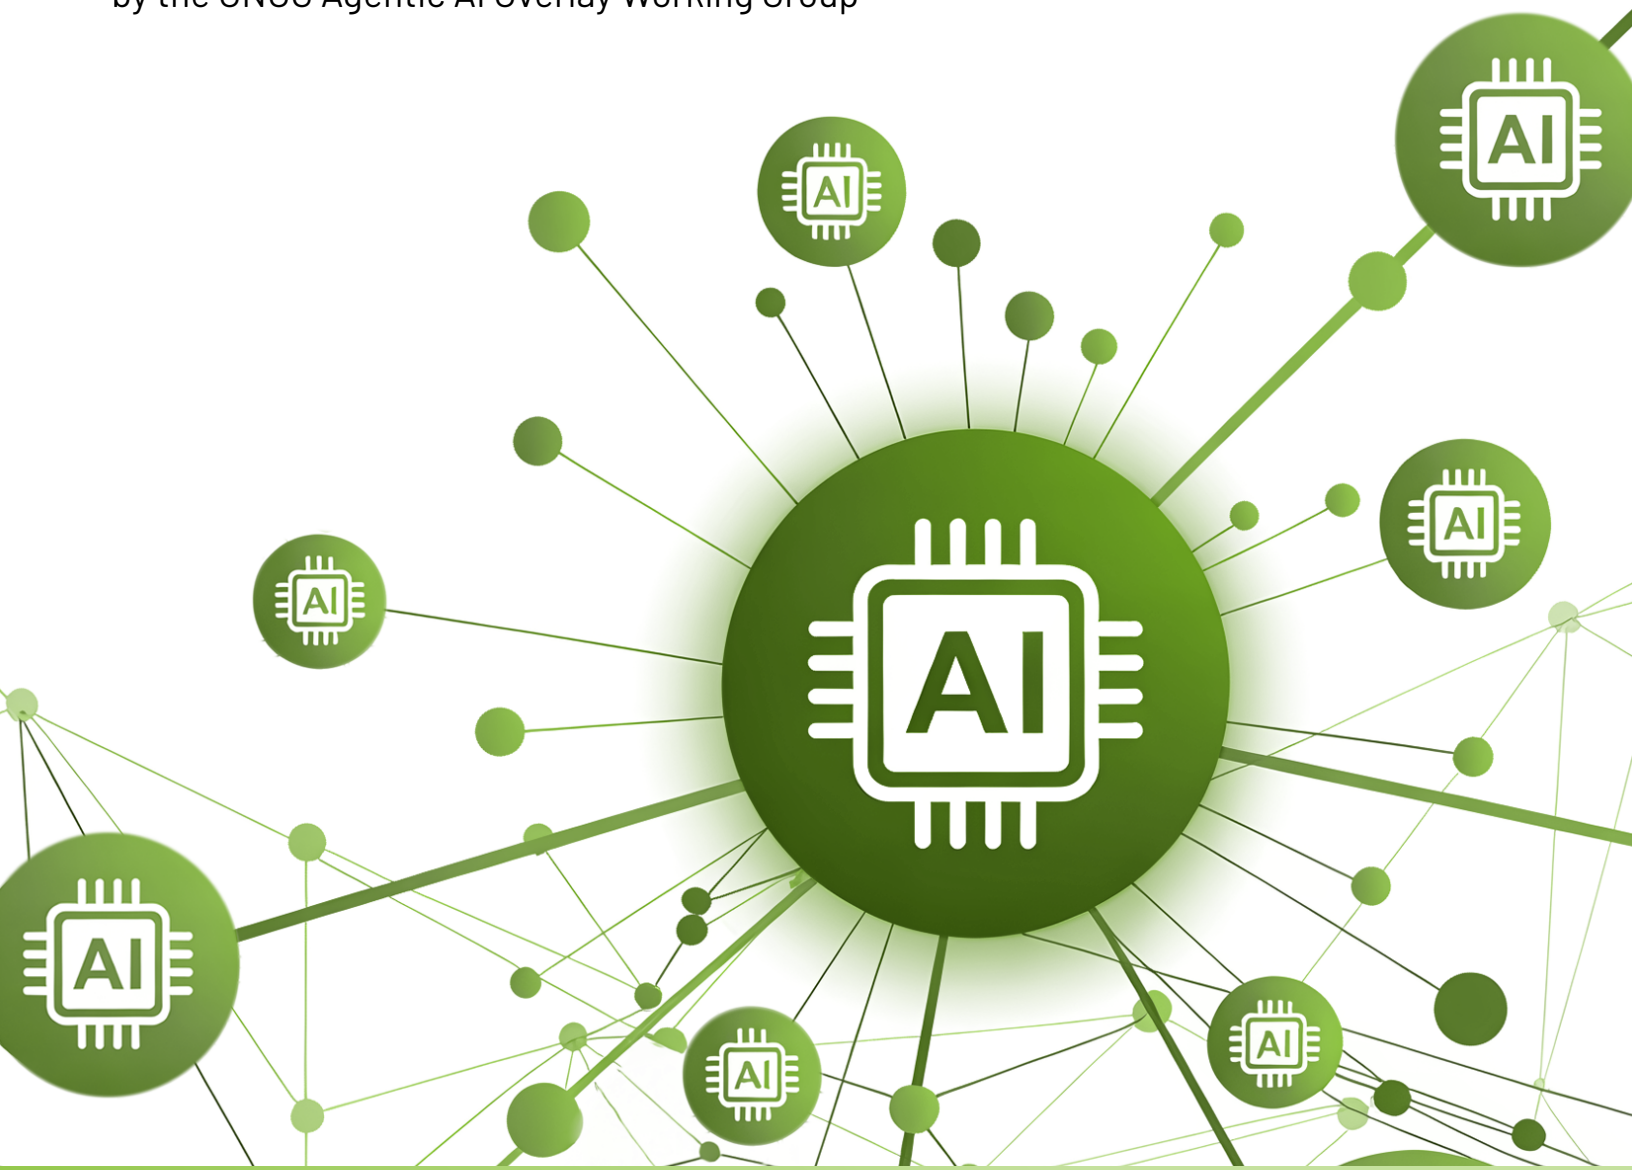
Large Enterprise Requirements to Scale Agentic AI

Part A: Agentic AI Overlay Architecture

by the ONUG Agentic AI Overlay Working Group



ONUG Collaborative | ONUG LLC
PO Box 8455, Naples, FL 34102 | info@onug.net

Executive Summary

The ONUG Agentic AI Overlay Working Group has produced a series of two papers to address large enterprise requirements for scaling agentic AI. This first paper describes a vendor-independent reference architecture aimed to assist in visualizing the various components and their interplay into one figure. The second paper is dedicated to scaling, deployment, and operationalization strategies.

Agentic AI represents a fundamental shift in how enterprises design, deploy, and operate intelligent systems. Unlike traditional AI applications or copilots, agentic systems are autonomous, persistent, and capable of coordinating actions across models, tools, data sources, and environments. The Agentic AI Overlay Reference Architecture responds to this challenge by defining a control-oriented, vendor-neutral approach for deploying agentic AI safely at scale. Rather than embedding trust and governance mechanisms inside individual agents or models, the overlay introduces an independent monitoring, control, and enforcement layer that spans private data centers, cloud environments, edge locations, and external services.

By reading these papers, enterprise IT leaders, architects, and decision-makers will gain a comprehensive understanding of the foundational architecture and practical strategies required to implement agentic AI with robust governance, enabling safe scaling while maintaining visibility, accountability, and compliance.

This work was created exclusively by IT executives from large enterprises, including the ONUG Board and its community, with leadership from eBay, Cigna, Bank of America, Indeed, Kraken, and others. It stems from discussions at the Fall 2025 AI Networking Summit and is validated by polling data from ONUG's 100,000+ member community.

This work will be used to compare and contrast the vendor community for their Agentic AI Overlay solutions. The ONUG Community encourages every vendor to become involved, as these are the issues keeping back Agentic AI in the large enterprise. This work is intended to help and support IT business leaders in the large enterprise make better decisions and speed their path to value of becoming enterprise AI companies.

Introduction

The Agentic AI Overlay Working Group is an ONUG Collaborative initiative bringing together enterprise IT executives, architects, and suppliers to define a vendor-independent reference architecture for AI-native networking. Our goal is to create an open, interoperable framework that enables secure, autonomous data exchange across on-premises data centers, public clouds, branch/edge environments, and operational technology (OT) domains.

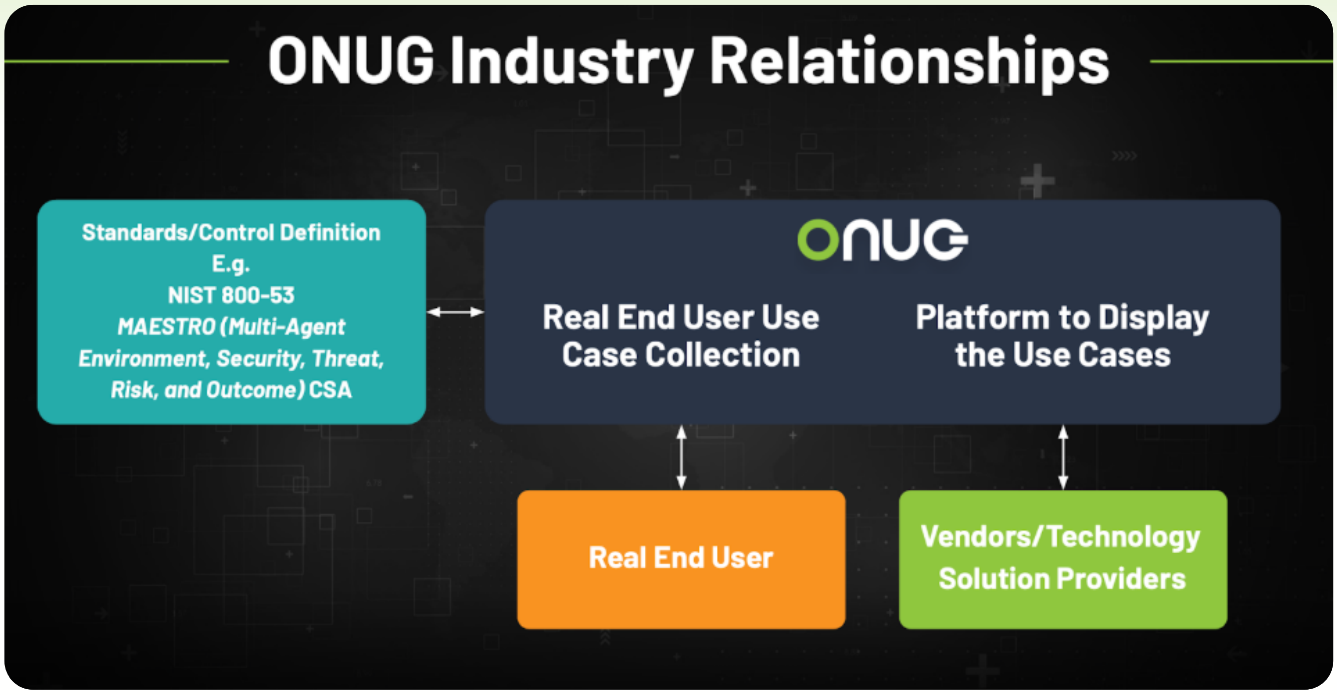
The Agentic AI Overlay defines the architectural framework required to securely deploy and operate autonomous AI agents within enterprise environments. This architecture focuses specifically on the structural components that enable agentic systems—including agent identity and attestation, agent-to-agent communication, capability directories, policy enforcement, secure connectivity, and runtime observability—as well as the communication flows between agents, tools, data sources, and infrastructure. Central to this design is the Agentic AI Overlay Control and Monitoring plane, which governs identity lifecycle, authorization, behavioral monitoring, rogue agent detection, and cross-domain trust enforcement.

This document concentrates on the governance, interoperability, and control mechanisms necessary to ensure distributed agentic systems operate securely and in alignment with enterprise policy across on-premises, cloud, edge, and OT environments.

The architecture provides:

- A defined set of core architectural components for agentic AI systems
- Standardized communication flows between agents, tools, models, and data sources
- A federated control and monitoring plane for identity, authorization, and runtime governance
- Cross-domain trust enforcement across hybrid and multi-cloud environments
- Continuous observability and behavioral monitoring of agent activity
- Rogue agent detection, isolation, and remediation mechanisms
- A vendor-neutral framework to guide supplier implementations and prevent ecosystem fragmentation

The Agentic AI Overlay is not a standards effort — it is an acceleration project. By working collaboratively, the enterprise community defines requirements and vendors align implementations to meet them. The image below illustrates ONUG's relationship with the vendor community and standards organizations.



The Problem We Are Solving

Current enterprise infrastructure is built on rigid, siloed systems and legacy protocols that cannot adapt to the traffic patterns, scale, and security demands of the AI era. AI workloads generate non-deterministic flows, unpredictable processing spikes, and multi-domain dependencies that strain traditional routing and management approaches.

Key challenges include:

- **No secure agent mesh** – Agents operating in different domains often lack a trusted, policy-enforced method for communication and coordination.
- **Data governance gaps** – Sensitive or regulated data may cross boundaries without proper classification, residency checks, or minimization.
- **Inefficient operational intelligence** – Centralized "haystack" data lakes delay incident detection and response.
- **Fragmented vendor ecosystems** – Enterprises struggle to integrate multiple proprietary agent frameworks without common protocols, directories, and policy enforcement.

Without a unifying overlay, enterprises risk deploying siloed AI solutions that cannot scale, interoperate, or meet compliance obligations.

How to Use This Document

This document is intended as both a **blueprint** and a **living reference**:

- **Architecture Readers** – Start with the labeled reference architecture diagram and use the component descriptions to understand each part's function, relationships, and enforcement points.
- **AI Security and Governance Teams** – Security principles and controls that should be embedded throughout the Agentic AI Overlay architecture.
- **Design and Planning Teams** – Map your current infrastructure against the architecture to identify where components already exist and where gaps remain.
- **Use Case Contributors** – Review the existing use cases and add your own ensuring they are specific, realistic, and tied to measurable business or operational outcomes.
- **Vendors & Suppliers** – Use the requirements and flows described here to align products and services with enterprise needs, ensuring interoperability with the overlay model.

This work is broken down into two papers; this paper or Part A, provides functional requirements for agent-to-agent communication, while the second paper, **Part B**, provides agentic AI scaling, deployment, and operationalization strategies. Part A describes a vendor-independent reference architecture aimed to assist in visualizing the various components and their interplay into one figure. **Part B** describes a "how to" or playbook for scaling agentic AI and the issues to be addressed primarily by the vendor community in terms of product features and functions around Agentic control and trust.

Use Cases

The working group is collecting practical, domain-specific scenarios that illustrate where an agentic AI overlay adds value. These use cases will drive requirements and guide the architecture's evolution. Most use cases are focused on agentic ***AI production enablement***. Current and example use cases include:

- **Agentic Wealth Advisory** – In wealth management, AI agents coordinate across market data feeds, CRM systems, and compliance engines to deliver personalized, compliant portfolio recommendations in real time.
- **Automated Incident Detection and Response** – Distributed agents at edge sites detect anomalies in telemetry, apply semantic analysis to classify the event, and coordinate remediation across Network Operations Center/Security Operations Center (NOC/SOC) systems without human delay.
- **Cross-Cloud Data Sovereignty Enforcement** – AI agents apply policy checks before data leaves a jurisdiction, ensuring compliance with GDPR/PCI while still enabling cross-cloud workflows.
- **AI-Enhanced Supply Chain Visibility** – Agents embedded in logistics and warehouse systems share context securely to optimize routing, reduce delays, and improve forecasting accuracy.
- **System of Record Modernization** – Replacing legacy Software-as-a-Service (SaaS) applications with agent-driven AI workflows that can self-adapt to changes in process, data schema, or regulation.

Your contributions are critical. If you operate AI agents today—or plan to—share your scenarios so we can ensure the reference architecture supports them.

Part A

Functional Requirements for Agent-to-Agent Communication

The figure below shows the ONUG Agentic AI Overlay reference architecture. Before we detail the diagram, here are a few notes. This reference architecture is a vendor-neutral/independent view of an agentic AI ecosystem performing work on behalf of a large complex enterprise. The reference architecture shows relationships between various components and trust domains. This architecture is a multi-agent/multi-domain structure, which represents nearly all large enterprise corporations' deployments. Note that a zero trust network infrastructure is mandatory and required to support an Agentic AI overlay.

Also, there are the components of the reference architecture that include agents tools, inference, data, etc., and a specific Agentic Overlay monitoring and control plane to manage its lifecycle and assure security.

Below is the reference architecture where we detail its various components.

The working group conducted a requirements survey of the ONUG community and responses overwhelmingly describe business use cases rooted in operational resilience, regulatory compliance, and controlled automation—rather than customer-facing or transformational AI initiatives.

The top 10 use cases are:

1 **Global Regulatory-Compliant AI Operations**

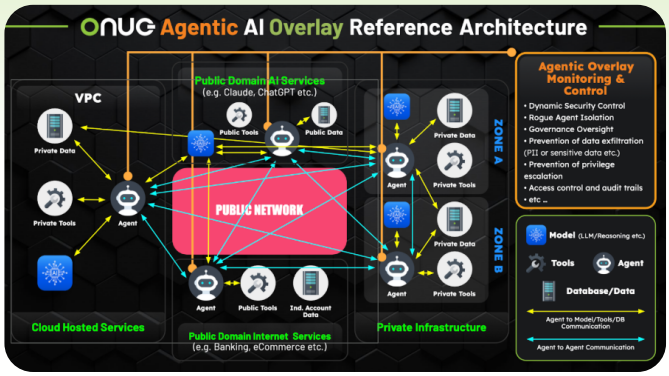
Enterprises deploy agentic AI across multiple countries while enforcing country-specific AI regulations, data residency rules, and governance requirements at runtime. Agents adjust behavior based on geography and regulatory context without requiring separate systems per jurisdiction.

2 **Autonomous Infrastructure Monitoring with Controlled Action**

Agents continuously observe infrastructure, applications, and networks to detect anomalies, risks, and failures while distinguishing between passive observation and actions that generate load or side effects. Automated actions are rate-limited, policy-governed, and recoverable to protect availability.

3 **AI-Driven Incident Response for Mission-Critical Systems**

During outages or security incidents, agents coordinate diagnostics, containment, and recovery steps across cloud and on-prem environments. Built-in failover, logging, and human override ensure incidents are resolved quickly without introducing new risk.



Agentic AI Overlay Reference Architecture – Detailed Description

This section describes the components and flows in the draft reference architecture so that the diagram can be read and fully understood.

Architecture Overview

Goal: a secure, autonomous overlay for AI-era operations that spans on-prem data centers, public clouds, branches/edge, and OT/IoT environments.

The overlay focuses on two core areas:

(1) Operational Intelligence and
(2) Data/Inference Orchestration.

It is vendor-independent, but allows vendor ecosystems to interoperate via common protocols, policy, and identity. A dotted-line "Monitoring & Control" service may be deployed centrally (as a service) or functionally distributed to fit resilience and governance needs.

4 Secure Agent Operations in Highly Regulated Environments

Financial services, healthcare, and other regulated enterprises deploy agentic AI only where full auditability, explainability, and policy enforcement are guaranteed. This enables AI adoption without violating fiduciary duty, client trust, or regulatory obligations.

5 Hybrid and Multi-Cloud AI Operations without Vendor Lock-In

Enterprises operate agentic AI consistently across hybrid and multi-cloud environments using loosely coupled, vendor-neutral architectures. This allows organizations to scale AI adoption while maintaining architectural flexibility and negotiating leverage.

6 Human-Governed Automation for High-Risk Actions

Agents execute routine operational tasks autonomously but require human approval for privileged, irreversible, or safety-critical actions. This preserves accountability and ensures humans remain in control as AI becomes more embedded in operations.

7 Enterprise-Grade AI Observability and Forensics

Organizations gain full visibility into agent behavior through human-readable logs that capture prompts, decisions, tool calls, and outcomes. This enables forensic reconstruction, compliance audits, and continuous improvement of AI systems.

How to Read the Diagram

- **Trust Domains:** The diagram shows multiple operational trust domains (Private DC, Cloud, Edge/OT/Branch). Agents live in each domain.
- **Control vs. Data:** Policies, identity, and directories form the control plane; telemetry and workload/inference flows comprise the data plane.
- **Agentic Overlay Monitoring and Control (AOMC) Service:** The orange "Agentic Overlay Monitoring & Control" indicates a service that enforces policy, provides a kill switch, and simplifies governance. AOMCs may be centralized, distributed, federated or some combination of each.
- **Vendor Enclaves:** "Agent-of-Agents" groupings illustrate vendor ecosystems (e.g., Cisco, Microsoft, OpenAI Frontier, Anthropic Claude Cowork, Gemini Agent, etc.) that must interoperate at some level through the common overlay. Blue lines show protocol messaging flows between agents.

8 Resilient AI Operations with Deterministic Failure Modes

Agentic systems are designed to fail safely through redundancy, execution limits, and operational kill switches. This ensures AI-driven workflows degrade gracefully rather than unpredictably in production environments.

9 Secure Agent Deployment Across Security Zones

Agents are segmented by enterprise security zones and network domains, preventing unrestricted access across systems or environments. Zero-trust enforcement ensures every interaction is authenticated, authorized, and continuously evaluated.

10 Responsible AI Adoption at Enterprise Scale

Enterprises treat agents as long-lived operational assets subject to the same onboarding, testing, approval, and lifecycle management as enterprise applications. This enables controlled, incremental adoption of agentic AI without over-engineering or stalling innovation.

Agentic Flow Types

Yellow Lines – Agent-to-Model/Tools/DB Communication

- Represent *direct calls* from an agent to an external service, such as:
    - An inference service (e.g., a Large Language Model (LLM) endpoint in a public AI service or an internal inference cluster)
    - A specialized tool (e.g., analytics, transformation, policy engine sub-service)
    - A database or structured knowledge store
- These flows are typically **point-to-point** over authenticated, encrypted channels and may cross security or compliance boundaries.
- **Governance/Policy** applies at the initiation of the request—meaning before the agent transmits data, it must be checked against identity, authorization, and compliance rules.

Blue Lines – Agent-to-Agent (A2A) Communication

- Represent peer-to-peer messaging between agents.
- Carried via protocols such as **A2A or yet to be developed protocols**, these flows allow agents to exchange context, pass subtasks, or share intermediate results.
- Blue line flows are **context-rich**—they don't just send raw data but also policy constraints, obligations, and metadata about the interaction.
- Every blue-line message should be monitored by Observability and Telemetry and evaluated by the Policy and Compliance layer.

Key Component Descriptions

Tools

- Discrete capabilities exposed to agents, such as APIs, automation scripts, Robotic Process Automation (RPA) tasks, or AI utilities.
- Tools are often tied to a specific domain (finance, observability, HR, manufacturing) and may require fine-grained authorization.
- Policy and Compliance layer ensures tools are invoked only under allowed contexts and data scopes.

Data

- Structured or unstructured information that agents retrieve, process, or share.
- Can be public, private, or regulated (PII, PHI, PCI).
- Data Protection & Compliance applies classification, residency, encryption, and minimization rules before data moves over any flow (blue or black).

Inference

- Execution of AI/ML reasoning or computation, such as running prompts through an LLM or submitting a query to an AI model.
- Can be internal (private inference cluster) or external (public AI API).

Operational Flows in the Agentic AI Overlay

The lifecycle of an agent within the Agentic AI Overlay begins with **enrollment and attestation**, establishing trust before any autonomy is granted. When an agent is introduced, it is registered in the Agent Directory with its declared capabilities, intended functions, and service-level objectives (SLOs). The overlay issues short-lived, non-human identity credentials—such as certificates, keys, or service principals—scoped precisely to the agent's authorized actions. Before activation, the platform may verify the agent's software bill of materials, supply-chain provenance, and runtime posture to ensure the agent has not been tampered with. Trust is not static; throughout the agent's operation, continuous verification is performed, and credentials can be revoked immediately if behavioral drift, policy violations, or compromise is detected.

Once active, agents communicate with other agents across trust domains using a **zero-trust, policy-mediated interaction model**. An initiating agent discovers a peer through the Agent Directory, evaluates applicable policies, and establishes mutual authentication before any exchange occurs. Agent-to-agent conversations follow defined protocols—such as A2A and MCP (Model Context Protocol), etc.—that govern message structure, delegation, and tool invocation. All interactions are observed by the overlay's observability fabric, which records inputs, outputs, and decision context. For higher-risk interactions, the policy engine may impose additional safeguards, including automated guardrails or human-in-the-loop approval. When data is exchanged, sovereignty, residency, and PII constraints are enforced, while stateless AI routing dynamically optimizes paths across public and private infrastructure without weakening trust boundaries.

To protect the enterprise from unintended or malicious outcomes, the overlay continuously performs **rogue behavior detection and containment**. The observability fabric applies both semantic analysis and statistical anomaly detection to agent outputs, tool usage, and interaction patterns. If an agent violates policy or exhibits suspicious behavior, the Monitoring and Control service can immediately quarantine the agent, sever agent-to-agent communication, or trigger a kill-switch to halt execution. At the same time, incident response workflows are initiated to capture forensic evidence, assess blast radius, and determine remediation steps. Service restoration occurs only under supervised conditions, ensuring compromised agents do not reenter the system without revalidation.

Resilience is built into the system through **coordinated failover and handover mechanisms**. If an agent becomes unavailable or fails to meet its SLOs, a peer agent may assume responsibility by inheriting approved context through MCP-mediated memory handover for example, subject to strict policy limits. The Agent Directory continuously evaluates health and performance signals, redirecting traffic as needed to maintain service continuity. In parallel, inference orchestration reallocates model capacity and execution resources to ensure workloads continue to meet operational objectives. Together, these mechanisms allow the Agentic AI Overlay to sustain reliable, governed autonomy even in the face of failures, scaling demands, or partial outages.

The Agentic AI Overlay Monitoring and Control

AOMC is the cross-cutting supervision plane for the overlay. It ingests telemetry ("emitters") from every agent and, where available, from their tools, and enforces guardrails during agent-to-agent and agent-to-tool/data activity.

Agentic Overlay Monitoring and Control - Reference Architecture Description

Agentic Overlay Monitoring and Control (AOMC) is a **logically centralized, policy-driven control plane** that governs all agent activity across public and private environments without being embedded inside any individual agent. AOMC operates outside the agent execution context, ensuring that trust, safety, and governance are enforced independently of agent self-assertions. Its purpose is to enable **secure, observable, and governable autonomy** at enterprise scale.

At runtime, every interaction—agent-to-agent, agent-to-model, agent-to-tool, and agent-to-data—is subject to **dynamic security control** by AOMC. Real-time policy checks are applied to all message exchanges and action requests, regardless of transport or protocol. These checks integrate with enterprise identity and zero-trust systems, the Policy and Governance Engine, and protocol mediation layers to determine whether a request should be allowed, denied, rate-limited, rewritten, or escalated. AOMC is protocol-aware but protocol-agnostic, inspecting and mediating interactions carried over A2A and MCP as well as other present or emerging protocols such as ACP (Agent Communication Protocol), ANP (Agent Network Protocol), and future agent control or messaging frameworks.

AOMC continuously monitors agent behavior to detect **rogue or out-of-policy activity**. Using both semantic analysis and behavioral telemetry, the observability fabric identifies deviations from an agent's declared capabilities, unusual access patterns, or anomalous decision flows. When violations are detected, AOMC coordinates with the Agent Directory, identity systems, and the overlay message bus to quarantine the agent, revoke credentials, or terminate sessions. Because AOMC operates externally, agents cannot bypass enforcement by self-attesting compliance.

Beyond transport-level enforcement, AOMC provides **governance oversight at the semantic layer**. Acting as a semantic firewall, it evaluates intent and content—not just API calls—blocking disallowed prompts, preventing unsafe or non-compliant outputs, and enforcing content policies across trust domains. These controls integrate with compliance frameworks (e.g., PCI, GDPR, HIPAA, NYDFS) to ensure that regulatory, residency, and data-handling rules are applied consistently across clouds, data centers, and edge environments.

AOMC plays a critical role in **preventing data exfiltration** in agentic systems. As messages and tool outputs traverse the overlay, AOMC applies classification, tagging, and data loss prevention controls to enforce strict boundary rules between security zones, VPCs, tenants, and public services. These controls are informed by the observability and telemetry fabric, enabling enforcement decisions that are both real-time and context-aware.

To prevent **privilege escalation**, AOMC validates that each requested action aligns with the agent's declared role, scope, and authorization. Attempts to acquire broader privileges mid-workflow are rejected or subjected to step-up authentication or human approval, enforced through integration with identity and zero-trust systems. This ensures that autonomy does not silently expand beyond enterprise intent.

All enforcement decisions are driven by **centralized access control**, combining Attribute-Based Access Control (ABAC) and Role-Based Access Control (RBAC) models to determine who an agent is, what tools and data it may access, and in which zone or tenant actions are permitted. These decisions are enforced consistently at agent-to-agent gateways, and at tool and data connectors, independent of the underlying protocol stack.

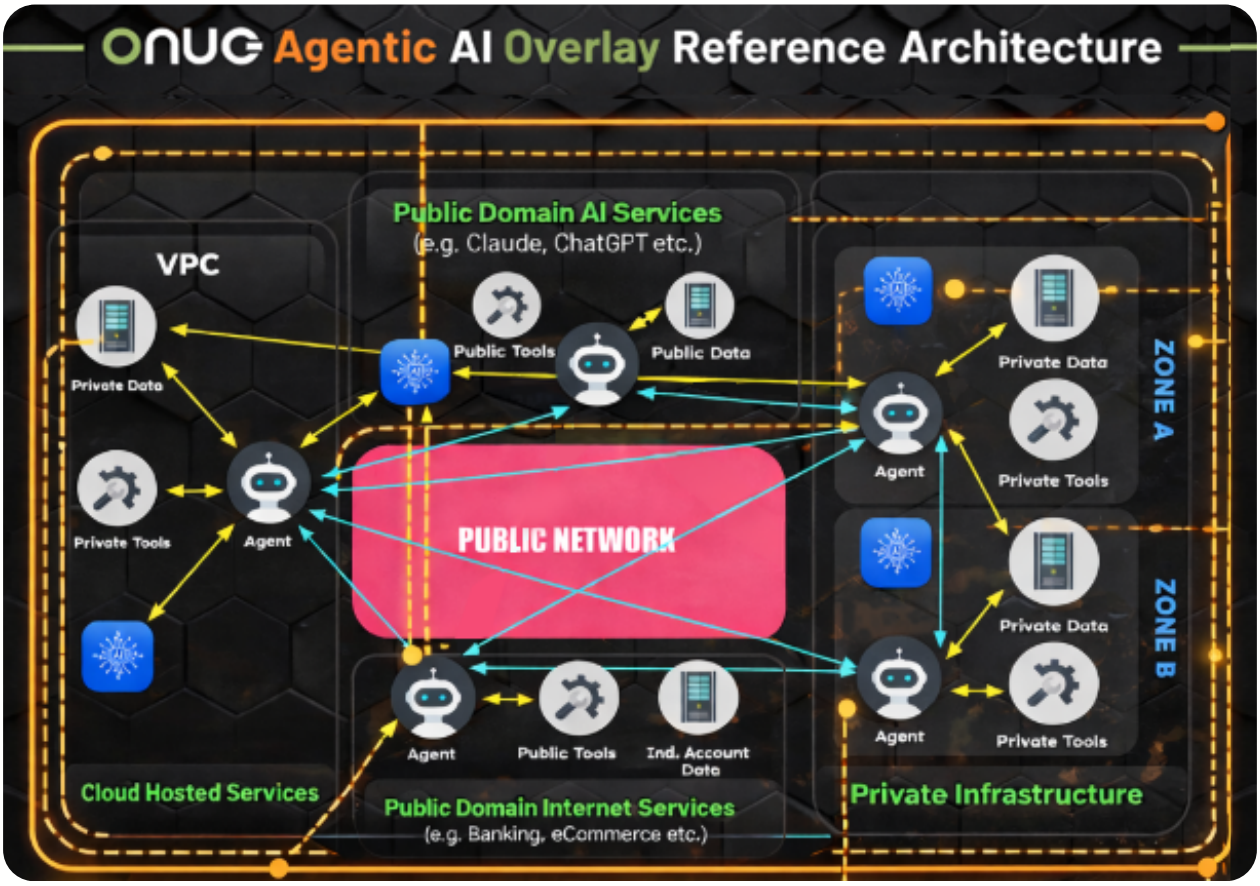
AOMC maintains **tamper-evident audit trails** capturing prompts, responses, tool invocations, policy decisions, and outcomes across private data centers, VPCs, and public services. This end-to-end traceability is foundational for regulated industries, post-incident forensics, and continuous governance improvement.

The **Agent Directory and registries** serve as authoritative sources of truth for agent identity, capabilities, allowed interactions, and endpoints. Agents consult the directory at connection time, while AOMC relies on it for policy evaluation, enforcement, quarantine, and lifecycle management. In practice, the directory operates alongside agent control-plane components, forming the backbone of discovery and trust establishment.

In aggregate, AOMC enables enterprises to scale agentic AI safely by enforcing **continuous verification, explicit trust, and policy-driven autonomy**, regardless of where agents run or which protocols they use.

The above AOMC description can be dissected into a set of call-outs or services, which the AOMC provides. The following ten call-outs are mapped into the reference architecture.

Use the following numbered call-outs are directly mapped on the Agentic AI Overlay Reference Architecture or as a legend.



1. **AOMC Control Plane (orange boundary/lines)**

   Logical control fabric spanning all trust domains; enforces policy, security, and governance independently of agents.

2. **Dynamic Security Control**

   Real-time allow/deny/rate-limit/rewrite decisions on all agent interactions (A2A, agent-to-tool, agent-to-data), integrated with Identity, Zero Trust, and Policy Engine.

3. **Protocol Mediation Layer**

   Inspection and enforcement across A2A, MCP, ACP, ANP, and future agent protocols without replacing them.

4. **Agent Directory/Registry**

   System of record for agent identity, capabilities, allowed tools/data, SLOs, and endpoints; used for discovery, policy lookup, and quarantine.

5. **Rogue Agent Detection and Isolation**

   External behavioral monitoring; quarantine, credential revocation, or session termination on drift or compromise.

6. **Semantic Governance Firewall**

   Intent- and content-aware enforcement blocking disallowed prompts, unsafe outputs, and policy violations.

7. **Data Exfiltration Prevention**

   Classification, tagging, and boundary enforcement across zones, VPCs, tenants, and public services.

8. **Privilege Escalation Protection**

   Continuous validation of agent actions against declared scope; step-up auth or rejection for elevated requests.

9. **Centralized Access Control**

   ABAC/RBAC enforcement at A2A gateways and tool/data connectors, independent of network location.

10. **Audit and Forensics Fabric**

    Tamper-evident logging of prompts, actions, decisions, and outcomes across all environments.

AOMC "wired into" the ONUG Agentic AI Reference Architecture

Across the architecture, **agents operate everywhere the enterprise runs compute** - in private data centers, cloud VPCs, and at the branch or edge. Regardless of location, agents do not act autonomously in isolation; every agent emits continuous telemetry to the AOMC plane and subjects its actions to real-time policy evaluation. Before messages traverse the **blue paths** (agent-to-agent communication) or **yellow paths** (agent-to-tool or agent-to-data access), AOMC evaluates identity, intent, and risk, ensuring that interactions remain within declared capabilities and approved trust boundaries.

On **yellow-line flows**, where agents invoke tools or access data stores, AOMC acts as the enforcement point for access control, data loss prevention, and auditability. Calls to databases, SaaS APIs, and internal services—especially those crossing zones, tenants, or regulatory domains—are inspected and logged, with payloads classified and constrained according to residency and sensitivity rules. In parallel, **blue-line flows** pass through A2A protocol gateways instrumented by AOMC, where identity assertions, governance checks, and protocol-level controls are applied. If policies change mid-conversation or risk conditions spike, AOMC can reshape, throttle, or terminate agent sessions without requiring agent cooperation.

Underpinning all flows, **Identity and Zero-Trust services feed AOMC with authentication and authorization context**, enabling continuous privilege validation and step-up authentication when warranted. The Policy Engine and Compliance Layer define what is allowed; AOMC is the execution arm that evaluates those decisions transaction by transaction and records tamper-evident evidence for PCI, GDPR, HIPAA, NYDFS, and similar regimes. At the same time, AOMC both consumes and produces signals for the **Observability and Telemetry Fabric**, correlating metrics, traces, semantic events, and audit logs to support detection, governance, and post-incident forensics. Even external or vendor-hosted "agent-of-agents" enclaves are treated as managed clusters; AOMC enforces policy at the enclave boundary and can quarantine the entire enclave if its aggregate behavior deviates from enterprise intent.

Note that the working group explicitly discussed leaving the door open for a **service-based AOMC**. The group felt that it could be easier for incident response (global shutdown) and simpler vendor integration—while still supporting a fully distributed model if preferred.

Component Descriptions, Assumed within the Architecture But Not Diagrammed Above

The following architectural components are not explicitly detailed in the ONUG Agentic AI Overlay Reference Architecture, but play important roles in the safe and secure operation of a large enterprise grade agentic AI system.

Agentic Overlay Monitoring and Control: An optional centralized service (represented as an orange box) that provides global governance: attestation, posture checks, behavioral policy enforcement, and emergency shutdown/quarantine for agents. It simplifies operations for highly regulated environments and incident response, while remaining compatible with a fully distributed architecture.

Identity and Zero Trust for Agents: A rethought identity model for non-human principals (agents). Provides authentication and authorization for every agent-to-agent/API interaction across domains. Implements short-lived credentials, certificates/keys or service principals, scope-limited permissions, and continuous verification to enforce least privilege.

Policy Engine and Governance: Encodes business logic, risk thresholds, regulatory constraints (e.g., PCI, NYDFS, GDPR), and data sovereignty rules. Policies determine which agents may talk, what data may flow, which tools may be invoked, and acceptable actions. Policies can adapt to changing conditions (risk, load, compliance events).

Agent Directory (Capabilities Registry): A dynamic registry advertising each agent's capabilities, trust level, health, SLOs, and location. Enables discovery, targeting, and load-aware selection of agents. Supports failover, blue/green updates, and graceful degradation when agents leave/join the mesh.

A2A Protocols (and similar frameworks): Secure, authenticated, and auditable messaging between agents across domains. Carries intents, tool calls, telemetry summaries, and coordination signals. Supports vendor ecosystems and cross-domain conversations with consistent enforcement of identity and policy.

MCP (and similar frameworks): A unifying context exchange and tool invocation layer used by agents, applications, and orchestration to preserve task context end-to-end. Facilitates federation across vendor ecosystems and simplifies plug-in models for tools, memory, and data sources.

Observability and Telemetry Fabric: Distributed collection and normalization of logs, metrics, traces, and events. Feeds agents and the policy engine with real-time situational awareness for anomaly detection, reasoning, validation of actions, and post-action audit. Enables semantic monitoring of agent outputs to detect brand-/safety-impacting behavior.

Inference and Data Orchestration: Matches AI inference capacity to demand across business units and domains, avoiding hot spots (100% utilization) and stranded capacity (0% utilization). Enforces data residency and access constraints while optimizing placement for latency and cost. Coordinates model selection, versioning, and rollout gates.

Stateless AI Router (AI-assisted Pathing): A routing/forwarding function that uses real-time signals (latency, congestion, cost, policy posture) to make rapid path decisions without heavy per-flow state. Targets millisecond-level reconfiguration for bursts, failures, and workload shifts; complements traditional routing where needed.

Vendor "Agent-of-Agents" Enclaves: Logical clusters of vendor-specific agents (e.g., cloud or networking stacks) represented as circles within domains. The overlay ensures these enclaves interoperate through common identity, policy, and A2A/MCP, etc., without lock-in.

Data Protection & Compliance Boundaries: Controls that constrain data movement and tool invocation (PII, IP, regulated records). Includes masking, redaction, tokenization, differential access per role/agent, and audit trails that support legal/compliance review.