Independent market research and competitive analysis of next-generation business and technology solutions for service providers and vendors

HEAVY READING

WHITE PAPER

Redefining Metro Networks for Scale, Simplicity, and Sustainability

A Heavy Reading white paper produced for Ciena

AUTHOR: STERLING PERRIN, SENIOR PRINCIPAL ANALYST, HEAVY READING

INTRODUCTION

Driven by cloud adoption, video, 5G, and the emerging edge, metro network traffic continues to increase. Metro network architectures must evolve to meet the challenges of scale and flexibility while also controlling costs. But communications service providers (CSPs) continue struggling to meet their conflicting network demands, and the status quo for metro networking no longer works for many.

This white paper identifies CSPs' metro networking challenges and describes how the challenges translate directly to their key network requirements for the future. The paper then discusses the enabling technologies and innovations operators need to address those requirements. Lastly, it describes three major hardware architectures for metro networking—traditional chassis, distributed disaggregated chassis (DDC), and hybrid chassis DDC—and details the benefits and drawbacks of each approach.

METRO NETWORKING CHALLENGES

Network traffic continues to increase in the range of 20–30% annually based on Tier 1 operator estimates and Heavy Reading's own survey data. Despite new revenue opportunities around public and private 5G, fiber broadband, and software-defined WAN/secure access service edge (SD-WAN/SASE), the overall revenue picture remains challenging. Omdia forecasts global CSP revenue will increase at a modest 2.8% CAGR from 2022 to 2027, while capex spend will increase at a 1.2% CAGR over the same period. This continuing mismatch between traffic growth and revenue growth puts tremendous pressure on available capex and opex spend, both of which are tied to revenue (see **Figure 1**).

Figure 1: CSP global revenue and capex, 2022–27

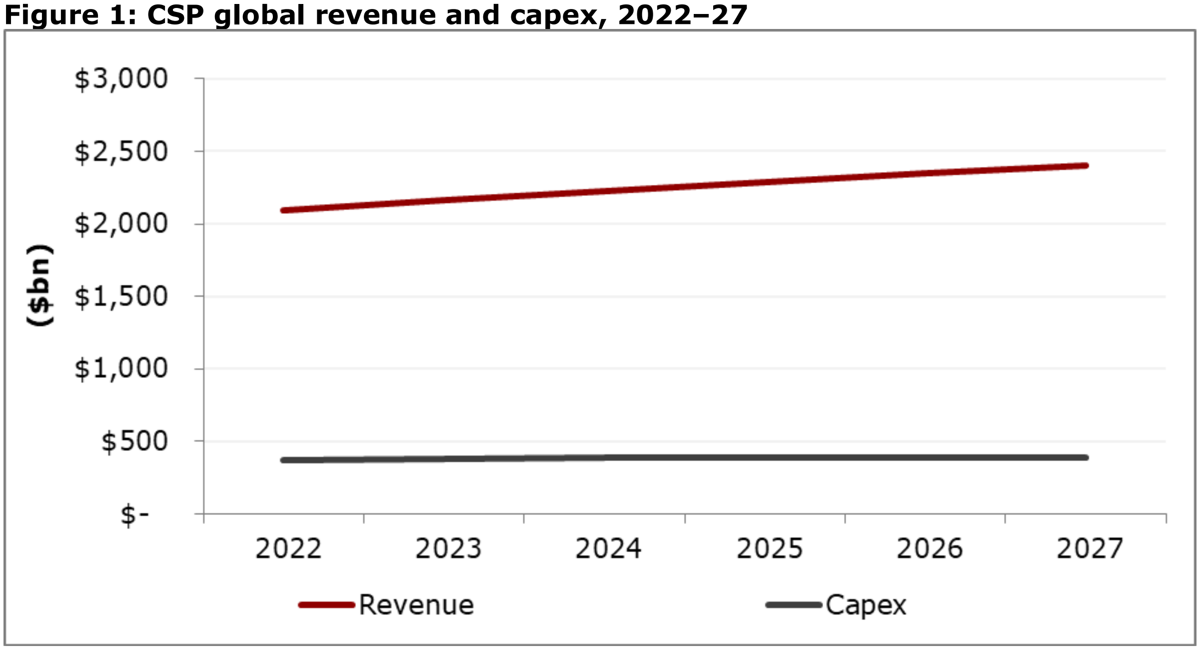

Source: Heavy Reading, Omdia (Communications Provider Revenue & Capex Forecast: 2022–27, September 2022)

While managing their capex, operators have also been on a mission for the last decade-plus to drive efficiency in their opex (whether by lowering opex or by reducing the ratio of opex to revenue). But on this front, recent Omdia research concludes that operators still have a lot of work to do. Omdia examined the financials for the last decade of 28 major operators globally. According to *The Future of Telcos and the Cloud: New Business Models and Paths to Growth for 2030*, only nine of them managed to reduce their opex over the time period. Of those nine, only three did so in an "efficient" manner by keeping their opex/revenue ratio below the 60% threshold, considered a sign of industry efficiency.

Enterprise adoption of the cloud model provides opportunities for network operators, including new network services such as SD-WAN/SASE, and new revenue-generating partnership opportunities with hyperscalers. The cloud model also promises opportunities for CSPs to reduce their opex by porting IT and network functions to the cloud.

But cloud services delivery also places new challenges on the network that go beyond capacity demand. Consistent with the moves to cloud and virtualization, server-to-server (or east–west) traffic will increase as a share of overall metro traffic. More network traffic will remain within the metro as data centers move closer to end customers. The migration to edge cloud will only increase metro traffic. CSPs will need to upgrade line rates as well as overall system capacities to accommodate the increasing metro-contained traffic volumes.

Sustainability requires producing goods and services with minimal environmental impact by lowering carbon footprints and moving to renewable energy sources wherever possible. In efforts to combat climate change, governments around the world are enacting sustainability regulations that directly affect production in major industries, including automotive, construction, oil & gas, mining, transportation, and many others. As large consumers of electricity in every country in which they operate, CSPs are not immune to government mandates. Already in 2023, the effects of government regulation are being seen, as CSPs are incorporating sustainability goals in their network requests for proposals (RFPs; particularly in Europe).

Beyond addressing government mandates, sustainability provides a crucial bottom-line benefit by reducing the power consumption contribution to opex. Opex costs have long been targeted by operators to reduce networking costs per bit (as noted above), but operational expenses have many contributors. Addressing the power contribution to opex has climbed the priority list as energy costs have skyrocketed over the past three years. It is not surprising that Europe—where energy costs increased 90% on average from October 2021 to October 2022 (according to the Household Energy Price Index by Energie-Control Austria, MEKH, and VaasaETT)—has been so aggressive about energy cost reductions.

METRO NETWORK REQUIREMENTS FOR THE FUTURE

Scale up and scale out with converged metro architecture

Fundamentally, next-generation metro networking equipment must scale to meet traffic demands, where scale has multiple dimensions. Operators need equipment that is future-proof with the ability to scale up compute and scale out the fabric via a converged metro architecture. These systems require 400Gbps interfaces today with the ability to add 800Gbps in the future. Multi-layer IP and optical coordination, as well as physical IP and optical integration, are key elements of convergence (both discussed in more detail below). Physical convergence means no capacity penalties for coherent optics and no power, footprint, or thermal limitations on the system.

Operators also require the ability to scale out with an extensible switch fabric that meets capacity needs for today and the future as well. New disaggregated routers called distributed disaggregated chassis (DDC) promise tremendous system capacity from low terabits per second up to 192Tbps per virtual system, with greater expansion built into the architecture.

Finally, the newest generation of routers disaggregates control and forwarding, giving operators the ability to scale up the compute function (x86 processing) completely separate from the route processor. Operators can independently scale fabric, compute, power, and operations, administration, and maintenance (OAM) protocols as needed.

Reduce operational complexity through coordinated multi-layer control

Software-defined networking (SDN) and automation can be used to converge control and management of the IP and optical layers, whether those layers are served by physically separate elements or are physically converged (i.e., integrated optics on routers). In a 2022 Heavy Reading global operator survey (*Open, Automated, & Programmable Transport Networks*), 52% of respondents selected multi-layer control and visibility across IP and optical layers as an end goal for managing their transport networks.

Operators seek several benefits in coordinating IP and optical layers and generally embark on a phased approach—moving from simplest to implement (e.g., least disruptive internally) to more sophisticated use cases—over time. As a starting point, operators use multi-layer visibility and analytics software to pull inventory data from the network itself and automatically correlate data between layers for a unified network view. Once a coordinated network view is built, it is used to feed existing management systems, significantly improving inventory accuracy.

Correlated data can also be used to analyze the network and identify optimizations and anomalies in its configuration. Anomalies are prevalent since the vision of the network as dictated by its planning team is often not fully implemented by the operations team or changes over time due to maintenance activities. This creates unplanned exposure to failures, inaccurate inventory data, and suboptimal service characteristics.

Emerging sophisticated use cases apply SDN control holistically across IP and optical layers. The application of SDN opens new use cases, including multi-layer network optimization to boost network utilization, coordinated multi-layer bandwidth-on-demand, and multi-layer network restoration.

Automate transport operations

A decade ago, the communications industry was focused on SDN. Now, the focus is on automation. The two trends are tightly interlinked, as SDN is the enabling technology for network automation. In a 2022 Heavy Reading survey, 39% of operators globally identified automation as a primary pillar of their next-generation transport network strategy (see **Figure 2**).

Figure 2: How important is automation for your next-generation transport network strategy?

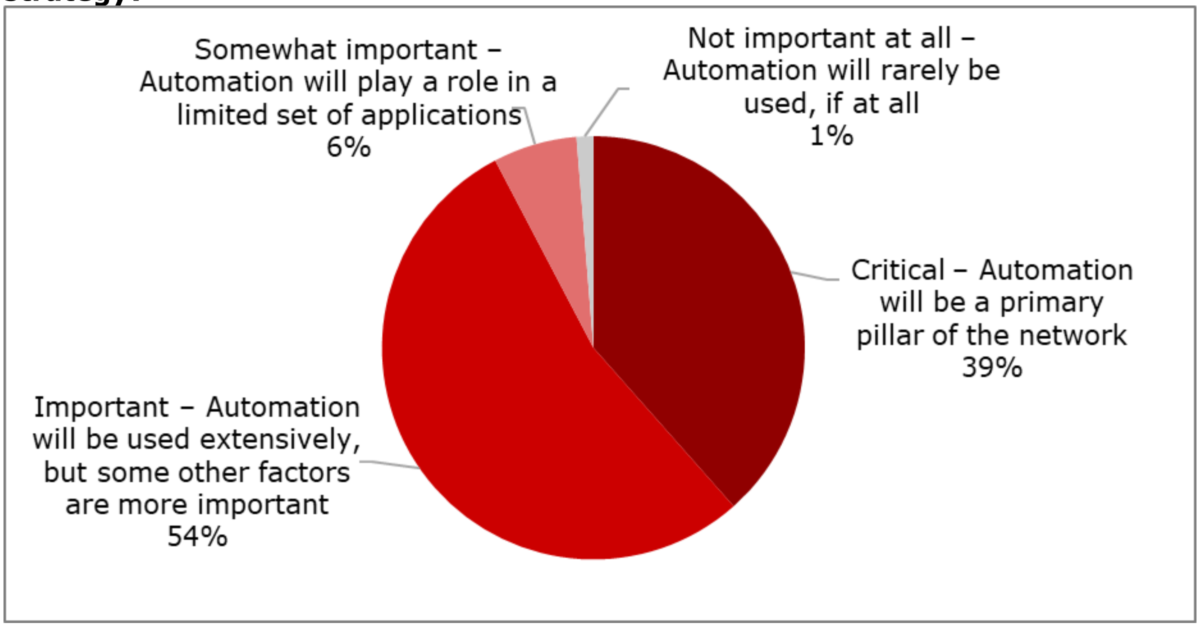

Source: Heavy Reading, Open, Automated, & Programmable Transport Networks: A 2022 Heavy Reading Survey, July 2022

Among the early transport automation use cases identified by operators are traffic engineering, service provisioning, activation and testing, network inventory and resource management, and network configuration compliance. Note that operators find value in transport automation in both single-layer and coordinated multi-layer implementations.

Reduce capex and opex with coherent routing

The concept of integrated DWDM optics on routers has been around for decades, though past generations never translated to large-scale deployments. The advent of 400Gbps coherent pluggable optics, however, is making this architecture a wide-scale commercial reality today. Integrated optics has moved from the IP over DWDM of the past to the coherent routed network of the future.

The primary appeal of the physical integration of optics on routers—and the reason it has been pursued for so long—is that inserting DWDM optics directly into the router eliminates the transponder shelf and optics between routers and DWDM systems. This includes the intermediate transponders as well as the gray optics and cabling that connect the shelves and routers. The elimination of transponder shelf hardware provides a capex benefit. It also saves space and power, which are two opex factors that have increased in importance in the terabit era.

Despite the promise, several challenges have stood in the way of wide adoption, the biggest of which has been the "faceplate tradeoff." At any given data rate, DWDM optics modules have always been significantly larger than client optics (typically by 2x). Thus, outfitting routers with DWDM optics used to mean sacrificing roughly half of the router's processing capacity and eliminating the capex and opex savings from integration.

With the commercialization of coherent pluggables in 2020 based on the OIF 400ZR Implementation Agreement, for the first time in history, the faceplate tradeoff was eliminated. OIF 400ZR DWDM optics fit the dominant 400GE client form factors using QSFP-DD or OSFP modules, including both space and power requirements. This means that there is no footprint or router capacity penalty imposed when integrating 400ZR optics on routers.

400ZR was developed for short reach data center interconnect (DCI) applications less than 60km, and the 150,000+ units shipped to date are almost exclusively for hyperscalers' deployments. CSPs require higher performance coherent pluggable optics that transmit over hundreds of kilometers, are capable of passing through multiple ROADM hops, and still fit in "client" form factors. Higher-performance coherent pluggables with higher transmit output power operating at 0 dBm and greater meet these CSP requirements and are coming to market from multiple suppliers.

To be clear, the benefits of multi-layer control (described in the previous section) can be had without physical integration, and not all operators are looking to migrate to coherent routed architectures. But the interest level is high. In a 2022 Heavy Reading survey (*Coherent Optics: 100G, 400G, & Beyond*), 39% of operators reported they intend to integrate 400G+ coherent pluggable optics into their switches and routers by 2024 (see **Figure 3**).

Figure 3: By the end of 2024, how do you see the distribution of 400G+ coherent optics in your network for various applications?

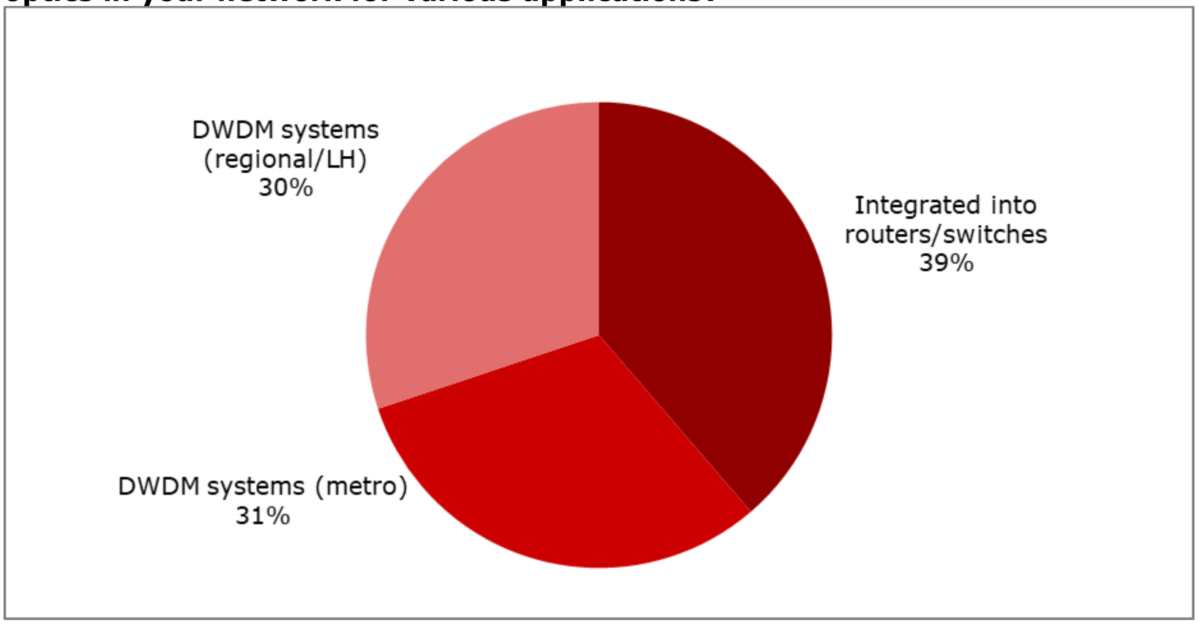

n=80
Source: Heavy Reading, Coherent Optics: 100G, 400G, & Beyond: A 2022 Heavy Reading Survey, October 2022

Heavy Reading believes the new generation of 0 dBm optics (which have marketing labels of ZR+ or sometimes ZR++) will unlock this large opportunity for coherent routing in CSPs.

Meet sustainability goals

Sustainability is likely not the primary factor for metro network modernization, but as operators move to their new next-gen architectures, sustainability will certainly be an important requirement. Multi-layer coordination, automation, and coherent routing all contribute to lower power consumption per bit and thus sustainability. In evaluating these technology introductions, operators will increasingly want to understand (and quantify) how the technologies directly contribute to overall power reduction. Additionally, operators will look for power-saving advances at the hardware level. These include pay-as-you-grow capacity expansion and power distribution and more energy efficient cooling mechanisms and cabling. Finally, CSPs want future-proof sustainability that yields multi-year and multi-generational benefits with fewer rip-and-replace upgrades.

Figure 4: Summary of metro architecture goals and benefits

Metro goal	Key benefits achieved	Requirements
Simplify and automate network operations	Opex	Multi-layer software control and automation of lifecycle operations
Fast-track Sustainability	Opex	Power and thermal flexibility
Scalability of IP and optical	Capex / opex	Scale up and scale out independently
Next-gen hardware – IP and optical	Capex	High-performance routing with high-power transport and integrated coherent optics without compromise

Source: Heavy Reading, 2023

METRO HARDWARE ARCHITECTURE OPTIONS AND TRADEOFFS

Traditional chassis architecture (status quo)

For decades, routers have been based on a monolithic chassis design with several key characteristics. Traditional systems are based upon a backplane or (more recently) a midplane or orthogonal direct connection, which provides connectivity to each slot in the system. These systems then use a set of switch fabric cards to move data across the backplane/midplane to various modules in the system.

Monolithic means that the router hardware, network operating system (NOS), and management software are supplied by one supplier that is responsible for the entire system build, integration, and support. The full capacity of the backplane/midplane is available on Day 1. This means operators also pay upfront for future capacity, though some vendors may introduce capacity licensing as a workaround. Operators fill the system with line cards until the backplane/midplane capacity is exhausted. For additional capacity, they purchase a new chassis and begin the filling process again.

Historically, the chassis architecture has been highly proprietary and closed, but over the past decade, it has adopted some open technologies. SDN introduced open APIs for a level of software programmability. On the hardware side, merchant silicon has diminished the reliance on proprietary ASICs for some applications. Additionally, pluggable optics bring greater supplier diversity for interfaces. Still, the chassis architecture is largely considered proprietary.

Advantages

The traditional chassis architecture is established and mature, and until recent times, it has met operator needs. Maturity means that the ecosystem is highly stable. Systems are fully integrated by the supplier and fully supported post-purchase. Operators understand this build and procurement model and know what to expect. There are not many systems suppliers to choose from, but the big suppliers are sophisticated companies that have decades of history.

Vendors continue to innovate on scalability within the chassis model. The evolution from backplane- to midplane-based is one example. Most recently, some vendors have adopted orthogonal direct designs that remove the physical midplanes for further scalability, but also with some challenges and restrictions that are tied to the chassis model.

Drawbacks

The primary drawback of the traditional chassis architecture is its lack of efficient scale (due in part to heat dissipation needs at multi-terabit capacities), high capex and opex costs, limited open programmability, and limited multi-chassis extensibility. The inability to scale to efficiently meet traffic demand drove AT&T to a new routing architecture in 2019. AT&T has been very vocal about its IP network transformation, but its requirements are not unique. In a Heavy Reading operator survey published in 2020 (*The Future of IP Networking: Time for Radical Change*), 46% of survey respondents reported that they expected to make radical changes in their IP network architectures, an astounding finding at that time.

Traditional chassis are being built with greater depth that fits data centers but may limit their application in telco environments (ETSI) that require 600mm in depth. The greater depths are a direct result of midplane (as well as orthogonal direct) architectures that boost density but add depth.

DDC architecture (white box)

An important alternative to the traditional monolithic chassis-based architecture is the DDC architecture based on a white box model. DDC is an approach to building routers where the different components of the system are separated into discrete modules that can be placed in different physical locations. Additionally, white box disaggregation decouples the NOS from the hardware, with the hardware based on merchant silicon and built to common specifications by contract manufacturers such as Edgecore Networks and UfiSpace.

The DDC was initially championed by AT&T for its own network architecture, but in 2019, AT&T submitted its specifications for DDC based on Broadcom Jericho2 processors to the open-source Open Compute Project (OCP). The move made DDC based on white box disaggregation of hardware and software a viable architecture for the global CSP community. Another prominent industry group, Telecom Infra Project (TIP), is working with OCP and incorporating DDC as part of its own work in disaggregated open routers (DORs).

In January 2023, AT&T announced that it had migrated more than 52% of all its production traffic to next-gen core routers using the DDC approach.

Advantages

DDC offers multiple benefits compared to the traditional chassis design. Chief among them, and the reason AT&T initially sought a radical new approach to routing, is scalability. DDC allows for flexible scaling of networking resources, as new line cards and switching modules can be added or removed as capacity requirements dictate. DDC scale has three dimensions:

- Capacity is not limited by the chassis backplane because there is no backplane in DDC. For example, the OCP's DDC specification states capacities up to 192Tbps (480 x 400Gbps ports).
- The modularity of the line cards and the switching modules allows for capacity to be increased incrementally (at the line card level), as opposed to the large stepwise increases of chassis-by-chassis additions.
- The architecture modularity allows scaling both up and down. This last point is important in the metro segment, where 400Gbps interfaces are needed, but aggregate capacities at the node level will be less than at the core.

Beyond scale, DDC offers other significant benefits:

- **Efficient resource utilization:** DDC architecture enables better resource utilization by separating the control, forwarding, and switching functions into different modules. This allows for more efficient use of resources and reduces waste.
- **Reduced downtime:** Because DDC architecture distributes the functions of the networking equipment across multiple modules, a failure in one module does not necessarily bring down the entire system.
- **Cost savings:** DDC architecture can cost less than traditional chassis-based networking equipment because it is based on commodity hardware with the ability to mix and match different components based on specific requirements.
- **Reduced vendor lock-in:** DDC based on white box can reduce vendor lock-in because software and hardware elements are sourced separately as opposed to by a single supplier. However, this benefit is predicated on a mature ecosystem in both hardware and software so that operators have multiple viable choices.

Drawbacks

The biggest challenges to adopting the DDC architecture come from its white box procurement model. Heavy Reading research shows that systems integration continues to be the biggest challenge in white box adoption. Many operators are not willing or able to take on this role or outsource it to a third party, which adds complexity to the purchase and maintenance. White box solutions are not optimized for managing cabling or fabric connectivity, thus driving the service provider to buy specialized cables and preplan for capacity growth and layout to fully utilize a DDC.

Beyond the integration issue, another challenge is that the ecosystem for white box DDC is immature for both the hardware and the NOS. As a result, there are currently few supplier choices, diminishing the white box value proposition of reduced vendor lock-in. Operators pursuing white box DDC hope that incumbent router vendors will ultimately be driven to adopt the model, but so far, this has not happened.

Hybrid DDC/traditional chassis

Today, the newest routing architecture to emerge is a hybrid option that combines elements of the DDC with elements of the traditional chassis. The aim is to take the best of each while avoiding the principal drawbacks.

In common with DDC, the hybrid router is based on disaggregated and modular hardware building blocks, including power, control, interface, and switch fabric modules with no backplane/midplane. Like DDC, hybrid disaggregates the control plane and forwarding plane functions. The hybrid architecture is also highly open in terms of software and hardware. For the NOS, this includes container-based applications, open APIs, and merchant silicon software development kit (SDK). On the hardware side, this includes the use of merchant switch fabrics (versus ASICs) and support for third-party coherent pluggable optics.

The most significant point in common with the traditional chassis is that the operating system and hardware are supplied by the same vendor. In other words, hybrid DDC is not white box. Thus, it can also support proprietary modes when needed. For example, it can support embedded coherent optics for applications where pluggable optics cannot meet performance specifications, like 800Gbps long distance links.

Advantages

Based on the DDC concept, the hybrid DDC/traditional chassis architecture yields many of the key benefits of DDC, including the following:

- **Flexible scale up and down:** Ciena's WaveRouter, the first product in this hybrid category, is based on 6Tbps interface modules that are stackable up to 8 per half rack unit for 48Tbps. With in-service extensible switch fabric, they can be deployed in non-adjacent rows for 192Tbps of virtual system capacity (matching the capacity specs of DDC in OCP).
- **Efficient resource utilization:** Control, switching, and forwarding are disaggregated. Disaggregation of control and forwarding planes provide the ability to scale up the compute function (x86 processing) outside the route processor in the product independently of any other function and can form a compute cluster.
- **Reduced downtime/higher availability:** Functions are distributed across multiple modules and can be placed in different cabinets, thus reducing the blast radius of a failure. Hardware efficiency is achieved by being able to buy only what is needed, as compared to a chassis, while providing high availability.

- **Cost savings:** The use of commodity hardware and merchant silicon—much of which is also used in white box products—is key to this value proposition. Coherent pluggable optics also contribute to lower costs (though Heavy Reading notes that traditional routers are also beginning to support coherent pluggables).
- **Faceplate flexibility:** No fixed backplane means no capacity tradeoffs if using both larger form factor interfaces (e.g., for high-performance coherent optics) and pluggable coherent optics.

Based on the traditional chassis model, the hybrid DDC/traditional chassis offers a fully integrated system of hardware and software with full support from the single supplier. As noted, systems integration is the biggest challenge in white box adoption. The traditional model also allows for vendor-proprietary features that may not be available through white box-specified hardware, such as embedded coherent optics.

Drawbacks

Single vendor NOS and hardware represent a double-edged sword. But the integrated approach will also be the primary drawback to the hybrid architecture for operators that are planning to migrate to a white box architecture in the near or medium-term. The hybrid model ties the operator to a single supplier for both hardware and NOS and will make it reliant on that vendor's roadmap for the lifecycle of the router. Open APIs and merchant silicon reliance can mitigate some limitations of proprietary systems, but the hybrid model will always be less open than the white box model, by definition.

Figure 5: Hardware architectures and tradeoffs

Architecture	Advantages	Drawbacks
Chassis	Well-established with a mature ecosystem and known procurement models	Lack of scale, inefficient use of capacity, highly proprietary, vendor lock-in
DDC	Scale, efficient resource utilization, cost savings, reduced vendor lock-in (potential)	White box systems integration complexity, lack of mature ecosystem
Hybrid DDC	Scale, efficient resource utilization, cost savings, known procurement model	Not a migration path to white box, vendor lock-in for hardware/NOS

Source: Heavy Reading, 2023

CONCLUSION

CSPs understand the network challenges they face, including handling the continuing mismatch between traffic growth and revenue growth, addressing massive traffic growth against capex/opex constraints, and adapting to cloud consumption models. But finding a solution has not been easy. For metro networks, operators have established a clear set of network requirements for the next decade, including the following:

- Scale up and scale out
- Reduce operational complexity (simplify)
- Reduce capex and opex
- Automate network functions wherever possible
- Address corporate sustainability goals

To meet these needs, many operators believe radical change is needed in their metro transport networks and architectures. Converged IP and optical transport networks (now known as coherent routing) represent one example of disruption that, when properly implemented, will touch every one of these requirements.

As they look to converged metros, operators are also investigating innovative, scalable hardware designs to supplant the traditional chassis. DDC based on white box is one alternative gaining traction. Most recently, a hybrid option has emerged that combines the benefits of chassis and DDC, while seeking to avoid the drawbacks of each. As they evaluate new metro technologies and architecture options, CSPs will keep an open mind and a strong focus on addressing their network requirements.

APPENDIX

This report draws upon data from the following research:

- Evan Kirchheimer, Richard Mahony, *The Future of Telcos and the Cloud: New Business Models and Paths to Growth for 2030*, Omdia, February 2023.
- Household Energy Price Index (HEPI) by Energie-Control Austria, MEKH, and VaasaETT, 2022 VaasaETT Ltd.
- Sterling Perrin, *Open, Automated, & Programmable Transport Networks: A 2022 Heavy Reading Survey*, Heavy Reading, July 2022.
- Sterling Perrin, *The Future of IP Networking: Time for Radical Change*, a Heavy Reading white paper produced for DriveNets, Heavy Reading, February 2020.
- Sterling Perrin, *Coherent Optics: 100G, 400G, & Beyond: A 2022 Heavy Reading Survey*, Heavy Reading, October 2022.