AWS Well-Architected Review — Sample Report

About this sample. This report was produced against a demonstration AWS environment built to show the structure, scoring, and depth of an Uptempo Cloud Well-Architected Review. It is not client work, and the "Northbound" company described below is fictional. Findings, evidence, and remediation guidance are representative of what a real engagement produces. Client reports additionally include an evidence appendix and a SOC 2 / security-questionnaire mapping worksheet.

Review window: four weeks · Environments reviewed: production and staging (2 AWS accounts) · Method: AWS Well-Architected Framework, all six pillars · Access: read-only · Prepared by Christopher Gutierrez, Uptempo Cloud LLC

1. Executive summary

Northbound is a 140-person B2B SaaS company running a multi-tenant platform on AWS: roughly $86,000/month across two accounts, EC2- and RDS-centred, with Terraform covering about half the estate. The platform is stable day to day, but its architecture has evolved faster than its safeguards: production and staging share an account boundary, several identity and backup controls assume people never make mistakes, and cost has grown 9% quarter over quarter without an owner.

The environment is sound at its core and none of the findings below require a rebuild. The risks concentrate in three places:

- **Blast radius.** Production and staging live in the same AWS account with broad IAM roles spanning both (findings SEC-01, OPS-03). One mistaken credential or script reaches everything.
- **Recovery assumptions.** RDS backups exist but restores have never been tested, and the recovery runbook references infrastructure that no longer exists (REL-01, REL-02). Recovery time is unknown, which for an enterprise security review is the same as "no."
- **Untended spend.** About $11,300/month (13%) is attributable to idle, oversized, or legacy-generation resources with no workload justification (COST-01 through COST-03).

Recommendation for the next quarter: execute the seven "Now" items in the roadmap (section 4). They are two to three engineer-weeks of work in total, none require downtime, and they close every finding scored 15 or higher — including everything an enterprise customer's security reviewer would flag in the first pass.

Top five findings by risk:

ID	Finding	Pillar	Risk score
SEC-01	Production and staging share one account; IAM roles span both	Security	20
REL-01	Database restores never tested; RTO/RPO undefined	Reliability	20
SEC-02	Nine IAM users with long-lived access keys, three unused 90+ days	Security	16

OPS-01	No organization-wide CloudTrail; API history has gaps	Operational Excellence	16
COST-01	$6,100/month in idle and oversized EC2/RDS capacity	Cost Optimization	15

2. Scope and method

- Accounts: northbound-prod (production + staging), northbound-shared (CI/CD, artifacts). Workloads: the multi-tenant API platform, one internal analytics stack.
- Data sources: read-only role with ReadOnlyAccess + AWSBillingReadOnlyAccess; CloudTrail, AWS Config, Cost & Usage Report, Cost Explorer, Trusted Advisor, VPC flow logs (sampled); three architecture interviews (~50 minutes each) with the platform lead and two senior engineers.
- Framework: AWS Well-Architected Framework, all six pillars, using the AWS Well-Architected Tool. Findings are scored Impact (1-5) × Likelihood (1-5) → Risk; Effort is S (≤1 engineer-day), M (≤1 engineer-week), L (multi-week).
- The client team keeps every artifact: this report, the scored findings register (CSV), and target-state diagrams for each high-risk finding.

3. Findings by pillar

Operational Excellence

OPS-01 · No organization-wide CloudTrail · Risk 16 (I4×L4) · Effort S

- Observation: CloudTrail is enabled per-account with 90-day event history only; no org trail, no S3 archive, no log-file validation. The shared account's trail was disabled during a cost cleanup in March (visible in the trail's own last events).
- Risk: security investigations and enterprise-customer audits depend on API history that currently has gaps and can be silently switched off. During an incident, the first question — "what changed?" — cannot be answered past 90 days.
- Recommendation: one organization trail to a dedicated, locked log bucket (S3 Object Lock, governance mode) with log-file validation on; Terraform module supplied in the findings register.

OPS-02 · Deploys depend on one person's laptop · Risk 12 (I4×L3) · Effort M

- Observation: CI runs tests, but the production apply step is run manually by the platform lead from a local Terraform workspace with admin credentials. State is in S3, but plan review is informal.
- Risk: single-person dependency for every release and an unauditable change path to production — the same gap the "key-person risk" question in vendor reviews is probing for.
- Recommendation: move applies into the pipeline with a plan-approval step; the deploy role assumes a scoped IAM role, humans lose standing write access. Pairs with SEC-02.

OPS-03 · Staging changes hit production dependencies · Risk 12 (I4×L3) · Effort M

- Observation: staging shares the production account, VPC peering, and two IAM roles whose policies name resources in both environments. A load test in staging in May throttled the production DynamoDB table it unknowingly shared.
- Risk: test activity with production blast radius.
- Recommendation: account separation (see SEC-01 target architecture); interim: split the shared roles and remove cross-environment resource ARNs — one day of IAM work.

Security

SEC-01 · Single account for prod + staging · Risk 20 (I5×L4) · Effort L

- Observation: one AWS account contains production, staging, and three developer sandboxes. Twelve IAM roles have resource access spanning environments. There is no AWS Organizations structure, no SCPs.
- Risk: maximum blast radius — one compromised credential, one wrong script, or one console mistake reaches every environment and every tenant's data. This is the finding an enterprise security review fails first.
- Recommendation: AWS Organizations with prod, non-prod, and security accounts; SCP guardrails (deny root use, deny region drift, protect the log bucket); migrate workloads by wave. Target architecture diagram included; typical execution is a 4-6 week fixed-scope project with zero-downtime cutover per workload.

SEC-02 · Long-lived IAM user keys · Risk 16 (I4×L4) · Effort S

- Observation: nine IAM users hold long-lived access keys; three unused for 90+ days; two keys are 3+ years old; one belongs to a contractor whose engagement ended last year. MFA is enforced for console but not API access.
- Risk: static credentials are the most common initial access vector in cloud incidents; unused keys are pure liability.
- Recommendation: delete the three dormant users today; migrate CI and human access to short-lived credentials (IAM Identity Center for humans, OIDC federation for CI); key age alarm via Config rule.

SEC-03 · Security groups open to 0.0.0.0/0 on management ports · Risk 12 (I4×L3) · Effort S

- Observation: two security groups allow SSH (22) from anywhere; one allows PostgreSQL (5432) from anywhere, currently attached to a staging replica. VPC flow logs confirm external scan traffic.
- Risk: internet-exposed management and database ports are the standing invitation every scanner looks for.
- Recommendation: close both rules; SSM Session Manager already in place makes SSH exposure unnecessary; add a Config rule to prevent recurrence.

SEC-04 · Unencrypted EBS volumes and snapshots · Risk 9 (I3×L3) · Effort S

- Observation: 14 of 41 EBS volumes (all pre-2024) and their snapshot lineage are unencrypted; account-level default encryption is off.
- Risk: fails the encryption-at-rest line item on every security questionnaire; snapshot sharing mistakes become data exposure.
- Recommendation: enable account default encryption now (one call); re-encrypt legacy volumes via snapshot-copy during the next maintenance windows.

Reliability

REL-01 · Restores never tested; RTO/RPO undefined · Risk 20 (I5×L4) · Effort M

- Observation: RDS automated backups are on (7-day retention) and snapshots replicate nowhere; no restore has ever been performed; no RTO/RPO targets exist in writing.
- Risk: backup success is measured at restore time. An untested backup is a hypothesis, and the current hypothesis includes an unknown multi-hour restore window during which every tenant is down.
- Recommendation: quarterly restore drill into an isolated VPC with a timed runbook (first drill scheduled as a "Now" item); define RTO/RPO per tier and let those targets drive retention and cross-region copy decisions.

REL-02 · Recovery runbook references deleted infrastructure · Risk 12 (I4×L3) · Effort S

- Observation: the DR document (last edited 14 months ago) names a bastion host, a Jenkins server, and subnet IDs that no longer exist.
- Risk: during a real incident the runbook actively misleads; the on-call engineer burns the first hour discovering that.
- Recommendation: regenerate the runbook from current state as part of the restore drill; store it next to the code it recovers, reviewed on every architecture change (checklist supplied).

REL-03 · Single-AZ production database replica gap · Risk 12 (I4×L3) · Effort S

- Observation: the primary RDS instance is Multi-AZ, but the read replica the API's read path depends on is single-AZ in the same AZ as two of three app instances.
- Risk: one AZ event removes the read path and most of the app tier simultaneously — the Multi-AZ primary survives but the platform still degrades.
- Recommendation: move the replica to a different AZ (minutes of work); rebalance the ASG's subnet spread.

Performance Efficiency

PERF-01 · Previous-generation instance families · Risk 9 (I3×L3) · Effort S

- Observation: 60% of compute runs on m4/r4/t2 families; the workload is x86 with no architecture constraint pinning it there.
- Risk: paying more for less: current-generation (and Graviton where compatible) families deliver better price-performance; t2 credit exhaustion has already caused two latency incidents (visible in CloudWatch).
- Recommendation: rolling family upgrade via launch-template change; Graviton evaluation for the stateless API tier (candidate for a follow-on sprint with load-test validation).

PERF-02 · No caching in front of the hottest read path · Risk 8 (I2×L4) · Effort M

- Observation: the tenant-config lookup (34% of all database reads, per Performance Insights) hits RDS directly on every API request; values change rarely.
- Risk: database load scales linearly with traffic for data that is effectively static; this is the first thing to fall over at 3-5x growth.
- Recommendation: ElastiCache (or DAX-style in-process cache) with short TTL on tenant config; expected 25-35% reduction in database load based on the observed read mix.

Cost Optimization

COST-01 · Idle and oversized capacity: ~$6,100/month · Risk 15 (I5×L3) · Effort S

- Observation: two stopped-but-provisioned RDS instances from a 2024 migration test ($1,900/mo), a staging EC2 fleet sized identically to production but averaging 4% CPU ($2,700/mo), and 23 unattached EBS volumes plus 400+ orphaned snapshots ($1,500/mo).
- Risk: pure waste — 7% of the monthly bill buying nothing.
- Recommendation: delete the test instances (snapshots retained 30 days first), right-size staging to 25% of production with schedules (off nights/weekends), lifecycle the orphaned storage. All read-only verified; the register lists each resource ID.

COST-02 · Zero commitment coverage · Risk 12 (I4×L3) · Effort S

- Observation: 100% of compute and RDS is on-demand; the baseline (steady-state ~$38,000/mo of it) has been stable for 11 months.
- Risk: paying the walk-in rate for a workload with a proven floor — roughly $9,000-$11,000/year left on the table at current usage (Compute Savings Plan + RDS Reserved, 1-year no-upfront, sized to 80% of the observed floor).
- Recommendation: commit conservatively to the floor after COST-01's right-sizing lands (order matters — commit after shrinking).

COST-03 · gp2 volumes and no S3 lifecycle · Risk 9 (I3×L3) · Effort S

- Observation: all 41 EBS volumes are gp2 (gp3 is ~20% cheaper at equal or better performance for this profile); application logs land in S3 Standard and stay there forever (14 TB and growing).
- Risk: slow leak that compounds with growth.
- Recommendation: gp2→gp3 migration (online, scriptable), S3 lifecycle to Infrequent Access at 30 days / Glacier at 180 for the log prefixes; register includes the exact bucket/prefix list.

Sustainability

SUS-01 · Idle capacity is also carbon · Risk 6 (I2×L3) · Effort S

- Observation: the same idle fleet from COST-01 plus previous-generation families from PERF-01.
- Recommendation: no separate work needed — the COST-01 and PERF-01 remediations are the sustainability wins; note them in the ESG section of enterprise questionnaires (customers increasingly ask).

4. Ranked remediation roadmap

Phase	Finding	Action	Effort	Risk closed
Now (weeks 1-3)	SEC-02	Delete dormant users; short-lived credentials for CI + humans	S	16
Now	OPS-01	Organization trail + locked log bucket	S	16
Now	SEC-03	Close 0.0.0.0/0 management/database rules	S	12

Now	REL-03	Move replica AZ; rebalance ASG	S	12
Now	COST-01	Delete idle capacity; schedule staging	S	15
Now	REL-01a	First timed restore drill + RTO/RPO definition	M	20 (partial)
Now	SEC-04	Default EBS encryption on	S	9 (partial)
Next (weeks 4-8)	SEC-01 / OPS-03	Organizations + account separation, wave by wave	L	20 + 12
Next	OPS-02	Pipeline-gated applies; remove standing admin	M	12
Next	COST-02/03	Commitments after right-sizing; gp3; S3 lifecycle	S	12 + 9
Later (quarter 2)	PERF-01/02	Family upgrades; cache the hot read path; Graviton eval	M	9 + 8
Later	REL-01b/02	Quarterly drill cadence; living runbook process	M	remainder

Dependencies: COST-02 commitments follow COST-01 right-sizing; SEC-01's account separation supersedes the interim OPS-03 role split if executed within the quarter.

5. 90-day follow-up

Included in the engagement: a 60-minute re-check at 90 days against this register — what closed, what moved, what new risk appeared. The register is delivered as a CSV your team can track in its own backlog tooling; nothing about the follow-up requires Uptempo access to be retained in the interim.

Uptempo Cloud LLC · hello@uptempocloud.com · uptempocloud.com · Sample report — produced against a demonstration AWS environment; not client work.