



RAPAX 2.0 • POSITION PAPER

The license is the cheapest part.
Enterprise-grade assurance no longer requires
an enterprise-grade operations team.

NOC-in-a-Box

How Rapax 2.0 Takes the Operating Cost Out of Service Assurance

One platform for network and infrastructure monitoring, hosted for you, and built so one NOC manager runs it in a couple of hours a week.

Shawn Ennis

Founder & CEO, Rapax

October 2026

Try it on your own network

Free 14-day evaluation, hosted for you.

rapax.app/evaluate

Summary

Most organizations don't fail at monitoring because their tools lack features. They fail because running the tools takes expertise they can't hire. In 2026, 52% of network teams struggle to hire and keep network engineers, up from 26% in 2022, and only 31% call their network operations strategy completely successful (EMA, 2026).

Rapax 2.0 was built for that reality. It installs from one script, updates in one click, and puts log collection, polling, availability, correlation, ticketing, SLA reporting and capacity forecasting in one platform. Our design target is that one NOC manager runs it in a couple of hours a week.

The rest of this paper shows what the old way costs, what changed in Rapax 2.0, and how to test it on your own network in 14 days, free.

I. The license is the cheapest part

I have spent 25 years building service assurance platforms. I founded a company that built one, sold it, and then watched customers run it. The problem I saw most often had nothing to do with features.

Organizations bought capable tools. Then they found the tools needed people they didn't have: engineers who understand polling, thresholds, correlation rules, upgrades, and the hundred small decisions that keep a monitoring platform trustworthy. The license was the cheapest part of the purchase. The expertise to run it was the expensive part, and usually the missing part.

Without that expertise, I watched the story end one of four ways:

- 1. A permanent struggle with vendor support.** Every upgrade, new device type and broken integration becomes a ticket. The vendor's support desk becomes the customer's operations team, and neither side wants that.
- 2. An expensive outsourcing agreement.** A third party takes over the tools. It works, at a price that rises every year, and the knowledge of how your network is monitored lives in someone else's building.
- 3. Quiet failure.** Upgrades slip, alerts get muted, and the team falls back to substandard monitoring: ping checks, an inbox of emails, and outages reported by users.
- 4. Your own people, pulled off their real jobs.** This is where most organizations end up by default.

Tool administration has always been a distraction for most teams and a necessary evil for the rest. So ask a direct question. Do you want your developers, who never signed up to learn assurance, or your operators, who should be running the network, spending their week keeping the monitoring platform alive?

We built Rapax 2.0 so they don't have to. Our design target is that a NOC manager runs the tools in a couple of hours a week, and the rest of the team's time goes back to the network.


II. The people problem is getting worse

The expertise gap is widening, not closing. Enterprise Management Associates (EMA) has asked network teams for years how hard it is to hire and keep network engineers. In 2022, 26% said it was difficult. In 2024, 41% did. In 2026, 52% do ([EMA Network Management Megatrends 2026](#)).


The expertise gap is widening

Share of network teams, EMA Network Management Megatrends survey

Struggle to hire and keep network engineers



Rate their network operations strategy a complete success



Sources: EMA Network Management Megatrends 2026 (n = 352); TechTarget and NetScout summaries of the same study.


Figure 1. More teams struggle to hire, and fewer call their operations a success.

The shortage is worst at the senior and mid-career levels, where monitoring expertise lives. The hardest skills to find are network security, networking for AI, and network automation ([NetScout summary of EMA 2026](#)).

Skills gaps also block the obvious way out. Asked what stops them automating day-two operations, 46% of teams named skills gaps inside the team. That ranks ahead of tool limitations (36%), data quality (32%) and budget (30%). The people who could automate the work are the people teams can't hire.

Skills, not tools, block automation

What stops teams automating day-two network operations (share naming each barrier)



Source: EMA Network Management Megatrends 2026, via NetScout summary. Respondents could name more than one barrier.

Figure 2. The barrier to automation is the same scarce expertise.

The outcomes follow. Only 31% of organizations call their network operations strategy completely successful, down from 42% two years earlier. Only 32% are completely satisfied with their monitoring tools ([TechTarget on EMA 2026](#)).

A note on the data: EMA's sample of 352 IT professionals leans toward large organizations, with more than 60% at companies of over 5,000 employees. Mid-market teams are smaller, so one departure takes a larger share of their expertise with it. We expect the gap there is wider, but that is our inference, not EMA's finding.

The usual answer is to hire your way out. The numbers say that road is narrowing. The other answer is to need less of the scarce expertise in the first place.

III. What the operating cost actually is

The labor to run monitoring costs more than the software. The U.S. Bureau of Labor Statistics puts the median wage for network and computer systems administrators at \$99,130 ([BLS, May 2025](#)). Benefits are 29.9% of private-sector compensation ([BLS ECEC, December 2025](#)), which makes the loaded cost about 1.43 times the wage.

That puts one administrator at roughly \$141,000 a year, before recruiting, training, tools or management. Now cover one seat around the clock:

Assumption	Value	Basis
Hours to cover per week	168	24 × 7
Hours per person per week	40	Standard week
People per 24×7 seat	4.2–5.0	168 ÷ 40, plus about 20% for leave and training
Loaded cost per person	~\$141,000	BLS median × 1.43
Cost of one 24×7 seat	~\$594,000–\$707,000 a year	Our model


That is one chair, staffed around the clock, with nobody senior behind it. It is why most mid-market organizations don't have a NOC at all.

Then add the care and feeding of the tools. The typical IT organization runs 4 to 10 network monitoring and troubleshooting tools, a number that has barely moved in a decade. Troubleshooting takes 29% of a network professional's day. Manual administration errors cause 28% of network problems, and only 37% of alerts point to a real problem (EMA, 2026).

Every tool in that stack has its own upgrades, certificates, backups, integrations and credentials. We estimate the administration of one self-hosted monitoring platform at 0.25–0.5 of a full-time person, or about \$35,000–\$71,000 a year in loaded labor. That is our estimate from field experience, not a survey figure, and it often exceeds the subscription.

What it costs to run monitoring, per year

Loaded labor only; software licenses excluded. Ranges show the low and high end of each estimate.



Sources: BLS median wage (\$99,130, May 2025) × BLS benefits share (29.9%, Dec 2025) = ~\$141K loaded. 24×7 seat and self-hosted admin labor (0.25–0.5 FTE) are our models. Rapax: 2 hours a week × 52 at the same loaded rate; a design target, not a measurement.

Figure 3. Running the tools, not buying them, is where the money goes.

This is what we mean by “the license is the cheapest part.” For a network under 1,000 devices, a year of Rapax costs less than one fully loaded administrator. The real question is not what the software costs. It is how much of your team the software consumes.

IV. The cost of getting it wrong

Running short-handed isn't free either. It moves the cost from payroll to outages.

Short-handed monitoring moves the cost to outages

Self-reported survey figures; treat them as a range, not a price tag



Figure 4. The cost of short-handed monitoring shows up as outages.

The tails are long: 41% of firms put an hour of downtime at \$1 million to over \$5 million ([ITIC 2024 Hourly Cost of Downtime Survey](#)), and one in five say their most recent major outage cost more than \$1 million ([Uptime Institute Annual Outage Analysis 2026](#)). The 73% figure on ignored or suppressed alerts comes from [Splunk State of Observability 2025](#).

These downtime figures are self-reported estimates, and ITIC notes that many organizations never calculate their own. Treat them as a range, not a price tag.

The pattern underneath them is not. When the tools are noisy and the people are stretched, real alerts get ignored. A muted alert is the moment substandard monitoring stops being cheap.

V. Why assurance got stuck

Most assurance platforms in service today run on 1990s technology carrying 2000s concepts. I know, because I built some of them. They were designed for carriers and the largest enterprises, and they assumed a dedicated platform team behind every install.

You can see that assumption everywhere in them:

- **Installation is a project.** It takes consultants, weeks, and a statement of work.
- **Upgrades are events.** They get scheduled, feared, and postponed, so patches lag for months.
- **Administration lives at the command line.** Config files, scripts, and tribal knowledge stand in for a user interface.
- **Root cause needs a specialist.** Correlation rules are written by the one person who knows the topology.
- **Every function is a separate product.** Logs, polling, tickets and reports each come with their own license, upgrade cycle and admin.

Those assumptions made sense for an operator with a hundred-person operations group. They fail a 40-person IT department, or an MSP watching 30 customers with five engineers.

The market knows it. 73% of organizations are at least somewhat likely to replace a monitoring tool within two years. 55% now call AI features a requirement, and 54% name AI-driven insight and automation as the top reason to replace a tool (EMA, 2026).

What changed is that capabilities once built only for the largest operators can now run on one server, for a network of a few hundred devices. The question is what to demand of them.

VI. The six-point test for modern assurance

Assurance should meet the same bar as the best enterprise platforms you already run. Score your current monitoring against six tests:

Test	What passing looks like	Your platform today
1. Cloud-native	Deploys from one script on any cloud VM or server, and upgrades in place	
2. Self-maintaining	One-click updates and patches that can be automated, with a snapshot before every change and rollback if it fails	
3. Point-and-click administration	Users, backups, certificates, components, collectors and reports managed from the browser	
4. AI-driven self-help	Answers grounded in your own runbooks and documents, with sources; no action without a person confirming	
5. Predictive, not just reactive	Baselines, anomalies and capacity forecasts that say what will break, and when	
6. Predictable at scale	Known capacity per server size, and collectors at every site without new infrastructure	

A platform that fails three or more of these is pushing the operating cost from Section 3 onto your team. Ten years ago, passing all six took a carrier's budget. That changed.

VII. What changed in Rapax 2.0: wider and simpler

Rapax 2.0, generally available since October 10, 2026, does two things at once. It covers more of the monitoring job, so you run fewer tools. And it takes less work to run, so the tools stop consuming your team.

Wider: one platform for network and infrastructure monitoring

Capability	What Rapax 2.0 does	Without it, typically
Log collection	Syslog and SNMP trap receiver, search, live tail, export	A separate log tool
Ping and SNMP polling	SNMP v2c and v3, latency and loss, a discovery wizard for new devices	A separate monitoring tool
Availability and SLA	Availability per device and service, SLA targets, error budgets, scheduled reports	A reporting tool or spreadsheets
Correlation and root cause	Parent/child, topology and service correlation, with root cause one click away	An event-management tool
Service impact	Maps every alert to the business services it affects	Often not done at all
Performance analytics	Baselines, anomaly detection and capacity forecasts	A separate analytics tool
Ticketing	Native incident management, a customer portal, an AI tier-1 agent, plus Zendesk, osTicket and Salesforce connectors	An ITSM tool
Configuration management	Collects device configurations, keeps their history and shows what changed, alongside inventory and discovery with duplicate review	A separate config tool or CMDB
Knowledge	AI answers grounded in your own runbooks, with sources	A wiki nobody reads
Multi-tenancy	Scoped views per customer or business unit (SOCs)	Separate instances

One caution, because the research supports it. EMA finds that the number of tools alone doesn't predict success, and that successful teams integrate rather than consolidate. The case for one platform is not fewer logos. It is one thing to administer, and correlation across logs, polling, alerts and services in one place, which separate tools can't do. Rapax integrates with the systems it doesn't replace.




Figure 5. Performance analytics: four devices over 24 hours with thresholds overlaid. Point and click, no query language.

Simpler: built so a NOC manager runs it in a couple of hours a week

- **One-script installer.** One shell script prepares the server and installs the version you choose.
- **Self-updating engine.** Install, update, upgrade and patch scripts are delivered by a containerized toolbox that updates itself.
- **One-click updates that can be automated.** Every change takes a snapshot first, verifies the result, and offers rollback if it fails.
- **Point-and-click administration.** Updates, backups, certificates and components run as audited, allowlisted jobs launched from the browser.
- **Certificates and backups handled.** Let's Encrypt certificates renew themselves; full backups, verification and restores run from Settings, with daily data snapshots kept 30 days.
- **Remote collectors.** A collector at each site connects over TLS and stores up to 24 hours of data if the link drops.
- **Guided setup.** A discovery wizard finds devices, and nothing is monitored until you review and approve it; a split-tunnel WireGuard VPN, set up from the browser, links Rapax to your network, or a remote collector runs inside it.
- **One modern interface.** Light and dark mode, a command palette, and the AI assistant on every page.

In a typical week, running Rapax 2.0 means:

- Applying updates and patches with one click
- Checking that backups and certificates show green
- Approving newly discovered devices
- Tuning a noisy threshold or rule
- Reading the SLA and capacity reports

Those hours cover running the tools, not working incidents. Incident work is what the tools exist to make faster.

Rapax 2.0 passed all 46 of its test suites at release: 4,475 checks, zero failures.




Figure 6. Updates and patches run from the browser, with live output and a full job history.




Figure 7. Device discovery: new devices flagged and a duplicate caught automatically. Nothing is monitored until you approve it. Internal hostnames blurred.

VIII. NOC-in-a-box: one incident and one forecast

NOC-in-a-box means everything real-time operations needs, in one hosted platform: monitoring, correlation, service impact, ticketing, notifications, knowledge and reporting, plus the AI agents and onboarding hours to put it to work. Your team still makes the calls. Rapax does the watching, sorting and paperwork, so fewer people with less specialist training can run a NOC.

Here is what that looks like when something breaks. A switch uplink fails at a branch office:

1. The site's remote collector reports it to Rapax over TLS.
2. Correlation nests the downstream alerts under one parent. In the example below, a router outage and its two failed interfaces show as one critical problem, not three separate alerts.
3. The alert drawer names the business service that is down and the device alerts behind it.
4. The knowledge assistant answers "what do we do?" from your own runbook, citing it.
5. The AI tier-1 agent updates the ticket and the customer portal.
6. An administrator confirms the fix and closes the loop: no command line, no specialist.

The screenshot shows the Rapax console interface. At the top, there's a navigation bar with 'Console / Operate / Alerts' and a search bar. Below this, a summary bar shows counts for various alert severities: Critical (1), Major (2), Minor (0), Warning (0), Information (0), Clear (1), Unknown (0), Ack'd (0), Assigned to me (0), and Unassigned (4). A table below lists the alerts. The first alert is 'Critical' (5s ago, remote-router, Device), with a description 'remote-router is unreachable'. It is expanded, showing two nested 'Major' alerts: 'GigabitEthernet0/1 on remote-router is down' and 'GigabitEthernet0/0 on remote-router is down'. A fourth alert, 'Cisco trap rules consumer running', is also visible.

Severity	Age	Device	Interface	Description	Category	Count	Assigned	Last occurred
Critical	5s	remote-router	Device	remote-router is unreachable	DeviceUpDown	1	-	2026-10-14 14:28:26
Major	3s	remote-router	GigabitEthernet0/1	GigabitEthernet0/1 on remote-router is down	InterfaceUpDown	1	-	2026-10-14 14:28:28
Major	3s	remote-router	GigabitEthernet0/0	GigabitEthernet0/0 on remote-router is down	InterfaceUpDown	1	-	2026-10-14 14:28:28
Clear	3s	cisco-trap-1	RulesConsumer	Cisco trap rules consumer running	Heartbeat	1	-	2026-10-14 14:28:28

Figure 8. One problem, not three: a router outage with its two failed interfaces nested beneath it.

The screenshot shows the 'Ask the KB' interface. It has a search bar and a list of questions. The selected question is 'How does alert correlation pick a parent alert?'. The answer is detailed, explaining the 'Parent Alert Selection' process, including 'Device Correlation' and 'Interface Correlation'. It also mentions 'Deduplication' and 'Status Filter'. The answer is sourced from 'rapax-help-correlation.html' and 'rapax-core-correlation.md'. A note at the bottom states 'Not Covered by the Knowledge Base'.

Ask the KB

Wade answers from the knowledge base (product documentation, runbooks and documents your team uploaded) inside your span of control, and lists the sources he used.

Wade is online. Answers are saved to your Assistants history.

Ask a question, for example: how does alert correlation pick a parent?

Sources: knowledge base (no web search) + New conversation Ask

How does alert correlation pick a parent alert? How do I add devices to discovery? How do I roll back an update? Which ports does Rapax use?

How does alert correlation pick a parent alert?

Wade, from the knowledge base

Parent Alert Selection

Parent alerts are selected by the correlation type. The knowledge base documents cover device and interface correlation.

Device Correlation

A device-level alert becomes a parent when all of the following are true:

- State is "Down"
- Interface is "Device"
- Category is "DeviceUpDown"

Its children are alerts with State = "Down", the same Device, and a different UUID from the parent. (Source: rapax-help-correlation.html)

Deduplication: If multiple DeviceUpDown:down alerts exist for the same device (DeviceKey), only the most recent one is used as the parent. This prevents older, possibly outdated alerts from competing with newer ones. (Source: rapax-core-correlation.md)

Interface Correlation

An interface-level alert becomes a parent when all of the following are true:

- State is "Down"
- Interface is "Device"
- Category contains "UpDown"

Its children are alerts with State = "Down", the same Device and Interface as the parent, and a Category that does not contain "UpDown". (Source: rapax-help-correlation.html)

Status Filter (Device and Interface)

Both device and interface correlation require the parent to have an active status: Critical, Major, or Minor. Cleared or informational alerts are never used as parents. (Source: rapax-core-correlation.md)

Not Covered by the Knowledge Base

Figure 9. Every answer cites its source, and says what the knowledge base doesn't cover.

And here is what it looks like before something breaks. The forecast agent fits a trend to each metric every hour. When a server's disk sits at 88.9% and grows 0.16 points a day, Rapax doesn't wait for 100%. It raises an alert that says the disk crosses capacity on December 13, in 59 days.




Figure 10. Capacity problems arrive as a date, not a surprise.

Some problems never cross a threshold. The anomaly agent learns a week of normal behavior for each metric and flags what breaks the pattern. When a router's ping latency moved from a normal average of 23 ms to 49 ms, Rapax raised a warning before any fixed threshold would have fired, and showed its work on the alert: the methods it used, 5 anomalous points out of 46, and a 111% deviation.




Figure 11. Anomaly detection learns what normal looks like and shows the evidence for every alert.

One example looks back to a root cause, one catches a pattern no threshold would, and one looks ahead to a date. None needed an expert to set up.

IX. What it means for you

IT and network operations teams. You get the coverage of several tools without the administration of several tools. The hours that used to go to upgrades, certificates and broken integrations go back to the network. For a growing team, that can mean a hire you don't need to make; for a stretched one, it means the projects that keep slipping.

MSPs and partners. Managed services are growing fast: the channel was forecast to reach about US\$595 billion in 2025, with roughly 341,000 partners delivering (Canalys/Omdia). Rapax 2.0 is built for that model. One platform serves many customers, each scoped to its own SOC, with its own reports, knowledge base and customer portal. Remote collectors reach customer sites without new infrastructure at each one.




Figure 12. Branding per customer, with live light and dark previews and built-in contrast checks.

Existing Rapax customers. Rapax 2.0 is a new installation, not an in-place upgrade from 1.x. Version 1.x remains maintained.

X. The honest numbers

There is no single savings number in this paper, on purpose. Here is every figure we used, where it came from, and how much weight it can bear.

Figure	Value	Source	Type
Teams struggling to hire network engineers	52% (2026), up from 26% (2022)	EMA	Survey

Figure	Value	Source	Type
Network operations completely successful	31%	EMA	Survey
Alerts that point to a real problem	37%	EMA	Survey
Network problems caused by manual admin errors	28%	EMA	Survey
Loaded cost of one administrator	~\$141,000 a year	BLS wage × BLS benefits share	Government data + our arithmetic
Cost of one 24×7 seat	~\$594,000–\$707,000 a year	Our model, Section III	Modeled
Admin labor for one self-hosted platform	~\$35,000–\$71,000 a year	Our model, Section III	Modeled (our estimate)
Weekly admin time with Rapax 2.0	About 2 hours	Rapax	Design target, not measured
Rapax admin labor at the design target	~\$7,000 a year	2 hours × 52 weeks at the ~\$141K loaded rate	Modeled from a design target
Traditional time to first alert	12–26 weeks	Our estimate from field experience	Modeled (our estimate)
Downtime cost, mid-size and large firms	Over \$300,000 an hour for 90%+	ITIC 2024	Self-reported
Release quality at GA	46/46 suites, 4,475 checks, 0 failures	Rapax	Our data

Where the time goes. If our admin-labor estimate fits your team, moving from 10–20 hours a week of tool administration to our two-hour design target returns 8–18 hours a week, or 0.2–0.45 of a person. Treat that as something to test during your evaluation, not a promise. We expect most teams to redeploy that time to projects and incident follow-up rather than reduce headcount. For a growing team, it shows up as a hire not made.

On AI. Gartner finds only 28% of AI use cases in infrastructure and operations fully pay off, and that success comes from building AI into existing workflows ([Gartner, April 2026](#)). That is why Rapax’s AI works inside the console, the ticket and the knowledge base, answers from your documents, and takes no action without a person confirming it.

About the numbers. Survey figures come from named third parties; EMA's 2026 sample leans toward large enterprises. Downtime costs are self-reported estimates. Staffing and admin-labor figures are our models, with every assumption shown. Rapax figures come from our own release testing; the two-hour week is a design target, not a measurement.

XI. Objections, answered

"Cheaper means less capable." It's the same code our largest deployments run. The price came down because the operating burden came down, not the feature set. The evaluation exists so you can check that on your own network.

"You're a small vendor. Will you be here in three years?" I've built and sold an assurance company before. Commitments are quarterly, not multi-year. You can back up your data and your Rapax configuration from the browser and download both at any time, so you are never locked in.

"Hosted means our monitoring data leaves our network." Collectors connect out over TLS, each limited to its own data, and nothing opens inbound to your sites. Secrets are never stored in plain environment variables. If you'd rather run it yourself, the same one-script installer works on your own server. Hosted systems run in the region you choose: the United States, Canada, Mexico or Europe. Every customer runs on dedicated hosting, so no data is shared between Rapax customers.

"AI will make things up, leak our data, or run up a bill." You choose the AI provider, including Ollama on your own network, so you control its data terms. Answers come from your own documents and say when those documents don't cover the question. Token use and cost are tracked per assistant and model, and nothing changes without a person confirming it.

"What about high availability and scale?" Each hosted system runs on its own server, sized to your device tier, with daily snapshots and full backups you can verify and restore from the browser. Remote collectors keep buffering if the central server is unreachable, and replay in order when it returns.

"We run ServiceNow / Jira Service Management." For network operations, you may not need them. Rapax includes native incident management, a customer portal and an AI tier-1 agent, so operational tickets open, update and close without a separate ITSM tool. Zendesk, osTicket and Salesforce connectors ship today for teams that keep one.

"We already have tools." Keep the ones that work. Rapax correlates across logs, polling, alerts and services, and integrates with what it doesn't replace.

XII. Getting started: days, not quarters

Doing more with less is the new reality for technology teams. The way you start with a monitoring platform should reflect that.

What you need to start

- 1. Your size.** How many devices do you monitor? Pick the tier that fits today; you can move up at any time.
- 2. Your devices.** A list of management IP addresses or DNS names, or the CIDR ranges you want discovered.
- 3. An AI provider.** Ollama on your own network, or OpenAI, Anthropic Claude, Gemini or Grok. You pick the provider on the form and connect it at first sign-in.
- 4. An email account for notifications.** Your mail provider and a sending account, which you also enter yourself at first sign-in.
- 5. Your payment preference.** ACH or credit card, for when you decide to keep it. Nothing is charged during the evaluation.

The request form never asks for passwords or API keys. Those go straight into your own system, encrypted, entered by you.

What happens next

Submit the request, and within two business days you receive a login to your own Rapax system. From there:

- 1. Follow the Get Started guide** to add devices, define services and set up notification profiles.
- 2. Connect your network** with a split-tunnel WireGuard VPN set up from the browser, so Rapax can poll your devices. If your firewall supports only IPsec or OpenVPN, a remote collector inside your network does the same job.
- 3. Ask the built-in assistant when you're stuck.** It walks you through setup and day-to-day tasks, and cites the documentation it draws on.
- 4. See your first correlated alert** on your own devices as soon as they're connected. Rapax starts polling, correlating and building baselines right away.
- 5. Decide by day 14.** Converting is a license install: everything you built stays, and your 10 setup hours carry over.

The alternative

The traditional path to an assurance platform looks like this:

- 1.** Get IT approval for new servers or a VPS.


2. Evaluate and hire a systems integrator to install it.
3. Hire a project manager to run the program.
4. Run a lengthy requirements analysis and design phase.
5. Write, issue and score an RFP.
6. Take the funding request to the board.
7. Hire and train a team of administrators to keep it running.

Side by side

	Rapax NOC-in-a-box	The traditional path
Approval to start	A request form	Infrastructure approval, an RFP and a board funding request
Who installs it	We do; you get a login	A systems integrator you evaluate and hire
Who runs the project	You, with the Get Started guide	A dedicated project manager
Requirements and design	Discovery and guided setup	A lengthy analysis and design phase
Who keeps it running	One NOC manager, a couple of hours a week (our design target)	A team of administrators you hire and train
Time to first alert	As soon as your devices are connected	After approval, funding, hiring and installation
Spend before you see value	Nothing during the 14-day evaluation	Integrator fees, a PM's salary, RFP effort and new hires

Time, effort and value: the two paths compared

Traditional path vs. Rapax NOC-in-a-box



Traditional time to first alert and admin hours are our estimates from field experience (0.25–0.5 FTE). Rapax login within two business days; first alerts as soon as devices connect. Rapax hours are a design target, not a measurement. Test both during your 14-day evaluation.

Figure 13. Faster to start, less to run, and the time comes back to your team.

That difference is the point of NOC-in-a-box. The traditional path spends months and a team before the first alert. Rapax gives you the alert first, and lets you decide.

Request your free 14-day evaluation at rapax.app/evaluate.

Want to talk it through first? Book 15 minutes at cal.com/shawn-ennis.

About the author

Shawn Ennis is the founder and CEO of Rapax. He has spent 25 years building network service assurance platforms, holds 12 patents in network management and service assurance, and previously founded Assure1, acquired by Oracle in 2021. He co-hosts the Transformation Leaders Podcast.

Sources

- [EMA Network Management Megatrends 2026, press release](#)
- [Network World on EMA 2026](#)
- [NetScout summary of EMA 2026](#)
- [TechTarget on EMA 2026](#)
- [BLS Occupational Outlook: Network and Computer Systems Administrators](#)
- [BLS Employer Costs for Employee Compensation, December 2025](#)
- [ITIC 2024 Hourly Cost of Downtime Survey](#)
- [Uptime Institute Annual Outage Analysis 2026](#)
- [Splunk State of Observability 2025](#)
- [Gartner, AI projects in I&O, April 2026](#)
- [Canalys/Omdia MSP Trends and Predictions 2025](#)