DRIVENETS

DriveNets Network Cloud

From Monolithic Routers to Cloud-native Networks

WHITE PAPER

Problem & Solution

As their networks grow more complex, tier-1 service providers (SPs) are facing continually rising operational costs, flat revenues, and shrinking margins. Traditional hardware-centric networks based on monolithic routers demand high upfront investment and specialized maintenance, driving up capital (CapEX) and operating expenses (OpEx) even as revenue flattens.

DriveNets was founded to tackle these commercial and operational challenges. Its vision is to fundamentally shift networking from proprietary hardware toward fully automated, cloud-based architectures, thereby reducing cost structures and enabling new growth. Instead of merely swapping one proprietary box for another, DriveNets supports a redesign of network architecture into an open software-centric model.

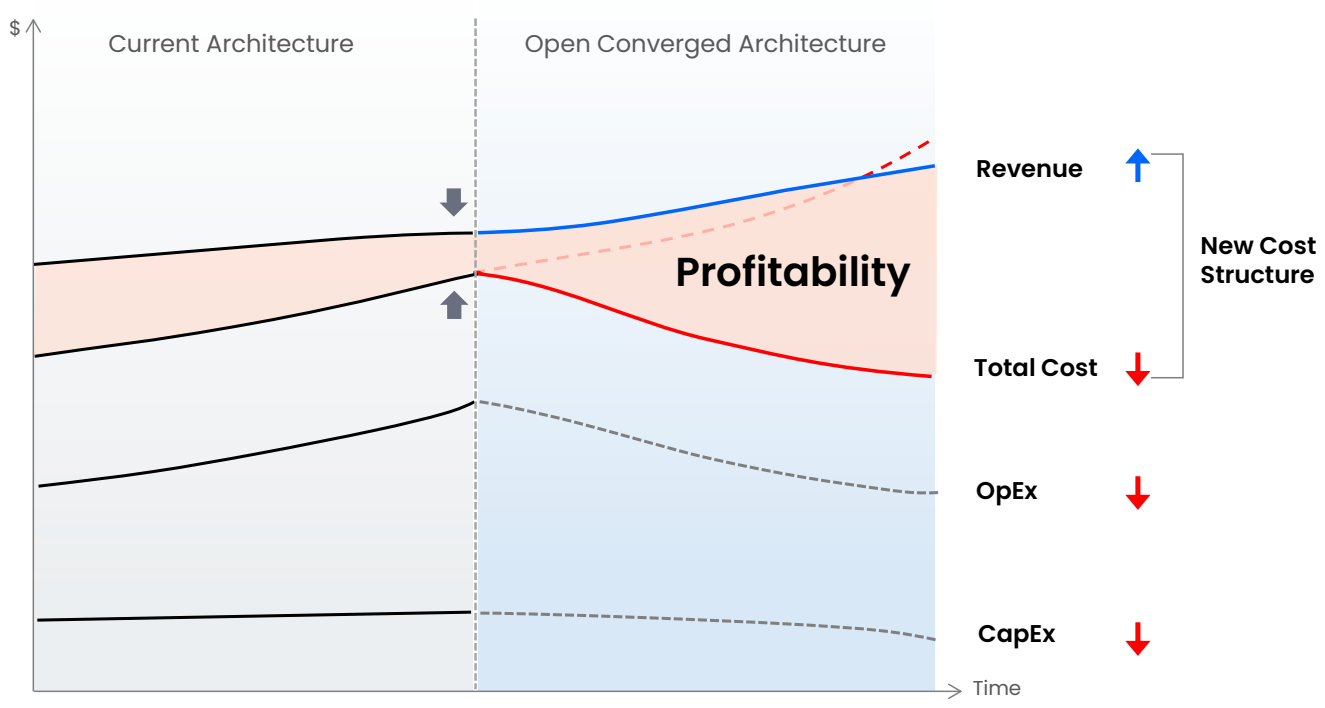

By separating software from hardware (disaggregation) and using proven cloud principles like software-defined virtualized infrastructure and elastic scalability, operators can achieve major efficiency gains across their networks. In fact, early adopters of DriveNets' approach – including top service providers in the US and Japan – have already saved hundreds of millions of dollars by transforming their networks this way.

Cloud-Native Architecture Attributes

DriveNets Network Cloud – DriveNet's cloud-native networking solution applies proven cloud principles to telecom networks. The result is a radical redesign of the router from a closed proprietary chassis to a scalable, open, cloud-native platform. In a DriveNets-based network, routing hardware is simplified into white-box building blocks, and all intelligence resides in DriveNets software.

Key attributes of this cloud-native approach include:

Hardware-software disaggregation

DriveNets runs on open, merchant-silicon white-box hardware, decoupling network software from underlying hardware. SPs can mix-and-match compliant white-box devices and optics, choosing best-of-breed components. This openness breaks vendor lock-in and frees operators to select hardware from multiple vendors, increasing bargaining power and supply-chain robustness.

Software-based architecture

The network operating system (NOS) is implemented in software microservices and containers rather than in fixed appliances. A single, unified NOS image runs across all platforms, which standardizes operations and enables high levels of automation and orchestration. This software-driven design means new features and updates can be rolled out like cloud software updates, accelerating innovation.

Elastic scale-out

The architecture scales elastically from a single 2.4Tbps router (one white-box device) up to a 921Tbps cluster composed of dozens of white boxes. Multiple units are interconnected in a Clos topology so that they function as a single, distributed routing entity. This horizontal scale-out approach (similar to cloud data centers) means capacity can be added incrementally, using the same software stack and hardware building blocks across the network from core to edge. Operators can start small and grow seamlessly, without forklift upgrades or architecture changes, simply by adding more white-box units to the cluster.

This cloud-native networking model transforms how networks are built and operated. By making the network flexible and modular, DriveNets Network Cloud supports a variety of deployment models and use cases under one architecture.

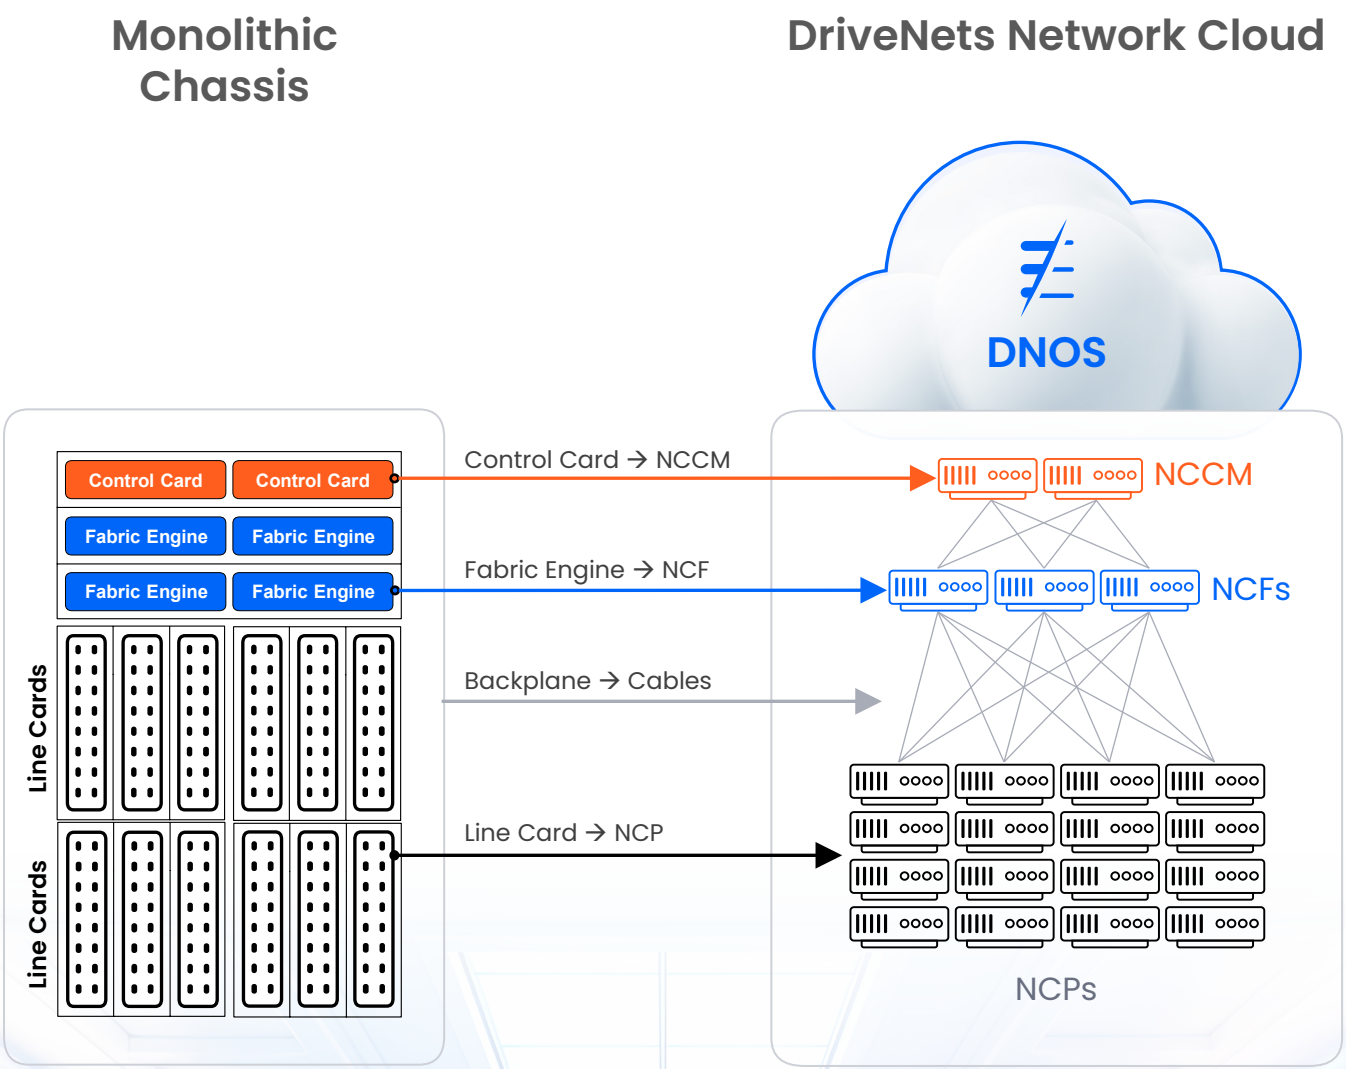

Solution Components

DriveNets Network Cloud is comprised of several modular software and hardware components, designed with a clear separation between the control plane (software intelligence) and the data plane (forwarding hardware). This logical separation allows each plane to scale and evolve independently, providing a high degree of flexibility and quick adaptation to new demands and technological innovations.

The main solution components include:

DriveNets Network Operating System (DNOS): DNOS is a distributed, cloud-native network OS that runs on the white-box cluster. Built on containerized microservices, DNOS enables rapid feature deployment and resilience. It provides all Layer 3 routing, management, and control-plane functions in software across the cluster.

DriveNets Network Orchestrator (DNOR): DNOR is the orchestration and management system for the DriveNets Network Cloud solution. DNOR treats a cluster of white boxes as a single entity, automating the network lifecycle – from initial provisioning to configuration, scaling, and upgrading. It coordinates the distributed components, handles software rollouts (with zero or minimal downtime), and integrates with external OSS/BSS for full automation.

Network Cloud Packet Forwarder (NCP): This is the data-plane workhorse, a high-performance white-box forwarding element built on merchant silicon (e.g., Broadcom ASICs). NCPs handle all packet forwarding and data traffic. In a cluster deployment, multiple NCP white boxes act like the line cards of a traditional chassis, providing the router's ports and throughput.

Network Cloud Fabric (NCF): NCF is the cluster interconnect fabric, essentially the switching fabric that links all NCP "line cards" in a distributed chassis arrangement. NCF white boxes create a non-blocking network among NCPs, ensuring that any port can reach any other port across the cluster with full bandwidth. The NCFs are similar to the backplane or fabric cards in a traditional chassis, but implemented with a cell-based technique.

Network Cloud Controller & Management (NCCM): The unified control-plane component, NCCM hosts the routing protocols, control logic, and synchronization services for the cluster. It ensures all NCPs operate as a single router, managing routing tables, distributing configurations, and maintaining state across the system.

DriveNets' solution aligns with the Open Compute Project's (OCP's) Distributed Disaggregated Chassis (DDC) specification and the Telecom Infra Project's (TIP's) Disaggregated Distributed Backbone Router (DDBR) specification. By building to open specs, DriveNets ensures interoperability and avoids proprietary traps. Conforming to frameworks defined by operator-led initiatives (OCP, TIP) reinforces the multi-vendor, open nature of the platform's hardware and software.

Deployment Models

Providing flexibility to address different network scales and scenarios, DriveNets Network Cloud supports two primary deployment models:

Standalone model:

In this simplest configuration, a single white-box NCP (plus its control-plane software) operates as a fully functional router. This is ideal for smaller-scale needs or remote locations; for example, a 2.4 Tbps standalone router can be used at a network metro edge or a small point-of-presence. The standalone model delivers all the benefits of DriveNets software on a compact footprint (one physical box). It uses merchant silicon hardware but is managed with the same DNOS and DNOR software as larger deployments. Standalone routers can support various interface speeds (from 1GbE up to 400GbE/800GbE ports, depending on the NCP hardware) and often are used for provider edge, aggregation, or peering use cases where a single-device solution suffices.

Cluster model:

For high-capacity needs (e.g., core, large aggregation sites), DriveNets employs a cluster of white boxes acting as one logical router. In cluster mode, multiple NCP units (forwarders) are interconnected via NCF fabric units in a Clos topology, effectively creating a distributed chassis router that can scale to extreme throughput. Cluster sizes can range from as small as 2-3 boxes up to configurations supporting hundreds of Tbps (e.g., a cluster up to 921 Tbps of capacity). Regardless of size, the cluster is managed as a single router by DNOS/DNOR. This model provides carrier-grade redundancy and scale: operators can start with a modest cluster and scale out incrementally by adding white boxes (linearly increasing capacity and port density). It offers "pay-as-you-grow" expansion, and new capacity can be added without replacing existing hardware.

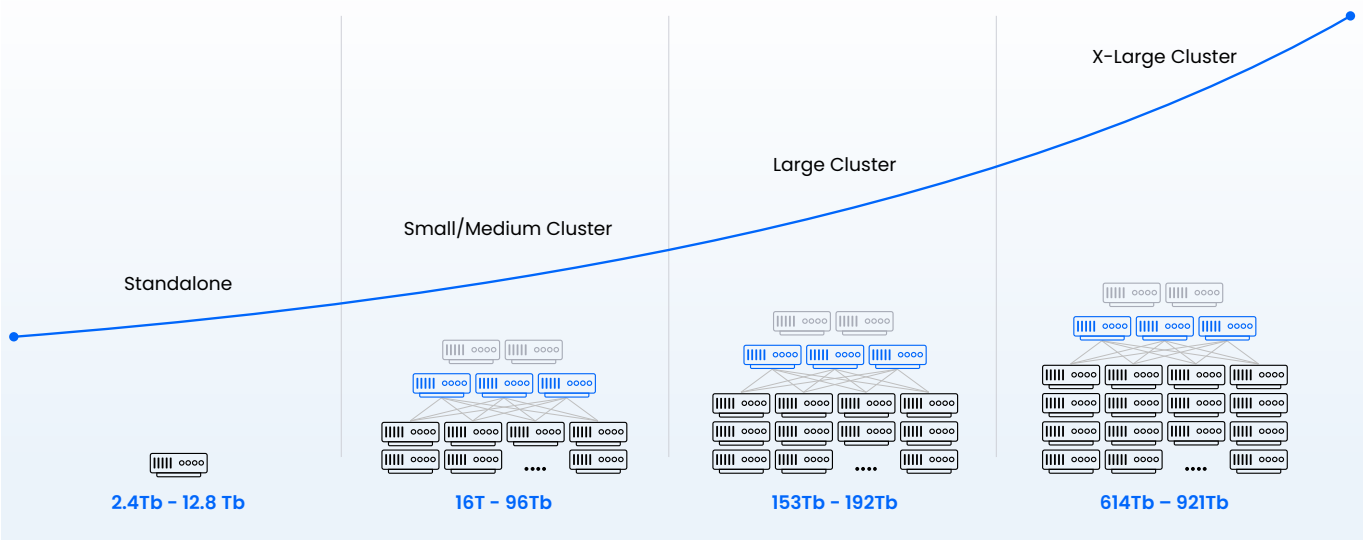

This cluster model enables DriveNets to serve even the largest core network roles, transiting [massive traffic volumes](#).

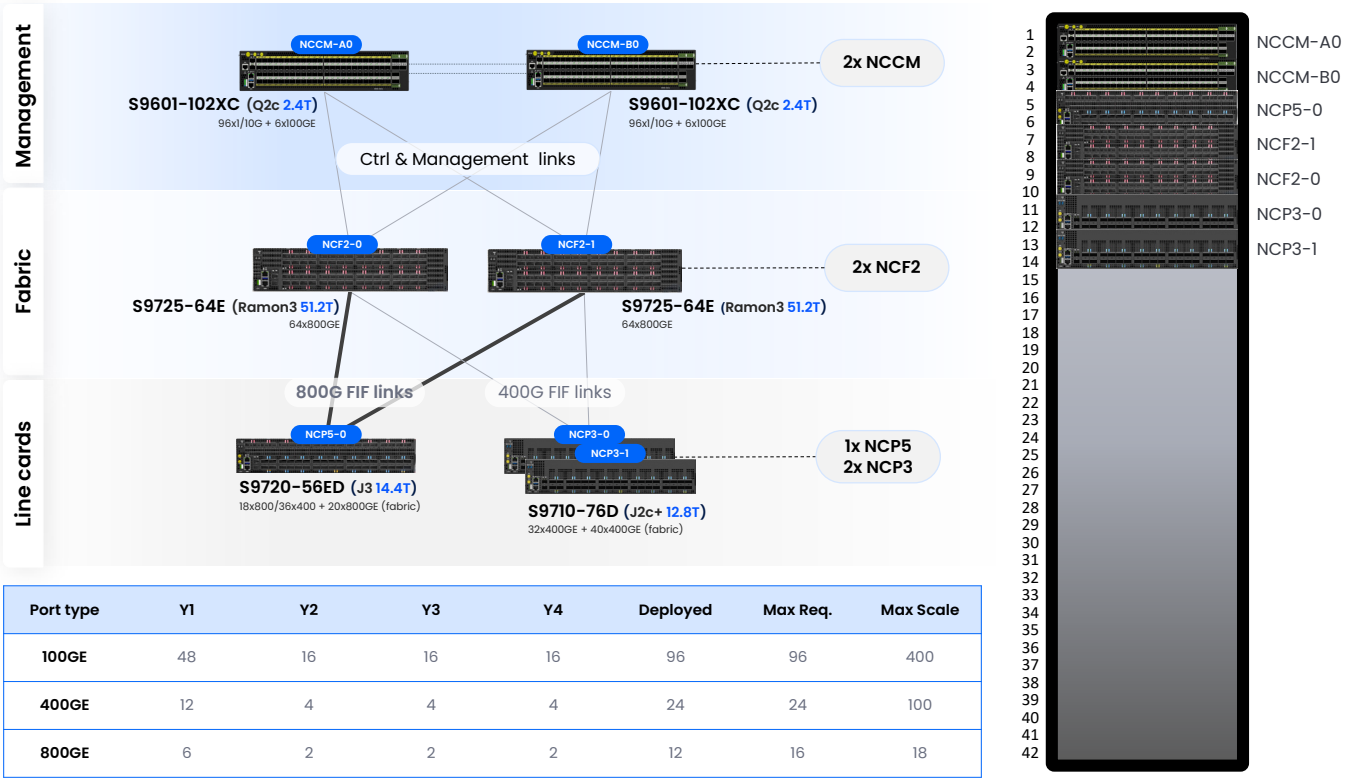

Port type	Y1	Y2	Y3	Y4	Deployed	Max Req.	Max Scale
100GE	48	16	16	16	96	96	400
400GE	12	4	4	4	24	24	100
800GE	6	2	2	2	12	16	18

Importantly, both models leverage the same software and hardware building blocks. An SP can deploy a uniform architecture network-wide, from core to aggregation to edge, all using the same two types of white-box devices (NCP and NCF) and the same NOS software.

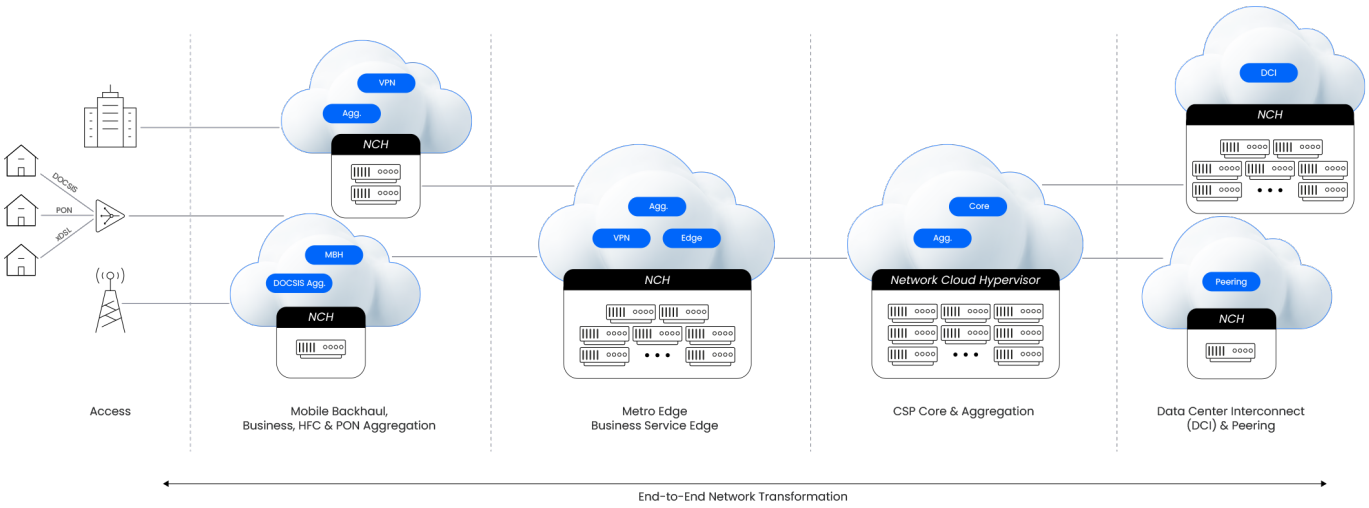

This consistency means use cases across the network (core routing, peering, data center interconnect, mobile backhaul, DOCSIS aggregation, etc.) can be served by the DriveNets Network Cloud platform using different sized deployments but identical technology.

Operational & Commercial Benefits

By adopting DriveNets' cloud-native model, operators gain a range of benefits including improved total cost of ownership (TCO), agility, and sustainability:

Lower cost structure (CapEx and OpEx)

Disaggregation fundamentally lowers capital expenditure by using commodity hardware. White-box routers based on merchant silicon can be an order of magnitude less expensive than traditional proprietary chassis systems; AT&T, for example, reported roughly 10× CapEx savings with white-box hardware compared to legacy routers. On the operational side, OpEx is reduced through network simplification and automation. With fewer hardware SKUs and one NOS across all sites, maintenance processes are streamlined and training overhead drops. The distributed architecture also improves availability as faults are isolated to individual units; as a result, a single failure has a minimal "blast radius" and does not take down an entire chassis. This increases overall uptime and lowers the cost of outages – a significant OpEx factor.

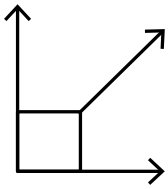

Elastic scalability and investment protection

The DriveNets model enables a "pay-as-you-grow" approach to network expansion. Instead of overspending on oversized chassis upfront or doing disruptive forklift upgrades, SPs can add capacity incrementally by inserting more white-box units. This not only defers capital spend until needed but also avoids service interruptions as new capacity comes online seamlessly. The architecture offers essentially "unlimited" horizontal scaling, so operators no longer are constrained by a vendor's fixed chassis size.

Upgradeability and sustainability are also greatly improved. When newer, higher-density ASICs or interfaces (e.g., 800GE, 1.6TE) become available, they can be introduced into the network cloud alongside existing hardware. There's no need to rip-and-replace entire routers; new technology can augment the cluster, protecting the initial investment. Existing white boxes can be repurposed to other roles or locations as needed, extending their life and avoiding stranded assets. This approach guards against technological obsolescence and drastically reduces the cycle of hardware churn, contributing to sustainability by cutting down on e-waste.

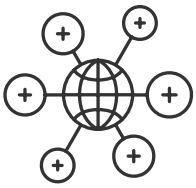

Vendor flexibility and open ecosystem

By separating hardware and software, DriveNets gives operators the freedom to choose best-of-breed components from multiple vendors, fostering a robust multi-vendor ecosystem. This breaks the traditional vendor lock-in where one supplier's proprietary chassis dictates the entire network. Instead, an SP could source white-box hardware from preferred ODM manufacturers and run DriveNets software (or even another compatible NOS) on top. This vendor diversity provides leverage in pricing and supply chain management: if one vendor faces delays or creates cost issues, alternatives can fill the gap. It also encourages innovation, as new hardware advancements can be integrated without waiting for a single vendor's roadmap. Standards compliance (with OCP DDC and TIP DDBR) further ensures that the platform adheres to industry-defined open interfaces, making it easier to collaborate across an ecosystem of hardware manufacturers, software providers, and integrators.

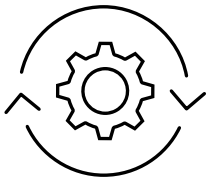

Operational simplicity and automation

All DriveNets Network Cloud deployments, large or small, run the same DNOS software and use the same two hardware types, which simplifies training, configuration, and troubleshooting. Inventory and spares management are greatly simplified – instead of stocking many different line cards and chassis variants for each network tier, an operator can stock just a few types of white-box units (which are interchangeable across roles).

Moreover, DriveNets embeds AI for operations (AIOps) capabilities for real-time analytics and root-cause analysis (RCA): the system can automatically collect and analyze logs/telemetry to detect anomalies and pinpoint root causes of incidents. Such automation speeds up troubleshooting and recovery, cutting down mean-time-to-repair (MTTR).

In-service software upgrades are another key advantage. The distributed router can be upgraded component by component, allowing for live updates with sub-second traffic impact. This means many maintenance tasks no longer require lengthy planned outages or maintenance windows, keeping services "always-on" and improving the customer experience.

These operational improvements have been validated by deployment experiences. KGPCo, a leading integrator working with DriveNets, noted up to a 50% reduction in network planning time, ~40% faster service turn-up, and ~30% faster installations and integrations when compared to traditional approaches.

In summary, DriveNets Network Cloud provides a dramatic TCO advantage (both CapEx and OpEx) over legacy routers. It empowers service providers to own their network economics and streamline operations, thereby lowering costs, avoiding vendor lock-in, automating processes, and scaling seamlessly like a cloud infrastructure.

TCO Parameter		Traditional Chassis-based Router	DriveNets Network Cloud Distributed Disaggregated Router
CapEx	Router cost	Chassis-based, monolithic architecture with a hefty premium in charges for branded network devices	Disaggregated low-cost standard networking white boxes from a choice of ODM vendors, with a cost-effective software licensing model
	Effective scale	Scale-up (high upfront investments), with limited chassis upgrade choices (+ OpEx cost of swapping out devices in "forklift" upgrades)	Pay-as-you-grow incremental growth model enabling scale-out (like cloud) and "unlimited" expansion flexibility with incremental horizontal scaling
OpEx	Operations and availability	Complex monolithic routers, with multiple hardware and software variations and diverse maintenance procedures, leading to compromised availability	Streamlined white-box devices that are less complex than traditional monolithic routers, resulting in lower failure rates, while their distributed architecture isolates faults (reducing blast radius)
	Ware house/ inventory	• Complex spare parts, with different chassis/card models and sizes used in different locations in the network • Operational complexity – due to inconsistency of physical infrastructure from high hardware variations, multiple software versions and varied maintenance procedures, which increases costs while slowing down network evolution and upgrades	• Simplified sparing – same white-box building blocks used for any sized router • Same DNOS software suite for all deployments – whether a single box or small, medium, or large clusters
	Vendor lock-in	Software and hardware from the same vendor, low probability of mix-and-match vendors	Mix and match vendors – hardware and software sold independently from different/ multiple vendors, resulting in supply chain robustness

Customer Use Cases

DriveNets Network Cloud is not just a theoretical solution – it has been proven in some of the world's most demanding networks. Tier-1 operators in the U.S. and Asia have embraced the solution, validating its scalability and benefits in live environments.

Notable tier-1 examples include:

→ Core backbone transformation:

AT&T, the world's largest telecom company, was an early adopter of DriveNets. Starting in 2020, AT&T began replacing its monolithic core routers with DriveNets Network Cloud as part of a strategic shift to open disaggregated architecture. Today, AT&T's next-gen IP core carries the majority of its backbone traffic on DriveNets-powered routers. In fact, as of 2025, over 80% of AT&T's core network traffic (around 840 petabytes per day) runs through DriveNets Network Cloud in a distributed disaggregated chassis (DDC) design. AT&T's success has provided a strong proof point that disaggregation can deliver carrier-grade results at massive scale, paving the way for other operators.

→ Multi-domain network upgrade:

Comcast is undertaking an end-to-end network transformation with DriveNets. In 2024 Comcast announced that it was implementing DriveNets Network Cloud as part of Comcast's "Janus" initiative to virtualize and embed AI innovation through the core of its network. The goal was to further enhance the flexibility, reliability, and resiliency of its connectivity offerings for both Xfinity and Comcast Business customers. In 2025 Comcast announced that it will use the DriveNets Network Cloud solution to expand the trial of Janus toward end-to-end network wide deployment.

→ Open core network deployment:

KDDI, one of Japan's top telecom operators, recently entered a strategic partnership with DriveNets to modernize its backbone network. Under this agreement, KDDI will deploy DriveNets Network Cloud software in its core routers across multiple major backbone sites, with commercial operation targeted by the end of 2025. This move is part of KDDI's broader effort to adopt DDBR designs.

These examples illustrate growing industry momentum toward cloud-native, disaggregated networks. DriveNets Network Cloud has proven itself in large-scale, mission-critical settings. After realizing strong operational and cost benefits in initial deployments, leading operators are choosing to redesign their entire networks end-to-end using DriveNets' disaggregated architecture. These savvy operators are building new, converged, and automated networks that are simplified, leaner, and fully softwarized.

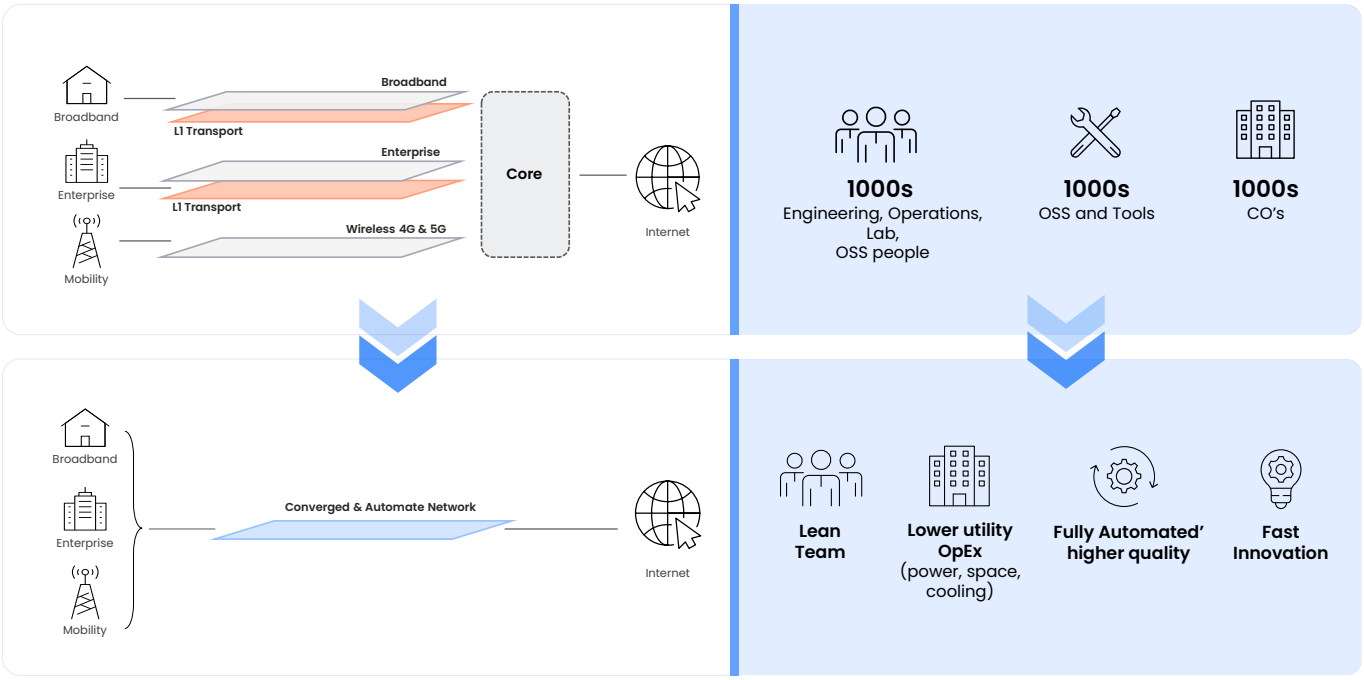

Conclusion

In conclusion, networks need to be built like clouds. Traditional, hardware-based systems are not built for the flexibility, scale, or cost-efficiency that modern services require. DriveNets Network Cloud takes a different approach: it decouples software from hardware and turns the router into a scalable and adaptable platform.

This shift toward software-based networks is already underway, with many service providers adopting disaggregated solutions. It's no longer a question of if this change will happen, but when. For many operators, that time is now.

DRIVENETS

DriveNets is a leader in high-scale networking solutions for service providers and AI infrastructures. The company created a radical new way to build networks that substantially improves the network's economic model and optimizes network utilization and efficiency.

DriveNets Network Cloud is a software-based solution that runs on standard white boxes and adapts the architectural model of cloud to high-scale networking, substantially increasing scale and lowering network costs. The solution is used by leading service providers globally.

DriveNets' Network Cloud-AI, which was introduced to the market in 2023, is based on the same technology, providing the highest-performance Ethernet-based AI networking solution and the best Ethernet alternative to InfiniBand in the market. The solution is used by Hyperscalers, NeoClouds, and Enterprises worldwide.

Learn more at **www.drivenets.com**