AI from the core to the edge

How communications service providers are building smarter networks and processes

Contents

Executive Summary

Artificial intelligence (AI) is creating vast new opportunities for communications service providers (CSPs), from faster, self-optimizing networks to intelligent chatbots that transform their customers' experience. With the number of AI workloads expected to substantially increase over the coming years, CSPs who can efficiently implement them stand to expand their market opportunities and generate significant competitive advantage. For example, a recent State of AI in Telecommunications survey by NVIDIA found:

84% of respondents said AI is helping to increase their company's annual revenue.[1]

77% reduced annual operating costs with the help of AI.[1]

65% plan to increase spending on AI infrastructure in 2025.[1]

The emergence of AI offers an intelligent solution to today's siloed and increasingly complex networks. CSPs currently face a "beyond human scale" challenge, where staff lack the resources to monitor, optimize and fix network and business functions in real time, 24/7. AI can automate these tasks, leading to improved network performance, stronger customer retention, operational efficiency gains and more.

This eGuide will show how CSPs can take advantage of the hardware and software needed to simplify and accelerate their path to AI innovation.

97% of respondents said they're adopting AI, while nearly half are already deploying it. Of respondents who are assessing or have deployed GenAI, **54% say they've already deployed their first GenAI service or application**.[1]

Take the lead with Telecom-optimized AI solutions

The challenge for CSPs is to develop infrastructure that can handle these AI workloads, whether in their data centers or at the edge. But producing this kind of optimized architecture requires an ecosystem of partners, hardware and solutions that can turn ideas into reality – and ROI – faster.
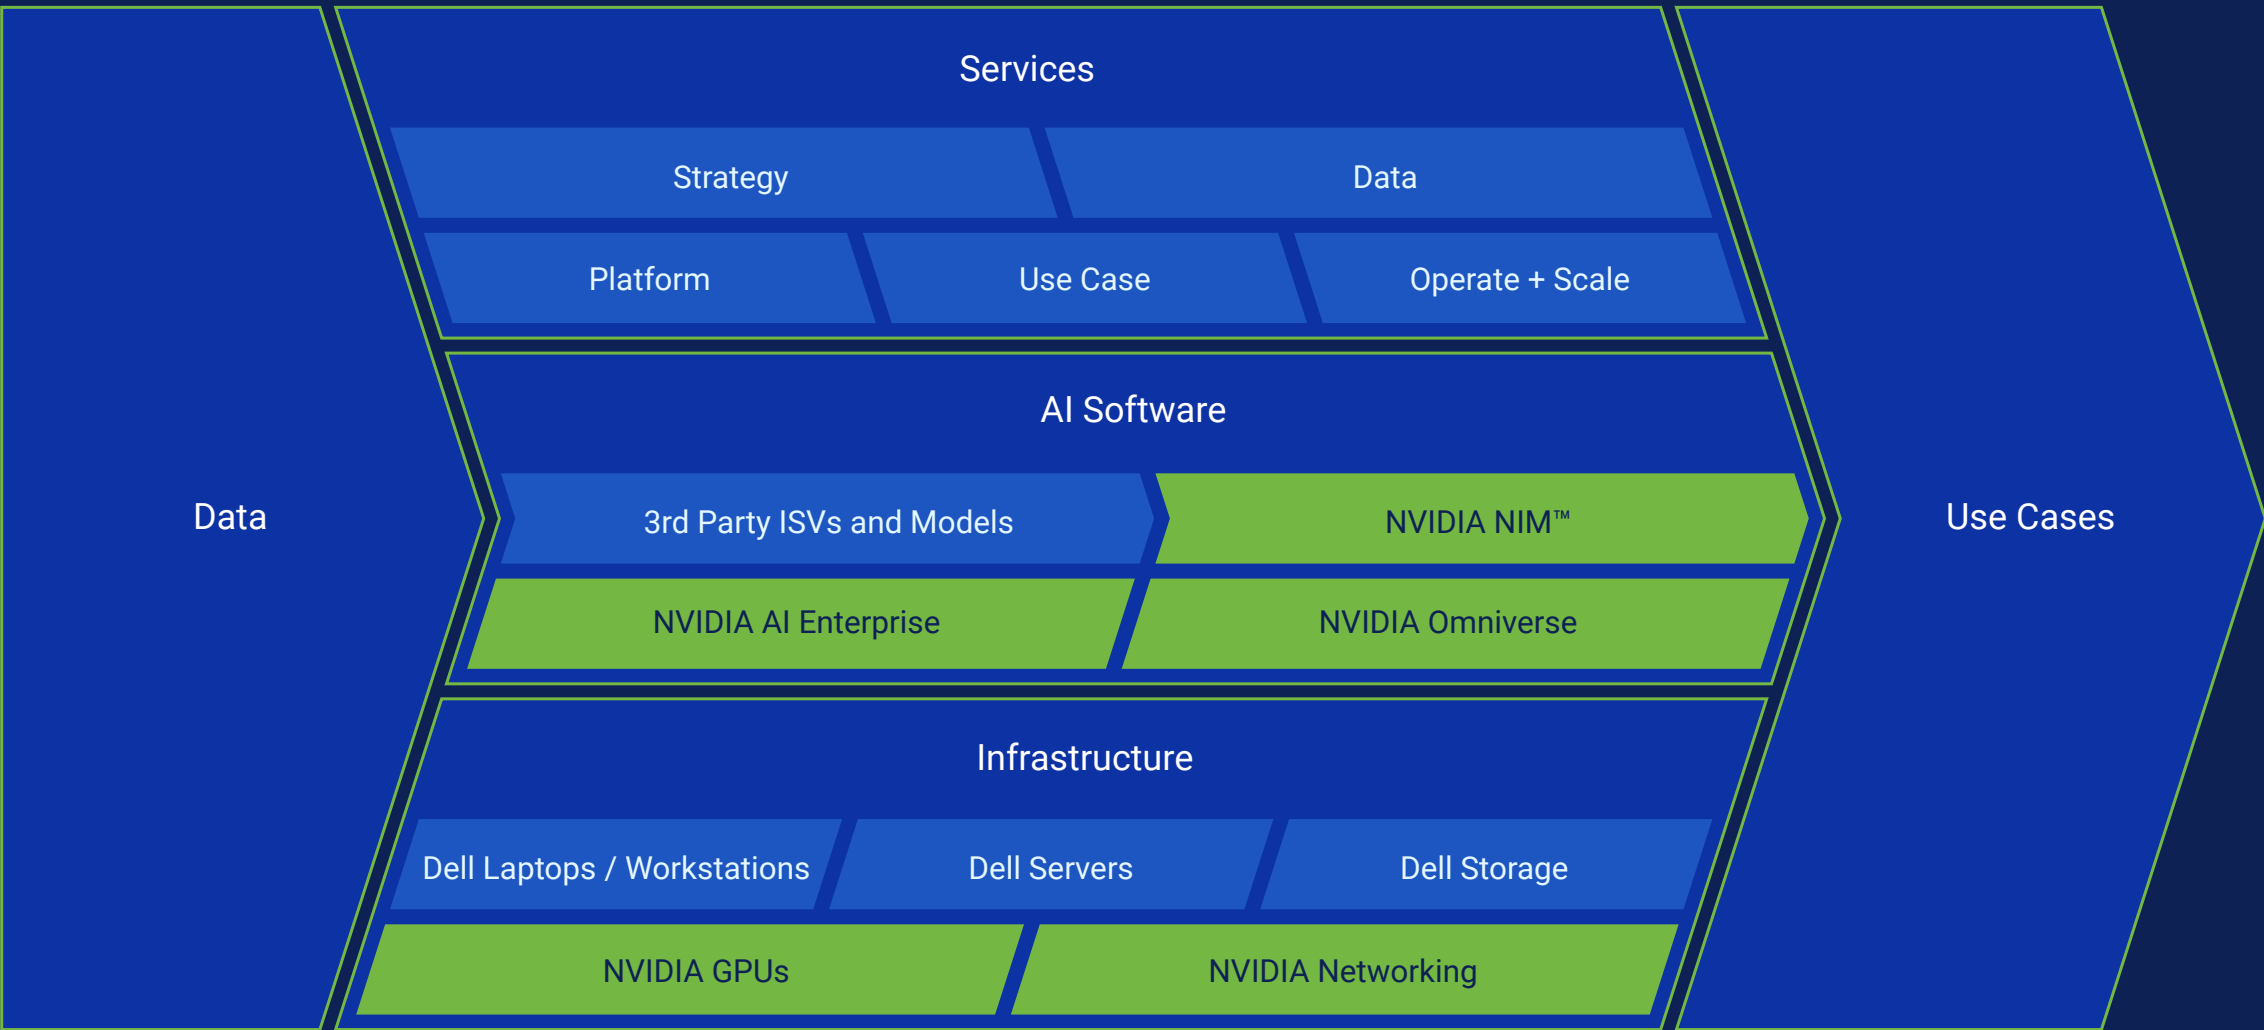
Dell AI Factory with NVIDIA

Dell AI Factory with NVIDIA gives CSPs the necessary building blocks to run any AI workload, anywhere. It combines Dell's broadest portfolio of solutions with expertise from NVIDIA and other industry-leading partners:

Dell Technologies
Servers, storage and data management technologies.

NVIDIA
Accelerated computing, networking, NVIDIA AI Enterprise, NVIDIA Omniverse.

Partner AI Software
Third party ISVs and models.

Data

Services

Strategy	Data
Platform	Use Case	Operate + Scale

AI Software

3rd Party ISVs and Models	NVIDIA NIM™
NVIDIA AI Enterprise	NVIDIA Omniverse

Infrastructure

Dell Laptops / Workstations	Dell Servers	Dell Storage
NVIDIA GPUs	NVIDIA Networking	

Use Cases

The Dell AI Factory with NVIDIA framework allows organizations greater control over their proprietary data and scales efficiently, providing a more affordable alternative to many public cloud solutions. It delivers a complete solution tailored to meet specific customer use cases or outcomes, including in the Telecom industry.

Innovating with AI from the data center to the edge

AI enables CSPs to maximize network performance at the lowest possible cost from data center to edge, improve customer service and loyalty, and deliver greater value at the enterprise edge.

Core data centers

- Capable of hosting more intensive AI workloads that can collect and model data to enhance network performance, energy efficiency, etc.

- Operating out of data centers can reduce costs associated with the training and inferencing of large language models (LLMs).

- Monetization of data center investment via the creation of compute services such as GPU-as-a-Service (GPUaaS).

Edge

- Ruggedized hardware enables CSPs to move AI workloads closer to where their data is created, including space-limited environments.

- Low latency environments are ideal for AI applications that require real-time processing.

- With AI at the edge, real-time decision-making can come to life, fuelling cost optimization and streamlining processes across all industries.

AI addresses challenges from the data center to the edge, where Dell Telecom-optimized hardware with NVIDIA accelerated computing can enable innovative, revenue-generating services.

Data center use cases

1. Amdocs AI Factory

Telecom use case

To address sovereign AI and enterprise-specific generative AI monetization opportunities, service providers need a robust generative AI value chain from infrastructure to monetization. The Amdocs AI Factory offers telcos full control over data and tenant separation, meets regulatory and data protection requirements, optimizes infrastructure, and simplifies the creation of market tailored models and services. Leveraging strategic collaborations with NVIDIA and Dell, it provides the full ecosystem needed to create innovative services, offering flexibility for emerging enterprise services.

Hardware/software stack

Dell PowerEdge XE Series Servers

PowerScale Storage platform

NVIDIA H200 Tensor Core GPUs

Dell ProDeploy, ProSupport

amAIz™ GenAI Suite

NVIDIA AI Enterprise

Amdocs AI Factory enables service providers to commercialize Gen AI across industries, in the form of GPU as a Service, S/LLM as a Service, vertically tailored applications, and more. This comprehensive end-to-end solution is designed to help service providers address the surging enterprise demand for AI-driven infrastructure and applications, enabling them to unlock significant new revenue streams and establish a competitive edge in the fast-growing market for generative AI services. Amdocs amAIz Suite unlocks the full potential of GenAI in the agentic area, utilizing a holistic yet modular approach that breaks down data silos – integrating AI/ML, data and analytics, together with GenAI platforms and apps.

Partnership benefits

- Trusted, Telecom-grade, enriched with telco skills.
- Combining cutting-edge technology with deep Telecom expertise drives innovation and efficiency.

Customer benefits

- **Improved Customer Service:** Rapid resolution of customer inquiries and issues with AI driven Copilots.
- **Enhanced Efficiency:** Increases productivity through automated decision-making and task management.
- **Time and Cost Savings:** Enhanced service offerings with AI apps that streamline operations and reduce operational costs.

"There's tremendous opportunity in the AI factory space for telecom providers. With NVIDIA NeMo and NIM microservices designed to speed the adoption of generative AI, the expansion of the new amAIz Gen AI platform to Amdocs AI Factory will help service providers deploy and monetize new business streams with greater ease."

– Christopher Penrose, Global VP of Business Development for Telco at NVIDIA

2. Amdocs amAIz Agents
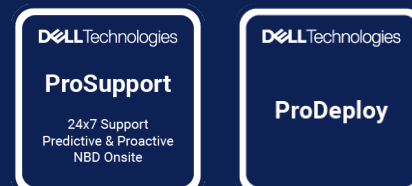
amdocs   **DELL**Technologies   **NVIDIA.**

Telecom use case

In key areas of the communications industry such as Care, Home & Network, Sales and Marketing, digital employees can be used to answer queries, streamline processes and improve efficiency – ultimately leading to enhanced operations and customer satisfaction.

Hardware/software stack
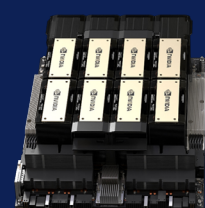
Dell PowerEdge XE Series Servers

Dell Technologies ProSupport 24x7 Support Predictive & Proactive NBD Onsite

Dell Technologies ProDeploy

Dell ProDeploy, ProSupport

Dell PowerScale Storage platform

Amdocs amAIz Suite

Amdocs amAIz Agents

NVIDIA H200 Tensor Core GPUs

NVIDIA.

NVIDIA AI Enterprise

Proven KPIs for amAIz Care Agents Business

- **63% reduction** in average handling time
- **49% improvement** in tNPS
- **50% reduction** in repeated calls

Performance

- **60% reduction** in tokens consumed
- Up to **30% improved accuracy**
- Up to **80% improved latency**

This Telecom-native GenAI platform offers amAIz Agents that can act autonomously, make telco-aware decisions, and learn and optimize processes while making intelligent recommendations. They seamlessly orchestrate among themselves, retrieving the data they need while ensuring the right task is routed to the relevant agent.

Partnership benefits

- Leveraging Amdocs: Telco terminology and ontology.

- Enabling proactive automations across care, sales, marketing, and network domains, by using Agents to optimize customer experiences and business operations.

- Dell's hardware and NVIDIA's technology optimize performance in essence of latency and token consumption with NIM microservices.

"Dell PowerEdge servers and our automation software create a flexible network foundation to boost telecom AI adoption. Through our collaboration with Amdocs and NVIDIA, communications service providers can unlock the potential of AI to drive industry innovation and positive business outcomes."

– **Dennis Hoffman, Senior Vice President and General Manager, Telecom Systems Business, Dell Technologies**

Customer benefits

- **Care Agents:** Enhance customer service by managing inquiries, interactions, disputes and billing with personalised experiences.

- **Sales Agents:** Facilitate product discovery and sales, boosting revenue and enabling upselling opportunities.

- **Marketing Agents:** Oversee campaign planning and personalization, competitive analysis, and market attribution to optimize marketing strategies.

- **Network Agents:** Analyze network resources and issues, proactively resolving problems to maintain network integrity.

- **Home Agents:** Unlock revenues with AI-based home insights by analyzing customer behavior at home, supporting tech-related inquiries, and enabling personalized upsell offers to resolve issues.

3. Network Troubleshooting and Analysis

kinetica **DELL**Technologies **NVIDIA.**

Telecom use case

CSP staff need optimized data analysis and management across 5G and OpenRAN applications, but a lack of SQL experience can prevent them from running complex queries on their large datasets. Natural language tool SQL-GPT helps engineers, support staff, management, finance, and marketing teams to design, test, debug, and troubleshoot network components such as CU, DU, RU, 5G Core, RIC and O-RAN interfaces.

Hardware/software stack

**Dell PowerEdge
XE Series Servers**

**PowerScale
Storage platform**

**NVIDIA H200 Tensor
Core GPUs**

DELLTechnologies
ProSupport
24x7 Support
Predictive & Proactive
NBD Onsite

DELLTechnologies
ProDeploy

kinetica

NVIDIA.

Dell ProDeploy, ProSupport

Kinetica SQL-GPT

NVIDIA AI Enterprise

An advanced conversational AI powered by Kinetica, SQL-GPT allows everyday language, rather than SQL commands, to be used for complex data queries. It offers unmatched speed, accuracy and scalability across 5G and OpenRAN environments, providing real-time data analysis that accelerates decision-making.

Partnership benefits

- **Enhance Integration:** SQL-GPT integrates seamlessly with various data platforms and tools, enhancing the capabilities of existing systems.

- **Collaborative Innovation:** Partners can use SQL-GPT's advanced features to develop new solutions and services, driving innovation in the Telecom sector.

- **Market Differentiation:** Offering SQL-GPT as part of a partner's portfolio can differentiate them in the market, providing a competitive advantage with cutting-edge AI-driven data management solutions.

- **Shared Success:** Partners benefit from joint marketing and sales efforts, expanding reach and driving mutual growth.

Customer benefits

- **Enhanced Efficiency:** Automates and accelerates data query and analysis, reducing the time and effort required for network management, and freeing up time for engineers.

- **Improved Accuracy:** Provides precise and reliable answers, minimizing errors in data interpretation.

- **User-Friendly:** Natural language processing allows users to interact with the system using everyday language, reducing the need for specialized SQL knowledge.

- **Scalable:** Capable of handling data from gigabytes to zettabytes, making it suitable for both small and large-scale deployments.

- **Secure:** Ensures data security by processing queries within the customer's premises or cloud perimeter.

- **Better Customer Experience:** Faster Time to customer resolution and improved NPS scores.

A major telco company reported a 40% reduction in troubleshooting time and a 30% increase in operational efficiency after implementing SQL-GPT.

4. Achieving Autonomous Network Operations with ISNA

Infosys Navigate your next

DELLTechnologies

NVIDIA.

Telecom use case

Telecom providers face rising customer expectations for seamless digital experiences, while also managing the escalating cost of network operations, which can account for up to 30% of total spend. Networks must evolve to support 5G and edge technologies, but the complexity of managing large-scale, heterogeneous systems increases operational risk. Traditional monitoring approaches are no longer sufficient to prevent disruptions or control costs. To address these challenges, CSPs need solutions that improve visibility, reduce manual intervention, and automate network operations while maintaining service quality and reliability.

Hardware/software stack

Dell PowerEdge R760 Servers

Dell PowerEdge XE Series Servers

NVIDIA H200 Tensor Core GPUs

DELLTechnologies
ProSupport
24x7 Support
Predictive & Proactive
NBD Onsite

DELLTechnologies
ProDeploy

Infosys Navigate your next

NVIDIA.

Dell ProDeploy, ProSupport

Infosys Smart Network Assurance (ISNA)

NVIDIA AI Enterprise

Platform, Software & Microservices

Ubuntu 22.4 LTS+, RHEL 8, Kubernetes, Java, Python, nodeJS, vLLM, etc.

Infosys Smart Network Assurance (ISNA) is a cloud-native AIOps platform that combines agentic AI with supervised and unsupervised AI/ML models. It detects anomalies, performs root cause analysis, and triggers remediation. Delivered as plug-and-play, ISNA reduces costs, accelerates autonomous operations, and strengthens service reliability across complex telecom environments.

Partnership benefits

- **Flexible integration:** Designed as a flexible AIOps platform to complement Telcos' existing OSS and network environments and address gaps and challenges without disruption.

- **Accelerated innovation:** Pre-built capabilities help partners guide customers toward an autonomous operations maturity of 4.

- **Differentiated offering:** Delivers a pre-verified "Network AI in a Box" that partners can package into portfolios.

- **Shared growth:** Joint go-to-market programs broaden reach and strengthen partner success.

Customer benefits

- **Rapid deployment:** Packaged solution integrates seamlessly with assurance stacks for quick adoption.

- **Autonomous operations:** Agentic AI studio pre-packaged with specialized AI agents supports fault detection, impact analysis, troubleshooting, triaging, root cause analysis, cross domain correlation, anomaly remediation, and self-healing.

- **Efficiency gains:** Reduces monitoring and fault remediation effort, lowering costs by 30–40%.

- **Faster resolution:** Automated root cause analysis, advanced Agentic AI features and low-code/no-code automation framework cuts MTTR by up to 30%.

- **Scalability:** Cloud-native microservices handle data at any scale with flexible on-prem, cloud, or hybrid deployment via standard APIs.
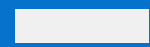
Europe's largest Cable MSO: Improved operator efficiency

Projected annual ROI of **39.5%**

20% improvement in operator efficiency

25% reduction in MTTR

Projected savings of **1M over 3 years** through consistent SLA adherence

Edge use cases

1. Computer Vision and GenAI at the Edge

DELLTechnologies   **NVIDIA.**

Telecom use case

CSPs need flexible and rugged compute at their AI platform edge locations. This enables their customers to implement computer vision and GenAI applications such as object detection, image classification and 3D mapping.
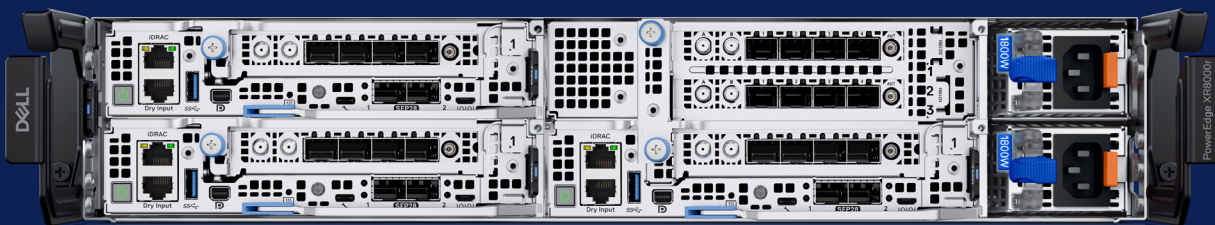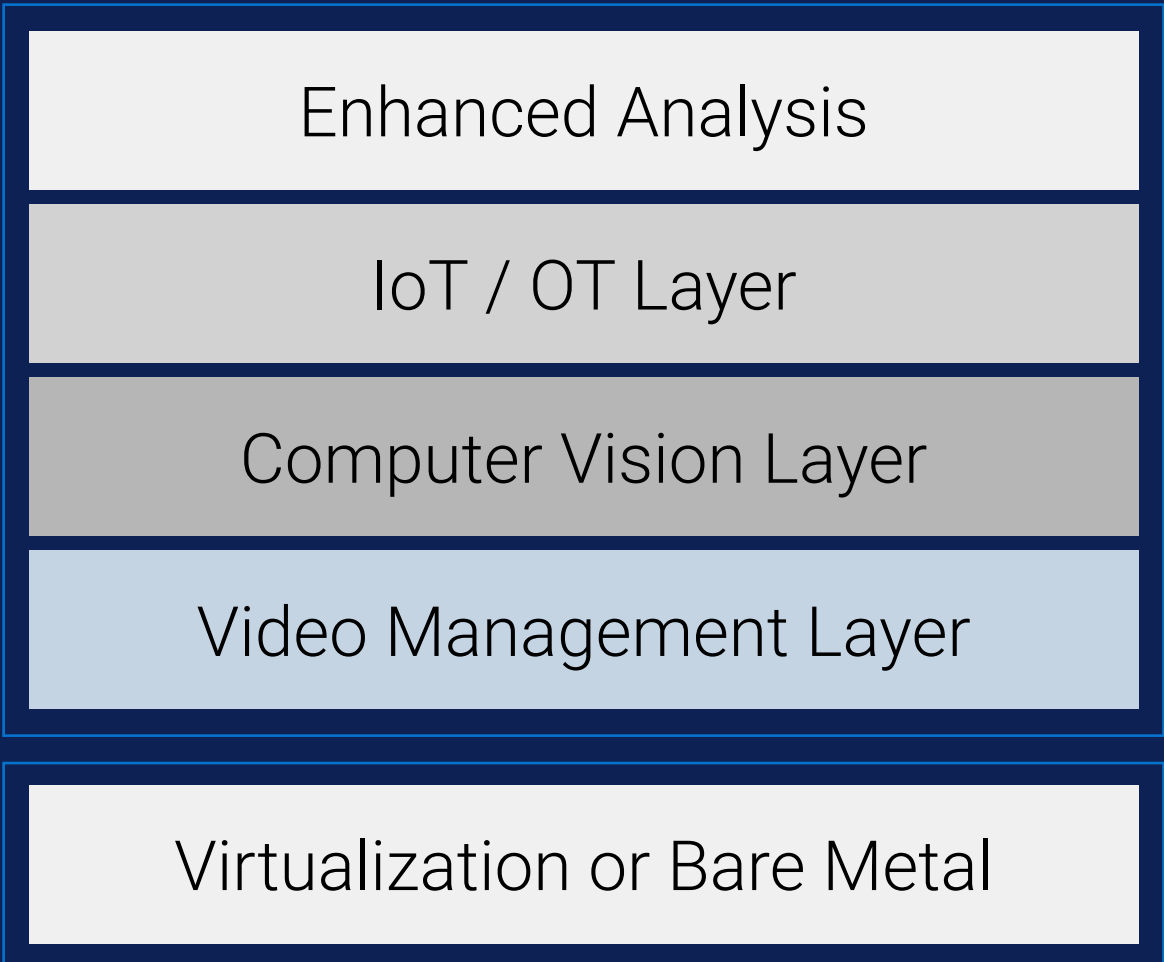
Hardware/software stack

NVIDIA L4 Tensor Core GPUs

DELLTechnologies
ProDeploy

DELLTechnologies
ProSupport
24x7 Support
Predictive & Proactive
NBD Onsite

Dell ProDeploy, ProSupport

NVIDIA.

NVIDIA Morpheus

Enhanced Analysis

IoT / OT Layer

Computer Vision Layer

Video Management Layer

Virtualization or Bare Metal

Dell PowerEdge XR8000 Server

The PowerEdge XR8000's 2U SLED based, multi-node, flexible and rugged compute platform is ideal for demanding computer vision, AI/ML and GenAI edge applications. Edge AI cybersecurity with the NVIDIA Morpheus framework supports real-time data monitoring with AI and threat detection.

Partnership benefits

- XR8000 is a carrier-grade flexible configuration server platform.

- Turnkey Purchase-to-Deploy.

- Dell ProDeploy and ProSupport.

Customer benefits

- **Flexibility:** 2U chassis architecture supports a mix of 1U and 2U computer sleds.

- **Designed for Limited-Space Environments:** Short depth (430mm) Class 1 server.

- **Charging Options:** Both AC and DC power supply options.

- **Rugged and Resilient:** NEBS Level 3 certification.

- **Temperature Ready:** Operates in environments between -20ºC and 65ºC.
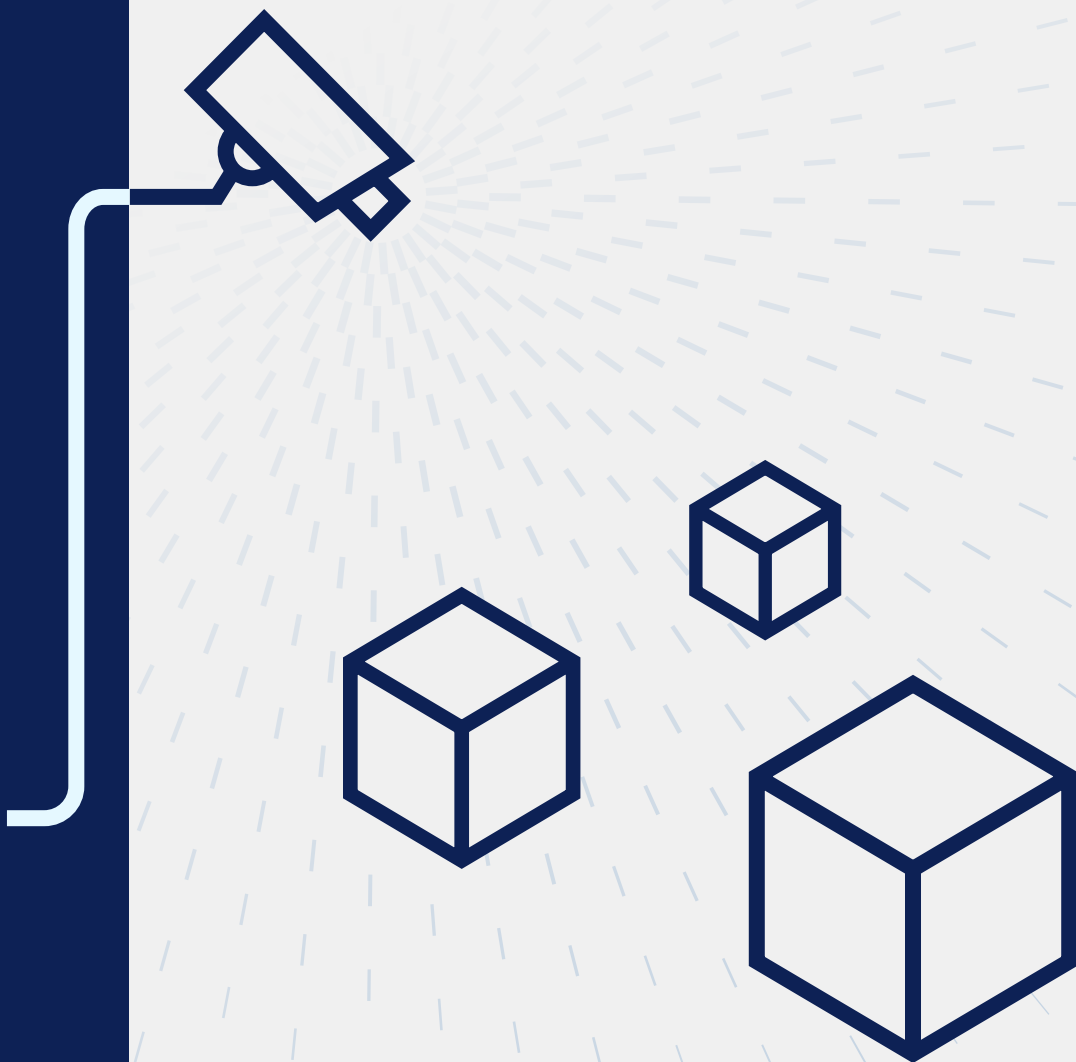
2. AI-Powered Production

Telecom use case

With AI-powered production, CSPs can combine innovative edge and AI solutions that run on private wireless networks to its enterprise customers. By incorporating AI workloads into the manufacturing floor, data can be acted upon closer to its creation, resulting in enhanced efficiency, worker safety and cost-savings. Within manufacturing chains, AI can detect, locate and address hazards faster, and enable predictive maintenance without losing production momentum.

Hardware/software stack

Dell PowerEdge XR8000 Server

NVIDIA L4 Tensor Core GPUs

Dell Edge Gateway

Dell ProDeploy, ProSupport

NVIDIA DeepStream SDK

NVIDIA DeepStream SDK is a complete streaming analytics toolkit that enables a faster, simpler way of developing AI vision-based imaging applications that can be deployed at the edge. IoT sensors capture, decode and process data on the manufacturing floor and transfer to AI inferencing models to produce real-time analytics and actions.

Partnership benefits

- PowerEdge XR8000 enables NVIDIA L4 Tensor Core GPUs to be deployed in edge environments for real-time AI video image processing.
- The XR8000 and NVIDIA L4 Tensor Core GPU combination can be applied to perform error detection, object detection and image classification in any manufacturing process.

Customer benefits

- **5G and AI Monitoring:** Enables immediate data exchange between different equipment, machinery and vehicles, facilitating quick response time for alert remediation.
- **Real-Time Feedback and Actions:** Results enable near real-time adjustments where needed, to get production back on track fast.
- **Streamlined Operations:** Frees up time and energy for innovation.

A manufacturing plant reduced its operating costs by **3%** and labor costs by **25%** using AI and computer vision.[2]

3. Dynamic Real-Time Management for Safer Sites

DELLTechnologies   **NVIDIA.**   **Khasm**   **Archetype**^{AI}

Telecom use case

CSPs need the ability to monitor security of their offices, cell sites and other urban locations, maintain the safety of people using them and respond to unfolding incidents in real time.
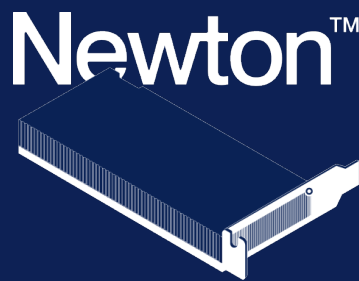
Hardware/software stack

Dell PowerEdge XR8000 Server

NVIDIA L4 Tensor Core GPUs

Newton™

DELLTechnologies
ProSupport
24x7 Support
Predictive & Proactive
NBD Onsite

DELLTechnologies
ProDeploy

**Newton AI Platform running on
NVIDIA L40s GPUs**

Dell ProDeploy, ProSupport

Khasm and Archetype AI are implementing a scalable solution with the City of Bellevue across multiple intersections. Archetype AI's advanced physical AI lens, trainable for various incident and safety management scenarios, enhances the solution's versatility, with pedestrian safety management as one key application. Khasm's solution uses GenAI to build intelligence about incident-heavy areas, delivering AI-driven insights and chat-based recommendations in response. L40S GPUs deliver the high-performance intelligence and processing needed to run the application, SLM and data transformation procedures, enabling organizations to create safer environments.

Partnership benefits

- PowerEdge XR7620 with NVIDIA L40S GPUs are an ideal combination for edge deployments, offering the versatility and power needed for real-time AI video image processing.

- The XR7620 and L40S GPU combination can be applied to monitor pedestrian and traffic behavior, creating automated alerts for high-risk situations.

- Seamlessly integrates with smart systems for enhanced decision making.

Customer benefits

- **Improved Safety:** Significant reduction in late crossings and near-miss incidents.

- **Enhanced Efficiency:** Optimized traffic flow during peak hours with adaptive signal timing.

- **Actionable Insights:** Data-driven decisions for urban planners and traffic managers.

- **Scalable Impact:** A model ready for expansion to other sites and infrastructure.

- **Continuous intelligence:** Edge AI powers a range of use cases scaling from facilities to cities, providing real-time insights that enhance decision-making.

4. BrewVerse: Modernizing Manufacturing with AI

XMPRO   **D∕∕LLTechnologies**   **NVIDIA.**

Telecom use case

Manufacturers are under pressure to boost efficiency, ensure worker safety, and manage costs, but traditional production systems lack the processing capacity to analyze and act on data in real time. Delays in detecting hazards, inefficiencies in workflows, and costly downtime often result from relying on reactive processes. CSPs can help overcome these challenges by helping extend AI workloads to the manufacturing floor over private wireless networks. Processing data closer to its source enables faster responses, predictive maintenance, and greater operational intelligence, giving enterprises a foundation for safer and more cost-effective production.

Hardware/software stack

Dell PowerEdge XR8000 Server

NVIDIA L4 Tensor Core GPUs

XMPRO

DELLTechnologies
ProDeploy

DELLTechnologies
ProSupport
24x7 Support
Predictive & Proactive
NBD Onsite

NVIDIA.

XMPro for Agentic AI   **Dell ProDeploy, ProSupport**   **NVIDIA Omniverse for Spatial Digital Twins**

AI-powered production combines NVIDIA Omniverse with agentic AI on private wireless networks to bring intelligence to the edge. Omniverse builds digital twins for real-time simulation and workflow optimization, while agentic AI enables autonomous decision-making, allowing systems to respond dynamically to changing conditions on the production floor. Together, they reduce downtime, enhance safety, and drive predictive maintenance for more efficient, cost-effective manufacturing.

Partnership benefits

- **Edge-ready infrastructure:** PowerEdge XR8000 enables NVIDIA L4 Tensor Core GPUs to be deployed in rugged environments for agentic AI applications.

- **Optimized performance:** XR8000 and NVIDIA L4 together deliver AI capabilities tailored for manufacturing settings, preventing failures and improving workflows.

- **Cutting-edge collaboration:** Omniverse empowers partners to offer real-time simulation and AI-driven collaboration, strengthening innovation and competitiveness.

Customer benefits

- **Autonomous optimization:** XMPro's agentic AI continuously observes, reflects, plans, and acts to improve operations.

- **Failure prevention:** AI agents anticipate risks, prevent downtime, and protect worker safety.

- **Industrial expertise:** Purpose-built AI agents combine generative AI with domain knowledge to solve complex manufacturing challenges.

- **Trustworthy Intelligence:** XMPro's agentic AI is designed for safe industrial use. Every decision is transparent, so teams can see how trade-offs were made. The reasoning is clear and auditable, not hidden in a black box. The system stays consistent under changing conditions, delivering predictable outcomes when it matters most. And with safety enforced by design, critical limits can never be bypassed. Together, this ensures autonomous operations earn operator confidence and meet enterprise standards.

A craft brewery faced the challenge of integrating real-time data from diverse equipment and legacy systems to optimize production. By deploying a digital twin powered by Dell edge compute, XMPro for streamlining sensor data, and NVIDIA Omniverse for advanced simulation and visualization, the brewery achieved seamless operations oversight, predictive analytics, and near instantaneous troubleshooting. As a result, production quality improved, inefficiencies were reduced, and teams collaborated more effectively across the brewing process.

Deliver innovative, revenue-generating services with Dell AI for Telecom

As CSPs seek to optimize their networks and maximize their ROI, they need compute infrastructure ready for AI workloads, from the data center to the edge.

Dell AI with NVIDIA helps CSPs to develop outcome-focused solutions for Telecom-specific needs, such as:

Network performance and optimization

Customer retention

Monetizing the enterprise edge

Working with NVIDIA and Telecom partners, Dell combines industry-leading technology, expertise, experience, and partnerships for CSPs to simplify and accelerate their AI adoption sustainably.

Learn how CSPs can realize the full potential of AI with our purpose-built Telecom and AI solutions – visit **Dell.com/Telecom-AI** today.

[1]   https://www.nvidia.com/en-gb/lp/industries/telecommunications/state-of-ai-in-telecom-survey-report/

[2]   https://infohub.delltechnologies.com/en-us/l/telecom-co-innovation/innovative-edge-ai-and-5g-solutions-to-simplify-manufacturing/

D&LL Technologies