



AI Adoption Framework for NetOps Teams

A Practical Guide to Piloting and Deploying AI-Driven Observability Workflows

Most NetOps teams are being asked the same question: what are you doing with AI? If your answer is, “we’re not sure where to start,” you’re not alone.

Network observability is a good place to start applying AI because it is central to incident response. With the right foundations in place, AI can help engineers correlate network data, investigate incidents conversationally, and surface evidence without manually pivoting between multiple dashboards.

This guide distills what we’ve learned taking AI-based observability from pilots to production: where it delivers real value, how to gauge your readiness, which guardrails matter, and how to sequence early use cases. It is written for the engineers and operations leaders who will drive the workflows, trust (or reject) the findings, and own the outcomes.

AI Empowers (Not Replaces) NetOps Engineers

Successful practitioners use AI as leverage: extending troubleshooting expertise across the team and reducing manual work. Senior engineers stop wasting cycles proving the network is innocent, L1 and L2 engineers find the root cause with fewer escalations, and application owners get an early diagnosis before the finger-pointing starts.

For NetOps teams we work with, the biggest “A-ha moment” is how AI can shrink the time and scope of incident resolution – from days of investigation across ten teams, to minutes with a handful of engineers. Let’s dive into how you can get there.



STEP 1:

Choose a Problem Worth Solving

AI projects get funded when they're tied to operational or business reasons. Knowing your "why" tells you which outcomes to measure. Nearly every directive rolls up into one of four business imperatives:

- **Protect the digital experience** – faster detection, triage, and root cause for service incidents that threaten your company's revenue, productivity, or customer experience.
- **Multiply scarce expertise** – scale the knowledge of senior engineers across the team and reduce manual toil.
- **Reduce operational complexity** – cut alert noise, unnecessary tool switching, and manual handoffs – all of which simplify and justify the stack.
- **Get ahead of failure** – move from reactive firefighting to proactive, coordinated operations.

With the business outcome in mind, determine the most important problems that AI can solve. Evaluate potential use cases by business impact, data readiness, and risk.

Low Risk [START HERE]

Incident summarization, ticket enrichment, knowledge retrieval, runbook recommendations, alert deduplication. Since AI is assisting investigations, rather than changing production systems, errors in these workflows are generally recoverable.

Moderate Risk

Root-cause hypotheses, change-risk scoring, capacity forecasting, configuration-drift analysis. Require stronger evidence because incorrect recommendations can steer an investigation or change decisions. It's best to earn trust with the low-risk jobs first.

High Risk

Automated remediation, configuration changes, policy enforcement, autonomous response. Treat these as later-stage use cases to deploy once the accuracy is proven, rollback controls, and governance guardrails are in place.

Define the baseline and before you start the pilot by measuring the current workflow – how long it takes, how many people are involved, where escalations happen, etc. Then outline your success criteria as improvements on the baseline. For example:

- Accuracy of AI findings against one or more incident types
- Time saved per ticket
- Share of recommendations your engineers accept
- Rate of false positives/negative (no additional alert noise)

When you're presenting results later, a compelling AI demo proves the technology can produce an answer, but a successful pilot shows measurable outcomes that your leaders will get behind.

Advice from NetOps Engineers

Start investigations with conversations and let the root cause reports surface the answers;

Insist on explainability and evidence, so you can validate, trust, and act on AI findings;

Keep your familiar GUI for the routine and periodic work it still does best.



STEP 2:

Assess AI-Readiness



Before you connect anything, take an honest look at four things:

Coverage – Do you have telemetry data from the environments involved in the use case?

Quality – Is the data accurate, complete, and timely?

Context – Can the system understand topology, dependencies, and baselines?

Workflow – Do insights fit into existing incident, ticketing, or change processes?

While all of these are important, data quality and context are essential. For many network diagnostic use cases, packet-based telemetry provides direct evidence of what happened on the network – complementing flows, metrics, logs, topology, and configuration context.

AI pilots don't necessarily require a centralized data lake. In many cases, a better starting point is controlled access so AI can reason over the data you already trust – with baselines, dependency maps, and discovery layered on top. But be realistic. Simply adding an MCP interface doesn't make you AI-ready. Make sure the right data is being collected and tuned first. Skipping that step is the most common way pilots stall – the model is only as good as its data foundation.

STEP 3:

Choose the Right Deployment Model

When deciding how to deploy AI, there are essentially three models – each with real trade-offs:



User endpoints

Fast and easiest to start but harder to govern and scale, due to security considerations. (ex: Claude Code, Codex)



On-prem centralized

Enterprise-hosted infrastructure provides full control and data sovereignty, but is more of a DIY build.



SaaS

Managed services can accelerate adoption, within your governance constraints. (ex: Microsoft Copilot)

Evaluate deployment models for security, accuracy, data exposure, latency, integration effort, cost, and change management. The acceptable accuracy threshold should rise with the consequence of a wrong answer.

STEP 4:


Ensure Governance and Human Accountability

Governance isn't a brake on the project – it's what lets you expand with confidence. A workable governance model spells out AI operating constraints and defines who is ultimately accountable for it. Align with project stakeholders to set up your governance model before you deploy or scale.

Key things to consider are:

- **Data & Infrastructure** – LLM constraints, data-residency requirements, tool handoffs, integration
- **Operational** – Workflows, approvals and escalations, rollback and disablement plans, change management
- **Security & Compliance** – Access controls, evidence retention and audit logs (required for HIPAA or PCI-DSS)

Humans should be at the center of early AI projects. For every AI-assisted workflow, name an owner for each of these:



Evidence, traceability, clear ownership, and defined approval points are what turn a promising pilot into an operational capability the organization can trust at scale.

→ How cPacket Applies the Framework

The good news for cPacket customers is that you don't need to rip and replace anything to add AI.

cPacket AI is enabling a new operating model for enterprise IT operations – built on trusted evidence, scaled through AI, and designed for the way network teams actually work. By applying AI reasoning to existing packet data, and equipping LLMs with deep troubleshooting skills, NetOps teams can get answers faster, without sacrificing accuracy or control. What's more, cPacket AI operates within a customers' environment to maintain data sovereignty and GRC compliance – even in air-gapped networks.

Packet Evidence

- Most accurate record of what happened on the wire
- Verify every AI finding with existing dashboards & PCAPs

cPacket MCP Server

- Connects AI platforms to network observability data and analytics (read-only)
- Runs within your environment – data stays private

Your LLM of Choice

- Use your enterprise-approved LLM to meet corporate AI, security, and compliance standards
- Includes support for cloud-based proxies, such as Azure CoPilot

cPacket Skills

- Encodes network knowledge and troubleshooting skills in generic models
- Repeatable queries and investigation patterns for 13 common incidents to ensure consistency for all users

Global finance, healthcare, and technology enterprises and government agencies rely on cPacket to ensure reliability, security, and business continuity in the world's fastest and most demanding networks.

To see what cPacket AI can do for you, visit www.cpacket.com/cpacket-ai.