QBurst | A High AI-Q™ Company

AWS Cloud Operations Transformation for Global eCommerce Reliability and Scale

Re-engineering AWS cloud operations with SRE and AIOps to ensure 99.9% uptime during peak sales while optimizing costs.

Overview

QBurst introduced an SRE operating model paired with AIOps anomaly detection and demand-driven auto-scaling to modernize the client's multi-million-dollar AWS estate.

- **Significant Cost Savings:** Generated over $1.5M in annualized infrastructure cost optimization across 10 AWS regions and reduced operational costs by 20%.
- **High Automation & Speed:** Automated 70% of repetitive operational tasks, leading to a 30–40% faster mean time to detect and resolve (MTTD/MTTR) incidents.
- **Peak Event Uptime:** Maintained a 99.9% uptime availability objective through peak commercial events while saving $45K–$50K in Black Friday prep costs.

Client Profile

The client is a leading Japanese global fashion retailer operating large-scale e-commerce and customer support operations across multiple international markets. Managing high volumes of product enquiries, order support, returns, and customer engagement interactions, the retailer embraces scalable AI-driven capabilities to improve operational efficiency and deliver consistent customer experiences across regions.

Challenges: Managing Peak-Demand Scale

- The client's legacy operational setup exposed severe technical and financial limitations:

- **Costly Manual Scaling:** 24/7 manual pre-scaling to prevent downtime during unpredictable sales events was financially unfeasible and operationally unsustainable.
- **Latency & Revenue Loss:** Unpredictable traffic surges caused sub-second latency spikes, directly impacting conversion rates and leading to immediate cart abandonment.
- **Reactive & Siloed Operations:** Lack of clear reliability ownership resulted in production issues reaching customers before being detected internally, constantly pulling L3 feature developers into routine support.
- **Alert Noise & Static Limits:** Static alert thresholds failed to adapt to varying seasonal baselines, triggering severe alert storms during incident events.

QBurst Solution: SRE & AIOps Implementation

We executed a comprehensive cloud operations transformation by implementing an SRE operating model anchored by shared ownership, focused observability, AIOps automation, and demand-driven capacity management.

- **SRE & Shared Ownership Model:** Established reliability KPIs, defined SLAs/SLIs, enforced production readiness reviews, and implemented blameless root cause analysis (Five Whys) to shift routine support away from L3 developers.
- **Standardized Observability:** Unified telemetry using Amazon CloudWatch (for AWS infra metrics) and Datadog (for APM, distributed tracing, and log analytics) across EKS, ECS, Aurora, RDS, ElastiCache, and MSK.
- **AIOps & Auto-Remediation:** Integrated ML-driven anomaly detection to set dynamic baselines, automated alert correlation to stop alert storms, and connected Amazon EventBridge with AWS Lambda for runbook-based auto-remediation.
- **Demand-Driven Scaling:** Automated infrastructure provisioning using Terraform and Jenkins, implementing dynamic Auto Scaling parameters paired with Gatling performance testing and Aurora connection pool tuning.

Technical Highlights

The platform architecture utilizes a modern AWS cloud-native microservices stack:

- **Compute & Containerization:** Orchestrated microservices on Amazon EKS with a phased migration path from Amazon ECS, utilizing AWS Auto Scaling and Kubernetes HPA.
- **Data & Caching Tier:** Powered by Amazon Aurora, Amazon RDS, and Amazon ElastiCache (Redis) to hold latency down during heavy database read paths.
- **Event Backbone & Routing:** Utilized Amazon MSK (Managed Apache Kafka) for asynchronous microservices decoupling, Amazon Route 53 for global health checks/failover, and Elastic Load Balancing (ELB) for regional traffic distribution.
- **Infrastructure as Code (IaC) & CI/CD:** Automated environment provisioning via Terraform and deployments through Jenkins pipelines.

Impact

- **Annualized Infrastructure Cost Optimization:** Achieved $1.5M+ in annualized savings across 10 regions by moving from cost-prohibitive 24/7 manual pre-scaling to dynamic, demand-aligned capacity management.
- **Operational Cost Reduction:** Lowered operational costs by 20% by eliminating manual intervention, streamlining shift handovers, and reducing shift overlap overhead.
- **Accelerated Incident Response:** Cut MTTD and MTTR by 30–40% using ML-driven anomaly detection and alert correlation to detect and resolve system anomalies before they impacted customer checkout flows.
- **High-Volume Process Automation:** Automated ~70% of repetitive operational tasks using AWS EventBridge and AWS Lambda auto-remediation runbooks.
- **Optimized Peak Event Performance:** Saved $45K–$50K in Black Friday preparation costs alone while maintaining platform stability and eliminating cascading database failures during high-traffic surges.

- **Guaranteed Uptime & Reliability:** Consistently achieved the 99.9% availability objective with zero customer-facing downtime during peak seasonal traffic.
- **Reclaimed Engineering Capacity:** Shifted routine support away from L3 developers, returning senior engineering capacity back to core product feature delivery and strategic innovation.