VERSATILITY

# Securing the Agentic Era: Versa's Innovations in SASE, Posture, and Security for AI

**Apurva Mehta**

Co-founder & CTO

# Innovating across all fronts

*Our goals for 2026 & beyond*

**Thank you for your continued trust, support, and guidance. We continue to improve and extend our lead with help from your suggestions**

- Enhancing different SASE services. Integration with more EDRs, UEMs, VTMs, and enterprise-specific threat feeds
- Extending our SD-WAN Leadership
- Lots of enhancements to performance and scale of Director, and Analytics. Further unification of SD-WAN and SSE capabilities in Concerto
- Advancements in Platform, SD-LAN, and IoT
- Posture Management: DSPM, and AI-SPM
- SASE for AI

VERSATILITY

# Enhancements: **Secure Services Edge**

- Support for more CASB applications
- Support for additional SaaS applications and IaaS for API based Data Protection
- Connectors to EDRs like CrowdStrike, MSFT Defender, Sentinel One; UEMs like MSFT Intune and Ivanti, and VTMs like Qualys, Tenable, and MSFT
- Enhancements to Clientless Access
- Support for Moderated flows for emails
- SASE Client and gateway support for endpoints where multiple users are active at the same time
- IPv6 support for SASE Clients
- Advanced Threat Protection, RBI, and UEBA support for air-gapped environments

# Enhancements: **AI for SASE**

- AI/ML-based DLP on the Versa SASE Gateways for source code, PII, image/document classification, PDF, MS Office extraction, and redaction
- Identifying security incidents and correlating security incidents to MITRE TTPs
- Enhancements with the identification of malware in JavaScript
  - Newer models
  - De-obfuscation and Sandboxing of JavaScript
- Model for ELF (Executable and Linkable Format)
- Explainable AI component built into VANI and UEBA
- MCP Server for Director and Concerto
- Zero Trust Verbo
- Use of GPUs and inference cards

# Enhancements: **SD-WAN**

- PQC support for traffic from branch to controller and branch-to-branch traffic
- Distributed software shaper which can scale up to 40+ Gbps
- Supported for metered circuits
- Support for DNS over HTTPS (DoH) and DNS over TLS (DoT)
- Support for Leo circuits
- Support for Tenant-specific Ingress Shaper

# Enhancements: **SASE Ecosystem**

- 23.1.1 replaces legacy Versa Director HA with a microservices-based architecture orchestrated by a Docker Swarm Multi-Node Cluster. This leads to increased resiliency and scale, as well as lower latency to commit configuration changes
- 13.1.1 Concerto provides a well-guided user interface for SD-WAN. It also allows security profiles to be shared between Secure Services Edge and SD-WAN
- Versa Data Lake Enhancements:
  - Integrates the Versa Advanced Logging Service with Lakehouse solutions for long-term analytics and reporting
  - Supports seamless streaming of archived data to third-party data lakes and platform

# Versa Director 23.1 Architecture

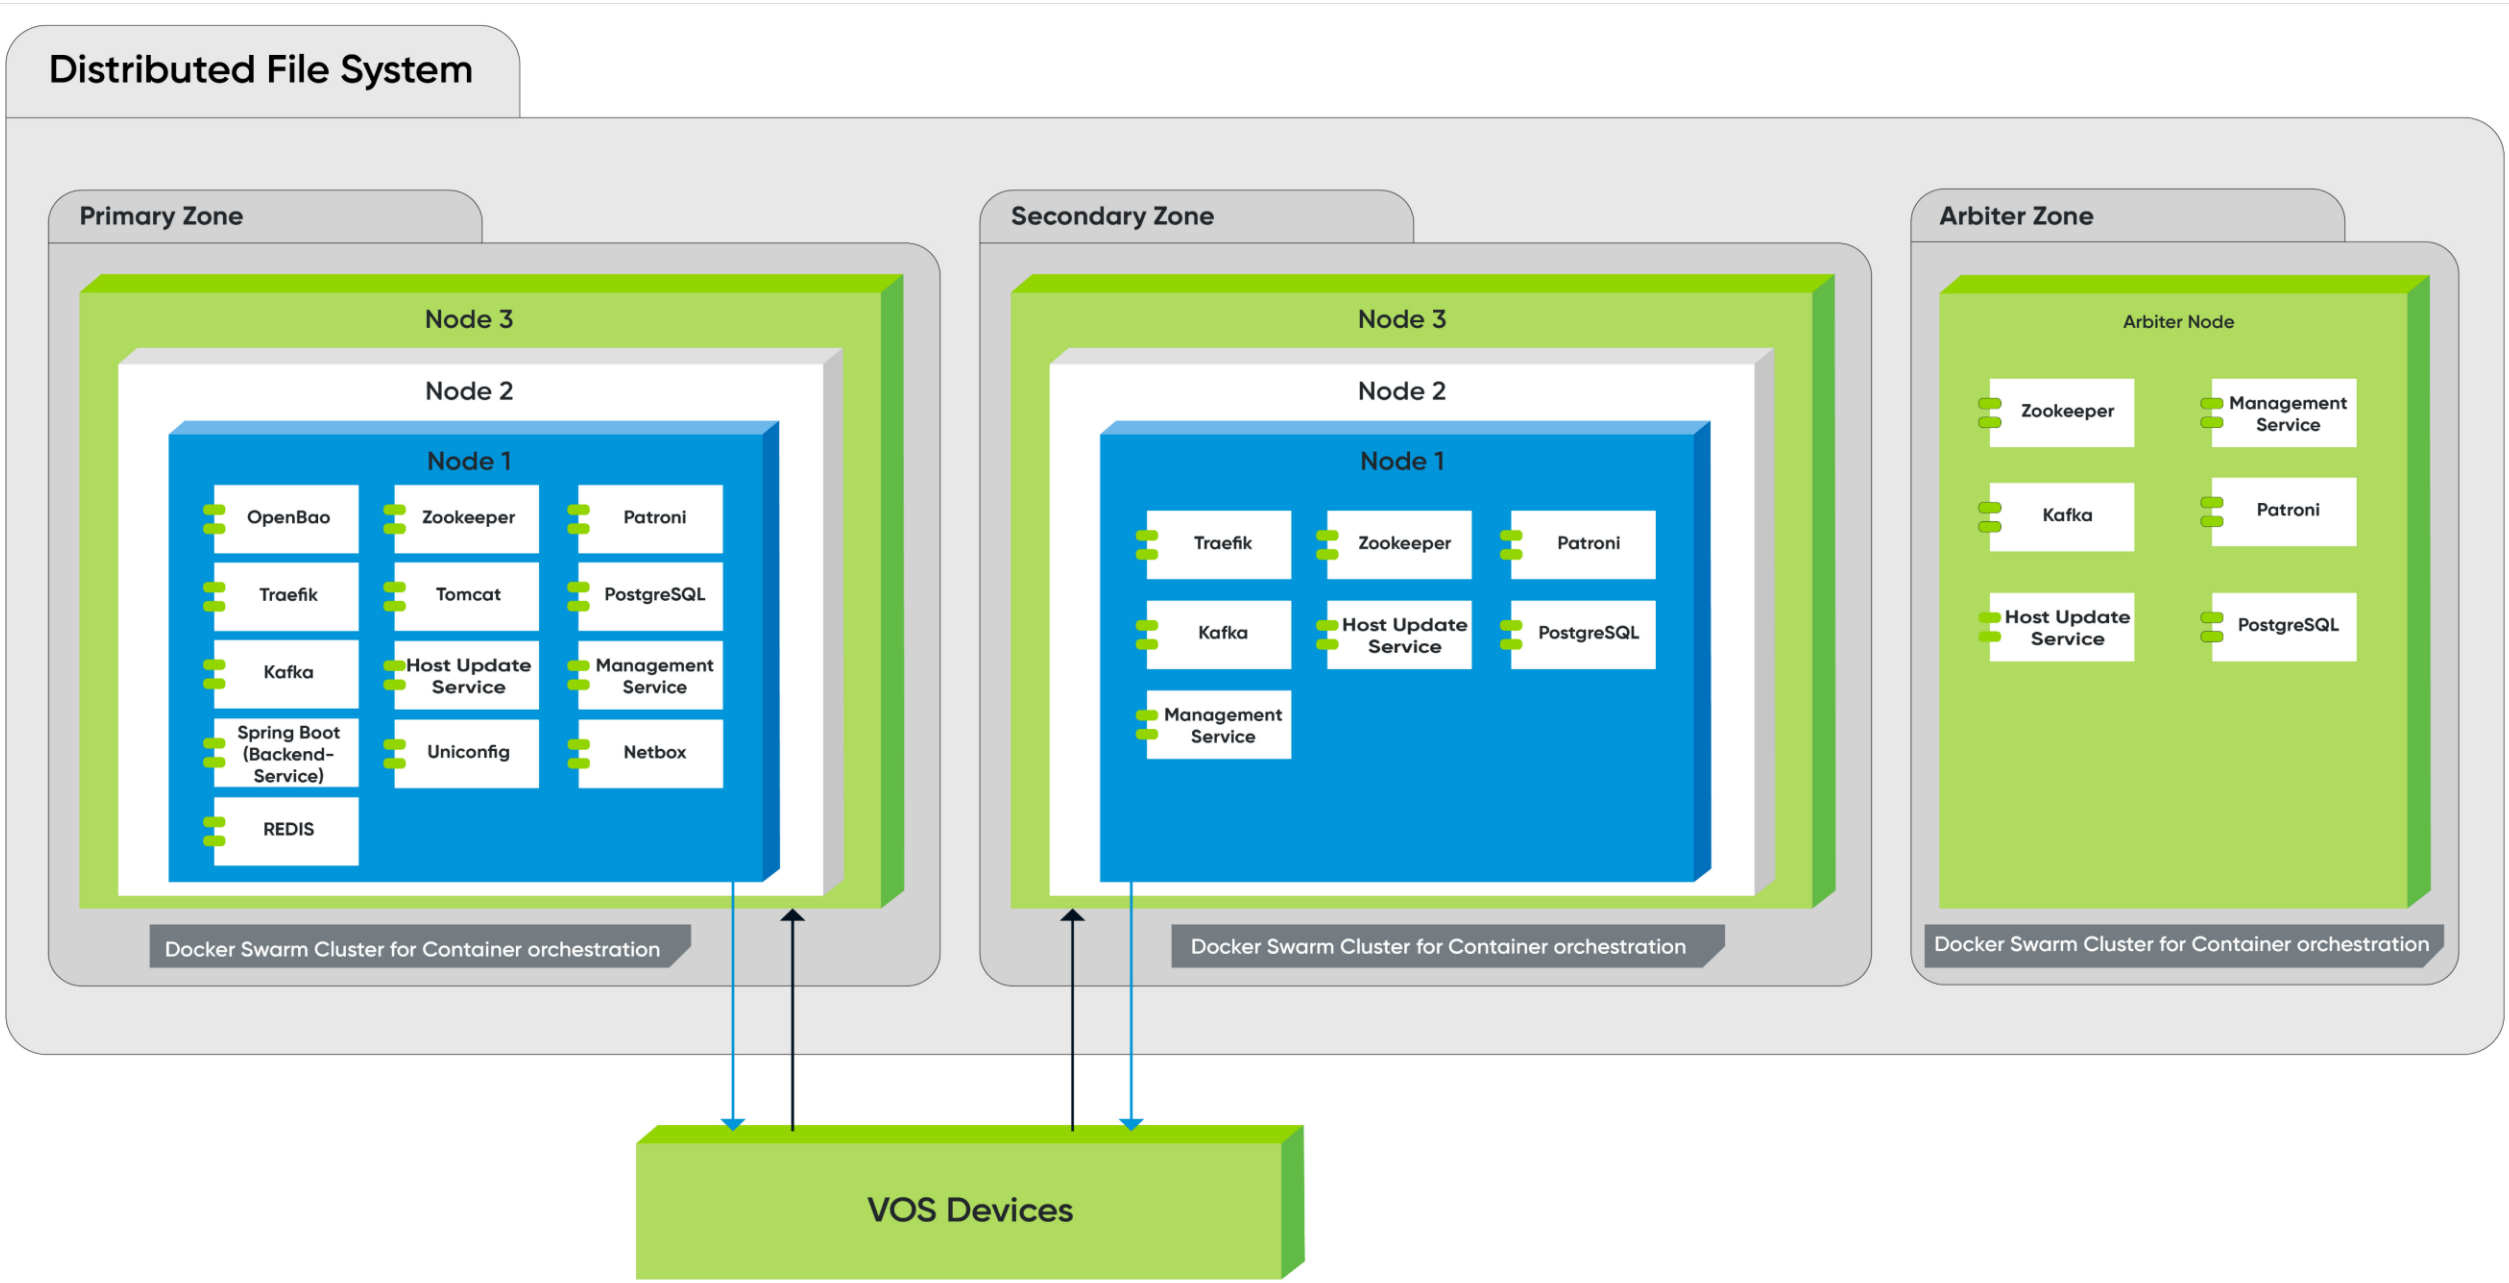

# Concerto for Unified SASE: *SSE, SD-WAN, Posture Management, and SASE on SIM*

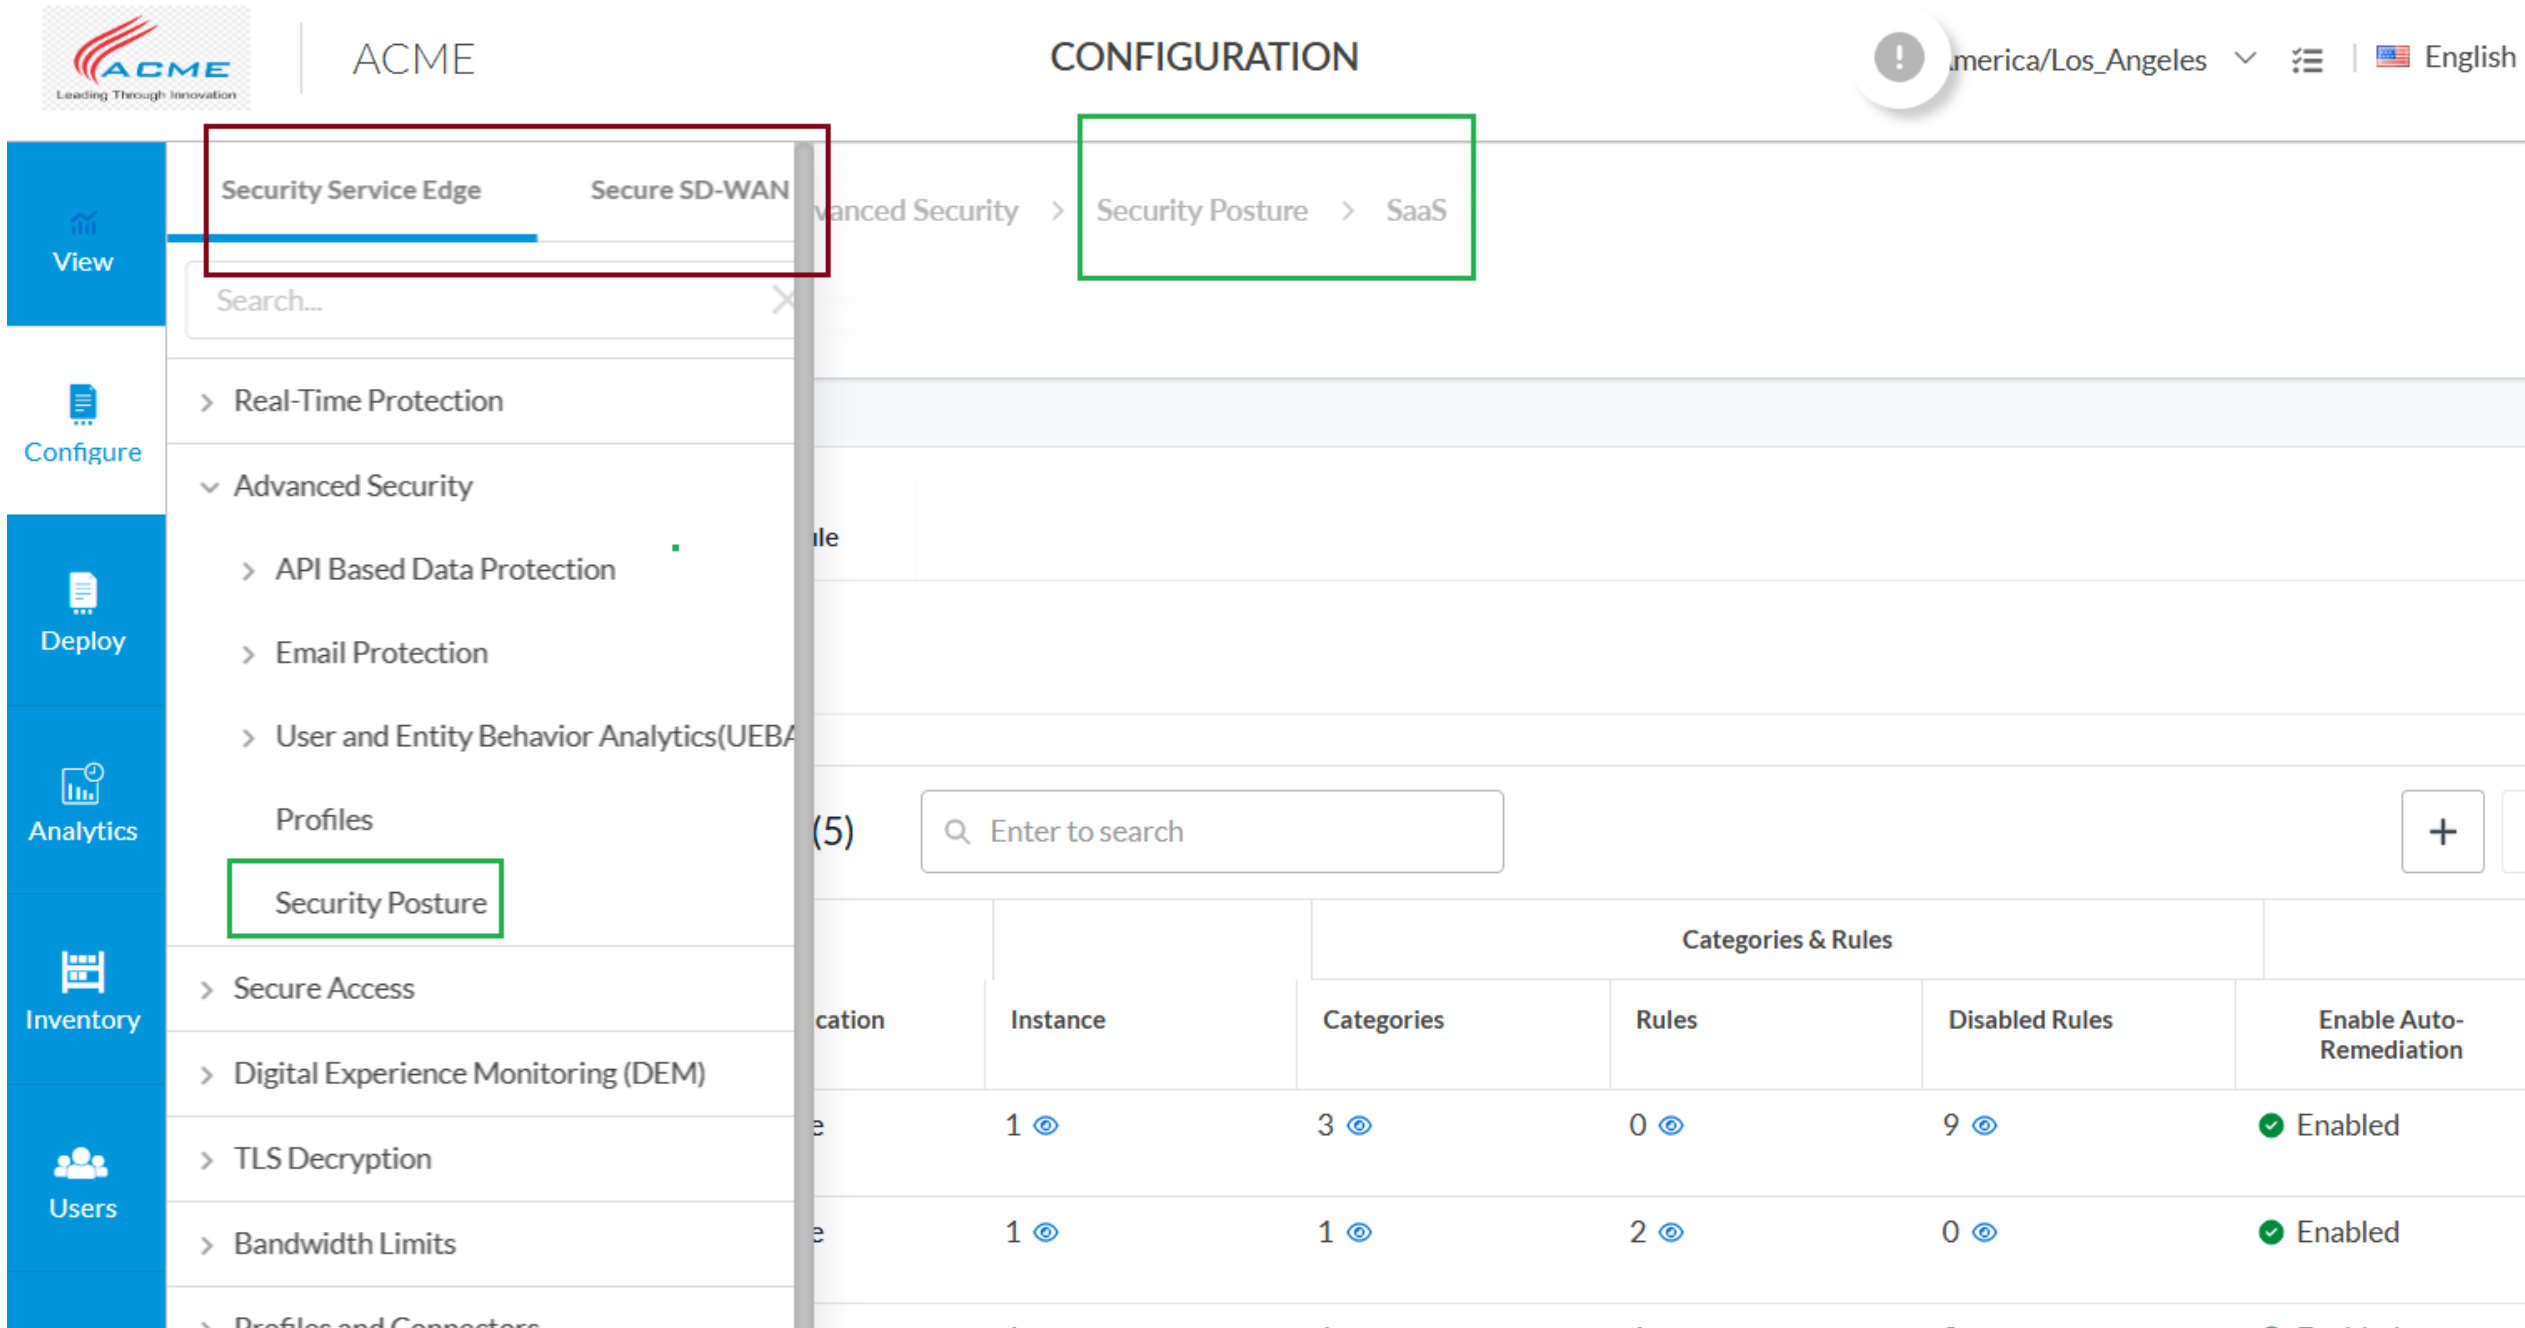

# Concerto: *Well Guided SD-WAN Configuration: Part-1*



## Edit Main Template: TestS2S_MT

1. Deployment Tier & High Availability
2. Network Interfaces
3. Topologies & Routing Protocols
4. Network Services
5. QoS, Traffic Steering & Traffic Monitoring
6. Authentication
7. Security
8. Servers & Settings

Deployment Tier | High Availability

### Deployment Tier

Select the deployment type (SDWAN or NGFW) and the corresponding solution tier based on the license you plan to apply to the appliance. The selected tier may restrict configuration of certain features.

Scope: Single Tenant

Solution Tier: (•) SDWAN ( ) NGFW

**Prime SD-WAN**

Core features of Prime SD-WAN include:

- Dynamic Overlays
- Rich Topology Options
- Application Identification
- Application Policy Based Traffic Steering

**Prime Secure SD-WAN**

Adds on top of Prime Secure SD-WAN with Next Generation Firewall includes:

- Application Firewall
- URL Feeds (Classification and Filtering)
- IP Address Feeds (Classification

**Premier Secure SD-WAN**

Adds on top of Premier Secure SD-WAN includes:

Application/User Optimization

- Forward Error Correction (FEC)
- Packet Shaping
- Packet Striping
- Packet Cloning

**Premier Elite Secure SD-WAN**

Adds on top of Premier Elite Secure SD-WAN includes:

Unified Threat Protection features such as,

- Known File has Feeds and Filtering
- IDS/IPS
- Network-based Anti-Virus and

# Concerto: *Well Guided SD-WAN Configuration: Part-2*

Configure > Secure SD-WAN > Edit Main Template: TestS2S_MT

1 Deployment Tier & High Availability | 2 Network Interfaces | 3 Topologies & Routing Protocols | 4 Network Services | 5 QoS, Traffic Steering & Traffic Monitoring | 6 Authentication | 7 Security | 8 Servers & Settings

Versa CSG770

NIC Port

None

Legend: Management WAN LAN WAN-LAN L2 Cross PPPoE AE T1E1 DSL

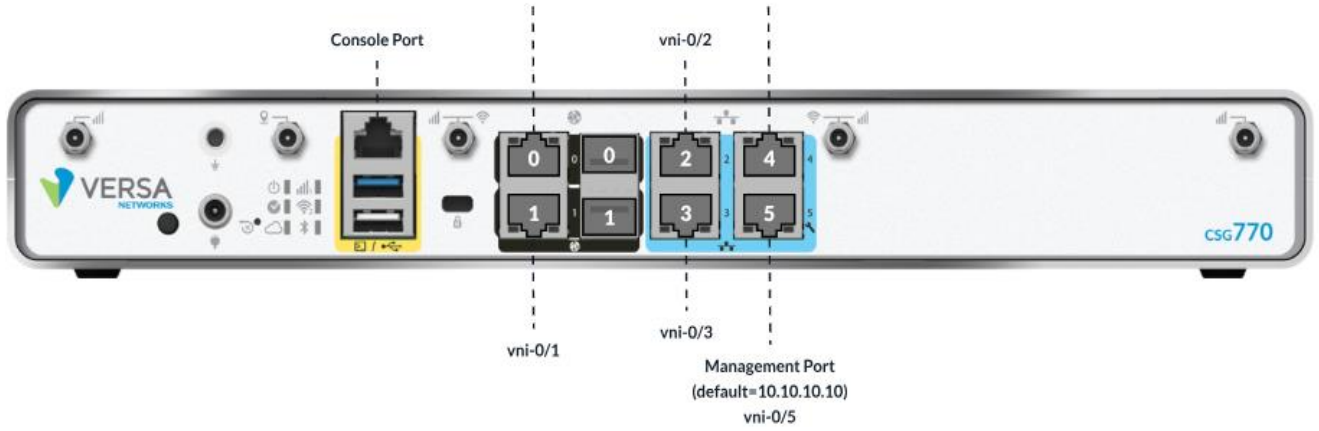


## Primary Device Interfaces

All Interfaces | WAN | LAN | Site-to-Site Tunnel | Loopback | Paired Virtual Tunnel

All Interfaces (1)

| Name | Version | Category | Sub-Category | Interface | Connection Name/Routing Instance | Variables | Status | Last Modified |
|---|---|---|---|---|---|---|---|---|
| TestS2S | 1 | Site-to-Site Tunnels | IPsec | | ACME-Enterprise | 6 | Enabled | 1/15/2026, 11:48: KhyatiV |

# Enhancements: **Platform**

- VOS support for 96-core appliances having tiled architecture
- Support for clustering of appliances for higher secure SD-WAN and SSE service scale
- Support for new appliances like CSM-64, CSG 450, CSX 2200, CSX 2300, CSX 1200, and CSX 1100
- Capability to do life cycle management of containers on VOS
- OT protocols such as Modbus, DNP3, ENIP (EtherNet/IP), and OPC UA Binary. All are over TCP/UDP
- Versa would also be releasing an WiFi7 Access Point in 2H 2026

# Life cycle management and service chaining of containers on Versa uCPE

Private or Virtualized DC

PoS Server

Versa Hub

SD-WAN Underlay

PoP or Data Center

Kubernetes Master

Versa Director

Versa Analytics

Versa Controller

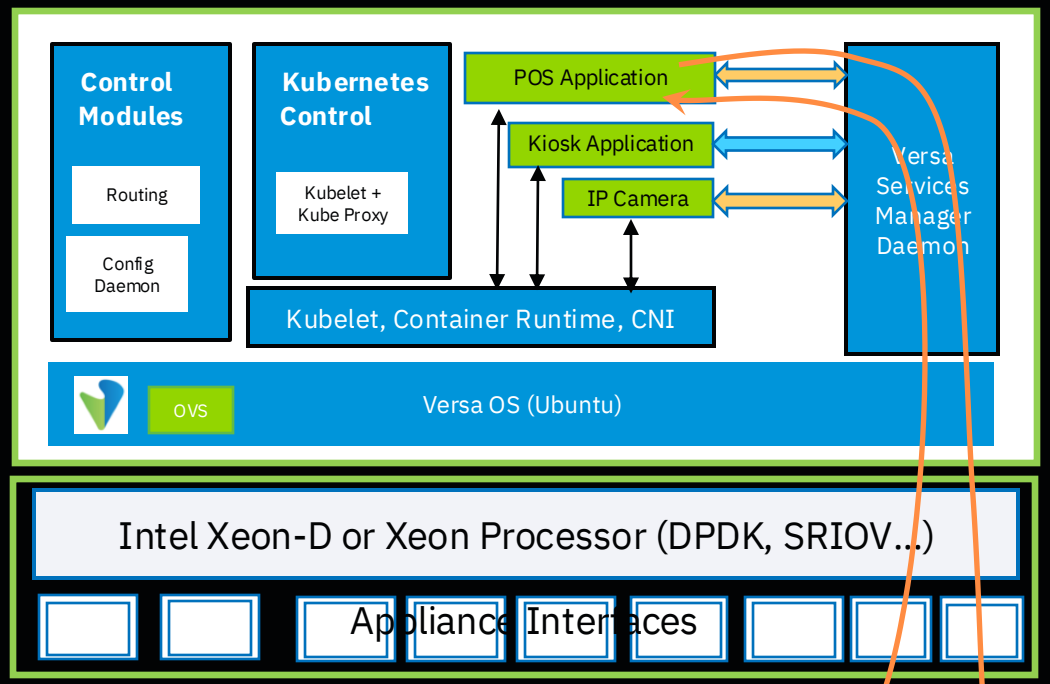


PoS Terminal

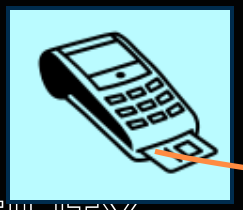

VERSATILITY

# Posture management: DSPM, AI-SPM

VERSATILITY

# Data Security Posture Management

Discover, classify, and continuously govern sensitive data across every cloud, SaaS, and on-prem data store

## Automated Data Discovery

Continuously scan GCP, AWS, Azure, OCI, Snowflake, MongoDB and SaaS stores to surface all sensitive data — including shadow stores

## Sensitive Data Classification

Auto-classify PII, PHI, PCI, financial and proprietary data with regulatory tagging for GDPR, HIPAA and PCI-DSS

## DSPM Risk Score

A single 0–100 posture score tracks infrastructure, data, user access and interaction risk

## Access Governance & Monitoring

Detect risky queries, over-privileged access, and abnormal user interactions with sensitive data in real time

## Data never leaves your cloud

The DSPM Agent runs analysis entirely inside your cloud environment. Sensitive data is scanned and classified in-place — nothing is extracted or transmitted externally — so you stay in full control of your data at all times

VERSATILITY

# Data Security Posture Management

Discover, classify, and continuously govern sensitive data across every cloud, SaaS, and on-prem data store

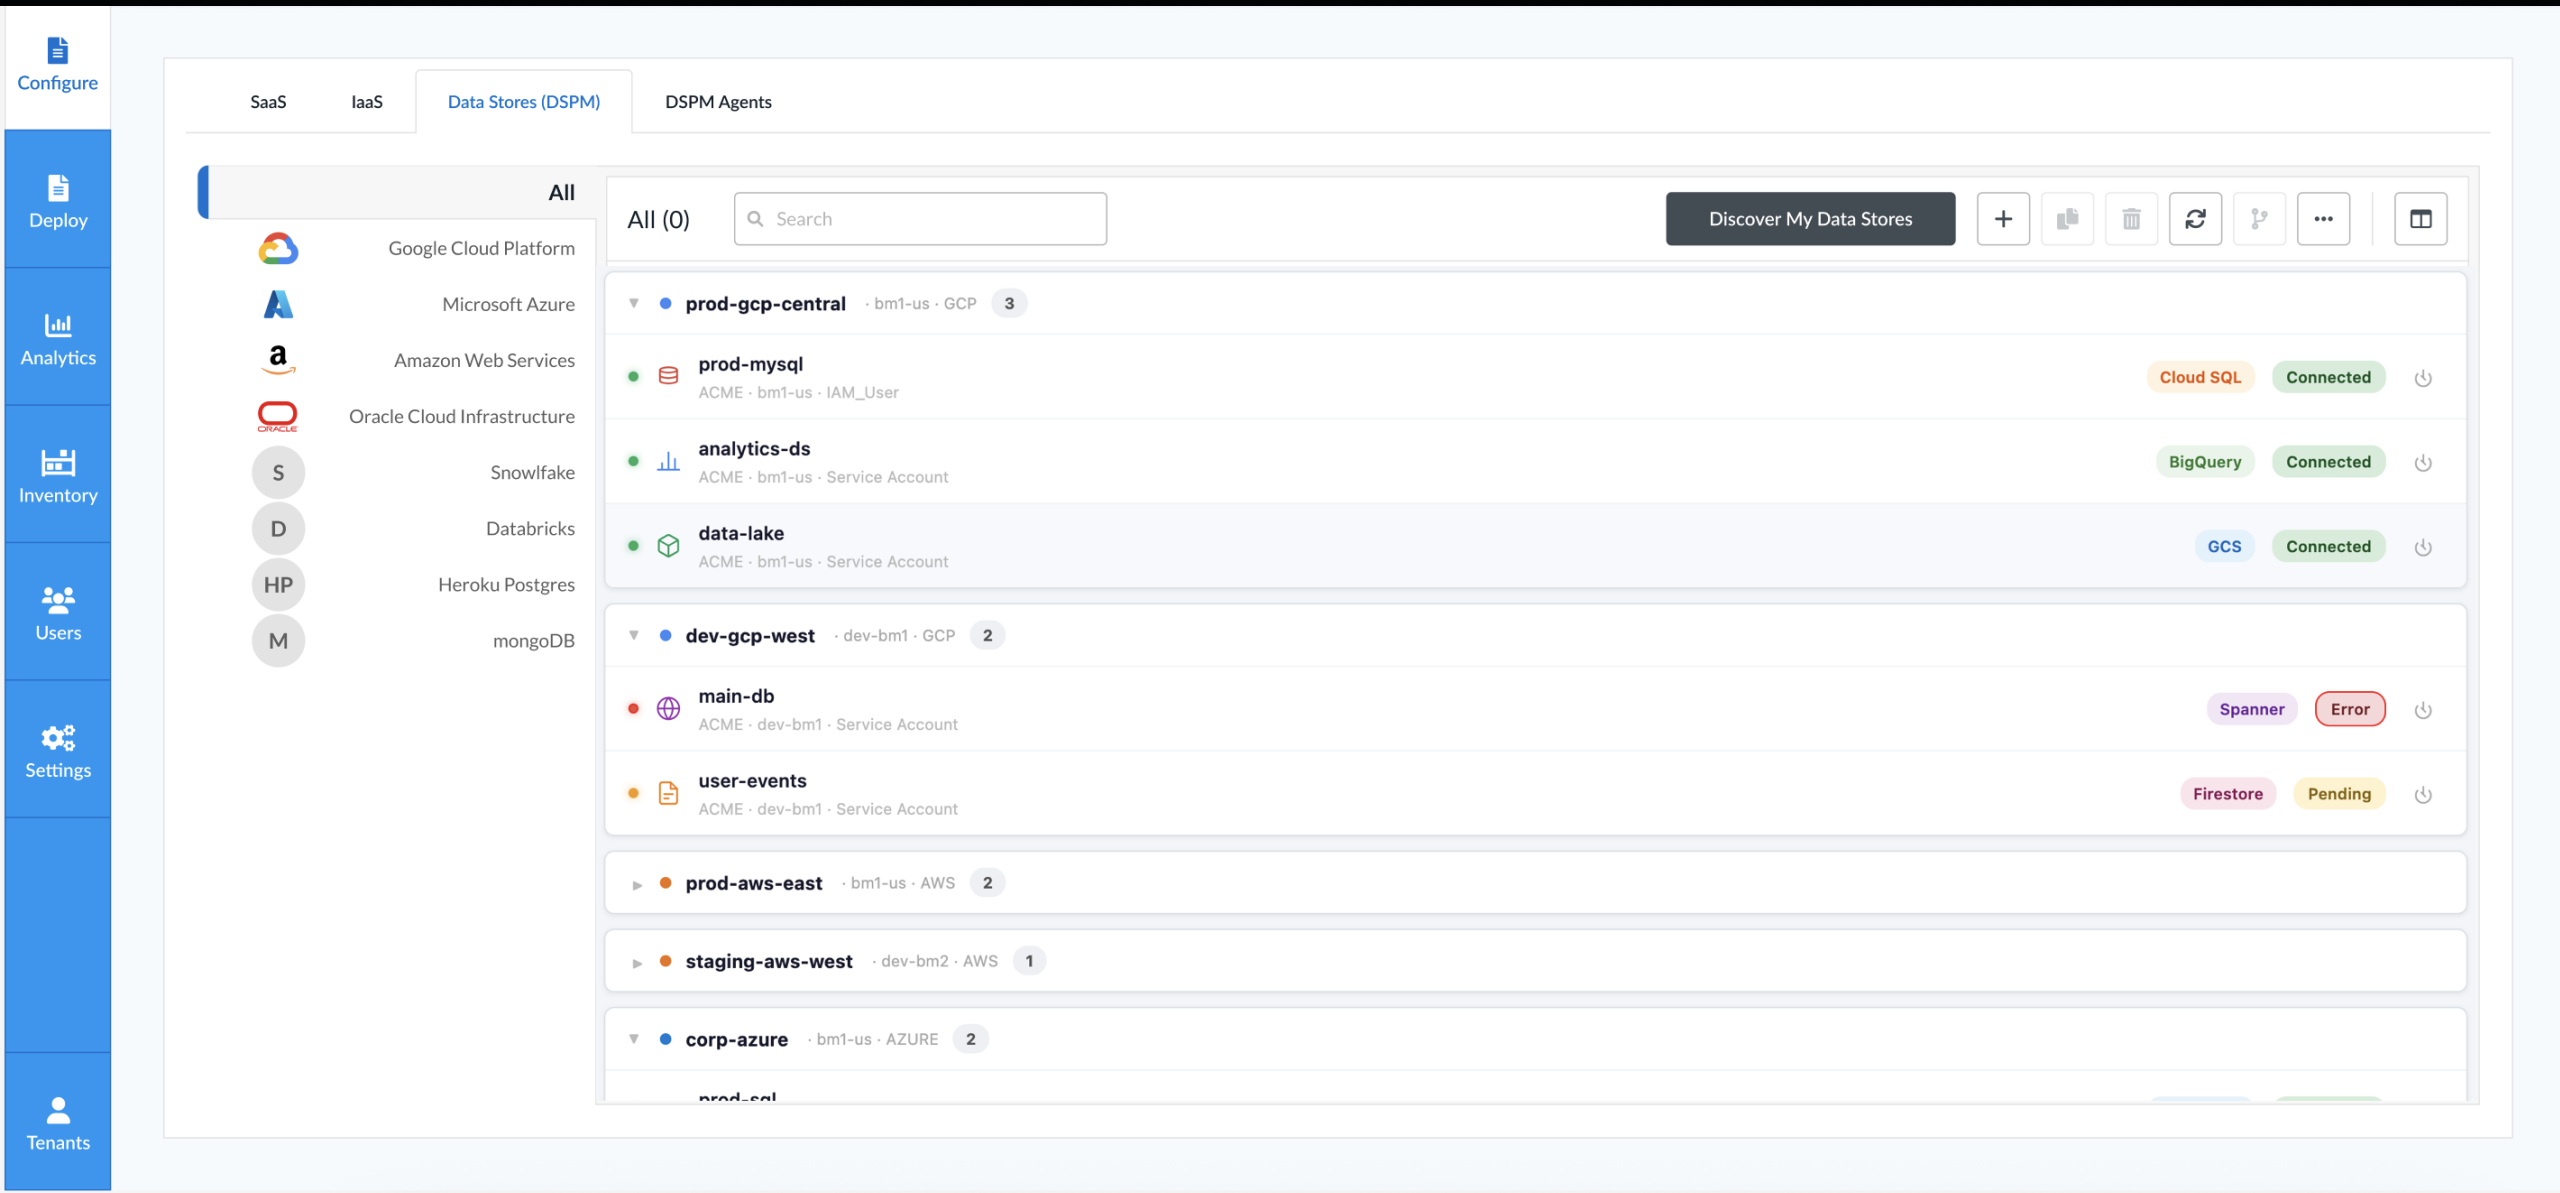

# Data Security Posture Management - Reports

## Unified risk score

The DSPM Score rolls infrastructure risk, sensitive data violations, and risky access into a single metric — giving security teams and executives one number to track posture over time

## Prioritized remediation queue

Findings are ranked Critical → Medium so teams act on the highest-impact issues first — ungoverned data stores, publicly exposed PCI data, and privacy-violating queries — without sifting through noise

## Integrated inside Versa Unified SASE

DSPM runs natively alongside CSPM and SSPM, correlating access risk, cloud posture risk, and data security risk in one platform — eliminating the blind spots that come from disconnected point tools

# Data Security Posture Management - Reports

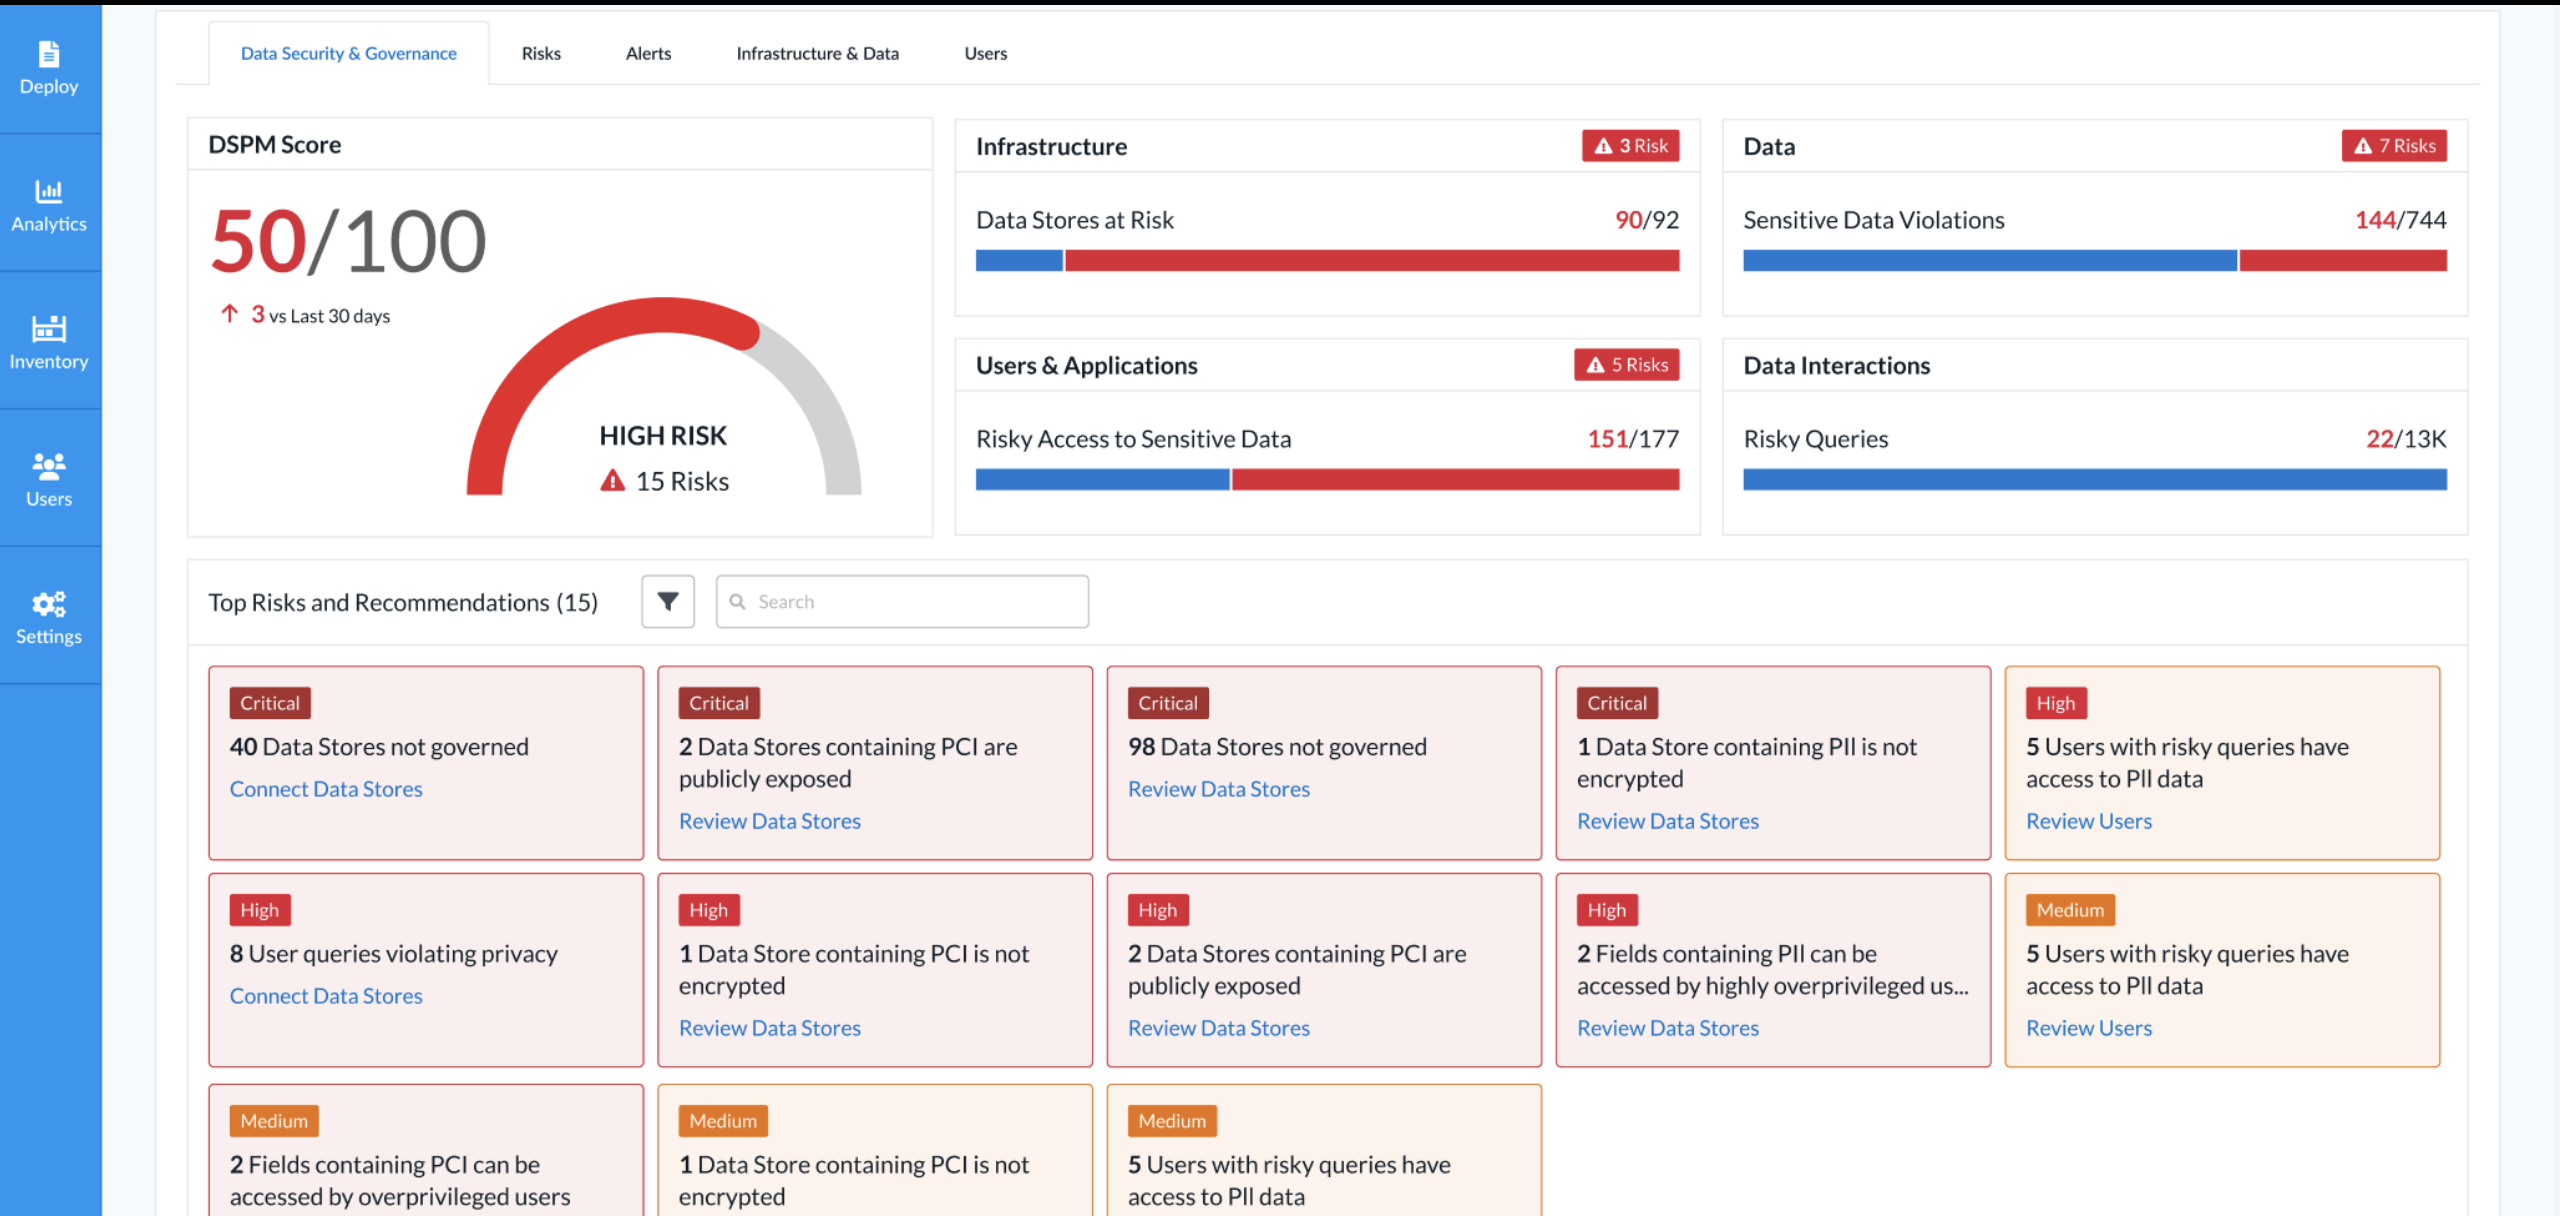

# SASE for AI

VERSATILITY

# Secure access for AI

- Versa SSE Gateways can prevent inappropriate upload and download of sensitive data using GenAI applications
- Audit log ingestion, usage telemetry, configuration posture monitoring by API integration with ChatGPT Enterprise, Claude Team/Enterprise, MSFT 365 Copilot, and others
- Agent governance in Slack and Teams via Versa's SSPM connectors
- Discovery, configuration assessment, and risk scanning for embedded AI models in SaaS applications (Salesforce Einstein – Prediction builder, Einstein GPT) via the SSPM connector framework, with AI-feature-specific posture rules
- Discovery, configuration assessment, and risk scanning for embedded AI agents in SaaS applications (Salesforce Agentforce, ServiceNow Now Assists agents) via SSPM connectors with agent-specific posture rules

# Versa SASE for AI

## Model Gateway

Single proxy for all LLM API traffic. Enforces auth, rate limits, token budgets, and provider failover. The agent never holds an API key

## MCP Gateway

Secures tool calls between AI agents and enterprise systems (databases, ERP, code repos, file systems). Enforces identity, access policy, and audit on every tool invocation

## LLM WAF

Dual-scanner prompt injection defense enforced inline via Model GW. Augmented by Versa DLP to prevent sensitive data leakage in prompts and responses

## A2A Gateway

Governs agent-to-agent communication. AuthN/AuthZ, rate limiting, and full audit trail for orchestrator-to-subagent traffic — east-west AI traffic never leaves ungoverned

# SASE for AI components on a Versa uCPE

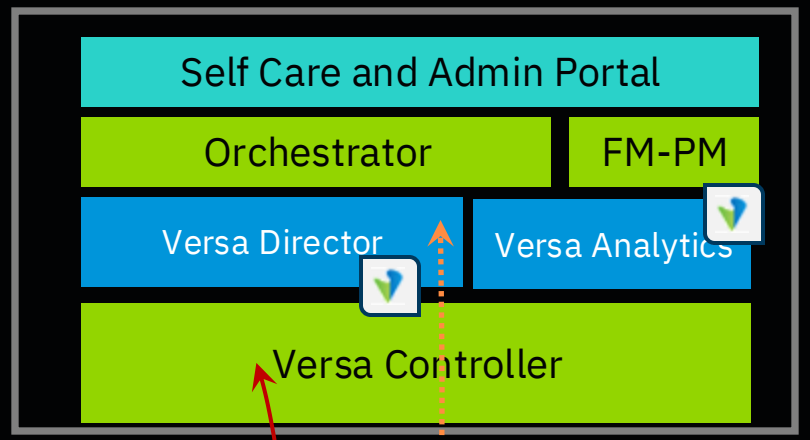


Data Center
or
PoP
Deployment

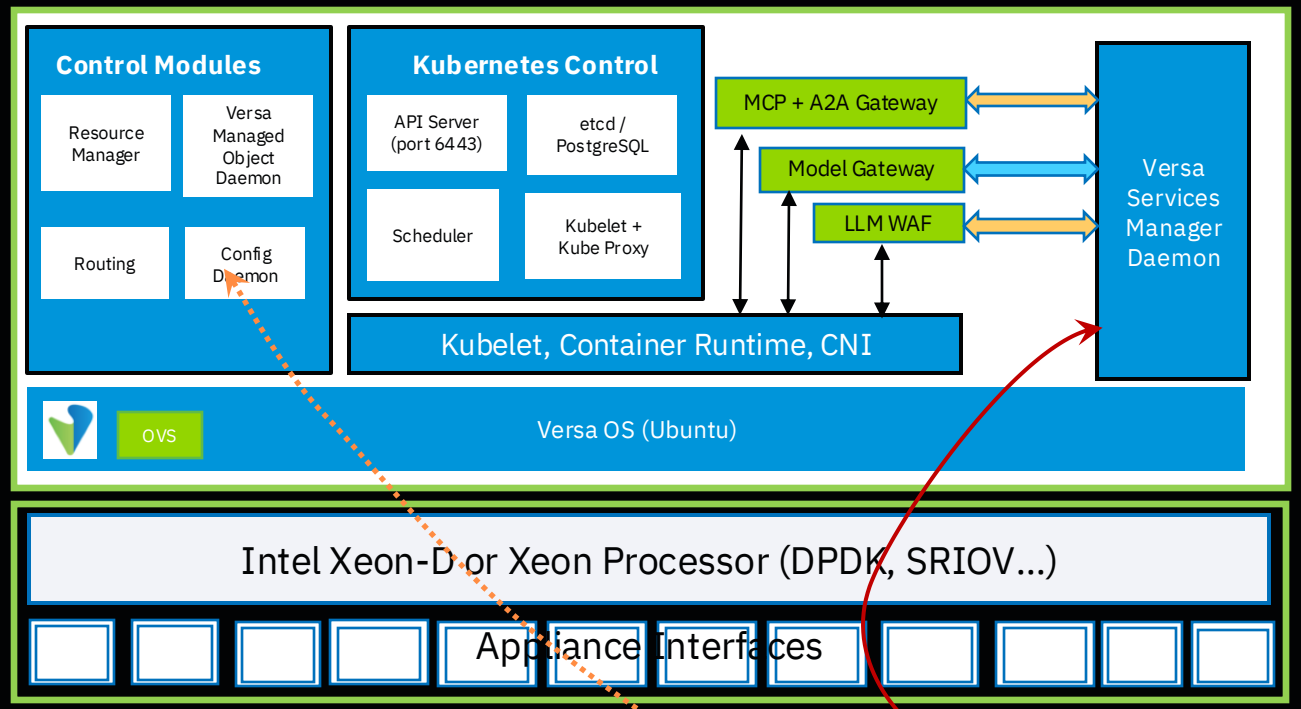


| | |
|---|---|
| **IKE Based IPSec** | ⟷ |
| NETCONF within ssh within IPSec | ⟷ |

VERSATILITY

# Agentic AI Security Flow on Versa SASE

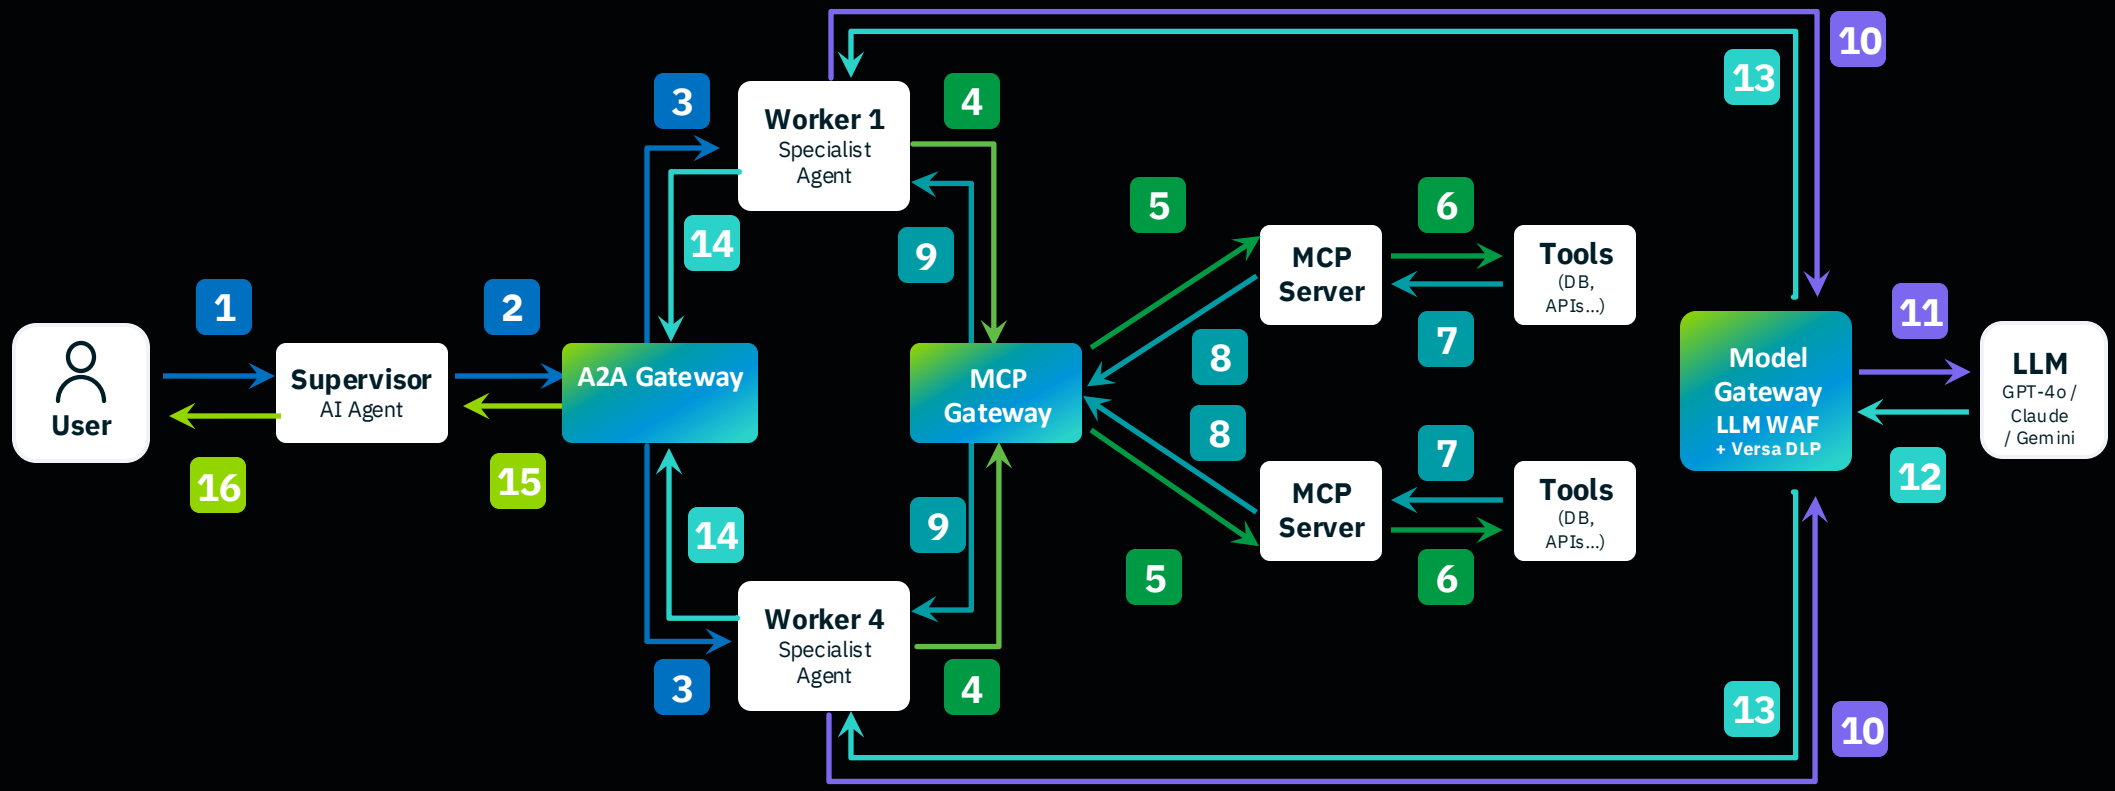


1 – 3
Agent Delegation

4 – 6
Tool Calls

7 – 9
Tool Response

10 – 11
LLM Request

12 – 15
LLM Response

15 – 16
Return to User

VERSATILITY

# Agentic AI Security Flow on Versa SASE

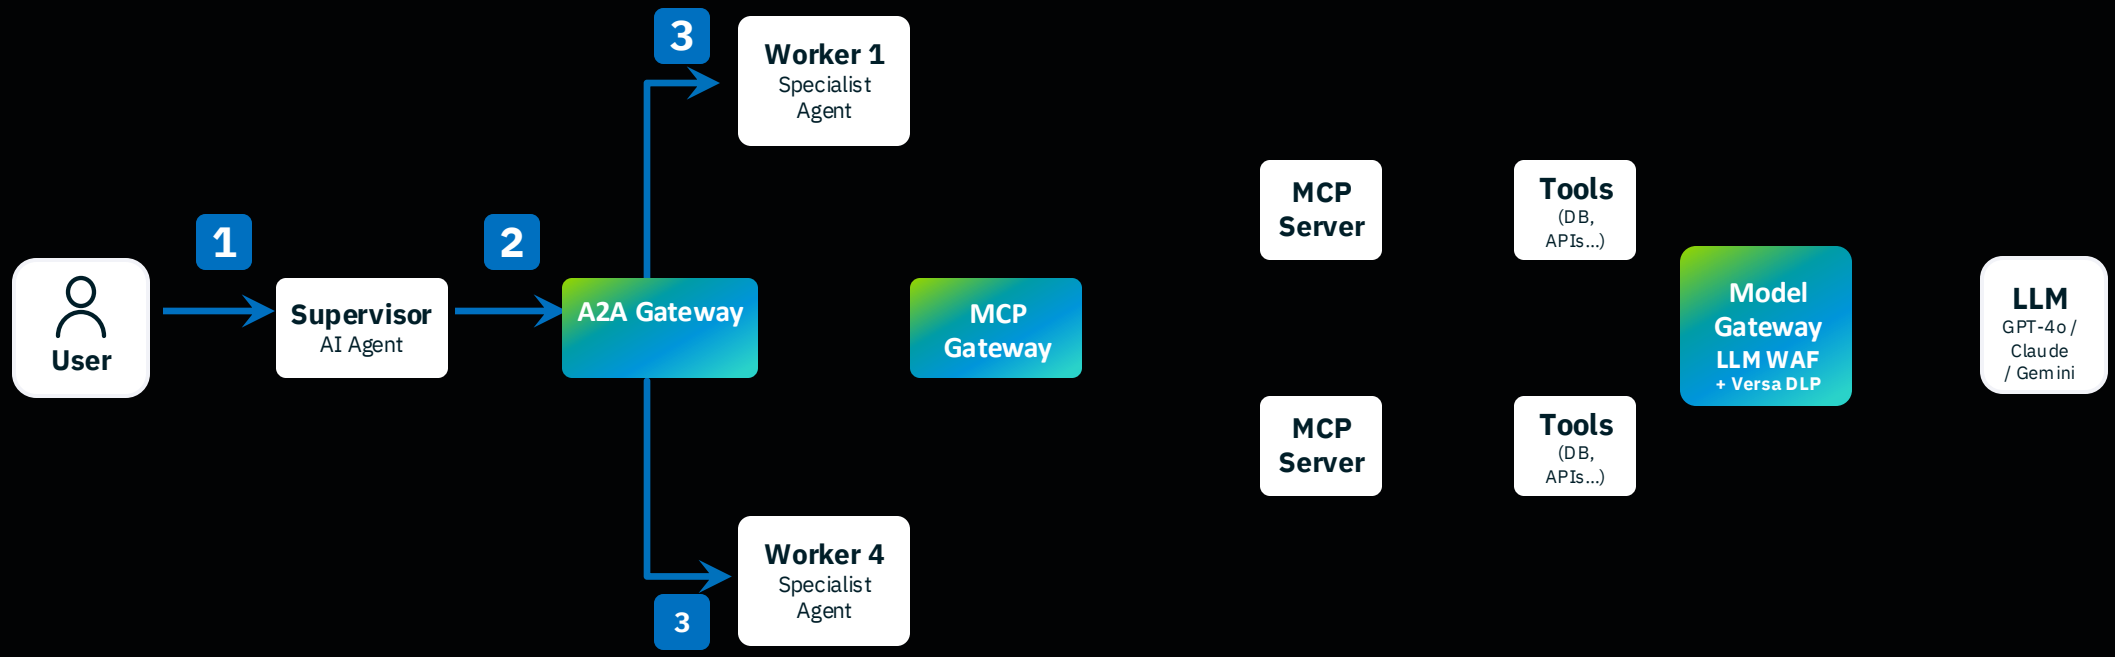


**1 – 3**
**Agent Delegation**

**4 – 6**
Tool Calls

**7 – 9**
Tool Response

**10 – 11**
LLM Request

**12 – 15**
LLM Response

**15 – 16**
Return to User

VERSATILITY

# Agentic AI Security Flow on Versa SASE

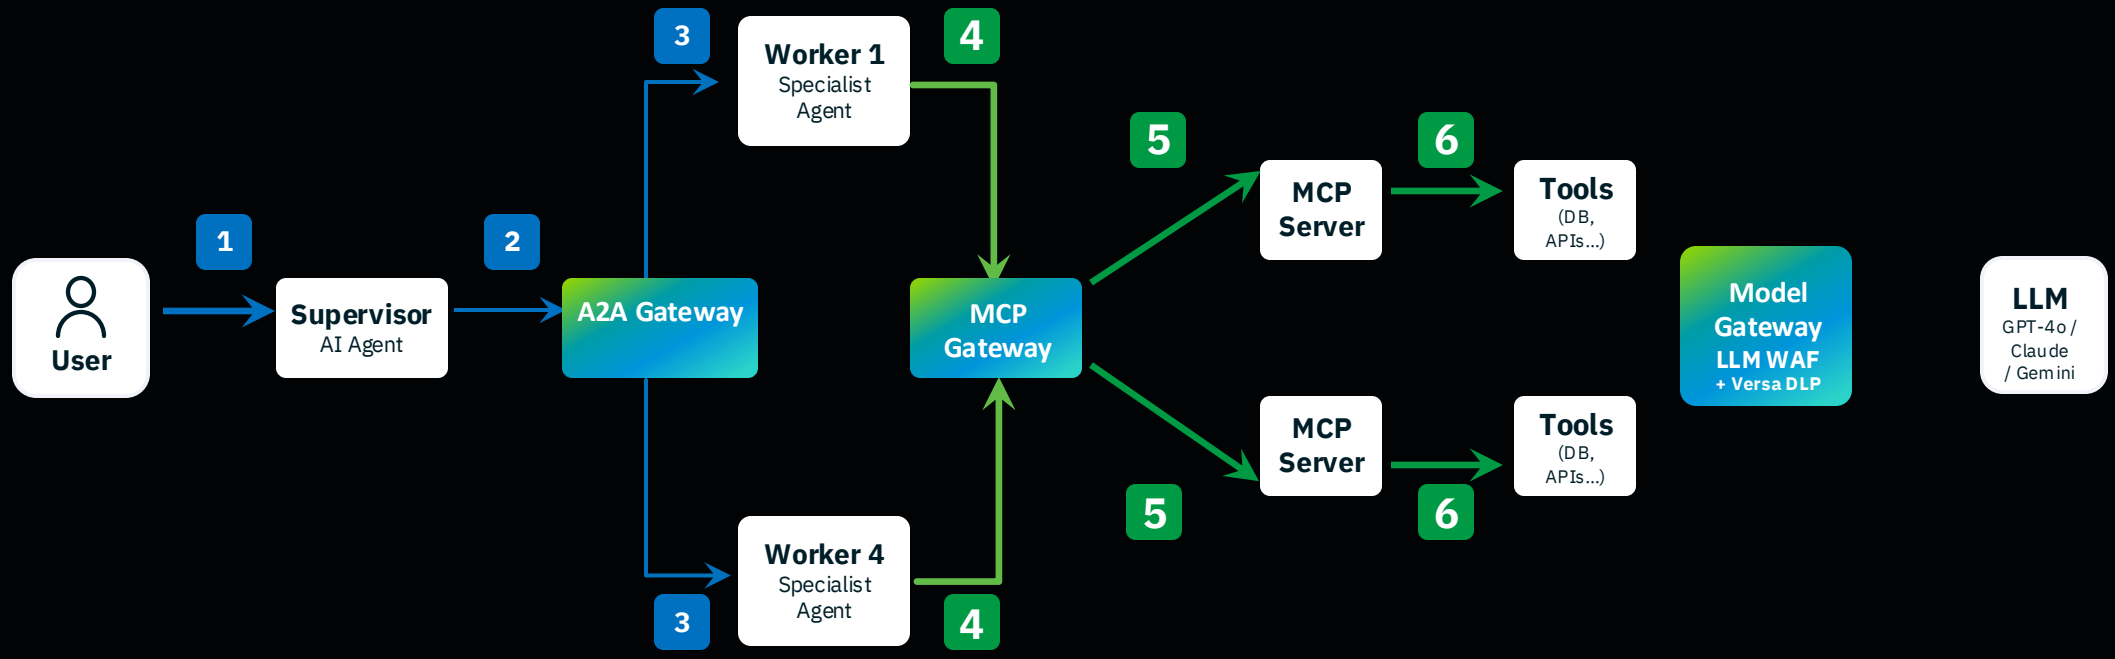


1 – 3
Agent Delegation

4 – 6
**Tool Calls**

7 – 9
Tool Response

10 – 11
LLM Request

12 – 15
LLM Response

15 – 16
Return to User

VERSATILITY

# Agentic AI Security Flow on Versa SASE

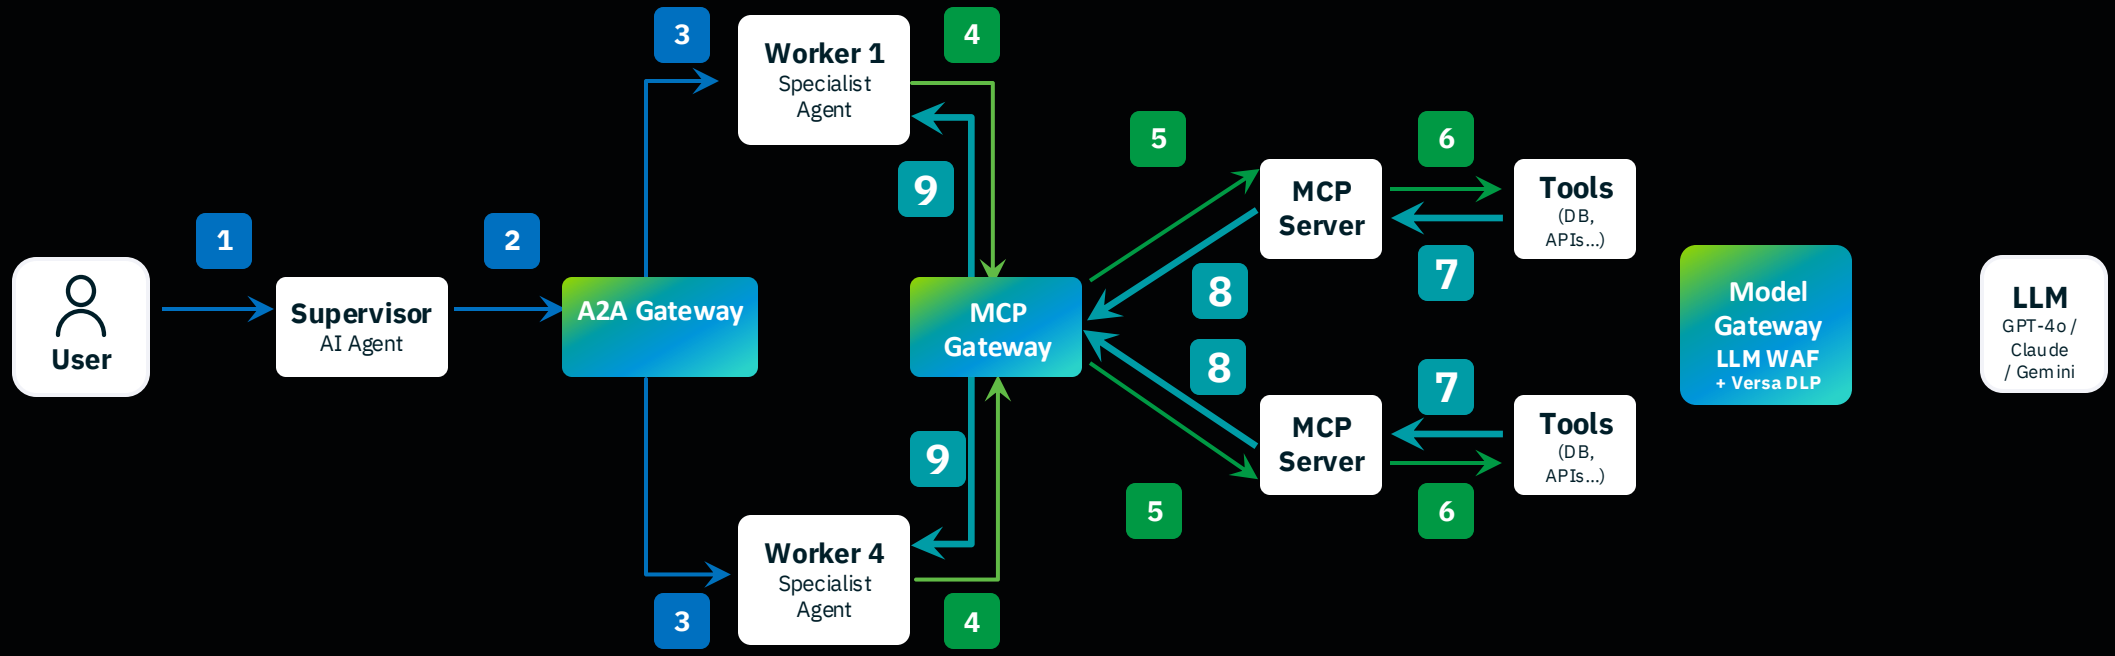


1 – 3
Agent Delegation

4 – 6
Tool Calls

7 – 9
**Tool Response**

10 – 11
LLM Request

12 – 15
LLM Response

15 – 16
Return to User

VERSATILITY

# Agentic AI Security Flow on Versa SASE

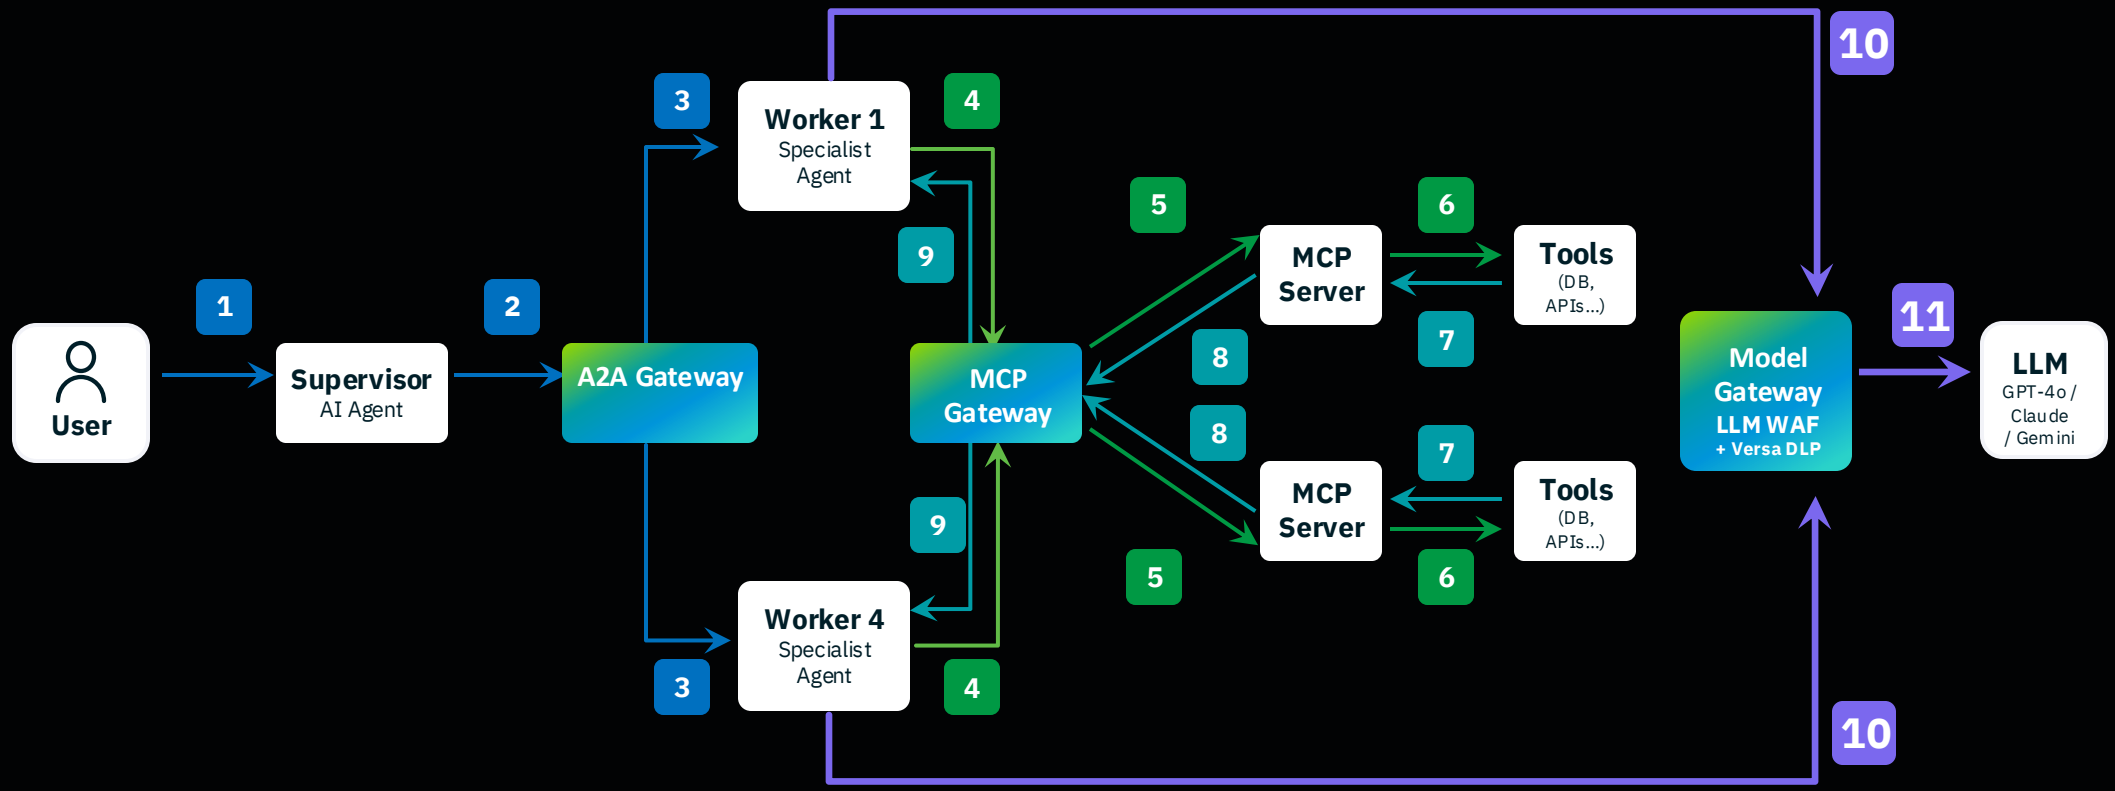


1 – 3
Agent Delegation

4 – 6
Tool Calls

7 – 9
Tool Response

10 – 11
**LLM Request**

12 – 15
LLM Response

15 – 16
Return to User

VERSATILITY

# Agentic AI Security Flow on Versa SASE

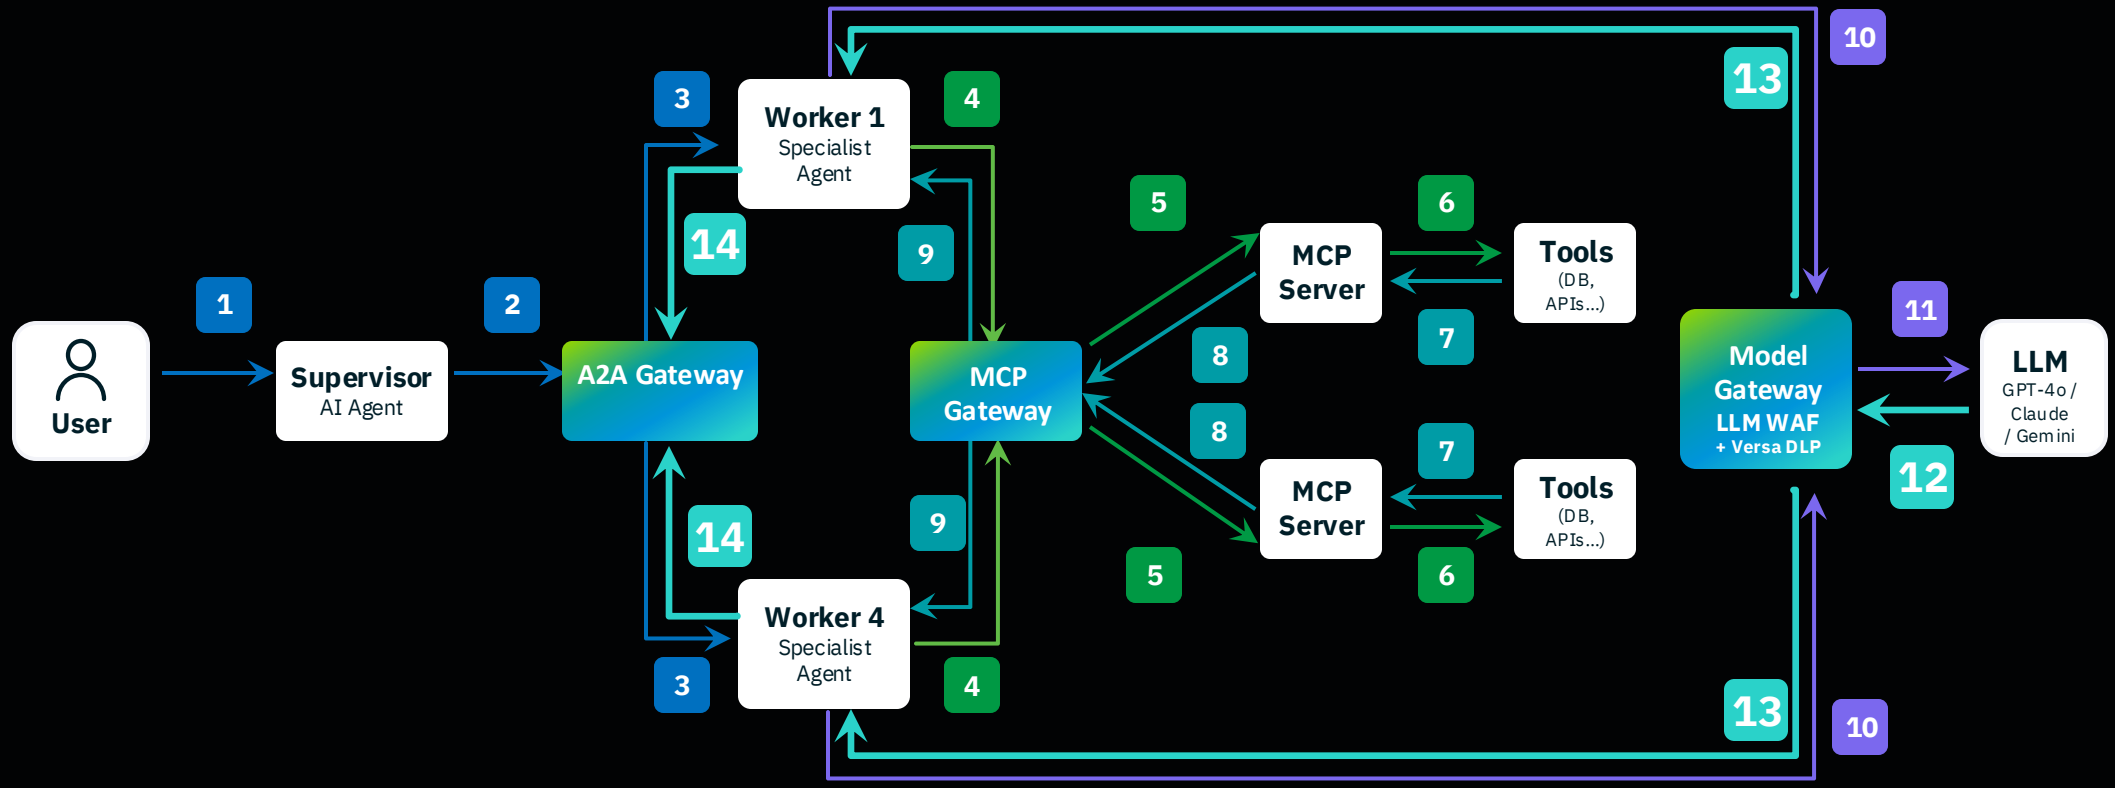


1 – 3
Agent Delegation

4 – 6
Tool Calls

7 – 9
Tool Response

10 – 11
LLM Request

12 – 15
**LLM Response**

15 – 16
Return to User

VERSATILITY

# Agentic AI Security Flow on Versa SASE

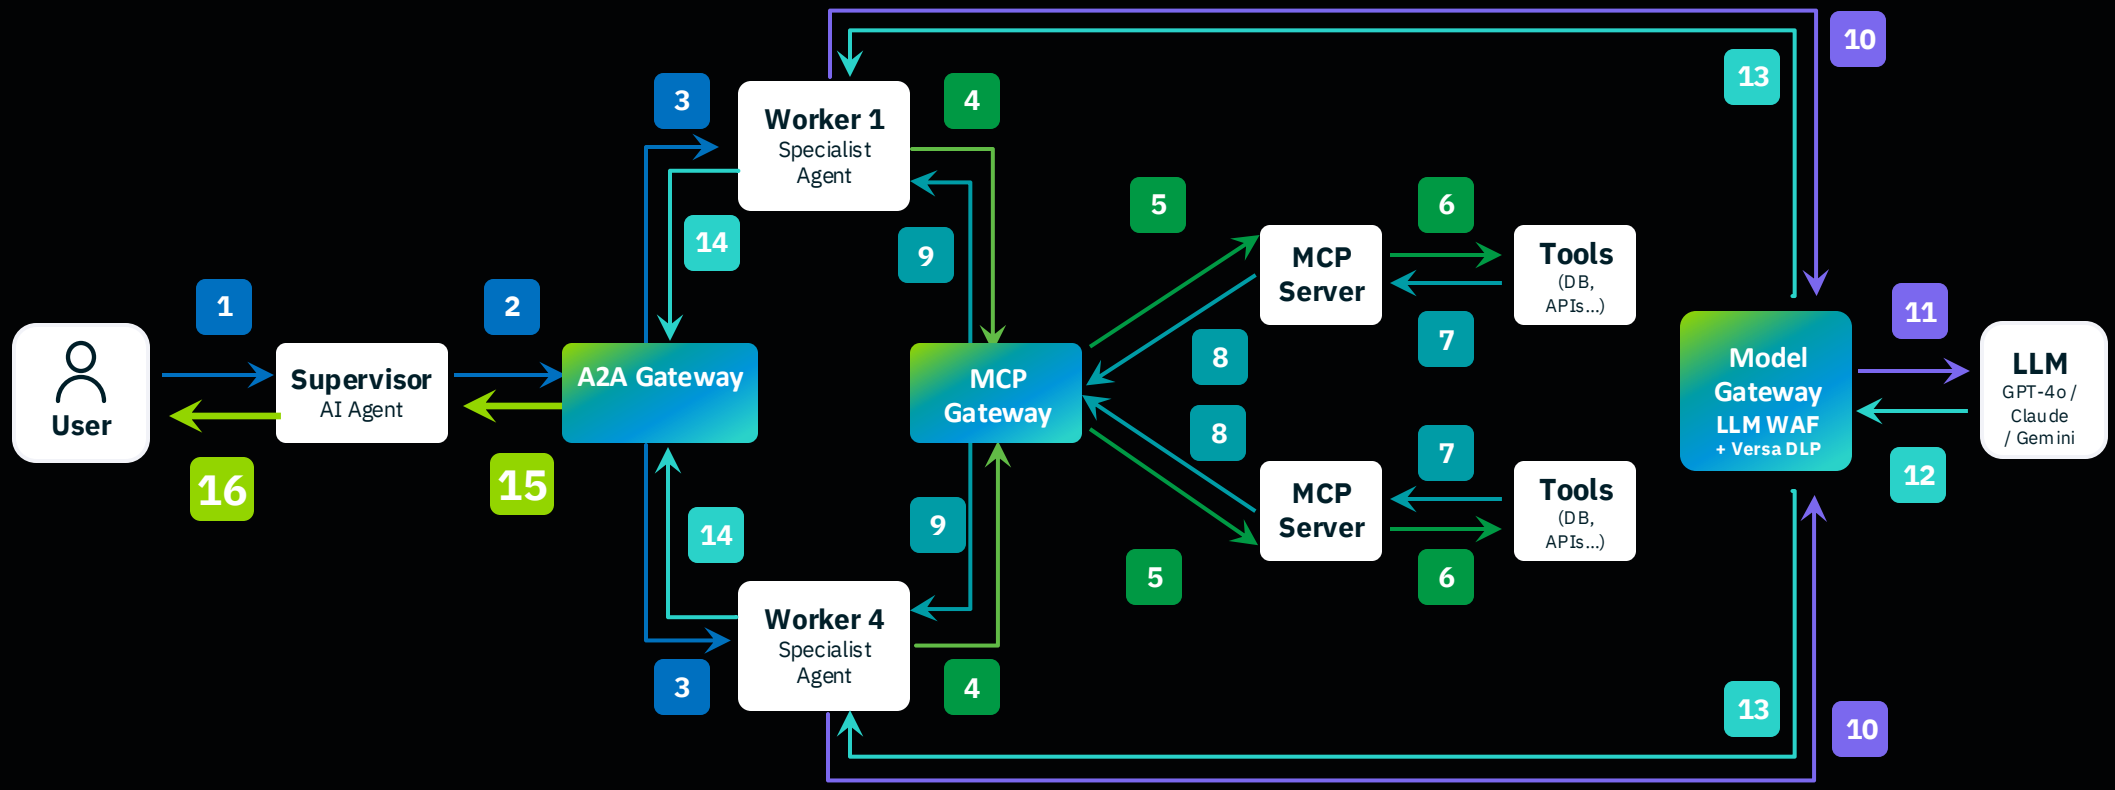


1 – 3
Agent Delegation

4 – 6
Tool Calls

7 – 9
Tool Response

10 – 11
LLM Request

12 – 15
LLM Response

15 – 16
**Return to User**

VERSATILITY

USE CASE

# Securing AI Copilots Accessing Enterprise Data

## Scenario

- An enterprise deploys an AI copilot (e.g. Microsoft Copilot, a custom LLM assistant)
- That accesses internal systems — HR records, ERP, customer databases, SharePoint
- Via MCP tool calls.

## Challenges

- Uncontrolled access
- No identity enforcement
- Scope creep
- Data exposure
- No audit trail
- Infrastructure sprawl
- Ungoverned egress

## Solution

**Versa MCP Gateway**

- **Identity-aware access:** only authenticated users/agents can invoke specific tools
- **Scope enforcement:** HR agent cannot call finance tools; read-only agents cannot write
- **DLP on tool responses**: sensitive PII/PHI/PCI in returned data is inspected inline before reaching the LLM
- **Full audit trail:** every tool call logged with user identity, timestamp, and payload summary

## Versa Advantage

- MCP Gateway runs on the on-prem VCG, co-located with the tools it governs.
- No data leaves the perimeter ungoverned.
- No new appliance needed — same VCG already enforcing ZTNA and SASE policies.

## Flow

User → AI Copilot → On-Prem VCG (MCP Gateway) → Internal Tools (DB, ERP, File Server)

USE CASE

# Governing Multi-Agent AI Workflows

## Scenario

- An enterprise runs an agentic AI workflow
- An orchestrator agent coordinates specialist subagents
- eg. security audit, network ops, HR policy, finance
- To complete complex tasks autonomously

## Challenges

- Unverified agents
- Permission escalation
- Prompt injection
- No east-west visibility
- Compliance gaps
- Ungoverned perimeter

## Solution

**Versa A2A Gateway**

- **Agent identity verification:** subagents must authenticate before receiving delegated tasks
- **Delegation scope:** orchestrator cannot grant subagents more permissions than it holds itself
- **Prompt injection defense:** Inspects every inter-agent message to detect injected instructions
- **East-west audit:** complete log of orchestrator-to-subagent calls, inputs, and outputs — essential for compliance

## Versa Advantage

- Agent-to-agent traffic never leaves the enterprise perimeter ungoverned
- A2A Gateway and MCP Gateway run on the same VCG instance
- One enforcement point for both inter-agent and agent-to-tool traffic

**Flow** User → Orchestrator Agent → VCG (A2A Gateway) → Specialist Subagents → VCG (MCP Gateway) → Enterprise Tools

VERSATILITY

USE CASE

# Protecting LLM API Traffic from Applications

## Scenario

- Enterprise developers and business applications call cloud LLM APIs directly
- eg. OpenAI, Anthropic, Gemini, Azure OpenAI
- Sensitive business data, source code, and customer PII is regularly sent in prompts — often without visibility or control

## Challenges

- Uncontrolled outbound data
- No prompt inspection
- Ungoverned consumption
- Provider lock-in risk
- Shadow AI usage
- Deployment friction

## Solution

**Versa Model Gateway + LLM WAF**

- **DLP on prompts:** detects and blocks PII, PHI, PCI, source code, and credentials before they reach the LLM provider
- **Prompt injection detection:** dual ML scanners block adversarial inputs attempting to hijack LLM behavior
- **Consumer-level controls:** per-application token budgets, rate limits, and model access policies
- **Provider abstraction:** applications call one endpoint; Versa handles failover across OpenAI, Azure, Anthropic

## Versa Advantage

- No endpoint agent or SDK change required
- Any app sending HTTPS to an LLM API is governed automatically by routing through the VCG
- The same enforcement point already handling SWG and CASB

## Flow

Application / Developer → VCG (Model Gateway + LLM WAF) → Cloud LLM API

# Why Versa: SASE-Native AI Security

**No New Product to Deploy**

Model GW, LLM WAF, MCP GW, and A2A GW run as service-chained containers on VOS. The same VCG already enforcing ZTNA, SWG, and CASB now also secures AI traffic — no separate appliance, no new management plane.

**Single Enforcement Point for All AI Traffic**

User-to-app, agent-to-tool, agent-to-agent, and app-to-LLM traffic all pass through the same VCG. One policy engine, one audit log, one place to enforce.

**SASE DLP Inside the LLM WAF**

Versa's inline DLP and CASB engines — already classifying documents, detecting PII/PCI/PHI — are applied to LLM prompts and responses. No standalone AI-specific DLP product needed.

**On-Prem and Cloud, Same Architecture**

MCP GW and A2A GW always run on-prem (tool traffic must stay inside the perimeter). Model GW and LLM WAF follow the LLM — cloud VCG for cloud APIs, on-prem VCG for private/local LLMs. Same containers, same policy model.

**Unified Management via Concerto**

AI security policies are configured and monitored in Concerto alongside existing SASE policies. No separate AI security console. No separate team required.

**Sovereign and Air-Gapped AI Support**

For government and regulated industries: all four components can run fully on-prem on Sovereign SASE or Private SASE deployments. AI governance without any cloud dependency.

VERSATILITY

# Thank you, and Questions

VERSATILITY