MindSOC 1.0

Mindwalker's AI-Powered SOC Automation Platform

Enterprise-ready alert correlation, deep attack context, and governed automation with proactive defense

MindSOC is a next-generation SOC automation platform that unifies alerts across security systems and applies AI-driven analysis to deliver deep attack context, cross-signal correlation, and root-cause identification. It autonomously investigates incidents using a five-stage LLM-driven pipeline with tool calling, integrating seamlessly with SIEM, EDR, ITSM, Threat Intel, and Identity providers. MindSOC delivers complete investigation verdicts directly into ticketing systems, reducing Mean Time to Respond (MTTR) by 87%, achieving over 95% verdict accuracy, and auto-closing more than 90% of false positives. Designed for enterprise resilience, it provides early warnings before escalation, reduces alert fatigue, and accelerates decision-making through governed automation with audit trails and analyst approvals.

Autonomous Investigation

5-stage AI pipeline processes incidents end-to-end in under 3 minutes

Non-Disruptive

Results appear directly in existing ticketing systems (Jira, ServiceNow), no context switching for analysts

Adaptive

Learns from every analyst override through a built-in feedback loop, continuously improving verdict accuracy

Multi-Source Enrichment

Queries SIEM, Threat Intel, EDR, and Identity providers in parallel during each investigation

Enterprise-Ready

On-premise/VPC deployment, encrypted credentials, RBAC, audit logging, SOC 2 Type II preparation

Deployment Model

MindSOC is deployed as a VPC/on-premise solution within the customer's infrastructure. All data remains within the customer's network boundary. Three tiers are available to match different SOC team sizes and operational requirements:

Essential
Small team production

Premium
Full automation

Enterprise
MSSP / mission-critical

Performance Highlights

INVESTIGATION TIME **20x faster**	VERDICT ACCURACY **>99%**	AUTO-CLOSE FP RATE **>90%**
THROUGHPUT **3x increased**	MTTR REDUCTION **<20 sec**	ROI **3x multiplied**

Tier Overview

MindSOC offers three deployment tiers designed to align with the scale and needs of your SOC team

	Essential	Premium	Enterprise
Target Customer	Small SOC teams (1-5 SOC team)	Mid-size SOC (5-20 SOC team)	Large SOC / MSSP (20+ SOC team)
Deployment	Single VPS Docker Compose	Single VPS (HA) Docker Compose	Multi-node Kubernetes High Availability
Use Case	Small team production SOC	Medium-size production SOC	Multi-tenant MSSP
AI Powered Incident Verdict	✓		
Automated Escalation	✗	✓	
Hi-Availability	✗	✗	✓

Technical Specifications

Specification	Essential	Premium	Enterprise
CAPACITY & LIMITS			
Incidents / Day	Up to 100	Up to 1,000	Unlimited (Infrastructure dependence)
Concurrent Investigation	5	25	100+
Knowledge Base Size	100 documents	1,000 documents	Unlimited (Infrastructure dependence)
Asset Inventory	500 end-points	5,000 end-points	Unlimited (Infrastructure dependence)
Data Retention	Based on storage capacity (VPC deployment)		
Concurrent Users	5	25	Unlimited
Use Cases / Playbooks	5	25	Unlimited
Support SLA	8x5 Business Day	24x7	24x7, 1 hour response + Deddicated TAM
PERFORMANCE SLA			
Investigation Time (Avg)	< 5 min	< 3 min	< 1 min
Investigation Time (P95)	< 10 min	< 5 min	< 3 min
API Response Time	< 500ms	< 200ms	< 100ms
Dashboard Load	< 3 sec	< 2 sec	< 1 sec
WebSocket Latency	-	< 500ms	< 100ms
Platform Uptime SLA	99.5%	99.9%	99.99%
AI & INVESTIGATION			
LLM Provider	Cloud Based LLM	Cloud Based LLM + Self-hosted SLM (optional)	Cloud Based LLM + Self-hosted LLM

Specification	Essential	Premium	Enterprise
Verdict Accuracy SLA	90%	95%	99%+
Auto-Close FP Rate	80%	90%	95%+
FP Confidence Threshold	Fixed (0.90)	Configurable per use case	Adaptive ML-tuned
MITRE ATT & CK Mapping	✓	✓	✓
Feedback Learning Loop	Basic (Only true or false)	Advanced (True or false + comment section)	Advanced + ML Retraining (Fully automated ML retraining)
Historical Similiarity	✗	✓	✓
Ask AI / RAG Chat	✗	✓	✓

Feature Comparison by Tier

Feature	Essential	Premium	Enterprise
CORE INVESTIGATION PIPELINE			
5-Stage AI Pipeline	✓	✓	✓
Incident Classification (LLM)	✓	✓	✓
Knowledge Base RAG Lookup	✓	✓	✓
Multi-Source Enrichment	3 connectors max	7 connectors	Unlimited (Infrastructure dependence)
Verdict Synthesis	✓	✓	✓
ITSM Write-Back	Jira only	Jira + Custom MCP	Jira + Custom MCP
Auto-Close False Positives	Manual only	✓	✓
Auto-Escalate True Positives	Manual only	✓	✓
SIEM Integration Compatibility		✓ Wazuh ✓ Google SecOps (Chronicle) ✓ Microsoft Sentinel ✓ IBM QRadar ✓ Custom SIEM Connector	
Threat Intelligence Compatibility		✓ VirusTotal ✓ MISP ✓ Shodan ✓ Recorded Future ✓ Custom TI Feed	
EDR Integration Compatibility		✓ SentinelOne ✓ CrowdStrike ✓ Microsoft Defender	
Identity Provider Compatibility		✓ Entra ID (MS Graph) ✓ Okta ✓ JumpCLoud	

Feature	Essential	Premium	Enterprise
VISUALIZATION & DASHBOARD			
Incident Queue	✓	✓	✓
Investigation Detail View	✓	✓	✓
KPI Dashboard	Basic (4 files)	✓ (Full)	✓ + Custom
Trend Charts (7/30/90d)	✗	✓	✓
Use Case Breakdown	✗	✓	✓
Real-time Websocket	✗	✓	✓
Executive Reports	✗	✓	✓ + Scheduled
Report Export (PDF/CSV)	✗	✓	✓
MANAGEMENT & SECURITY			
User Management	5 users	25 users	Unlimited
RBAC Roles	Admin, Analyst	Admin, Analyst, Viewer	All 4 roles + Custom
SSO (SAML / OIDC)	✗	✗	✓
Audit Logging	Based on storage	Based on storage	Based on storage + Export
Asset Inventory	500 assets	5,000 assets	Unlimited
API Rate Limitting	50 req/min	100 req/min	500 req/min + burst
Multi-Tenancy	✗	✗	✓
MSSP Dashboard	✗	✗	✓
Data Residency (Multi-region)	✗	✗	✓
SOC 2 Type II Compliance	✗	Preparation	✓

Enterprise Tier Scalability

- Incident Volume
  Up to 10,000+ incidents/day sustained
- Concurrent Investigations
  100+ parallel investigations
- Concurrent Users
  Unlimited
- Knowledge Base
  Unlimited documents with HNSW vector index auto-optimization
- Multi-Region
  NA, EU, APAC deployment zones
- Horizontal Scaling
  Kubernetes auto-scaling for API, Worker, and Redis
- Database
  PostgreSQL cluster with read replicas and pgvector HNSW tuning
- LLM
  Multi-provider failover (primary + fallback mechanism)

Deployment Options

	Starter	Professional	Enterprise
Docker Compose	✓	✓	✓
Kubernetes (Helm)	✗	Optional	✓ (Required)
Managed Cloud	✗	✗	✓
On-Premise	✗	✓	✓
Air-Gapped	✗	✗	✓
Nginx Reverse Proxy	✓	✓	✓
TLS Termination	✓	✓	✓ + mTLS

Infrastructure Requirements

	Essential	Premium	Enterprise
API Server	Bundled in subscription	2 nodes (HA): 8 vCPU / 16 GB RAM 200 GB SSD each	4+ nodes (HA): 16 vCPU / 32 GB RAM 500 GB SSD each
Kubernetes (Helm)	Bundled in subscription	2 nodes (HA): 8 vCPU / 16 GB RAM 100 GB SSD each	4+ nodes (HA): 16 vCPU / 32 GB RAM 200 GB SSD each
Managed Cloud	Bundled in subscription	2 nodes (Primary/Replica): 8 vCPU / 32 GB RAM 1 TB SSD each	2+ nodes (Cluster): 16 vCPU / 64 GB RAM 2 TB SSD each
On-Premise	Bundled in subscription	2 nodes (Sentinel): 4 vCPU / 8 GB RAM each	3+ nodes (Cluster): 8 vCPU / 16 GB RAM each
Air-Gapped	Not Available	500 GB	5 TB+
Nginx Reverse Proxy	Not required (Cloud LLM only)	Optional (SLM acceleration)	Recommended NVIDIA A10G / A100
TLS Termination	Bundled in subscription	2 nodes (HA): 8 vCPU / 16 GB RAM 200 GB SSD each	4+ nodes (HA): 16 vCPU / 32 GB RAM 500 GB SSD each

Contact
+62 21 2880 2137
PT. Reka Citra Solusi
Ruko Grand Centro Bintaro Blok A19
JL. Kodam Bintaro Raya, Pesanggrahan
Jakarta Selatan
Jakarta 12320
Indonesia

Website
https://www.mindwalker.ai