

# AI automation: Red Hat, NVIDIA, and Palo Alto Networks collaborate to deliver an integrated, security-first foundation for AI-native telecommunications

[Blog](#) | [Gurpreet Singh](#) | 3-minute read

AI-driven network operations are increasingly becoming central to how service providers manage network complexity, optimize performance, and deliver services faster and more efficiently. AI is no longer just an application-layer technology, it's becoming foundational to how modern service provider's networks function. However, to benefit from this shift, service provider infrastructure must become AI-native, which requires embedding AI directly into network operations, service platforms, and security models across the core, edge, and network. With this in mind, Red Hat is teaming up with NVIDIA and Palo Alto Networks to deliver an optimized architecture, designed for AI from the ground up—creating 2 architectural pillars by using Palo Alto Networks' security capabilities running on Red Hat OpenShift and Red Hat OpenShift AI, powered with NVIDIA AI infrastructure and AI solutions.

## **An integrated, security-first foundation for AI-native service providers**

AI-native service providers require an architecture designed to operationalize AI as part of the network itself across centralized datacenters, edge environments, and network locations.

At the core of this architecture is Red Hat OpenShift and Red Hat OpenShift AI, which serves as the common cloud-native platform for AI-native telecommunications (telcos). Red Hat OpenShift and OpenShift AI provide a consistent, enterprise-grade Kubernetes platform to deploy, manage, and govern both network functions and AI workloads across core, edge, and far-edge environments. This operational consistency allows service providers to introduce AI-driven capabilities without fragmenting operations or creating isolated environments.

AI workloads are powered by NVIDIA-accelerated computing systems deployed across the network. [NVIDIA RTX PRO Servers](#) provide high-performance, energy-efficient AI acceleration for datacenter and edge AI workloads, while the [NVIDIA Aerial RAN Computer](#) (ARC) family of accelerated servers enables AI-driven RAN and network-edge intelligence. Together, these platforms deliver the performance and efficiency required for real-time inference, operational analytics, and AI-native network services.

By pairing Prisma AIRS with OpenShift and [NVIDIA BlueField](#), we enable a powerful synergy of centralized and distributed security enforcement—leveraging [NVIDIA DOCA](#) steering for hardware-level precision—specifically optimized for the demands of AI-driven telecom workloads. The combination is accelerated by [NVIDIA ConnectX](#) to deliver real-time threat detection and policy enforcement. By operating at the infrastructure layer, security is applied consistently across core, edge, and network environments without compromising performance or increasing latency.

# AI automation: Red Hat, NVIDIA, and Palo Alto Networks collaborate to deliver an integrated, security-first foundation for AI-native telecommunications

Rather than relying on loosely coupled components, this approach integrates AI lifecycle management, AI infrastructure, and security-first operations into a single, cohesive platform that operates at scale. With an integrated architecture, service providers can benefit from a consistent foundation to deploy AI-driven automation, analytics, and services wherever they deliver the most value, while maintaining predictable performance, operational simplicity, and built-in security across the distributed network footprint.

Together, Red Hat, NVIDIA, and Palo Alto Networks aim to deliver an integrated, AI-native telco stack where cloud-native software, accelerated computing, and zero trust security are designed to operate as one system. This enables service providers to evolve beyond hardware-centric networks and build a foundation for scalable AI-driven operations, new services, and future 6G-era networks.

## Advancing toward AI-native telco networks

Those service providers that embed AI into network operations and service platforms early can gain lasting advantages in efficiency, agility, and service differentiation.

At Mobile World Congress (MWC), service providers can connect with Red Hat, NVIDIA, and Palo Alto Networks to explore how AI-native architectures can accelerate operational transformation and service innovation. These conversations can help you to identify high-impact use cases such as AI-driven RAN optimization, predictive maintenance, energy efficiency, and automated security enforcement, and show how cloud-native platforms, accelerated computing, and built-in security come together to support distributed AI workloads.

The transition to an AI-native service provider is a journey, not a single deployment milestone. By making AI a core part of how networks are designed and operated, service providers can position themselves as AI-native businesses—capable of delivering operational excellence, new services, and long-term growth in an AI-driven economy.

## About the author



### [Gurpreet Singh](#)

Principal Product Manager

Experienced Product Manager renowned for driving innovation in network and cloud technologies. Expertise in strategic planning, competitive analysis, and product development. Skilled leader fostering customer success and shaping transformative strategies for progressive organizations committed to excellence.

Highly engaged in the pioneering stages of ETSI NFV standards and OPNFV initiatives, specializing in active monitoring and ensuring reliability in virtualized network deployments.

[Read full bio](#)