


Tier-1 Managed Service Provider Transforms IT Operations with Fabrix.ai Agentic AIOps



Situation

A Tier-1 global Managed Service Provider (MSP) operates one of the most complex managed hosting, monitoring, and security environments in the industry — serving 1,900+ enterprise customers across 100,000+ managed assets, supported by 30+ heterogeneous management tools. As scale grew, their home-grown integration and correlation engine could no longer keep pace, creating siloed visibility, swivel-chair operations, and mounting alert noise that threatened service quality and SLA compliance.

Impact

The lack of a unified operational platform drove significant inefficiencies: siloed data blind spots led to missed outages, growing alert and incident volumes overwhelmed engineering teams, and long customer onboarding lead times — driven by high integration costs — slowed new service delivery. The inability to aggregate asset visibility across tools created compliance gaps and operational risk at scale

Resolution

Fabrix.ai's Agentic AIOps Platform replaced the home-grown correlation engine within six months, integrating with 30+ tools to deliver cross-domain correlation, alert noise reduction exceeding 90%, and full-stack service mapping. Building on this AIOps foundation, the provider expanded into Agentic AI — deploying AI Agents for autonomous incident response and root cause analysis. Critically, this Agentic AI deployment runs entirely on-premises against a 120B open-source LLM via Ollama on dedicated NVIDIA GPU hardware, ensuring complete data sovereignty without reliance on external cloud AI services

The client is a Tier-1 global MSP providing managed hosting, monitoring, and security services to 1,900+ enterprise customers with 100,000+ managed assets. Operating across 30+ management tools in a complex multi-cloud and on-premises environment, the provider expanded the Fabrix.ai deployment from its initial managed hosting division into additional internal departments.

The Need

With 30+ management tools in silos and a home-grown correlation engine at capacity limits, the provider needed a unified AIOps platform that could scale without heavy custom integration for every new tool or customer. Beyond consolidation, they needed to dramatically reduce alert noise, accelerate customer onboarding, and build the foundation for autonomous, AI-driven operations — while meeting strict data governance requirements that precluded routing operational data to external cloud AI services.

Business Objectives

- ▶ Replace home-grown correlation platform
- ▶ Reduce alert fatigue and missed outages
- ▶ Accelerate new customer onboarding
- ▶ Enable proactive, autonomous operations
- ▶ Full data sovereignty with on-prem AI inference

Technical Needs

- ▶ Integrate with 30+ existing tools at deployment
- ▶ Full-stack topology and service mapping
- ▶ On-premises LLM with dedicated GPU hardware
- ▶ Extensible platform for new use cases and divisions

The Solution

Fabrix.ai's Agentic AIOps Platform was deployed to replace the existing home-grown platform within six months, integrating with 30+ tools spanning alerts, metrics, logs, inventory, topology, and ITSM. Deployment covered security tooling (ArcSight, Securonix, LogRhythm, Palo Alto, Fortinet), cloud platforms (AWS, GCP, Azure via CoreStack), infrastructure (VMware vSphere, Kubernetes, Docker, Red Hat), compute and storage (Dell, HPE, NetApp, EMC), backup and DR (Commvault, Zerto, Double-Take), and ServiceNow for ITSM and CMDB bidirectional sync. Ansible was integrated for automated remediation workflows.

As a managed services environment serving 1,900+ customers, multi-tenancy was a core platform requirement. Fabrix.ai's native multi-tenant architecture provides strict data isolation per customer, role-based access controls, and a customer self-service portal giving end customers direct real-time visibility into their own assets, incidents, and service health — without requiring MSP engineer involvement for routine queries.

For Agentic AI, the platform was extended with role-based AI Personas, enabling different user types — executives, operations engineers, and customer teams — to interact with AI through experiences tailored to their context and responsibilities. AI Projects allow teams and departments to define scoped AI workspaces with their own data, agents, and objectives. AI Observability provides cost insights, budget tracking, and spend controls across AI workloads, ensuring governance over LLM consumption. Fabrix.ai's dynamic tooling capability enables AI Agents to work directly with third-party ITOM tools even where MCP servers are present, without requiring custom integration per tool. Hyper-connectivity is delivered through a library of 1,900+ pre-built bots, enabling rapid onboarding of new tools, customers, and workflows at scale.

The on-premises GPU inference stack — an HPE physical server with 128 CPUs, 512 GB RAM, and 4× NVIDIA L40S GPUs (180 GB total GPU memory) running a 120B parameter open-source LLM via Ollama, fronted by Nginx over HTTPS — ensures all AI execution remains within the customer's own infrastructure. The success of the initial managed hosting deployment led to expansion into additional internal divisions, extending Fabrix.ai's footprint across the broader organization.

Key Agents

- RCA Agent
- Incident Assignment Agent
- Remediation Agent
- Capacity Insights Agent
- CMDB / SACM Agent
- AIOps Reporting Agent
- Network Config Compliance Agent
- AI Observability Agent

Data Sources & Integrations

- Containers: Kubernetes, Docker
- OS: Windows, Red Hat Linux
- Virtualization: VMware vSphere
- Compute: Dell, HPE
- Storage: NetApp, EMC
- Network: Cisco, F5, HPE
- Security (SIEM): ArcSight, Securonix, Log-Rhythm
- Security (Network): Palo Alto, Fortinet, Arbor, Pearl DDoS
- Cloud: AWS, GCP, Azure (CoreStack)
- Resiliency: Commvault, Zerto, Double-Take, Perpetuuti
- Automation: Ansible
- ITSM / CMDB: ServiceNow (bidirectional sync)
- Monitoring: Prometheus, Zabbix, HPE OneView
- On-Prem LLM: 120B OSS model via Ollama / NVIDIA L40S GPUs

Use Cases Delivered

- Alert correlation, deduplication & noise reduction
- Full-stack service mapping & topology ingestion
- Predictive insights & anomaly detection
- Root cause analysis acceleration (RCA Agent)
- Automated remediation via Ansible integration
- Incident triage & assignment prediction
- CMDB reconciliation and bidirectional sync
- Multi-tenancy with per-customer data isolation
- Customer self-service portal for real-time service visibility
- Role-based AI Personas for executives, ops, and customer teams
- AI Projects for team- and department-scoped AI workspaces
- AI Observability — cost insights, budget tracking & spend controls
- Dynamic tooling — direct integration with 3rd-party ITOM tools
- Hyper-connectivity via 1,900+ pre-built bots in the bot library
- On-prem Agentic AI — 120B LLM on NVIDIA L40S GPUs

Highlights

- Replaced home-grown platform in under 6 months
- 90%+ alert noise and ticket volume reduction
- Faster MTI / MTTR with AI-driven RCA Agents
- Customer onboarding cut from 6 weeks to 2 weeks
- Native multi-tenancy with customer self-service portal
- Role-based AI Personas and scoped AI Projects per team
- AI Observability with spend controls across LLM workloads
- Dynamic tooling enables direct ITOM integration without MCP dependency
- 1,900+ bots for hyper-connected, rapid tool and customer onboarding
- On-prem GPU LLM stack ensures full data sovereignty

Quantified Benefits - \$1M+ in Operational Savings; 100% SLA Compliance

KPI	Result	Industry Benchmark
Alert noise reduction	90%+	75%
Incident resolution time	50% faster	35%
Customer onboarding time 6 wks → 2 wks	6 wks → 2 wks	—
FTE savings	10 FTEs	—
SLA compliance	100%	100%
Platform expansion	Multiple divisions	—

“

CloudFabrix’s AIOps platform has already delivered significant value to Tata Communications. Correlating and deduplicating multiple events from disparate sources and triggering actionable alerts has helped us improve efficiency and productivity of our engineers. We are excited to further explore their new agentic AI capabilities and enhance our operations. — Robert A, Vice President, Tier-1 Managed Service Provider