theCUBE research

theCUBE Research White Paper

AI-Native Routing: Enterprise Private WANs for the AI Era

Date: February 2025 **Author:** Bob Laliberte

ABSTRACT:

Modern enterprises operate in distributed and complex IT environments spanning private data centers, public clouds, and edge locations, with wide area networks (WANs) playing an increasingly critical role. The rise of generative AI (GenAI) initiatives is driving a shift back to on-premises data centers and campus/branch locations. This puts more pressure on WANs to support growing AI traffic and existing critical business applications. These private WANs are essential for industries like finance, healthcare, research, and federal or state governments, with more constantly being added to the list. Juniper Networks' AI-Native routing and automation solutions provide enterprises with a clear pathway to accelerate the deployment and operation of efficient, self-driving WAN environments. These solutions optimize performance, enhance security, and reduce costs while supporting the demands of AI initiatives.

Modern enterprise IT environments are highly distributed

To remain competitive in a global economy, enterprises build and deploy applications across private data centers, multiple public clouds, and edge locations. As a result, modern IT environments have become highly distributed and reliant on WANs to ensure secure connectivity and positive experiences.

While much of the focus over the last five years has been on shifting applications to one or more public clouds, the introduction of GenAI initiatives is driving applications and computing back to private data centers. Why the shift? One key reason is that enterprises still have most of the data in their own data centers that needs to be used for training and inference to train the AI models and bringing the AI to the data is more secure. Research[1] also indicates that more than half (58%) of organizations focus on creating GenAI training and inference environments across on-premises data centers.

Enterprises are also focusing their IT budgets on creating these new AI environments. The ETR Technology Spending Intentions Study from October 2024 highlighted that 56% of enterprises reported increased spending on AI, surpassing even support for modern applications (container orchestration and platforms) and cloud computing.

These AI initiatives drive growth across on-premises data centers and in campus and branch locations that generate or house the data required for AI model training. This drives the need to ensure robust, secure connectivity between all

[1] Source: The Impact of AI on the Network, September 2024 Report Name

locations. For enterprises operating across the metro, region, nation, or globe, this places significant pressure on the WAN to deliver highly performant and secure private WANs to connect private data centers, public clouds, and edge locations involved in the training and inferencing of AI workloads. Most commonly, private WANs have been deployed in financial services, manufacturing, federal or state government, education, and research facilities that need to transfer data quickly or gain critical real-time business insights from data, regardless of where it is generated. It is not limited to just these industries. As AI adoption grows, local municipalities, healthcare, and many others will also need private WANs to securely transfer data to and from locations where it is generated or processed.

Distributed environments create challenges for enterprises

Unfortunately, these highly distributed AI environments inevitably create more complex network environments. Research[2] reveals that 80% of enterprise organizations state that the network is more or much more complex than it was two years ago (see Figure 1).

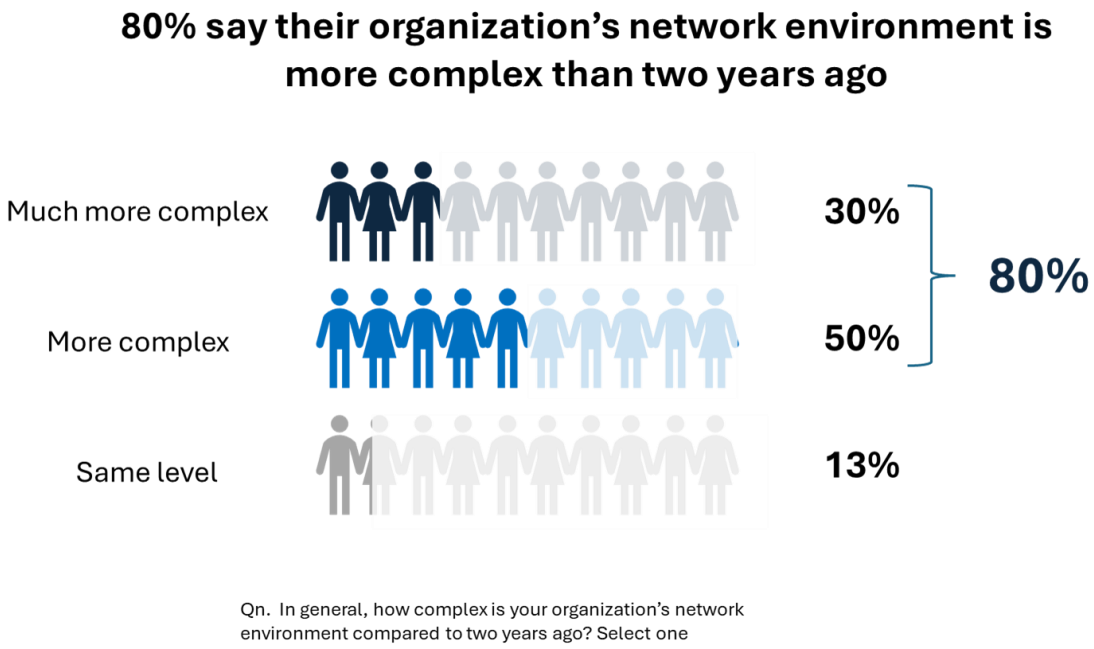

Figure 1 – Network complexity is increasing

This complexity makes it more challenging for network teams to ensure availability and performance across locations and support the business. It also delays network operations teams trying to find and fix issues, making it difficult to ensure a highly available and secure environment.

When we asked what was driving overall network complexity, respondents[3] stated that the top three factors were increased traffic, a distributed network, and multiple different management tools, with the requirement to support highly dynamic modern apps following closely behind in fourth place.

The top challenge, increased traffic, has more implications for a business, ranging from the requirement to upgrade or modernize network infrastructure to ensuring that critical business traffic is prioritized over less critical but potentially

[2] Source: Research Report The Impact of AI on Networking. tjeCUBE Research and ZK Research July 2024
[3] Ibid

higher volumes of data. Also, organizations must now deal with all the new AI workload traffic that needs to be moved for training and inferencing. As a result, network teams must ensure that critical business traffic is not impacted as AI workload and overall traffic continues to increase. This becomes more difficult as organizations try to impose different performance objectives for specific applications across a WAN in highly distributed environments. This is especially true for highly dynamic, modern application environments (see Figure 2).

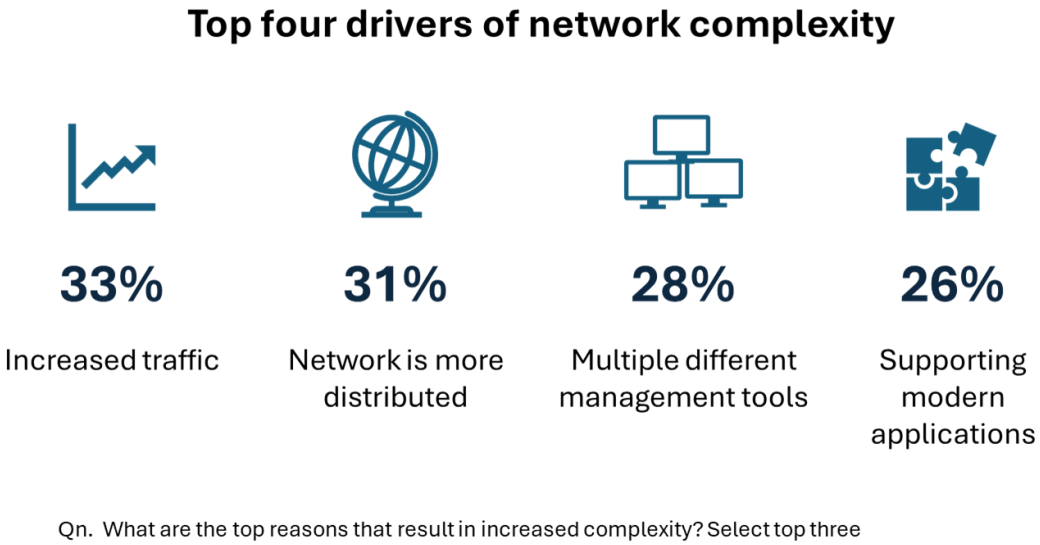

Figure 2 – Top four drivers of network complexity

Given the increased complexity, many organizations must consider using AI and automation to simplify networking. However, research[4] highlights that challenges also exist there (see Figure 3), including:

- Lack of a unified AI framework: Many network vendors have domain-specific AI that lacks a comprehensive or extensible framework. Different AI interfaces for each network environment only add complexity and can hinder remediation times, potentially limiting the solution's value.
- Comfort with AI technology: This is a big challenge as operations teams are conservative (and rightfully so). Before a new technology is introduced, it needs to be appropriately validated. Operations need time with the solutions, manually confirming the conclusions the AI tool has reached before trusting it to launch any automated remediation.
- Aversion to risk or change from the ops team or executives: Cultural change is often more challenging than technological change. As a result, organizations' operations and executive teams can be a barrier to adopting AI and automation. Executives may view new technology as a risk and seek to mitigate it, and operations teams may be more comfortable with existing processes and resist implementing new technology.
- Efficacy of AI tools: Closely tied to the comfort level with AIOps technology, organizations are wary of the performance of immature AI solutions and will need to verify the alerts and recommendations. Both vendors and organizations need to take steps to shorten the time to validate the proficiency of this technology.

[4] Ibid

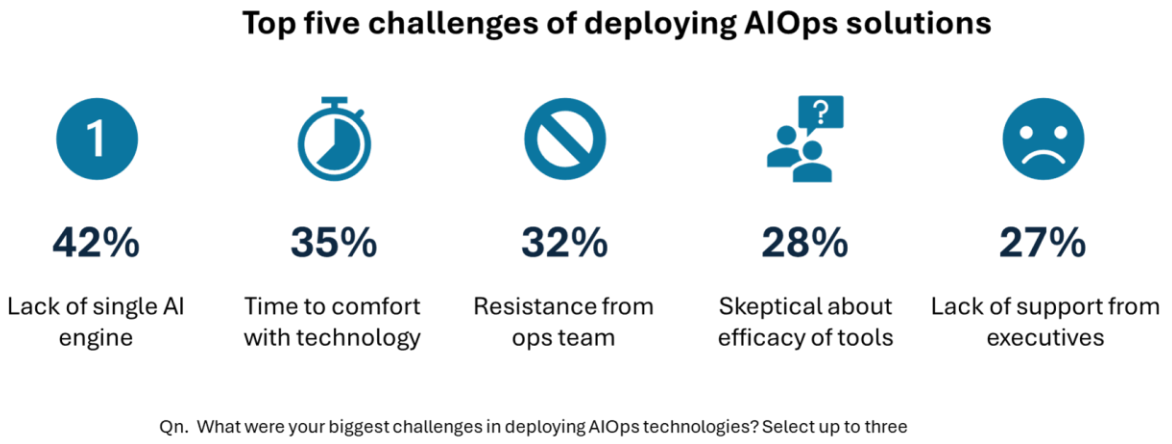

Figure 3 – Top five challenges of deploying AIOps solutions

In addition, organizations are emphasizing remaining highly available while actively managing power constraints. This is especially true for organizations with GenAI deployments, which require massive amounts of power. Organizations are looking for locations with significant, sustainable power sources (nuclear, wind, solar, etc.) to build AI data centers. Because of this, there is more scrutiny on the power efficiency of all the IT infrastructure deployed in a data center or edge location.

Why private WAN for enterprises?

Virtually every enterprise utilizes WAN connectivity. Historically, these networks have been contracted services with fixed lines provided by a telecommunications company. SD-WAN technologies enabled enterprises to create virtual overlay networks that abstract multiple internet broadband underlay connections for their WAN with best-effort service levels, not guaranteed performance levels or SLAs. This can be challenging for organizations with business-critical, latency-sensitive applications.

Building out a private WAN can yield better results for organizations that require higher levels of control, visibility, and performance for critical business applications. This is more than deploying SD-WAN and broadband or MPLS links. Private WANs enable organizations to own and control the IP routing infrastructure at all locations where it is deployed. This ensures organizations have total visibility into the network environment and can deliver highly secure, highly performant connectivity that can prioritize business-critical applications and AI workloads as needed. Enterprises control the end-to-end infrastructure (routers for data center interconnects and branch and campus connectivity) and can precisely define the performance requirements for all applications and workflows (see Figure 4). As GenAI initiatives rapidly expand, most organizations will adopt a hybrid approach for network connectivity that balances both private WAN and SD-WAN to meet the appropriate business requirements.

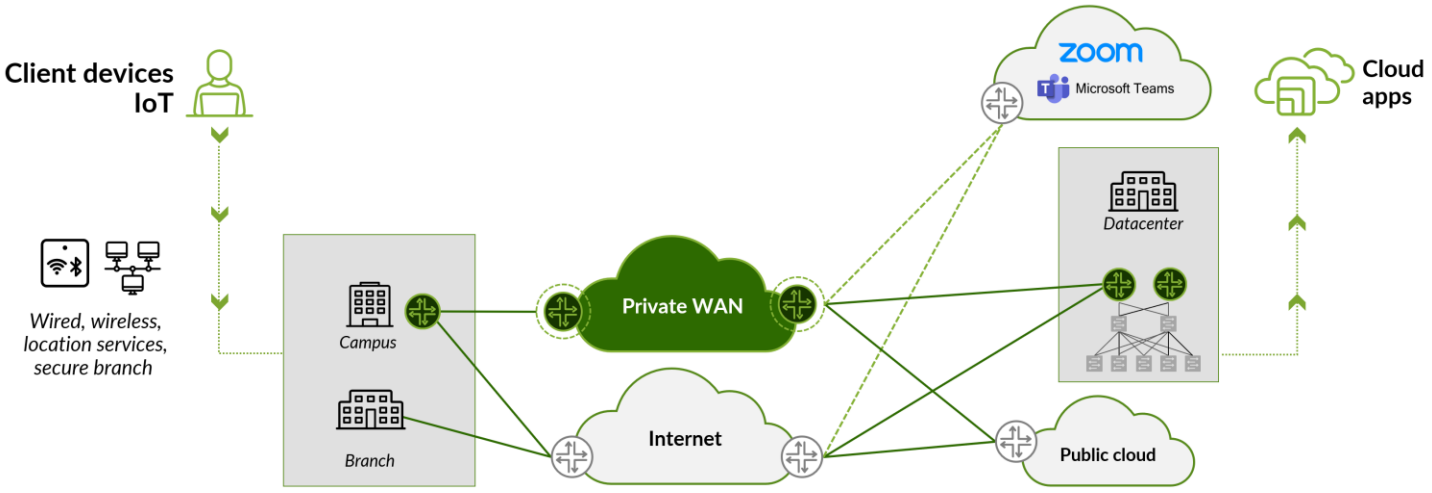

Figure 4 – Private WAN and SD-WAN hybrid environments Source: Juniper

Why is getting the right networks in place so important? Because in a highly distributed environment, the network has become a critical enabler for business. In fact, more than nine out of 10 organizations (93%) report[5] that the network is now more or much more important in meeting business goals (see Figure 5).

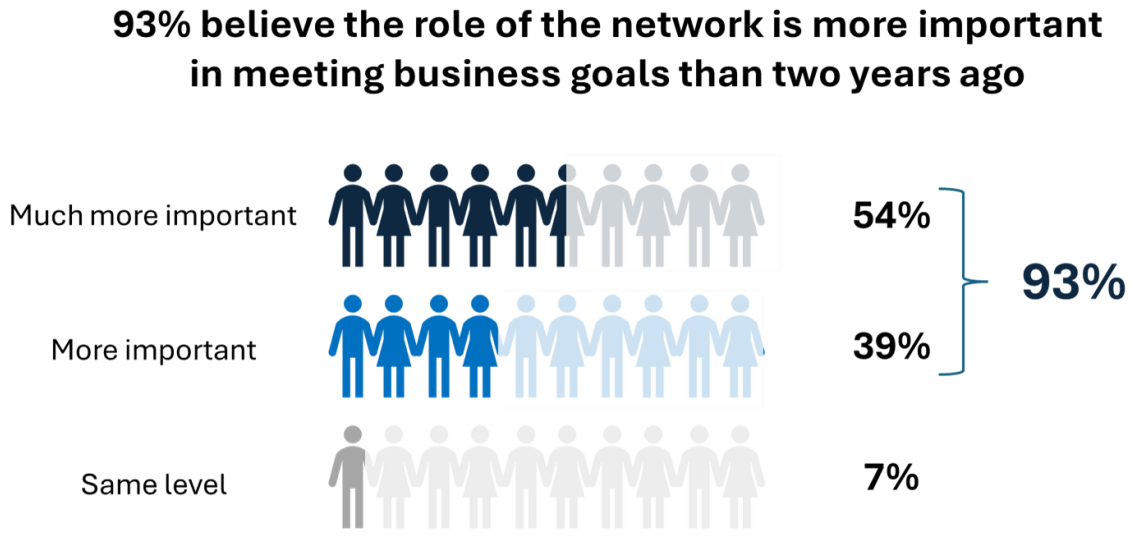

Figure 5 – The importance of the network in meeting business goals

As a result, organizations that want greater control, security, and customization of application performance will leverage private WAN environments.

[5] Ibid

Juniper's vision for AI-Native routing and self-driving WANs

Ensuring optimized performance across distributed and complex environments requires a solution that heavily leverages AI and automation. Juniper has extensive WAN experience, as the company initially developed the routers to power the internet and communication service providers.

Juniper has also pioneered the use of AIOps and automation technologies, what it refers to as AI-Native Networking, for enterprise network environments—data centers, campuses, SD-WANs, and branch networks. Mist AI has spent almost a decade developing and optimizing an extensible AI engine and Marvis Virtual Network Assistant (VNA) to cover all the networking environments Juniper supports, including AI capabilities for routing assurance.

It should be noted that Juniper has also made similar investments in the network service orchestration domain, adding intent-based networking and automated closed-loop assurance, a core capability of the Juniper Paragon Automation offering. Juniper also has deep experience with software-defined networking (SDN) controllers in transport networks (with Pathfinder evolving into integrated components within Paragon Automation for closed-loop automation scenarios).

As a result of this experience and expertise, Juniper has created a vision to enable AI-Native routing for private WAN environments. Juniper believes private WANs will serve many enterprise use cases, including Campus and Branch Edge, Enterprise WAN Edge, private WAN backbone, internet edge peering, Data Center Interconnects (DCI), DCI Edge, Industrial Metro, and AI/ML clusters (see Figure 6).

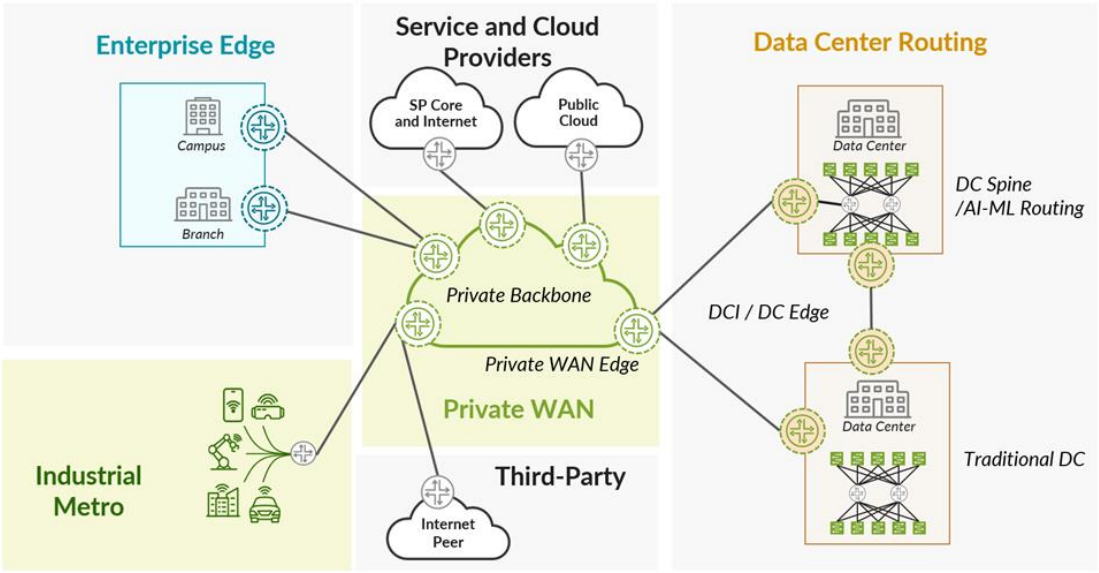

Figure 6 – Application insights leveraging AI and analytics – Source: Juniper Networks

Juniper understands the requirements for enterprises to build a private WAN and has a portfolio of solutions from which it can draw to build comprehensive AI-Native routing solutions for private WANs. This includes capabilities from Mist AI, Paragon, Juniper routers (ACX, MX, PTX), and optics. It also leverages custom silicon, Zero Trust principles, and adaptive power solutions (see Figure 7).

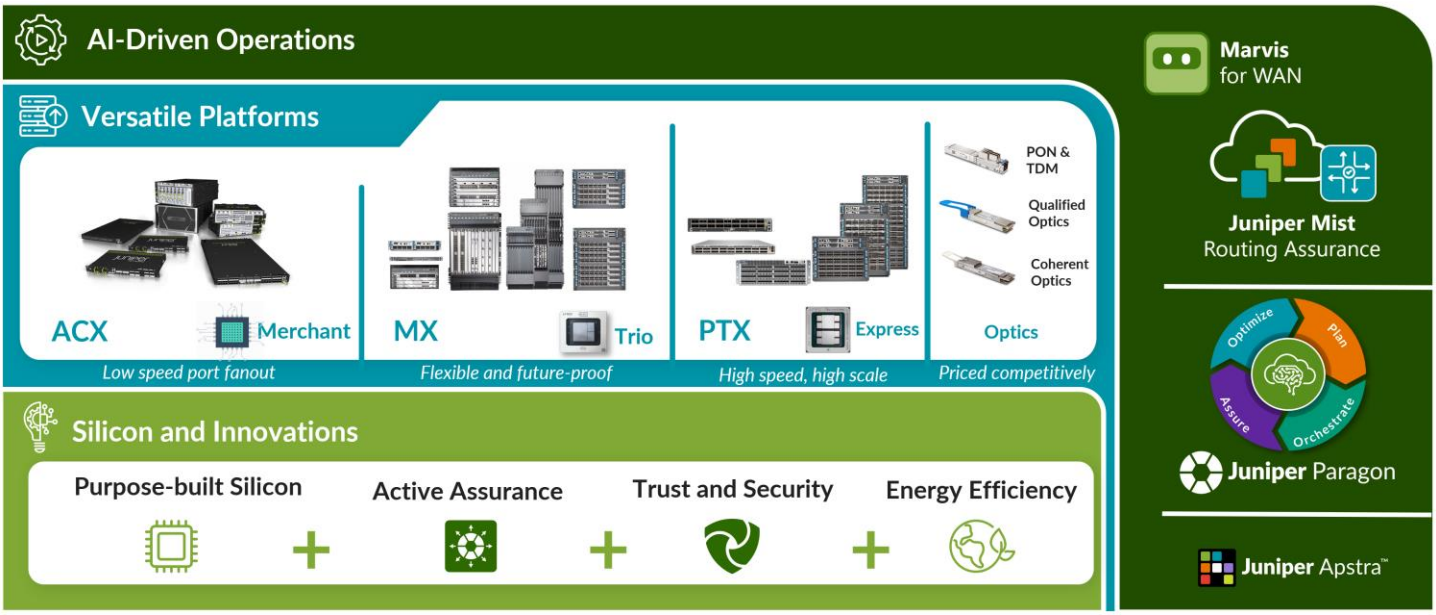

Figure 7 – Juniper Enterprise Private WAN portfolio of solutions – Source: Juniper Networks

Juniper's routing platforms: The Juniper family of routers, which includes the ACX, MX, and PTX Series, are designed to deliver high performance, scalable, and secure solutions for several different networking needs that span service provider core networks to enterprise edge deployments. All router series leverage Juniper's proven operating system, Junos OS. The MX Series is a power-efficient series typically deployed for edge and cloud-ready routing while the PTX Series is focused on ultra-high-capacity core routing with 100 and 400Gb widely deployed, and the series is 800Gb-ready. The ACX Series provides compact, energy-efficient access and aggregation with Zero Trust security for enterprise edge environments (Branch/Campus, Metro, DC, WAN).

Juniper's routing platforms are equipped with segment routing, EVPN, and AI-driven insights, ensuring seamless operations, low latency, and optimized performance for various environments.

Purpose-built and merchant silicon: Juniper's approach to silicon is about more than just power; it's about providing the right solution for each unique network need. Juniper's custom-designed chips are optimized for specific tasks like high-speed packet processing, advanced traffic engineering, and robust telemetry, and they enable intelligent power management while maintaining low latency and high throughput.

The MX Series leverages the trio chipset for the delivery of advanced complex network services while the PTX Series takes advantage of the Express 5 ASICs to deliver the highest throughput efficiency. The ACX Series uses merchant silicon to balance edge performance and cost with Broadcom DNX™ technology for robust, innovative solutions.

Intent-based orchestration: Juniper's dynamic Paragon Automation platform delivers intent-based orchestration in the WAN, translating high-level business objectives into automated, enforceable policies to optimize network performance and reliability. Paragon automates the life cycle of WAN services, from provisioning to assurance, using AI-driven insights, real-time telemetry, active assurance, and closed-loop feedback to adjust configurations and maintain alignment with business intent. It enables multi-domain orchestration across core, edge, and cloud environments, ensuring scalability and efficient resource utilization while meeting SLAs. This approach reduces

operational complexity, accelerates service delivery, and enhances customer experiences through proactive management and consistent network performance.

Trust and security: In addition to inline MACsec encryption and inline IPsec, Juniper platforms extend state-of-the-art infrastructure security everywhere. Juniper routers have a unique device identity, using the Trusted Platform Module (TPM) 2.0 framework to automatically validate hardware and OS integrity when booting. They can also embed advanced security services like Corero SmartWall® Threat Defense Director (TDD) for distributed denial of service (DDoS) protection.

Energy efficiency: Juniper's adaptive power solutions provide several energy-efficient features across its router portfolio. These include Dynamic Power Management, which adjusts power use based on network demand, and efficient thermal designs to reduce cooling costs. The routers all have high port density to enable fewer devices to be deployed with the ACX line, which supports compact designs for power- and space-limited environments. The routers can also minimize redundant power use or idle time.

Marvis Virtual Network Assistant (VNA): This AI-driven solution is designed to simplify and optimize network operations by providing conversational interfaces, actionable insights, and automated troubleshooting across the entire Juniper networking portfolio (WAN, LAN, WLAN, and DC). Powered by Mist AI, Marvis uses machine learning to analyze vast amounts of telemetry data from Juniper's network infrastructure, enabling real-time anomaly detection, root cause analysis, and proactive issue resolution. It enhances user experiences by predicting and addressing network performance issues before they impact operations. It also empowers IT teams with natural language queries for fast, intuitive access to network insights. By streamlining workflows and reducing manual intervention, Marvis VNA increases operational efficiency, improves network reliability, and ensures seamless connectivity.

Juniper Mist Routing Assurance: This solution uses AI-driven insights and automation to optimize routing performance across the network. Integrating Juniper's Mist AI platform simplifies network troubleshooting by continuously monitoring routing behavior, detecting anomalies, and validating network performance against intent-based policies. It provides granular visibility into routing paths, identifies potential misconfigurations or performance bottlenecks, and offers actionable recommendations or automated corrections to maintain network health. This proactive approach ensures reliable application delivery and reduces downtime, allowing IT teams to optimize routing efficiency while meeting business-critical performance and availability requirements.

Juniper accelerates the journey to AI-Native routing for the AI era

Enterprises recognize the need for intelligent, sustainable solutions to help with increasingly complex and power-constrained network environments. As the research indicates, they adopt AI and automation for their network operations environments as they validate and become more comfortable with technology. The end goal is to be able to utilize AI-driven automation across all network domains.

For enterprise networks, more than half (55%) of the respondents[6] stated their network automation tool leverages AI/ML extensively in data centers, campus and branch, and SD-WAN environments, driving greater availability and operational efficiency. Organizations are now realizing they also need to adopt AI to drive intelligent alerts, recommendations, and fully automated solutions for their WAN networks (see Figure 8).

[6] Source"The Impact of AI Impact on Networking, theCUBE Research and ZK Research July 2024

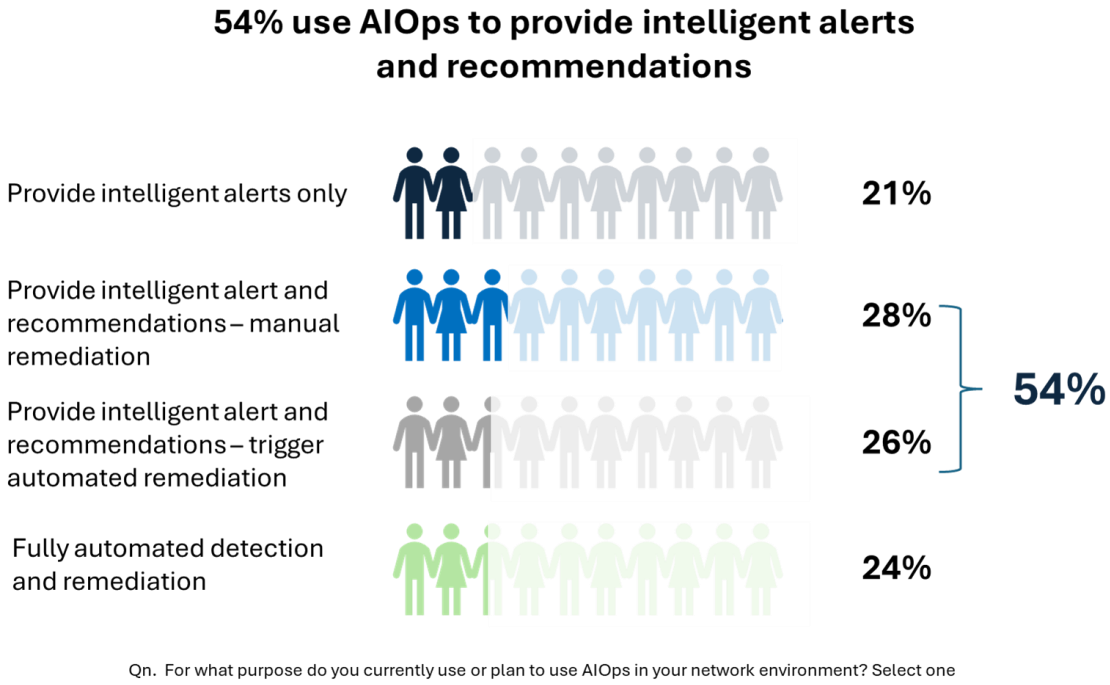

Figure 8 – AIOps usage in the network environment

Juniper AI and automation covers the entire network life cycle (Day 0, 1 and 2) to help accelerate the time to value for Juniper private WANs. As Juniper continues to develop AI-Native routing and automation technologies to accelerate the journey to a self-driving, self-healing, energy-efficient private WAN environment, it is worth highlighting their latest innovations:

- AI-driven anomaly detection and troubleshooting: This includes the ability to auto-detect performance-impacting issues that humans cannot predict or detect using open or closed-loop systems that can locate routing issues and accelerate troubleshooting using AIOps.
- Integration with multiple LLMs: Juniper provides APIs to enable organizations to access their private or public large language models (LLMs) to retrieve internal or public documents. This capability provides greater context for the environment and awareness of specific business processes.
- Intent-based network optimization: Using intent-based policies as a guidepost, organizations can resolve or mitigate issues in open or closed-loop systems by recalculating network paths and reconfiguring network services end-to-end in case a networking event impacts performance or quality.
- Greater levels of power efficiency: Leveraging advanced telemetry and Junos OS capabilities to monitor energy consumption and suggest configurations will improve energy efficiency without a noticeable impact on end-user experience.

To help their customers drive sustainability initiatives, Juniper has also built a robust Circular Economy Portfolio with programs that include Try and Buy, Certified Pre-Owned, and Take Back. This enables organizations to mitigate the risk of new investments, gives them the option to reuse equipment, and ensures that any equipment turned in will be handled appropriately.

All those investments are paying off for Juniper customers who can benefit from a proven and highly performing WAN platform designed with ease of automation in mind and offered with a lower TCO than other vendors. To that point, Juniper commissioned a report[7] to investigate the capital and operational expense (CapEx and OpEx) improvements when utilizing Juniper routers. The study found an 83% improvement in CapEx and a 61% improvement in OpEx while lowering an organization's carbon footprint as compared to another leading vendor. Another report determined TCO over five years, showing 27% additional network savings[8] when using its autonomous capacity optimization.

Our ANGLE

Modern IT and application environments are highly distributed and complex. In addition to connecting to multiple public cloud providers and edge locations, organizations invest in on-premises AI data centers and deploy AI workloads at the edge.

The key to enabling and unifying these distributed environments is the WAN. Private WANs could be the difference between success and failure for organizations in regulated industries and those looking for a competitive advantage in the emerging AI era.

Enterprises need private WAN solutions to ensure the performance of latency-sensitive applications, the quick and secure movement of all data required for GenAI initiatives, and the migration of modern application workloads from public clouds to private data centers as needed. Most importantly, enterprises need private WAN solutions that heavily leverage AI and automation, are easy to deploy and operate, and guarantee performance quality over complex networks.

Juniper AI-Native routing enables enterprises to quickly deploy and efficiently operate private WANs using proven AI and automation technology to ensure energy-efficient solutions, optimized network performance, and differentiated user experiences. Juniper private WANs provide the fastest path to a self-driving network and dramatically reduce upfront and ongoing costs.

2000829-001-EN Jan 2025

[7] https://www.juniper.net/us/en/forms/2024/an-economic-and-environmental-comparison-of-juniper-networks-routers-with-a-leading-competitor.html,

[8] The Economic Benefits of Automating Capacity Optimization in IP NetworkS