Omdia
by informa techtarget

MAY 2026

Harnessing Network-Derived Telemetry to Strengthen Security in the AI Era

Rik Turner, Chief Analyst, Cybersecurity

Abstract: The widespread adoption of large language models (LLMs) to underpin a generative AI (GenAI) user experience in corporate applications, as well as their imminent expansion to support agentic deployments, heightens the need for reliable data on which to base decisions from a performance and, perhaps even more critically, a security perspective.

While the event logs collected in a SIEM and the alerts generated by EDR tools are key sources of information, they tell only part of the story. Network-derived telemetry, including packet data, flow records, and application metadata, provides the critical context that SIEM logs and EDR alerts lack. Combining these data sources enables deep observability, delivering a more complete understanding of performance and security behavior across hybrid cloud environments.

Agents, LLMs, and hybrid clouds make network telemetry a must-have

With the number of actors on corporate networks set to proliferate thanks to the entry of agents into the fray, such real-time visibility and context will become critical. Agents are non-deterministic non-human identities (NHIs) that can, among other things:

- Elevate their privileges.
- Create their own child agents.
- Seek to access data sources well beyond their original remit.

All of these capabilities demand an immediacy of response that only network telemetry can provide.

Many of the applications into which LLMs are being integrated will be cloud-native, residing in public, private, and often hybrid cloud environments. To effectively monitor the movement behavior of humans, agents, and systems in this scenario and respond to potential threats and exploits in a timely manner, the comprehensive nature of network telemetry becomes a must.

The addition of network telemetry into the mix of security operations (SecOps) data sources enables security professionals to move from discrete, point-in-time snapshots of "what happened" (logs) to continuous, high-fidelity streams of behavioral data ("how it is happening"). In other words, while event logs report that a server logged a user in, network telemetry provides the full context of that session, including traffic volume, latency, and packet-level behavior. Brought together, therefore, these data sources enable deep observability, providing a more complete understanding of system behavior for more effective detection, investigation, and response.

East-west traffic evades security tools on the perimeter

Yet another complicating factor in modern cybersecurity derives from broader changes in enterprise technology. As applications have moved to microservices architectures (containers, serverless functions, etc.), ever more traffic on enterprise networks, including data centers, is so-called east-west (lateral) traffic, which means it travels between microservices, as well as APIs, GPUs, and datasets, never actually leaving the corporate perimeter.

According to data from content delivery network Akamai, as much as three-quarters (**76%**) of all data center traffic is now east-west,[1] rendering it essentially invisible to security tools that were designed to inspect north-south traffic (i.e., in- and outbound traffic going to and from corporate endpoints and servers). This makes it essential to have a telemetry pipeline that sits in the network to collect and deliver telemetry to perimeter-based tools such as firewalls that have no visibility into east-west traffic. And while internal firewalls can be deployed to deliver network segmentation, which inevitably requires some level of east-west traffic visibility, it is neither best practice nor cost-effective to implement such devices throughout a corporate infrastructure.

Figure 1. More than three-quarters of data center traffic is now east-west

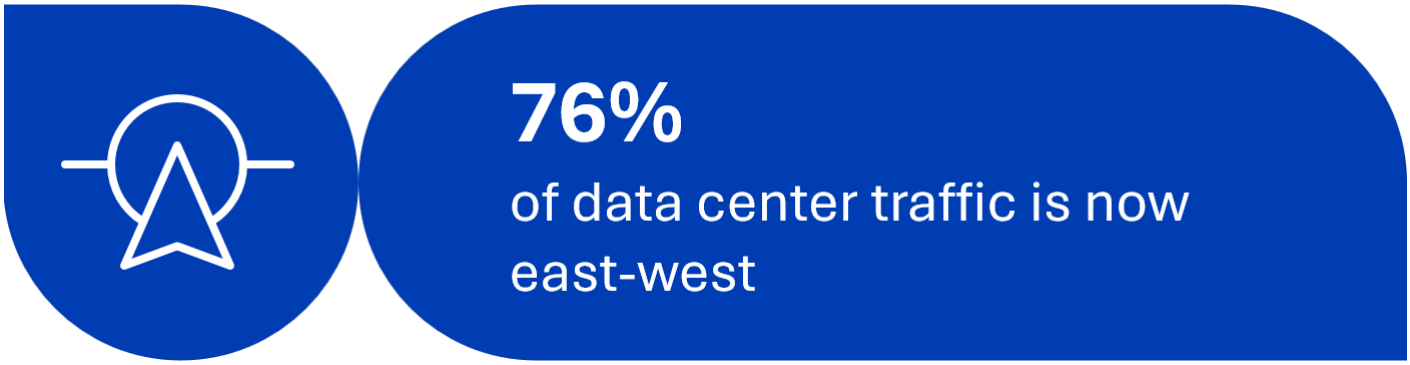

Source: Akamai

The FSI dimension

Cybersecurity is, of course, a requirement in all sectors of the economy, from non-profits, through the private sector, to government, and, as such, any and all of them can benefit from investing in a deep observability capability. That said, there are one or two verticals where there are extra benefits in play—perhaps none more so than financial services. This is because the rapid evolution in AI is enabling the creation of ever more convincing deepfakes to underpin financially motivated attacks. As such, the cyber-adjacent activities of fraud detection and risk management will only be able to keep up by harnessing network telemetry and making it an essential input for their analytics.

Indeed, the financial sector provides a prime example of the importance of network telemetry. Not only are financial service institutions (FSIs) prime targets for cyberattacks, but they are in the forefront of AI adoption, seeking to incorporate LLMs across their businesses. From the contact center to their CRM and HR systems,

[1] Source: Akamai blog, *East-West Is the New North-South: Rethink Security for the AI-Driven Data Center*, January 2026.

as well as the private apps enabling investment and corporate banking and even the core banking systems that sit at the heart of these entities, the use of the new forms of AI is rapidly expanding.

In other words, the accelerated adoption of AI, both within their organizations and by their criminal adversaries, makes deep observability, based on network-derived telemetry, an essential component of FSI security going forward.

Addressing the Big Data and encryption challenges

While logs provide the equivalent of a snapshot of a recent event, network telemetry data delivers a recording of what is actually going on at that moment. Furthermore, there are entire classes of devices such as network elements (switches, routers, load balancers, etc.), IoT nodes, and operational technology (OT) endpoints that simply cannot support the software clients required by security platforms such as EDR. Even in those classes that can carry such a client, there are devices that fall outside the purview of a corporate security team, such as the unmanaged laptops and smartphones used by contractors or partners.

As such, the inclusion of network telemetry in any analysis carried out for security purposes is clearly an invaluable addition. The problem, then, becomes one of separating the wheat from the chaff, not only because of the "three Vs" of network data (i.e., its volume, velocity, and variety) but also because security teams frequently need to respond quickly to a situation.

Network telemetry is Big, Big Data

It is common practice for organizations to collect, curate, and analyze network telemetry for performance and security purposes. This, nowadays, entails collection across complex, hybrid infrastructures, comprising on-premises assets and a range of cloud and edge environments, as well as the devices of numerous employees working remotely.

With 400Gbps backbones finding their way into enterprise networks over the last couple of years, this was already a challenging task. Now more so than ever, network telemetry involves a lot of data. As such, the key—and the knack—is in knowing how to manage it.

Efficient access is, thus, just one part of the requirement, with others including:

- The efficient transformation of packets into metadata.
- Filtering and selection.
- Efficient ingest and storage regimes.

However, the situation has now been complicated even further by the rapid adoption of generative and, going forward, agentic AI. These technologies drive a dramatic increase in the three Vs of network data.

Most traffic is now encrypted

Additionally, ever more traffic on corporate networks, not to mention the public internet, is now encrypted. Estimates suggest the figure is now as high as 90%,[2] presenting a further obstacle for organizations seeking to

[2] Source: European Union Agency for Cybersecurity, *Encrypted Traffic Analysis: Use Cases & Security Challenges*, November 2019.

gain a comprehensive understanding of their security posture, their attack surface, and the present threat landscape. In a report issued in December 2024, cloud-based security vendor Zscaler found that **87.2% of all cyberthreats that year were delivered over encrypted channels (TLS/SSL)**, driving a 10.3% increase year-over-year in threats over HTTPS.[3] That figure will only increase.

Figure 2. Nearly 90% of cyberthreats are now delivered over encrypted channels

Source: Zscaler

Decrypting all this traffic to inspect it is a Herculean—nay, impractical—task that would add an unacceptable degree of latency on security devices themselves, not to mention consume significant processing cycles that could and should be devoted to other work.

The right approach

In light of these challenges, a network telemetry pipeline that can collect, aggregate, triage, and forward to an organization's security tools the relevant information from its network must have a series of features and functions, including:

- The ability to streamline the data, focusing on and forwarding only what matters (i.e., actionable context delivered as metadata).
- A centralized decryption management capability, given the resource constraints involved in trying to decrypt traffic on individual devices such as endpoints or firewalls.
- The automatic discovery of applications, both authorized and unauthorized, including GenAI apps in operation on the network. The larger the number of applications for which a telemetry pipeline has signatures for identification purposes, the better.
- Visibility into GenAI usage for informed policy decisions and data governance. This is of particular importance because of the ease with which GenAI services delivered in SaaS mode can be accessed, constituting the very real risk of a "shadow AI" existing alongside sanctioned GenAI platforms.
- The ability for security professionals to drill down into individual applications, application components, and protocols to identify vulnerabilities.

[3] Source: Zscaler ThreatLabz, *2024 Encrypted Attacks Report*.

- The ability, using simple flow control, to pass only the specific, relevant application traffic to the security tooling for analysis, dropping irrelevant traffic. In other words, the platform must be able to address the volume challenge.
- Contextual information generation. This capability is the most critical of all. After having extracted and summarized information from data packet data and flow records, the platform must then generate relevant contextual information, including metadata, about the applications and protocols in use on the network. It must then feed this data into a SIEM or data lake of the organization's choosing, exposing it to a range of AI tools for analysis. The metadata should cover multiple attributes across as large a number of apps as possible, enabling the security tools to identify suspicious activities and performance issues.

NetFlow's blind spots in modern networks

Many tools in both the security and performance/observability fields attempt to rely on classic network logs (5-tuple) and NetFlow for the identification and tracking of individual communication sessions across a network. This approach has the following limitations:

- Port numbers do not always accurately identify specific applications or application versions because modern apps can use dynamic port assignments or tunnel multiple services through standard ports like HTTP/HTTPS.
- Reliability of the classic log data may be questionable, given that they are sourced from devices that may themselves be compromised.
- Complete traffic characterization becomes more difficult when the traffic is encrypted, for, while the 5-tuple remains visible, application-layer content requires additional techniques like deep packet inspection (DPI) or metadata analysis.

Ideally, a comprehensive network telemetry pipeline, like the Gigamon Deep Observability Pipeline, would provide many more parameters than just the 5-tuple or NetFlow, such as the type of networking gear in use, the identity of any appliances in use on the network, and whether there are any OT devices present in the environment, while also generating high-fidelity metadata to address the 5-tuple's data encryption shortcoming.

Harness AI to streamline investigations and remediation

While SecOps teams face ongoing challenges due to their organizations' adoption of both generative and agentic AI (more data to triage, more identities to monitor and manage, etc.), the good news is that AI can also be leveraged to help address these issues. AI is ideally suited to sift through large volumes of information, select and prioritize the most important parts, perform analysis, and provide actionable guidance in a timely fashion.

When selecting a pipeline approach to collecting and curating network telemetry, organizations should investigate not only the essential features and functions of any proposed solution but also its ability to deliver an AI overlay to streamline the actions of the security team. In addition to working with the network and application metadata derived from the telemetry, it should ideally have integrations into SIEM and observability platforms, as well as the ability to handle cloud logs from the leading cloud service providers, accelerating investigations without the need for analysts to comb through dashboard data manually. And as a

GenAI capability, the ideal telemetry pipeline should enable users to ask questions, query the metadata, and receive insights and recommended actions, enriched with context, within the analytical tools they already use.

Conclusion

Modern networks demand greater visibility

Network telemetry is an essential supplement to other forms of performance and security data. This is particularly the case now that corporate networks have become more complex, incorporating both private and public cloud assets alongside those that reside on an organization's premises.

There is a need to overcome these challenges, not only of the three Vs of such telemetry (volume, variety, and velocity), which have themselves been turbocharged by AI adoption in the enterprise, but also of achieving visibility in the face of increased encryption and east-west traffic.

Fortunately, while it compounds the Big Data problem, AI can also be harnessed to provide at least part of the solution. AI can perform the analysis of contextual metadata from network packets, then go on, in a SecOps context, to provide recommendations for remedial action.

Building a strategic foundation for AI-era security

It must be noted that a technology has never evolved at such a speed as AI, making the provision of security for AI a rapidly moving target. Indeed, AI architectures are evolving so fast that the end state, if such a thing even exists, remains unclear. This, in turn, makes perfect upfront design unrealistic.

Instead, organizations should prioritize lasting fundamentals that remain valuable no matter how AI develops, such as full visibility into network traffic and well-planned telemetry management. Building this foundation early makes it easier to adapt, scale securely, and avoid expensive rework as AI technologies and use cases evolve.

Going forward, organizations should include network telemetry as a key data source for their security tools, since it complements what tools such as SIEMs, EDRs, etc. provide, while also offering information about exposing possible blind spots. Telemetry management is strategic and, as such, should be both built in and budgeted for early.

A network telemetry pipeline needs to capture and filter the data by relevance to the specific incident and efficiently deliver it to the security tools used. It should also be able to generate metadata to add context to the analysis carried out by security platforms. Additionally, organizations should consider whether the provider of the network telemetry pipeline can also deliver an AI capability to streamline the work of their SecOps team.

A network telemetry pipeline with AI to enhance outcomes

Omdia considers network telemetry to be an essential complement to event logs and other metrics if SecOps teams are to gain a comprehensive view of their organization's security posture and its current threat landscape. A telemetry pipeline deployed to provide such data should not only be able to monitor east-west

traffic but also handle the growing volume of encrypted network traffic, generating metadata from raw packets and delivering it into downstream security tools for analysis. As a further enhancement, the provider should be able to deploy modern AI technology to accelerate and streamline the work of the SOC team.

Together, these capabilities enable a more complete and context-rich view of system and network behavior—deep observability—supporting improved security and performance outcomes.

Get in touch: www.omdia.com askananalyst@omdia.com