

MindOps 2.0

Mindwalker's AI-Native Operational Intelligence Platform

From watching systems, to understanding them

MindOps is an AI-native operational intelligence platform that unifies metrics, logs, and traces into a single observability layer, with Mindy — an agentic AI SRE — automatically identifying root cause through hybrid vector-graph correlation and ML-based anomaly detection. Available as a complete observability platform or as a focused module that extends the APM you already have, MindOps deploys on-premise, on-cloud, or fully air-gapped — keeping all processing, storage, and AI inference inside your network.



Adaptive Baseline

Learns normal behavior for every metric automatically, including daily and weekly patterns, so you don't need manual thresholds.



Causal Root Cause Analysis

Uses your system's topology to find the exact root cause of a problem instead of just listing symptoms.



Problem Correlation

Groups related events into one record, preventing alert floods during cascading failures.



Mindy

Agentic AI that investigates and acts — as APM, SRE Agent, Log Analyzer, or Infra Monitoring.



Dashboard

Interactive BI-style dashboard for trend & SLO analytics, customizable with JSON import and custom metric cards.



Modularity

Five modular packages, from Mindy SRE AI Agent alone up to full Unified Observability, with flexible licensing and 8x5/24x7 SLA.

How It Works

1

Reads

Pulls topology, metrics, and events from your infrastructure or existing APM's REST API, based on deployment mode

2

Analyses

Baselining, fault-tree RCA, and correlation run locally; Mindy handles natural-language queries and investigation.

3

Returns

RCA findings and actions show up in the Dashboard, flow to ITSM, and, in intelligence-layer mode, back to your APM.

Data Sovereignty Guarantee: All processing, storage, and AI inference occur exclusively within your network. No telemetry, metrics, logs, or LLM queries leave your environment.

Performance Highlight



Operational Efficiency

- ✓ 50% MTTR reduction
- ✓ 30–40% cost reduction
- ✓ 60% senior engineer time savings



Improved Reliability

- ✓ < 2 min alert time
- ✓ < 30 sec RCA time
- ✓ 99.99% platform availability
- ✓ 85% RCA accuracy



Business Impact

- ✓ 40% reduction in customer impact
- ✓ 100% audit compliance
- ✓ 90% noise reduction



Data Sovereignty

- ✓ 100% on-premise processing
- ✓ Zero external data transfer
- ✓ Air-gapped deployment support

Feature List

Feature	Description
CORE AI CAPABILITIES	
Adaptive Baselining	Automatically learns normal behavior for every metric with seasonal patterns (hour-of-day, day-of-week) — no manual thresholds needed.
Causal Root Cause Analysis	Traverses topology using fault-tree analysis to pinpoint a single root-cause entity.
Problem Correlation	Groups all events sharing the same root cause into unified Problem records.
Mindy (Agentic AI)	Natural language interface powered by a local or hosted LLM — investigates app, infra, and data, recommends remediation actions, with human approval and full audit trail for high-risk actions.
APPLICATION PERFORMANCE MONITORING (APM)	
Service Performance	See latency (P50/P90/P99), request rate and error rate for every service in one table.
Service Deep-Dive	Drill into one service's latency, throughput, errors, DB calls and external API calls over time.
Operation Ranking	Rank a service's top-level operations by latency and call volume.
Service Dependency Map	A live dependency graph showing call paths, throughput and error hotspots between services.
DISTRIBUTED TRACING	
Trace Explorer	Query spans by any attribute, aggregate them, and switch between list, time-series and table views.

Feature	Description
End-to-End Trace View	Waterfall and flamegraph of a full distributed trace with per-span timing and attributes.
Journey Funnels	Build funnels across spans to measure conversion and drop-off between steps of a request flow.
Saved Trace Queries	Save frequently used trace queries as named, shareable views for the whole team.
LOG MANAGEMENT	
Log Search & Analysis	Full log search with the query builder, live histogram, and list/table/time-series views.
Live Tail	Tail logs live as they arrive, with your filters applied.
Log Pipelines	Build processing pipelines to parse, rename and enrich log fields before storage.
Contextual Log View	Jump to the surrounding log lines for any entry to get full context.
Saved Log Queries	Save common log searches as named, shareable views.
METRICS	
Metrics Catalog	A searchable summary of all ingested metrics with metadata and cardinality.
Metrics Explorer	Query and visualise metrics with the query builder, formulas and PromQL.
Saved Metric Views	Save metric queries as named, shareable views.
INFRASTRUCTURE MONITORING	
Host Monitoring	CPU, memory, disk and network for every host, out of the box.
Kubernetes Monitoring	Cluster, node, pod and workload health across your Kubernetes estate.
EXCEPTIONS & ERRORS	
Error Tracking	Every exception grouped by type with counts, trends and affected services.
Error Investigation	Full stack trace, occurrences and the linked trace for a specific exception.
MESSAGING QUEUES	
Kafka Monitoring	Consumer lag, throughput and partition health for Kafka.

Feature	Description
Celery Monitoring	Task rate, latency, failures and retries for Celery workers.
Queue Health Overview	One overview across all monitored messaging queues.
API MONITORING	
Outbound API Monitoring	Latency, error rate and volume for the external/third-party APIs your services call.
DASHBOARDS & VISUALIZATION	
Custom Dashboards	Create and organise dashboards from any signal: metrics, logs and traces.
Interactive Dashboards	Template variables, time-range control and drill-downs on every dashboard.
Panel Builder	Time-series, bar, table, value, pie, histogram and more, built with the query builder.
Public Dashboard Sharing	Publish a read-only dashboard via a public link, no login required.
ALERTS & NOTIFICATIONS	
Alert Rule Management	One place to view, enable/disable and organise every alert rule.
Threshold Alerts	Alert on metrics, logs or traces when a threshold is crossed.
AI Anomaly Alerts	Alert when a metric deviates from its learned normal behaviour, no manual threshold needed.
Alert History	Timeline of every firing, with the data that triggered it.
Alert Routing	Send alerts to Slack, email, PagerDuty, Opsgenie, MS Teams or a webhook.
QUERY & ANALYTICS ENGINE	
Unified Query Builder	The same visual query builder works across metrics, logs and traces.
Raw Query Support	Drop to ClickHouse SQL or PromQL when you need full power.
Formulas & Functions	Combine queries with formulas and apply functions for advanced analysis.
AI & ASSISTANT	
AI Observability Assistant	A conversational AI assistant that answers questions across your telemetry.

Feature	Description
AI/LLM Observability	Monitor LLM and AI-application workloads, including token usage.
ONBOARDING & INTEGRATIONS	
Guided Setup	Step-by-step onboarding for APM, logs, infra, AWS and Azure.
Integrations Catalog	One-click integrations for common data sources and tools.
Cloud Account Integrations	Pull telemetry directly from AWS/cloud accounts via the gateway.
ACCESS CONTROL & IDENTITY	
Team Management	Invite users, assign roles and manage membership.
Role-Based Access	Admin / Editor / Viewer roles control what each user can do.
Custom Fine-Grained Roles	Define custom roles with granular, per-resource permissions.
Enterprise SSO	Log in via SAML / OIDC SSO with your identity provider.
Programmatic Access	Service accounts and API keys for automation and ingestion.
PLATFORM ADMINISTRATION	
Organisation Settings	Manage org name, members and org-wide preferences.
Personal Settings	Each user manages their own profile and preferences.
Ingestion Management	Ingestion keys and limits for sending telemetry in.
Custom Domain	Serve the platform on your own custom domain.
MCP Server for AI Tools	Expose observability data to AI tools via a built-in MCP server.
Keyboard Shortcuts	Power-user keyboard shortcuts across the app.
Version & Health	See build version and system status.
METERING, BILLING & SUPPORT	
Usage Explorer	Usage explorer breaks down ingested data volume by signal and source.
Data Metering	Detailed data-metering explorer for chargeback and quota tracking.
Billing & Plans	Manage subscription, invoices and plan from inside the app.

Feature	Description
License Management	Apply and manage the Enterprise license key.
In-App Support	Chat and premium support channels built into the product.

Modules

Module	Adds On Top of Base	Included Scope
Mindy SRE AI Agent	✓ AI Assistant ✓ Access Control & Identity ✓ Alerts & Notifications ✓ Platform Administration ✓ Metrics ✓ Onboarding & Integrations ✓ Metering/Billing & Support	First 10 instances
+ Log Analyzer	✓ Dashboards & Visualization ✓ Distributed Tracing ✓ Exceptions & Errors ✓ Log Management	+ 100 nodes (from existing APM)
+ Infra Monitoring	✓ Infrastructure Monitoring ✓ Query & Analytics Engine	+ 100 infrastructure nodes
+ Application Observability	✓ API Monitoring ✓ Application Performance Monitoring (APM) ✓ Messaging Queues ✓ Query & Analytics Engine	+ 100 application nodes
Unified Observability	All of the above combined	10 instances + 100 infra nodes + 100 application nodes

Expansion Add-Ons

Add On	Adds On Top of Base
Additional 50 Nodes — Log Analyzer	Module 2 or 5
Additional 50 Infrastructure Nodes	Module 3 or 5
Additional 50 Application Nodes	Module 4 or 5
Additional 1 Instance — Mindy SRE AI Agent	Any module

Licensing & Support Options

License Type	Adds On Top of Base	SLA Options	Notes
Perpetual	One-Time Charge (OTC)	8x5 or 24x7	Includes 1st-year support
Annual Technical Support (ATS)	Annually	8x5 or 24x7	Support renewal for an existing Perpetual license (year 2 onwards)
Subscription	Monthly or Annually	8x5 or 24x7	Support included for the active subscription period

Support & Deployment Terms (common across all modules and license types)

- Response Time: 4 hours
- Support Tier: L1 and L2
- Support Channel: Online and Onsite
- Deployment: On-Premise, Cloud, or Air-Gapped
- Training: Product + M&O (Maintenance & Operation) training included with Perpetual and Subscription licenses.

Not included: deployment service (quoted separately via Professional Services SOW) and infrastructure/hardware provisioning.

Software Pre-Requisites

Component	Version	Notes
Operating System	Ubuntu 24.04 LTS	Or Compatible Linux distribution
Container Runtime	Docker 24+ / containerd	Required for K3s
Orchestration	K3s (Kubernetes)	Lightweight self-hosted K8s
Service Mesh	Istio 1.21+	mLTS, zero-trust network policy

Network Requirements

Connection	Protocol	Purpose
MindOps → APM Managed	HTTPS:8443	Or Compatible Linux distribution
Operator → Dashboard/API	HTTPS:443	Internal network access
MindOps → ITSM Systems	HTTPS/SMTP	Alert dispatch (ServiceNow, Jira, email, Slack)
Internet/External	None	No external connectivity required or permitted

Contact

+62 21 2880 2137
PT. Reka Citra Solusi
Grand Centro Bintaro
Jakarta, Indonesia

Website

www.mindwalker.ai