Infinera

WHITE PAPER

THE CASE FOR DISAGGREGATED ROUTING IN 5G AND DAA TRANSPORT NETWORKS

The Case for Disaggregated Routing in 5G and DAA Transport Networks

Two key applications are driving major transport network upgrades. For mobile networks, 5G promises a major increase in bandwidth to users and in capacity per unit area, while also enabling new applications based on ultra-low latency and high availability as well as offering massive Internet of Things (IoT) scale. Meanwhile, Distributed Access Architectures (DAA), along with fiber deep, will enable cable multiple-systems operators (MSOs) to deliver broadband internet speeds of up to 10 gigabits per second (Gb/s) and 4K/8K video, with converged interconnect networks (CIN) also providing a platform for new residential, enterprise, and wholesale services. IP routing has an important role to play in both these key network evolutions. However, traditional IP routers based on closed, proprietary hardware, with the chassis-based architectures required for scaling, result in significant vendor lock-in, leading to slowed innovation and high costs.

Traditional Router Limitations

Traditional routers provide a monolithic package of hardware and software from a single vendor. While in the past this model has provided benefits like guaranteed hardware/software interoperability and a single point of contact for purchasing, service, and support, today it creates a number of challenges for network operators. The significant vendor lock-in of proprietary, closed routers and the small number of router vendors result in limited pricing pressure and high costs. Innovation is dependent on the capabilities of the single selected vendor rather than an ecosystem, and network operators are prevented from leveraging different rates of hardware and software innovation independently, or selecting best-of-breed hardware and software independently.

Network planners also face an unenviable choice when selecting a traditional router for a specific location in the network, as shown in Table 1. Do they take the short-term approach of selecting a device optimized for the more immediate capacity requirements, which may need to be replaced as traffic and services evolve? Or do they select a larger device that will have a longer life but suffer the short-term pain of high up-front capital expenditure (CapEx) and a large footprint for a partially filled chassis with initially underutilized fabrics and backplane?

Fully Utilized Device	**Partially Filled Chassis**
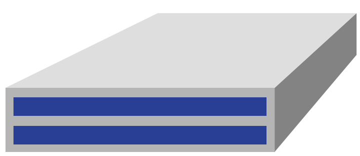	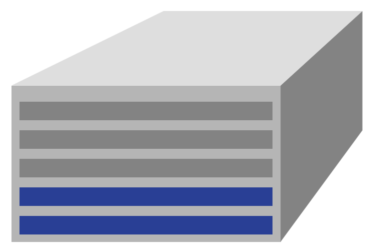
✓ Minimized up-front CapEx ✓ Minimized footprint and power	✓ More future-proof (upgradeable/scalable)
✕ Limited upgradeability/scalability ✕ Higher long-term CapEx (replacing device) ✕ OpEx cost of swapping out device	✕ High up-front CapEx ✕ Large footprint ✕ Higher power consumption

Table 1: The Traditional Router Selection Dilemma

The Carrier Network Disaggregation Trend

Disaggregation is not a new concept. In fact, in economic literature it is a natural consequence of industry maturity and has been successfully adopted in a wide range of industries in addition to IT and telecommunications. One recent carrier network example is the disaggregation of optical transport systems into open line systems (OLSs) and open transponders/muxponders. Another

Scope	Disaggregation	Example Industry Initiatives
Optical Transport	Disaggregation of WDM systems into OLSs and open transponders/muxponders	Open ROADM Multi-Source Agreement (MSA) ONF's Open Disaggregated Transport Network project Telecom Infra Project's Open Optical Packet Transport (OOPT) project
Network Functions	Network functions virtualization (NFV) with VNFs running on servers powered by commodity general-purpose CPUs	ETSI Open Source NFV Management and Orchestration (MANO) The Linux Foundation's Open Network Automation Platform (ONAP) and Open Platform for NFV (OPNFV)
Central Office	Central office/cable headend/mobile edge disaggregation based on data center principles	ONF's CORD project ETSI's MEC initiative

Table 2: Carrier Network Disaggregation Examples

is the virtualization of network functions such as firewalls, broadband network gateways, and the mobile core, disaggregating proprietary integrated systems into software-based virtual network functions (VNFs) running on server hardware powered by commodity general-purpose CPUs. The network edge is also being disaggregated based on data center principles, with initiatives such as the Open Networking Foundation's (ONF) Central Office Re-Architected as a Datacenter (CORD) and European Telecommunications Standards Institute's (ETSI) Multi-Access Edge Computing (MEC).

Disaggregated Routing Overview

This disaggregation trend is now coming to routers, but what is router disaggregation? At its most basic level, disaggregated routing separates router software and hardware, as shown in Figure 1. The hardware becomes a white box based on merchant silicon, also known as a bare metal switch. The software becomes a network operating system (NOS). While the hardware and NOS can come from different vendors, it is not possible to just take any white box and load any NOS.

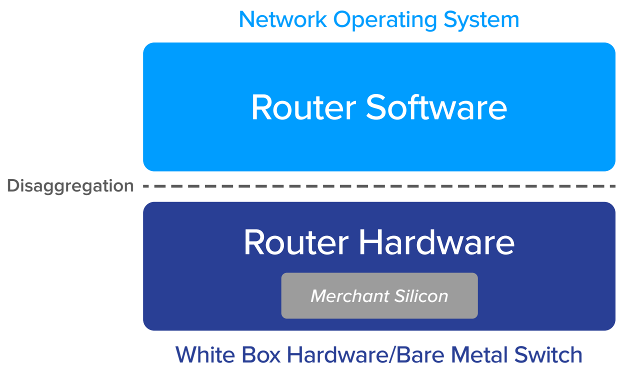

Figure 1: Hardware/Software Separation

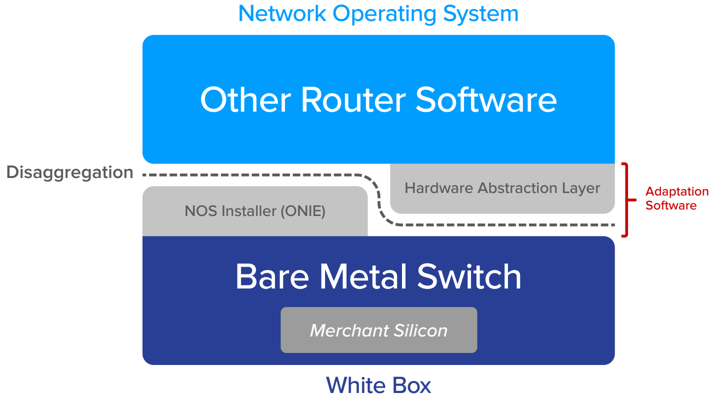

Figure 2: White Box/NOS Adaptation

Software is needed to adapt the NOS to the specific white box and its merchant silicon, as shown in Figure 2. This software comprises two elements, the first being the software that comes on the white box to load the NOS. Here, Open Network Install Environment (ONIE) provides an open environment for installing a compliant NOS onto third-party white box hardware and is fast becoming a de facto standard. A hardware abstraction layer in the NOS that adapts it to run on the specific merchant silicon in the white box is the second element.

However, disaggregation is not binary but rather exists on a spectrum, as shown in Figure 3. A first step from traditional routers with command line interface (CLI)-based configuration and/or network management system (NMS)-based management is the transition to software-defined networking (SDN) and open application programming interfaces (APIs) based on NETCONF, which lowers operational expenditure (OpEx), increases programmability, and provides the option to disaggregate some control plane functionality from the network element to the SDN controller/orchestrator. Beyond the basic separation of software and hardware, the different hardware elements of a chassis-based router can also be disaggregated to enable horizontal scaling. Coming from the data center world, horizontal scaling means scaling by adding more "machines," or in this case white boxes, as opposed to vertical scaling, upgrading the power of existing machines.

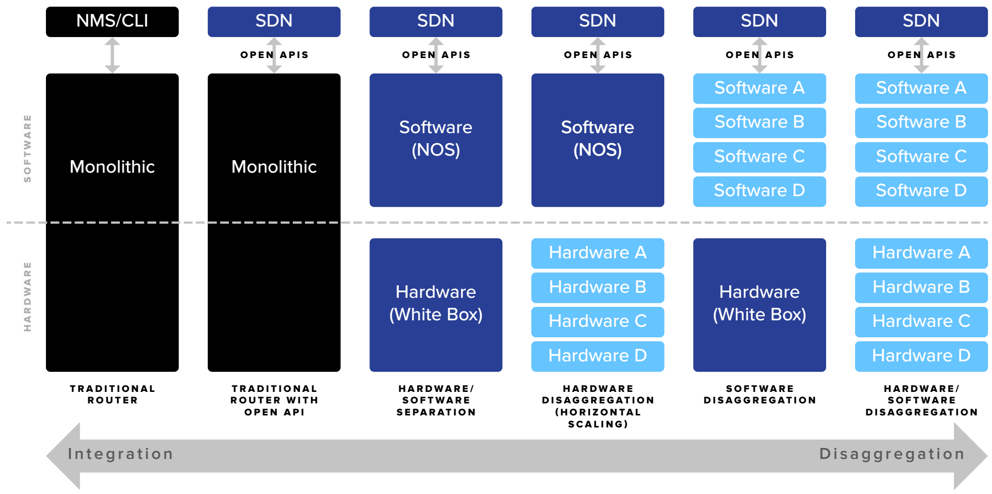

Figure 3: The Router Disaggregation Spectrum

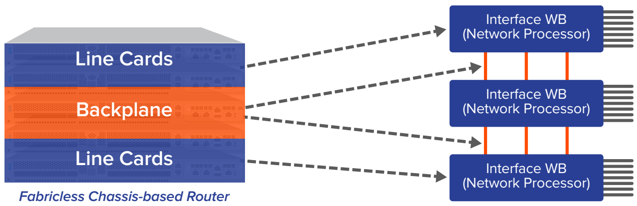

Figure 4: Multi-unit Scaling: Stacking

Stacking provides the simplest option for horizontal scaling, with the line cards becoming white boxes and the backplane becoming external interconnects, as shown in Figure 4. Another option when even more scalability is required is to put the fabric silicon into "fabric" white boxes. The line cards become "interface" white boxes, with the backplane again replaced by external interconnects, as shown in Figure 5.

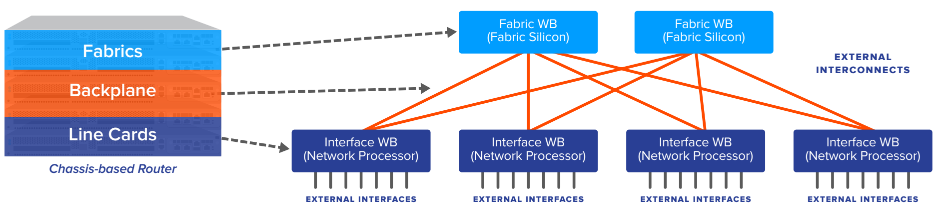

Figure 5: Multi-unit Scaling: Fabric-based

A point of delivery (POD) for MEC, as shown in Figure 6, provides another example of horizontal scaling, with spine and leaf switches providing a data-center-like fabric for interconnecting servers running VNFs and leaf routers providing IP routing for connectivity to the rest of the network.It is also possible to combine this example and the previous multi-unit scaling example with the leaf routers in Figure 6, actually multi-unit nodes as shown in Figures 4 and 5.

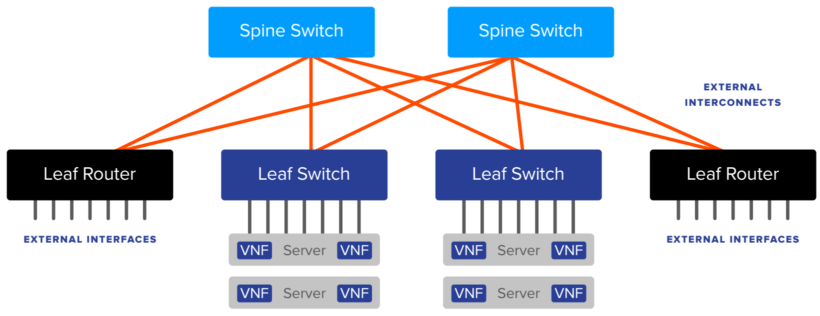

Figure 6: Leaf-Spine for Multi-access Edge Compute POD

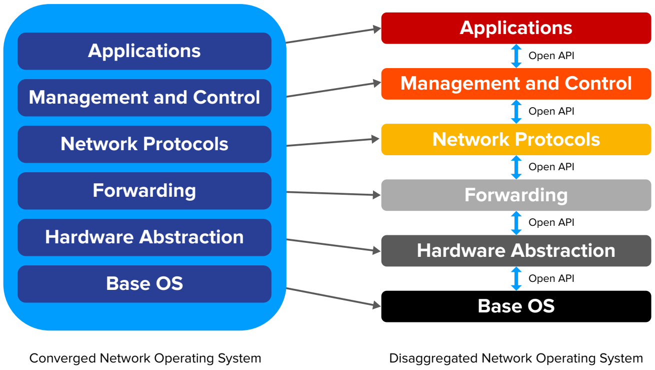

Figure 7: NOS Disaggregation

Another possibility, shown in Figure 7, is to disaggregate the NOS itself, breaking it into specific functional blocks with open APIs between blocks. These blocks could include applications, routing protocols, data plane forwarding, hardware abstraction, management and control, and the base OS kernel. This would enable a network operator that wants to develop an innovative application or routing protocol to do so without building the complete NOS from scratch. However, leveraging this form of NOS disaggregation requires significant expertise, and is likely to remain a niche use case confined to only the largest and most adventurous network operators.

Industry Maturity Enables Disaggregated Routing

But why now? As the router industry has matured, merchant silicon and open source software have greatly reduced the barriers to router disaggregation. High-performance silicon is now available from a number of merchant vendors and is already widely used by leading router vendors in many of their traditional router product lines. Open Network Linux (ONL) provides a powerful platform for network operating systems, enabling developers to leverage a library of basic open source functions while focusing development resources on more advanced functions and differentiation. ONIE is another enabler. As discussed previously, it provides an open environment for installing a compliant NOS onto third-party white box hardware. Cost-effective high-speed 100 Gigabit Ethernet (GbE) and 400 GbE pluggables are a fourth enabler and can be used for external interconnects as an alternative to chassis backplanes.

Merchant silicon	High-performance silicon, including network processors and fabric chips, available from merchant silicon vendors.
ONL	Provides a platform for NOS development leveraging a library of basic router functions.
ONIE	An open source initiative that defines an open install environment for white box switches.
High-speed Ethernet pluggables	100 GbE and 400 GbE. Optical and direct attach copper (DAC) cable assemblies. Cost-effective alternative to chassis backplanes.

Table 3: Key Router Disaggregation Enablers

Leading Operators Embrace Router Disaggregation

As with many trends, the leading internet content providers (ICPs) were the first to embrace disaggregated routing by building their own routers with merchant silicon and a combination of in-house or open source software. More recently, leading network operators have begun to embrace the concept of disaggregated routing. AT&T acquired the vRouter software from Brocade in 2017, and in 2018 announced plans to deploy 60,000 disaggregated cell site routers.

Organization	Project	Lead Operators
The Linux Foundation	DANOS project	AT&T
Open Compute Project	Out Side Plant (OSP) Cell Site Router	AT&T
TIP	OOPT project	Facebook
	DCSG	Vodafone, Telefónica, TIM Brasil, Orange, BT

Table 4: Router Disaggregation Industry Initiatives

The Disaggregated Network Operating System (DANOS) project, hosted by The Linux Foundation and building on AT&T's disaggregated router software (dNOS), is another proof point, while TIP's Open Optical & Packet Transport project also addresses the need for disaggregated packet and optical systems, with the Voyager reference design combining white box packet switching hardware and coherent optics. More recently, TIP began a sub-group in collaboration with Vodafone, Telefónica, TIM Brasil, Orange, and BT to define the Disaggregated Cell Site Gateway (DCSG) that connects base stations to the rest of the transport network, while AT&T released the hardware specification for its disaggregated 5G cell site router to the Open Compute Project.

The Benefits of Disaggregated Routing

Lower Capex and Opex

Enabled by reduced vendor lock-in, disaggregated routing provides benefits including competitive pricing, more choice, faster innovation, and cost-effective scaling, which all ultimately have the potential to significantly lower both CapEx and OpEx, as shown in Table 5.

<table>
<tr><th></th><th></th><th>Lower CapEx</th><th>Lower OpEx</th></tr>
<tr><td rowspan="3">Reduced Vendor Lock-in</td><td>Competitive Pricing</td><td>✓ Competitive pricing and selection
✓ Competitive upgrade pricing</td><td>✓ Competitive service and support pricing</td></tr>
<tr><td>More Choice</td><td colspan="2">✓ Mix and match best hardware and best software = more optimal price/performance</td></tr>
<tr><td>Faster Innovation</td><td>✓ Reduced cost per bit
✓ More efficient routing</td><td>✓ Smaller footprint
✓ Lower power consumption</td></tr>
<tr><td></td><td>Cost-effective Scaling</td><td>✓ Reduced up-front CapEx
✓ Extended solution lifespan
✓ Simplified sparing</td><td>✓ Optimized space and power
✓ Fewer part numbers
✓ Simplified sparing</td></tr>
</table>

Table 5: Disaggregated Routing Lowers CapEx and OpEx

Reduced Vendor Lock-In

Traditional monolithic routers force operators to purchase a bundle of hardware and software from a single vendor, then lock them in for the duration of the router's lifecycle in the network. With a small number of vendors dominating the service provider router market, competitive pricing pressures are muted. By lowering the barriers to entry for both hardware and software vendors, disaggregation has the potential to significantly increase price competition. Reduced vendor lock-in is also a key enabler for more choice and faster innovation.

More Choice

With traditional routing, network operators are forced to purchase a bundle of hardware and software from a limited number of vendors. As shown in Figure 8, if vendor A has the best software and vendor B has the best hardware, running vendor A's software on vendor B's hardware is not an option, leaving the operator with a set of suboptimal choices that ultimately results in higher costs and/or less than ideal performance.

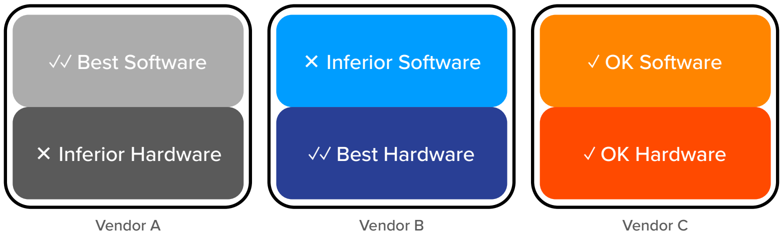

Figure 8: Traditional Routing Can Result in Suboptimal Solutions

Disaggregated routing enables network operators to optimally mix and match the best hardware and the best software for a given use case. Not all white boxes are the same. Interfaces, capacity, buffering, latency, features, and functionality will differ from white box to white box. Carrier-class white boxes, such as the Infinera DRX Series, also need to address requirements like synchronization, redundancy, environmental hardening, sub-300-millimeter depth, and stacking interfaces. Likewise, not all NOSs are the same. Different use cases will require different features and protocols. A carrier-class NOS, such as Infinera's Converged Network Operating System (CNOS), will also require a degree of protocol "hardening" that may not be required for an enterprise-class NOS, as well as high-availability features such as hitless software upgrades and multi-unit scaling. However, with at least 15 vendors currently providing multiple white boxes and at least 10 current NOS options, disaggregated routing provides an unprecedented level of choice with the flexibility to mix and match white boxes and NOSs to best match the requirements of a particular operator/use case, as shown in Figure 9.

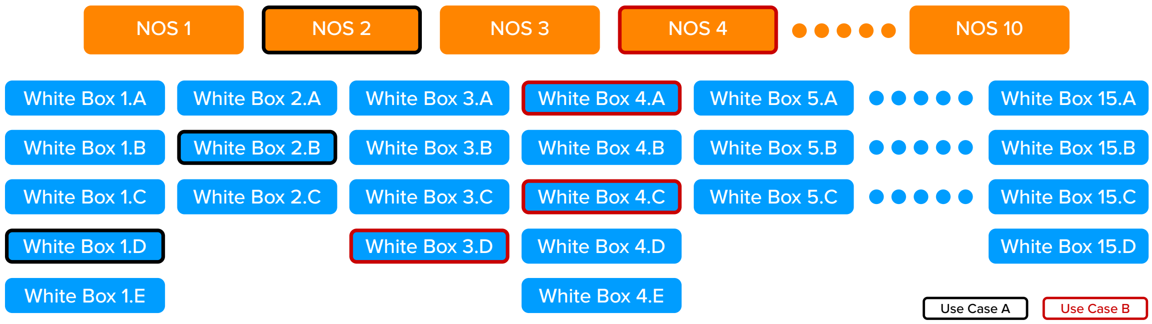

Figure 9: More Choice = More Optimal Solutions

Faster Innovation

Innovation is key to driving down CapEx and OpEx, scaling performance, and enabling new services. With a traditional monolithic router, innovation is constrained by the capabilities of the single selected vendor. The disaggregated router approach enables network operators to leverage the innovation capabilities of the entire ecosystem, selecting best-in-class vendors and upgrading based on the innovation cycle of each component of the disaggregated solution. As discussed previously, with reduced barriers to entry, new and smaller vendors will now be able to compete, bringing their innovation capabilities to the market.

Cost-effective and Flexible Scaling

Disaggregated routing lets operators avoid having to choose between a fully utilized device that meets initial requirements but will not scale to meet future requirements and an underutilized chassis with high up-front CapEx, footprint, and power consumption. Starting with the low up-front cost of a single white box, horizontal scaling, including stacking and fabric-based multi-unit scaling, enables capacity to be added very incrementally, as shown in Figure 10. With this approach, disaggregated routing can also address the capacity requirements of different parts of the network, all leveraging a very limited set of individual white boxes and a common NOS.

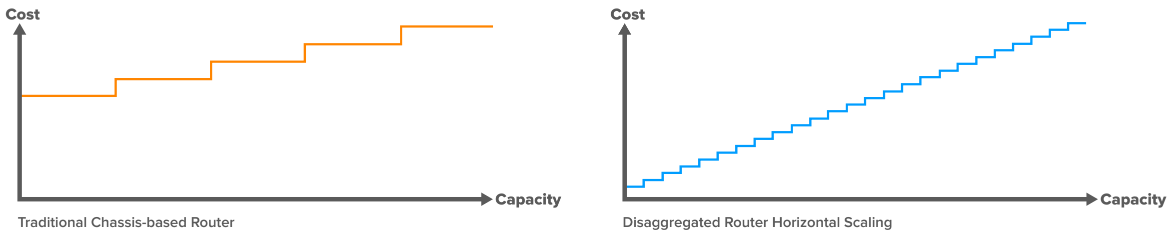

Figure 10: Cost-effective Horizontal Scaling

Use Case 1: 5G Transport

5G, which provides up to 10 Gb/s to users and a thousandfold increase in the capacity per unit area, as well as lower latency, higher availability, and sliceability, requires a major upgrade of the transport infrastructure. And while the scalability of IP routing has long made it the technology of choice for mobile backhaul, it has an even stronger role to play in 5G as traffic patterns become more meshed and less predictable with the adoption of multi-access edge computing, as shown in Figure 11.

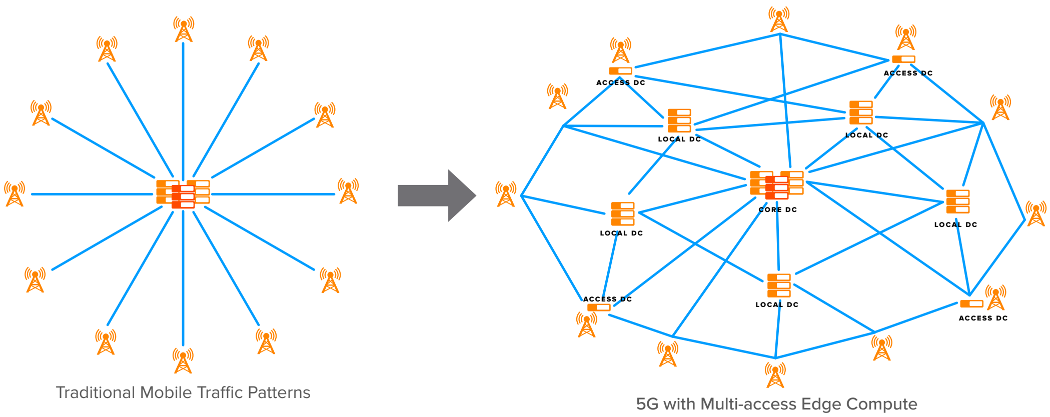

Figure 11: Meshed Traffic Patterns of 5G Require IP Routing

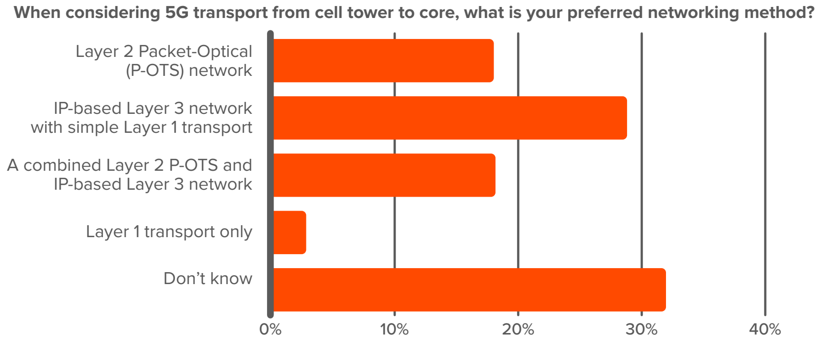

Figure 12: 5G Network Architecture Survey Question on Preferred Networking Method

This adoption of MEC is required to meet the low latency requirements of 5G's ultra-reliable low-latency communications (URLLC), to scale for massive machine-type communications (mMTC), and to support emerging applications such as augmented reality. In a recent Infinera-sponsored report, "5G Network Architectures: From X-haul transport to Multi-layer Service Slicing" by Mobile World Live, IP-based Layer 3 was the most popular network choice for 5G transport, as shown in Figure 12.

5G therefore represents a key use case for disaggregated routing, as evidenced by the TIP DCSG and OCP Cell Site Router initiatives discussed previously and the results of another recent Infinera-sponsored report, "Operator Strategies for Disaggregation in 5G Transport Networks" by Mobile World Live, shown in Figure 13. In this survey the majority of respondents indicated they were "very likely" or "somewhat likely" to include disaggregation in their 5G transport.

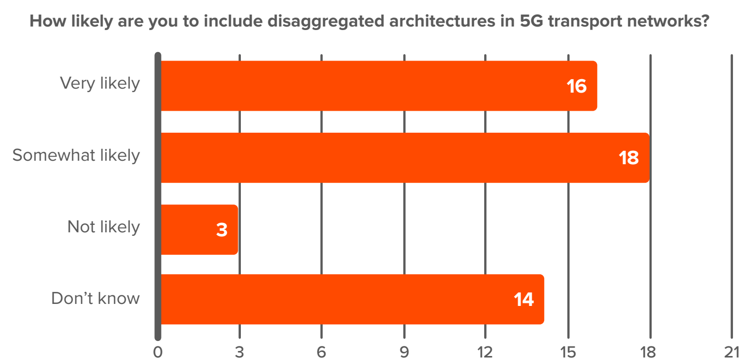

Figure 13: Disaggregation and 5G Transport Survey Question on Disaggregated Architectures

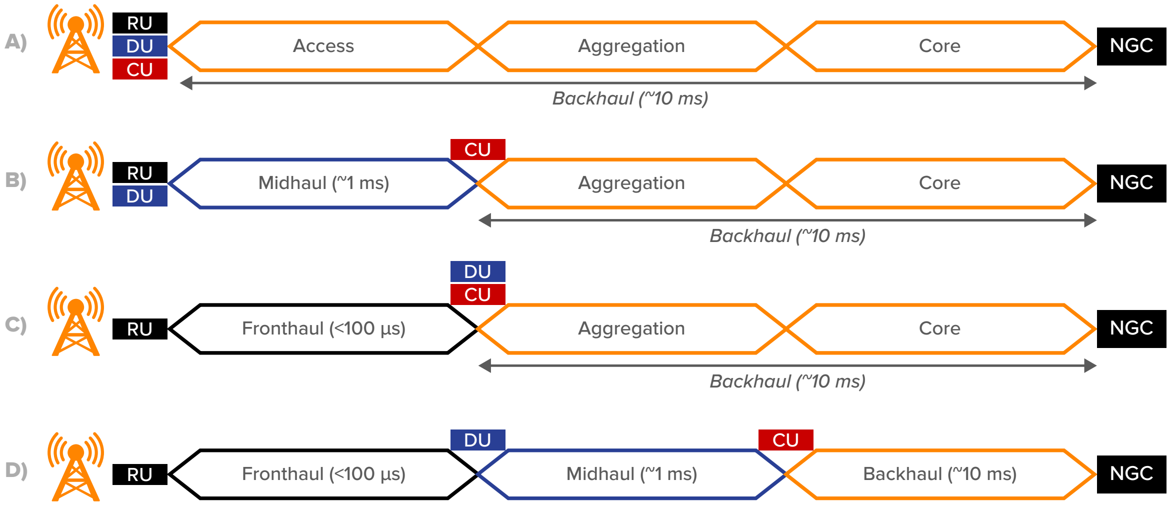

Figure 14: 5G X-haul Options

The question, though, is which parts of the 5G transport network can disaggregated routing address? To answer this question, we need to look at the options for how to split the functionality of the base station, as shown in Figure 14. At a high level, it splits into the radio unit (RU), the distributed unit (DU), and the centralized unit (CU). However, there are a number of options for how to split or combine these.

The RU, DU, and CU can be colocated like in a traditional base station, as shown as Option A. Option B combines the RU and DU to create a midhaul network and a backhaul network, while Option C combines the DU and CU to create a fronthaul network and a backhaul network. Option D completely separates the RU, CU, and DU with fronthaul, midhaul, and backhaul networks. In a recent 5G survey, Option C was the most popular selection, while Option D was just ahead of Option A. However, early non-standalone 5G deployments focused on enhanced mobile broadband (eMBB) are likely to be weighted in favor of Option A, evolving to the more distributed options (B, C, and D) as mobile operators move to standalone 5G and extend their service portfolios to URLLC and mMTC.

While disaggregated routing can provide a solution for all aspects of 5G X-haul, there are some specific requirements regarding scalability, availability, latency/jitter, and environmental hardening depending on the type of X-haul and location in the network, as shown in Table 6. Environmental hardening is likely to be required for access locations, while multi-unit scaling is likely to be required for the aggregation and core parts of the backhaul network, with stacking in the aggregation nodes and fabric-based scaling in the core. Regarding latency, while there are some inconsistencies around the exact requirements, as a rule of thumb backhaul requires maximum latency of around 10 milliseconds (ms), while midhaul is an order of magnitude lower at 1 ms, and fronthaul is another order of magnitude lower at 100 microseconds (µs). With merchant network processors able to deliver latencies of under 10 µs, and even under 3 µs, disaggregated routing can provide a good fit for the access, aggregation, and core network parts of a backhaul network and for midhaul.

Location	X-haul	Option	Requirements
Core	Backhaul	All	Stacking and fabric-based multi-unit scaling.
Aggregation	Backhaul	A, B, C	Higher capacity. Stacking.
	Midhaul	D	Higher capacity. Stacking. Lower latency.
Access	Backhaul	A	Lower capacity. Environmental hardening.
	Midhaul	D	Lower capacity. Environmental hardening.
	Fronthaul	C, D	Lower capacity. Environmental hardening. Must meet stringent latency, jitter, and packet loss requirements.

Table 6: Disaggregated Routing and 5G X-haul

Fronthaul, however, is where it gets interesting. While the eCPRI protocol used for 5G fronthaul can be transported over Ethernet or UDP/IP, according to IEEE 802.1CM (Time-Sensitive Networking for Fronthaul), one-way latency must be less than 100 μs , which, assuming 10 kilometers (km) at 5 μs per km, leaves 50 μs for transport equipment and 5 μs or 10% of latency for jitter, while packet loss must be better than 10^{-7}. Of these requirements, jitter is probably the most challenging and requires a network processor that delivers very flat latency. 802.1Qbu preemption, which can pause the transmission of a lower-priority packet that has already begun transmission when a higher-priority packet arrives, can help, especially on 10 GbE or 25 GbE interfaces.

Use Case 2: Distributed Access Architecture for Cable MSOs

Cable MSOs are evolving to distributed access architectures, distributing the PHY or MAC/PHY function of the Converged Cable Access Platform (CCAP)/cable modem termination system (CMTS) to the fiber node, with digital packet-based transport replacing the analog transport that previously extended to the fiber node, as shown in Figure 15. This, together with fiber deep, which pushes the fiber node closer to the home, and DOCSIS 3.1/Full Duplex DOCSIS 3.1 are enabling cable MSOs

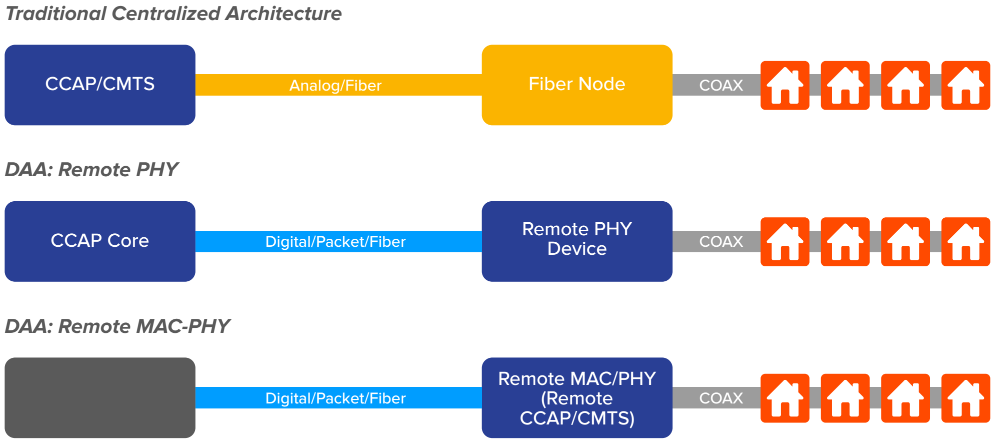

Figure 15: Distributed Access Architecture

to scale access network capacities to 10 Gb/s. Furthermore, cable MSOs are building converged interconnect networks as platforms for residential, business, and wholesale services, including 5G transport. This all requires a major upgrade of the transport infrastructure.

Layer 3 (IP Routing, IP/MPLS, Segment Routing)	**Layer 2** (Carrier Ethernet, MPLS-TP)
✓ Scalability ✓ Simplified operations (one end-to-end control plane) ✓ Multi-failure resiliency ✓ Resiliency less dependent on topology	✓ Simplified design and configuration ✓ Less expensive?

Table 7: Layer 3 vs. Layer 2

Layer 3 and Layer 2 both provide options for this packet transport. Layer 3 offers a more scalable solution with the ability to support a much larger number of nodes. Layer 3 can also protect against multiple failures, and its resiliency mechanisms have less dependency on network topology than Layer 2 protection mechanisms such as G.8032 Ethernet Ring Protection and Spanning Tree Protocol. Layer 2 technology is often perceived as simpler and more cost-effective. However, while the initial network design and configuration for Layer 2 might be simpler, once the IP network has been designed and configured, ongoing operations are greatly simplified by the IP control plane. And while Layer 2 has historically been lower in cost than Layer 3, with merchant silicon network processors and disaggregated routing, the price delta between the two is quickly diminishing.

While cable MSO networks and the terminology used to describe the different sites vary significantly, Figure 16 shows a generic network structure from the PoP/data center to the optical node. First is the PoP/data center, the functions of which may include broadcast headend, video on demand content hosting, internet peering, and connectivity to the long-haul network. Next is the primary hub, where the CCAP core would typically be located in a Remote PHY deployment. Then comes the secondary hub, which provides a conditioned aggregation point closer to the optical node. Pre-DAA, the CCAP or CMTS and edge quadrature amplitude modulation (QAM)

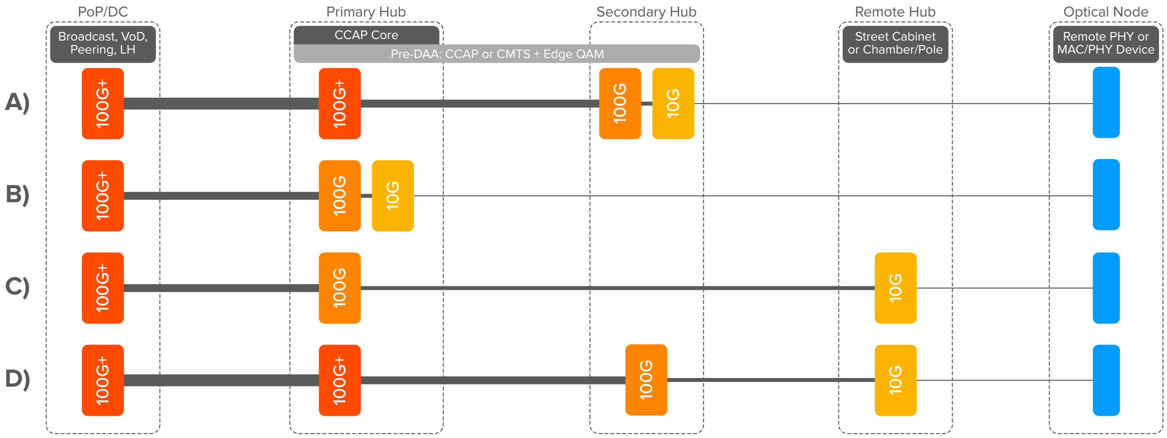

Figure 16: DAA Aggregation Scenarios

would have typically been located at either the primary or secondary hub. The remote hub provides an outside aggregation point closer to the optical node, in a street cabinet or chamber or on a pole. Finally, the optical node itself is where Remote PHY or Remote MAC/PHY devices would be located and where the fiber network meets the coax network, as shown previously in Figure 15.

Four common scenarios for building a CIN are shown. Scenario A includes 100 gigabit (100G)+ aggregation at a primary hub and both 100G and 10G aggregation at secondary hubs. Scenario B runs 10G all the way from the primary hub to the fiber node, with no intermediate aggregation. Scenario C moves the 10G aggregator to the remote hub. Scenario D provides a variant on scenario C with 100G aggregation at a secondary hub.

Location		Scenarios	Requirements
PoP/Data Center: 100G+ Aggregator		All	Fabric-based multi-unit scaling.
Primary Hub	100G+ Aggregator	A, D	Higher-capacity white box, stacking.
	100G Aggregator	B, C	High-capacity white box, stacking.
	10G Aggregator	B	Lower-capacity white box, stacking.
Secondary Hub	100G Aggregator	A, D	High-capacity white box, stacking for high availability.
	10G Aggregator	A	Lower-capacity white box, stacking for high availability.
Remote Hub	Street Cabinet	C, D	Temperature-hardened 10G aggregator.
	Chamber/Pole	C, D	Lowest capacity. Temperature-hardened. Compact for weatherproof clamshell enclosure.

Table 8: Disaggregated Routing and DAA

As shown in Table 8, disaggregated routing can provide an ideal solution for packet aggregation in the primary, secondary, and remote hubs, as well as the PoP/data center. In the primary hub, stacking may be required for both scaling and high availability, while the PoP/data center is more likely to require fabric-based multi-unit scaling. Stacking may also be used for high availability at the secondary hub. Temperature hardening is required at remote hubs, and in the case of chamber/pole deployment, a compact form factor that will fit in a weatherproof clamshell enclosure is also required.

Disaggregated Router Challenges

In the same Infinera-sponsored report referred to previously, "Operator Strategies for Disaggregation in 5G Transport Networks," the largest number of respondents saw the potential of disaggregated routing "if issues are resolved," as shown in Figure 17. These issues include

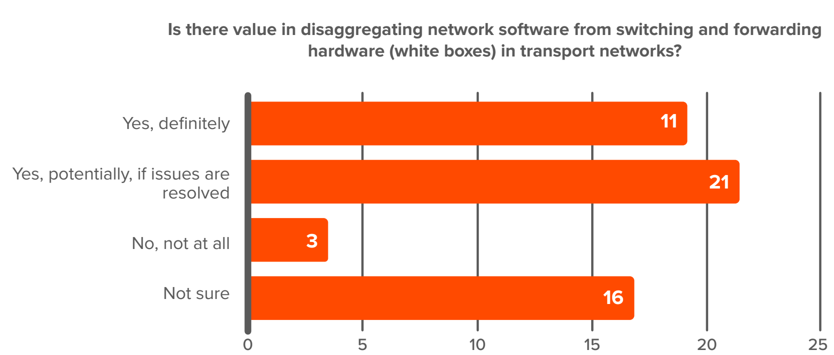

Figure 17: Disaggregation Value Survey Question

integration and interoperability between the NOS and white box, service and support, and perceived added management complexity.

Integration, Service, and Support

With a traditional monolithic router, a single vendor is responsible for integrating all hardware and software, as well as providing service and support. As shown in Figure 18, multiple models are available to address requirements for integration, service, and support for disaggregated routing.

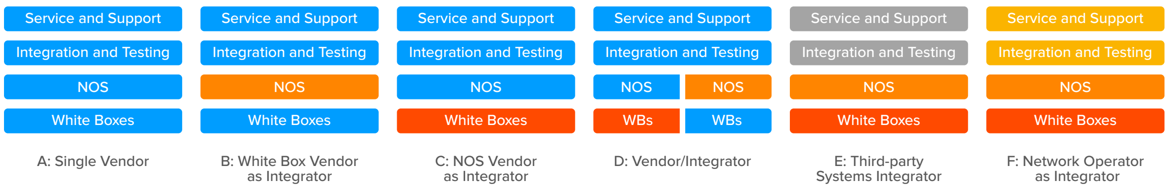

Figure 18: Options for Testing, Service, and Support

A single vendor (A) can provide both the white boxes and NOS, which reduces risk while providing the option to migrate to third-party NOSs or white boxes in the future. The white box vendor (B) or NOS vendor (C) can also take responsibility for integrating and testing the solution and providing service and support. In addition, it is possible to mix these options, with a single vendor (D) offering its own NOS and white boxes as well as third-party white boxes and/or NOSs to optimally address a broader set of requirements and use cases. Alternatively, a third-party systems integrator (E) or the network operator itself (F) can fill the integrator role, taking responsibility for service, support, integration, and testing. As a white box and NOS vendor with an experienced global services organization, Infinera is able to support all these options.

Management

A monolithic chassis-based router provides a single network element to manage. Providing the same level of routing capacity for a disaggregated router leveraging horizontal scaling could result in many more individual white boxes to manage. One example of a solution to this issue is for a carrier-class NOS, such as Infinera's CNOS, to make a multi-unit configuration behave as a single network element. A second example is for a function within the network/control system, such as Infinera Transcend's vPOD, to abstract a collection of "independent" white box switches, white box routers, servers, and VNFs as a single network element (NE) to the user, hiding complexity, as shown in Figure 19.

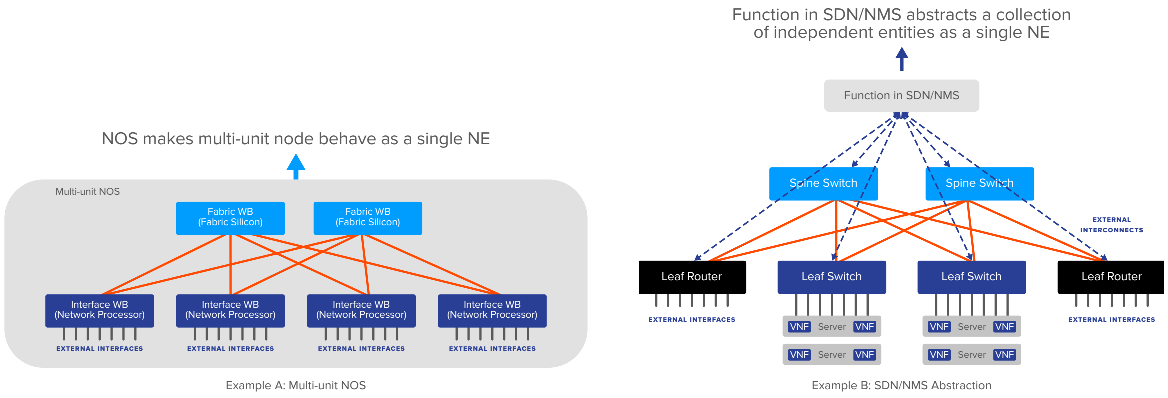

Figure 19: Managing Multi-unit Nodes as a Single NE

In addition, network management and control systems, such as Infinera's Transcend, support traditional routers, white boxes, and NOSs from multiple vendors and enable network operators to seamlessly transition from traditional routers to disaggregated routing. These management and control systems can also operationalize the new disaggregated router environment and provide high levels of automation for both traditional and disaggregated routers.

Summary

5G and DAA are both driving major transport network upgrades, with IP routing playing a key role due to its scalability, flexible resiliency, and the operational simplicity of a single end-to-end control plane. Providing an alternative to traditional monolithic architectures, disaggregated routing enables network operators to gain these benefits of IP routing while also benefiting from minimized vendor lock-in, competitive pricing, more choice, faster innovation, and cost-effective scaling, all with the potential to significantly lower both CapEx and OpEx.

Infinera®

Infinera Corporation
140 Caspian Court
Sunnyvale, CA 94089 USA
Telephone: +1 408 572 5200
Fax: +1 408 572 5454
www.infinera.com

Have a question about Infinera's products or services?
Please **contact us**.

Specifications subject to change without notice.

Document Number: 0195-WP-RevA-0619