

CLOSING THE TRUST GAP ***IN AI-DRIVEN NETWORKING***

**How zero trust principles can make AI-native
networks more secure and easier to trust**

Executive Summary

AI is changing enterprise networking at a fast speed and in drastic ways. It can spot anomalies earlier, automate routine decisions, improve traffic management and help security teams respond before issues spread. For senior leaders of security, networking and infrastructure, the appeal is more than obvious.

As is the hesitation.

Many are still uncertain about handing meaningful control to systems that can be hard to interpret and difficult to audit, especially at the scalable impact they make. This isn't resistance to innovation; it's sound judgment. The same technologies that make networks more adaptive can also introduce new risks. Bad models or weak data can turn automation into exposure while increasing the governance load. Besides, attackers are also using AI to accelerate reconnaissance, refine social engineering and increase infiltration speed. The [World Economic Forum](#) found that 94% of cyber leaders identify AI as the defining force in cybersecurity in 2026, yet roughly one-third of organizations still lack any process to validate AI security before deployment.

That's where we see the trust gap widening. Enterprises note the value of AI-native networking but aren't fully convinced it can be adopted without introducing new operational and security problems.

The answer to this dilemma? It isn't kicking the AI can down the road, nor is it always mistrusting the tech. The better path is to place it inside a security model built on continuous verification, constrained access, clear policy and strong oversight.

In other words, zero trust.

Zero trust often gets framed as an "identity" strategy. But at its core, it's a discipline that requires constant verification, explicit enforcement and never assuming anything is safe just because it sits inside the perimeter. Those principles map cleanly to AI-native networking. More important, they give enterprise teams a way to govern AI without giving up the benefits that made it attractive in the first place.

The real question, then, is not whether AI-native networks can be trusted in general. It's whether they're designed to verify and continuously validate AI-driven actions. When they are, AI stops being a leap of faith and becomes a manageable part of the control plane.

THIS PROVIDES THREE PRIMARY CONCLUSIONS:

1

AI-native networks can improve security and operations by detecting anomalies sooner, adapting policies faster, and reducing the lag that comes with manual response

2

Zero-trust principles provide the guardrails that make those systems safer, more auditable, and easier to operationalize

3

Trust in AI comes from visibility, governance, accountability and proof that the system operates within clear boundaries. NIST defines trustworthy AI in terms of security, resilience, accountability, transparency, explainability and privacy enhancement.

1. INTRODUCTION: THE TRUST GAP IN AI-DRIVEN NETWORKING

Enterprise networks no longer sit in one place. They don't follow one pattern or serve one user type. They stretch across branches, cloud environments, remote endpoints, partner ecosystems, APIs and machine identities. Workloads move and policies shift while attack surfaces keep expanding. In this context, manually managed infrastructure starts to show its limits.

That's why AI has moved so quickly into networking. It can process telemetry at a scale people just can't, pulling out patterns previously buried in noise and supporting near-real-time decisions about performance and security. We're talking massively less manual triage with fewer blind spots and faster response.


And yet the case for AI resilience and functionality isn't the same as the case for trusting it.

Many remain wary of systems that make consequential decisions without their logic being fully visible. There's reasonable worry about whether an AI model can tell the difference between a harmless anomaly and an early-stage compromise. Or what happens when a model acts on flawed inputs, or when an automated response affects a business-critical service. The stakes change once AI moves from observation into enforcement.

There's also a harder reality: defenders are not the only ones using AI. [ENISA has noted the dual role of AI in cybersecurity](#). It can strengthen defense, certainly, but it can also be manipulated or used in adversarial ways. So, enterprises must rely more heavily on AI at exactly the moment attackers are following suit.

That tension sits at the center of the trust gap in AI-driven networking. And it requires an answer grounded in design.

Zero trust happens to be that answer because it starts from a pressure-cooker premise: trust is never assumed. It is earned, verified and checked again.



2. THE FOUNDATIONS OF ZERO TRUST IN AI-NATIVE NETWORKS

“Never trust, always verify” has become a familiar line as it captures the essence of zero trust better than most longer explanations.

A zero trust model does a few basic things well. It verifies users, devices, workloads and applications based on identity and context rather than location alone. It limits access to what’s actually needed. It treats trust as temporary, not permanent. And it depends on broad visibility because policy without observation is guesswork.

This makes it a natural fit for AI-native networking.

Since AI can absorb the flood of signals modern environments generate, continuous verification becomes more practical. User behavior, device posture, traffic patterns, workload communication, access timing, lateral Movement, configuration drift, policy exceptions. A human team can review some of that; an AI-enabled system can analyze all of it around the clock.

The same is true for micro-segmentation. For years, many enterprises knew segmentation would reduce risk but struggled to implement it in maintainable ways. Policy sprawl, operational overhead and poor visibility made the idea sound better than the reality. AI can help classify assets, identify expected communication patterns, recommend narrower trust boundaries and support faster issue isolation.

So, zero trust sets the policy logic while AI helps enforce it at speed and at scale.

This is where the combination becomes powerful. Zero trust defines what must be protected, who should have access and under what conditions. AI dynamically evaluates those conditions instead of relying on static assumptions. This combo moves the enterprise from periodic trust decisions to adaptive trust decisions.

[CISA’s Zero Trust Maturity Model](#) reinforces that broader view. Its framework spans identity, devices, networks, applications and workloads, data, visibility, and automation. AI-native networking is most useful when it works across those pillars rather than sitting off to the side as a performance feature.



3. ADDRESSING THE TRUST GAP

BUILDING CONFIDENCE IN AI DECISION MAKING

Trust is not built by calling a system intelligent. It's built by making its behavior legible and answerable to human oversight.

Transparency and explainability

The real problem isn't that AI might make the wrong call. Any system can do that. The deeper issue is a system making the wrong call without leaving behind a clear reason.

[NIST places transparency, accountability, explainability, and interpretability](#) among the core characteristics of trustworthy AI. In networking, that principle needs to become operational. Teams should be able to see what signals were considered, what policy condition was triggered, how confident the model was, what action followed and whether a person approved or overrode it.

That does not require turning every network leader into a machine learning specialist. Just enough visibility to understand what happened, review whether it was appropriate and improve the system over time.

Accountability and governance

AI should never become an ownerless layer in the environment. Someone has to define the guardrails and approve where automation begins and ends. Evidence must be logged, exceptions must be made and model performance must be reviewed.

[NIST's AI Risk Management Framework](#) is useful here because it treats governance as a foundational requirement, not a post-deployment step. For CIOs and infrastructure teams, good governance usually comes down to a few practical decisions:

- What can the AI recommend, and what can it enforce automatically?
- Which actions require human approval?
- What must be logged for every automated decision?
- How will model drift, poor data quality, or degraded performance be detected?
- Where does authority sit when the system and the operator disagree?

These can help determine whether AI becomes useful infrastructure or an unexamined risk threshold.

Proactive threat mitigation

The clearest value of AI in networking is shortening the distance between weak signals and meaningful action. While attackers are moving faster than traditional workflows were built to handle, AI can correlate events across identity, network, application and device telemetry far more quickly than manual review. It can surface issues earlier, prioritize them more intelligently and support rapid response.

[The World Economic Forum's 2025 cybersecurity findings](#) emphasize that reality. AI is expected to reshape cyber risk significantly, yet many organizations still lack even basic processes for assessing the security of the tools they deploy. There's a second main benefit here: time. When AI-native networking is paired with unified visibility and clear policy, teams spend less energy chasing noisy alerts, reconciling overlapping tools or moving manually across vendors and consoles.



4. THE BUSINESS IMPACT OF AI-NATIVE NETWORKS WITH ZERO TRUST

This is often framed as a security issue. Yes, it is one. But the business case is broader and, in many organizations, more persuasive.

Enhanced security

Continuous verification narrows the opportunities for implicit trust to become exploited trust. AI strengthens that model by improving anomaly detection and enabling faster, more targeted response. Instead of waiting for an analyst to piece together signals from different domains, the system can surface a credible risk pattern as it forms.

Operational integrity

Security teams are accountable for more than breach prevention. They are expected to keep services available at peak performance and support change without creating constant operational friction.

AI-native networks can help by spotting signs of degradation early, correlating incidents across environments and supporting predictive maintenance. Within a zero trust framework, automation becomes easier to trust because it's constrained by policy rather than left to roam.

Cost efficiency

AI is not magic, and network modernization doesn't pay for itself overnight. But real efficiencies are available as automation reduces repetitive work and earlier detection can lower incident cost. Better segmentation limits the blast radius. More coherent policy and visibility can reduce overlapping toolsets and some of the waste that comes with vendor sprawl.

Still, the most important return may be organizational rather than purely financial. After all, experienced engineers are expensive and hard to replace. When they aren't buried in manual triage, they can focus on higher-value work.

5. OVERCOMING BARRIERS TO ADOPTION

Yet even when the logic is sound, adoption is rarely straightforward.

The first barrier is control. Teams can worry that once AI is embedded deeply enough, they will not be able to predict or contain its behavior. While taking this concern seriously, the solution isn't to insist that the system is smarter than the people using it. Instead, you must define where automation is appropriate, where it is not and how exceptions are handled.

The second barrier is complexity. Adopting zero trust and AI-native networking doesn't mean replacing everything at once. In most cases, a phased approach is more practical and credible.

A sensible path often starts with visibility. Bring together telemetry, identity context and asset posture before expanding enforcement. Then focus on the highest-value use cases: privileged access, critical applications, sensitive east-west traffic and high-risk user populations.

Keep automation bounded at the get-go. Let AI recommend actions or allow it to execute only low-risk responses until the organization has enough evidence to trust broader use.

Build governance alongside deployment, not as an afterthought. Logging, review processes, accountability, rollback procedures and exception handling should be established early on.

Training is important, too. Teams need to understand how decisions are made, what signals matter and where human judgment still belongs. External partners can help, but only if they make the system easier to understand and govern.



6. CONCLUSION: **THE FUTURE OF SECURE AI NETWORKING**

The reality of the trust gap in AI-driven networking reflects a healthy skepticism about handing critical responsibilities to systems that can act quickly and opaquely if left unchecked.

But manual operations are not enough anymore. The true choice is no longer between human judgment and machine intelligence. It's between unmanaged automation and governed intelligence. Zero trust provides the discipline necessary to bridge the gap while AI provides the speed and analytical scope. Together, they can produce networks that are more resilient and accountable than either approach on its own.

WILL YOUR ORGANIZATION ADOPT AI ON TERMS THAT PRESERVE CONTROL AND EARN TRUST OVER TIME?

That starts with a hard look at where your environment still relies on implicit trust. From there, you must identify where AI can improve verification and response without allowing for uncontrolled autonomy. In the end, organizations that do this well won't be focused on the fastest automation but on the strongest governance.

ABOUT GTT

GTT is a leading networking and security as a service provider for multinational organizations, simply and securely connecting people and agents to data and applications anywhere in the world. We serve thousands of organizations, bringing together the right people, partners and technology to reduce the burden on IT teams and solve the most pressing networking and security challenges. Built on our top-ranked global Tier 1 network, GTT Envision is a single global technology platform to connect, orchestrate, virtualize and automate enterprise networks, enabling customers with consumable solutions to achieve business missions and meet ongoing demand when, where and how needed. Our portfolio includes SASE, SD-WAN, security, internet, voice and other connectivity options, complemented by a suite of professional services and sales and support teams in local markets around the globe. We partner with our customers to deliver Greater Technology Together. For more information, visit www.gtt.net.

ManServ_Closing the trust gap_AIWtPaper_GGMCM-4487_v1_08042026