



TEMPLATE

Amazon EC2 Instance Recovery and Failover Template

A practical PDF template for deciding when to recover an EC2 instance versus fail over, with validation checkpoints, decision criteria, and production readiness checks.

Amazon EC2 Instance Recovery and Failover Template

Use this template to define, document, and validate your EC2 recovery and failover approach before an incident occurs.

1) Purpose

Workload / application name: _____

Environment: ? Dev ? Test ? Staging ? Production

Owner/team: _____

Date created: _____ Last reviewed: _____

Objective:

- ? Restore the same workload on the same logical service path
- ? Preserve availability by failing over to another instance or environment
- ? Support both recovery and failover depending on failure type

Business impact if unavailable:



2) Workload profile

Service characteristics

- State model: ? Stateless ? Stateful ? Mixed
- Single instance or clustered: ? Single ? Clustered ? Autoscaled
- Traffic entry point: ? Direct ? Load balancer ? DNS ? API gateway ? Other: _____
- Recovery tolerance: ? Minutes ? Hours ? Same day ? Other: _____
- Can it tolerate brief downtime? ? Yes ? No
- Can the instance be trusted after failure? ? Yes ? No ? Unknown

Critical dependencies

Mark all that apply and document the source of truth.

- ? Attached EBS volume(s)
- ? Database / datastore
- ? Object storage
- ? Session store
- ? Internal DNS
- ? External DNS
- ? Load balancer target group
- ? Security group / firewall rules
- ? IAM role / instance profile
- ? Secrets manager / vault



- ? Certificates / keys
- ? Third-party API
- ? License server
- ? Monitoring / alerting

Notes on dependency constraints:

3) Failure classification guide

Use this guide to decide whether recovery or failover is the safer path.

Failure type	Likely path	Notes
OS crash, hung process, minor app fault	Recovery	Restart or replace from a known-good template may be sufficient
Full filesystem, local config error, bad deployment	Recovery	Validate rollback and config state before re-entry
Storage corruption or inconsistent data	Case-by-case	Verify snapshots, replicas, and integrity checks before restoring
Host impairment or repeated instance instability	Failover	Prefer a clean target over reusing an unreliable host path
Network reachability issue	Case-by-case	Determine whether the issue is instance-local or broader
Security incident or trust issue	Failover	Treat the original instance as untrusted until proven otherwise
Unknown root cause	Failover or manual review	Preserve evidence and use a controlled cutover

Default rule for this workload:

4) Recovery decision criteria



Recovery is allowed only if all conditions are true

- ? Failure is understood and isolated
- ? Instance and data are trusted to return to service
- ? Required state can be restored consistently
- ? Dependencies are available and verified
- ? Validation steps are documented and pass
- ? Rollback path is available if validation fails

Failover is required if any condition is true

- ? The instance may be compromised or untrusted
- ? Data integrity is uncertain
- ? Recovery would take longer than the defined RTO
- ? The workload cannot safely resume on the same instance path
- ? A standby or secondary environment is available and tested

Decision owner: _____

Approval required from: _____

5) Restoration path selection

Choose the preferred path for each failure mode.

Scenario	Primary path	Secondary path	Notes
Local application failure	?	Restart ? Replace ? Fail over	_____ _____



OS / OS-level instability | ? Restart ? Replace ? Fail over | _____ | _____

Host impairment | ? Replace ? Fail over | _____ | _____

Data corruption | ? Restore from snapshot ? Fail over | _____ | _____

Security incident | ? Fail over ? Isolate ? Rebuild | _____ | _____

Planned maintenance | ? Replace ? Fail over ? Drain and shift | _____ | _____

Launch template / image reference: _____

Golden configuration source: _____

Infrastructure as code location: _____

6) Pre-return validation checklist

Complete these checks before allowing traffic to return.

System health

- ? Instance is reachable by the expected management path
- ? Operating system boots cleanly
- ? Required services are running
- ? CPU, memory, and disk usage are within normal bounds
- ? No unexpected errors in system logs

Application health

- ? Application process is active
- ? Expected ports are listening



- ? Application health check returns success
- ? End-to-end test transaction succeeds
- ? Background jobs or schedulers are in the correct state

Data and storage

- ? Attached volumes are present
- ? Mount points are correct
- ? Permissions are correct
- ? Data matches the expected recovery point
- ? Database connectivity is confirmed
- ? Integrity or checksum checks passed, if applicable

Network and routing

- ? Security groups / firewall rules are correct
- ? Route tables / network paths are correct
- ? DNS target is updated if needed
- ? Load balancer registration is correct
- ? Health checks are passing from the load balancer or monitoring layer

Identity and secrets

- ? Instance profile or service identity is attached correctly
- ? Secrets are available and readable by the workload
- ? Certificates and keys are valid
- ? External auth providers are reachable



Operational checks

- ? Monitoring alerts are set to normal thresholds
- ? Logging is functioning
- ? Backup schedule is active
- ? Rollback plan is ready
- ? Change owner has approved return to service

Validation performed by: _____

Validation time: _____

Result: ? Pass ? Fail ? Conditional pass

Notes:

7) Cutover and rollback plan

Cutover trigger

Describe when traffic moves:

Rollback trigger



Describe when the recovery/failover attempt is reversed:

Communication plan

- Incident channel: _____
- Stakeholders to notify: _____
- Status update frequency: _____
- Customer-facing message owner: _____

8) Recovery runbook fields

Fill in the operational steps required for your environment.

- Detect and classify the issue
- Trigger:
- Symptoms:
- Suspected failure type:
- Stabilize the environment
- Isolate instance if needed:
- Preserve logs / evidence:
- Freeze changes:
- Select the path
- Recovery method:
- Failover destination:
- Manual approvals needed:
- Execute restoration



- Step 1:
- Step 2:
- Step 3:
- Validate
- Health checks:
- Data checks:
- Traffic checks:
- Return to service
- Cutover action:
- Monitoring window:
- Backout plan:

9) Testing and rehearsal record

Test schedule

- Last test date: _____
- Next test date: _____
- Test type: ? Recovery ? Failover ? Both
- Test environment: _____

Test results

Item	Expected	Actual	Pass/Fail	Notes
Instance recovery				



Failover cutover | | | |

DNS / routing update | | | |

Data consistency | | | |

Application health | | | |

Rollback | | | |

Lessons learned

10) Production readiness checklist

Before declaring this plan production-ready, confirm the following:

- ? Failure modes are documented
- ? Recovery and failover triggers are explicit
- ? Dependencies are mapped and owned
- ? Validation checks are executable and repeatable
- ? RTO and RPO are defined
- ? Backup and restore procedures are tested
- ? At least one recovery drill has been completed
- ? At least one failover drill has been completed
- ? Rollback criteria are clear
- ? Security and trust assumptions are documented
- ? Monitoring and alerting support the plan



- ? Change approval path is known

Production readiness status: ? Ready ? Not ready ? Ready with exceptions

Exceptions:

11) Quick decision summary

Use this compact decision rule during an incident review:

- Recover when the failure is local, understood, and the instance can be trusted again.
- Fail over when service continuity matters more than preserving the original instance identity.
- Validate first: compute, storage, network, identity, and application state must all be confirmed.
- Rehearse both paths before production depends on them.

12) Sign-off

Prepared by: _____

Reviewed by: _____

Approved by: _____

Approval date: _____

Review cycle: _____