QBurst | A High AI-Q™ Company

AI-Powered Edge & Enterprise Platform for Connected Operations

A multi-product, cloud and edge-native platform enabling real-time intelligence, secure operations, and scalable enterprise management across remote environments.

Client

The client is a technology-driven organization focused on building advanced platforms that leverage data, AI, and automation to transform enterprise operations. They deliver cloud-native solutions across areas such as risk management, security, compliance, and operational intelligence.

Engagement Overview

QBurst partnered with the client to design and deliver a connected, multi-product platform integrating cloud, edge infrastructure, AI, and enterprise systems.

The engagement spans platform engineering, AI-driven applications, cloud-native architecture, identity and access management, and observability systems, enabling real-time intelligence and resilient operations across distributed environments.

Engagement Scale & Reach

- **40+** member cross-functional engineering team
- **15+** interconnected platform modules delivered across AI, edge, cloud, and enterprise systems
- **100+** microservices and infrastructure components orchestrated across distributed environments
- Multi-cloud and edge deployments supporting real-time telemetry, video analytics, and operational intelligence

Key Engagement Pillars

- Cloud-native, microservices-driven platform architecture
- Edge computing and connected infrastructure enablement
- AI-powered operational and safety solutions
- Centralized identity and access management

- Monitoring and observability for real-time insights
- Scalable enterprise web and application platforms

Business Impact

- Unified architecture reduced platform fragmentation and enabled 3–5x faster onboarding of new services and use cases.
- Centralized observability improved issue detection time by 60% and reduced mean time to resolution (MTTR) by 50%.
- Cloud and edge-native deployments enabled high-volume telemetry and video workloads, with 45% improvement in processing latency compared to traditional architectures.
- Automated analysis and real-time processing reduced manual monitoring efforts by 75%, while enabling near real-time decision-making.
- Centralized identity and access management reduced administrative overhead by 40% and enabled consistent policy enforcement across services.
- Modernized enterprise interfaces improved internal adoption and reduced time-to-insight by 30%.

Key Project Snapshots

AI-Powered Safety Intelligence Platform

Objective

Develop an AI-driven solution capable of detecting safety risks in video data across cloud and edge environments, enabling real-time alerts and improving workplace safety.

Our Role

Designed and delivered a scalable AI platform integrating real-time video processing, cloud infrastructure, and edge deployment capabilities for safety monitoring.

Solution Highlights

- AI processing unit built in Python for detecting PPE compliance, fire hazards, and unsafe actions
- React-based user interface for monitoring, alerts, and reporting
- Spring Boot-powered backend services for API orchestration
- High-volume data streaming managed using Kafka and GPU resources
- Cloud infrastructure implemented using Azure services, including storage, event streaming, database, and Kubernetes
- GitOps-driven automated deployments to self-hosted Kubernetes across distributed edge locations, enabling reliable processing in remote environments

Impact

- Reduced manual safety monitoring efforts by 70% through automated video analysis
- Enabled near real-time incident detection with significantly lower latency (50% improvement)
- Scaled to handle high-throughput video streams across distributed environments
- Improved responsiveness to safety events through automated alerts and reporting

---

Azure Service Onboarding Automation for FedRAMP Compliance

Objective

Design and deliver an automated infrastructure provisioning framework to onboard microservices onto Azure Public and Government Cloud (AGC) in compliance with FedRAMP Moderate requirements — eliminating manual configuration, enforcing security controls consistently, and enabling auditable, repeatable deployments at scale.

Our Role

Architected and implemented a declarative, policy-driven service onboarding platform that provisions all required Azure infrastructure for each microservice across commercial and government cloud environments, integrating GitOps-based workflows, identity federation, and compliance-by-default resource configuration.

Solution Highlights

- Built a centralized, declarative onboarding framework where new services can be added through simple configuration files instead of custom engineering work, making onboarding faster and more consistent.
- Automated the setup of all required Azure infrastructure through reusable **OpenTofu/Terraform** Infrastructure-as-Code (IaC) templates, including networking, identities, databases, security controls, and access management.
- Designed the platform to work seamlessly across both Azure Public Cloud and Azure Government Cloud using a single shared framework, reducing duplication and operational complexity.
- Implemented GitOps-driven, approval-based deployment workflows so all infrastructure changes are automatically reviewed, validated, and tracked before going live.
- Replaced stored credentials and manual secret management with secure **OIDC-based federated identity** authentication, significantly improving security and reducing operational overhead.
- Enforced security and compliance standards by default, including private endpoints , centralized secret management, role-based access controls (RBAC), and secure database authentication.
- Isolated infrastructure deployments for each service and environment, ensuring issues in one deployment do not impact other applications or teams.

Impact

- **90%+** reduction in service onboarding time, from days to under two hours, with **30+** Azure resources auto-provisioned per service, including identity, networking, secrets, and databases.

- **150+** manual provisioning steps eliminated per onboarding, replacing fragmented portal-based setup (including access management, network configuration, security rules, and credential provisioning) with a fully automated, configuration-driven onboarding model.
- **100%** environment parity across dev, staging, and production, eliminating configuration drift via a single declarative source of truth validated at every deployment.
- **95%+** reduction in infrastructure deployment errors through policy-enforced IaC replacing manual provisioning workflows.
- Eliminated long-lived CI/CD credentials, with full migration to OIDC-based federated identity and removal of stored secrets and key rotation overhead.
- Achieved FedRAMP Moderate compliance consistency across all services, with fully logged, traceable, and audit-ready provisioning actions.
- Zero-dependency service onboarding enabled for engineering teams, allowing new services to be provisioned via configuration after initial platform setup.