VERSA

WHITE PAPER

Zero Trust MCP: Securing Agentic AI for Autonomous Network and Security Operations

Contents

Introduction

The rise of Agentic AI is transforming how enterprises operate their networks. Unlike traditional automation, which relies on predefined workflows and human intervention, Agentic AI introduces autonomous agents that can monitor, decide, and act in real time. These agents can correlate telemetry, diagnose issues, and execute changes, shifting network operations from reactive troubleshooting to proactive, intelligent operations.

However, as AI agents become more deeply integrated into enterprise infrastructure, a fundamental challenge emerges. These agents are no longer passive observers; they actively interact with systems, APIs, and control planes. Without proper governance, this introduces significant risks, including unauthorized access, privilege escalation, and cross-tenant exposure. This is where Zero Trust MCP becomes essential.

The Evolution of Network and Security Operations

For decades, network operations have relied on manual processes supported by incremental automation. While automation has improved efficiency, it has not eliminated the need for human intervention. Engineers still spend significant time troubleshooting outages, correlating logs, and managing complex environments.

Agentic AI represents a paradigm shift. Instead of requiring humans to initiate actions, AI agents operate continuously, guided by policies and objectives. These agents can:

- Correlate logs and telemetry across network and security domains
- Detect anomalies and predict failures
- Validate configurations against network and security policies
- Optimize traffic and performance dynamically
- Check enforcement against threats with threat monitoring capabilities

This evolution shifts operations from “humans in the loop” to “humans on the loop,” where engineers supervise AI-driven systems rather than execute every task manually.

Challenges with Current Approaches

Despite the promise of AI, current implementations fall short in several areas:

- Alert Overload: Organizations face overwhelming volumes of logs, alerts, and events. Manual correlation is time-consuming and error-prone, leading to delayed response times and increased operational risk.
- Primitive Co-Pilots: Existing AI assistants rely on static knowledge bases and lack real-time context. They provide recommendations but cannot safely execute actions, limiting their value.
- Security Gaps: AI systems often lack secure communication pathways. Direct API access can expose infrastructure to unauthorized actions, increasing the risk of misuse.
- Lack of Governance: Without strong controls, AI-driven actions can lead to unintended consequences, including configuration errors, privilege escalation, and policy violations.

These challenges highlight the need for an architecture that combines AI-driven automation with strong security and governance.

Zero Trust MCP: What It Is

Zero Trust MCP is a security-first architecture designed to govern how AI agents interact with enterprise infrastructure. It ensures that every AI-driven action is authenticated, authorized, and auditable.

At its core, Zero Trust MCP introduces a brokered execution model. AI agents do not directly execute API calls or interact with infrastructure. Instead, all requests are routed through a centralized control system that enforces policies, role-based access control, and tenant isolation.

This approach ensures that AI remains assistive and governed, rather than fully autonomous without oversight.

Core Architectural Components

1. **Agentic Architecture**

   Zero Trust MCP is built on an agentic framework where specialized AI agents perform distinct functions such as routing queries, troubleshooting issues, and executing workflows. These agents are context-aware and operate based on user roles and system state.

2. **Brokered Execution Model**

   All AI-generated actions are executed through a management plane rather than directly by the AI agent. This prevents unauthorized access and ensures that every action is validated before execution.

3. **Zero Trust Enforcement**

   Every interaction is governed by strict policies. Access is role-based, context-aware, and continuously validated.

4. **Tenant Isolation**

   Actions are scoped within defined boundaries, ensuring that data and operations remain isolated across tenants.

5. **Auditability and Governance**

   All actions are logged and auditable, enabling compliance and governance across the environment.

Key Architectural Principles

Zero Trust by Design: Every AI action is governed, auditable, and role restricted. This ensures compliance and reduces risk.

Context-Aware Decision Making: AI agents operate with real-time awareness of network state, user roles, and operational context, enabling accurate and secure execution.

Separation of Control and Execution: AI agents generate intent, while execution is handled by a controlled system. This separation prevents direct access to infrastructure.

Policy-Driven Operations: All actions are enforced through centralized policies, ensuring consistency across the environment.

Use Cases

1. **Network Health Monitoring**

   AI agents continuously monitor distributed infrastructure across WAN, LAN, and cloud environments, providing real-time visibility into network performance, availability, and anomalies. By correlating telemetry from multiple sources, they can proactively identify degradation in latency, packet loss, or application performance before it impacts users, enabling faster detection and continuous operational assurance.

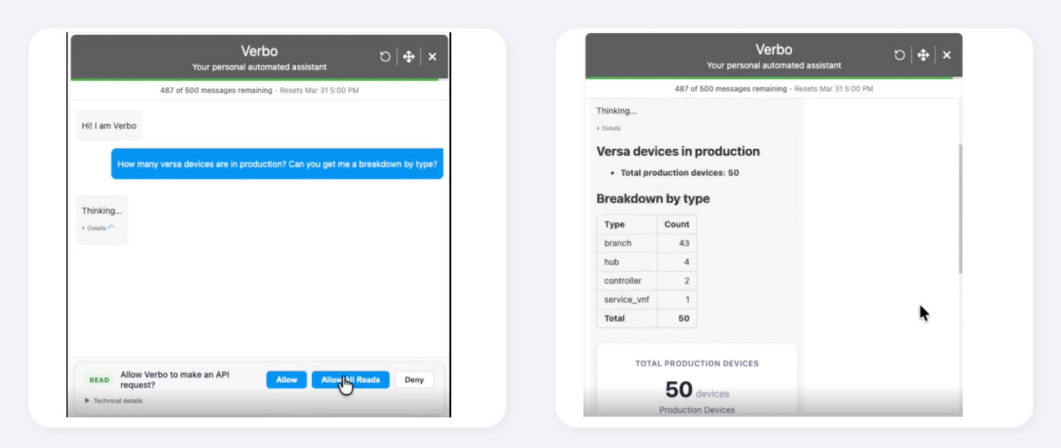

2. **Troubleshooting and Incident Response**

   AI agents execute complex, multi-step diagnostic workflows in seconds by correlating logs, flow data, and performance metrics across the network and security stack. They can isolate root causes, validate configurations against policies, and recommend or trigger guided remediation actions, significantly reducing mean time to resolution (MTTR) and minimizing operational disruption.

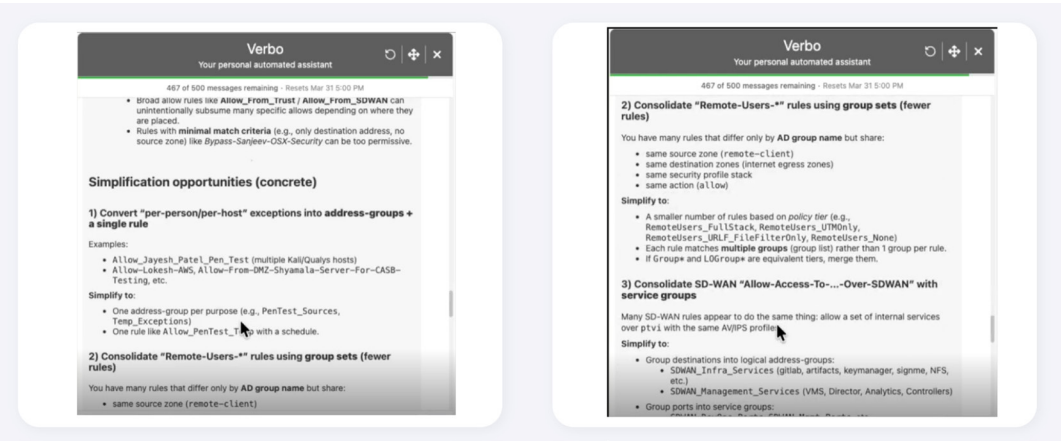

3. **Security Enforcement**

   Zero Trust MCP enforces granular, identity- and context-aware policies across all interactions. Every AI-driven action is validated through role-based access control (RBAC), tenant isolation, and policy enforcement layers, ensuring that access is tightly controlled and unauthorized lateral movement is prevented. This enables consistent and auditable security enforcement across the entire infrastructure.

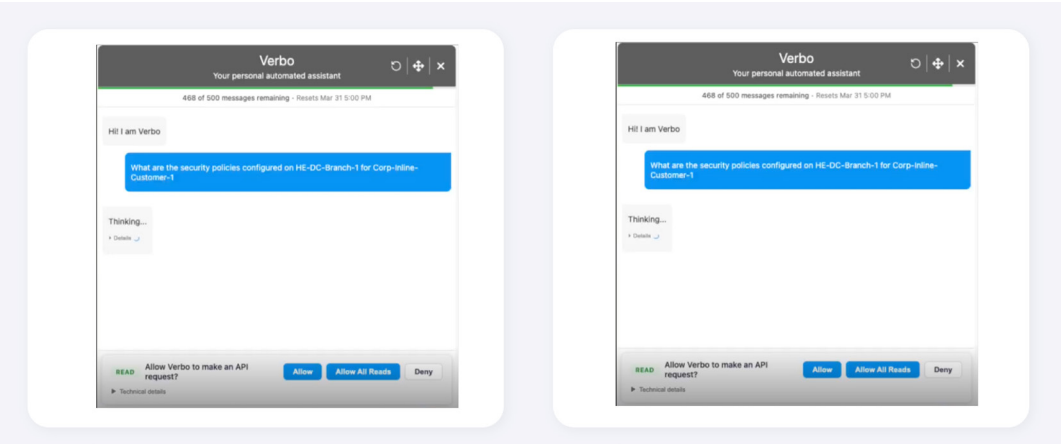

4. **User Experience Optimization**

Through integrated digital experience monitoring, AI agents analyze user-level performance across applications, SaaS services, and collaboration tools. They can identify performance bottlenecks, correlate them with underlying network conditions, and dynamically adjust policies or resource allocation to ensure optimal user experience across distributed environments.

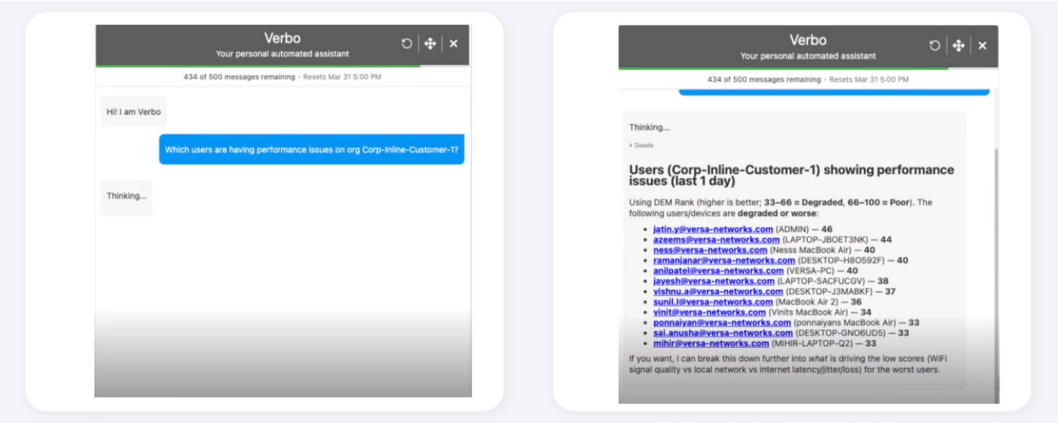

5. **Threat Monitoring**

AI-driven monitoring continuously analyzes network and security telemetry to detect anomalies, suspicious behaviors, and potential threats in real time. By integrating security context with network insights, agents can enforce containment actions, reduce dwell time, and limit the spread of threats across the environment.

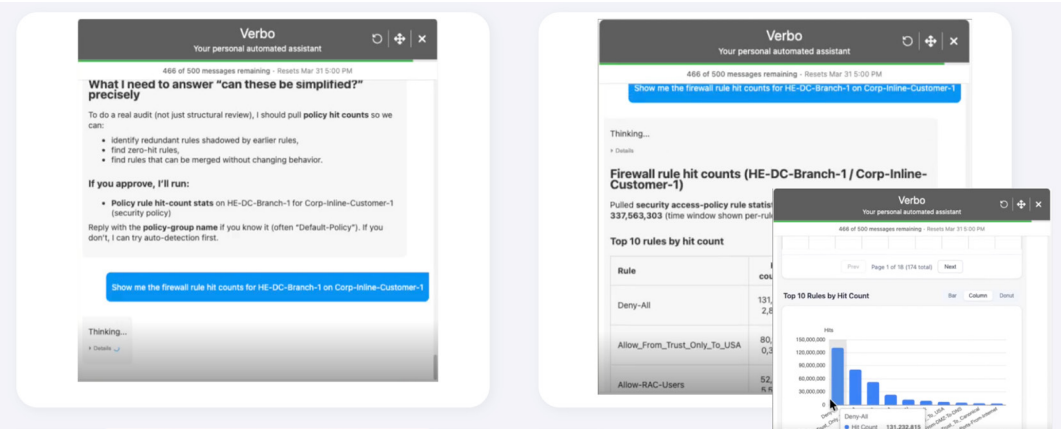

6. **Compliance and Governance**

Zero Trust MCP provides continuous auditing and validation of security policies, configurations, and AI-driven actions. Organizations gain full visibility into who performed what action, under which policy, and when, ensuring compliance with regulatory requirements and internal governance standards while maintaining operational transparency.

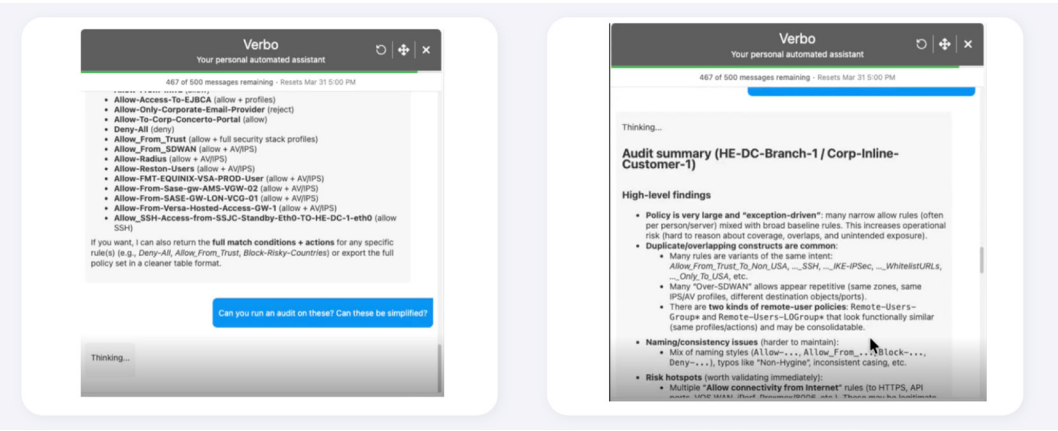

Business Outcomes

- Reduced MTTR
- AI-driven troubleshooting and guided workflows significantly reduce mean time to resolution.
- Improved Uptime
- Proactive monitoring and anomaly detection ensure continuous availability of services.
- Enhanced Security
- Zero Trust MCP enforces strict controls, reducing the risk of unauthorized access and lateral movement.
- Operational Efficiency
- Automation reduces manual effort, allowing teams to focus on strategic initiatives.
- Simplified Operations
- A unified platform integrates observability, security, and automation, reducing complexity.

Why This Matters Now

The shift to Agentic AI is inevitable. As networks become more complex and distributed, manual operations will not scale. AI-driven systems offer the ability to operate at speed and scale, but only if they are governed securely.

Zero Trust MCP provides the foundation for this transformation. By ensuring that every AI-driven action is controlled, validated, and auditable, it enables enterprises to adopt autonomous operations with confidence.

Conclusion

Agentic AI represents a fundamental shift in network operations, enabling autonomous, intelligent, and adaptive infrastructure. However, without proper governance, this shift introduces new risks.

Zero Trust MCP addresses these challenges by combining AI-driven automation with Zero Trust principles. Through brokered execution, policy enforcement, and tenant isolation, it ensures that AI operates securely and effectively.

As enterprises move toward autonomous networks, adopting Zero Trust MCP will be essential to achieving secure, scalable, and resilient operations.

VERSA

Versa Networks, Inc
2550 Great America Way, Suite 350
Santa Clara, CA 95054
Tel: +1 408.385.7660
Email: info@versa-networks.com
www.versa-networks.com

About Versa
Versa, a global leader in SASE, enables organizations to create self-protecting networks that radically simplify and automate their network and security infrastructure. Powered by AI, the VersaONE Universal SASE platform delivers converged SSE, SD-WAN, and SD-LAN solutions that protect data and defend against cyberthreats while delivering a superior digital experience. Thousands of customers globally, with hundreds of thousands of sites and millions of users trust Versa with their mission critical networks and security.

Part# WP_ZTMCP-01.0