gruve

SOLUTION BRIEF - MANAGED SERVICES

AI Managed AESM & AEBA Services: Always-on AI agent & endpoint security monitoring and behavioral analytics

01

Business problem

Enterprises are rapidly deploying AI-enabled applications and autonomous agents to accelerate business operations. These systems introduce a new class of security and governance challenges that traditional SOC platforms, processes, and tooling are not designed to detect or manage.

AI agents interact directly with enterprise APIs, SaaS platforms, and internal systems through Model Context Protocol (MCP) servers and AI gateways. These interactions are non-deterministic, highly dynamic, and driven by autonomous decision logic. As a result, organizations face critical risks including:

AI-specific security threats

- Unauthorized or unintended business actions executed by AI agents
- Sensitive data exposure through agent tool calls and API responses
- Business logic abuse and automated misuse of backend services
- Insecure or unmanaged MCP servers creating direct backdoors into production systems
- Agent jailbreak attempts and prompt injection attacks
- Credential replay and session hijacking by AI-driven bots

Operational and governance challenges

- Lack of visibility into agent behavior and AI workflow execution
- Inability to demonstrate compliance with emerging AI governance regulations
- High alert volumes with 40%+ false positive rates overwhelming security teams
- No dedicated AI security expertise or capacity for 24/7 monitoring
- Ungoverned developer-created MCP servers creating security blind spots
- Inability to detect subtle data exfiltration patterns in AI traffic

The core problem: Security teams struggle to operate these environments because existing SOC tools focus on network, endpoint, and identity telemetry—not agent behavior, AI tool usage, or MCP-based workflows. Internal teams also lack the specialized skills required to continuously tune AI security policies, interpret AI-specific threats, and manage AI infrastructure at scale.

Without a dedicated operational model for AI security, organizations risk silent data breaches, uncontrolled AI behavior, regulatory violations, and operational instability.

02

Why now

The urgency for AI-native security operations has reached a critical inflection point that demands immediate action.

AI agents moving to production
AI agents are transitioning from pilot environments into business-critical production workflows. This shift dramatically increases the impact of security failures. Organizations deploying multiple AI applications accessing sensitive systems (databases, CRM, financial platforms) cannot afford to operate without continuous monitoring and threat detection capabilities.

Regulatory pressure intensifying
Global frameworks such as the EU AI Act and ISO 42001 require organizations to demonstrate continuous governance, monitoring, and accountability for AI systems. Organizations subject to these regulations face upcoming audits requiring AI governance demonstration, board and executive pressure for AI compliance, and the need for continuous monitoring evidence.

AI-native threats emerging faster than security programs can adapt
AI-native threats such as agent jailbreaks, automated abuse of APIs, and large-scale data extraction through agent workflows are emerging faster than traditional security programs can adapt. Adversaries leverage AI to compress attack timelines while defenders struggle with manual processes.

The operational capacity gap
Organizations recognize that building internal AI security capability takes 6-12+ months. The cybersecurity talent shortage has grown from 3.4 million to 4.8 million professionals in the past year. Internal SOC teams are overwhelmed with current workloads and lack both the capacity and specialized expertise required for AI security monitoring.

Organizations must move immediately from ad-hoc AI security controls to a managed, operational AI SOC capability that continuously protects AI systems and their underlying integrations.

03

Solution overview

Gruve's Managed AI Security Services (AESM & AEBA) deliver a fully operated, 24×7 AI Security Operations capability for organizations running AI applications and autonomous agents.

This service provides continuous monitoring, detection, response, governance, and lifecycle management for AI gateways, MCP servers, and agent-to-application interactions through **AI Agent & Endpoint Security Monitoring (AESM)** and **AI Agent & Endpoint Behavioral Analytics (AEBA).**

How it works:

Gruve integrates customer AI environments directly into its dedicated AI SOC, where expert analysts and automated workflows:

- Monitor every AI interaction in real-time
- Identify abnormal agent behavior through behavioral analytics
- Detect business logic abuse and data exfiltration patterns
- Enforce security guardrails and compliance policies
- Manage MCP server lifecycle and integrations
- Provide rapid threat containment and incident response

Unlike traditional managed SOC services, this offering is purpose-built for AI workloads and agentic systems, providing deep visibility into:

- **AI agent tool usage:** Every tool call and API invocation
- **API invocation patterns:** Normal vs. anomalous behavior baselines
- **MCP server activity:** Comprehensive monitoring of all MCP traffic
- **Session behavior:** User-agent interactions and session binding
- **Policy enforcement outcomes:** Real-time guardrail effectiveness
- **Data access patterns:** Detection of sensitive data exposure

This service ensures AI systems remain secure, compliant, and operational as organizations scale their agentic workloads.

04

Service tiers

Service tier	Monthly investment	Description
Tier 1: Managed monitoring & incident response	$25,000 - $45,000	24×7 monitoring of AI gateway and MCP traffic, alert triage, investigation of suspicious agent behavior, automated and manual containment actions, and bi-weekly operational reporting. Core "eyes-on-glass" protection integrating your AI Gateway directly into Gruve's Security Operations Center. Designed for organizations requiring immediate AI threat detection and response capability.
Tier 2: Managed governance, compliance & lifecycle	$45,000 - $85,000	Full operational ownership of the AI gateway ecosystem including all Tier 1 capabilities PLUS policy lifecycle management, MCP server integrations, patching, configuration management, advanced threat hunting, compliance mapping, and quarterly strategic posture reviews. Designed for enterprises running mission-critical AI workloads at scale with regulatory compliance requirements.

05

Operating methodology

Tier 1 – Managed monitoring & incident response

Core capabilities:

24/7 AI security monitoring
- Continuous surveillance of AI-API and MCP traffic for anomalies and indicators of compromise
- Real-time monitoring of every interaction between AI agents and backend applications
- Behavioral baseline establishment and deviation detection

AI-driven alert prioritization and noise suppression
- Advanced analytics to suppress false positives (40-60% reduction)
- Intelligent correlation of AI-specific threat indicators
- Risk-based alert prioritization highlighting high-impact threats

Analyst-led investigation of agent behavior anomalies
- Expert analysis distinguishing authorized tasks from malicious intent
- Investigation of jailbreak attempts and prompt injection attacks
- Deep-dive analysis of suspicious data access patterns

Automated and manual containment
- Automated playbooks to disconnect rogue MCP servers or block suspicious IP addresses
- Manual intervention for complex threat scenarios
- SOAR-integrated response workflows for rapid containment

Regular operational reporting
- Bi-weekly detailed summaries of blocked threats, agent performance, and Gateway health
- Trend analysis and emerging threat intelligence
- Metrics on detection effectiveness and response times

Typical use cases:

- **Agentic threat hunting:** Detecting AI agents being "jailbroken" to perform unauthorized database queries
- **Bot mitigation:** Stopping malicious bots that mimic legitimate human-AI interactions to scrape sensitive data
- **Data exfiltration prevention:** Identifying and blocking sensitive data exposure through agent tool calls
- **Business logic abuse detection:** Recognizing patterns of automated misuse of backend services
- **SOC load reduction:** Outsourcing complex AI monitoring to free internal teams for core security operations

Key outcomes:

You gain immediate, around-the-clock peace of mind knowing that every AI transaction is being scrutinized by security experts specifically trained in agent behavior analysis, significantly reducing the risk of a silent breach through your AI interfaces.

Tier 2 – Managed governance, compliance & lifecycle

Core capabilities:

All Tier 1 Capabilities PLUS:

Policy lifecycle management

- Continuous tuning of security guardrails and AI policies
- Regular reviews and updates of RBAC, session binding, and data masking rules
- Policy versioning and change management
- Optimization based on evolving business requirements

MCP server and application integration management

- Full lifecycle management of MCP server deployments
- Coordination of MCP integrations with application changes
- Discovery and governance of developer-created MCP servers
- Centralized MCP registry and governance enforcement

Systems management and maintenance

- Patch management and configuration control
- Platform health and performance monitoring
- Configuration backups and disaster recovery
- Proactive maintenance preventing downtime

Compliance telemetry mapping and audit reporting

- Continuous mapping of AI telemetry to regulatory requirements (EU AI Act, ISO 42001, NIST AI RMF)
- Audit-ready documentation with evidence of continuous governance
- Regular compliance reports for regulatory examinations
- Board-level compliance presentations

Advanced threat hunting

- Proactive, analyst-led searches for hidden vulnerabilities within AI-to-Application integrations
- Quarterly threat hunting campaigns
- Emerging threat intelligence integration
- Strategic risk identification and mitigation

Quarterly strategic posture reviews

- Deep-dive sessions with Gruve experts aligning security posture with AI roadmap
- Strategic recommendations for emerging AI use cases
- Risk assessment for planned AI deployments
- Continuous improvement planning

Typical use cases:

- **Regulated AI adoption:** Ensuring AI agents in healthcare or finance strictly follow data privacy and PII masking rules
- **Multi-cloud governance:** Maintaining unified security policy for AI Gateways across hybrid cloud environments (AWS, Azure, GCP, On-prem)
- **Operational modernization:** Shifting from reactive security posture to proactive, managed model anticipating AI-specific risks
- **MCP server governance:** Centralizing management of ungoverned developer-created MCP servers
- **Compliance demonstration:** Providing continuous evidence of AI governance for auditors and regulators

Key outcomes:

Your organization receives a turnkey governance framework that scales with your AI ambitions. You move from "securing a tool" to "managing a mission-critical AI capability," ensuring long-term resilience, operational continuity, and audit-readiness.

06

Key benefits

24×7 AI SOC coverage
Continuous protection of AI gateways and agent interactions with expert analysts trained specifically in agent behavior analysis, MCP architecture, and AI workflow security. Maintains monitoring during nights, weekends, and holidays without requiring internal staffing increases.

Reduced operational overhead
Organizations typically reduce internal operational effort for AI security and MCP management by 30–50%. Eliminates need to hire specialized AI security analysts (6-12 month timeline). Frees internal SOC team to focus on traditional security threats rather than AI-specific monitoring.

Noise suppression and analyst efficiency
AI-driven alert prioritization suppresses false positives by 40-60%, allowing security teams to focus only on high-risk AI-related threats. Intelligent correlation and behavioral analytics dramatically reduce investigation time and analyst burnout.

Continuous compliance readiness
Automated policy enforcement and reporting aligned to enterprise governance and emerging AI regulations (EU AI Act, ISO 42001, NIST AI RMF). Continuous evidence generation for audits and regulatory examinations. Quarterly compliance posture reviews ensuring ongoing alignment.

Secure AI infrastructure operations
Enterprise-grade management of MCP servers, integrations, and platform health. Full lifecycle ownership including patching, configuration management, and performance monitoring. Proactive maintenance preventing downtime and security incidents.

Rapid threat containment
Immediate blocking of malicious agent activity and abnormal API usage through SOAR-integrated response workflows. Automated playbooks for common threat scenarios combined with expert manual intervention for complex attacks. Mean time to contain reduced by 50-60%.

Deep behavioral analytics (AEBA)
Advanced behavioral analytics establishing baselines for normal agent activity and detecting subtle deviations indicating compromise, jailbreak attempts, or data exfiltration. Machine learning models trained specifically on AI agent behavior patterns.

Complete visibility
100% visibility into agent behavior, API usage, MCP server activity, and data access patterns. Comprehensive audit trails of all AI interactions. Real-time dashboards and historical reporting for security posture assessment.

07

What makes Gruve's managed AI security unique

Purpose-built AI SOC
Gruve operates a dedicated AI SOC that monitors AI behavior and agent workflows rather than only infrastructure or user activity. Our analysts are specifically trained in:

- Agent behavior analysis and anomaly detection
- MCP architecture and security implications
- AI workflow security and threat modeling
- Business logic abuse patterns in autonomous systems
- AI-specific attack techniques (jailbreaks, prompt injection)

We operate what others only deploy
The service takes full ownership of operational responsibilities including:

- 24/7 monitoring, detection, and response
- Patching and configuration management
- Policy tuning and optimization
- Threat hunting and vulnerability identification
- Lifecycle management of MCP servers and integrations

Most providers sell you tools and leave you to operate them. We operate the complete AI security capability on your behalf.

Deep AI domain expertise
Our team combines:

- Security operations expertise from managing 24/7 SOCs for 100+ global enterprises
- AI-specific knowledge understanding how agents work, how they can be compromised, and how to detect threats
- Integration experience across all major security platforms (SIEM, SOAR, EDR/XDR)
- Compliance expertise mapping AI security controls to regulatory requirements

Integrated automation platform
Automated containment and remediation workflows are integrated with Gruve's SOAR capabilities for rapid response:

- Pre-built playbooks for common AI threat scenarios
- Custom automation tailored to your environment
- Integration with existing security tools and workflows
- Continuous playbook optimization based on threat intelligence

Compliance-first operations

- Continuous mapping of AI telemetry to regulatory and internal governance requirements ensures audit readiness at all times:
- Evidence collection for EU AI Act, ISO 42001, NIST AI RMF
- Audit-ready documentation and reporting
- Quarterly compliance validation and gap analysis
- Board-level presentations on AI security posture

Rapid time-to-value

- Monitoring active: Within 24 hours of deployment
- Full governance & reporting: Comprehensive visibility in 7 days
- Systems management: Integration and management of MCP servers in 7–14 days

No lengthy implementation projects. Immediate protection and value delivery.

08

Common operational use cases

Agent threat detection and response

- Agent jailbreak detection and containment
- Prompt injection attack identification
- Unauthorized data extraction through AI tool calls
- Business logic abuse through autonomous workflows
- Credential replay and session hijacking by AI-driven bots
- API misuse and rate limit violations

MCP server security and lifecycle

- Insecure MCP server discovery and shutdown
- Centralized governance of developer-created MCP servers
- MCP server patching and configuration management
- Monitoring MCP server activity for anomalies
- Integration management for application changes
- Performance optimization and capacity planning

Data protection and privacy

- Real-time detection of PII/PHI exposure in agent responses
- Enforcement of data masking policies
- Monitoring for unauthorized database access
- Detection of sensitive data exfiltration patterns
- Compliance with data protection regulations (GDPR, HIPAA)

Compliance and governance

- Continuous monitoring evidence for EU AI Act compliance
- Audit-ready documentation for ISO 42001 certification
- Mapping AI security controls to NIST AI RMF requirements
- Regulatory examination support with compliance reports
- Board presentations on AI security posture

Operational efficiency

- Reducing internal SOC workload from AI alert investigation (30-50% overhead reduction)
- Eliminating need to hire specialized AI security analysts
- Providing 24/7 coverage without staffing increases
- Freeing internal teams to focus on traditional security threats
- Accelerating time to production for new AI applications

09

Expected service outcomes

Tier 1: Managed Monitoring & Incident Response

Security effectiveness

- 40-60% reduction in false positive alerts from AI systems
- 50% improvement in mean time to detect AI-specific threats
- 24/7 coverage with expert analyst triage and investigation
- Immediate containment of confirmed threats

Operational efficiency

- 30-40% reduction in internal SOC workload for AI monitoring
- Elimination of analyst burnout from AI alert volume
- Bi-weekly reporting providing full visibility into AI security posture
- No requirement to hire specialized AI security analysts

Compliance and governance

- Continuous monitoring evidence for regulatory requirements
- Audit trails of all AI interactions and security events
- Baseline governance capability meeting essential compliance needs

Tier 2: Managed Governance, Compliance & Lifecycle

Comprehensive security and compliance

- 50-70% reduction in false positive alerts across all AI systems
- 60-70% improvement in mean time to detect threats
- 50-60% reduction in mean time to respond and contain incidents
- Complete visibility across entire AI infrastructure with minimal blind spots

Audit-ready compliance

- Complete compliance telemetry mapped to EU AI Act, ISO 42001, NIST AI RMF
- Audit-ready documentation with evidence of continuous governance
- Quarterly compliance validation and gap analysis
- Board-level presentations demonstrating AI security posture

Strategic AI security capability

- Full lifecycle management preventing MCP server sprawl and security gaps
- Proactive threat hunting identifying vulnerabilities before exploitation
- Continuous policy optimization adapting to evolving AI use cases
- Quarterly strategic reviews aligning security with AI roadmap

Operational excellence

- 40-50% reduction in total operational overhead for AI security
- Zero unplanned downtime from security incidents or maintenance
- Enterprise-grade platform management and performance optimization
- Complete internal team enablement and knowledge transfer

10

Time to value

Unlike lengthy security transformation projects requiring 6-12 months of implementation, Gruve's Managed AI Security Services deliver immediate value:

Week 1:
- Consultation and scoping (2-3 hours)
- Contracting and onboarding planning (1-2 weeks)
- Environment assessment and telemetry integration
- Monitoring activated within 24 hours

Weeks 3-4:
- MCP server inventory and governance (Tier 2)
- Advanced threat hunting initiation (Tier 2)
- First strategic posture review (Tier 2)

Week 2:
- Behavioral baselines established
- Custom containment playbooks deployed
- Full operational capability with bi-weekly reporting

Ongoing:
1. Continuous monitoring, detection, and response
2. Regular policy tuning and optimization
3. Quarterly reviews and strategic planning (Tier 2)

No lengthy implementation. Immediate protection. Measurable value from day one.

11

Service delivery model

Onboarding and integration (week 1)

Environment assessment
- Inventory of AI applications and agents
- MCP server discovery and cataloging
- Integration point identification
- Baseline traffic analysis

Telemetry integration
- Secure connectivity establishment
- Log forwarding configuration
- API integration with AI gateways
- SIEM/SOAR platform integration

Behavioral baseline establishment
- Normal agent behavior profiling
- API usage pattern analysis
- User-agent interaction modeling
- Risk scoring calibration

Containment playbook customization
- Environment-specific playbook development
- Integration with existing security tools
- Escalation workflow configuration
- Communication protocol setup

Stakeholder enablement
- Kickoff session with security leadership
- Training on reporting and escalation procedures
- Portal access and dashboard orientation
- 24/7 contact information distribution

Ongoing operations

24/7 monitoring and detection
- Continuous surveillance of all AI traffic
- Real-time behavioral analytics and anomaly detection
- Alert triage and prioritization
- Expert investigation of suspicious activity

Incident response and containment
- Automated response for known threat patterns
- Manual intervention for complex scenarios
- Coordination with internal security team
- Post-incident analysis and lessons learned

Regular reporting

- Bi-weekly operational reports (Tier 1)
- Monthly trend analysis and metrics (Tier 2)
- Quarterly strategic reviews (Tier 2)
- Ad-hoc executive briefings as needed

Continuous improvement

- Policy tuning based on operational feedback
- Playbook optimization and development
- Threat intelligence integration
- Emerging threat research and adaptation

Lifecycle management (Tier 2)

- MCP server patching and updates
- Configuration management
- Application integration coordination
- Performance monitoring and optimization

12

Investment and pricing

Tier 1: Managed Monitoring & Incident Response

Monthly investment: $25,000 - $45,000

Pricing based on:

- Number of AI applications and agents under management
- Volume of AI traffic and transactions
- Number of MCP servers being monitored
- Integration complexity with existing security infrastructure

Tier 2: Managed Governance, Compliance & Lifecycle

Monthly investment: $45,000 - $85,000

Pricing based on all Tier 1 factors PLUS:

- Number of MCP servers requiring lifecycle management
- Complexity of compliance mapping requirements
- Frequency and depth of threat hunting activities
- Number of integrations requiring ongoing management

Engagement Model

- Month-to-month service agreements (no long-term lock-in)
- Quarterly business reviews with expansion/reduction flexibility
- Fixed monthly pricing with transparent scope definition
- No hidden fees or surprise charges

Typical ROI: Organizations reduce operational overhead by 30-50% while preventing high-cost data breaches (average cost: $4.45M) and regulatory fines (EU AI Act fines up to €30M or 6% of global revenue).

Return on investment typically achieved within 3-6 months through combination of operational efficiency, prevented incidents, and compliance cost avoidance.

13

Establish continuous AI security operations

Engage Gruve's Managed AI Security Services to protect your AI applications, agents, and MCP infrastructure with 24×7 monitoring, response, and governance.

This service enables organizations to safely scale AI adoption while maintaining full visibility, security, and regulatory compliance.

Next steps:

Step 1: Schedule Consultation (2-3 hours)
- Connect with Gruve's AI security specialists
- Assess your current AI deployment and security posture
- Identify specific monitoring and governance requirements
- Review service tiers and investment levels
- Define success criteria and engagement scope

Step 2: Contracting and Onboarding Planning (1-2 weeks)
- Finalize statement of work and service scope
- Execute master services agreement
- Schedule onboarding activities
- Provision secure connectivity and access
- Mobilize Gruve AI SOC team
- Establish communication protocols

Step 3: Onboarding and Integration (Week 1)
- Conduct environment assessment
- Configure telemetry integration
- Establish behavioral baselines
- Customize containment playbooks
- Stakeholder kickoff and training
- Activate 24/7 monitoring

Step 4: Begin Realizing Value (Week 2 and Beyond)
- Continuous monitoring, detection, and response
- Bi-weekly operational reporting
- Regular policy tuning and optimization
- Quarterly strategic reviews (Tier 2)
- Ongoing threat intelligence and adaptation

Website: https://www.gruve.ai/

Email: info@gruve.ai

Request consultation: https://www.gruve.ai/contact

Prerequisites:
- AI agents in production or pilot-to-production transition
- AI gateway or MCP infrastructure deployed or planned
- Executive sponsorship for managed AI security investment

Partner technology

Cequence AI Gateway with Cequence Security
Gruve's Managed AI Security Services are delivered in partnership with Cequence Security, providing enterprise-grade AI gateway infrastructure with comprehensive visibility, control, and protection capabilities for AI workloads and agentic systems.