SOLUTION BRIEF

Ataya Edge

Connect. Compute. Secure. At the Site Edge.

Enterprises deploying AI-driven operations face a growing architectural challenge: how to run real-time AI workloads where they are needed most — at the site edge — while maintaining secure, reliable connectivity across diverse device environments. Traditional approaches force organizations into painful trade-offs: ship data to the cloud for inference and accept latency that makes real-time decisions impossible, or deploy isolated point solutions for networking, compute, and security that multiply operational complexity and cost.

Ataya Edge is a unified platform that resolves this tension. It converges private 5G connectivity, edge AI compute orchestration, and Zero Trust security into a single, cohesive architecture deployed at the site edge. This brief describes the Ataya Edge platform architecture, its deployment models, and the use cases it enables across manufacturing, logistics, healthcare, and other industrial environments.

The Challenge: AI at the Point of Action

Modern industrial operations depend on devices that must act instantly and intelligently: autonomous mobile robots navigating a factory floor, drones conducting infrastructure inspections, vision AI cameras detecting safety hazards, and AGVs executing precise logistics workflows. These use cases share a common requirement — millisecond-class latency and uninterrupted connectivity that cloud-dependent architectures cannot reliably deliver.

At the same time, enterprise IT and OT teams are under pressure to reduce operational complexity and vendor sprawl. Managing separate networking infrastructure, AI compute clusters, and security stacks — each from different vendors, each with its own management plane — creates fragility, slows deployment, and inflates total cost of ownership.

Ataya Edge addresses both dimensions simultaneously: it puts AI inference, network control, and security enforcement at the site edge, unified under a single management architecture.

Platform Architecture

Ataya Edge is structured around three tightly integrated layers: the Management and Control Plane, the Site Edge Layer, and the Connectivity Layer.

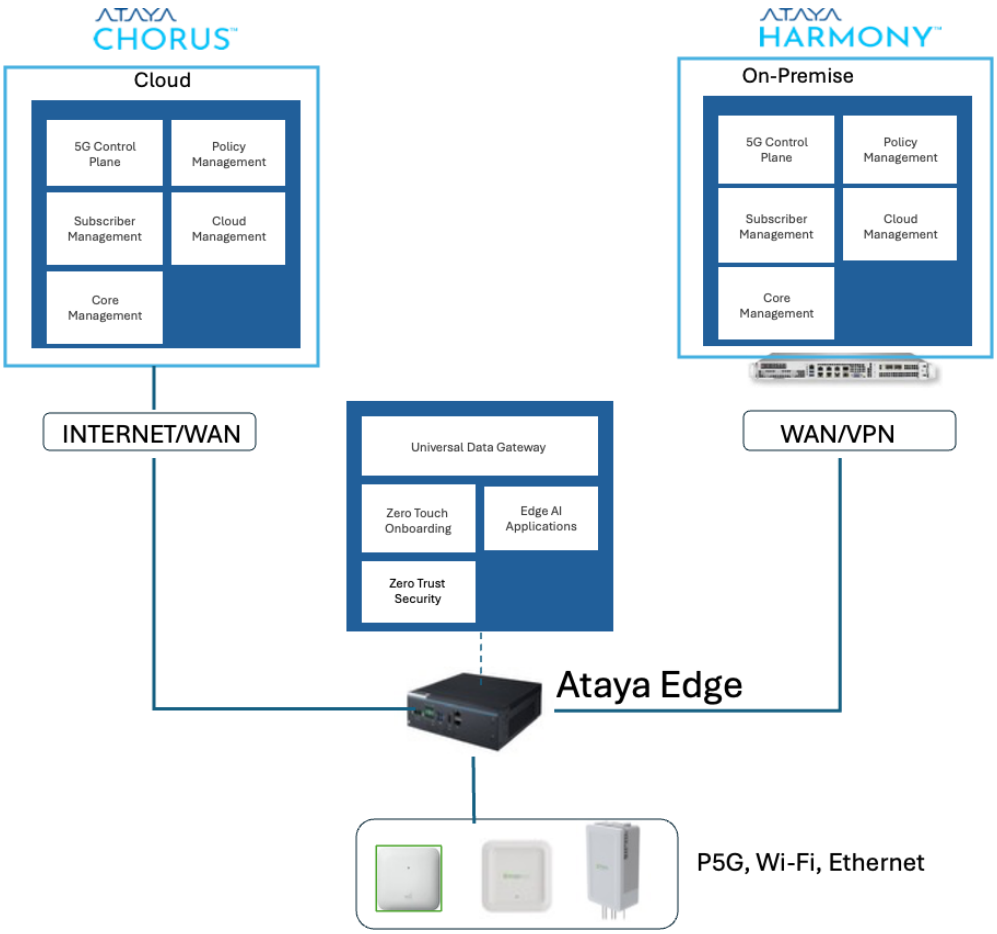

Fig 1: Ataya Edge connected to either Ataya Chorus or Ataya Harmony

Management and Control Plane

Ataya Edge supports two deployment models for the 5G control plane, giving organizations the flexibility to match their operational requirements and IT policies.

Chorus Cloud is a cloud-managed private 5G platform running on AWS. The 5G Control Plane functions — including AMF, SMF, and subscriber policy — run as cloud-native services on AWS EKS. Zero-Touch Onboarding allows new edge sites to self-register and pull configuration automatically upon connection, eliminating the need for on-site technical personnel during deployment. Chorus Cloud is optimized for organizations that want simplified operations, elastic scalability, and centralized management across multiple distributed sites.

Harmony is the on-premise alternative, hosting the 5G Core, subscriber management, and policy enforcement entirely within the customer's own facility. Control plane traffic never leaves the local network, making Harmony the preferred choice for organizations with strict data sovereignty requirements or regulatory constraints on cloud connectivity. Harmony connects to edge infrastructure via WAN/VPN.

Both deployment models share the same site edge architecture described below, allowing customers to choose or migrate between management approaches without replacing edge hardware.

Site Edge Layer — The Core of Ataya Edge

The site edge layer acts as the Policy Enforcement Point for the site — applying QoS rules that prioritize traffic by application and device type, and enforcing security policies that govern what each device and workload is permitted to access. All enforcement happens locally, ensuring that policy is effective even when WAN connectivity to the cloud or central core is degraded or unavailable.

Three integrated functions run on edge compute hardware deployed at each site:

Universal Gateway serves as the unified data gateway for all access technologies at the site edge — not just 5G, but also Wi-Fi and Ethernet-connected endpoints. For 5G devices, the UPF performs local traffic breakout so that user plane data (N3 interface) is processed entirely on-site, with only N2 control plane signaling traversing the WAN. For Wi-Fi and Ethernet endpoints, the same gateway function applies consistent policy, QoS, and traffic handling regardless of how a device connects. The result is that AI inference, video analytics, and real-time control messages all operate on sub-10ms timescales, whether the originating device connects over 5G, Wi-Fi, or Ethernet.

AI Orchestration enables deployment and management of AI workloads — from large-scale GPU inference servers to compact edge inference modules — through a single orchestration framework. AI applications running at the edge include computer vision for quality inspection, anomaly detection on sensor streams, predictive maintenance models, and autonomous navigation for robots and AGVs. The orchestration layer schedules workloads, manages inference pipelines, and allocates compute resources dynamically based on application priority and available edge capacity.

Zero Trust Security enforces identity-based access, micro-segmentation, and continuous threat detection without requiring a separate security stack. ZTNA is implemented natively within the edge layer, ensuring that every device, application, and user is authenticated and authorized before accessing network resources. AI applications run in dedicated namespaces and are isolated from one another through explicit ingress and egress traffic rules — a compromised inference workload cannot reach adjacent applications or network segments. Because security enforcement happens locally at the edge, policy remains effective even if WAN connectivity is interrupted.

Connectivity Layer

Ataya Edge unifies three access technologies under a single management and QoS framework:

- **Private 5G** (Indoor and Outdoor radios) provides wide-area, high-density wireless connectivity with SIM-based device authentication, deterministic QoS, and sub-10ms latency. Supported bands include n48 (CBRS), n77, n78, and n79. RedCap devices are supported for IoT sensors and wearables.
- **Wi-Fi Access Points** extend connectivity to devices and areas where 5G is not optimal or where legacy Wi-Fi devices must be supported.
- **Ethernet Switching** connects fixed infrastructure — production equipment, cameras, servers — into the same policy-managed network fabric.

Application-aware QoS policies are applied uniformly across all three access technologies, ensuring that latency-sensitive workloads like autonomous navigation or real-time video receive prioritized treatment regardless of connection type.

Key Platform Capabilities

Edge AI App Enablement

Ataya Edge provides a unified runtime for deploying AI workloads across the full range of edge compute — from high-performance GPU inference servers to compact, power-efficient inference accelerators. A single orchestration plane manages workload scheduling, model versioning, and resource allocation, eliminating the need to maintain separate AI infrastructure management tools. Organizations can bring their own models and inference frameworks and deploy them to any site without re-architecting the underlying infrastructure.

Always-On Connectivity

Private 5G, Wi-Fi, and Ethernet are unified within the Ataya Edge management plane, with consistent QoS policies applied across all access types. Sub-10ms local latency is built into the architecture through local UPF breakout rather than being bolted on as an afterthought. Network health, device connectivity, and policy enforcement are monitored centrally, with anomaly detection alerting operations teams before issues affect production.

Built-In Zero Trust Security

Unlike traditional approaches that treat security as an overlay added after the network is designed, Ataya Edge embeds Zero Trust from the ground up. 5G's native SIM-based authentication (5G-AKA) ensures only provisioned devices can connect — credential sharing and rogue device risks that affect Wi-Fi are eliminated by design. Micro-segmentation isolates traffic flows between device groups, applications, and network zones. AI applications run in dedicated namespaces with explicit ingress and egress traffic rules, ensuring that workloads are isolated from one another and cannot communicate beyond their defined scope.

10x Faster Deployment

Zero-Touch Onboarding allows new sites to self-configure upon connection — the edge node discovers the cloud infrastructure, retrieves its configuration, and registers with the 5G core automatically. Standardized, pre-configured hardware at the edge eliminates the need for specialized on-site commissioning. Organizations report up to 10x reduction in deployment time compared to traditional private 5G or enterprise network build-outs.

40% Lower TCO

By converging networking, AI compute orchestration, and security into a single platform and management system, Ataya Edge eliminates redundant infrastructure, licensing, and operational overhead. The integrated stack replaces point solutions for private 5G, Wi-Fi management, edge compute orchestration, and network security — each of which previously carried independent hardware, software, and support costs.

Vendor Agnostic

Ataya Edge is designed around open standards and interoperability. It supports multiple radio vendors, standard compute hardware, and integrates with third-party AI frameworks and applications. Customers are not locked into proprietary ecosystems for either connectivity or compute and can evolve their hardware choices independently as technology and market conditions change.

Industry Applications

Manufacturing and Industry 4.0

Factory floors running autonomous robots, AGVs, and vision AI inspection systems require deterministic, low-latency connectivity that Wi-Fi alone cannot consistently deliver. Ataya Edge provides the 5G connectivity and local AI inference needed to support closed-loop automation — where sensor data, AI decisions, and actuator commands must complete within milliseconds. A single edge node can simultaneously support robotic navigation, machine vision quality control, worker safety monitoring, and IoT telemetry collection without separate infrastructure for each application.

Logistics and Warehousing

Modern distribution centers depend on fleets of autonomous vehicles, handheld scanning devices, and real-time inventory systems operating in concert. Ataya Edge unifies connectivity across these device classes with QoS policies that prioritize AGV control signals over background inventory synchronization traffic. AI workloads for path planning and collision avoidance run locally, ensuring that autonomous operations continue reliably even if WAN connectivity is temporarily degraded.

Healthcare and Life Sciences

Healthcare campuses require secure wireless connectivity for a diverse mix of clinical devices, mobile workers, and IoT-connected medical equipment. Ataya Edge's Zero Trust architecture provides the security posture that healthcare environments demand, with SIM-based device authentication and micro-segmentation isolating clinical device traffic from general-purpose networks. Edge AI enables real-time processing of patient monitoring data and imaging workflows without transmitting sensitive data to external cloud infrastructure.

Construction and Temporary Sites

Large construction sites present a recurring operational challenge: delivering enterprise-grade connectivity and AI-assisted monitoring across a site that changes in layout, often lacks fixed infrastructure, and must be stood up and torn down within weeks or months. Ataya Edge's rapid deployment model — days rather than months — and vendor-agnostic hardware support allow project teams to establish 5G coverage, drone inspection workflows, and site safety monitoring quickly, then relocate the infrastructure to the next project when work is complete.

Critical Infrastructure

Energy facilities, water treatment plants, and transportation hubs require secure, reliable connectivity for SCADA systems, environmental sensors, and remote monitoring applications. Harmony's on-premise deployment model ensures that control plane traffic never leaves the facility network, meeting data sovereignty requirements common in these environments. Zero Trust micro-segmentation enforces strict separation between OT device traffic and IT management systems, reducing the attack surface for cyber threats targeting critical infrastructure.

Education and Campus Environments

Large campuses — universities, school districts, and corporate facilities — need flexible, scalable wireless coverage across both indoor and outdoor environments. Ataya Edge supports seamless roaming across 5G and Wi-Fi access points, enabling continuous connectivity for mobile devices, IoT sensors, and video applications as users and devices move across campus. Centralized management

allows IT teams to apply consistent policies across all access types without managing separate infrastructure for each technology.

Deployment Architecture Summary

The following table summarizes the key architectural characteristics of each Ataya Edge deployment model:

Characteristic	Chorus Cloud	Harmony On-Premise
5G Control Plane	AWS (cloud-native)	On-premise server
WAN Dependency	Required for control plane	Required for management only
Zero-Touch Onboarding	Supported	Supported
Data Sovereignty	User plane local	Full local control
Scalability	Elastic (AWS EKS)	Capacity-planned
Best For	Multi-site, distributed	Multi-site, regulated

In both models, the site edge layer — Universal Gateway, AI Orchestration, Zero Trust Security — is deployed identically on-site, ensuring that latency, security, and AI execution characteristics are consistent regardless of which control plane model is selected.

Conclusion

Ataya Edge represents a fundamental architectural shift for industrial connectivity and edge AI. By converging private 5G, multi-access networking, edge AI orchestration, and Zero Trust security into a single platform — and deploying that platform at the site edge rather than in a central data center or cloud — Ataya enables organizations to run AI workloads at the speed and location where they create value.

The result is measurably faster deployment, lower total cost of ownership, and a security posture that is built into the network rather than layered on top of it. For enterprises navigating the transition to AI-driven industrial operations, Ataya Edge provides the foundation to connect, compute, and secure at the site edge.

Architecture and Acronym Reference

Term	Description
UPF / Universal Gateway	User Plane Function — handles user data traffic (forwarding, routing, NAT) across 5G, Wi-Fi, and Ethernet endpoints. In Ataya Edge, deployed at the site edge for local breakout across all access technologies.
AMF	Access and Mobility Management Function — manages device registration, connection, and mobility in the 5G core.
SMF	Session Management Function — manages data sessions, IP address assignment, and QoS policy.
gNB	Next Generation Node B — the 5G base station (radio) providing wireless connectivity to user devices.
UE	User Equipment — any device connecting to the 5G network (phones, sensors, robots, AGVs, etc.).
ZTNA	Zero Trust Network Access — identity-based access control model where no device or user is trusted by default.
N2 Interface	Control plane interface between gNB and AMF. Carries signaling for device registration and session setup.
N3 Interface	User plane interface between gNB and UPF. Carries actual user data traffic.
N6 Interface	Data network interface between UPF and external networks or applications.
QoS	Quality of Service — policies prioritizing traffic types to meet application performance requirements.
RedCap	Reduced Capability (3GPP Release 17) — lighter 5G device category for IoT sensors and industrial monitors.
CBRS	Citizens Broadband Radio Service — shared spectrum band (n48) enabling private LTE/5G deployments in the US.
5G-AKA	5G Authentication and Key Agreement — native 5G protocol for authenticating devices using SIM credentials.
EKS	Amazon Elastic Kubernetes Service — managed Kubernetes platform used to run Chorus cloud-native 5G core functions.
SCADA	Supervisory Control and Data Acquisition — industrial control systems for monitoring and managing critical infrastructure.