APPLICATION NOTE · AUGUST 2026

Internet Peering and Data Center Gateway with Exaware Open Networking

High-scale BGP, fast convergence and flexible policy control for service provider and IXP environments.

6M / 6M
IPv4 / IPv6 routes (FIB)

12M
BGP path scale

1-2RU
Converged router + DWDM

Up to 50%
TCO reduction

Executive Summary

Internet peering and data center gateway routers sit at the busiest, most operationally critical point in a service provider, cloud, or IXP network. Beyond routing, the same site also needs a wavelength-level handoff to the outside world — for data center interconnect, metro transport, or a link back to a carrier backbone. Meeting both requirements with separate routing and optical infrastructure is expensive and operationally complex.

Exaware addresses this with a single open networking architecture: ExaNOS routing software running on merchant-silicon platforms, with the option to converge IP routing and DWDM optical transport onto the same device.
The result is carrier-grade BGP peering, DC gateway, and routed interconnect capability, at up to 50% lower total cost of ownership than traditional, multi-vendor infrastructure.

This application note walks through that architecture: the peering, DC gateway, and routed interconnect use cases it supports; the technical capabilities that make it carrier-grade; its validation in real production networks; the economics behind the TCO reduction; and the hardware platforms available today.

6M / 6M	12M	1-2RU	Up to 50%
IPv4 / IPv6 routes (FIB)	BGP path scale	Converged router + DWDM	TCO reduction

THE CHALLENGE

The Peering and DC Gateway Challenge

Every service provider, cloud provider, and IXP depends on a small number of routers to connect its network to the rest of the internet. The demands on these routers have grown steadily as internet routing tables, peer counts, and traffic volumes have expanded — and the rise of AI-driven workloads and modern, high-density data centers is adding further pressure on gateway routers to move more traffic, at higher port speeds, without a proportional increase in cost or rack space.
Whether the primary driver is a growing peering footprint, a new IXP presence, or a data center fabric that needs a resilient path to the outside world, operators are evaluating peering and DC gateway infrastructure against the same core requirements:

- **Full internet-scale BGP RIB/FIB support, with headroom for continued table growth**
- **Fast, predictable convergence when peers, links, or paths fail**
- **Flexible, maintainable policy control for inbound and outbound traffic engineering**
- **Route origin validation (RPKI) to guard against hijacks and misconfiguration**
- **A converged path to DWDM optical transport where interconnect distances require it, without adding a second infrastructure layer**
- **A cost structure that scales with traffic growth rather than forcing large upfront chassis investment**

Open networking — separating the routing software from the underlying hardware — addresses these requirements directly. It lets operators run carrier-grade BGP, and optionally converged optical transport, on merchant-silicon platforms from established manufacturers, avoiding single-vendor lock-in while matching or exceeding the performance of traditional, multi-box infrastructure.

Exaware Open Networking Architecture

Exaware delivers ExaNOS V8.7, a carrier-grade routing operating system, on open networking platforms built around Broadcom DNX silicon (Qumran MX, Q2A, Q2C, J2C+, Q3D). The same software image runs across every certified platform, from 800G to 25.6T, giving operators one operational model regardless of which platform fits a given site — and, on 400G/800G-capable platforms, the same image also drives converged IP-over-DWDM transport (see Converged IP Routing and DWDM, page 6).

ExaNOS was not adapted for peering after the fact — peering was the first application built on the platform. The BGP control plane was designed from the ground up to handle internet-scale route churn efficiently, and it has been proven over more than a decade of carrier deployment, including three years of production operation in Comcast peering sites (see Proven in Production)

Built on a cloud-native architecture

ExaNOS uses a distributed, process-isolated software architecture: a fault in one routing process does not affect forwarding or other protocols. Non-stop forwarding and graceful restart are supported across all routing protocols, and NETCONF/YANG interfaces allow SDN controller integration and automation.

What Operators Get

- Best-in-class BGP convergence, independently validated in production peering environments
- Full internet routing table support with room to grow
- Open APIs and standard NETCONF/YANG for automation and third-party integration
- Pay-as-you-grow licensing instead of large upfront chassis investment
- The option to converge routing and DWDM optical transport on the same device
- A single software image across the full hardware portfolio — no per-platform training or tooling

Peering and DC Gateway Use Cases

Exaware open networking platforms support every standard peering and DC gateway deployment model on a single ExaNOS software image — from a single-homed transit connection to a fully meshed, multi-homed peering footprint:

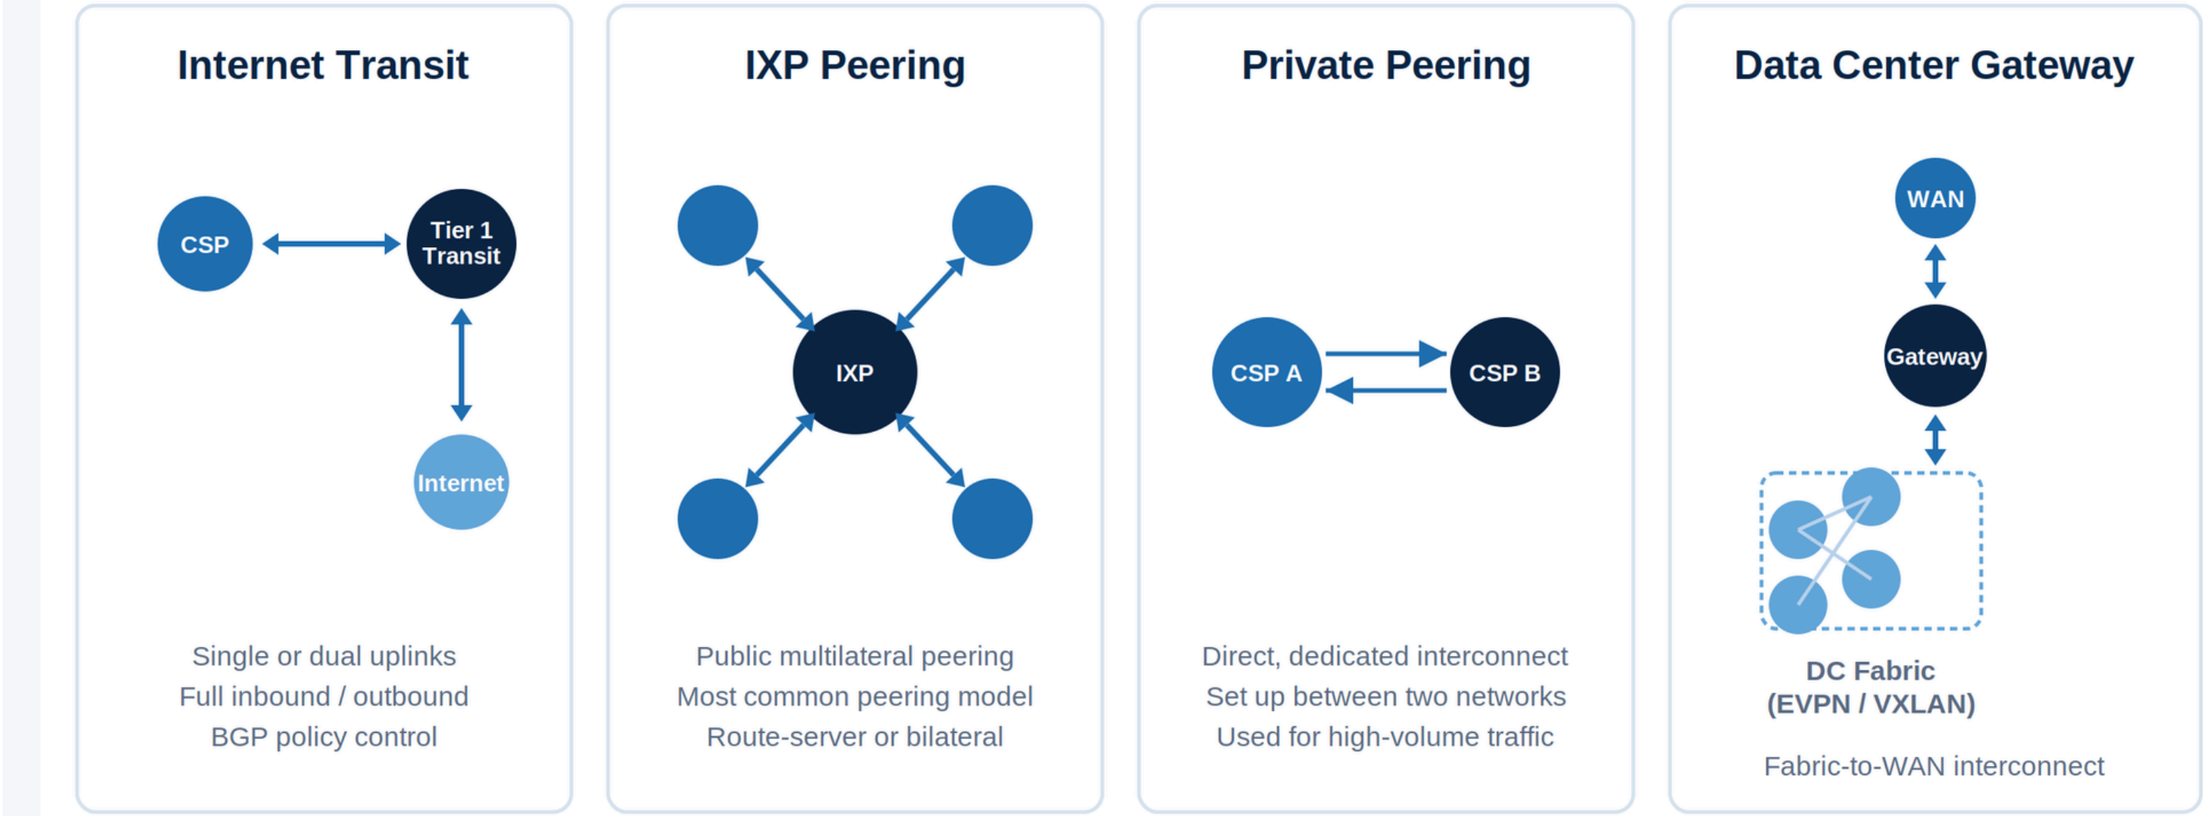

Figure 1. Exaware peering and data center gateway deployment scenarios.

Many of these deployments don't stop at routing. A data center gateway or peering site increasingly needs a wavelength-level handoff as well — a routed data center interconnect (DCI) between sites, or a metro/long-haul link back to a carrier hub.

Exaware's technical capabilities and converged IP-over-DWDM option, covered next, are what make a single platform able to serve both roles.

Key Technical Capabilities

The capabilities below are the ones that matter most in a production peering, DC gateway, or routed interconnect deployment.

SCALE & CONVERGENCE

BGP Scale and Convergence

ExaNOS redesigned the BGP best-path algorithm around per-prefix batching rather than per-peer, per-prefix processing, avoiding the exponential recalculation that slows legacy implementations at scale.

Insertion rate is hardware-limited rather than software-limited: 16,000 IPv4 / 8,000 IPv6 prefixes per second, with a typical 1M-route IPv4 table inserted in roughly 10 seconds.

BGP path scale reaches 12 million paths today, with ongoing development targeting 30 million.

BGP-PIC and BFD

When an IGP next hop changes, ExaNOS moves the entire dependent group of BGP prefixes to its backup path in a single hardware write, under 20 ms with a pre-computed Fast Reroute path, and ~300 ms without FRR, versus several seconds under legacy behavior.

BFD adds millisecond-level failure detection, improving PIC-protected convergence to ~50 ms. Both PIC-Core and PIC-Edge are supported, for eBGP and iBGP sessions alike.

Route Origin Validation (RPKI)

BGP is a trust-based protocol; a misconfigured or malicious peer can advertise prefixes it does not own. RPKI validates the binding between a prefix and its originating AS, and is increasingly a hard requirement at major IXPs.

ExaNOS supports RPKI Route Origin Validation with proven interoperability against common RPKI server implementations.

CUPL — Communication Unified Policy Language

CUPL is Exaware's structured policy framework for both control-plane (BGP, IGP) and data-plane (ACL, QoS) policy, in a single language.

Named prefix/community groups, action templates, and global or local parameters can be defined once and reused, with policy chaining and BGP multi-path load balancing across up to 300 iBGP or eBGP peers.

DC Gateway & Routed Interconnect

For sites acting as a data center gateway or routed interconnect, ExaNOS supports BGP Label Unicast (Inter-AS Option C), EVPN, and VPNv4/VPNv6 — bridging a BGP/EVPN data center fabric to an MPLS/IP carrier backbone, or extending routing across a routed DCI link.

Converged IP Routing and DWDM (ZR+)

Many peering, DC gateway, and routed interconnect sites also need a wavelength-level handoff — a routed DCI between sites, or a metro/long-haul DWDM link back to a carrier hub. Traditionally, meeting that requirement has meant deploying an entirely separate layer of infrastructure.

The Traditional Two-Layer Problem

A conventional interconnect stacks an IP routing system and a dedicated DWDM optical transport platform — typically two vendors, separate rack space, power, cooling, and element management, physically joined by short-reach grey optics.

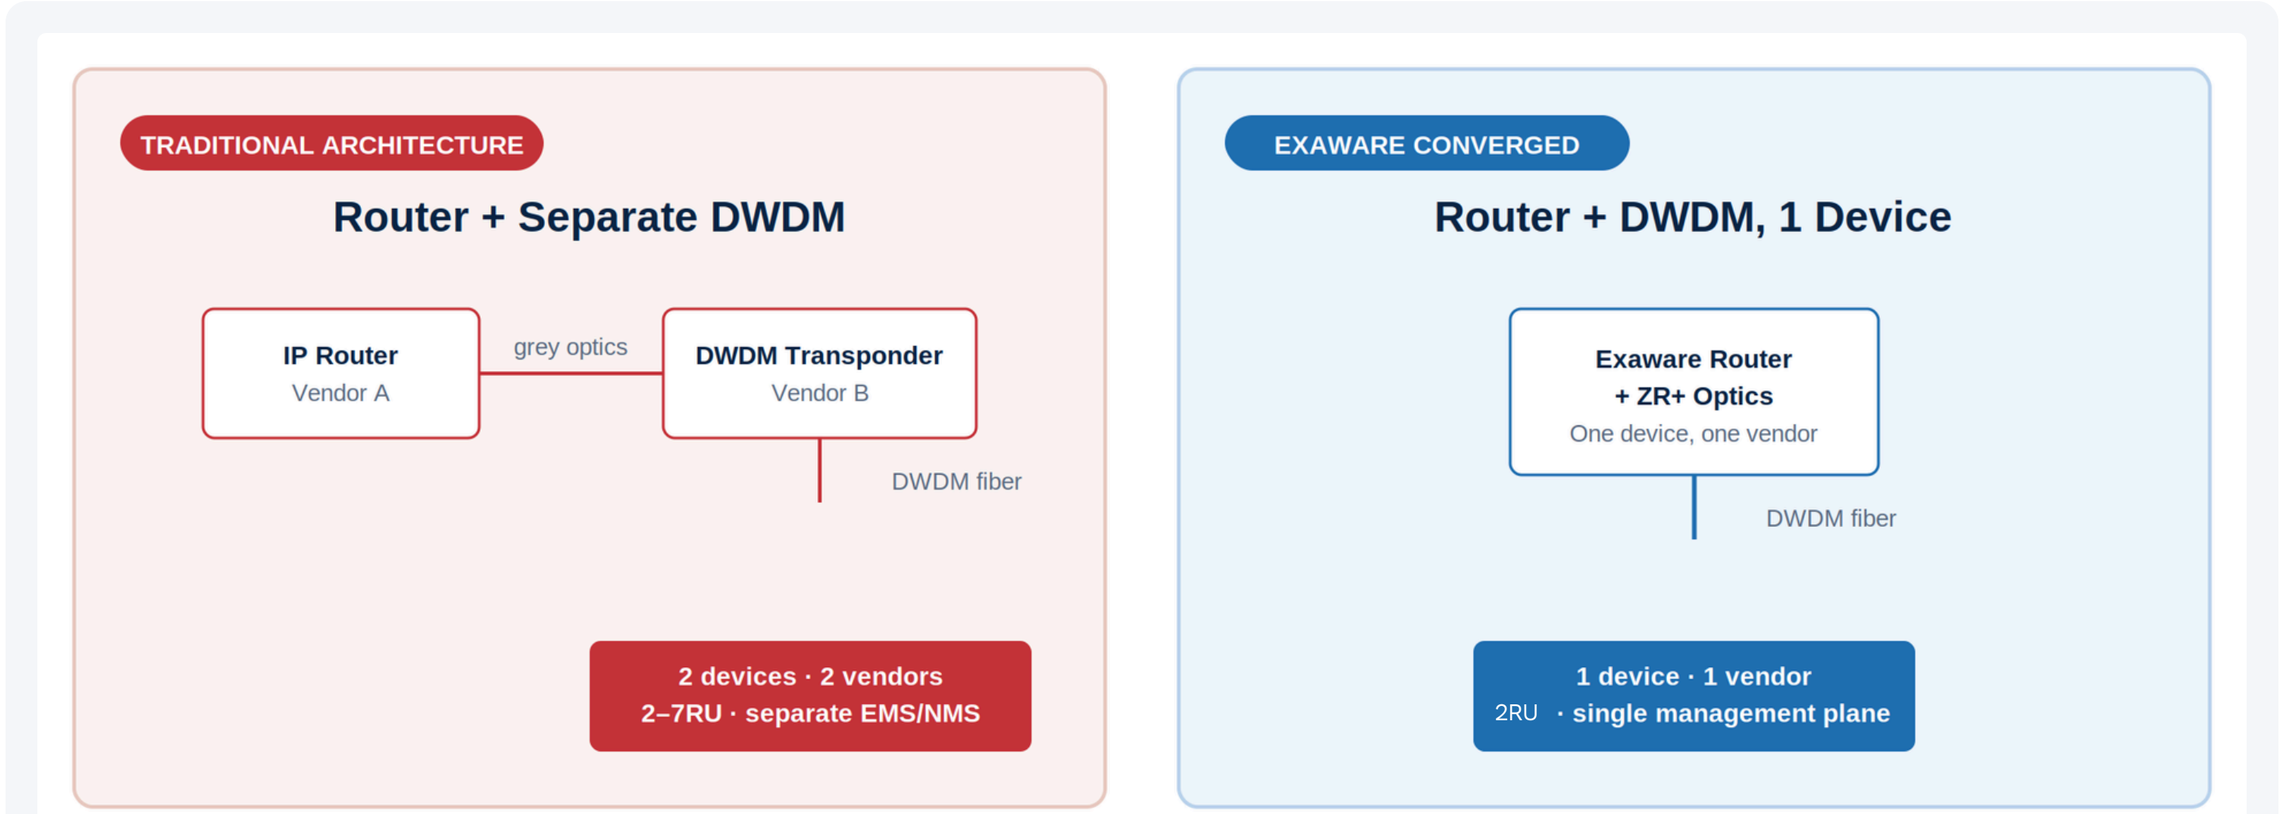

Figure 3. Traditional two-layer router-plus-DWDM architecture versus Exaware's converged single-device architecture.

Converged IP over DWDM

Exaware platforms with 400G/800G ports — including the EXA2400-30X — take ZR+ tunable coherent transceivers (OpenZR+ MSA) directly in the router's QSFP56-DD cage, C-band, 75 GHz spacing, no intermediate transponder, up to 130 km per wavelength.

Dimension	Traditional: Router + Separate DWDM	Exaware Converged with ZR+
Devices per site	2 (router + transponder/muxponder shelf)	**1 (router + DWDM in the same device)**
Rack space	2–7RU	**2RU**
Interconnect cabling	Grey optics patched per wavelength	**None — ZR+ drives directly onto the fiber**
Management	Separate EMS/NMS per layer	**Single ExaNOS management plane**
Vendor relationships	Two vendors, two contracts	**One vendor, one contract**
Reach	Set by transponder spec, EDFA managed separately	**Up to 130 km per wavelength**

Key Benefits: Operational simplification through a single management plane and faster fault isolation · up to 85% less rack space (2RU vs. 2–7RU) · one vendor relationship instead of two · no optics lock-in, with support for third-party ZR+ transceivers. Because converged IP over DWDM eliminates an entire device class rather than simply trimming its cost, it is one of the largest single contributors to Exaware's total cost of ownership advantage — removing transponder hardware, a second rack footprint, and a second vendor relationship in one step (see Economics, page 8).

PROVEN IN PRODUCTION

Proven in Production

Peering and DC gateway claims are only as credible as the environments they have been tested in. Exaware's BGP implementation has been validated in two of the most demanding categories of production network: a Tier-1 carrier peering footprint, and a top-tier global internet exchange.

TIER-1 CARRIER

Comcast — Multi-Year Production Deployment

Peering was the first application built on the technology that became ExaNOS, developed originally by Compass Networks. The resulting BGP implementation was installed and operated for three years in production Comcast peering sites — validating scale, convergence performance, and interoperability under real, sustained peering traffic rather than lab conditions alone.

GLOBAL IXP

DE-CIX — 2023 Proof of Concept, Frankfurt

Exaware completed a Proof of Concept with DE-CIX, one of the world's largest internet exchanges, running from March to July 2023 in DE-CIX's production peering environment in Frankfurt. The test used a full BGP routing table across both IPv4 and IPv6, with Layer 2 and Layer 3 switching provided at the productive peering LAN.

Over the course of the PoC, DE-CIX found that route-table insertion time was notably fast, and observed latency typically in the range of 1–2 milliseconds. DE-CIX also confirmed the platform was well suited to both dedicated peering tables and full internet routing tables.

The platform tested at DE-CIX was Exaware's EXA5300-800G, running ExaNOS on an Edgecore AGR130 (AS5916-54XKS) 800G aggregation router with large external TCAM. Edgecore's leadership noted at the time that the router's large-TCAM design, combined with the Exaware solution, supported roughly 3 million prefixes in the FIB. Exaware's CEO characterized the results as having “consistently demonstrated excellent performance,” reflecting the broader pattern seen across the company's peering deployments.

> ***"Every millisecond counts" for peering-grade latency; results that were "impressive in terms of performance."***
>
> **— Dr. Thomas King, CTO, DE-CIX · 2023 Proof of Concept announcement**

Platform naming note: the EXA5300-800G platform tested in the 2023 DE-CIX Proof of Concept is represented today under Exaware's current naming as the **EXA800-54XK**.

Economics and Hardware Portfolio

Chassis-based routers, paired with a separate DWDM layer where optical transport is needed, require large upfront capital investment sized for peak future scale, spread across two vendors. Exaware's open networking approach lets operators match hardware and licensing spend to actual traffic growth, on merchant-silicon platforms priced independently of the routing software running on them — with the option to remove the separate optical layer entirely.

Across hardware, licensing, and operations, Exaware's open networking approach delivers up to 50% TCO reduction versus a traditional chassis-based router at equivalent scale — with converged IP-over-DWDM adding further savings wherever a site also needs optical transport.

Up to 50%

reduction in Total Cost of Ownership vs. traditional chassis-based infrastructure

Where the Savings Come From

Lower initial CapEx — open networking hardware at commodity pricing, with carrier-grade software licensed separately

Pay-as-you-grow licensing — scale incrementally with traffic rather than over-provisioning upfront

One license includes all features, with no hidden costs or increases over time

No optics lock-in — standard SFP+, QSFP28, QSFP56-DD, and QSFP-DD support third-party optics

Reduced operational overhead — a single software image across the portfolio simplifies training and automation

Where needed, converged IP + ZR+ optical transport removes a second device, rack footprint, and vendor relationship entirely

Open Networking Hardware Portfolio

Exaware works with major established open networking hardware manufacturers using Broadcom DNX chipsets (Qumran MX, Q2A, Q2C, J2C+, Q3D). All certified peering and DC GW platforms include external TCAM for internet-scale route lookup, high-performance CPUs, and redundant power supplies and fans. A single ExaNOS software image runs across all platforms, simplifying operations and minimizing upgrade complexity.

The table and product images below summarize the current and near-term hardware options applicable to peering and DC gateway deployments:

EXA800-54XK

Access peering, regional ISPs, IXP edge

6x40/100G QSFP28 · 48x1/10G SFP+ · 800G

EXA2400-74XK

Full internet routing table peering

10x40/100G QSFP28 · 64x1/10/25G SFP28 · 2.4T

EXA2400-30X

Converged IP + DWDM, routed DCI

4x400G QSFP56-DD · 18x100G QSFP28 · 2.4T

EXA7200-36XK

High-scale peering and DC gateway

12x400G QSFP56-DD · 24x100G QSFP28 · 7.2T

EXA25600-32XK/OK

(4Q26) Next-gen, large-scale peering, AI-ready

32x800G QSFP-DD · 25.6T

ExaNOS BGP Feature Set (v8.6, March 2026)

Parameter	Value
Max IPv4 / IPv6 routes (FIB)	**6 million / 6 million**
BGP path scale	**12M paths (30M in progress)**
iBGP / eBGP peer scale	**300 iBGP / 100 eBGP**
IPv4 / IPv6 insertion rate	**16,000 / 8,000 prefixes per second**
Time to insert 1M IPv4 routes	**~10 seconds (typical deployment)**
BGP-PIC convergence	**<20 ms with FRR · ~300 ms without · ~50 ms with BFD**
ECMP	**Two-level (IGP + BGP)**

Built for What Peering Sites Need Today

Internet peering is a demanding proving ground for a routing platform — increasingly paired with a routed interconnect or optical transport requirement a routing-only platform can't address alone. Exaware's open networking architecture meets both: ExaNOS V8.6 delivers internet-scale BGP RIB/FIB capacity, best-in-class convergence, and flexible policy control, while the same hardware can converge IP routing with DWDM optical transport, on a single software image spanning 800G to 25.6T.

That architecture is proven where it matters: three years of production operation in Comcast peering sites, and a 2023 DE-CIX Proof of Concept validating full-table BGP performance and sub-2-millisecond latency at one of the world's largest internet exchanges. Combined with single-vendor simplicity and up to 50% lower TCO, Exaware gives operators an architecture built for what these sites need today — and ready for the converged interconnects they will need next.

Capability Area	Key Points
BGP Scale & Convergence	6M IPv4 + 6M IPv6 FIB; 12M path scale; hardware-limited insertion; BGP-PIC in hardware for <20 ms convergence
Peering, DCI & Converged Optical	Transit, IXP, private peering, DC gateway, routed DCI, with optional IP+DWDM (ZR+) convergence on one 1RU device
Policy Control & Security	CUPL unified policy language for BGP/QoS/ACL, plus RPKI ROV, prefix filtering, control-plane ACL, HW CPU policing
Hardware & TCO	800G to 25.6T platforms; external TCAM on key models; 10G / 25G / 100G / 400G / 800G port options to match deployment needs; up to 50% TCO reduction; pay-as-you-grow licensing
Proven Deployment	Three years in Comcast peering production; 2023 DE-CIX Proof of Concept at Frankfurt

About Exaware

Open networking for AI, data centers and carrier networks. Exaware delivers open building-block solutions for end-to-end AI and data center networking, built on SONiC, ExaNOS, and DCI technologies — providing a complete networking stack with up to 50% lower total cost of ownership. One perpetual license. All features included. No vendor lock-in. Founded in 2008. More than $100M invested. Field-proven across telecom, broadband, and data center networks.

Talk to Exaware

Contact Exaware for a personalized analysis of your peering, DC gateway, or routed interconnect environment, and to discuss validating performance in your own lab prior to deployment.
www.exaware.com/contact-us