KORE-TEK STRATEGIC PERSPECTIVES, NO. 2

The Co-Managed NOC

Rethinking the Sourcing Spectrum for Enterprise Network Operations

WRITTEN FOR

CTOs, CIOs, and senior infrastructure leaders evaluating where their network operations should sit on the modern sourcing spectrum.

PUBLISHED BY

Kore-Tek | Optical Engineering | Network Infrastructure | Managed NOC Services

kore-tek.com/insights

Executive Summary

For two decades, enterprise leaders have framed the network operations question as binary: build it ourselves, or hand it to a managed services provider. That binary no longer reflects how the market actually works, and the cost of pretending it does is becoming visible in budgets, talent retention numbers, and incident response times.

Three structural shifts have collapsed the old framing. The platforms underlying NOC operations have commoditized, with toolsets now available on subscription. AIOps has matured to the point that alarm correlation, ticket enrichment, and incident summarization are no longer differentiators but baseline expectations. And the network engineering talent market has tightened to the point that hiring 24x7 coverage at acceptable quality has become structurally difficult for most enterprises.

The old either-or choice is gone. Instead of picking between building a NOC in-house or handing it entirely to a provider, organizations now have five operating models to choose from, ranging from fully in-house at one end to fully provider-delivered at the other. Where any given organization should sit is shaped by five factors: operational maturity, talent strategy, compliance posture, cost structure, and required time-to-value.

This paper provides a framework for that decision. It is written for senior infrastructure leaders evaluating where their network operations should sit, and what would have to be true to move from one position to another. The framework is deliberately model-agnostic: the goal is alignment, not advocacy for any single approach.

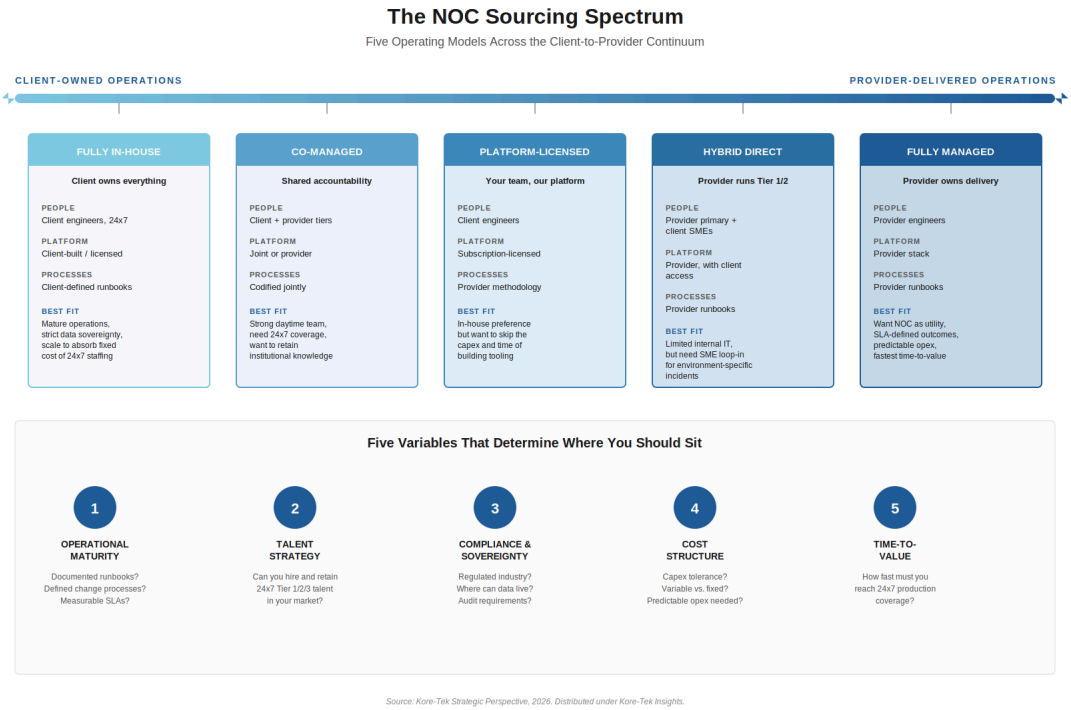

Figure 1: The NOC Sourcing Spectrum and the five variables that determine sourcing alignment.

The Build-vs-Buy Binary Has Quietly Disappeared

The old framing made sense when it emerged. Building a NOC required hiring engineers, selecting and integrating monitoring tools, designing escalation processes, and writing runbooks, each of which took years. Outsourcing meant handing all of that to a provider in exchange for an SLA. The two were genuinely different, and you really had to choose one or the other, because the people, the tools, and the processes all came bundled together.

That coupling has come apart. Three forces have driven the unbundling.

First, NOC platforms have become off-the-shelf products. The integrated monitoring, ticketing, and automation tools that once required years of internal development are now available as subscription services. The market reflects this: industry analysts size the global NOC-as-a-Service market at $3.67 billion in 2025, growing at a 10.2% compound annual rate to $8.83 billion by 2034.[1] More tellingly, hybrid deployment models , those that don't fit neatly into “in-house” or “outsourced” , are growing at over 8% CAGR.[2]

Second, AIOps has matured. What was experimental three years ago is now operational. Gartner projects that generative AI will account for over 25% of initial network configuration by 2027, up from less than 3% in 2024.[3] Mid-market deployments report 20–40% alert volume deflection after AIOps

correlation tuning.[4] This matters for the sourcing question because it dissolves one of the historical justifications for full outsourcing , the idea that only a provider could afford the tooling investment to operate efficiently. That is no longer true.

Third, and most consequential, the talent equation has shifted. Industry analysts project more than 1.2 million unfilled network engineering roles globally in 2025.[5] Bureau of Labor Statistics data shows network engineering employment growing 12% from 2024 to 2034, faster than the average for all occupations, yet supply is not keeping pace.[6] The Linux Foundation's 2025 State of Tech Talent report found that 65% of organizations are understaffed in cybersecurity and compliance roles, with similar gaps in cloud and platform engineering.[7] At the operational tier, Foote Partners' Q3 2025 data shows declining premiums for Level-1 NOC monitoring as AIOps automates routine work, while architecture and engineering roles command an 18% pay premium and remain difficult to fill.[8]

Taken together, these forces have done something specific: they have separated the people, the platform, and the processes. An enterprise can now choose any combination, its own engineers operating a licensed platform with provider-defined methodology, or a provider's engineers operating a client-defined process. The combinations are real, the market is delivering them, and the old either-or question no longer makes sense.

The Sourcing Spectrum: Five Operating Models

A more useful way to think about NOC sourcing is as a continuum across three axes, who owns the people, who owns the platform, and who owns the processes, and as a position along that continuum where each combination becomes a coherent operating model.

Five such models are recognizable in the market today. They are not arbitrary; each represents a stable equilibrium where the ownership of people, platform, and processes aligns with a particular set of operational and strategic conditions.

1. Fully In-House

The client owns the engineers, the technology, and the operating procedures. Network operations are treated as a strategic capability, not a service to be purchased. This model fits organizations large enough to absorb the fixed cost of round-the-clock staffing, with the operational maturity to define and govern their own procedures, and with compliance or data-residency requirements that rule out third-party involvement. It is also the most expensive option, the slowest to scale, and the most vulnerable to network engineering talent shortages.

2. Co-Managed

Accountability is shared. Typically, the client retains daytime operations and senior engineering ownership while a provider extends coverage into off-hours, weekends, and surge events. The platform may be jointly operated. Runbooks are codified through a structured handoff process. This model fits organizations with a strong daytime team that needs 24x7 coverage without doubling the headcount, and that strives to retain institutional knowledge while accessing provider scale.

3. Platform-Licensed (NOC Platform-as-a-Service)

The client's engineers operate the technology, but the technology itself is licensed from a provider. It comes pre-built with the operating procedures, dashboards, alert filtering, and automated workflows, refined across many environments, but the people running it are the clients. This model has carved out its own place in the market because it lets organizations keep their operations team in-house while avoiding the multi-year capital investment of building the technology from scratch.

4. Hybrid Direct Engagement

The provider handles the front-line monitoring and initial response work, while the client's specialists step in directly when an incident requires deep knowledge of the client's specific environment. Rather than escalating problems through tickets and email, the client's experts work side-by-side with the provider's team inside the same platform. This model fits organizations with smaller internal IT teams, especially those that don't run their own ticketing system, but that still need their senior people involved when something complex breaks.

5. Fully Managed

The provider runs everything from end to end: the people, the technology, and the processes. The relationship is governed by service-level agreements with measurable performance targets, not by joint operations. This model fits organizations that want their NOC to work like a utility, predictable monthly cost, defined results, fastest path to 24x7 coverage, and that are willing to trade hands-on operational control for that simplicity.

These five models are not ranked. None is better than the others on its own merits. Each is the right answer for a specific set of circumstances, and the most common mistake organizations make is choosing a model that doesn't match their actual situation.

The C-Suite Decision Framework

Choosing where to land among the five models is a decision with several moving parts. The five variables below are, in our experience, the ones that actually drive the right answer. Gut feelings or institutional habits can override them in practice, but they should not.

Variable 1: Operational Maturity

The honest test is whether the operation could survive the loss of any single person. If the answer requires consulting the senior network engineer's email archive, the maturity is lower than the org chart suggests. Mature operations have documented operating procedures, defined change-management processes, measurable performance targets, and a current inventory of every system and how it connects to the others. Less mature operations do not, and Uptime Institute analysis identifies configuration and change management failures as the cause of 45% of network outages, with human error involved in 66–80% of all downtime incidents.[9] Less mature operations should lean toward Hybrid Direct or Fully Managed; more mature operations have viable options across all five models.

Variable 2: Talent Strategy

Two questions matter here: can the organization hire and retain the network engineering talent it needs in its market, and is doing so the right use of that talent? The 1.2 million-role global shortfall is not evenly distributed; some markets and some specializations are dramatically harder than others.[5] Organizations with a strong daytime team but no realistic path to 24x7 staffing are natural candidates for Co-Managed. Organizations whose engineers are specialized to their environment but spend their nights triaging Tier 1 alerts are misallocating expensive talent, a condition that points toward Hybrid or Fully Managed.

Variable 3: Compliance and Data Sovereignty

Regulated industries, financial services, healthcare, government, face structural constraints on where data can live and who can touch it. Some constraints rule out specific delivery models entirely; others require contractual structures that some providers cannot accommodate. The decision here is not philosophical but contractual: the question is what your auditor will accept, not what your engineering team prefers.

Variable 4: Cost Structure

Building a 24x7 NOC requires capex (tooling, platform integration, facility) and creates fixed opex (24x7 staffing, training, retention). Subscription models convert that into variable opex. The right structure depends on the organization's preference for predictability versus flexibility, its ability to absorb capital expenditure, and the rate at which its operational footprint is changing. Organizations growing rapidly or undergoing M&A often benefit from variable-cost models; stable, mature operations may not.

Variable 5: Time-to-Value

Building a NOC in-house takes 18 to 36 months from decision to a smoothly running operation. Provider-delivered models can be up and running in 60 to 120 days. The question is not which approach is better; it is whether the organization can afford to wait. Companies under regulatory or commercial pressure to show they have mature operations often have no choice but to start with a

provider-delivered model and bring it in-house over time as their internal team matures, a deliberate path along the spectrum, not a permanent state.

Where Most Enterprises Get This Wrong

Three patterns of misalignment are worth naming directly, because they account for most of the operational pain we observe in the market.

The “we built it ourselves” trap

Organizations that built their NOC five or ten years ago often stick with that model long past the point where the conditions that made it the right choice still hold. The platform is dated, the engineers who built it have moved on, the documentation no longer matches how the team actually works, and the cost per incident has become difficult to defend. No one ever really says why the organization is staying, it tends to live in the assumption that "it's working." A useful test: if the CFO asked tomorrow for a defensible cost-per-incident number, could the organization produce one? If not, the operation isn't really working in financial terms, it just hasn't been examined closely enough to fail yet.

The “outsource everything” trap

Organizations that handed their NOC to a provider years ago often discover they have lost the institutional knowledge required to govern the relationship effectively. They cannot challenge SLAs they don't understand, cannot meaningfully participate in post-incident reviews, and cannot make informed architectural decisions about their own networks. The fix is not to bring it back in-house, that ship has often sailed, but to move toward a co-managed structure where the client retains enough operational visibility and engineering presence to govern the relationship as a peer, not as a customer.

The “we'll figure it out” trap

This is the most common and the most expensive. Organizations operate without documented operating procedures, without defined performance targets, and without a clear ownership model. Operations succeed because individuals make them succeed, often through informal heroics. This pattern looks fine until it doesn't, such as the senior engineer leaves, the monitoring system raises an alert nobody knows how to interpret, or the auditor asks for evidence that changes to the network are being properly controlled. Industry analysis underscores the cost: ITIC's 2024 data shows that more than 90% of mid-size and large enterprises now lose over $300,000 per hour of unplanned downtime, with 41% reporting hourly losses between $1 million and $5 million.[10] Splunk's 2024 analysis estimated $400 billion in annual downtime costs across the Global 2000.[11] A meaningful portion of that cost is preventable through the kind of basic discipline, written procedures, clear ownership, defined performance targets, that comes with deliberately choosing any one of the five models.

The honest test in each case is the same: was the current operating model deliberately chosen, evaluated against the five variables, and reviewed within the last 24 months? If not, the model is inherited, not chosen, and inherited models are where misalignment compounds quietly.

Where the Spectrum Is Headed

Three trajectories are visible in the market and worth factoring into any sourcing decision being made today.

The center of gravity is moving toward co-managed and hybrid models. Hybrid deployment growth at over 8% CAGR, outpacing both pure in-house and pure outsourced, reflects an honest assessment by enterprises that neither extreme is the right long-term answer for most of them.[2] The market is voting for shared accountability.

AIOps is becoming a baseline requirement, not a differentiator. Within the next 24 months, expect AIOps-enabled alarm correlation, ticket enrichment, automated incident summarization, and predictive maintenance to be table-stakes capabilities for any provider, and a basic governance expectation for any in-house operation. The customer-owned insight layer becomes a parallel concern, addressed in our previous Strategic Perspective on *AIOps and Managed NOC Convergence*.

Self-service controls are becoming part of the contract. On-call routing, maintenance window suppression, runbook governance, and self-service notification preferences, capabilities that were, until recently, "nice to have", are increasingly being written into operational requirements. The reason is straightforward: when a downtime hour costs $300,000 or more, the gap between detecting an issue and reaching the right person who can actually fix it is no longer tolerable as an operational rounding error.[10] Self-service controls close that gap, and the market is moving quickly in that direction.

The implication for organizations making sourcing decisions today: the spectrum will not stay static. The right sourcing model for an organization three years from now is unlikely to be exactly the right model today, and the best decisions are made with an explicit migration path in mind.

Closing the Loop

The build-versus-buy framing is finished. The right question is no longer whether to outsource; it is which of the five operating models fits the organization today, what conditions support that fit, and where the organization should be heading as those conditions change.

Most organizations have not made that decision deliberately. Their current operating model is inherited from a context that no longer applies, and the operational and financial cost of that misalignment is being absorbed quietly, until it isn't.

The framework above is offered without advocacy for any particular model. The right answer for any individual organization is shaped by its own circumstances, and the goal is alignment, not preference.

ENGAGE WITH KORE-TEK

If this perspective resonates , or if it raises questions about where your organization actually sits on the spectrum versus where it should , we welcome the conversation.

The Kore-Tek NOC Solutions Team works with enterprises across regulated and unregulated industries to evaluate operational maturity, assess sourcing alignment, and design migration paths across the spectrum. Our discovery conversations are structured around the five variables in this paper and produce a clear, defensible view of where you are, where you should be, and what would have to be true to move.

To schedule a discovery conversation, reach the NOC Solutions Team through **kore-tek.com/contact** or through your existing Kore-Tek relationship. Initial conversations are 55 minutes, structured against the framework above, and require no prior preparation on your part.

About This Paper

This paper is a Kore-Tek Strategic Perspective, part of an ongoing series intended as framework-level reference material for senior infrastructure leaders. The series is published in our Insights library at kore-tek.com/insights, alongside earlier papers including *The AIOps and Managed NOC Convergence*.

Kore-Tek delivers optical engineering, network infrastructure, and managed NOC services to enterprises, communications service providers, and public-sector clients across North America. The company's NOC Solutions practice supports clients across all five operating models described in this paper.

Sources and References

1. Straits Research, Network Operations Center as a Service Market Report, 2025–2034 forecast. Global market sized at $3.67B (2025) reaching $8.83B (2034) at 10.2% CAGR.
2. Global Market Insights, Network Operations Center as a Service Market Analysis, 2024–2034. Hybrid deployment models growing at over 8% CAGR.
3. Gartner, Generative AI in Network Operations Forecast, 2024. Projection: GenAI to handle 25%+ of initial network configuration by 2027 (up from <3% in 2024).
4. MarketsandMarkets, Network Operations Center as a Service Market Report, 2025–2030 forecast. AIOps deployments showing 20–40% alert volume deflection.
5. Industry analyst projections cited in network engineering workforce reports, 2025. Estimated 1.2 million+ unfilled network engineering roles globally.
6. U.S. Bureau of Labor Statistics, Occupational Outlook Handbook, Computer Network Architects and Engineers, 2024–2034 employment projection.
7. Linux Foundation, 2025 State of Tech Talent Report. 65% of organizations report understaffing in cybersecurity and compliance; similar gaps across cloud and platform engineering.
8. Foote Partners, IT Skills and Certifications Pay Index, Q3 2025. Network architecture commands 18% pay premium; Level-1 NOC monitoring premiums declining as automation expands.
9. Uptime Institute, Annual Outage Analysis 2024. Configuration/change management failures cited in 45% of network outages; human error implicated in 66–80% of downtime incidents.
10. Information Technology Intelligence Consulting (ITIC), 2024 Hourly Cost of Downtime Survey. 90%+ of mid-size and large enterprises lose over $300,000 per hour of unplanned downtime; 41% report $1M–$5M+ hourly.
11. Splunk, The Hidden Costs of Downtime Report, 2024. Estimated $400 billion in annual downtime costs across the Global 2000.

, END ,