

WHITE PAPER

AI-Native Requirements for Modern Networks

By Bob Laliberte, Principal Analyst
Enterprise Strategy Group

January 2024

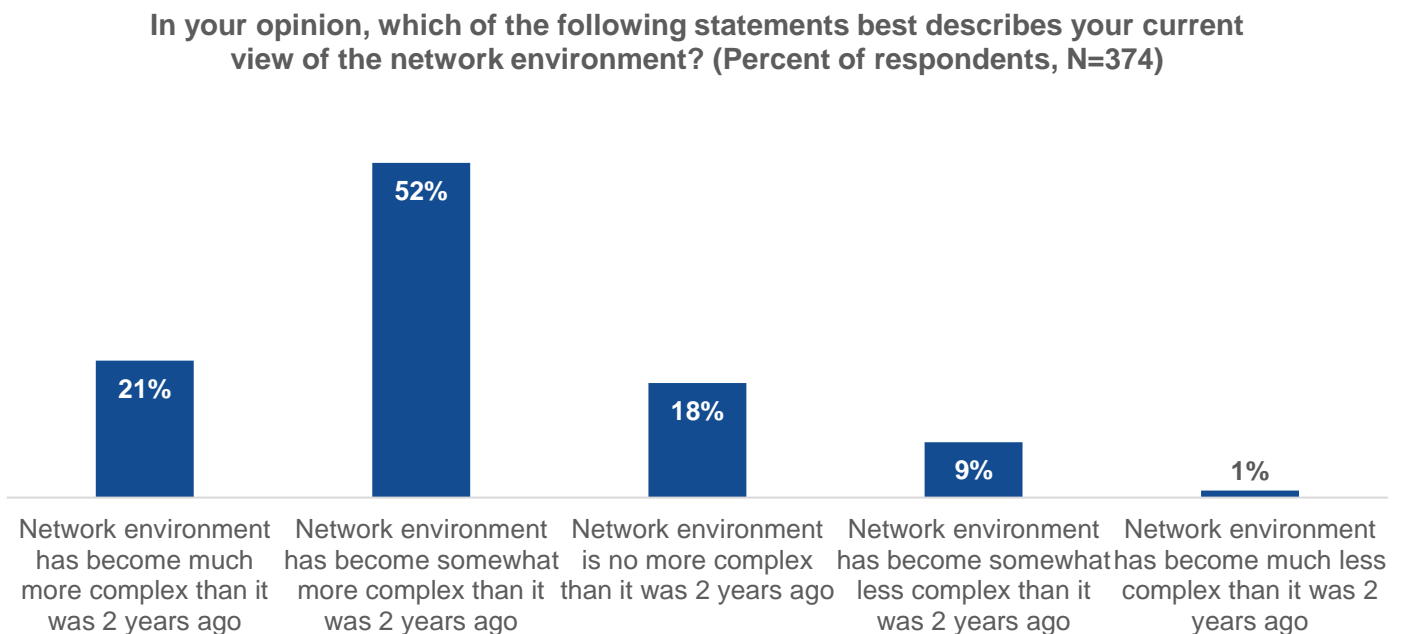
Contents

Modern Network Environments Are Complex	3
Challenges to Implementing AI	4
Technical Challenges	4
Cultural Challenges	4
Process Challenges	6
5 Key Criteria to Drive AI Adoption.....	6
Juniper Networks AI-Native Networking Platform.....	7
Organizations Need to Embrace AI	10

Modern Network Environments Are Complex

There have been multiple drivers in play that have led to the creation of modern network environments. Businesses need to be more agile and responsive to rapidly changing market demands and are increasingly deploying cloud-native application architectures in public clouds and private data centers. Plus, edge computing environments are growing in order to gain real-time business insights, and hybrid work initiatives remain popular. As a result, the network required to support these initiatives has become increasingly distributed and complex. In fact, research from TechTarget’s Enterprise Strategy Group (ESG) shows that almost three-quarters of organizations (73%) stated their network environment has become somewhat or much more complex than it was just two years ago (see Figure 1).¹

Figure 1. Network Complexity Over Time



Source: Enterprise Strategy Group, a division of TechTarget, Inc.

Essentially, the modern network must support every connection across this environment. This would include connecting multiple private data centers, all campus locations, multiple different public cloud providers, as well as potentially hundreds or thousands of branch/edge locations and home offices. As a result, network operations teams are inundated with network data and increasing alarm storms, resulting in alarm fatigue, delays in remediating issues, and poor experiences. Unfortunately, these modern networks create so much data it is virtually impossible to manually correlate and understand what is happening in real time.

This is where network teams need to leverage artificial intelligence (AI) and machine learning (ML) tied to automation to drive operational efficiency, create better experiences, and shift from reactive to proactive or even predictive management. It is important to note that this isn’t about AI/ML replacing the network team but about the network operation team embracing the technology to be more effective. When built correctly, it will dramatically

¹ Source: Enterprise Strategy Group Research Report, [A Network Perspective on SASE and SD-WAN](#), November 2023.

reduce, if not eliminate, alarm fatigue and accelerate the transition to proactive and predictive management of the network environment.

The hype and interest around generative AI (GenAI) is certainly driving more attention to the use of AI, with TechTarget website properties showing a 909% increase in search activity for the topic in 2023.² While this can be seen as a “rising tide” to help drive adoption, there is also a lot of confusion related to this topic and the large language model (LLMs) being used to support it.

Similar to the rise of public cloud computing and the cloud-native environments (e.g., microservices architectures, Kubernetes), we are now seeing the pioneers in the AI space for network operations deliver AI-native solutions. This is an important distinction: AI for network operations should not be confused with the network infrastructure used to support GenAI LLMs (see Ultra Ethernet Consortium³). The latter is about high-performance networking to enable models to be built, while the former is about leveraging AIOps in the network space.

Challenges to Implementing AI

While the benefits of leveraging AI/ML tools are appealing, there are still a number of challenges that can arise in creating the models and limiting adoption in customer environments. The challenges are grouped into three main categories related to technical, cultural, and process considerations.

Technical Challenges

- Some of the most important parts of an AI model are the quality and quantity of the data. Abundant, high-quality data is crucial for training AI models. Unfortunately for the network operations teams, the network operations data has been siloed in customers’ on-premises management systems historically and cannot be shared.
- Adding to the challenge of siloed network operations data is the fact that many network management solutions have not been unified, so data center network data is separate from campus networks, creating more data puddles of unshared data. Without access to end-to-end data, it will be difficult to isolate network events in context to the domain in which they occur. As a result, network vendors might not have a large volume of quality network data to use to create models, despite being in the business for a long time. Even more important is that the vendors are focused on the right data. This goes beyond just “up” or “down” status and being able to detect changes in experience.
- Many organizations are also very concerned about leveraging the benefits of public GenAI models without exposing themselves to risk. While the open nature of the public model enables access and drives growth, it might not be appropriate for sensitive or proprietary data, so organizations require solutions that don’t expose their data.
- Lack of convergence of AI/ML and network automation is another obstacle. While network operations teams have been leveraging network automation solutions for some time, the rise of AI/ML creates a technical challenge of ensuring the technologies can work together across all the network domains. According to ESG research, just over one out of 10 organizations (12%) have tied network AI solutions to their automation capabilities.⁴

Cultural Challenges

- One of the biggest challenges with AI/ML technologies is a cultural hurdle. After becoming experts in correlating network data and determining the root cause of problems over years, many network operations team members are reluctant to trust these models. Much of this is ingrained into the conservative nature of

² Source: TechTarget, *2024 Media Consumption Study: North America*, December 2023.

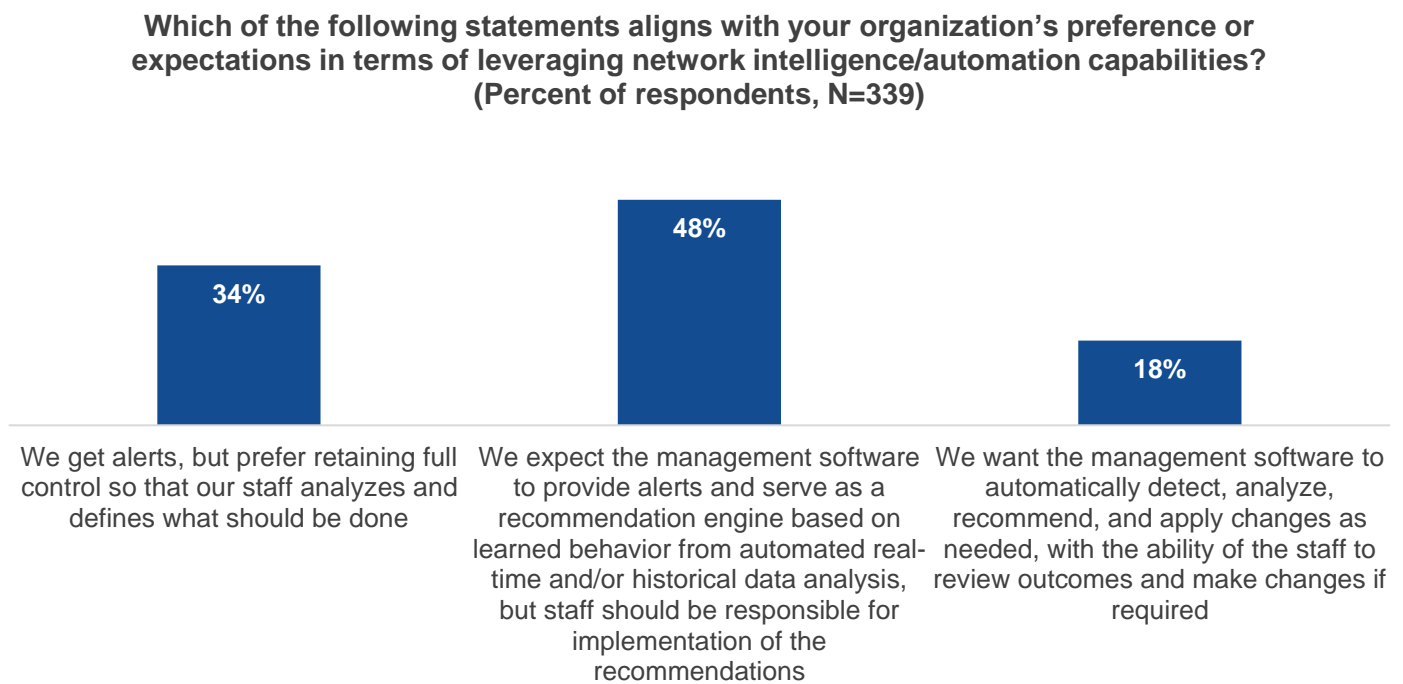
³ [Ultra Ethernet Consortium](#), 2023.

⁴ Source: Enterprise Strategy Group Research Report, [End-to-end Networking Visibility and Management](#), April 2023.

operations teams. When the company is relying on the network to conduct business, the appetite to deploy innovative new technologies will always be met with healthy skepticism. As a result, operations teams need time to become comfortable with the technology and validate that the AI/ML results are in line with their trusted, albeit manual, methods and years of experience.

- If the AI/ML technology is not transparent about how it derived its results, that could further delay full acceptance. Inherently, operations teams don't want to trust the "magic happens here" approach. So, the lack of an explainable AI solution (i.e., one that explains how AI arrived at an alert or recommendation) could delay adoption and the time to become comfortable with the technology.
- Closely tied to the lack of an explainable AI is the lack of a closed-loop control system. Essentially, this makes the network operations team part of the vendor's AI team. Without a closed-loop system, experienced network operators will not be able to provide feedback in a timely manner. This is problematic for the network operations team, as not having this system in place would call for the creation of separate tickets that require more time and effort to resolve. It would also slow down the ability to refine algorithms in a timely manner or, conversely, the ability to confirm and validate that an algorithm is correct. Plus, this type of system is critical for demonstrating operations have returned to normal, regardless of who fixed it.
- AI adoption is progressing slowly. Organizations are beginning to understand the value but want to stay in control of remediation efforts. As ESG research shown in Figure 2 highlights, organizations want to leverage AI for alerts and recommendations, but full reliance on AI to automate a fix is still limited.⁵ (Note: As highlighted above, organizations are in the early stages of integrating AI/ML into automation. This creates a potential intermediate step, where operations teams can manually engage automation to fix a known issue prior to letting it be changed automatically.) Interestingly, organizations interviewed by ESG also cited a lack of AI/ML maturity as the top challenge to using more network automation.⁶

Figure 2. AI Adoption Preferences



Source: Enterprise Strategy Group, a division of TechTarget, Inc.

⁵ Ibid.

⁶ Ibid.

Process Challenges

- These highly distributed and complex network environments mean more frequent and larger alarm storms when an error occurs. As a result, many operations teams or individuals are suffering from alarm fatigue, resulting in the shutting off of alarms, which ultimately can impact experience and create a more reactive environment.
- Many legacy network management systems require knowledge of CLI commands to initiate problem resolution and identify the root cause of a problem. This requires vendor-specific training to become proficient and knowledgeable for each vendor/domain.
- Skilled IT resources are in high demand, and the network domain is no different. Organizations struggle to find skilled resources or lack the time and/or money to get them trained.
- Virtual assistants that leverage natural language processing (NLP) instead of CLI commands are in very limited supply. While industries such as contact centers have been working on improving the technology for years, it is still early for network operations to leverage these tools.

Organizations need to overcome these challenges and work with AI-native network vendors that can deliver real value today.

5 Key Criteria to Drive AI Adoption

For organizations actively researching and investigating AI operations solutions for their network environments, it is important to understand which criteria should be evaluated to determine the efficacy and value of a solution.

Here are five key criteria that should be considered when evaluating AI-native network solutions:

1. **Cloud-based and cloud-native.** While technically these are two criteria, their value is amplified when leveraged together. The network management system should be cloud-based, not on premises. This enables network vendors to collect all the necessary network data (both real and synthetic) across all network domains and all customers (anonymously). It also provides vendors with a wealth of the *right* data, as opposed to an LLM that scrapes the internet for data, thereby supplying network data that can be trusted and used to create more accurate algorithms. Being cloud-native is important, as it shows the solution is not a lift-and-shift operation from on premises to the cloud. Having a management solution that is cloud-native ensures rapid innovation and scale, in addition to relieving operations teams of the burden of lifecycle management, security patches, and bug fixes.
2. **A platform approach to provide end-to-end context.** This is also very important, as the value of the AI solution will increase with each new technology domain added. Having a platform approach eliminates data silos and puddles, ensuring that all the data collected from wired and wireless technology, software-defined WANs (SD-WANs), switches, and routers deployed across the business environment (including data centers, home offices, as well as all campus, branch, and cloud networks) can be used to get both end-to-end visibility and, more importantly, context across these environments and beyond just the network (see criterion No. 5).
3. **Virtual assistant with conversational interface.** Given the shortage in skilled resources and the need to focus on strategic initiatives, network operations teams should understand how virtual assistants with conversational interfaces can provide benefits, including reducing the time and money required to learn vendor-specific CLIs and the ability of virtually any team member to query information quickly, easily, and accurately with efficient human language communication. In addition to assisting network operations team members, these virtual assistants enable organizations to drive convergence with security, helping developers to enable secure connectivity and accelerate development velocity. Plus, C-level executives can quickly get status reports.

4. **Use of trusted data at the granular session layer or application layer.** This also has two components, as AI models require not only a significant volume of relevant data but also that the right questions are asked of the data. The first part is empowered by cloud-based management and ensures that network AI solutions are using data collected from real user network environments and across all domains of wired and wireless environments. Because the vendor is using the telemetry data anonymously, there is no risk to the business. Next, it is imperative that algorithms are focused on the right problems—for example, not just “up” and “down” status but digging deeper into connection times and other factors that could affect a user’s experience. This requires contextual data and, in some cases, might even require a combination of real and synthetic data generated by a network vendor to ensure optimal configuration.

Another key factor to ensure AI generated results are valid is the ability for experienced users to provide feedback. This would include incorporating a closed control loop to help refine algorithms based on feedback from operations teams. Ideally this would be more than just a “yes” or “no” response and enable skilled network members to describe the actual root cause of the problem as well as the verified fix. This is important, as it shows that the network vendors want to work in partnership with the operations team and value their feedback to improve the system. Ultimately, the faster an operations team becomes comfortable with the recommendations the AI solution is generating, the faster it can transition to tightly integrated and automated solutions.

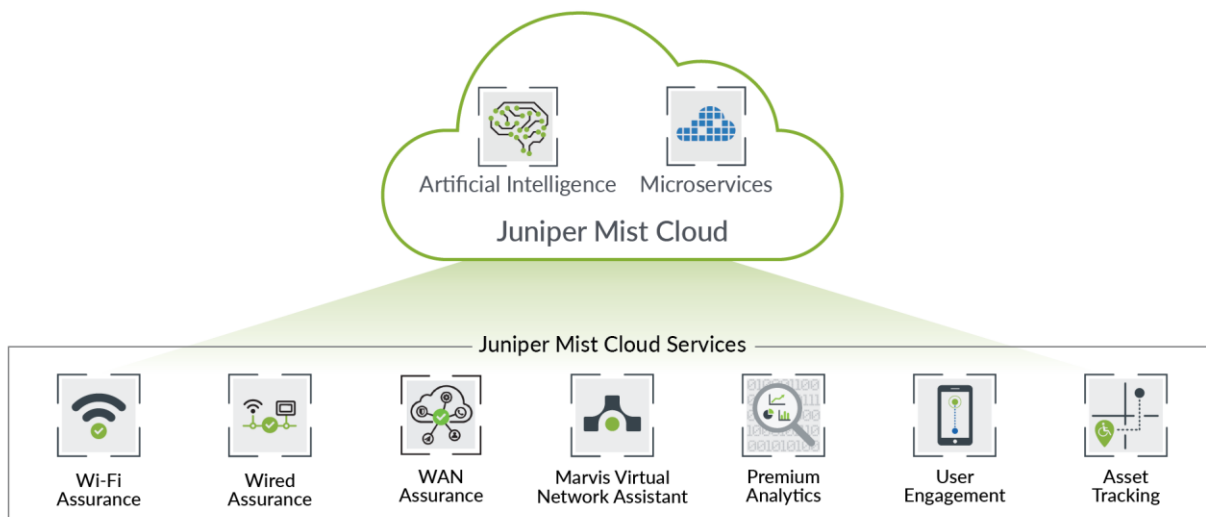
5. **A bidirectional API ecosystem.** Having context across the entire network is important, but that value can be amplified by collected data from adjacent domains, such as collaboration applications or even private 5G networks. Both of these can be accomplished by creating APIs that allow for bidirectional data flows. This not only could help to improve the experience when using voice and video apps but also be used to drive even greater operational efficiencies by tying into workflow automation and security tools.

Juniper Networks AI-Native Networking Platform

Juniper is driven by its mantra of “experience-first networking,” and it has recognized the importance of AI and automation in delivering the best possible experiences across an organization’s end-to-end network.

To accomplish this, Juniper ensures that any network is ideated, architected, and built to drive optimal outcomes by leveraging AI. These are the foundational elements that comprise being AI-native.

Figure 3. Juniper Mist Cloud AI-Native Architecture

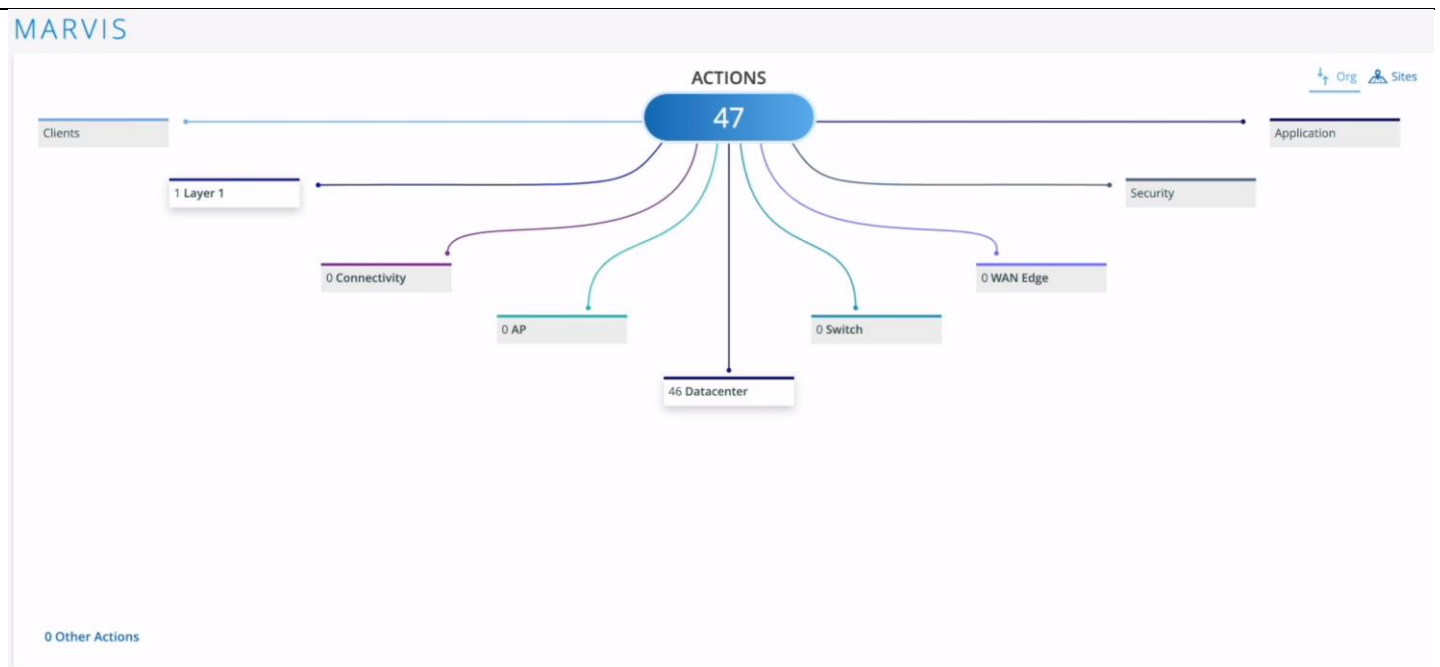


As shown in Figure 3, to deliver against this vision, Juniper has brought together several key cloud-native innovations and technologies to market, including:

- **The Mist-ification of Juniper.** By making the strategic Mist acquisition several years ago, Juniper acquired the talent and technology that would become the building block for a cloud-native, cloud-based AI management platform. In addition to the Wi-Fi technology, Mist now has seven-plus years of experience with AI engines and virtual network assistants (VNAs). Marvis, which is the product name for their VNA, provides Juniper with NLP expertise and extensive network data to build and refine its AI models.
- **A comprehensive, end-to-end, AI-Native Networking platform.** The cloud-native, cloud-based platform started with wireless access points (APs) and quickly incorporated Juniper wired switches, firewalls and SD-WAN technology. It is now expanding to include data center networking, with plans to cover all routers in the WAN environment. This platform has a number of benefits, including:
 - **The ability to collect trusted network data from each additional network domain.** This provides operations teams with additional context and value. Essentially, with each new domain added, the AI model becomes more effective at isolating root causes and ultimately delivering even better experiences. Juniper refers to this as the *flywheel effect*. Juniper's AI engine covers data center, campus, branch, SD-WAN, and (soon) WAN environments.
 - **A customer-focused culture.** The Juniper Mist approach is to have the customer success team tightly integrated with the data sciences team. This ensures that all major issues are pushed to the Marvis AI platform and the data sciences team, ensuring they spend their energy and focus solving customers' top issues. This "experience-first" culture ensures the most important problems are always addressed in a timely manner.
 - **The use of both real and synthetic network data collected from on-premises and cloud environments.** This includes data collected via routers, switches, firewalls, APs, and other sources. The data is analyzed in a cloud-native platform capable of rapidly scaling to analyze massive amounts of data, deriving real-time insights and ensuring enhanced experiences:
 - Granular data collection goes beyond up/down network status, including application- or session-level data to ensure users have the best experience for each interaction. This reflects Juniper's focus on quality of experience versus simply quality of service.
 - Such granular data collection enables organizations to shift from reactive to proactive or, in some cases, predictive management, finding and fixing faults before they create issues. For example, a large public university in Massachusetts shared how its incumbent wired and wireless solution would generate over 200 trouble tickets a semester, leaving the operations team in mostly reactive mode. However, after deploying Juniper, that number dropped to only two tickets per semester. Furthermore, it shared that the AI engine was finding and fixing issues that neither the students nor staff were aware of, ensuring optimized experiences and a far more proactive approach to managing the network. The staff learned of these changes when the closed control loop shared what it found and the actions it had taken, as well as validating that it was running in an optimized state.
 - Juniper is now expanding its comprehensive data collection capability by adding Marvis Minis. Deployed virtually within its wireless APs, this technology simulates users or devices to test the Wi-Fi environment in order to expedite problem resolution or ensure it is ready for business and delivering an optimized experience. Leveraging the synthetic data from the Marvis Minis, organizations can accelerate the adoption of a self-driving network that takes advantage of AI and automation. Juniper Minis are effectively digital twins of users and devices that enable the system to understand the impact of changes without affecting actual users.

- Simplicity of operations.** By incorporating the data center and eventually the WAN, Juniper will provide its customers with a single portal and cloud-based AI platform to cover the end-to-end network environment. This could dramatically simplify network operations, as it eliminates swivel-chair management of multiple different portals and reduces errors from manual correlation. Plus, the AI engine provides rapid insights with context to expedite problem resolution. In addition to the single portal, the Marvis Virtual Network Assistant has also been extended to the data center, reducing the learning curve, and enabling operations teams to take advantage of its conversation interface (see Figure 4). Organizations can now get an end-to-end view and ensure optimized experiences from the application to the user. To further aid network operations teams, Juniper has incorporated its technical documents and manuals into a GenAI model integrated into the portal to further assist users in getting up to speed. This enables Juniper customers to use natural language to find out how to perform specific tasks and provides a simplified experience for the operations teams, reducing the amount of time required to become proficient and enabling them to focus on strategic initiatives.

Figure 4. Marvis Portal



Source: Juniper Networks

Organizations Need to Embrace AI

Modern IT and application environments are driving more complex and distributed network environments. As a result, operations teams can be quickly overwhelmed by alarms and remain stuck in reactive, firefighting mode. To overcome this complexity, organizations require end-to-end visibility and contextual intelligence to transition to a more proactive approach and drive greater operational efficiencies.

Juniper Networks has steadily improved its cloud-native, cloud-based platform to incorporate additional network technologies and domains to drive greater context and efficacy for its AI engine. Today the platform covers data center, campus, branch, and SD-WAN environments, with more planned. It encompasses wired and wireless technologies and leverages both real and synthetic data. The virtual network assistant Marvis, AI engine, and ability to leverage a single portal for end-to-end visibility will dramatically simplify operations and enable operations teams to harness the capabilities of AI and network automation in order to become more proactive and predictive.

Juniper's experience-first approach to technology acquisition and development has enabled it to deliver a cloud-native, cloud-based platform that leverages AI-native capabilities in order to ensure that both operators and users have optimized experiences. Juniper's AI experience and expertise can help any organization accelerate the time to adopt, validate, and recognize value using its comprehensive AI-Native Networking Platform. Organizations need to embrace AI now to streamline network operations, deliver enhanced experiences, and add value to their business.

©TechTarget, Inc. or its subsidiaries. All rights reserved. TechTarget, and the TechTarget logo, are trademarks or registered trademarks of TechTarget, Inc. and are registered in jurisdictions worldwide. Other product and service names and logos, including for BrightTALK, Xtelligent, and the Enterprise Strategy Group might be trademarks of TechTarget or its subsidiaries. All other trademarks, logos and brand names are the property of their respective owners.


Information contained in this publication has been obtained by sources TechTarget considers to be reliable but is not warranted by TechTarget. This publication may contain opinions of TechTarget, which are subject to change. This publication may include forecasts, projections, and other predictive statements that represent TechTarget's assumptions and expectations in light of currently available information. These forecasts are based on industry trends and involve variables and uncertainties. Consequently, TechTarget makes no warranty as to the accuracy of specific forecasts, projections or predictive statements contained herein.

Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of TechTarget, is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact Client Relations at cr@esg-global.com.

About Enterprise Strategy Group

TechTarget's Enterprise Strategy Group provides focused and actionable market intelligence, demand-side research, analyst advisory services, GTM strategy guidance, solution validations, and custom content supporting enterprise technology buying and selling.

 contact@esg-global.com

 www.esg-global.com