

Collapsing the OSS Stack to Achieve  
AN-4 Autonomous Networks

The logo for Axon Maestro, consisting of a black cloud shape with the 'AN' monogram inside, followed by the text 'AXON Maestro™' in a large, bold, sans-serif font.

**AXON Maestro™**

By **Kumar Vishwanathan**  
Chief Product Officer  
& **Siddhartha Dattagupta**  
Chief Technology Officer

# 1. EXECUTIVE SUMMARY

Telecommunications networks have evolved into highly distributed, software-driven ecosystems spanning Fiber, xDSL, HFC, 5G FWA, Wi-Fi, edge compute, and cloud-native cores. Yet the Operations Support Systems (OSS) that manage these networks remain architecturally fragmented.

Traditional OSS stacks – built as siloed layers of inventory, configuration, network management, service orchestration, and assurance – are structurally incapable of supporting real-time, customer-experience-driven, autonomous operations.

**The consequences are systemic:**

- Manual provisioning and reconciliation
- Configuration drift
- Reactive fault management
- Disconnected capacity planning
- High OPEX
- Slower service innovation

AXON Maestro transforms this model.



**It collapses legacy OSS layers into a unified Autonomous Network Operating Fabric, powered by:**

- A real-time, customer-aware Digital Twin with continuous data assimilation and reconciliation rendering the Digital Twin as the source of truth.
- Closed-loop automation from CPE devices at the edge to the core
- Cross-domain orchestration across network and access technologies
- Agentic AI as an operational control plane

**The result:** a practical and scalable path to TM Forum's AN-4 (see Section 5 below for a complete definition) autonomous network maturity, enabling predictive, self-healing, experience-driven operations. This facilitates an Operator-as-a-Service (OaaS) business model with complete autonomy of network operations delivering the business intent.

## 2. THE CTO IMPERATIVE: OSS MUST EVOLVE

Today's telecommunications networks have become structurally unlike anything their original OSS architectures were designed to manage. Network functions once running on dedicated hardware are now virtualized and containerized, deployed dynamically across distributed edge nodes and cloud-native cores. Multi-vendor environments are the norm, not the exception, spanning fiber access, HFC, xDSL, 5G FWA, and Wi-Fi within a single operator footprint. Service boundaries no longer align with technology domains—a single customer session may traverse multiple access technologies, aggregation layers, and core segments simultaneously, each with its own performance characteristics and failure modes. SLA commitments have evolved from simple availability metrics to experience-based guarantees that measure what the customer perceives, not just what the network reports. And the pace of service innovation has compressed timelines from months to weeks, requiring operators to launch, modify, and retire network services with a speed that legacy provisioning models cannot support (more on this later).

Each of these shifts independently strains traditional OSS. Together, they expose a fundamental incompatibility: the operational complexity of modern networks has scaled non-linearly, while the OSS architectures managing them have not.

- Services cross the boundaries of multi-variate complex network segments while maintaining deterministic KPIs and SLAs.

However, traditional OSS remains layered as:

- Configuration Management (CMDB/Inventory)
- Network Inventory
- Element & Network Management (FCAPS)
- Service Fulfillment
- Capacity Planning

Traditional telecom OSS architectures were designed as discrete, layered systems – inventory, configuration management, fault management, service orchestration, and assurance – each operating with its own independent data model. While this structure enabled modular evolution over time, it fundamentally limits real-time responsiveness. Correlation across systems is weak because data synchronization is periodic rather than continuous. Reconciliation between logical and physical states is often manual or batch driven. Automation, where present, is typically rule-based and reactive, triggered after thresholds are breached rather than driven by predictive intelligence. In such an environment, no single system has a complete, continuously updated understanding of the network’s operational state, making true autonomous behavior structurally impossible.

Autonomous Networks require something fundamentally different: a collapsed, AI-driven OSS architecture built around a real-time Digital Twin that models the network end-to-end. A unified Digital Twin maintains a continuously reconciled representation of topology, services, configurations, telemetry, and customer experience from CPE to core. This shared intelligence layer eliminates data silos, enables instantaneous correlation of events, and provides the contextual foundation required for Agentic AI to observe, decide, and act in closed loop. Only with a real-time, holistic network model can the system react dynamically to changing conditions, distinguish between localized and systemic issues, and execute predictive remediation. This contrasts with the model of other software vendors. This contrasts with the model of other software vendors. Unencumbered by legacy software that must be maintained and supported like other vendors in the market, AXON Networks was able think outside-of-the-box and drive a new architecture with a clean slate. In short, autonomous operation is not achievable through layered OSS integration – it requires architectural collapse around a living Digital Twin that becomes the brain of the network. Below, we will explore some of the limitations that will be experienced by telecoms who fail to make the leap to a unified, real-time Digital Twin based architecture.

## 2.1 Data Integrity Risk

In a layered OSS architecture where each functional domain—inventory, configuration management, fault management, service orchestration, and assurance—maintains its own database, data integrity becomes a systemic risk rather than an isolated issue. Each system evolves its own schema, update cadence, and reconciliation logic, resulting in multiple “versions of truth” across the organization. Logical inventory may not match physical deployment, service topology may not reflect real-time configuration changes, and performance data may

be interpreted without full contextual alignment. Synchronization between systems is often batch-driven or manually validated, introducing latency and human error into what should be deterministic processes. Over time, configuration drift, stale records, and inconsistent correlations erode operational confidence, making automation unreliable and root cause analysis increasingly complex. In such an environment, decision-making—whether human or AI-driven—is based on fragmented and potentially outdated data, directly undermining the foundation required for closed-loop, autonomous network operations.

## 2.2 Operational Latency

In a layered OSS environment, operational latency is an inherent structural consequence of cross-system handoffs. Fault resolution and service provisioning workflows traverse multiple independent platforms—inventory, configuration management, orchestration, assurance—each requiring data validation, synchronization, and often manual intervention before the next step can proceed. Events detected in one system must be correlated in another, verified against inventory in a third, and executed through yet another control layer. These sequential dependencies introduce delays that compound under scale, particularly in multi-vendor, multi-domain networks. Even when automation exists, it is frequently confined within system boundaries, forcing inter-system coordination through APIs or ticket-driven workflows. The result is slower mean time to repair (MTTR), delayed service activation, and limited ability to respond dynamically to network conditions—directly constraining the real-time responsiveness required for autonomous network operations.

## 2.3 Planning Disconnect

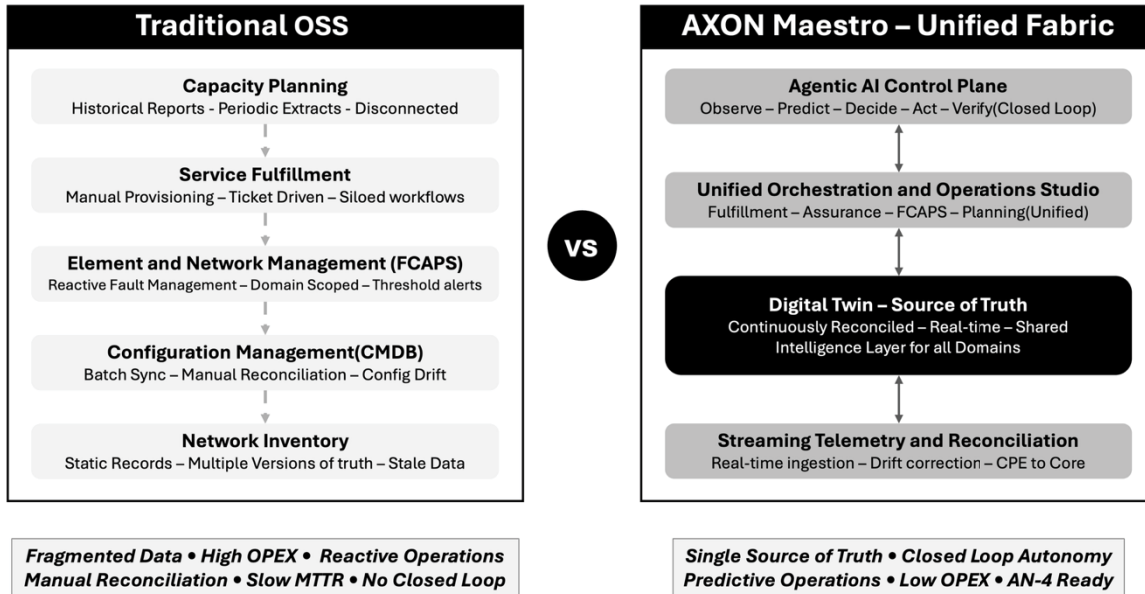
In a traditional layered OSS environment, planning functions—such as capacity forecasting, network expansion modeling, and investment prioritization—operate largely disconnected from real-time operational data. Forecasts are typically based on historical reports, aggregated performance counters, or periodic data extracts rather than continuously validated network state. Because planning systems do not share a live, reconciled model of topology, configuration, and service relationships, assumptions may not reflect current utilization, transient congestion patterns, or configuration drift. This disconnect creates risk: over-provisioning where capacity appears constrained but is misconfigured, or under-provisioning where localized saturation is masked by aggregated metrics. Without real-time validation against an accurate Digital Twin, planning remains a backward-looking exercise, limiting the operator's ability to proactively optimize resources and align capital investments with actual network dynamics—an essential requirement for autonomous network evolution.



## 3. AXON OPERATOR-AS-A-SERVICE (OaaS) VISION

The risks inherent in a layered OSS model are structural, not procedural. Data inconsistency, operational latency, and planning disconnect are not problems that can be fully resolved through

incremental automation or point integrations. Adding AI or workflow tools on top of fragmented systems only accelerates existing inefficiencies; it does not eliminate them. As long as multiple databases, siloed domains, and cross-system handoffs remain, true real-time autonomy is unattainable. To achieve Autonomous Network maturity, the operating model itself must evolve—from loosely integrated layers to a unified, AI-driven architecture centered around a real-time, continuously reconciled Digital Twin.



AXON Maestro doesn't automate the old model. It replaces it. A unified autonomous fabric that thinks, acts, and self-corrects at network speed – with humans setting the intent, not managing the complexity.

Rather than layering automation onto legacy systems, AXON Maestro collapses and unifies them into a single operating fabric built on three pillars:

1. Digital Twin Data Foundation
2. Unified Orchestration & Operations Studio
3. Agentic AI Control Plane

### 3.1 Pillar 1: Digital Twin Data Foundation

The Digital Twin is the foundational layer of the entire AXON Maestro architecture. Unlike traditional CMDBs or inventory systems that are periodically updated, the Digital Twin maintains a continuously reconciled, real-time representation of every element in the network—physical topology, logical configurations, service relationships, telemetry streams, and customer experience context from CPE to core. It is not a reporting tool or a passive model. It is a living state engine: the single authoritative source of truth against which all orchestration decisions are validated, all fault events are contextualized, and all AI actions are grounded. By eliminating the fragmented, multi-database model of traditional OSS, the Digital Twin makes cross-domain correlation native, configuration drift detectable in real-time, and closed-loop automation structurally possible.

## 3.2 Pillar 2: Unified Orchestration & Operations Studio

The Operations Studio collapses the historically fragmented workflow layers—service fulfillment, fault management, configuration management, capacity planning, and assurance—into a single, integrated operational environment. Rather than requiring operators to navigate between separate platforms and reconcile conflicting data, the Operations Studio provides a unified interface grounded in the Digital Twin’s real-time state. Service provisioning workflows execute against a consistent, validated model of the network. Fault management correlates events across all domains simultaneously rather than escalating through siloed ticketing chains. Capacity planning operates on live utilization data rather than periodic extracts. This convergence eliminates the cross-system handoffs and sequential dependencies that are the primary source of operational latency in traditional OSS environments, directly accelerating MTTR, service activation, and change velocity.

## 3.3 Pillar 3: Agentic AI Control Plane

Agentic AI in Maestro operates as a control plane, not a diagnostic assistant that adds diagnostics on the layer like other vendors. Because it is grounded in the continuously reconciled Digital Twin, the AI does not need to reconcile inconsistent data before acting—it observes accurate, contextualized network state in real-time and can reason, decide, and execute at machine speed. Specialized AI agents handle distinct operational domains: anomaly detection and root cause analysis, predictive congestion management, service assurance, and autonomous remediation. These agents operate in closed loop, with humans retaining supervisory control and configurable authority boundaries that determine when autonomous action is permitted versus when human approval is required. This model delivers the intent of AN-4 autonomy while maintaining the governance and auditability that production network operations demand.

Together, these three pillars form a mutually reinforcing architecture:

1. the Digital Twin supplies the intelligence layer,
2. the Operations Studio delivers the unified workflow environment, and
3. the Agentic AI provides the autonomous decision and execution engine.

No pillar is effective in isolation—it is their integration around a single, continuously reconciled model of the network that enables Maestro to deliver on the OaaS vision: an operator platform where network operations run with complete autonomy, aligned to business intent, and governed by the operator rather than constrained by system fragmentation.



**AXON Digital Twin™**

## 4. DIGITAL TWIN AS THE SOURCE OF TRUTH

Autonomous Networks cannot be built on fragmented, loosely synchronized data systems. They require a single, authoritative, continuously reconciled representation of the network’s state. This

is the role of the Digital Twin – not as a static model, but as a real-time operational intelligence layer that becomes the source of truth for the entire operating fabric.

## 4.1 From Static Inventory to Living Model

Traditional OSS architectures rely on Configuration Management Databases (CMDBs) and inventory systems that are updated periodically and reconciled manually. These systems store representations of devices, services, and topology, but they do not maintain continuous alignment with the live network. Over time, configuration drift, undocumented changes, and asynchronous updates create multiple “versions of truth.” Automation built on top of such systems inherits their inconsistencies and leads to delays, errors, and unnecessary costs in operations and performance.

A true Digital Twin replaces this static model with a continuously synchronized, dynamic representation of the network. It integrates telemetry, configuration state, topology relationships, service dependencies, and performance metrics into a unified data model that reflects the live operational environment.

The Digital Twin is not simply a reporting tool. It is the authoritative state engine for the network.

## 4.2 End-to-End Network Representation

For a Digital Twin to function as the source of truth, it must model the complete service path, including:

- Physical and logical topology
- Device configurations and policies
- Service-to-resource relationships
- Traffic flows and performance metrics
- Dependency mappings across domains
- Customer-experience context
- Network Element resource modeling with dependencies to network functions and states
- Device, environment and user behavioral contexts.

This end-to-end modeling ensures that any change in one domain – access, aggregation, core, or CPE – is immediately contextualized within the broader operational framework.

CPE-level intelligence is critical to making the Digital Twin truly representative of customer experience rather than merely network infrastructure. A significant percentage of perceived “network issues” originate not in the core or transport layers, but within the home environment—Wi-Fi interference, poor CPE placement, local congestion, or limitations of subscriber devices. Traditional OSS tools lack visibility into this last-mile domain and therefore cannot reliably distinguish between a true core network fault and localized CPE Wi-Fi degradation. This blind spot leads to false truck rolls, unnecessary escalations, extended mean-time-to-repair (MTTR), and ultimately customer dissatisfaction. By incorporating real-time CPE telemetry and Wi-Fi interface intelligence into the Digital Twin, AXON eliminates this ambiguity, enabling precise fault differentiation, faster resolution, and experience-centric autonomous operations.

Without this holistic representation, fault isolation, SLA validation, and capacity optimization remain domain-scoped and reactive.

## 4.3 Continuous Reconciliation

The defining characteristic of a true Digital Twin is continuous reconciliation.

**Rather than relying on batch updates or manual audits, the Twin:**

- Compares intended vs. actual configurations
- Detects configuration drift
- Identifies missing or inconsistent resources
- Validates service topology integrity
- Correlates performance data with structural dependencies

This constant validation maintains data integrity and operational confidence. It ensures that orchestration engines and AI agents operate on accurate, real-time information.

In an autonomous network, decisions are only as reliable as the data foundation that supports them. Continuous reconciliation makes the Digital Twin trustworthy.

## 4.4 Real-Time Telemetry Integration

To serve as the source of truth, the Digital Twin must ingest telemetry in real-time. Poll-based, periodic data collection is insufficient for dynamic network environments.

**High-fidelity telemetry enables:**

- Immediate detection of anomalies
- Microburst and transient congestion analysis
- Real-time SLA monitoring
- Cross-layer performance correlation

By integrating streaming telemetry and, where possible, on-device agents for granular visibility (including CPE and Wi-Fi layers), the Digital Twin reflects not only structural state but experiential reality.

This transforms the model from descriptive to actionable.

## 4.5 Contextual Intelligence Layer

**The Digital Twin becomes the shared intelligence layer for:**

- Orchestration systems
- Planning and simulation engines
- Assurance platforms
- Agentic AI control planes

Because all domains reference the same underlying model, correlation becomes native rather than inferred. Fault events are evaluated in full-service context. Capacity planning is validated against real utilization. Configuration changes are simulated before deployment.

The Digital Twin eliminates the need for inter-system reconciliation because there is only one system of truth.

## 4.6 Enabling Closed-Loop Autonomy

Autonomous Networks require the ability to:

1. Observe
2. Analyze
3. Decide
4. Act
5. Verify

The Digital Twin underpins each step:

- **Observe:** ingest telemetry and structural state
- **Analyze:** correlate across topology and services
- **Decide:** provide contextual input to AI agents
- **Act:** validate execution pathways
- **Verify:** confirm outcome and update state

Without a unified Digital Twin, closed-loop execution cannot be safely automated at scale. Decisions would rely on incomplete or inconsistent data, increasing risk rather than reducing it.

## 4.7 Planning and Simulation Grounded in Reality

A Digital Twin as the source of truth also bridges the gap between operations and planning. Capacity forecasts, expansion models, and service introduction scenarios can be simulated against the live, reconciled network model.

This enables:

- Accurate “what-if” analysis
- CAPEX optimization
- Proactive congestion mitigation
- Risk-free validation of configuration changes

Planning ceases to be a historical analysis exercise and becomes a real-time, data-validated function integrated into operations.

## 4.8 Foundation for AI-Driven Operations

Agentic AI systems require structured, contextualized, and accurate data to function effectively. When AI operates across fragmented systems, it must reconcile inconsistencies before making decisions. This increases latency and error potential.

## A unified Digital Twin provides:

- Structured topology context
- Service dependency awareness
- Accurate configuration state
- Performance baselines
- Customer experience mapping

This allows AI agents to operate as a control plane rather than a diagnostic assistant.

## 4.9 Single Source of Truth

The Digital Twin is not an enhancement to traditional OSS – it is the architectural foundation for Autonomous Networks. By replacing fragmented databases with a continuously reconciled, real-time operational model, it eliminates data integrity risk, reduces operational latency, and bridges the planning disconnect.

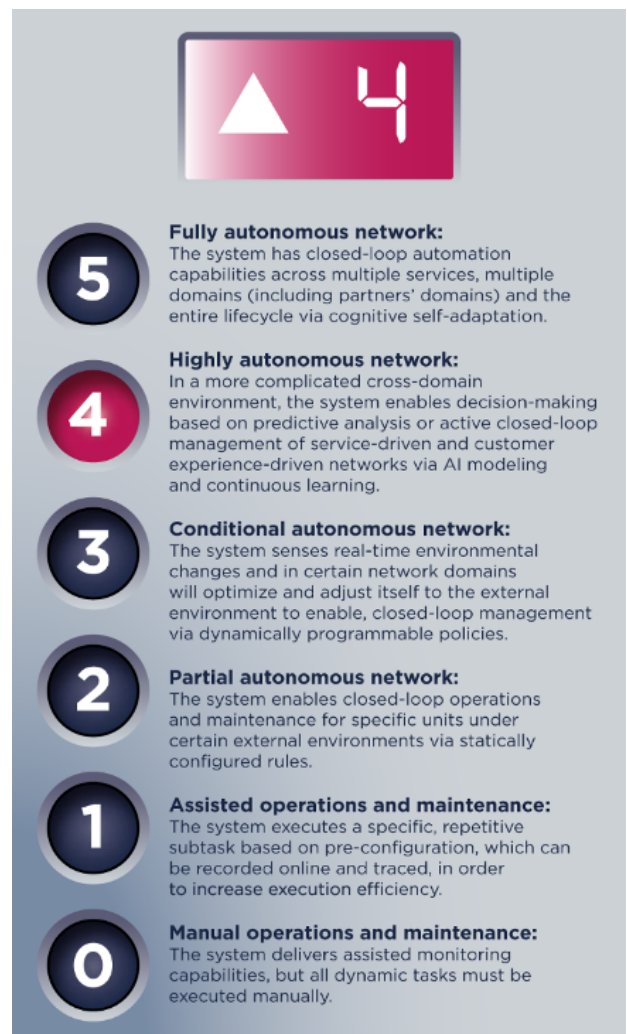
In an autonomous network, there can be no ambiguity about state. There must be a single, living source of truth. The Digital Twin fulfills this role, becoming the brain of the network and the enabling layer for closed-loop, AI-driven autonomy.

## 5. THE PATH TO AUTONOMOUS NETWORKS

Autonomous network maturity is not a single capability—it is a progression. The TM Forum AN-4 model defines a state in which the network continuously observes its own condition, predicts degradation before it occurs, executes corrective action without human initiation, and verifies that outcomes align with business intent.

Reaching AN-4 requires more than deploying AI tools or automating individual workflows. It requires a unified architectural foundation from which closed-loop intelligence can operate reliably, at scale, across all network domains simultaneously.

A new collaborative project from TM Forum will help simplify CSPs' path to achieving level 4 of network autonomy by identifying the most efficient autonomous network use cases.



Source: TM Forum Autonomy Framework

## 5.1 The AN-4 Autonomy Framework

At lower maturity levels (AN-1 through AN-3), operators progressively automate discrete tasks, introduce closed-loop remediation within individual domains, and begin integrating AI-assisted decision support. These stages have real value, but they do not eliminate the structural constraints of layered OSS. Event correlation still requires cross-system handoffs. Planning still operates on lagging data. Fault management still escalates through siloed workflows. AN-4 is the level at which these constraints are resolved—not by improving integration between fragmented systems, but by eliminating the fragmentation itself. AN-4 is also the level at which operators begin to see exponential impacts to their OPEX savings and service agility. Achieving AN-4 is a major leap, but essential for future operator success, and the foundation for the breakout AXON OaaS business model.

## 5.2 How Maestro Enables AN-4

AXON Maestro's architecture directly maps to each AN-4 requirement. The Digital Twin provides the real-time, end-to-end network model that makes accurate observation possible. Streaming telemetry integration, from core infrastructure to CPE and Wi-Fi endpoints, ensures that the observed state reflects what customers experience—not just what network elements report. Continuous reconciliation detects configuration drift and topology inconsistencies before they affect service, enabling proactive remediation rather than reactive response.

The Agentic AI Control Plane translates observation into action. Because agents operate on the same continuously reconciled model as the orchestration and assurance layers, they can distinguish localized CPE degradation from core infrastructure faults, prioritize remediation by customer impact, and execute corrective actions—configuration changes, traffic re-routing, resource reallocation—within operator-defined authority boundaries. Outcomes are verified against the Digital Twin, and the model is updated to reflect the new state. The loop closes without manual intervention.

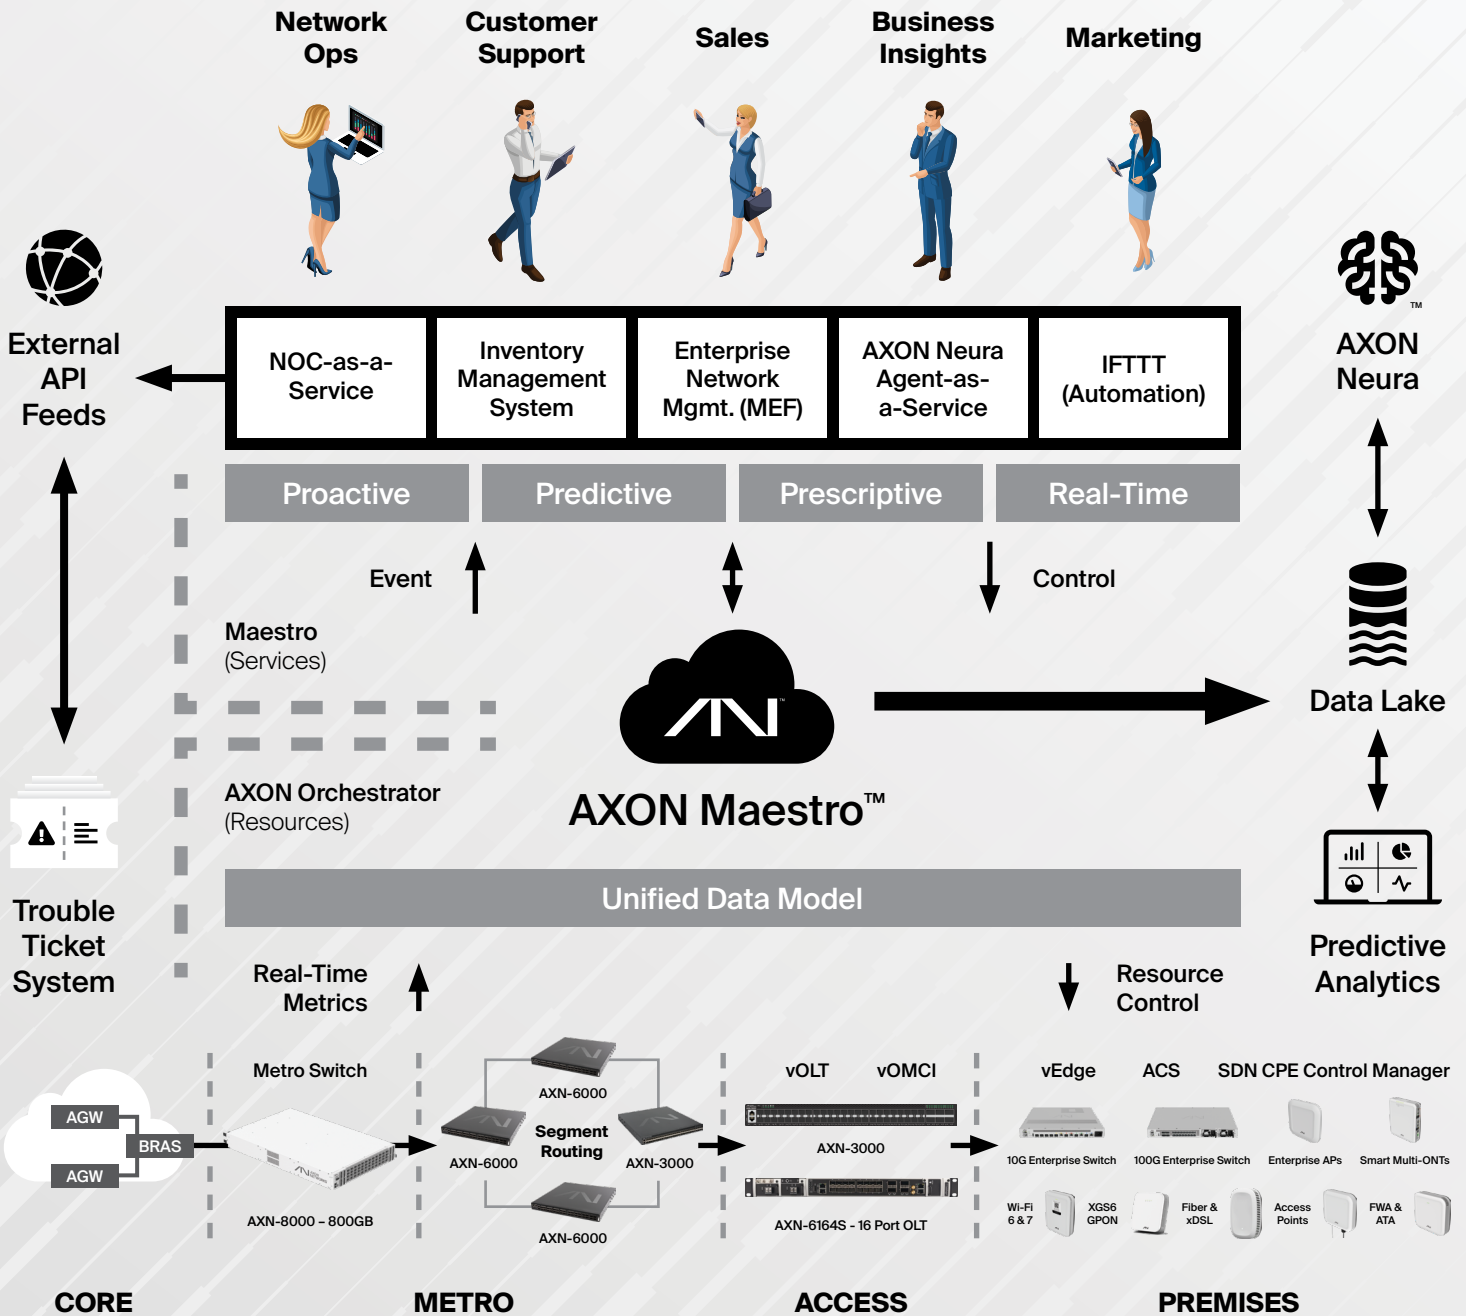
## 5.3 Operator-as-a-Service (OaaS): The Business Outcome

Reaching AN-4 autonomy is not simply an operational milestone—it is the technical foundation for the AXON OaaS business model. When network operations run with structural autonomy, the operator is no longer constrained by manual intervention cycles, OSS fragmentation, or reactive fault management. OPEX scales down as network complexity scales up. Service launch timelines compress because provisioning is driven by intent rather than manual configuration across siloed systems. Customer experience SLAs become enforceable in real-time because the system has both the observability and the agency to act before degradation reaches the subscriber. The AXON OaaS business model is transformational for the operator, positioning them to become a winning player in the AI economy.

AXON Maestro provides a practical, phased path to this outcome. Operators do not need to replace their entire OSS stack in a single transformation. Maestro is designed to ingest existing infrastructure data, layer the Digital Twin progressively across domains, and expand the scope of autonomous operation as confidence and coverage grow. The destination is AN-4 autonomy—a network that understands, adapts, and optimizes itself in alignment with business intent. The path is structured, measurable, and achievable. The impacts are exponential, reflecting an entirely new operating model for the operator going forward, and a service agility that operates in days and weeks, instead of months or years.

# Leverage Intelligent Cloud Orchestration Driven by AXON Maestro with Real-Time Digital Twin Telemetry to Win the AI Economy

The Power of API-First Design Enabled by AXON Digital Twin Technology



Contact our sales team to schedule a custom demo: [sales@axon-networks.com](mailto:sales@axon-networks.com)

