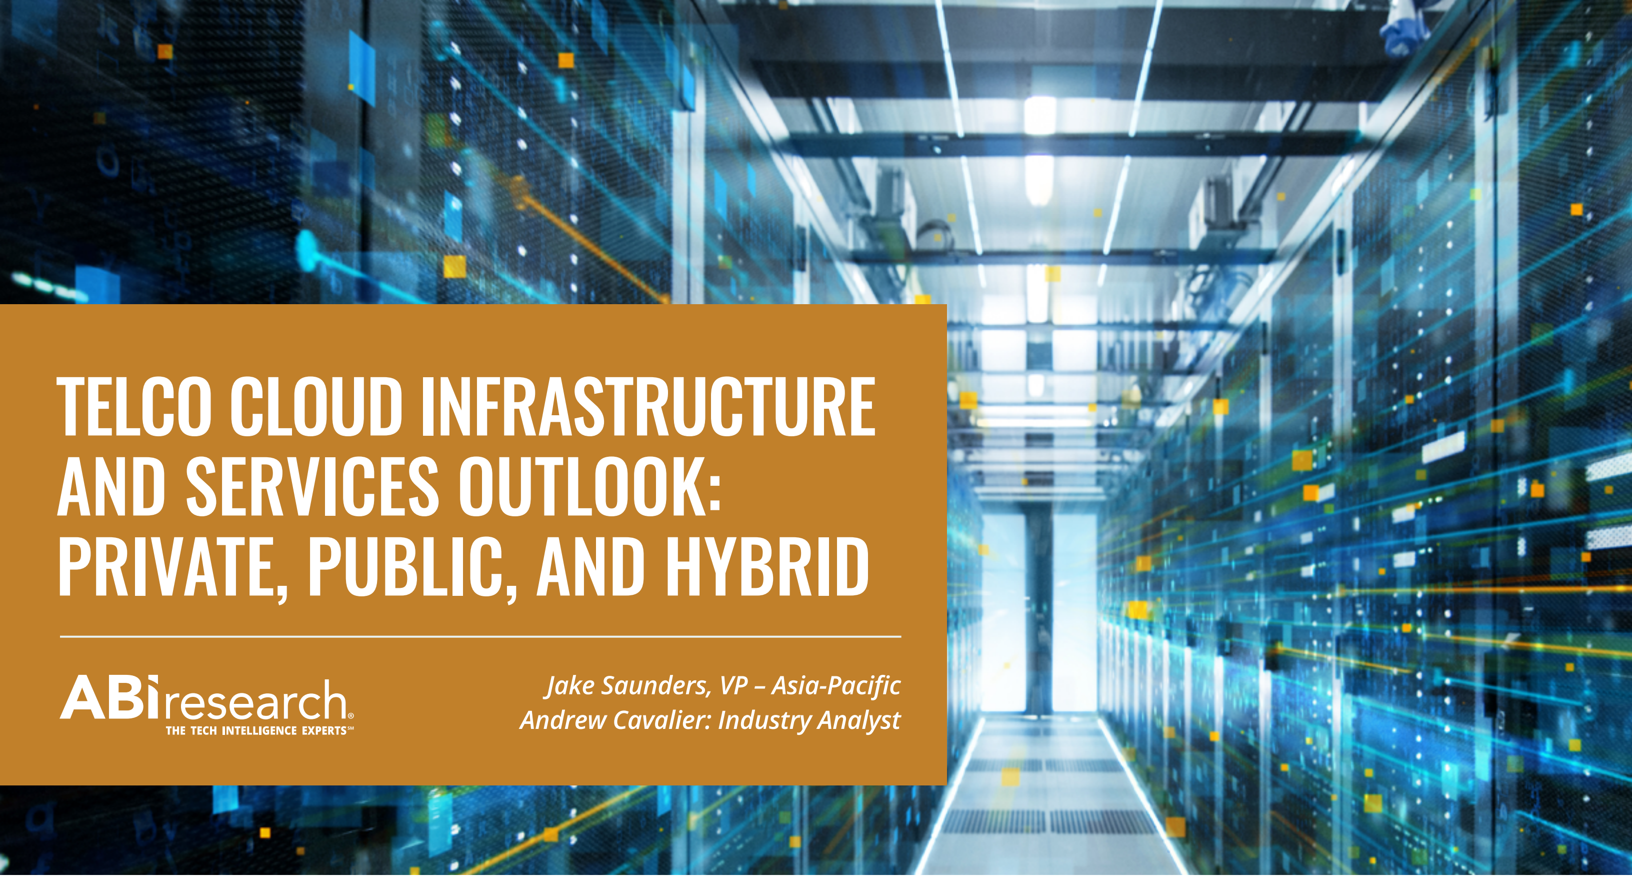

TELCO CLOUD INFRASTRUCTURE AND SERVICES OUTLOOK: PRIVATE, PUBLIC, AND HYBRID

ABIresearch
THE TECH INTELLIGENCE EXPERTS™

Jake Saunders, VP – Asia-Pacific
Andrew Cavalier: Industry Analyst

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Cloud computing has been one of the central technologies dramatically transforming the telecommunications industry. Communication Service Providers (CSPs) have worked together with Network Equipment Vendors (NEVs) to untether network applications and functionality from specialized proprietary hardware to Commercial-off-the-Shelf (COTS) commodity hardware. With NEVs integrating Virtual Network Functions (VNFs), Virtual Network Functions Managers (VNFMs), and Virtual Infrastructure Managers (VIMs) into the underlying hardware, CSP cloud architecture has traditionally been highly secure and well optimized with vertically-integrated products or components. While this telco cloud architecture has been tried and tested, the entrance of public cloud services geared for CSPs has changed the outlook and potential for architectural modularity and service flexibility.

With the introduction of public cloud services that are based on pay-as-you-go service models and modular lightweight applications called microservices, the options for CSP cloud architecture have expanded substantially. In this respect, CSPs can now leverage shared virtualized external storage, computing, and processing provided by a public cloud service provider at great speed and low costs. This has prompted some operators to reevaluate their business model and deliver their network services with a more flexible and modular infrastructure that taps into the public cloud. While some greenfield operators are showing what can be done with the public cloud, the option to move network workloads to the public cloud is not so clear-cut for brownfield operators.

Now, with the advent of 5G and the transition to a post-COVID-19 era, there has been a clear acceleration of the digitalization and cloudification of infrastructure and many CSPs now face a plethora of options when undergoing cloud transformation. Crucially, CSPs have critical aspects to consider, such as network security, control, and Total Cost of Ownership (TCO) when exploring the options for infrastructure cloudification. Alongside this, the size of the operator, cost/pricing structure, geography, regulations, network design, and deployment scenarios for each CSP has a bearing on the appropriateness of the cloud infrastructure model. In this regard, CSPs have an opportunity to leverage their inherent strengths provided by many years of infrastructure development, network optimization, and customer service to deliver a cloudified telecommunications network that best suits the needs of modern enterprises and consumers.

Based on the analysis of cloud architectures and TCO scenarios leveraging on-premises telco cloud or public cloud infrastructure, the telco cloud has shown many long-term cost and deployment advantages for the Data Center (DC), edge, and private networks. In this respect, achieving high link capacities and packet rates, strong network security and automation, and network control and optimization with the public cloud can often prove more costly than deploying on-premises. To this end, public cloud services in the telco space can be effective for reducing short-term costs and quickly expanding network capabilities, but are not a suitable long-term solution or replacement for most CSPs' cloud solutions.

OUTLOOK FOR CLOUD-BASED SERVICES

The core network is the backbone of a telecommunications network and provides the technology and infrastructure necessary for voice, text, video, and audio communications to anywhere in the world via the Internet, phone, airwaves, cables, and wires. This core network consists of high-capacity lines that manages aspects of network traffic and services for customers, such as subscriber authentication, user data management, Internet Protocol (IP) address allocation, device positioning, and all interactions between devices with the Internet and phone network. With the advent of 5G and standards set by The 3rd Generation Partnership Project (3GPP), the core network is set to evolve with new Network Functions (NFs) associated with the packet core and user data management domains.

EVOLUTION OF APPLICATIONS

Cloudification has been at the core of digital transformation for the telecommunications industry and is one of the key elements for transforming CSPs into Digital Service Providers (DSPs). The beginning of this transformation leveraged a virtualization layer (using hypervisors) for the abstraction of physical resources (at the DC) into virtual ones (on the network). This formed the basis for telco Network Function Virtualization (NFV), which are made up of Physical Network Functions (PNFs) that have been decoupled and deployed as software equivalents or VNFs that are packaged as Virtual Machines (VMs) running on x86 hardware. A number of CSPs have implemented NFV, but it now looks like technology and competitive forces are seeing that virtualization migrate to the cloud in the form of Cloud-native Network Functions (CNFs).

CLOUD-NATIVE INFRASTRUCTURE

With the advent of 5G and the evolution of NFVs to CNFs using containers, operational complexity has increased with the proliferation of VMs and centralized core environments for new services and use cases like Fixed Wireless Access (FWA), cloud gaming, Extended Reality (XR), and other vertical industries. To accommodate the increasing level of service and network complexity, telco cloud deployments have begun to embrace newer technologies like Linux containers (containers), microservices, and flexible and homogenous architectures leveraging public cloud services. Critical in this evolution has been the growing interest in going cloud-native, with NFs designed and deployed from the onset as software to run on cloud servers.

VNFS TO CNFS

A distinguishing feature of this approach is the transition from using VNFs deployed as VMs to CNFs deployed inside containers. While VNFs abstract hardware to run Operating System (OS) instances, containers leverage a container engine to abstract the OS and divide a single OS instance into isolated environments that can run applications. This allows for containerization, with software applications, functions, or microservices packaged with all the OS libraries and dependencies required to run code on the cloud. Despite this evolution, telco-grade networks were never designed to run on the cloud and NFV coexistence remains vital as operators adopt virtual Evolved Packet Core (vEPC) for 4G and, eventually, 5G.

SPLITTING THE MONOLITH

As a part of the transition to a cloud-native architecture, CSPs have started to leverage microservices to rapidly build out new applications and independently iterate app components through parallel development. Deployed inside containers, microservice-based apps have an ideal self-contained environment, enabling easy container orchestration and hardware utilization. The transition from the hardware-dependent, stateful qualities of what were effectively monolithic applications to cloudified, stateless cloud-native applications with microservices and containers, has enabled dynamic and flexible network infrastructure and processes that can be deployed in any cloud. CSPs can deploy these technologies on existing on-premises infrastructure (private cloud), on public cloud resources leased from cloud service providers, or leveraging on-premises infrastructure supported with resources from cloud service providers (hybrid cloud).

RISE OF PUBLIC CLOUD PLATFORMS

While there is speculation about the role the public cloud will play for telcos, with the likes of DISH deploying their RAN and mobile core on AWS Cloud, Deutsche Telekom piloting 5G standalone with Google Cloud, and AT&T deploying their 5G core on Microsoft Azure, there is growing debate and interest in telco's migrating some or even all of their CNFs and applications to the public cloud. Among these 3 cases, only DISH has been commercialized and has yet to reach carrier-class stability and large-scale commercialization. In this respect, "large-scale commercialization", is lacking a clear threshold to classify deployments, for example, if millions of subscribers run stably for one year with the network. Moving forward, ABI Research believes that VNFs, MANO, and CNFs will become the primary infrastructure considerations influencing the growth and adoption of the public cloud for CSPs. Alongside these factors, addressing the Total Cost of Ownership (TCO) associated with priming infrastructure for 5G and transitioning to digital service-based business models will drive adoption. While not the highest priority, businesses and enterprises have been considering migrating workloads to the public cloud and CSPs have to carry out a careful assessment before migrating on-premises workloads, including OSS/BSS, to the public cloud.

AUTOMATION

With the growing prevalence of 5G, CSP network architectures need to be flexible and highly automatable to accommodate the many new services that require low-latency and high-bandwidth capabilities. In this respect, networks are becoming more complex and increasing the use of automation to manage network resources. Network automation and virtualization have been integral for Software-Defined Networking (SDN) and NFV technologies, and are effective in reducing costs and speeding up the delivery of network-based services. To this end, automation in the cloud has become attractive for helping manage resources across all parts of a network lifecycle.

CHALLENGES

TCO: There are TCO challenges tied to hosting large NFs on the public cloud. The size of NFs is typically larger than enterprise applications and has more stringent performance requirements which can 1) cause issues related to migration between different cloud platforms, and 2) inflate costs from hosting them on the public cloud for communications services. In this respect, new NFs and customer-customized requirements can drive up costs to meet service-level performance requirements.

Reliability & Performance: Alongside higher costs associated with running larger telco workloads continuously on the cloud, there are stringent Service Level Agreements (SLAs) and service uptime agreements for network services that telcos must consider. In this respect, the telco cloud guarantees quick and effective network fault recovery that can potentially reach a reliability of 99.999%, while public cloud vendors often can only ensure a much lower level of reliability that can be as low as 99.9%. Telcos have to spend more to increase reliability.

O&M: With public cloud Operations and Management (O&M) platforms CSPs also face challenges with a divided O&M, as telcos must now disaggregate network control and interface with public cloud vendors. Not only does this create challenges in cross-layer invoking and demarcation, but network monitoring time can increase from milliseconds to seconds.

Rapid Iteration: CSP networks require fast and continuous iterative upgrades to get new functions based on newer technologies to market. Separate operational environments and disaggregated O&M architecture can make the network more complex, slow the iteration process down, and impact network maintenance and automation.

Data Sovereignty & Security: Moving IT and core workloads to the public cloud draw security and sovereignty concerns for CSPs. In this respect, once these workloads are on the public cloud, certain region-specific data containment measures may not be possible.

REQUIREMENTS

Service: Support of communications services, especially voice and data calls, are key services that have stringent performance, reliability, and security parameters that require telco-specific engineering arrangements.

Application Type: Telecom services are forwarding-intensive services and require large-scale traffic forwarding on the user plane. The public cloud has no cost advantage in heavy traffic forwarding. The telco cloud can meet forwarding performance requirements in open mode and dedicated software and hardware consume less resources, enhance forwarding performance, and use the user plane to forward large-scale traffic. Virtualization for telcos is more demanding than those based on the public cloud, as user-plane applications have stringent high-performance requirements (for latency, jitter, and packet loss) and requires dedicated hardware or software. By trying to meet these performance and resource needs, total decoupling may be challenging.

Operability: CSPs often have to manage network elements and software from several vendors and require a platform to drive synergy and operational excellence between operational silos, equipment, and software.

Platform Observability: New service deployments require monitoring, logging, and tracing capabilities, with component standardization and ecosystem interoperability aiding in the End-to-End (E2E) observability of the platform.

Deployment Environment: CSPs continue to invest in DC infrastructure and require cloud environments that support DC IT workloads and edge use cases where CSPs have a physical presence.

Business Environment: With the introduction of hyperscalers (*e.g.*, AWS, Google Cloud Platform (GCP), and Microsoft Azure) as major stakeholders in telco-specific ecosystems, there is a requirement for System Integrators (SIs) with human capital and expertise to span both IT and communications.

Human Capital: The workforce will need to become adept at working and innovating on the cloud supported by governance that gets the best out of the ecosystem.

TCO CONSIDERATIONS

Reducing the Capital Expenditure (CAPEX) associated with building out and running DCs is one of the advantages of the public cloud. In this respect, unlocking greater service flexibility by running on-premises IT workloads on the public cloud has enabled operators to focus on networking and service diversification. Inherent features of the public cloud, such as inexpensive scale-on-demand and pay-per-use capabilities, automation of the provisioning and maintenance of services, enhanced application deployment versatility, and unlocking an increased pace of innovation, have been critical. Furthermore, all these features that come with deployment on the public cloud have little to no risk to the CSP.

DATA CENTER

CAPEX and Operational Expenditure (OPEX) considerations for telco DCs (high-performance NFs and CNFs, region build-out costs, Operations, Administration, and Maintenance (OAM) costs, *etc.*) are impacted by migration onto the public cloud.

	Telco Cloud	Public Cloud
Strengths	■ Cost-effective for telco-specific NFs, CNFs, and core network ■ Unified cross-layer O&M can save OAM costs	■ Additional compute resources do not require build-out of DCs (if covered by cloud region/zone) ■ Pay-as-you-go model
Weaknesses	■ High CAPEX and OPEX for additional DC build-out	■ Divided O&M can increase OAM costs ■ Carrier asset loss caused by reduction of Public Cloud CAPEX
Opportunities	■ Deploying at the edge for high-performance compute ■ Leveraging existing DCs can save on construction costs for edge compute	■ Deployment of IT workloads with low-cost microservices ■ Architectural modularity can decrease CAPEX and OPEX
Threats	■ Cyberattacks and infrastructure attacks can increase costs and shift telco operational priorities ■ Vendor lock-in (business strategy induced)	■ Vendor lock-in (from platform integration challenges)

The public cloud does offer many benefits with a flexible microservice-based architecture and pay-as-you-go service model that can reduce CAPEX for DC build-out, provided the CSP is in a covered public cloud region/zone. By moving all workloads to the public cloud, CSPs will effectively be switching to an OPEX-based model, as operations and service delivery are tied to resources managed by a cloud vendor. While certain applications like telco IT, OSS/BSS, and management plane functions can be migrated to the public cloud, other NFs (like user/data plane and control plane) are more cost-effective, secure, and manageable on-premises. To this effect, it is unlikely that all telco workloads can be migrated to the public cloud, especially in the context of "brownfield" operators.

Chart 1: Data Center Yearly Accumulated TCO: Telco Cloud versus Public Cloud

(Source: ABI Research)

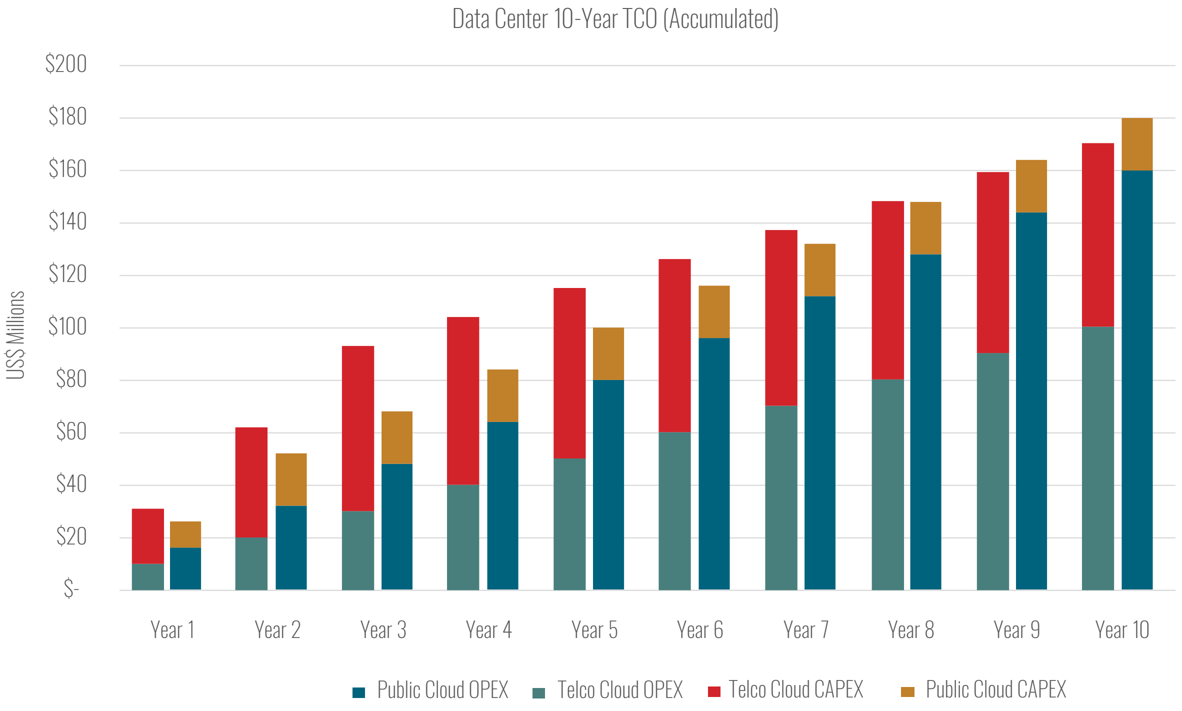

The 10-year TCO to build, deliver, install, and run a DC (US$169 million) is lower in the long run than leveraging mostly hyperscaler infrastructure and resources for operations (US$180 million). For a DC, annual OPEX is roughly US$10 million a year and maintenance CAPEX is roughly US$2 million to US$3 million a year. Initial CAPEX for construction is roughly US$62 million for an average telco capacity building at 10 Megawatts (MW) and 10,590 Square Meters (m^2). Assuming the CSP leveraged an hourly on-demand or monthly subscription model, OPEX is approximated to US$16 million a year to compensate for the larger-size and higher-performance workloads that CSPs require 24 hours 7 days a week. Based on these parameters, a 10-year TCO analysis reveals that leveraging hyperscaler equipment and deploying high-performance packet processing stacks to their cloud will outstrip the costs of running workloads purely on-premises. This scenario demonstrates the potential negative long-term impact that running a higher volume of NFs and high-performance workloads (*i.e.*, high link capacities at 100 Gigabits per Second (Gbit/s) and packet rates up to 150 Mega Packets per Second (MPPS)) continuously on the public cloud can have on TCO. Under the conditions of using hourly on-demand charges or monthly rental charges, deploying telco-grade workloads on the public cloud can become expensive and outstrip the OPEX of running the DC in the long haul.

EDGE CLOUD

Distribution of processing capabilities at the edge presents many opportunities for revenue generation and reducing expenses, with robust infrastructure being a distinct advantage that CSPs can leverage for an edge cloud strategy.

	Telco Cloud	Public Cloud
Strengths	▪ Mature existing infrastructure ▪ Reduces data processing OPEX and can free up resources ▪ Data sovereignty & security	▪ Reduce OPEX spending tied to data processing ▪ Reduce CAPEX spending tied to DC/server build-out ▪ High enterprise interest
Weaknesses	▪ Overallocation can reduce available network resources ▪ Limited developer/ enterprise interest	▪ Data processing costs for telco-specific workloads (high OPEX) ▪ Limited available infrastructure
Opportunities	▪ Leveraging existing DCs and RAN for edge co-location, IaaS, and SI revenue streams ▪ Reducing backhaul to lower interconnect and network costs	▪ Cutting-edge applications and services can provide new revenue streams ▪ Hyperscalers need infrastructure for edge: infrastructure revenue opportunities
Threats	▪ Increasing land acquisition costs ▪ Increasing energy costs ▪ Government regulation ▪ Distributed Denial of Service (DDoS) attacks	▪ Cybersecurity (sniffing attacks, ransomware, etc.) ▪ Internet of Things (IoT) device security ▪ Data sovereignty

The public cloud can provide several cost-benefits tied to data processing and building out additional DCs. In addition, hyperscalers need to invest significantly in real estate, deploying infrastructure, and managing distributed edge servers to distribute their processing capabilities—something telcos are well-positioned to address. Many telcos already have the significant infrastructure, real estate, and edge locations, so deployment costs can be lower than the above-referenced scenario. Alongside this, the potential for extreme distribution of processing capabilities will require many nodes, something public cloud NFs could support.

Chart 2: Edge Deployment Yearly Accumulated TCO: Telco Cloud versus Public Cloud

(Source: ABI Research)

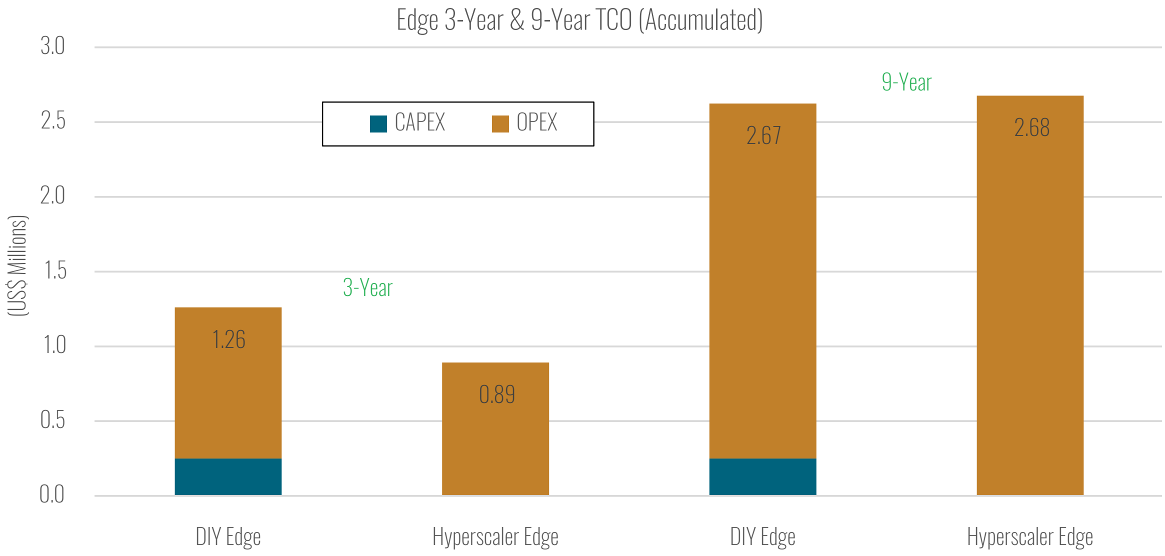

Given a memory-optimized large unit with local storage for high-performance compute (rack) that is in North America, the TCO for hyperscaler edge deployments cost US$892,000 for a 3-year term. Assuming all expenses remain consistent, the cost to run a hyperscaler edge will outstrip a Do-It-Yourself (DIY) edge in the long run (by year 9, in this example). The total cost to build, deliver, and install a micro-edge DC (DIY edge) has been estimated at between US$250,000 and US$1 million+ (depending on build-out, location, *etc.*), while OPEX to run the center is estimated at between US$220,000 and US$380,000 a year, with costs going down each year as processes and efficiency improve. To this end, DIY edge proves to have a lower 9-year TCO of US$2.67 million than using hyperscaler edge solutions with a 9-year TCO of US$2.68 million.

CAMPUSES AND INDUSTRIAL SITES

Campuses and industrial sites have network latency, security, and bandwidth requirements that are well-suited for on-site telco cloud infrastructure. Key features of the network like Enhanced Mobile Broadband (eMBB), Ultra-Reliable Low Latency Communications (URLLC), and Massive Machine-Type Communications (mMTC) can support use cases like enterprise IoT, automation, XR, and network slicing.

	Telco Cloud	Public Cloud
Strengths	■ Mature existing infrastructure ■ Ability to support larger telco-grade and industrial workloads ■ Manage network security and resources	■ Simplified operating model (if all on-premises DC operations are replaced) ■ Good for specific DC workloads ■ Absorbs CAPEX of server build-out
Weaknesses	■ Overallocation can reduce available network resources ■ Deployment costs for micro DCs and edge infrastructure (high CAPEX)	■ Certain NFs may require vendor configs ■ Restricted by geography—edge site locals ■ Low data sovereignty and network security
Opportunities	■ Enterprise/Industrial IoT (IIoT) services for private network deployments	■ 5G edge use cases and revenue streams (slicing, XR, IIoT, automation)
Threats	■ Government regulations, increasing energy costs, semi-conductor shortages	■ Vendor lock-in ■ Cybersecurity (sniffing attacks, ransomware, etc.) ■ Government regulations

While public cloud edge sites could theoretically provide a more simplified operating model, industrial sites need to leverage on-premises infrastructure for high-reliability and low-latency use cases (like networking slicing, automation, machine vision, and XR). Therefore, migrating all DC workloads to public cloud infrastructure is unlikely given the network requirements for industrial settings and the value proposition to move everything to the public cloud. Realistically, some industrial sites may opt for public cloud solutions, while other enterprises opt for on-premises solutions or indeed opt for a network slicing strategy.

Chart 3: Factory Private Network Yearly Accumulated TCO: Telco Cloud versus Public Cloud

(Source: ABI Research)

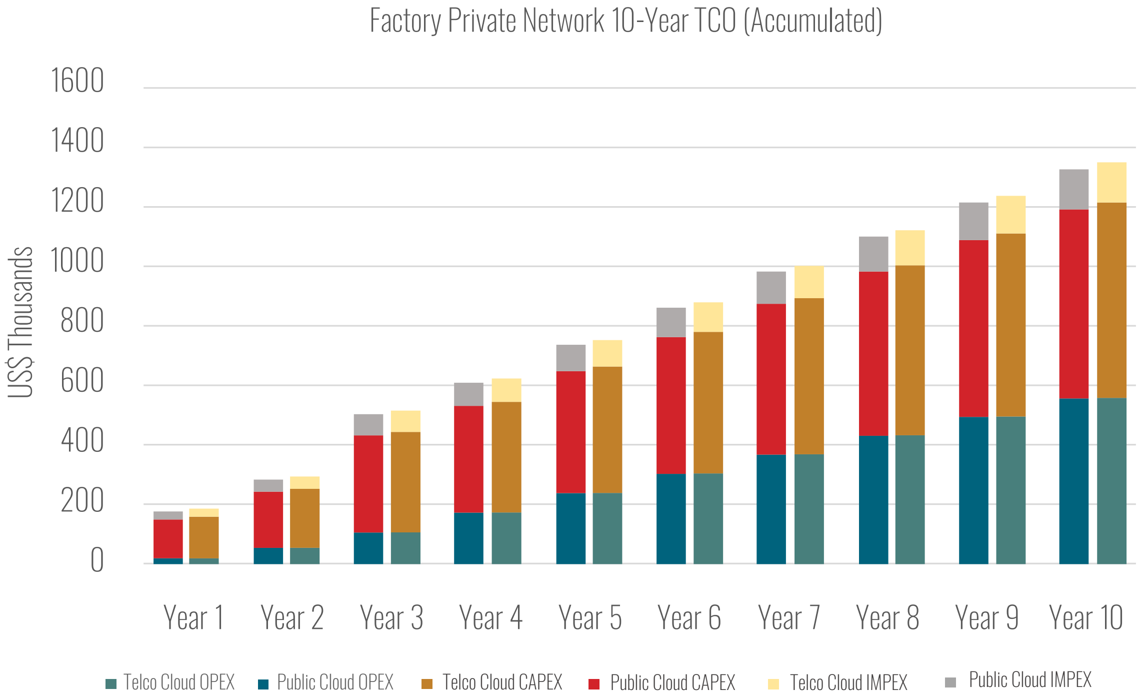

Two TCO scenarios (pure on-preises or hybrid cloud) were analyzed for private 5G network deployments in a 25,000 m^2 factory. At US$10,500 per edge server and US$6,000 for each initial set of access points, with an access point for every 450 m^2, the 10-year TCO for a fully isolated 5G network at the factory using on-premises equipment totaled US$1.34 million. In comparison, the hybrid approach, which offloads the 5G control plane, sees a slightly lower 10-year TCO of US$1.32 million for network deployment at the same factory. While this analysis demonstrates the comparative cost advantages of offloading certain NFs to the public cloud, this does not account for other factors, such as reduced security and control of the network.

CLOUD-BASED SERVICE REVENUE FORECASTS

Revenue for cloud services is broken down into three categories: public cloud (*i.e.*, revenue earned from the public cloud), telco cloud (*i.e.*, revenue earned from CSP cloud spending), and hybrid cloud (*i.e.*, revenue generated from rendering services for CSPs on the public cloud).

PUBLIC CLOUD

Driven by software services, ABI Research forecasts that global public cloud revenue (consisting of Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), Software-as-a-Service (SaaS), and Network-as-a-Service (NaaS)) will grow from US$324 billion in 2020 to over US$778 billion in 2025, at a Compound Annual Growth Rate (CAGR) of 19.1%.

Chart 4: Public Cloud Revenue Forecast (IaaS, PaaS, SaaS, NaaS)
World Markets: 2020 to 2025

(Source: ABI Research)

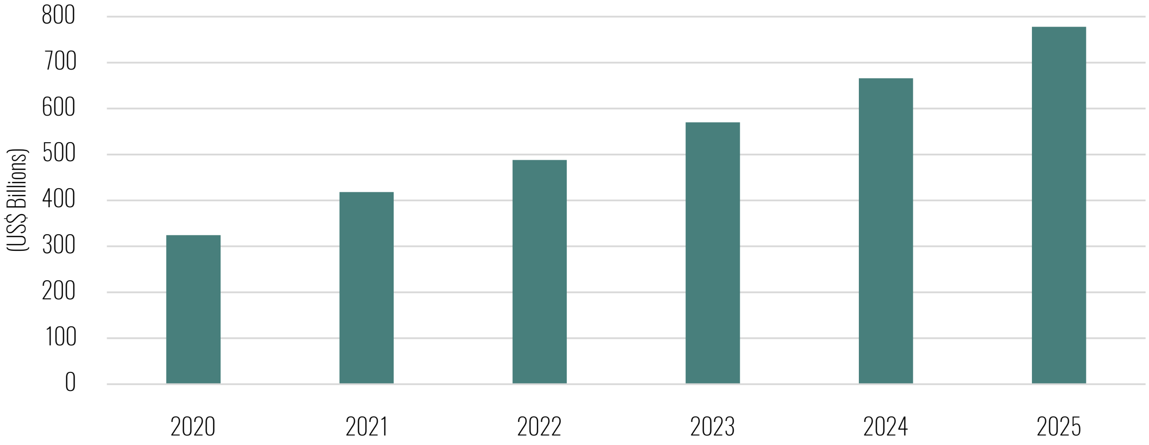

Global adoption of shared public cloud services has been growing at a rapid pace because of enterprises' interest in modernizing operations, expanding product and service portfolios, and providing more value for customers. In this respect, the pandemic, geopolitical, and economic pressures have created a need to adapt and make these changes at speed and for lower costs. Public cloud services are projected to grow at an aggressive pace as enterprises digitalize operations and reduce spending on capital.

TELCO CLOUD

ABI Research forecasts that global telco cloud revenue (generated from infrastructure) will grow from US$8.7 billion in 2020 to US$29.3 billion by 2025, at a CAGR of 27%.

Chart 5: Global Telco Cloud Equipment & Software Spending by Network Function Forecast
World Markets: 2020 to 2025

(Source: ABI Research)

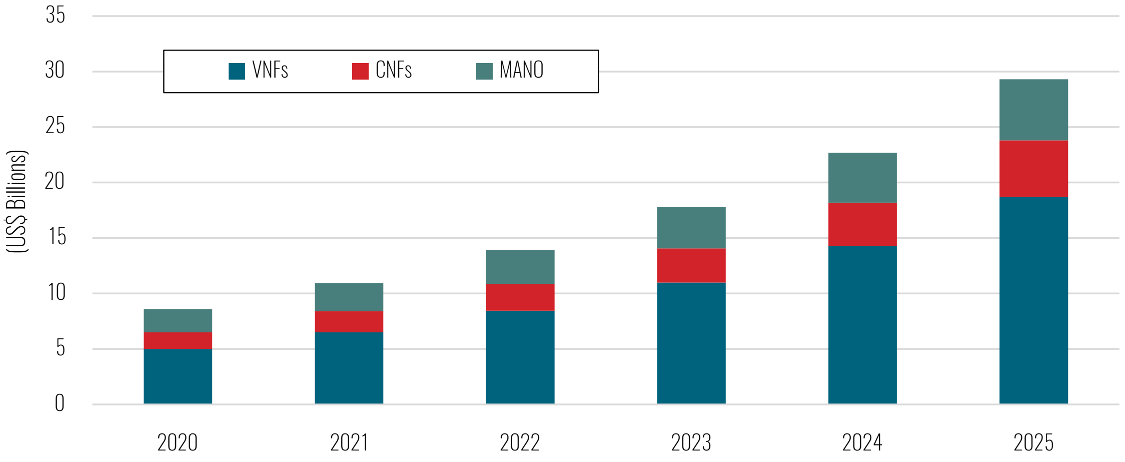

Global telco cloud revenue refers to the CSP spending on network virtualization and cloud infrastructure. Many CSPs are preparing for the move to 5G and are using their investments in virtualization and physical infrastructure to build out their own cloud platforms. At the same time, NEVs are providing many solutions for CSPs to deploy their own telco cloud at low costs and are driving a strong value proposition for retaining NF deployments on a private cloud. As a result, many CSPs are showing a preference for on-premises deployments (especially with telco-grade workloads) and will continue to drive revenue.

HYBRID CLOUD

ABI Research forecasts that global hybrid cloud revenue (generated from telco spending on public cloud equipment and services) will grow from an estimated US$5.2 billion in 2021 to over US$8.5 billion in 2025, at a CAGR of 13.06%.

Chart 6: Hybrid Cloud Equipment and Software Spending Forecast
World Markets: 2021 to 2025

(Source: ABI Research)

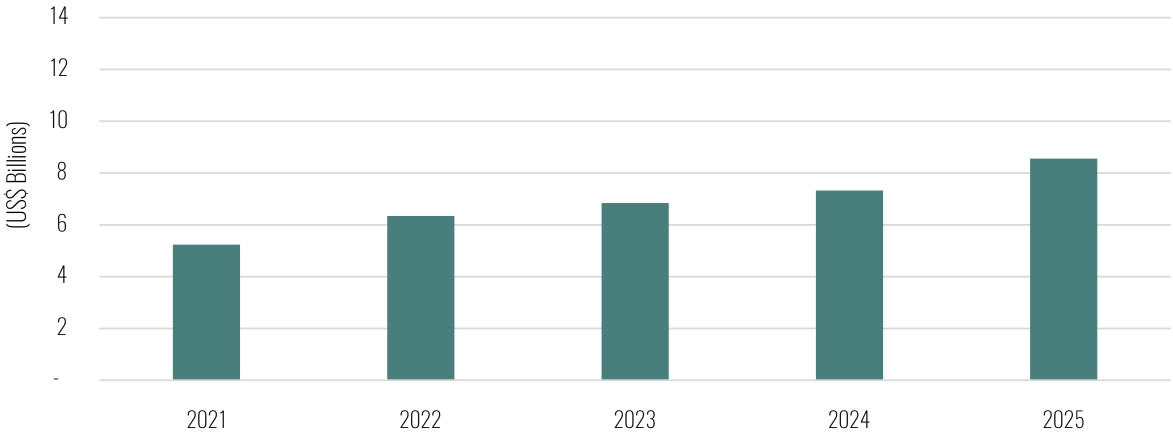

Telcos have started to partner with public cloud providers in varying degrees, with some working with a single or multiple providers for infrastructure, storage, data analytics, and more. This forecast estimates the proportion of public cloud revenue generated from CSP deployments. ABI Research believes that while some CSPs are leveraging public cloud services, many are showing a preference for investing in their own private cloud. In this respect, telco spending on public cloud services is expected to increase slowly as CSPs accelerate the use of their own infrastructure for cloud deployments and services.

SUMMARY AND CONCLUSIONS

While there is some debate about the level of public cloud's involvement in CSP's cloud strategy moving forward, a closer relationship with public cloud vendors is likely to develop in certain markets as well as depending on the CSP's level of cloud-native expertise, government oversight and access to public cloud solutions. Several greenfield operators and industry disruptors (e.g., DISH) have started to use the public cloud and demonstrate that CSPs are evaluating the public cloud model. While these early adopters and use cases do exist, CSPs still have many critical operational, performance, and data aspects that are not fully addressed with the public cloud.

CUSTOMIZED CLOUD STRATEGY

There are many TCO considerations, based on deployment, and challenges tied to network reliability, performance, data sovereignty, and security that should give CSPs pause and urge careful evaluation. Operators stand to gain much from leveraging on-premises telco cloud alongside the public cloud for innovation and agility. Despite this opportunity, not all operators have the same market position, geographical location, and growth strategy, and will not see the same benefits. Therefore, CSPs should consider their unique position and the benefits the public cloud could provide to network operations.

ROLE OF TELCOS IN THE CLOUD SERVICES EXPERIENCE

APPLICATIONS SUITED TO A TELCO CLOUD

The telco cloud is effective for VNFs/CNFs that handle user plane/data plane and control plane functions (directory services, routers, firewalls, traffic control, load balancers).

CONTAINERS

Containers have been at the core of transforming CSP cloud workloads and architecture, as they can be leveraged to help with the migration of certain NFs to cloud-native environments. Containers are particularly attractive as they provide unique benefits with lower Central Processing Unit (CPU) and memory overhead than VMs, can be instantiated and deallocated quickly, and move applications between environments, while retaining full functionality. The advent of 5G has been an accelerator for container adoption, as the ability to deliver lightweight scalable NFs is a foundational element for 5G network capabilities and on-demand network slicing services. To this end, containers are highly suited for the telco domain, especially when it comes to microservices and 5G edge computing, as they prioritize low latency, resiliency, and portability. Depending on the use cases and applications, CSPs can opt for deployment on Bare-metal containers or VM containers and smoothen the transition to cloud-native architectures.

KUBERNETES

Kubernetes (or K8s) is an open-source platform that automates the deployment, scaling, and management of containerized workloads. In this respect, Kubernetes provides a framework to run distributed systems resiliently and abstracts away some of the underlying internal networking complexities of applications. Kubernetes consists of clusters that contain a control plane and one or more compute nodes that can be run over virtualized infrastructure or bare-metal infrastructure.

VM APPROACH

Containers deployed on top of VMs in a CNF architecture, where the hypervisor overhead is present, can benefit from a separate lifecycle and start, recover, and upgrade services substantially faster than on bare metal. To this end, VM-based containers can enable a quicker route to cloud-native deployments (relative to bare metal), as they do not suffer from the portability limitations of CNFs across different environments with a different OS.

BARE-METAL APPROACH

Containers running on bare-metal servers benefit from no compute overhead or hypervisor and can be deployed at a higher density than on VMs. Since containers benefit from a lightweight footprint, the speed of instantiating or recovering services is optimized and the reduced overhead allows for more containers to run on the server. To this end, containers have better performance and higher efficiency when running on bare-metal infrastructure.

MICROSERVICES

Applications built and deployed using cloud-native methodologies emphasize microservices—a software development approach that breaks apart an application into highly-cohesive and loosely-coupled services. Antipode to coarse-grained models, such as a Service-Oriented Architecture (SOA), which uses large applications and "imperative/prescriptive" programming languages, a microservice is a "declarative" service-based architectural design pattern. Between these ends of the spectrum is a Service-Based Architecture (SBA), a hybrid model that leverages microservices architecture, while increasing service granularity in the core network.

Chart 7: Network Design Architecture (Microservices, SBA, and SOA)

(Source: ABI Research)

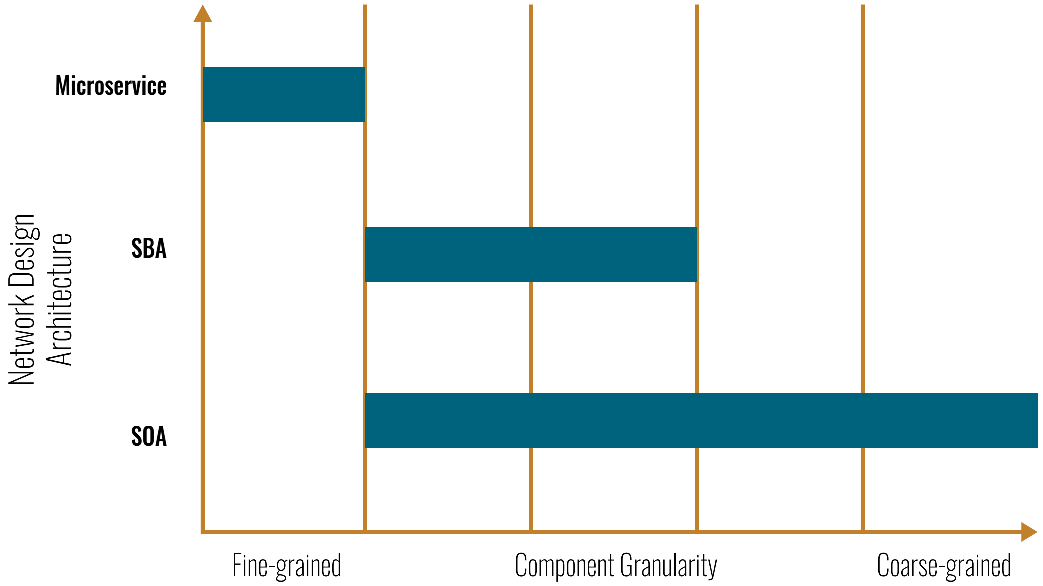

Choosing between monolithic or microservice applications is a preference, with both being valid under the correct circumstances. Monolithic architectures often can create much simpler developer workflows, such as monitoring and troubleshooting, while activities like E2E testing can also be greatly simplified. However, as the complexity of the application grows, it limits the ability of developers and architects to fully understand and iterate upon it. Microservices reduce complexity by having independent development, high modularity, and the advantage of being technology agnostic. In this respect, so long as services can communicate with each other via a network, CSPs can mix and match technology stacks to best suit their needs. It is likely that CSPs will have to strike a balance between fine-grained and course-grained architecture via automation and human interaction.

SERVICE MESHES

Service meshes are a configurable infrastructure layer for microservice applications that facilitates service-to-service communications and automation between microservices. As deployments of distributed services grow in size and complexity, service meshes can help understand and manage requirements, such as discovery, load balancing, and failure recovery, for more complex operational requirements like A/B testing, canary deployments, encryption, and E2E authentication. Without the service mesh, the network does not understand the traffic or make decisions based on the traffic being sent over.

Istio is an open-source service mesh that layers transparently onto existing microservices and provides a uniform and efficient means to secure, connect, and monitor services. With the data plane used for service-to-service communications, Istio's control plane takes the desired configuration and dynamically programs proxy servers to update as rules or the environment change.

Chart 8: Service Mesh: Istio Deployment Model

(Source: Istio)

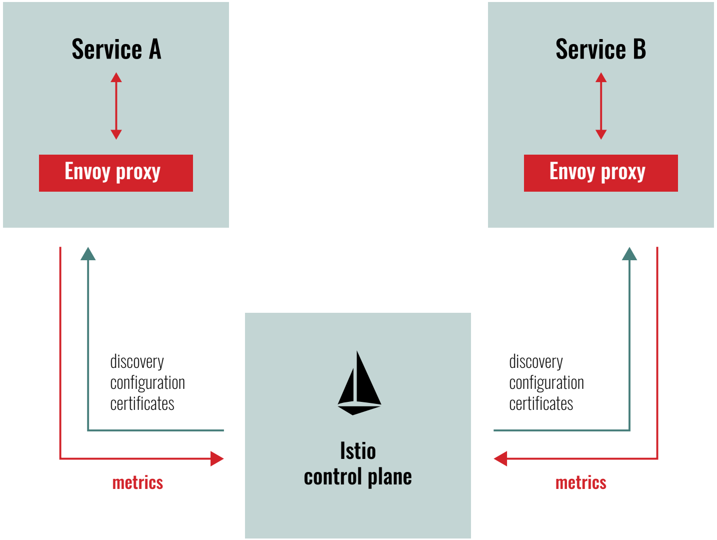

CSPs leveraging service meshes, such as Istio, can benefit from consistent service networking without adding developer overhead, swiftly integrate service-to-service security, and improve application performance via performance insights.

EXPECTATIONS AND REQUIREMENTS FOR THE TELCO CLOUD

BENEFITS OF THE TELCO CLOUD

Telco-Grade Applications: There are TCO benefits from hosting telco-grade NFs with existing on-premises cloud infrastructure. Due to the nature of traffic traversing throughout CSP networks, many NFs are overall better suited for telco cloud, like those on the user/ data plane. Since these NFs are typically larger than enterprise applications, the costs for migration, and public cloud vendors' compute and storage costs to run these functions can demonstrate a higher TCO profile than using an on-premises telco cloud. To this effect, telco cloud is cost-effective for applications with stringent requirements like low latency, large amounts of storage, and specialized compute resources, such as VNFs, RAN, and edge deployments.

Telco-Grade Reliability and Performance: CSPs provide a reliable service, encapsulated in SLAs with performance guarantees such as latency and bandwidth requirements alongside 99.999% for service layer disaster recovery features. This remains a distinct advantage of telco cloud architecture stacks as they are often optimized for performance.

Unified O&M: Streamlined and unified O&M is another benefit of the telco cloud, as vertically integrated stacks benefit from a single or lower number of vendors, reducing the complexity and interworking required with existing solutions and vendors. Depending on the deployment, divided O&M can also inflate costs, as coordinating between organizations can be time-consuming and create synergy issues with larger telco workloads.

Rapid Iteration and Automation: As network complexity increases with NFs moving to VNFs/CNFs in a cloud environment, automation and rapid iteration of services become a necessity for CSPs. With the telco cloud, CSPs often have fewer vendors and can rely on fine-tuned VNFs, hypervisors, and NFV Infrastructure (NFVI) layers that can streamline E2E automation and service iterations.

High Network Security: With a lower number of vendors and network resources managed by the CSP, the private telco cloud reduces access points for cybersecurity threats and provides an ideal environment for organizational security policies. In this respect, the public cloud is predicated on a shared network with multiple organizations, which increases the risk of data leakage and inhibits CSPs' meeting strict regulatory compliance standards. Telco Cloud Architecture Considerations

SINGLE-VENDOR CLOUD AND INTERDEPENDENT ARCHITECTURES

Interdependent architectures consist of vertically-integrated products or components that cannot be created independently of another part. NEVs like Ericsson, Huawei, and Nokia excel at providing these vertically integrated stacks and normally upgrade products and do not require customers to change operations. Naturally, these architectures support a single-vendor cloud strategy with one vendor integrating its VNFs, VNFM, and VIM with the underlying hardware. Due to the VNFs, hypervisor environment, and NFVI layer being fine-tuned by the supplier, this can reduce complexity and cost of operations, while augmenting management of multiple standards and products at the VIM and NFVI.

MULTI-VENDOR CLOUD AND MODULAR ARCHITECTURES

Modular architectures leverage a diverse set of components (VNFs, CNFs, etc.) provided by different vendors that need to be chained together and managed as one coherent platform. While modular components do make some compromises in performance, they do offer benefits with enhanced convenience and flexibility. The pool of suppliers expands beyond NEVs to include pure-play software vendors (like Red Hat, Affirmed Networks, and Mavenir) and hyperscalers. This diversity in suppliers mandates a high level of integration for individual VNFs/CNFs and can create compliance challenges as well as increased OPEX due to fragmentation and complexity in the cloud stack. Despite this, a multi-vendor approach can have lower upfront costs for a cloud platform and could be attractive for new horizontal B2B value creation.

TOWARD CLOUD-NATIVE

Established “brownfield” telco operators seeking a cloud-native architecture need to consider the interoperability of legacy network environments with new cloud-native ones. Aspects of this include VNF migration, management and automation practices, cloud mix (purely on-premises, purely public cloud, or a mix of both), and ecosystem partners. Therefore, the infrastructure required to run cloud-native applications is different, as responsibilities that infrastructure used to handle have started to move into applications. This is key when migrating network architecture to the cloud, as a holistic approach is needed.

INTEREST IN THE HYBRID CLOUD

The hybrid cloud is an architectural approach that combines on-premises infrastructure with the public cloud and private telco cloud through common data management for a single, unified computing environment. A distinctive feature of this architectural approach is a flexible distributed computing environment in which workloads can run and scale on the most appropriate computing model—effectively future-proofing CSP needs. The ability to host certain IT and telco workloads on a public cloud vendor’s resource can help save CAPEX and enhance scale-out capabilities during high-traffic periods with pay-as-you-go services.

Figure 1: Hybrid Cloud Model

(Source: ABI Research)

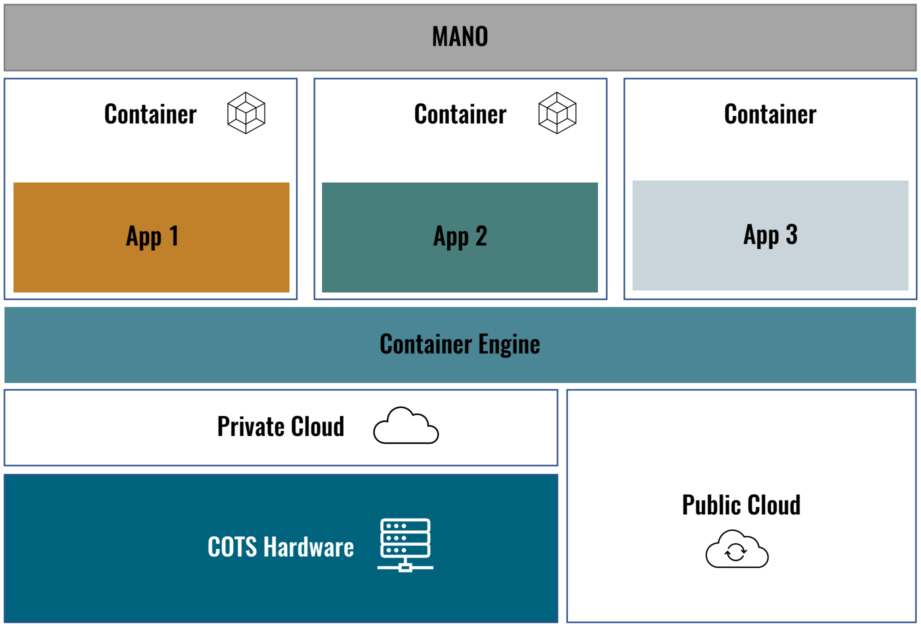

This hybrid approach shows an effective way for CSPs to enhance their service flexibility, reduce costs, and test new computing environments for different workloads. Public cloud vendors do not require CSPs to install and deploy specialized hardware to run services and applications, so the business strikes a balance between CAPEX- and OPEX-focused models.

PROS AND CONS

PROS

- CSPs can unlock c**ost reductions** with better resource utilization and network management, reducing the DC footprint, and lowering the TCO by leveraging optimized telco and public cloud services.
- **Greater network flexibility** with the ability to scale up and down quickly. The operator can also mix and match services, and even cloud providers, to meet their specific needs.
- **Enhanced business agility** with the ability to rapidly respond to customers and create new services on the fly.

CONS

- A hybrid approach can **increase complexity** in the cloud stack due to a mix of different vendors from the physical infrastructure and private cloud (NEV-side) and public cloud (hyperscaler-side).
- The **implementation of network capabilities** (refactoring NFs and software), storage, and servers can be challenging (especially when interfacing between different VNFMs and equipment).
- **Security** remains a prevalent concern in cloud environments, and this is especially the case when managing network and data security between different cloud environments.

SUMMARY AND CONCLUSIONS

CSPs have much to gain from cloudification. Moving network functions from hardware provided by NEVs to software comes with a host of benefits such as faster time-to-market and agility for services, enhanced network flexibility, and reduced operational and capital expenses. The advent of 5G has only increased the interest in virtualization in the telco realm. In this respect, CSPs have many aspects of cloudification that need to be put into perspective and help drive decision-making for how to best approach moving network functions to the cloud.

CLOUD STRATEGY FOR TELCOS

Indeed, CSPs have many aspects to consider when evaluating their cloud strategies. There are service performance and security thresholds, cost considerations, and integration with existing network infrastructure and NFs that must be evaluated when developing a cloud strategy. In this respect, the telco cloud remains an attractive option for moving NFs to the cloud. While a pure public cloud or hybrid cloud approach does exist, there remain many potential drawbacks for "brownfield" operators. Alongside this, a CSP's market positioning, infrastructure, and geographic location can make some deployment models more attractive than others. Given this perspective, CSPs should not only consider their unique position when evaluating a cloud strategy, but also the maturity and track record of said strategy. The on-premises private cloud stands not only as the most mature and tested cloud strategy for telcos, but the premiere option for meeting the special performance, security, and criteria telcos require.

ABIresearch
THE TECH INTELLIGENCE EXPERTS℠

Published October, 2022

157 Columbus Avenue 4th Floor
New York, NY 10023
+1.516.624.2500

About ABI Research

ABI Research is a global technology intelligence firm delivering actionable research and strategic guidance to technology leaders, innovators, and decision makers around the world. Our research focuses on the transformative technologies that are dramatically reshaping industries, economies, and workforces today. ABI Research's global team of analysts publish groundbreaking studies often years ahead of other technology advisory firms, empowering our clients to stay ahead of their markets and their competitors.

ABI Research is an independent producer of market analysis and insight and this ABI Research product is the result of objective research by ABI Research staff at the time of data collection. The opinions of ABI Research or its analysts on any subject are continually revised based on the most current data available. The information contained herein has been obtained from sources believed to be reliable. ABI Research disclaims all warranties, express or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.