 minutes. The alternate should be identified and documented before a disaster occurs, and should participate in all DR testing exercises.

Employee Notification

When a disaster is declared, all employees must be notified as quickly as possible. The notification should include: (1) confirmation that a disaster has been declared, (2) whether employees should report to work or work remotely, (3) which systems are affected, (4) estimated time to restoration if known, and (5) where to find ongoing updates. Notification channels in priority order:

1. Automated alert to the DR team
2. SMS broadcast to all IT staff through the approved notification tool
3. Phone tree for executive leadership

Multiple notification channels are configured to provide redundancy. If the primary channel is unavailable (e.g., Slack is down during an internet outage), the DR Coordinator should immediately switch to the next channel in the priority list. At least one notification method should not depend on the organization's own IT infrastructure — SMS and phone calls continue to work when corporate email and chat are down.

Escalation Procedures

The following escalation timeline should be followed once a disaster is declared:

Escalation Timeline

TIMEFRAME	ACTION	RESPONSIBLE PARTY
0-15 minutes	DR Coordinator notified and assesses the situation	On-call personnel
15-30 minutes	DR team assembled; initial impact assessment completed	DR Coordinator
30-60 minutes	Recovery operations begin; first status update to leadership	DR Coordinator
Every 30 minutes	Status updates to leadership and affected stakeholders	DR Coordinator
4 hours	If RTO at risk, escalate to executive team for resource decisions	DR Coordinator
RTO exceeded	Executive briefing; evaluate alternate recovery options	Executive leadership

External Communication

External communication — to customers, partners, regulators, and the public — requires careful coordination. All external communications during a disaster must be approved by the DR Coordinator or a designated communications lead before release. Premature or inaccurate external communication can cause reputational damage, contractual liability, or regulatory complications.

Customer / client communication approach: Status page updated within 30 minutes of declaration. Direct email to enterprise clients within 1 hour. Public website banner if outage exceeds 4 hours. All external communications approved by DR Coordinator and Legal.

External communications should be factual, avoid speculation about root cause (which may not be known during recovery), provide a realistic timeline for resolution, and include a point of contact for follow-up questions. Retain copies of all external communications as part of the post-incident record.

Critical Vendor Contacts

Vendors providing services that may be needed during disaster recovery. The role/service is documented; provider names and contact details are filled locally.

Critical Vendor Directory

VENDOR	SERVICE	CONTACT	PHONE	SLA
-	Cloud Infrastructure	Priority support desk	-	< 15 minutes for critical events
-	Backup & Recovery	Priority support desk	-	< 30 minutes for severity 1
-	Primary ISP	Network operations center	-	4-hour on-site for circuit issues

Phone and email columns are completed locally on the exported copy to preserve the hosted working-copy privacy contract.

Recovery Completion Communication

When recovery is complete and all systems have been validated, the DR Coordinator issues a formal recovery completion notice. This notice should confirm: (1) all critical systems have been restored and validated, (2) normal operations have resumed, (3) any temporary workarounds still in effect, and (4) the date and time of the post-incident review meeting.

CIS Controls v8.1: Incident Response Communication `cis-v8`

CIS Control 17 requires structured incident communication, post-incident review, and severity thresholds that trigger the disaster recovery plan.

- CIS Control 17.1: Designate personnel responsible for incident handling during a disaster, with clear authority to make recovery decisions.
- CIS Control 17.6: Establish out-of-band communication channels for use when primary infrastructure is unavailable during a disaster.
- CIS Control 17.7: Conduct routine disaster recovery exercises at least annually, including communication and coordination components.
- CIS Control 17.8: Perform post-incident reviews after every DR activation to identify process improvements and update the plan.
- CIS Control 17.9: Define security incident thresholds (severity levels, impact criteria) that trigger DR plan activation.

RECOMMENDATION: POST-INCIDENT REVIEW

CIS Control 17.8 requires post-incident reviews after DR events. Establish a standard template for lessons-learned documentation and schedule reviews within 72 hours of incident resolution.

Cloud Disaster Recovery

Cloud disaster recovery has failure modes that have no on-premises equivalent. IAM policies that look correct fail under recovery conditions because the recovery region was never in the policy's resource scope. Service quotas that were adequate for the last exercise are insufficient at production scale. Cross-region replication that was configured years ago has silently stopped. These failures are invisible until a recovery is attempted, which is why the readiness of the recovery environment is as load-bearing as the recovery procedure itself — and why both are documented below.

Preconditions — Before Beginning Cloud Failover

The DR Coordinator clears the following preconditions before any cutover action begins.

A missing precondition pauses the cutover (see Recovery Discipline).

- Recovery environment identified — the recovery region, account, subscription, or project the failover targets is named in this plan and verified reachable from the responders' locations
- Out-of-band console access confirmed — credentials reach the recovery environment through a path that does not depend on the primary environment being up
- Infrastructure-as-code artifacts current — the templates or scripts that will provision or scale the recovery environment are the versions under review, and their last successful dry-run is recent enough to trust
- Replication and snapshot freshness verified — the data the recovery environment will depend on is within the documented RPO window; a stale replica is treated as a data-loss decision, not absorbed by the responder
- Quotas and capacity pre-checked — service quotas in the recovery region cover the production footprint, not the exercise footprint, and any pre-negotiated quota-increase path is in hand
- Billing alert posture set — cost-management alerts are set against DR spend before activation so the event does not produce a surprise bill weeks later; cost-management detail is documented below
- Cloud-failover record opened — the recovery environment, responders, start time, and the decision to enter cutover are recorded in the recovery record before the first cutover action runs

Recovery-Environment Readiness

The readiness of the recovery environment is verified as an ongoing invariant, not only during DR exercises. Each row below names a dimension that must be verified, the failure mode that the organization is guarding against, and how verification is performed. Readiness gaps in any dimension are treated as open findings against this plan until closed:

Recovery-Environment Readiness

DIMENSION	READINESS QUESTION	MOST COMMON FAILURE MODE	VERIFIED BY
Region / zone parity	Is the recovery region or zone capable of running the workloads being failed over at the volume required?	Recovery region has a narrower service catalog or smaller machine-type selection than primary; certain workloads cannot be provisioned	Pre-recovery inventory of required services and machine types against the recovery region, re-verified whenever the workload mix changes

DIMENSION	READINESS QUESTION	MOST COMMON FAILURE MODE	VERIFIED BY
Account / subscription / project parity	Are the recovery account, subscription, or project configured with the tenancy, billing relationship, and organization-policy boundaries the workloads require?	Recovery environment belongs to a different billing account or organizational unit; organization policies block provisioning or cross-tenancy access	Recovery-environment identifier documented in this plan; a provisioning dry-run is executed against it on a rolling schedule
IAM and role scope	Do the IAM roles used during recovery grant the permissions required in the recovery region without granting permissions they should not carry during steady state?	Role policies reference only the primary region in resource scopes; permission checks pass in testing but fail when the recovery region is actually used	Explicit test of each recovery-role permission against the recovery-region resources during the exercise cadence owned by Testing & Maintenance
Network reachability	Can responders and dependent systems reach the recovery environment — VPN, peering, or dedicated interconnect — from the locations they will operate from during the event?	Peering or dedicated-interconnect configuration is primary-region-only; the recovery region is reachable by console but not by the client or partner networks that need to consume it	Network-path test from each responder location and each in-scope partner network to the recovery environment
Data replication and snapshots	Is the data the recovery environment depends on actually present in the recovery region within the documented RPO window?	Cross-region replication was configured years ago, silently stopped, and is discovered to be stale only when failover is attempted	Replication-freshness check is a continuous signal, not a one-time verification; its absence is itself a high-severity finding
Service quotas and capacity	Are the service quotas in the recovery region sufficient to provision the workloads at production scale, not just at test scale?	Quotas were adequate for the exercise footprint but block the production footprint when real failover is attempted; quota-increase requests take hours to days	Quota-versus-footprint inventory documented in this plan; quota-increase requests are pre-negotiated where the provider requires them
DNS and traffic routing	Are the DNS records and traffic-routing policies cut-over-ready — low TTLs, health checks, and automated or documented manual cutover paths?	TTLs are long enough that cutover propagation exceeds RTO even when the cutover action itself is fast; health checks target endpoints that do not exist in the recovery environment	TTL posture is reviewed against RTO as part of the exercise cadence; health checks and routing policies are tested in the exercise scope, not only in production

Cutover Sequencing

Within the cloud failover, the cutover itself is sequenced so that each layer is healthy before the next layer depends on it. This intra-cloud ordering is executed inside the broader tier sequence owned by the Recovery Procedures section of this plan:

1. Network cutover — routing, peering, and any dedicated-interconnect paths into the recovery environment are brought up and tested from each responder and dependent-system location before identity or data actions begin
2. IAM and identity cutover — recovery-scope IAM roles are assumed by responders and service identities as documented in the plan; the Identity and Directory Recovery section owns identity restoration itself, and this step only re-points operations at the already-trusted identity tier

3. Data cutover — the recovery environment is pointed at the replicated data tier or the restored snapshot set; the Data Recovery Procedures section owns the acceptance rules for each restore path and its hold-state rule applies here
4. Application cutover — application instances are provisioned from infrastructure-as-code or pre-staged images against the recovered identity and data tiers; each application passes its pre-defined recovery smoke test against its recovered dependencies before traffic is sent to it
5. DNS and traffic cutover — traffic-routing changes are made last and only after the steps above are healthy; propagation time against the documented TTL is part of the effective RTO and is captured in the recovery record

Each provider-scoped procedure below is executed against the recovery environment identified in the Preconditions section and proceeds through the cutover sequencing above.

Amazon Web Services (AWS) Recovery Procedures

Actual step timings are captured in the recovery record so variance against the documented RTO is analyzable in the after-action review:

1. Authenticate to the secondary cloud region using break-glass credentials
2. Run the approved infrastructure template to provision compute, database, and load-balancing services
3. Restore the managed database from the most recent replicated snapshot
4. Deploy application services from the approved artifact registry
5. Update failover routing and health checks
6. Verify monitoring dashboards and alerts are active

Recovery-Target Validation

Recovery is measured against the documented RTO and RPO for the workloads in scope. Measurements are taken during the event itself — start time from declaration, restore points from backup and replication metadata, and end time from the first production-trusted transaction — rather than reconstructed afterward. Failure to meet a target is a finding captured in the after-action record, not a signal to lower the target:

1. Compare the restored database recovery point to the RPO target
2. Measure time from declaration to first successful health check against the RTO
3. Run the application test suite to validate data integrity
4. Document actual recovery metrics for post-incident review

Validation results are recorded after every exercise and every actual event. Over time, these measurements are the primary evidence that recovery times are improving, stable, or degrading, and are the basis for infrastructure-investment decisions rather than replacing that evidence with estimates.

Cost Management During DR Activation

Running production and recovery environments simultaneously can compound cloud spending quickly. Without active cost controls, a multi-day recovery event can generate a bill comparable to the incident it recovered from. The controls below are in place and are applied from activation through the decommissioning of the recovery footprint:

1. Set a cost alert at 200% of normal daily spend
2. Use reserved baseline capacity with burst controls during failover
3. Tag all DR resources for cost tracking
4. Schedule an automatic scale-down review at the 48-hour mark
5. Notify Finance of DR activation for accrual purposes

Acceptance — From "Recovery Environment Up" to "Trusted for Production Traffic"

A recovery environment is "up" when provisioning completes without error. It is "trusted for production traffic" only after the checks below pass.

The recovery environment stays restricted until acceptance passes (Recovery Discipline).

- Provisioning completed without error — infrastructure-as-code runs, console actions, and automated scripts all produced expected exit status; warnings or partial-success signals are investigated, not absorbed
- Identity plane authenticated against the recovery environment — responders and service identities are signing in through the trusted identity tier rather than through emergency shared credentials
- Data plane readable — restored databases, object stores, and caches return representative queries successfully; row counts and sizes reconcile against the last known-good baseline within a defensible tolerance
- Application smoke tests pass — representative transactions execute end-to-end through the recovered stack, not just a health-check endpoint
- Integration points healthy — dependent systems outside the recovery environment (partners, payment processors, customer identity providers, analytics sinks) can authenticate to and transact against the recovered stack
- Monitoring and alerting re-pointed — logging, detection, and alerting for the recovered workloads flow to the organization's SIEM or equivalent before traffic is admitted
- Owner acceptance recorded — the accountable system owner acknowledges the result against a pre-defined recovery smoke test, and the DR Coordinator captures the acceptance in the recovery record before cutting over user traffic

Failed cutover checks follow the three outcomes in Recovery Discipline.

In the cloud context, the re-provision outcome means re-provisioning from an earlier known-good infrastructure-as-code configuration point rather than patching the current environment in place.

Failback to Primary Environment

Once the root cause of the disaster has been addressed and the primary environment is confirmed stable, workloads are migrated back from the recovery environment. Failback is often more complex than failover because the recovery environment has accumulated data and state during the event that must be reconciled back to the primary without loss. The following invariants govern failback:

- Primary verified stable — the primary environment has passed the same acceptance checks applied to the recovery environment above, not merely "booted and appears healthy"
- Data divergence handled explicitly — the dataset written in the recovery environment during the event is reconciled back to the primary through a defined merge or authoritative-copy rule, not through ad-hoc comparison
- Cutback sequenced in reverse — traffic-routing, application, data, IAM, and network steps are reversed in the opposite order used during cutover; each layer is verified before the next layer's cutback proceeds
- Observation window — production traffic on the primary is observed for a documented window before the recovery footprint is scaled down, so a regression on the primary surfaces while the recovery environment is still usable
- Controlled decommissioning — the recovery footprint is scaled down or decommissioned on a documented plan with cost and audit record, not by silent cleanup; artifacts created during the event (logs, snapshots, configuration exports) are preserved for the after-action review

On-Premises Disaster Recovery

On-premises disaster recovery carries physical preparedness that cloud DR does not. There is no equivalent of provisioning a new host in ninety seconds when a server fails; power, HVAC, facility access, and hardware supply all have to be verified before they are needed, not assessed during an event. The readiness of the facility itself is as load-bearing as the recovery procedure executed inside it, which is why both are documented below. If the facility is itself compromised — fire, flood, structural damage, loss of utility power beyond the UPS-and-generator envelope — the plan's cloud or alternate-site path is used instead of this one.

Preconditions — Before Beginning On-Premises Recovery

The DR Coordinator clears the following preconditions before on-premises recovery work begins.

A missing precondition pauses the recovery (see Recovery Discipline).

- Facility determined safe to enter — no active environmental hazards (water intrusion, fire damage, structural damage, air-quality concerns) that would make the server room unsafe; confirmed by on-site observation or by the facility-services contact rather than assumed
- Power envelope verified — utility-power status is known, UPS runtime is sufficient to either complete recovery or perform an orderly shutdown, and the generator is operational and fueled if utility power is out
- Thermal envelope verified — HVAC is functioning or the equipment load has been reduced to a level the degraded cooling capacity can sustain; rack temperatures are within equipment-safe range
- Facility access paths confirmed — responders can reach the server room through the documented access path (primary key, bypass key, or after-hours facility contact), not just the normal weekday path
- Out-of-band management reachable — baseboard management, KVM, or remote-console access to critical hosts is usable over a path independent of the production network, and credentials are reachable from a location that survives the outage
- Hardware supply known — on-site spare inventory or vendor-response contracts are in hand for the components the recovery is expected to replace; end-of-life compatibility gaps are resolved before the recovery action, not during it
- Recovery record opened — the facility state, responders present, start time, and decision to enter execution are recorded before the first recovery action runs

Facility and Physical Readiness

Facility readiness is verified as a continuous invariant, not only during exercises. Each row below names a readiness dimension, the question the organization guards against, the most common way that dimension fails silently, and how verification is performed. Readiness gaps in any dimension are treated as open findings against this plan until closed:

Facility and Physical Readiness

DIMENSION	READINESS QUESTION	MOST COMMON FAILURE MODE	VERIFIED BY
Facility access	Can responders reach the server room at the time of the event — including after-hours, when electronic access systems are down, and when the primary key-holder is unavailable?	Electronic access control depends on a system that is itself part of the outage; emergency bypass keys are held by a single person who is unreachable	Documented bypass key custody with a named primary and backup; after-hours facility contact path; access-path test performed on the exercise cadence owned by Testing & Maintenance
Power	Will the facility carry equipment through utility-power interruption long enough to either complete recovery or perform an orderly shutdown, and is the generator operational and fueled?	UPS battery runtime is shorter than assumed because units were never load-tested; generator fuel is not checked against the cadence the plan assumes	UPS runtime is load-tested on a documented cadence, not assumed from nameplate; generator fuel level and start-test results are captured in the plan's maintenance log
Thermal / HVAC	Does the server room stay within equipment-safe temperature range when primary HVAC is degraded, and how long can equipment run before thermal shutdown becomes a risk?	HVAC degradation is not observable until rack temperatures reach alert thresholds; equipment shuts down in uncoordinated order, complicating recovery	Environmental monitoring is in place and alerts through a channel that does not depend on the affected systems; HVAC failure runbook is documented separately from general DR

DIMENSION	READINESS QUESTION	MOST COMMON FAILURE MODE	VERIFIED BY
Out-of-band management	Can responders reach baseboard management or equivalent console-level access on critical hosts without physical presence, over a path independent of the production network?	Out-of-band management was configured once and is reachable only from the primary office network; credentials are stored only in the system that is offline	Out-of-band management is accessible from responders' remote locations over a dedicated path; credentials are stored in the secrets-management path that survives the primary being down
On-site spares and vendor supply	Does the organization have the replacement components required to recover within the documented RTO, either on-site or from a vendor whose response is under contract?	Hardware is end-of-life and the vendor no longer stocks replacement parts; on-site spares drift out of compatibility with deployed hardware generations	Spare-parts inventory and vendor-response-time documentation in this plan are re-verified on a documented cadence and whenever the hardware generation changes
Network and ISP reachability	When on-prem recovery completes, can the recovered systems reach users, partners, and any cloud-side dependencies the organization operates against?	ISP circuits are single-homed; recovery completes the server side but the connectivity side is assumed rather than verified	Primary and backup ISP paths, internal routing, and any required partner-network reachability are included in the exercise scope — not only internal server restoration
Responder access and after-hours coverage	Will enough qualified responders reach the facility and the recovery surface during the hours the outage occurs, including nights, weekends, and holidays?	After-hours coverage is assumed rather than documented; the one responder with the key is the one responder on vacation	On-call and after-hours contact paths are documented in the Contact Directory and exercised at a documented cadence; multiple responders can perform each critical recovery action

Restoration Layer Sequencing

Inside the facility, recovery advances through layers in the order below. Each layer's advancement check must pass before the next layer depends on it; a layer that fails advancement is held in a restricted state until the gap is closed. This intra-facility sequence runs inside the broader DR tier sequence owned by the Recovery Procedures section of this plan:

1. Facility and power — the physical environment is verified per the preconditions above; equipment is safe to power up and temperature is stable at equipment-safe range
2. Network fabric — core switching, routing, firewalls, and internal DNS are brought up and verified from responder workstations and out-of-band management before hosts depend on them
3. Hardware and hosts — physical servers and any hardware that failed are replaced, posted, and confirmed healthy at the BMC / out-of-band level before operating systems boot onto them
4. Virtualization layer — hypervisors and cluster management are restored and verified; VM configuration is present before the data tier is pointed at VM storage
5. Storage and data tier — shared storage, SAN/NAS, and backup-source media are brought up; the Data Recovery Procedures section owns the per-source restore rules and acceptance criteria that apply here
6. Services and applications — identity services (per the Identity and Directory Recovery section), then line-of-business applications, then secondary and supporting services — each against the acceptance gate below before it takes user traffic

Hardware Failure Response

Hardware-failure response differs depending on whether the failure is a single replaceable component (disk, power supply, memory, network card) or a full-system failure that requires replacement and a bare-metal restore. The operator procedure below executes within the preconditions above and the restoration-layer sequence; actual step timings are captured in the recovery record so variance against the documented RTO is analyzable in the after-action review:

1. Assess the failed component and determine whether the issue is isolated or a full system failure
2. For component failure: replace from on-site spare inventory and restore from backup if needed
3. For full system failure: provision a replacement server from the spare pool
4. Perform a bare-metal restore using approved recovery media
5. Verify the system boots and rejoins core services
6. Run application-specific post-restore validation

Hyper-V Recovery

The Hyper-V-scoped procedure below is executed within the Restoration Layer Sequencing above; cluster or standalone host recovery completes before VMs are imported, and VMs start in the dependency order documented in the Server Restoration Sequence:

1. Access the virtualization management console
2. Import the virtual machine from backup to an available host
3. Verify virtual switch and network adapter configuration
4. Start the virtual machine and confirm it registers with core services
5. Validate application services start correctly

Server Restoration Sequence

Restoring servers out of dependency order is one of the most common and costly mistakes in on-premises recovery — an application server that boots before its database or before the identity tier it authenticates against fails on first request, wasting the effort spent restoring it. The sequence below reflects the dependency analysis from the organization's Critical Systems Inventory. Responders follow this sequence as ordered and record variance (including any conscious re-ordering) in the recovery record:

1. Active Directory with DNS and DHCP (identity and name resolution for all tiers)
2. Primary database cluster
3. Core transaction platform
4. Customer Portal
5. Messaging suite (SaaS, provider-side; verify access)
6. Monitoring and logging infrastructure

This sequence is kept accessible offline — in an offline document store, a printed copy in the server room, or an equivalent location that survives the outage of the systems being recovered. The exercise cadence owned by the Testing & Maintenance section validates that the sequence is still correct against current dependencies whenever the environment changes.

Physical Access Requirements

Physical access arrangements supporting this recovery are documented below. Where the arrangement relies on an emergency-bypass path (sealed envelope, facility-services hand-off, or equivalent), that path is exercised on a documented cadence so responders know it works when needed, not only in concept:

Server room access requires badge and secondary verification. Emergency bypass instructions and keys are stored in the approved offline facilities packet. After-hours access uses the documented security escalation path. A printed copy of the DR plan is stored with the local continuity binder.

Spare Hardware and Emergency Procurement

Spare inventory and emergency-procurement arrangements supporting this recovery are documented below. Stocked spares are inventoried on a documented cadence so compatibility with deployed hardware generations stays current, and emergency-procurement paths (vendor response contracts, reseller accounts, hand-carry arrangements) are verified before they are needed:

1. Two spare rack servers are stored in the primary equipment room
2. Spare disk shelves and SSDs are maintained for replacement use
3. Spare network switches are available for emergency swap
4. Emergency hardware procurement uses the approved rapid-replacement agreement
5. Inventory is audited quarterly

Acceptance — From "Recovered" to "Trusted for Production Traffic"

A system is "recovered" when it boots and presents its service. It is "trusted for production traffic" only after the checks below pass.

The restored system stays restricted until acceptance passes (Recovery Discipline).

- Hardware clean — POST, vendor self-tests, and firmware versions reconcile against the baseline the plan documents; component-level telemetry (disk, memory, thermal) is inside expected ranges
- Operating system healthy — the host boots without errors, system logs are clean of unresolved warnings, and the OS has joined the domain or authenticated to its identity source as expected
- Identity plane reachable — responders and service identities sign in through the trusted identity tier (see the Identity and Directory Recovery section), not through emergency shared credentials
- Data tier validated — application-level consistency checks pass and the acceptance rules in the Data Recovery Procedures section have been cleared for the datasets the system depends on
- Application smoke tests pass — representative end-to-end transactions execute successfully, not just a port-open or login-OK check
- Integration points healthy — dependent internal and external systems can reach and transact against the recovered host; monitoring and alerting for the host flow to the organization's SIEM or equivalent before traffic is admitted
- Owner acceptance recorded — the accountable system owner acknowledges the result against a pre-defined recovery smoke test, and the DR Coordinator captures the acceptance in the recovery record before user traffic is cut over

Failed recovery checks follow the three outcomes in Recovery Discipline.

Recovery & Restoration

Identity and Directory Recovery

Identity and directory services are the prerequisite to every downstream recovery activity — IT staff cannot administer systems, applications cannot authenticate to databases, and users cannot reach any recovered resource until authentication works. Identity recovery also carries a unique security dimension: an identity service that is technically restored but still contains attacker-planted accounts, modified trust relationships, or weakened policies becomes a persistent backdoor. This section treats "online" and "trusted" as separate states, and the application tier does not advance until the identity tier has passed the trust gate below.

Crossing the identity trust gate is defined in Recovery Discipline.

Preconditions — Before Beginning Identity Restore

The DR Coordinator clears the following preconditions before the identity restore begins.

A missing precondition pauses the identity restore (see Recovery Discipline).

- Network and edge tier green — DNS, routing, firewalls, and site-to-site connectivity from the tier-1 Recovery Procedures section are healthy; identity services cannot be restored to an unreachable target
- Break-glass access confirmed — the emergency-access path for each identity stack in scope is reachable from a survival-mode location (credentials stored offline, phishing-resistant MFA where applicable)
- Most recent identity backup verified present — backup location, backup software version, and decryption material are all accessible; a backup that has never been restored is assumed broken until a restore proves otherwise
- Compromise posture captured — the DR Coordinator records whether the triggering event is suspected to include an identity-compromise component, because that classification changes whether trust artifacts (tokens, tickets, service-account secrets) are rotated during restore
- Restore target isolated from primary where primary may carry compromised state — restore to a recovery domain, recovery tenant, or scratch instance rather than into a live primary that may re-infect the restored environment through replication or trust
- Concurrent administrative activity stopped — routine identity-admin actions (user provisioning, policy changes) are paused during restoration so the recovery record is a clean audit trail rather than interleaved with normal operations
- Identity-restore record opened — the target, source backup timestamp, executing responder, and start time are recorded before the first restore command runs

Bootstrap and Break-Glass Access

Identity restoration has a bootstrap problem: administering the identity service normally requires the identity service to be working. Every identity stack in scope for this plan has a documented bootstrap path that reaches the admin surface without depending on the service being recovered. The DR Coordinator confirms the bootstrap path is usable before any per-stack restore begins:

- For on-premises directory services — an administrative path into the first domain controller being restored that does not depend on the directory being online (for example, a local-console administrative mode on the restored host) with its credential documented and reachable from a location that survives the disaster
- For cloud-managed identity providers — at least two emergency-access accounts that do not depend on the organization's normal conditional-access or MFA policy for sign-in, with phishing-resistant authenticators, stored securely offline, and exercised at a documented cadence so the accounts are known to work when needed
- For hybrid stacks — the bootstrap path into each of the on-premises and cloud sides is independently usable, so a failure in one stack does not prevent access to the other
- Bootstrap credentials are treated as high-sensitivity artifacts — their existence, custody, and last verification date are tracked in the plan, and any use is logged and reviewed after the event

Identity Recovery Sequence

More than one identity stack is in scope for this plan. The stacks are restored in an order that does not create a circular bootstrap dependency — the first stack restored must be reachable through its own bootstrap path rather than through another stack that is itself still being recovered. Where the stacks are federated or synchronized, the side that other sides depend on for their source of truth is restored first, followed by re-establishment of the sync or federation relationship. Each stack completes its own per-stack procedure and passes the trust gate below before the next stack advances.

On-Premises Active Directory Recovery

On-premises directory recovery follows a forest-recovery pattern: one controller is restored first from known-good backup in an administrative mode that does not depend on the directory being online, the restored controller is brought to a known-consistent state before anything else replicates from it, and additional controllers are either restored separately or allowed to replicate outward from the known-good controller once it is healthy. The exact procedure, consistency commands, and shared-key rotation steps are documented in the operator's directory runbook and captured in the operator-authored steps below:

1. Identify the last known-good AD backup (system state backup, < 24 hours old)
2. Restore the PDC emulator first using Directory Services Restore Mode (DSRM)
3. Seize all FSMO roles to the restored DC
4. Verify AD replication health with dcdiag and repadmin
5. Reset KRBTGT password twice with a 12-hour interval
6. Restore remaining DCs from backup or allow replication from the restored PDC

Microsoft Entra ID (Azure AD) Recovery

Cloud-managed tenant recovery focuses on regaining administrative access to the tenant using the bootstrap path above, reviewing what changed during the outage window (sign-in activity, conditional-access policy modifications, privileged-role assignments), and re-establishing any hybrid-directory synchronization the organization depends on. Recovery is scoped to the tenant surface the organization owns; platform availability is the provider's responsibility and is not restored by this plan:

1. Access the cloud identity admin center using a break-glass account with phishing-resistant MFA
2. Review sign-in logs for suspicious activity during the outage period
3. Verify conditional access policies remain enforced and unchanged
4. Re-establish directory synchronization if on-premises identity was also affected
5. Verify federation trust and SSO configuration
6. Notify users to re-authenticate after sessions are revoked during recovery

Privileged and Service Account Recovery

Privileged and service-account recovery is sequenced after the identity stack is back online but before the trust gate below — a credential rotation landed into a healthy-but-untrusted directory has to be re-examined for post-compromise artifacts in any case. Service-account credentials are rotated only with the dependency map in hand; a rotation without the map tends to break application dependencies in non-obvious ways and turn an identity recovery into an application-tier incident:

1. Reset all domain admin and enterprise admin passwords
2. Rotate service account credentials per the dependency map
3. Verify the privileged access vault is accessible and credentials are current
4. Review admin group membership and remove any unauthorized additions
5. Audit recent privileged access for suspicious activity during the outage

MFA Device Loss and Recovery

MFA reset is among the highest-abuse social-engineering vectors because resetting an MFA method on a privileged account grants full access to whoever completes the reset. The procedure below applies to every MFA reset executed under this plan, with no exception for urgency claims, authority pressure, or executive status. The schema hint for this field is repeated here as the operating rule: strong identity verification is required before any MFA reset, even for executives:

1. Verify identity through video call with the user's direct manager
2. Remove all existing MFA methods from the account
3. Re-enroll with phishing-resistant method (FIDO2 security key)
4. Confirm enrollment and test sign-in from a known-clean device
5. Log the MFA reset in the security incident tracker

> **WARNING: NO SELF-SERVICE MFA RESET DURING RECOVERY**
>
> Self-service MFA reset paths are disabled during active recovery. A reset proceeds only after the identity of the requester is verified through an out-of-band channel that does not rely on the credentials under reset (for example, a video call with a managerial confirmer plus a government-issued identity document). The reset is logged in the recovery record and reviewed after the event regardless of whether it succeeded.

Acceptance — From "Online" to "Trusted for Downstream Use"

An identity service is "online" when authentication flows succeed. It is "trusted" only after the checks below pass. Downstream application-tier restoration does not advance against an identity service that is online-but-untrusted.

The identity tier stays restricted until acceptance passes (Recovery Discipline).

- Administrator and privileged account review — every admin and privileged-role assignment is accounted for; any that are not recognized or expected are disabled, not simply noted
- Trust and federation review — external trust relationships, federation configurations, and external identity-provider connections are verified against the documented baseline; unexplained entries are disabled pending investigation
- Conditional access and MFA policy verification — authentication and access policies are active and enforced at the levels documented in this plan, and modifications made during the outage are explained or reverted
- Authentication-log review — sign-in and admin-action logs for the outage window are reviewed for unexpected successful authentications, unusual sources, or privilege changes
- Service account audit — service accounts used by applications are functional, unmodified, and unchanged in membership and permission since the last known-good state; the dependency map is treated as authoritative
- Credential rotation if the event is suspected to include security compromise — passwords for privileged accounts, shared service-account secrets, and any identity-backed tokens whose material may be exposed are rotated before trust is granted
- Replication and synchronization health — for multi-controller directories and hybrid stacks, replication and sync are green and consistent before any stack in the set is treated as trusted

Failed identity checks follow the three outcomes in Recovery Discipline.

Data Recovery Procedures

Data recovery procedures must match the backup method exactly. A team that backs up using one tool and documents recovery against a different tool discovers the mismatch at the worst possible moment. Each procedure below names the actual backup location, the specific software or console path, and the validation that confirms the restore was successful. "Restore from backup" is not a procedure. Recovery-time estimates are drawn from measured test restores rather than assumed from backup size.

Preconditions — Before Beginning Restore

The DR Coordinator clears the following preconditions before any data restore begins.

A missing precondition pauses the restore (see Recovery Discipline).

- Dependency tiers are green — network and identity tiers are already restored per the Recovery Procedures section; restoring data to a target the team cannot authenticate against produces stalled downstream restarts
- Restore target confirmed — the restore lands in a clean, known target (alternate instance, recovery volume, scratch database) rather than overwriting a primary that may still contain partial data
- Backup verified restorable — the most recent backup is confirmed present, readable, and within the documented RPO window; a backup that has never been restored is assumed broken until a restore proves otherwise
- Credentials and keys accessible — backup-service credentials, encryption keys, and any decryption passphrases are reachable from a survival-mode location, not only from the primary environment
- Target capacity sized — free storage at the target, network bandwidth between source and target, and remaining time within RTO are all sized against measured restore throughput rather than assumed
- Concurrent writes stopped — application services that write to the restore target are stopped (or pointed elsewhere) before the restore begins; an in-place restore with live writers produces data divergence and usually a second outage
- Restore record opened — the target, source backup timestamp, executing responder, and start time are recorded in the recovery record before the first restore command runs

Restore Source and Path Selection

The restore paths documented in this plan are summarized below. The DR Coordinator selects the path based on the scope of the event, the freshness of the available backup relative to the RPO, and the restore time each path implies against the RTO:

Restore Source and Path Selection

RESTORE SOURCE	WHEN TO USE	OPERATOR NOTES
On-premises backup with offsite copies	Primary path when backups are held on local media with a copy at an offsite facility or custodian	Offsite retrieval time is part of the effective RTO; if offsite is unreachable, the local copy is used first and the offsite path is re-attempted once access is restored

On-Premises Backup Restore Procedure

On-premises restores operate against local backup media with an offsite copy as the resilient tier. The offsite-retrieval time is part of the effective RTO and is documented as a separate step so the responder knows when data becomes available, not just when the restore command runs:

1. Access the backup platform console
2. Select the most recent valid restore point for the target system
3. Choose the restore destination (original location or alternate host)
4. Initiate the restore and monitor progress
5. Run post-restore consistency checks
6. Rejoin the server to core services if performing a full virtual machine restore

Database Recovery Procedure

Database recovery is the most technically complex data operation. Databases carry transactional state that must be internally consistent — a partially restored database produces application errors, silent data corruption, or both. The restore is executed against a recovery instance, not the primary, and is not exposed to applications until the engine-specific consistency check documented by the operator for the database in scope has passed. Row counts for critical tables are compared against the last known-good baseline as a second, independent integrity signal:

1. Stop application services that depend on the database
2. Restore the PostgreSQL base backup to the data directory using pg_restore
3. Configure the point-in-time recovery target, then start the cluster so WAL segments replay automatically to the target
4. Run pg_isready and confirm the cluster has exited recovery and is accepting connections
5. Compare critical table row counts against the last known baseline
6. Restart application services and verify connectivity

The exact consistency command, transaction-log replay syntax, and point-in-time-recovery procedure for each database engine the organization runs are maintained in the operator's database runbook; this plan points to that runbook rather than restating engine-specific commands that drift over time.

SaaS Application Data Recovery

SaaS data recovery depends on the vendor's native retention and on any independent backup the organization maintains. Native vendor retention is typically time-boxed and frequently does not support granular restore of a single record or a precise point-in-time; independent backup closes that gap where it matters. The steps below describe the paths documented for this organization:

1. Restore collaboration-suite mailboxes and document repositories using the approved SaaS backup workflow
2. Restore CRM data from the scheduled export or backup recovery console
3. For other SaaS tools, contact the provider support desk and reference the documented retention commitment
4. Verify recovered data completeness by spot-checking recent records
5. Notify affected users when their data is restored

> **INFO: SAAS SHARED RESPONSIBILITY**
>
> SaaS vendors are responsible for platform availability; the organization is responsible for the availability and recoverability of its own data. Vendor service-level agreements typically limit remedies to service credits rather than data-recovery assistance. For any SaaS data whose loss would be material, maintain an independent backup through vendor-native export or a third-party SaaS-backup capability reviewed against recovery objectives.

Acceptance — From "Restored" to "Trusted for Production Use"

A completed restore is not the same as a system ready for production use. The checks below are cleared before any restored data is exposed to users or to dependent systems.

The restored dataset stays restricted until acceptance passes (Recovery Discipline).

- Restore completed without error — backup-software logs are clean; warnings and partial-success signals are treated as failures until investigated
- Data completeness — file counts, row counts, or storage volumes reconcile against a last-known-good baseline within a tolerance the responder is prepared to defend in the recovery record
- Data integrity — checksums, hashes, or application-level consistency checks pass against the restored dataset; engine-specific consistency checks complete without error
- Application functionality — the application starts, authenticates, and executes a pre-defined recovery smoke test covering representative transactions rather than only a login check
- Integration points healthy — dependent systems can authenticate to the restored data tier and read from it; downstream consumers (reporting, analytics, billing) are verified before opening user access
- RTO and RPO variance captured — measured time-to-restore is recorded against the documented RTO and measured data-loss point against the documented RPO; variances are captured in the recovery record rather than reconstructed from memory
- System owner acceptance — the accountable business owner acknowledges the result against the recovery smoke test, and the DR Coordinator records the acceptance before any user is re-enabled against the restored system

Failed restore checks follow the three outcomes in Recovery Discipline.

Testing & Maintenance

Testing and Maintenance

An untested disaster recovery plan is fiction. The failures organizations discover during real events — backup restores that take three times the estimated time, rotated credentials the plan still references, DR procedures that assume access to a system that is itself unavailable — cannot be discovered from a document review. They require hands-on exercises, and they require the output of those exercises to be tracked as material until closed.

This section defines the testing cadence, validation criteria for declaring recovery complete, after-action record required for each exercise, the conditions that trigger an out-of-cycle plan update, and the artifacts that are updated when a test or real event uncovers a gap. Annual testing is a floor, not a ceiling; organizations whose environment changes at a faster rate test at that rate.

Plan Governance

The parameters below establish who maintains the plan and at what cadence the plan is exercised and reviewed. The plan owner is accountable for each of these — if the plan drifts out of alignment with the environment, the owner is the single point of closure.

Plan Governance

PARAMETER	VALUE
Plan owner / maintainer	IT director role
DR testing frequency	Quarterly tabletop exercises, annual simulation (partial failover), biennial full failover test
Plan review frequency	Semi-annually, and within 5 business days of any material change
Date of last DR test	2026-02-15: Simulation (partial failover) of the secondary cloud region

Recovery Validation and Return to Service

Before a recovery is declared complete — whether following a real disaster or a full-failover exercise — the following checks are cleared by the plan owner and the accountable system owners. A system that fails any check is held in a restricted state until the gap is closed.

- Restoration sequence observed — dependencies were restored in the order documented in Recovery Procedures, and no downstream tier was returned to service before its prerequisites were healthy
- Data integrity verified — restored data passes application-level integrity checks or reconciles against a known-good baseline; anything that cannot be verified is quarantined rather than exposed to users
- Identity and authentication restored — directory services, SSO, MFA, and privileged-access paths are operational and re-authenticated before downstream systems accept traffic
- Monitoring and alerting re-established — logging, detection, and alerting for each restored system are confirmed flowing to the SIEM or equivalent before the system takes production traffic
- Recovery objectives measured — actual time-to-restore is recorded against the documented RTO and actual data loss against the documented RPO, and variances are captured in the after-action record rather than smoothed over
- System owner sign-off — the accountable business owner confirms the system behaves as expected against a pre-defined recovery smoke test, and the plan owner records the sign-off in the exercise or incident record

Testing Program

The organization conducts the exercise types below. Each exercise validates a different layer of the plan; a mature testing program uses more than one type over a rolling window so that the gaps each type cannot detect are caught

by another.

Testing Program

EXERCISE TYPE	WHAT IT VALIDATES	OPERATIONAL RISK
Tabletop exercise (discussion-based)	Team familiarity with the plan, clarity of roles and decision authority, and procedural gaps that only become visible when the plan is walked end-to-end	None — no systems are touched; the exercise cannot validate that technical procedures actually work
Walkthrough / checklist review	Accuracy of the plan document: contact information, credentials, system references, runbook steps, and cross-references to dependent artifacts	None — maintenance activity rather than simulation
Simulation (partial failover)	That backups can be restored, DR infrastructure is functional, and the recovery procedure can be completed within the documented RTO for the scoped systems	Moderate — a scoped subset of production is exercised; schedule during low-activity windows and pre-stage rollback
Full failover test	End-to-end recovery capability, including cross-system integration, identity and data-tier interaction, and the organization's ability to operate from DR for a bounded window	High — requires executive approval, pre-staged rollback, and explicit return-to-primary criteria; carries the highest potential for production disruption

Post-Test Review and After-Action Record

Every DR exercise — regardless of type — produces a written after-action record. The record is the auditable evidence that the exercise occurred and is the input to the update loop described below. At minimum, the record contains:

- Date, scenario, scope, and duration — what was exercised, for how long, and within what boundaries
- Participants and roles — who attended in what capacity, including observers and absent role holders
- Step-by-step outcomes — each documented step marked pass, fail, or partial, with the actual time required versus the documented target
- RTO / RPO variance — for simulation and failover exercises, the measured time-to-restore and data-loss values against the documented objectives
- Issues and root cause — findings encountered during the exercise, with a short root-cause analysis for each rather than a symptom list
- Corrective actions with named owner and due date — each finding becomes a concrete owned action, not a standing aspiration
- Sign-off — the plan owner records acceptance of the after-action record and the open-finding list

Corrective actions identified in the record are tracked to closure. Unresolved findings from DR tests are known risks — they are reported to management and reflected in the organization's risk register until resolved, rather than allowed to sit in an exercise report that is never reopened.

Plan Update Triggers

In addition to scheduled reviews, the following events trigger an immediate out-of-cycle review of this plan. The plan owner initiates the review within five business days of the triggering event, and tracks the resulting changes against the artifact list below.

- Major infrastructure change (new data center, cloud migration, vendor switch)
- Organizational restructuring affecting DR team membership
- After any actual disaster recovery activation

- After any failed DR test revealing procedural gaps
- Regulatory or compliance requirement changes

Each update is recorded in the Revision History appendix with a description of what changed and why. DR team members are notified of material changes and acknowledge receipt of the updated plan so that the plan responders are working from matches the plan the owner has maintained.

Artifact Update Integration

Findings from exercises, real events, and triggered reviews are only useful if they land in the artifacts responders will reach for next time. The plan owner is accountable for integrating each finding into the following artifacts before the next plan revision cycle closes:

- This Disaster Recovery Plan — updated roles, decision authority, validation criteria, or escalation paths
- Critical Systems Inventory — dependency corrections, RTO/RPO revisions, and added or retired systems
- Backup Strategy — cadence, retention, storage location, or encryption posture corrections uncovered by restore failures
- Recovery Procedures — per-platform runbook steps, credentials, and prerequisite checks
- Identity and Directory Recovery — IdP bootstrap steps, break-glass account handling, and re-authentication dependencies
- Data Recovery Procedures — per-platform restore order, validation queries, and application-level integrity checks
- Communication Plan — internal channels, customer-facing templates, and the notification matrix when external obligations shift
- Contact Directory — names, roles, reachability paths, and retained-provider details

Changes made in response to findings are cross-referenced in the plan's revision history so that auditors, insurers, and future responders can trace which artifacts evolved from which tests or incidents. The update loop is considered closed only when every artifact affected by a finding has been updated and the owning role has acknowledged the change.

Appendix: Revision History

VERSION	DATE	DESCRIPTION
1.0	2026-06-12	Initial document generation

Appendix: Glossary of Terms

TERM	DEFINITION
Business Impact Analysis (BIA)	A systematic process to determine and evaluate the potential effects of an interruption to critical business operations as a result of a disaster, accident, or emergency.
CIS Controls	A prioritized set of cybersecurity best practices published by the Center for Internet Security, organized into 18 control families with specific safeguards.
Cloud Computing	The delivery of computing services — including servers, storage, databases, networking, software, and analytics — over the internet, enabling flexible capacity and reducing dependency on on-premises infrastructure.
Communication Plan	A structured framework defining how an organization will communicate with employees, customers, regulators, media, and other stakeholders during and after a business disruption.
Encryption	The process of encoding data so that only authorized parties holding the key can read it. Used to protect data at rest (for example full-disk encryption) and in transit (for example TLS) against unauthorized access.
Implementation Group (IG)	One of three self-assessed categories (IG1, IG2, IG3) in the CIS Controls framework that help organizations prioritize safeguard implementation based on their risk profile and resources.
Multi-Factor Authentication (MFA)	Authentication method requiring two or more verification factors (something you know, have, or are) to gain access to a resource.
Recovery Point Objective (RPO)	The maximum acceptable amount of data loss measured in time, determining the minimum frequency of backups.
Recovery Time Objective (RTO)	The maximum acceptable length of time that a system or business process can be down after an incident before the impact becomes unacceptable.
Safeguard	An individual security action or control within a CIS Controls family, identified by a numeric ID (e.g., 17.1) and mapped to an Implementation Group.
Tabletop Exercise	A discussion-based exercise where key personnel walk through a simulated disruption scenario to validate the business continuity plan, identify gaps, and improve team readiness without activating actual recovery procedures.

Appendix: Identified Gaps and Improvement Opportunities

The questionnaire responses are complete: no missing-prose or empty required sections were found. Contact and identifier fields are intentionally left for local completion and are marked where they appear. Control implementation gaps and partial safeguards, if any, are listed in the maturity summary and remediation roadmap sections. The following governance and operational items should be completed before this plan is considered approved for operational use:

- Obtain approval signatures from all designated signatories before the document is approved for operational use.
- Complete any local-fill placeholders in the exported copy only — do not re-enter named contacts or company identifiers into the hosted workspace.
- Complete the local-only coordinator, escalation-contact, and emergency-phone-path placeholders in the exported copy.
- Conduct an initial tabletop exercise to validate team readiness and identify procedural gaps not visible in written form.
- Confirm incident and breach notification timelines and recipients with legal counsel for all applicable jurisdictions.
- Confirm legal and regulatory applicability with counsel for any obligations referenced in this document.
- Schedule the first annual review date and assign the review owner.

These items represent standard governance follow-up for a newly generated plan and do not indicate deficiencies in the plan content itself.

Appendix: CIS Controls v8.1 Regulatory Control Mapping

CONTROL	SECTION	DESCRIPTION
CIS Control 11.1	executive_summary	Establish and maintain a data recovery process : document and maintain a current data recovery process covering scope, prioritization, and testing.
CIS Control 11.2	backup_strategy	Perform automated backups : perform automated backups of in-scope enterprise assets on a regular schedule.
CIS Control 11.3	backup_strategy	Protect recovery data : protect recovery data with equivalent controls as the original data, including encryption and access control.
CIS Control 11.4	backup_strategy	Establish and maintain an isolated instance of recovery data : maintain an isolated instance (air-gapped or offline) of recovery data.
CIS Control 11.5	testing	Test data recovery : test backup recovery quarterly or more frequently to verify the recovery process.
CIS Control 17.1	communication	Designate personnel to manage incident handling : assign and train personnel responsible for incident response.
CIS Control 17.6	communication	Define mechanisms for communicating during incident response : establish out-of-band communication channels.
CIS Control 17.7	testing	Conduct routine incident response exercises : perform tabletop or functional exercises at least annually.
CIS Control 17.8	testing	Conduct post-incident reviews : perform a lessons-learned review after every major incident to improve future response.
CIS Control 17.9	executive_summary	Establish and maintain security incident thresholds : define severity thresholds that trigger the disaster recovery plan.

Appendix: Internal Review and Local Completion Checklist

This document was generated from questionnaire responses and requires internal validation before it is approved for operational use. The following checklist should be completed by the designated review team.

- Complete any local-fill placeholders in the exported copy only: do not re-enter named contacts or company identifiers into the hosted workspace.
- Confirm legal and regulatory applicability: engage legal counsel to verify notification timelines, jurisdictional obligations, and contractual requirements referenced in this document.
- Verify tool names, commands, and procedures: confirm that the specific platform commands, console/portal steps, and operational procedures in this document (for example, your identity provider's session-revocation step, firewall rule changes, or backup-restore commands) match your current production environment.
- Conduct at least one tabletop exercise: walk through a simulated incident scenario using this plan to identify procedural gaps not visible in written form.
- Review and validate all severity thresholds and escalation timelines: confirm these reflect current organizational risk tolerance and operational capacity.
- Confirm backup and recovery targets: validate that stated RTO, RPO, and backup frequencies are achievable with current infrastructure.
- Obtain approval signatures from all designated signatories before the document is approved for operational use.
- Schedule the first annual review date and assign the review owner.

Local completion items for the final distribution copy:

- Replace role placeholders with the approved local titles or personnel for the final distribution copy.
- Add direct contact details such as shared mailboxes, emergency phone paths, vendor support numbers, and named coordinators only in the local copy.
- Insert company-specific legal names, locations, domains, internal system identifiers, and account references only after export.

Once all checklist items are complete, update the document status from "Internal Review Draft" to "Approved" and re-export.