




REASON

The On-Prem AI Agent for Deep Reasoning

From generating tickets to actually resolving tickets.
LLM agnostic & GPU independent

What is REASON?



REASON is an on-premises AI deep reasoning and response platform that sits between your existing alerting stack and your ITSM workflow.

Unlike conventional AIOps tools that are cloud-first, data-hungry and platform-dependent, **REASON** thinks through incidents the way your own engineers would.

Whether it's about handling new situations, incomplete data, or multi-variable failures, **REASON** applies deep multi-step reasoning and turns solutions into remedial actions. No data migration, workflow redesign, or model training. Simply plug in and turn AIOps from alerts to actions.

? Why Isn't Your AI Strategy Working?

Most enterprise AI initiatives are burning money without making a dent to the business bottom line. But companies that adopt AI beyond improving productivity also faces multiple challenges, including:



Data Privacy

Public models expose sensitive operations data to leakage and regulatory risk. Compliance frameworks in regulated industries treat this as a hard procurement blocker.



Cost Control

Inference costs on public cloud platforms are unpredictable and escalate with event volume. Hosting large models in-house demands infrastructure budgets most teams cannot justify.



Talent Gap

Specialized AI engineers capable of building, maintaining, and fine-tuning custom AIOps pipelines are scarce, expensive, and stretched across competing priorities.



Model Training

Foundational model training requires data volumes, compute, and timelines that most enterprise IT teams simply don't have — killing initiatives before they reach production.

What can REASON do?



On-Premises Deployment

Data and response context stays entirely within your environment. Designed for regulated industries where data sovereignty is a hard requirement.




Tree of Thought Deep Reasoning

Multi-step reasoning that decomposes incidents into verifiable hypotheses, cross-validates evidence, and generates structured conclusions, even for novel, incomplete, or multi-variable scenarios.



Alert-to-Ticket Closed Loop

Ingest signals from your alerting platform, perform deep reasoning, and write traceable, auditable, complete recommendations directly into your ITSM workflow.



LLM & GPU Agnostic Architecture

Swap language models freely as the AI market evolves, start minimal hardware and scale GPU capacity only when you need it. No vendor lock-in at any layer.



Human Feedback & Self-Improvement


Capture engineer corrections and confirmations as a feedback loop. Similar incidents become progressively faster and more consistent to resolve. The platform continues to get smarter from your team's expertise.



Lightweight Integration via MCP, API, A2A

Connect to existing monitoring, SIEM, and ITSM systems through controlled, permission-bound protocols. No data migration required.

Use Cases



Business Transformation You Can Measure



Reduced MTTR

Incident context is assembled and reasoned automatically, eliminating the manual signal-gathering that consumes 60–80% of resolution time for on-call engineers.



MSP Delivery Acceleration

Managed service providers use REASON to compress PoV timelines, reduce delivery risk, and productize incident response practices into repeatable, margin-positive templates.



Consistent & Complete Tickets

Every incident produces a structured, traceable ticket with rationale, evidence, and next steps, that every level of engineers can understand.



Compounding Accuracy Over Time

The self-improvement feedback loop means REASON's accuracy on your specific incident types improves continuously — creating an institutional knowledge asset that grows with every resolved event.



No Infrastructure Overhaul Required

Validate value on 1–2 live incident types before any significant investment. Start small, prove outcomes, then scale.

Other Enterprise AIOps vs REASON

The leading enterprise AIOps platforms were built for a world of abundant cloud, unlimited GPU, and data that flows freely. If your world looks different with regulated data, constrained infrastructure, scarce talent, or sovereign requirements, you need a different architecture.

	Other Big Vendors	Low-Code Platforms	REASON
LLM Dependency	Locked to proprietary AI engine - model upgrades on vendor's timeline.	API-dependent. Tied to OpenAI / cloud model pricing and availability.	Fully LLM agnostic. Swap models freely. No inference lock-in.
GPU Requirements	Requires significant cloud GPU or on-prem GPU farms to run at scale. High capex or variable opex.	Offloaded to cloud. Hidden GPU cost baked into usage pricing.	Minimum GPU investment. Scale only when event volume demands it.
Data Sovereignty	Cloud-first. Data leaves your environment. Compliance and audit controls are limited.	Data flows through third-party automation platforms. Limited governance.	On-prem native. Data never leaves your environment. Built for regulated industries.
Reasoning Depth	Alert correlation and pattern matching. Struggles with novel, incomplete, or multi-variable incidents.	Rule-based logic. Breaks on exception scenarios. No feedback loop.	Tree of Thought deep reasoning. Handles unknown incidents the way a senior engineer would.

CONTACT US

WEBSITE reform.today/
 PHONE +65 9350 2204
 EMAIL enquiry@reform.today / sales@cyberview.com.tw

LEARN MORE

About REFORM reform.today/
 About REASON reform.today/product-reason/
 f Reform.todayFB in reformtoday y Reform.todayYT

Request a **FREE** product demo now!