Perspective

Making the case for the true costs, benefits and risks of disaggregation in CSP IP networks

August 2021

Gorkem Yigit, Gilles Monniaux and Caroline Chappell

RESEARCH

Contents

List of figures

1. Executive summary

In this perspective, we discuss the key findings of Analysys Mason's total cost of ownership (TCO) study that analyzes the capex and opex requirements for the implementation of a modular chassis (scale-up) and a range of disaggregated routing architecture for communications service provider (CSP) IP core networks. We built a 5-year TCO model based on the real-life cost and traffic parameters of a large Tier-1 CSP's network in a developed market. We also assessed the operational processes and cost implications of a range of network disaggregation scenarios based on our interviews with several CSPs and ecosystem players. Figure 1.1 summarizes the results of our TCO model and Section 3 provides more details about the financial analysis (capex, opex and cumulative TCO with depreciation and WACC) of each scenario.

Figure 1.1: TCO for different IP core routing architecture scenarios for a Tier-1 CSP in a developed market[1]

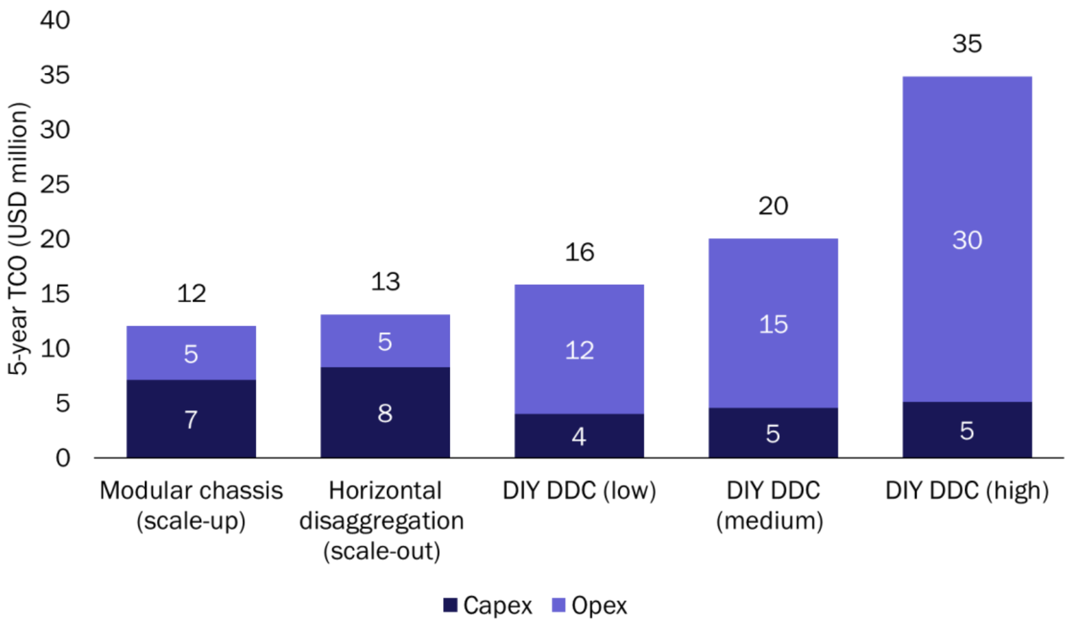

Source: Analysys Mason, 2021

Our research and TCO model indicate that network disaggregation is not currently inherently better than other architectural approaches in terms of TCO. CSPs' architectural choices depend on many factors, including the particular network domain under consideration and the CSP's larger business and operational objectives.

1 DIY DDC stands for do-it-yourself Disaggregated Distributed Chassis.

1.1 Modular chassis platforms are currently the most cost-effective option for CSPs that want to keep network costs down against the backdrop of increasing traffic and capacity demands

Our TCO model shows that modular chassis-based routers (scale-up) with pay-as-you grow pricing provide the optimal 5-year TCO for the IP core; it is around 8% lower than that for the horizontal disaggregation (scale-out) scenario. More significantly, the same chassis-based router approach provides a 24–66% lower TCO than the do-it-yourself (DIY) white-box disaggregation scenarios using the Disaggregated Distributed Chassis (DDC).

The DDC approach combines horizontal scale-out with vertical (or software) disaggregation. It uses merchant-silicon-based white-box platforms and independent network operating system (NOS) and routing stack software from different vendors, with a proprietary cell-based interconnect fabric between the platforms, as explained further in Section 2.

The favorable TCO for chassis-based routers is largely due to the high density, massive scale and performance benefits of these platforms, combined with the flexible pricing for capacity expansions and CSPs' existing familiarity with their operations.

1.2 Horizontal scale-out enables CSPs to get the granular scalability and improved reliability benefits of disaggregation without incurring high opex and additional risks

The horizontal disaggregation (scale-out) scenario, which uses both merchant- and custom-silicon-based platforms, has a slightly higher TCO than the modular chassis approach, mainly due to the costs of additional interconnections/cabling. However, this model provides several operational benefits over modular chassis systems, including more-granular capacity upgrades, improved network reliability (thanks to the smaller blast radius) and IP-based interconnect fabric interoperability. Having a vendor partner that can carry out all of the pre- and post-deployment activities with economies of scale in engineering and operations provides further benefits, while significantly reducing the costs and risks of disaggregation.

1.3 The high operational costs and complexities of reaggregation reduce the financial viability of DIY DDC disaggregation and increase the risk

The DDC scenarios, which use merchant-silicon-based white boxes and NOS software from separate vendors, offer potentially more choice than the other options (thanks to a broader vendor ecosystem), as well as capex savings (due to lower hardware costs). However, our research shows that the price gap between white boxes and traditional routers is narrowing. Incumbent vendors have responded to the evolving market conditions by building out new portfolios of competitive fixed form factor platforms, which use the same merchant silicon as white-box routers but do not require DIY integration. However, the overall TCO for all three DDC scenarios was significantly higher than that for the scale-up and horizontal disaggregated scale-out models, despite the capex savings. This is because of the higher opex for the complex and expensive reaggregation processes and operational change requirements of the DIY disaggregation approach, which is discussed in Section 3.2. The breadth and scope of these additional costs, including the specific calculations and parameters used in our model, are based on numerous interviews with vendors and CSPs for this particular project, as well as additional research and the aggregate expertise gained by Analysys Mason through relevant analytical work.

1.4 Disaggregation comes with new implications for network reliability, multi-vendor management, compatibility, openness and futureproofing

We have identified several important operational factors that CSPs need to consider, in addition to TCO, when they are evaluating different architectural options (Section 3.3). One of these factors is openness. This is a clear goal of disaggregation in order to expand the vendor choice for CSPs and extend the ecosystem of partners to co-innovate with. However, disaggregation solutions based on proprietary, cell-based fabrics (such as DDC) can limit the hardware/silicon choices and increase vendor lock-in. Disaggregation solutions based on IP-based fabric architecture, on the other hand, are grounded in interoperability and facilitate multi-vendor networks through their support of a range of options.

Disaggregation can also introduce network reliability issues, which will require new approaches to network management and automation. For example, out-of-band (OOB) management for the DDC model poses new Mean Time Between Failures (MTBF) challenges. Also, both the DDC and horizontal disaggregation approaches bring the additional complexity of isolating and troubleshooting failures due to the increased number of externalized interconnections. In addition, there are several new risks and responsibilities that CSPs need to take on with DIY DDC disaggregation, including managing chipset upgrade cycles, backward/forward compatibility and migration.

2. CSP IP networking is evolving via a mix of architectural approaches

CSPs are faced with the growing cost pressure of supporting today's IP network traffic demands (driven by video, gaming and remote working) and future 5G traffic volumes, while their revenue remains largely flat. They have high ambitions for new 5G-, edge computing- and network slicing-enabled service opportunities, but these services will require new investments in the wide-area/transport network in order to connect numerous cells and edge locations and address the needs for ultra-low latency, high reliability and increased capacity. Enterprises' connectivity demands are also changing rapidly as companies increasingly shift mission-critical applications and workloads to multiple public and hybrid clouds and demand dynamically scalable and highly performant and secure connectivity solutions. Furthermore, competition in the connectivity market is becoming more intense, both within the telecoms industry and from the new types of agile, alternative service providers.

All of these trends and market forces are spurring CSPs to rethink and rearchitect their legacy IP networks and operations in ways that can reduce costs while meeting future capacity and service demands. Transforming these networks into more cost-effective, open, scalable, agile and automated environments is becoming a strategic imperative for CSPs in order to deliver differentiated, digital connectivity services and customer experiences while limiting costs (particularly opex, the largest component). Some CSPs are considering disaggregation as a potential solution to support these requirements while reducing costs, but network disaggregation can mean different things to different people and may have a significant impact on the total cost of ownership.

2.1 Network disaggregation is much more than just implementing white-box hardware

Network disaggregation is a highly discussed topic in the telecoms industry. CSPs are investigating whether network strategies deployed by hyperscalers can be emulated in their own network transformation initiatives. Hyperscalers pioneered disaggregated networking with large-scale implementations in their data centers in order to enable mix-and-match software and hardware for perceived cost advantages, as well as for service innovation and automation, but it should be noted that they are not typically taking the same approach in the routed WAN. Several CSPs are exploring the feasibility of adopting disaggregation approaches to transform their transport (IP/MPLS, optical) and access (RAN, BNG, OLT) networks. This is driving efforts by vendors, CSPs and industry initiatives such as Telecom Infra Project (TIP) and Open Compute Project (OCP).

Network disaggregation is often defined as the decoupling of network element hardware and software and is equated with the use of merchant-silicon-based white boxes and independent network operating system (NOS) and routing stacks (also known as vertical or software disaggregation). This is one example of disaggregation, but in reality, disaggregation is a much broader concept that takes on several forms, as illustrated in Figure 2.1.

Figure 2.1: Overview of the disaggregation approaches that are possible in the main layers of a network component

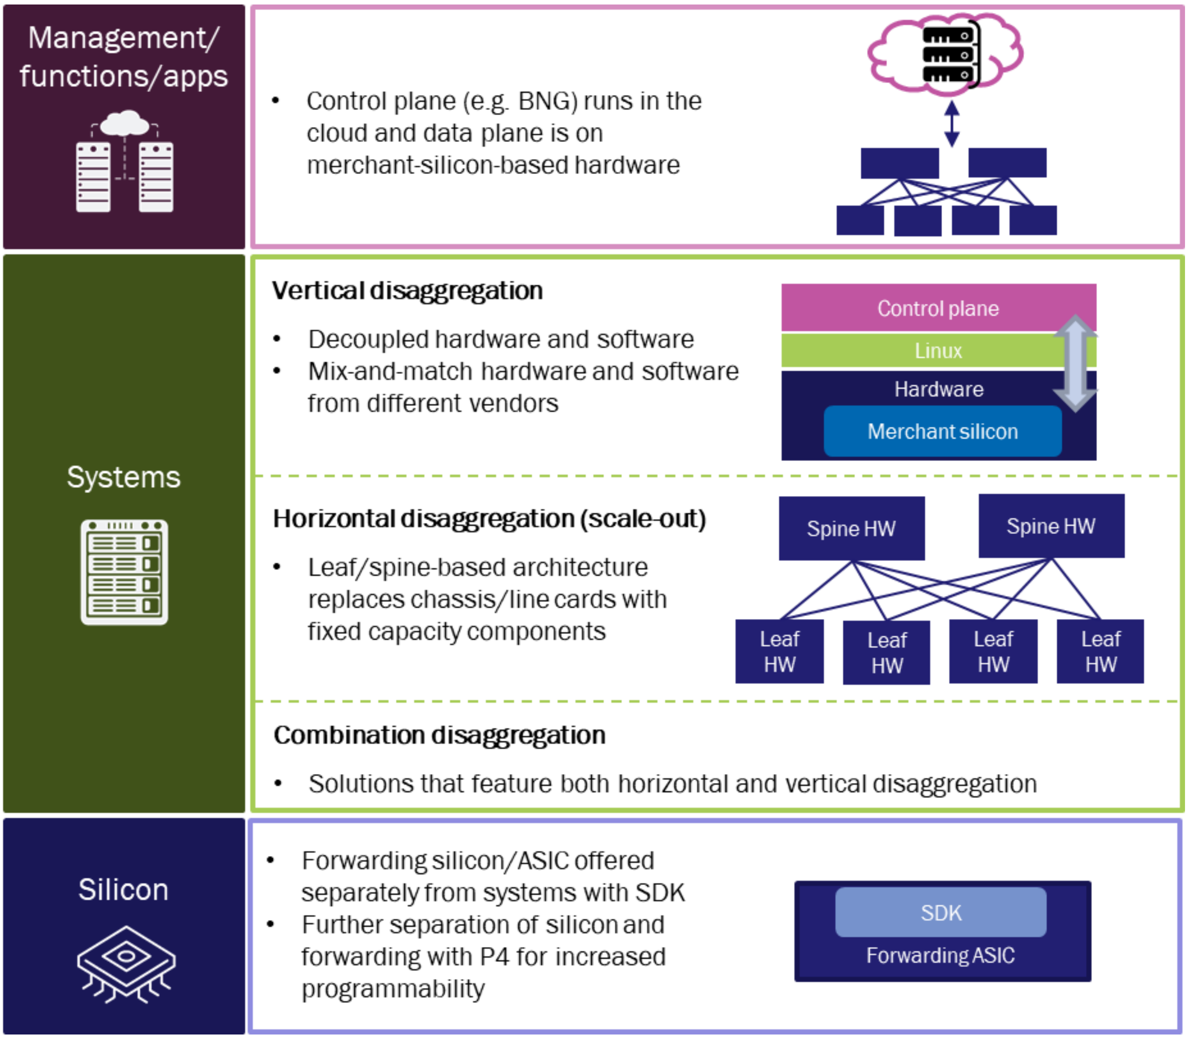

Source: Analysys Mason, 2021

Any integrated subcomponent of a network router can potentially be disaggregated across its various layers in the same way, whether these layers are part of the system (hardware and software) architecture, forwarding silicon or management, control and data planes. Before embarking on a disaggregation project, it is important for CSPs to clearly identify what operational and/or business challenges they are aiming to solve. These could include:

- reducing costs through price transparency, lower-cost hardware and increased competition
- standardizing/streamlining hardware architecture on merchant silicon
- improving network reliability with more-modular and granular architecture
- eliminating vendor lock-in and introducing new vendors into the ecosystem.

Disaggregated networking can potentially deliver benefits in all of these areas, but it is important for CSPs to develop a clear understanding of the costs, risks and operational implications of network disaggregation and to assess which type of disaggregation solution, if any, is the most beneficial approach for them. In this perspective, we focus on the different approaches to systems disaggregation (including horizontal and vertical (software) disaggregation), because this is the area that is currently receiving the most attention and it is imperative that CSPs thoroughly analyze their options.

2.2 CSPs must choose the right routing platform architecture for their IP core and edge networks

Growing traffic volumes due to 5G, cloud and digital services mean that CSPs must increase their capacity and implement more cost-effective network architecture, especially in their IP core and edge networks. They need to align their choice of routing platform and architecture with their business objectives and use case/network domain requirements, as well as their operational readiness. Figure 2.2 provides an overview of the main options for CSPs.

Figure 2.2: Main routing system options for CSPs' IP core and edge networks

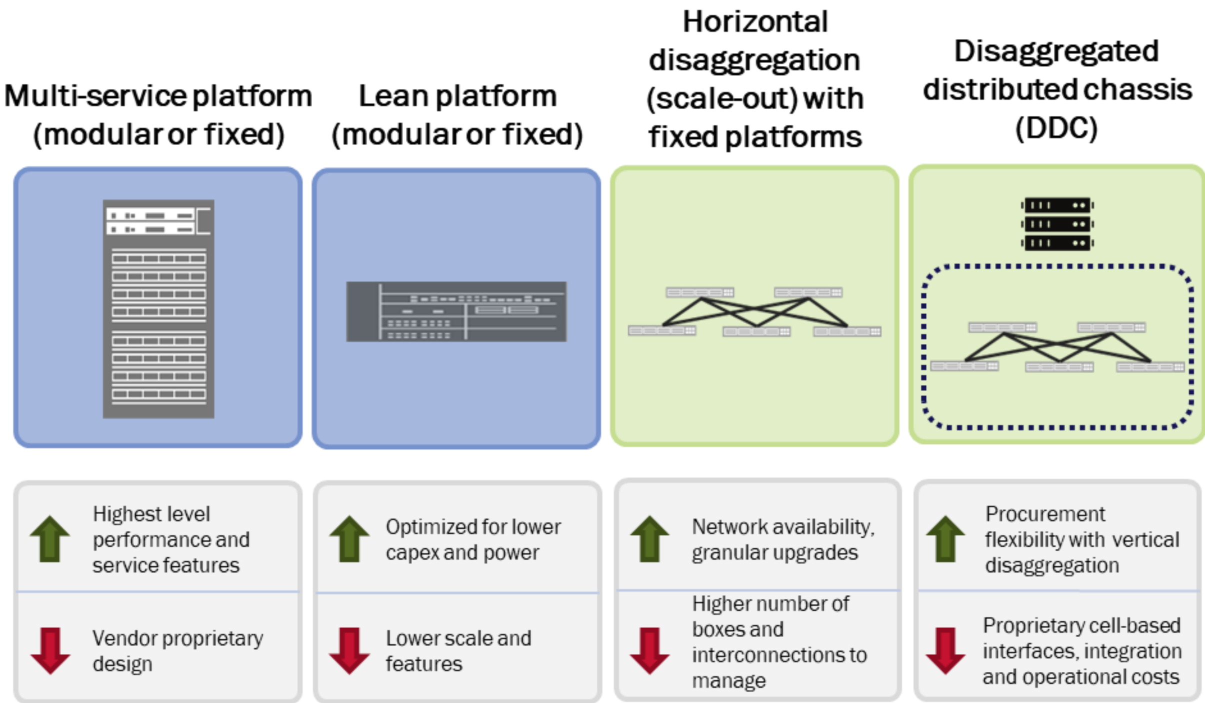

Source: Analysys Mason, 2021

Modular chassis routing systems, which can be based on custom or merchant silicon and use pluggable line cards to scale up capacity and service features, form the foundation of CSPs' IP core networks today. These scale-up routing platforms provide CSPs with the highest level of performance, service and automation features, but they have traditionally offered limited granularity for scale and expansion. In addition, these products typically command a premium up-front price. However, we note that many traditional vendors have recently started to offer ‘pay as you grow’, per-port pricing and additional software licensing flexibility in order to alleviate the previous issues of high up-front costs and coarse capacity and feature upgrades. Vendors have also introduced leaner platforms that are based on either custom or the same merchant silicon as white-box routers, and can offer lower costs by optimizing for lower scale, fewer features and/or less flexibility for expansion. These platforms are best-positioned for simpler deployment scenarios (peering, data center interconnect (DCI) and aggregation), as discussed in Section 3.4, but can also be deployed in leaf and spine components in an IP Clos architecture.

An alternative approach to the modular chassis route is the scale-out (horizontal disaggregation) model. This model breaks apart the chassis system into a leaf and spine Clos architecture; this has been a key pillar in data center designs perfected by hyperscalers. CSPs are now trying to take these concepts and apply them to the routed WAN. Scale-out designs are often based on fixed form factor platforms with merchant or custom silicon and lean feature sets, which offer many of the benefits of vertically disaggregated architecture. In addition, these platforms can provide greater vendor diversity provided that they use non-proprietary IP based interconnects. The main driver for this architecture is to enable more-granular upgrades and capacity expansion by distributing the interconnections between many smaller routers. This also can improve network availability thanks to a smaller ‘blast radius’ for outages. However, the connection of leaf and spine components comes at a cost and there is potentially more operational complexity due to the increased number of components to manage.

The Distributed Disaggregated Chassis (DDC) model is a relatively new approach for IP routing. It combines the scale-out model with vertical disaggregation by using merchant-silicon-based white boxes and NOS software from separate vendors. The NOS creates a cluster from the distributed white-box components of the DDC system that functions as a single router instance instead of a number of independent devices. The vertical disaggregation of this model offers supply chain flexibility and potentially more innovation because it provides a broader vendor ecosystem on top of the architectural benefits of the above-mentioned horizontal disaggregation approach. However, the reaggregation of these components and the associated operational changes can be highly complex and costly for most CSPs. AT&T has been spearheading this approach and has support from an ecosystem of vendors and industry initiatives (OCP and TIP OOPT), but it is unclear how much momentum the DDC approach for the core has beyond AT&T.

There is no one size fits all; each option has its advantages and limitations. CSPs are presented with the challenge of deciding how to mix and match these various options to maximize their financial and operational viability in the near term. Analysys Mason built a robust model based on a real-life Tier1 CSP’s network in a developed market to analyze the TCO of these various approaches and to help CSPs with their IP network architecture design and investment decisions.

3. TCO model results

3.1 Introduction and methodology

Analysys Mason, in collaboration with Juniper Networks, carried out a TCO study to examine the capex and opex required for the implementation of the main architectural choices for CSPs’ IP core networks (discussed in Section 2.2). The TCO model analyzes the following deployment options in a greenfield network:

- traditional, single-vendor modular chassis (scale-up)

- single-vendor, horizontal-only disaggregation (scale-out) with fixed form factor merchant- and custom-silicon-based platforms
- Disaggregated Distributed Chassis (DDC) using white-box hardware and independent NOS software with DIY integration by the CSP.

Analysys Mason developed a comprehensive 5-year TCO model for a hypothetical CSP that reflects the network traffic, design and cost parameters of real Tier-1 CSPs in developed markets. We interviewed several CSPs and ecosystem vendors about their experiences with disaggregated network deployments to collect quantifiable information and to verify our modelling and scenario assumptions.

Key modeling assumptions

The key modeling assumptions are provided in Figure 3.1.

Figure 3.1: Key modeling assumptions

Attribute	Assumption(s) made
CSP profile	• Large Tier-1 CSP that operates in a Western European country (60 million population). • Provides residential broadband (FTTH, xDSL), business broadband (FTTH, xDSL), business dedicated connections (<100 Mbit/s, 100 Mbit/s, 1 Gbit/s) and VPN services.
Network traffic	Network traffic data for each of the connectivity services is calculated based on the real-life values observed in the UK, France and Western Europe (see annex for details).
Network design	Network assets (ports, routers/switches and cables) are dimensioned across metro (MSANs and OLTs), edge and core nodes based on the traffic assumptions.
Capex parameters	Hardware, software licenses, system integration, installation and depreciation.
Opex parameters	Headcount, power and space and support and maintenance.
Length of analysis	5 years
Weighted average cost of capital (WACC)	10%

Source: Analysys Mason, 2021

We used Juniper Networks's routers and switches to model the TCO in the scale-up and scale-out scenarios with conservative discounts on the list prices. We used Edgecore white boxes for the DDC scenarios with optimal port density and estimated discounts on the list prices (Figure 3.2).

Figure 3.2: Network components used for each scenario

Architecture	Assets modeled
Traditional chassis (scale-up)	• Juniper PTX10008-PREM3 (100/400 GE) • JunOS
Scale-out	• Leaf: Juniper PTX10001-36MR-AC (100 GE) • Spine: Juniper QFX5220-32CD (400 GE) • JunOS
Distributed disaggregated chassis (DDC)	• Leaf: Edgecore AS7926-40XKE (100 GE) • Spine: Edgecore AS9926-24D (400 GE) • Generic x86 COTS hardware and Ethernet switches • NOS software

Source: Analysys Mason, 2021

Independent NOS vendors do not generally publish standard price lists and pricing can vary by CSP type and deployment scope. As such, we developed some benchmark estimates for NOS license and support costs for the DDC scenarios based on our conversations with industry stakeholders, vendors, customers and system integrator (SI) partners. In addition, we assumed that an SI partner would be needed to handle the end-to-end lifecycle management (design, integration, logistics, support and upgrades) of the DDC deployment. Our interviews with CSPs also revealed that a significant number of additional full-time employees (FTEs) are needed to carry out the necessary internal processes such as architecture design, multi-vendor management and operational change (network planning, software automation and physical operations). We tested all of these cost parameters for sensitivity, as shown in Figure 3.3.

Figure 3.3: Overview of the DDC model sensitivity analysis

Parameter	DDC (low) scenario	DDC (medium) scenario	DDC (high) scenario
NOS license cost as a percentage of the cost of white-box hardware	50%	75%	100%
NOS support cost as a percentage of the cost of the software license	15%	50%	100%
FTEs needed for architecture design (over 5 years)	2	2	2
FTEs needed for vendor management (over 5 years)	5	5	5
FTEs needed for initial set-up/support (over 3 years)	0	0	10

Source: Analysys Mason, 2021

3.2 The scale-up model provides the lowest TCO, even though the DDC models offer lower capex

Figure 3.4 shows the cumulative TCO for each of the five scenarios (scale-up, scale-out, DDC (low), DDC (medium) and DDC (high)). Figure 3.5 summarizes the TCO results, which include asset depreciation and WACC calculations.

Figure 3.4: Cumulative network TCO for each scenario over a 5-year period

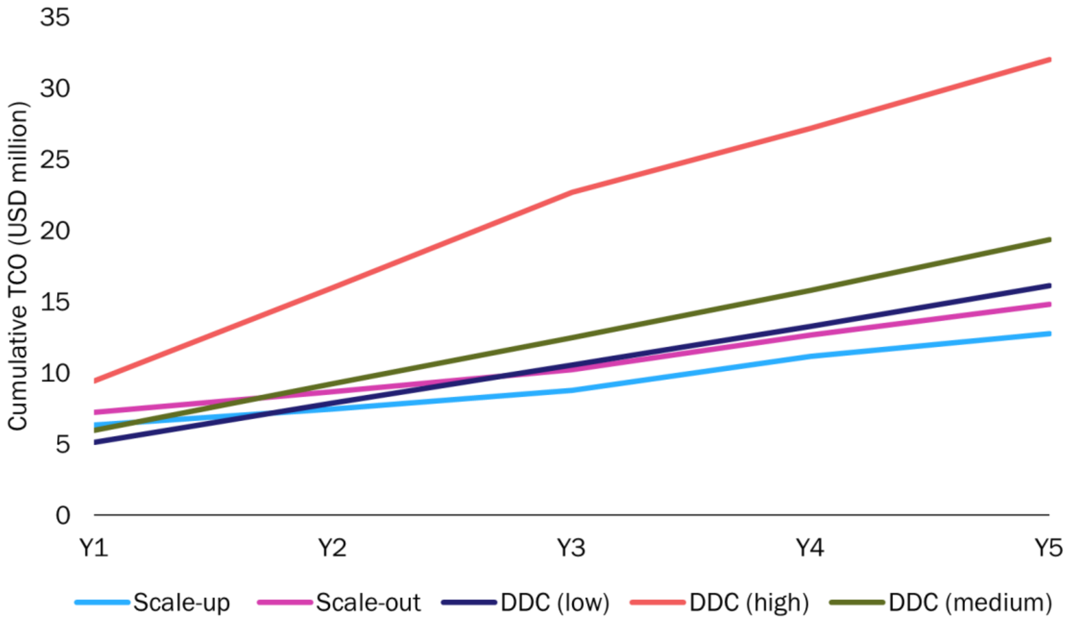

Source: Analysys Mason, 2021

Figure 3.5: TCO results overview

	Scale-up	Scale-out	DDC (low)	DDC (medium)	DDC (high)
Capex (USD)	7 084 844	8 276 600	4 003 478	4 549 162	5 094 846
Opex (USD)	4 929 732	4 822 075	11 786 627	15 468 043	29 744 939
Total TCO (after 5 years)	12 014 576	13 098 675	15 790 105	20 017 205	34 839 785
Difference in TCO from the scale-up scenario	0	+9.0%	+31.4%	+66.6%	+190.0%

Source: Analysys Mason, 2021

The scale-up scenario provides the lowest 5-year TCO of the five modeled scenarios. This shows that the higher unit costs of the scale-up platforms are justified over time. CSPs can reap the benefits of the high density, massive scale and strong performance while traffic volumes grow, and the flexible, pay-as-you grow pricing models enable more-granular capacity upgrades. CSPs are also already familiar with the operations of these platforms, which offers further advantages. Interconnection costs across the local switch fabric are another important factor: the scale-out and DDC designs require high-cost, external interconnections, and optics and fiber is required between the boxes. The horizontal scale-out scenario yields a comparable TCO to the scale-up scenario but with a slightly higher capex (+17%), mainly due to the additional fiber cables needed to connect the leaf and spine components. We also used both custom-silicon and merchant platforms in the scale-out model.

Our results for DDC approaches support the common industry belief that CSPs can achieve capex savings (for example, between 28–43% compared to the scale-up model) by using lower-cost, merchant-silicon-based white boxes and independent NOS solutions. However, it should be noted that it may not always be straightforward to realize these capex savings due to the non-transparent pricing for independent NOS software. Incumbent OEMs can provide additional discounts to narrow the gap, plus they can offer platforms that use the same merchant silicon as the white boxes, but with a pre-integrated NOS.

One of the key findings of our TCO study is that CSPs may face much higher opex and indirect capex when taking the DDC approach rather than the scale-up or scale-out approach. This may be detrimental to the business case. A multi-vendor, white-box-based disaggregation approach shifts several, complex network component lifecycle costs and responsibilities (Figure 3.6) onto CSPs, few of whom possess the DIY capabilities that are needed to handle them. All of the pre-deployment activities are carried out by incumbent OEMs with large R&D investments and economies of scale in engineering and operations in the traditional chassis and horizontal disaggregated (scale-out) models, and the associated costs are already included in the OEM's prices. Furthermore, OEMs provide a single point of contact in terms of support and upgrades, and offer the benefits of scale since these costs are shared and amortized over a large customer base. All of these factors also make the pre-integrated scale-out scenario more economically appealing than the DDC approach if a CSP wants the benefits of horizontal disaggregation without incurring high opex and additional risk.

Figure 3.6: Additional opex and indirect capex of DIY disaggregation

Step	Costs and operational requirements
Solution design, validation and integration	• It can take at least 6 months to develop a multi-vendor, white-box-based disaggregated solution from scratch (depending on a CSP's customization requirements). This process involves skilled architects and engineers. • Extra validation, certification, homologation and integration work must be done for every different white-box vendor, every generation of device (even from the same vendor) and every NOS. • Vendors and SIs are continuously building portfolios of solution blueprints, which can reduce the time and effort required, but they will need to have sizeable customer bases in order to significantly reduce costs.
Logistics	This includes the cost of supply chain management, sourcing and white-box/NOS installation. These are highly volume-driven, manual processes that need economies of scale to reduce costs.
Deployment, configuration and management	CSPs need to make the following upfront investments to prepare their operations for multi-vendor, white-box-based disaggregated networks: • software automation (templates, data models, configurations and NETCONF) • staff training, re/upskilling and hiring • changes in network planning and physical operations (such as pre-wiring).
Support and upgrades	CSPs are likely to face high support costs and long problem resolution times due to the more unusual combinations of vendors and hardware and software configurations. The risk of intractable problems due to multi-vendor finger-pointing may also increase. In addition, there is a limited pool of software developers with network disaggregation expertise, especially for white box and independent NOS software.

Source: Analysys Mason, 2021

CSPs with considerable in-house resources and advanced skills may decide to own the disaggregation lifecycle processes, but they will need to make large upfront investments in automation, planning and organizational change, and must achieve cost benefits of scale in order to amortize these investments. CSPs can also work with professional services partners who can assume the responsibilities and risks of managing the white-box-based disaggregation lifecycle and vendor ecosystem, but this will have a considerable negative impact on TCO.

The DDC model and ecosystem are still nascent, and the associated additional opex and indirect capex may decrease as the value chain matures and CSPs gain more experience. It should be noted that the DDC approach provides a control and data plane separation that enables transport network functions and services to be managed in a more cloud-like and orchestrated manner; this benefit is not considered in the current TCO study. This is a radical change from the present mode of operations that brings certain advantages, but also limitations (as discussed in the next section). The potential operational benefits and costs of this approach should be examined in detail and compared to incumbent vendor control/data plane separation approaches in order to have a more complete opex assessment for the DDC scenarios.

3.3 There are several key operational considerations beyond TCO

There are several important operational factors (such as reliability, compatibility, openness and ease of technology upgrades) that are not quantified in our TCO model, but that CSPs should take into account when evaluating the IP routing architecture and platforms for their networks.

Network reliability. A distinct architectural feature of the DDC model is that the control plane, network services and applications typically run on x86 COTS servers in an out-of-band (OOB) network, rather than being tightly coupled with the data plane inside the network elements. With the right hardware abstraction and resource orchestration mechanisms, this promises software-based IP network operations that are similar to those of IT clouds over a common, shared pool of white-box hardware resources. However, the DDC model is still in development and there are some network reliability concerns that are associated with the OOB management and control network. The OOB network requires that the Ethernet management port is always operational, thereby creating a new point of failure in the network and bringing additional reliability and MTBF challenges. Another reliability challenge that applies to both the DDC and horizontal-only disaggregation (scale-out) approaches is related to the externalization of the internal fabric traces from the traditional modular chassis via a large number of connections. This can make it more challenging to identify and troubleshoot potential failures. CSPs will therefore need network automation and management solutions that are capable of mitigating these risks.

Openness. One of the main reasons for CSPs to consider moving away from the traditional modular chassis to disaggregated solutions is to reduce their reliance on a limited number of vendors and to create more-open network environments. Using scale-out architecture with standard IP/BGP connectivity instead of proprietary interfaces (as modeled in our TCO analysis) can help to address these issues. The DDC approach appears to offer a broader, open vendor ecosystem that can foster competition and innovation in the network value chain. However, there are some aspects of the current DDC design that may actually limit openness and operational flexibility. For example, DDC interfaces between leaf and spine elements use proprietary, cell-based interconnect technology in order to meet capacity and scalability requirements. This can limit spine and leaf choices to Broadcom-based platforms, whereas IP-based architecture support a more diverse range of leaf and spine options. In addition, DDC leaf and spine elements have a physical limitation of 100 meters to prevent delays. As such, CSPs would need to build a completely new fabric if they need to extend the fabric to another building.

New risks and responsibilities. It is critical for CSPs to choose an approach that allows them to futureproof their networks. CSPs have traditionally outsourced hardware innovation to their routing vendors that have extensive R&D experience and investments and assume responsibility for hardware upgrades. Silicon and system-level disaggregation can enable CSPs to increase their level of control over innovation, but they can also bring major responsibilities and risks. These include carrying out their own chipset upgrade cycles, ensuring backward/forward compatibility and creating smooth migration paths. An example of the kind of scenario that a CSP might face involves upgrading the fabric spine of the DDC architecture with a Broadcom Ramon 2 chipset. This chipset is not backwards-compatible with existing infrastructure and could potentially require significant downtime and effort to swap out.

3.4 There are several benefits of partnering with a vendor who can provide both traditional and disaggregated platforms

There are clear advantages that CSPs can gain from working with established vendors, particularly if they are concerned about opex and overall TCO and lack the skills to adopt the new IT automation paradigms associated with IP disaggregation models. Some of these advantages are as follows.

- Established vendors can offer CSPs both traditional, scale-up designs and newer, disaggregated scale-out models, thereby giving experience of the latter and supporting mixing and matching in different parts of the network.

- The pre-integration that these vendors can offer significantly improves the cost economics of disaggregated networking compared to DIY approaches, including simpler platform certification and homologation.

- CSPs already trust these vendors and do not need to risk working with a nascent and evolving vendor ecosystem that may have immature technology and financial instability.

- The broad portfolios, which include merchant-silicon-based platforms, mean that there is pricing flexibility, without the need for DIY integration.

Another key advantage of this approach is that it enables CSPs to standardize and streamline their software architecture with a single, market-proven, carrier-grade platform and NOS for operational consistency across multiple hardware and network domains. For example, CSPs may prefer to deploy different types of routing platforms in their IP edge networks to optimize network costs (similar to the IP core), but this may increase operational complexity if each platform uses different operating models, interfaces and hardware. Our analysis, shown in Figure 3.7, demonstrates that:

- multi-service systems can provide optimal TCO in networks with a large number of business connectivity and VPN services and high QoS requirements thanks to the scale, flexibility and programmability provided by these platforms
- CSPs with lower service QoS and VPN requirements may save costs by using lower-capex, lean platforms that are merchant-silicon-based and optimized for lower power.

Figure 3.7: TCO for multi-service systems and merchant-silicon-based lean edge platforms for different QoS and VPN demands

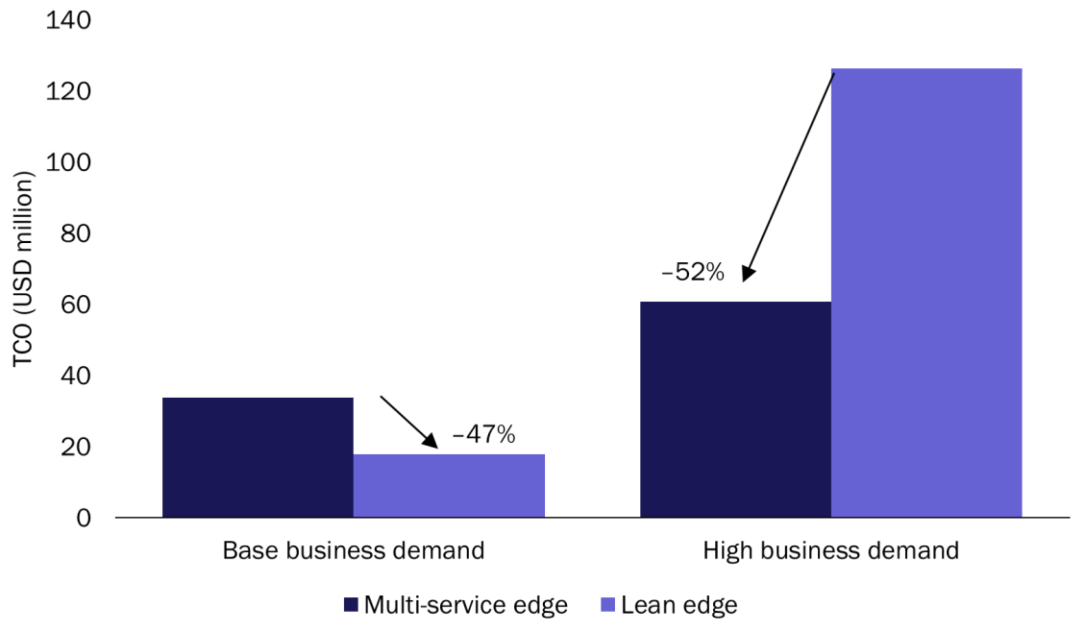

Source: Analysys Mason, 2021

Having a vendor partner that can provide a broad range of platform options with different costs, feature sets and scale optimization levels across the IP network domains can help CSPs to achieve operational simplicity and consistency with common software architecture.

4. Conclusions and recommendations

We provide the following recommendations for CSPs that are planning to embark on a network disaggregation journey, based on our research and our key TCO findings.

- **CSPs should be very clear about what they want to achieve with disaggregation and should identify which type of disaggregation, if any, is the most suitable for their goals.** CSPs' interest in network disaggregation has been growing, driven by a variety of motives such as cost reduction, the removal of vendor lock-in and service innovation. However, network disaggregation is not a panacea for all CSP network and operational challenges, and it may not always deliver cost savings as demonstrated in our TCO model. The optimal architectural choice for IP networks depends on specific domain and service requirements, as well as CSPs' commercial and operational objectives.

- **CSPs that are considering DIY disaggregation should take into account the costs, risks and operational implications.** The lifecycle of disaggregated networking hardware and software involves a series of activities that most CSPs are not familiar with. These activities can result in high opex and indirect capex when taking a DIY approach due to the upfront organizational investments in reaggregation processes and operational change. In addition, there are new operational risks and responsibilities that CSPs need to manage when using disaggregated architecture, as discussed in Section 3.3. Certain proprietary aspects of the DDC model may even limit CSPs' choices. It is therefore important for CSPs to look beyond the short-term capex savings and do a complete TCO assessment based on their operational readiness and in-house capabilities.

- **CSPs should consider partnering with an established vendor who can support multiple architectural options and help with operational change.** CSP IP networks consist of a mix of routing architecture, including both traditional and disaggregated options, across different network domains and services. Partnering with an established vendor who also has a balanced, open proposition can minimize the costs and difficulties associated with disaggregation, and can provide many benefits to CSPs as they continue to evolve their networks. It will be critical for CSPs to ensure that any vendor that they partner with has:
  - an open platform (for example, one that uses open interfaces and data models such as OpenConfig and NETCONG/YANG and complies with industry standards)
  - the ability and willingness to integrate with third-party hardware and software
  - real disaggregation deployment experience with hyperscalers and other CSPs
  - a comprehensive set of network automation capabilities and professional services support.

5. Appendix: TCO model details

This section provides details of the modeling assumptions and parameters of the TCO analysis.

5.1 Connections and traffic

Figure 5.1: Types and numbers of connections for the CSP modeled

Connection type	Year 1	Year 2	Year 3	Year 4	Year 5
Residential broadband	8 377 408	8 479 090	8 533 442	8 588 024	8 650 020
Business broadband	1 087 894	1 096 316	1 104 935	1 114 344	1 123 270
Business dedicated connections	54 185	54 734	55 086	55 241	55 216

Source: Analysys Mason, 2021

Figure 5.2: Average busy-hour internet throughput per broadband connection, by domain

Traffic (Mbit/s)	Year 1	Year 2	Year 3	Year 4	Year 5
MSAN/OLT to metro	1000	1000	1000	1000	1000
Metro to edge	200	200	200	200	200
Core	2.3	2.9	3.6	4.3	5.1

Source: Analysys Mason, 2021

Figure 5.3: Split of business dedicated line connections by capacity

Capacity	Year 1	Year 2	Year 3	Year 4	Year 5
Below 100 Mbit/s	19 165	17 720	16 080	14 310	12 377
100 Mbit/s and above, but below 1 Gbit/s	25 029	26 236	27 427	28 589	29 753
1 Gbit/s and above	9991	10 778	11 579	12 342	13 086

Source: Analysys Mason, 2021

5.2 Network design and architecture

Figure 5.4: Overview of the network design assumptions for the main TCO scenarios

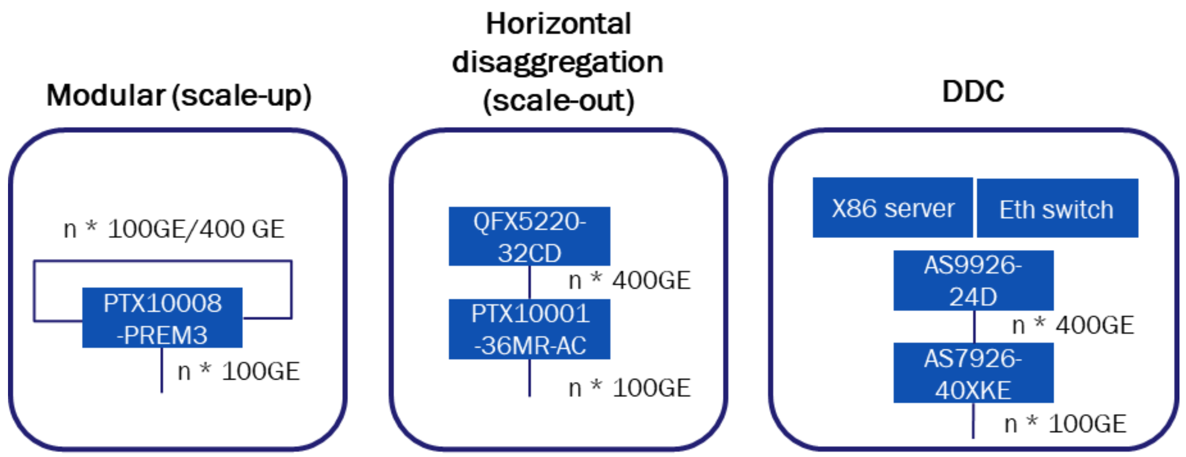

Source: Analysys Mason, 2021

Figure 5.5: Number of node sites in the core, edge and metro domains

Network domain	Number of nodes
Metro	2500
Edge	25
Core	5

Source: Analysys Mason, 2021

5.3 Cost parameters

- The costs for automation and provisioning software, training and OSS/BSS integration are excluded in all scenarios.
- The SI professional services fee is assumed to be 50% of the total software and hardware cost for the DDC models.

- We assumed additional FTEs for the following areas for the DDC models:
  - architecture design (5 years)
  - vendor management (5 years)
  - initial set-up/support (3 years)

  These labour costs are based on a market benchmark for the full salary of a qualified network engineer FTE.

Capex

Figure 5.6: Capex components for each scenario

Scenario	Capex components
Modular (scale -up)	• Core router chassis • Core router line card • PAYG software license
Pre-integrated scale-out	• Core leaf router • Core leaf router PAYG software license • Core spine router • Core spine router PAYG software license • Core cables
DDC	• Core leaf router chassis • Core leaf router software license (annual, flat) • Core spine router chassis • Core spine router software license (annual, flat) • Core cables • Ethernet switch • Ethernet switch non-PAYG software license • x86 server

Source: Analysys Mason, 2021

Opex

Figure 5.7: Summary of labor cost parameters

Parameter	Value
Hourly cost of labour (USD)	130
Cost trend of labour	2.0%
Number of working hours per year	1,725

Source: Analysys Mason, 2021

Figure 5.8: Summary of floorspace cost parameters

Parameter	Value
Rack units per standard rack	45
Rack floorspace (square meter per rack)	4
Annual floorspace cost (USD per square meter)	3000
Cost trend of floor space	3%
Annual floorspace cost (USD per rack)	12 000

Source: Analysys Mason, 2021

Figure 5.9: Summary of power cost parameters

Parameter	Value
Power (at a consumption of 48V per kWh) (USD)	0.10
Cost trend of power	3%
Air conditioning power consumption as a share of total power consumption	90%
Power outlet cost (USD)	250

Source: Analysys Mason, 2021

6. About the authors

Gorkem Yigit (Principal Analyst) is the lead analyst for the Cloud Infrastructure Strategies and the Edge and Media Platforms research programmes. His research focuses on the building blocks, architecture and adoption of the cloud-native, disaggregated and programmable digital infrastructure and networks that underpin the delivery of 5G, media and edge computing services. He also works with clients on a range of consulting projects such as market and competitive analysis, business case development and marketing support through thought leadership collateral. He holds a cum laude MSc degree in economics and management of innovation and technology from Bocconi University (Milan, Italy).

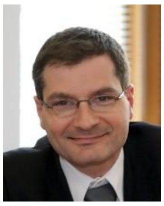

Gilles Monniaux (Principal) is a Principal in Analysys Mason's Cambridge office. He joined Analysys Mason in 2007 and specialises in network economic cost modelling, spectrum valuation and regulatory analysis. Gilles has significant cost modelling experience, having been involved in the cost modelling of fixed and mobile networks in Ukraine, Norway, France (including overseas departments), the UK, Belgium, French Polynesia, Brazil, Malta, Mexico, New Zealand and Australia. Gilles has presented techno-economic cost modelling issues in open industry and operator meetings during these projects.

Caroline Chappell (Research Director) heads Analysys Mason's Cloud research practice. Her research focuses on service provider adoption of cloud to deliver business services, support digital transformation and re-architect fixed and mobile networks for the 5G era. She is a leading exponent of the edge computing market and its impact on service provider network deployments and new revenue opportunities. She monitors public cloud provider strategies for the telecoms industry and investigates how key cloud platform services can enhance service provider value. Caroline is a leading authority on the application of cloud-native technologies to the network and helps telecoms customers to devise strategies that exploit the powerful capabilities of cloud while mitigating its disruptive effects.

This perspective was commissioned by Juniper Networks. Analysys Mason does not endorse any of the vendor's products or services.

Analysys Mason Limited. Registered in England and Wales with company number 05177472. Registered office: North West Wing Bush House, Aldwych, London, England, WC2B 4PJ.