Mysterious Code.

AWS SELECT TIER PARTNER
UK & EU · SINCE 2014

SAMPLE REPORT · AWS WELL-ARCHITECTED REVIEW

Production-Readiness Snapshot

A focused, read-only review of your AWS account against the Well-Architected Framework, sized to your team and budget.

the subject "Initech" is fictional. Findings are composited and anonymised from real UK and EU AWS engagements, no real client is shown. Percentages illustrate the pattern, they are not a quoted result.

WHAT A SNAPSHOT IS

A **Production-Readiness Snapshot** is a focused, read-only review of your AWS account against the AWS Well-Architected Framework, run by the architect who would do the work. Over 3 to 5 working days we go through your account and your CloudFormation templates, then deliver this report: the findings that affect a production launch, ranked, each with a fix, a rough effort and its cost impact.

It is the Well-Architected review, **sized to your team and budget**. We tell you what to fix now, what can wait, and where a simpler option is good enough for a company your size, rather than hand you a long list of enterprise practices you do not need yet. That judgement is the point of the review.

It is not a compliance audit or a certification, and we leave nothing running in your account afterwards. The review uses read-only access and AWS's own tools (Trusted Advisor, Service Quotas, IAM Access Analyzer, Cost Explorer). **"Production-ready"** means the High Risk Issues are cleared: the things that cause an outage, a data leak or a runaway bill at launch or under load.

EXECUTIVE SUMMARY

Verdict: not yet production-ready. Initech has a solid base, the application runs on ECS Fargate and the core infrastructure is defined in CloudFormation, but three High Risk Issues need clearing before the v2 scales onto it.

The most serious will not show up on a dashboard: a **production Aurora snapshot is restored into staging without anonymisation**, so live customer data sits in a lower-trust environment that most of the team can reach. On top of that, the **Aurora cluster is a single writer in one Availability Zone with an untested restore**, and **long-lived IAM access keys with broad permissions** are still in use.

Fix the staging data first. It affects the most and costs the least to put right. Beyond the High Risk Issues, a few medium-risk ones only show up under launch traffic, and the easiest to miss is **service-quota headroom** (Finding 6).

WELL-ARCHITECTED RISK SCORECARD

PILLAR	HIGH RISK	MEDIUM RISK
Operational Excellence	0	1
Security	2	1
Reliability	1	2
Performance Efficiency	0	1
Cost Optimization	0	1
Sustainability	*not in scope*	
Total	**3**	**6**

The six findings below are the ones we would action first. The remaining Medium Risk Issues are listed in the full report with the same detail.

Effort S = up to ~2 days · M = a few days to ~2 weeks · L = multi-week
Cost impact Saves (reduces spend) · Neutral (no material change) · Invest (a justified increase)

PRIORITISED FINDINGS

1 Production data restored into staging without anonymisation

HIGH RISK

Security · Effort M (masking) · Cost Neutral

FOUND Staging is refreshed from a **production Aurora snapshot with no masking step**, and the **staging database credentials are shared across the team**. Live customer data therefore sits in your lowest-trust environment.

RISK No monitoring will catch this, so you find out when it becomes an incident, not before. A leaked credential, a compromised dependency or a simple mistake in staging then exposes real customer data, with the GDPR problem that brings.

FIX The quick fix is to add a masking step to the staging refresh, so PII columns are masked or replaced with fake data as the snapshot is restored. The bigger fix is splitting production and non-production into separate accounts under AWS Organizations. That is a project rather than a launch blocker, so we would do it after launch instead of holding things up. Neither needs paid tools.

2 Aurora cluster is single-writer, single-AZ, with an untested restore

HIGH RISK

Reliability · Effort S to M · Cost Invest

FOUND The Aurora cluster runs a **single writer instance in one Availability Zone**. Automated backups and point-in-time recovery are enabled, but **no restore has been exercised**, so recovery is unproven and RPO/RTO are undefined.

RISK The database is a single point of failure for the whole platform. An AZ impairment or instance failure means full downtime, with recovery depending on a restore path nobody has run.

FIX For a launch with real customer data, the minimum we would accept is automatic failover to a second Availability Zone. That means an **Aurora read replica in a second AZ**, plus a restore you have actually run, so you know it works and how long it takes, and an agreed RPO and RTO. A replica is a second instance of the same size, so it roughly **doubles the cluster's compute cost**. If the business can live with a longer recovery time, staying single-AZ with a tested restore is a cheaper option that is good enough. Most launches cannot, so this is where we would spend and save elsewhere.

3 Long-lived IAM access keys with broad permissions

HIGH RISK

Security · Effort M · Cost Neutral

FOUND Several IAM users authenticate with **static access keys over 18 months old**, and at least two carry AdministratorAccess or equivalent.

RISK A single leaked key is a full account compromise, and broad policies remove any containment. Static keys are rarely rotated, so they tend to persist indefinitely.

FIX Pulling IAM users out entirely is a lot of upheaval for a small team, so the practical move is to **limit what a leaked key can reach** instead. Using the account split from Finding 1, give engineers a **lower-permission role they assume to work in production, with MFA required**, so their everyday keys no longer hold admin. Then rotate and review the long-lived keys regularly. IAM Access Analyzer is worth running alongside those reviews, but only if someone acts on what it finds. The stronger end state, once the team is ready for it, is **IAM Identity Center** (it is free): short-lived SSO sessions replace the long-lived keys completely, for a bit more setup up front.

4 Non-production runs 24x7 at production sizing, with no cost guardrails

MEDIUM RISK

Cost Optimization · Effort S · Cost **Saves**

FOUND Three non-production environments run **continuously on Fargate at production task sizes**, and the account has **no AWS Budgets or Cost Anomaly Detection** configured.

RISK You pay for capacity that sits idle outside working hours, and a cost spike first shows up on the invoice instead of as an alert.

FIX

- *Non-production:* schedule these environments to scale down overnight and at weekends, which takes roughly **30%** off their cost, and move suitable tasks to Fargate Spot.
- *Whole account:* turn on **AWS Budgets and Cost Anomaly Detection for the whole account**, not just non-production, so a cost spike shows up the same day rather than on next month's invoice.

5 Deployments are manual and non-reproducible

MEDIUM RISK

Operational Excellence · Effort M to L · Cost Invest

FOUND Services are deployed **from engineer workstations** using local credentials. There is no pipeline, no environment promotion and no automated rollback.

RISK Deploys are not reproducible, configuration drifts from what is in version control, and recovering from a bad release is a manual scramble during an incident.

FIX Move deployments into a CI/CD pipeline that deploys your CloudFormation templates from version control, with promotion between environments and a one-step rollback, so every change is reviewed, repeatable and reversible. That is all it needs to be. A team your size does not need a GitOps stack or a dedicated platform team for this. How much work it is depends on what has to be built first (container images, AMIs, the promotion flow), and the pipeline needs somewhere to run, either managed (CodePipeline and CodeBuild) or self-hosted runners, so it carries a small ongoing cost.

6 No service-quota headroom ahead of launch

MEDIUM RISK

Reliability · Effort S · Cost Neutral

FOUND The account is **close to several default service quotas**, including Fargate On-Demand vCPUs, Elastic IPs and ALB rules per listener, and nothing monitors quota utilisation.

RISK Service quotas are account-level hard limits, not something autoscaling can grow past. If a launch or a campaign spike needs more Fargate tasks than your vCPU quota allows, those tasks simply will not start, and an increase can take hours or days to come through. So the limit stops you right when traffic is peaking. This is the kind of failure load testing rarely catches, because tests rarely run at full production scale.

FIX Ask for the increases you will need for launch early, as they can take days to come through. For ongoing cover, set CloudWatch alarms through Service Quotas on the limits that matter, both are free and need no Support plan. These alarms are one per quota, so there is no single "any quota near its limit" alarm to rely on, set them on the quotas that matter for launch. Trusted Advisor can watch a wider set of limits, but only on a paid Support plan, which a team your size rarely buys, so we keep an eye on this for you instead.

IMPROVEMENT PLAN

Ordered to clear the **highest-risk, lowest-cost work first**, and to leave time for the quota increases to come through:

1. **Straight away, the two zero-cost High risk Issues:** add the masking step to the staging refresh (Finding 1), then limit what a leaked key can reach with a scoped production role, MFA and key rotation (Finding 3).
2. **In parallel, ask for the quota increases you will need (Finding 6).** The risk only lands at launch, but the increases can take days to come through, so ask early.
3. **Next, close the reliability gap:** add the second-AZ read replica and test a restore (Finding 2).
4. **Before launch:** set up the deployment pipeline (Finding 5), and put non-production scheduling and account-wide cost alerts in place (Finding 4).

Done in this order, the three High Risk Issues are cleared early and the platform is ready for the v2.

NEXT STEPS

This Snapshot is the map, not just what to fix, but what you can safely leave for now.

The next step is a defined piece of remediation work to clear the High Risk Issues, scoped from the findings above. Once that work is defined, AWS Proof-of-Concept funding may be available to help cover it (it applies to defined improvement projects, not to the review itself). Want a Snapshot like this for your own AWS account? Start with a free second opinion, a 30-minute call with our founder, no account access needed. Book a time or email us using the details below.

web mysteriouscode.com

mail hello@mysteriouscode.com

book savvycal.com/mysteriouscode/intro