Publication date:
November 2022
Author:
Rémy Pascal

Cloud-Native Networks are the Future of the Telecom Industry

Brought to you by Informa Tech

Omdia commissioned research, sponsored by Google Cloud

Contents

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Omdia commissioned research, sponsored by Google Cloud

Communication service providers are under pressure

Consumer demand and mobile traffic continue to increase worldwide, requiring communication service providers (CSPs) to continuously augment the capacity and capabilities of their networks. Even in places where traffic growth is slowing down, network-based competition is fierce, and CSPs need to invest massively in 5G to remain competitive and relevant.

Mobile networks are expensive to build, and as they continue to expand to handle more traffic, they are consuming more spectrum, more technologies, and more nodes. The result is a highly complex mixture of generations of networks that are becoming more expensive to manage, operate, and maintain, especially due to the lack of automation. In short, CSPs are not only facing high network capital expenditures (capex) but also high network operational expenditure (opex).

In the meantime, besides the well-established mobile broadband use case, there has been limited adoption of immersive AR/VR and multi-access edge computing (MEC) applications until today. No 5G killer applications means no significant new revenue in sight for CSPs in their traditional consumer business.

There is no shortage of potential new and exciting use cases and applications in the enterprise market but tapping into these opportunities requires networks that can address enterprise requirements. It also requires the combining high-performance connectivity with other technologies such as distributed computing, data analytics, machine learning (ML), closed loop automation, computer vision artificial intelligence (AI), and cloud-like operations.

Today, networks are built largely in the same way they have been built for decades, using purpose-built equipment and integrated solutions. The networks have become inflexible, costly to upgrade, and hard to tailor to the specific needs of their users.

In this context, mobile CSPs need to become more flexible, agile, and timely to bring new differentiated services to market. They also need to modernize and automate their operations, reduce their costs, and ensure the kind of reliability that human intervention cannot consistently guarantee.

This white paper looks at the adoption of cloud-native networks as one of the critical steps in this network transformation. In the context of this white paper, cloud-native networks refer to an operating model where, depending on the requirements and the CSP’s assets, the deployment can be on a centralized region of the hyperscaler’s cloud, on customer premises, or a combination of the two, but always using a consistent operating model.

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Omdia commissioned research, sponsored by Google Cloud

CSPs need to build more efficient and agile networks

To grow their business and improve their profitability, CSPs should focus on three main objectives:

- Modernizing their network
- Monetizing their network
- Lowering their total cost of ownership.

Modernizing the network

As mentioned before, the only way for CSPs to generate significant new revenue is for them to address new industry vertical markets. But it is not possible to address these new markets with the existing rigid, integrated networks that have been built for consumer mobile and fixed broadband.

One of the key difficulties for service providers is rolling out and extending the availability of their services across their entire territory of operation, which is usually a country. Multinational enterprise clients are not confined within the borders of the country in which they operate, they also often need to extend their services and operations to other geographies. Enterprise clients want a continuous, reliable, and seamless experience, and they want their CSPs to accompany them where their business happens, both domestically and internationally. By extending networks over a distributed public cloud with global reach, CSPs will be able to expand their roaming services rapidly, and far beyond the limits of their legacy physical infrastructure and points of presence.

Besides international markets, for most service providers there are still untapped opportunities closer to home in their domestic market, particularly with small and medium-sized businesses, which are not always as well served as larger enterprise clients.

Such cloud-native infrastructure can be leveraged to run the core as well as the access network functions for both any-G mobile and fixed networks and to support a flexible delivery of services and meet on-demand elasticity, local breakout, and quality-of-service (QoS) guarantees where and when needed. Learn more about the status of delivering 5G Core and RAN in this video.

Monetizing the network

Once the underlying cloud-native infrastructure, consistency, and availability is in place, the next challenge for CSPs is to design and offer attractive services to generate enterprise demand and revenue, and ultimately make a return on their investment. 5G-enabled edges can be designed to meet use cases in low-latency, security, and data residency with an on-premises solution, in contrast to more tolerant solutions delivered from the network edges of the CSP and private data centers.

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Regardless of the deployment model chosen, a cloud-native infrastructure can bring significant operational and cost benefits. The cloud-native network is indeed built on a common horizontal infrastructure that serves not only for 5G connectivity but also for AI, computer vision, or any other workloads that need to be processed in order to deliver the desired service. A common infrastructure means shared resources, higher utilization and efficiency, common control and management, and economies of scale. This also maximizes design, development, and portability in application modernization.

CSPs able to combine various technologies and capabilities will be more competitive than those focusing on a single isolated connectivity service. But because use cases are almost infinite and requirements almost always unique, CSPs cannot address them all by themselves. They should instead build an ecosystem of partners and developers to expand and enrich their offerings at the service layer. Again, the cloud-native model is better suited to achieve that, as it enables open network APIs to the ecosystem to build joint solutions.

Lowering the total cost of ownership (TCO)

The other side of the coin is the CSP's internal cost structure. To operate a profitable business, CSPs not only need to sell more and to new clients, but they also need to control their costs and adjust them in a flexible manner and as close as possible to real time.

This is the beauty of the public cloud elastic consumption model, which enables teams to add or remove capacity and features where and when needed, instead of having to build for peak demand that only occurs a fraction of the time, or to future-proof the infrastructure on the assumption of hypothetical future demand. These cloud-native technologies, security, and operational fabric are extended to the edge for that desired consistency and agility.

The on-demand model is equally useful to right-dimension the core network capacity and access network capacity, according to the precise number of users and the volume of traffic, and their constant fluctuations.

Such flexibility and scalability are made possible by the containers-as-a-service (CaaS) model, which allows the onboarding and efficient management of the lifecycle of network functions and applications in near real time, and the ability to scale up and down capacity in minutes instead of weeks or months before.

Another key benefit of the cloud-native approach is the advanced level of automation possible with integrated ML and AI, and software-centric operations. It significantly lowers the risk of human error, reduces the number and therefore cost of human interventions, and also frees up skilled staff to accomplish more strategic tasks than daily operations, troubleshooting, and maintenance. Generally, CSPs migrating their network workloads to a cloud-native infrastructure should expect less site visits, less truck rolls, and less alarms to manage.

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Omdia commissioned research, sponsored by Google Cloud

Understanding open and cloud-native networks

The cloud is a culture and set of principles

To fully benefit from the cloud, CSPs should not think of it as just a new type of underlying infrastructure to run their network functions on. They should instead think of the cloud as a new paradigm for designing, deploying, and operating networks as well as a set of concepts, guiding principles, methodologies, processes, and business models – one could say that all that shapes a culture.

This cloud culture encompasses many aspects. **Table 1** elaborates on some of the fundamental principles of flexibility and scalability, efficiency, reliability, and speed.

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Omdia commissioned research, sponsored by Google Cloud

Table 1: The many aspects of a cloud culture

Flexibility and scalability	Cloud-native flexibility and scalability are possibly the most important concepts of all. Flexibility to select infrastructure by choosing either an on-premises edge cloud or a more centralized public cloud region, depending on the service requirements. Scalability to adjust the capacity to meet the fluctuations in demand, whether to serve a few users or a few million. Cloud-native enables resources to be delivered through an "as-a-service" model.
Simplicity	Simplicity refers to deployment-agnostic consistency in application distribution, modernization, continuous integration/continuous delivery (CI/CD), compute, networking, operations, and management. Open standards across the cloud-native and solutions stack with API abstractions as key contracts.
Efficiency	Efficiency refers to making the most of available resources, including network resources, but also spectrum resources, computing resources, and human resources. One of the key enablers is advanced and systematic automation, which in turn leads to better cost efficiency and a lower TCO.
Reliability	Reliability refers to the quality of delivering trusted and consistent desired results. It is also closely associated with automation and the operational model. The point of automation is to improve the speed of repetitive tasks and to minimize errors. This is accomplished by limiting the human factor in operations. With less human intervention, we achieve less errors and more speed. One practical implementation of this principle is site reliability engineering (SRE). SRE is a Google mindset and a set of engineering practices leveraged to run reliable production systems, solve problems, and automate operations tasks.
Speed	When it comes to bringing to market new services and applications fast while maintaining the lowest cost possible, the cloud's superiority has been proven. Speed matters but it should not be improved at the expense of reliability; it therefore involves necessary methodologies and processes. DevOps is a philosophy and a set of practices and tools that enables the quick development and deployment of applications and services by combining software development and lifecycle management processes.

Source: Google Cloud and Omdia

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Omdia commissioned research, sponsored by Google Cloud

Cloud-native networks have building blocks in common

The cloud-native approach combines a set of fundamental principles (as seen above) with practical building blocks and technologies, as shown in **Table 2**.

Table 2: The building blocks of a cloud-native approach

Common platform	At the base of a cloud-native architecture is a common horizontal platform to support the needs of different domains. For telecom service providers, it means the possibility to run radio access networks (RAN), core, enterprise applications, and OSS/BSSs on the same cloud platform. Some of the key benefits of using a common platform are that it enables the sharing of computing and network resources via a common management layer and tools, thus enabling the benefits of hybrid cloud-edge distribution and scale.
Vertical and horizontal disaggregation	Vertical disaggregation refers to the separation of the hardware and software that enables the network functions to run as software on a cloud infrastructure, instead of using purpose-built appliances. Horizontal disaggregation is the separation of network functions into smaller sub-functions that can run and scale independently from each other as needed and can also be positioned in different locations. In the RAN, for example, disaggregation enables the centralization and pooling of the baseband central unit resources and the processing of the corresponding workloads on the public cloud while the distributed unit function remains at the cell site.
Openness and multi-vendor	An open framework simply means a framework open to various participants, extending across the infrastructure, application, services, and monetization model with stringent security guardrails. It should be combined with open and common interfaces in order to facilitate the interoperability between these participants and their respective solutions. A good example is the open RAN's open fronthaul interface, which enables the radio unit from one vendor to connect to the distributed unit of another vendor. Application programming interfaces (API) also enable third parties such as application developers to interface and communicate with the CSP's network. They help incubate ecosystems and test ideas that will ultimately enhance the value of the network for the operator itself and for their clients and users.

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Omdia commissioned research, sponsored by Google Cloud

Micro-services	Micro-services involve a software approach where a monolithic service is broken into smaller autonomous services. Such micro-services run on containers, making it easier to manage their lifecycle and scale independently.
Intent-based automation	Declarative automation means the possibility to define a desired behavior using intent-based policies and configurations – typically for a network operator, offering the best possible customer experience. Such semi-supervised constructs enable the system to automatically translate the intent into real-time actions, with data and model-driven decisions and closed loop automation in an accelerated manner.
End-to-end security	A cloud-native infrastructure should guarantee network security and at-rest security (when data is on disk rather than flowing). It should also enable customers to manage encryption. Besides infrastructure security and cybersecurity, the multi-vendor cloud-native model enables another type of security: supply chain security. The diversification of the vendor pool indeed allows reduction in the dependence on a small number of vendors.

Source: Google Cloud and Omdia

Omdia's views on the virtual and cloud-native core and RAN markets

Even though when 4G was first introduced, CSPs only used purpose-built appliances, the 4G LTE packet core, known as the EPC, which was one of the first groups of network functions to be virtualized by CSPs almost a decade ago. Today, the virtualization of 4G core networks is well advanced around the world, and most new investments made to modernize 4G core or expand their capacity are directed toward virtual and cloud EPC solutions.

The 5G core is cloud native in essence and, with a few exceptions, all 5G cores are deployed on private or hyperscalers public clouds, and almost no purpose-built equipment is used.

As a result, when considering the entire market for 4G and 5G packet core networks, Omdia estimates that, as of 2020, already more than 63% of the total packet core investments made by CSPs globally were on virtualized or cloud-native packet core solutions. Omdia forecasts that this will reach 99% by 2026 (see **Figure 1**).

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Omdia commissioned research, sponsored by Google Cloud

Figure 1: Packet core revenue, 2020–26

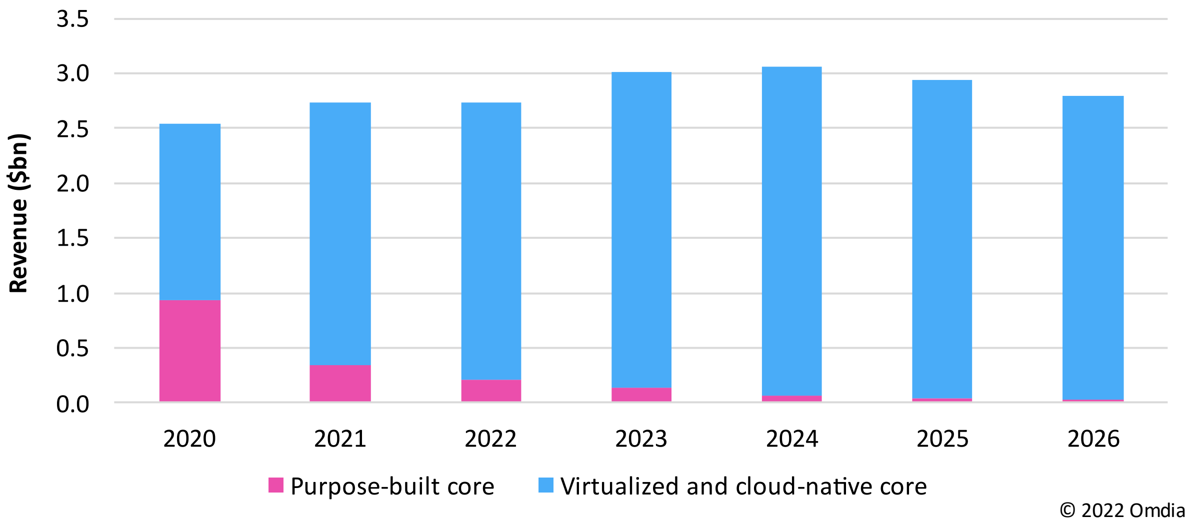

Source: Omdia

In the RAN, the journey toward virtual and cloud solutions started more recently. CSPs wanted to make sufficient progress in their cloudification journey in other network domains (i.e., OSS/BSS and core) before getting to the RAN, which is perceived as the last frontier. The RAN indeed presents specific challenges related to its distributed nature and stringent latency requirements, as well as business and operating models.

Nonetheless, virtualized RAN (vRAN) started gaining popularity in the last few years, initially driven by greenfield CSPs such as Rakuten Mobile in Japan and DISH Wireless in the US, and some leading brownfield CSPs such as Verizon, which is now deploying vRAN at scale in its 5G network.

The virtualization of the RAN represents a massive effort, so adoption will take time, but vRAN will be the fastest-growing segment of the RAN market in the next few years. Omdia estimates that open vRAN will grow from less than 1% of the total RAN market in 2020 to almost 16% in 2026, as shown in **Figure 2** (this forecast does not include non-open vRAN deployments, which represent an additional portion of the total market).

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Omdia commissioned research, sponsored by Google Cloud

Figure 2: RAN revenue, 2020–26

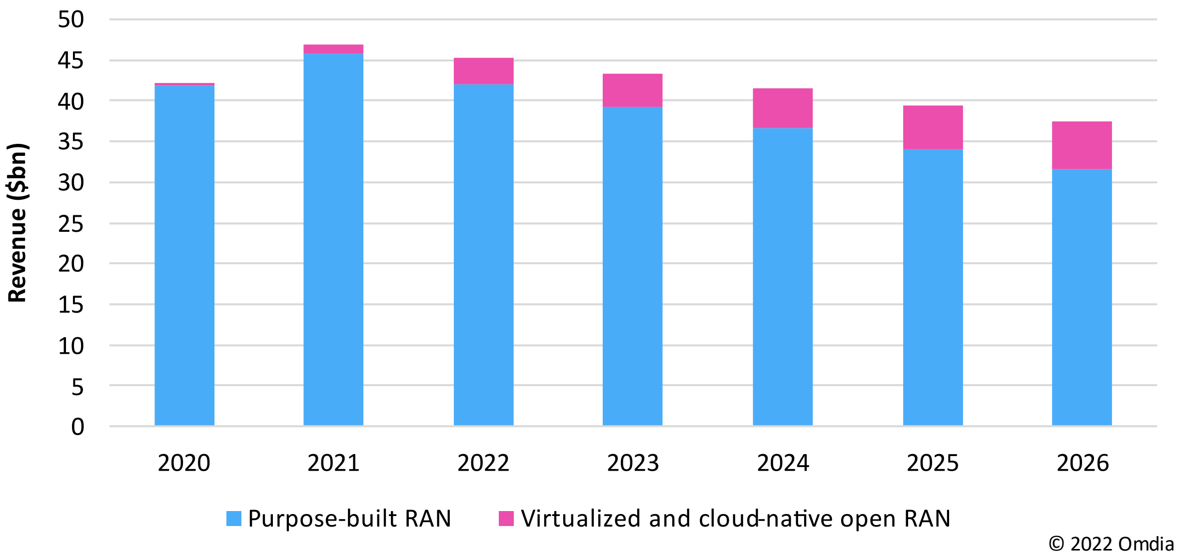

Source: Omdia

These forecasts focus on the market for network functions and therefore do not differentiate between basic virtualization and cloudification, but there is an important distinction that should be understood. While virtual network functions (VNFs) have been the first step in this journey, the ultimate goal is to use containers as they offer many benefits over virtual machines including higher efficiency, better scalability, and easier upgrades.

The above essentially corresponds to capex but, globally, capex only represents about 25% of all mobile service providers' expenditure, with the rest being composed of opex.

Cloud-native, highly automated networks can be designed, deployed, operated, and maintained very efficiently, and by smaller teams of engineers and technicians than traditional networks, which enable service providers to lower their opex. Network cloudification, when combined with baseband resource pooling, also enables the simplification of cell sites, thereby reducing their footprint and resulting in lower site rental and electricity costs. To learn more about 5G core networks and RAN, listen to Roberto Kompany, Principal Analyst, Omdia Research: Delivering cloud-native 5G Core and RAN.

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Omdia commissioned research, sponsored by Google Cloud

Google Cloud for CSPs' RAN and core networks

While Google's contributions to the cloud and open source ecosystems, like Kubernetes (K8s), are probably better known, Google has also been a key participant in the telecom ecosystem and various industry organizations for almost a decade. Google, for example, has made more than 400 contributions to the 3GPP 5G standards since 2015, and it has also been a key participant and contributor to the TM Forum.

In June 2021, Google Cloud joined the O-RAN Alliance whose mission is to reshape RAN to be more intelligent, open, virtualized, and fully interoperable. In April 2022, Google Cloud and 24 other organizations co-founded Nephio, a project to build a cloud-native, intent-based network function automation framework for CSPs.

Google Cloud is not a network function provider, but a technology provider that offers a unified horizontal platform for operators to run their network and IT workloads on, across domains (RAN, core, apps) and across locations (data centers, edge, cell sites). Google Cloud also provides a range of tools for software deployment, auditing and repairs, and configuration management, as well as databases, data analytics, ML, storage, and third-party services.

Google Distributed Cloud Edge (GDC Edge) is Google Cloud's fully managed hardware and software offering, which extends Google Cloud's infrastructure and services to the edge and into CSPs' and enterprises' data centers. In connected mode, it offers full flexibility for edge-to-cloud distribution of workloads and seamless integration with Google Cloud services (see **Figure 3**).

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Figure 3: Google Distributed Cloud Edge

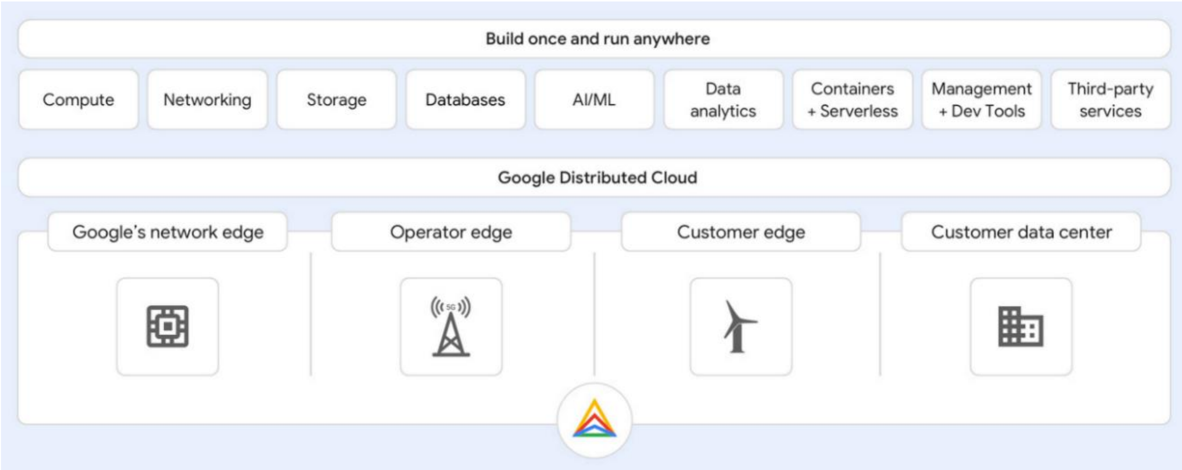

Source: Google Cloud

Depending on what is required, Google Distributed Cloud can run on:

- Google's own edge network with over 173 locations worldwide, connected via a high-performance transport network
- CSP's own edge infrastructure
- Enterprises' data centers or on-premises edge locations for converged application and cellular private network deployments.

Figure 4 below shows an example of high-level architecture for a 5G network hybrid cloud deployment. It leverages Google Cloud Regions for OSS/BSS, IMS and other non-latency-sensitive functions, GDC Edge for the 5G control plane and MEC applications, the CSP's far edge location for the user plane function (UPF) and RAN central unit (CU), and, finally, the RAN distributed unit (DU), which is located at the cell site. However, what they all have in common is that they use Anthos. Anthos is Google's open-source platform that unifies the management of infrastructure and applications across the public cloud, edge, and on-premises.

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Omdia commissioned research, sponsored by Google Cloud

Figure 4: 5G deployment using Google Distributed Cloud Edge and Anthos

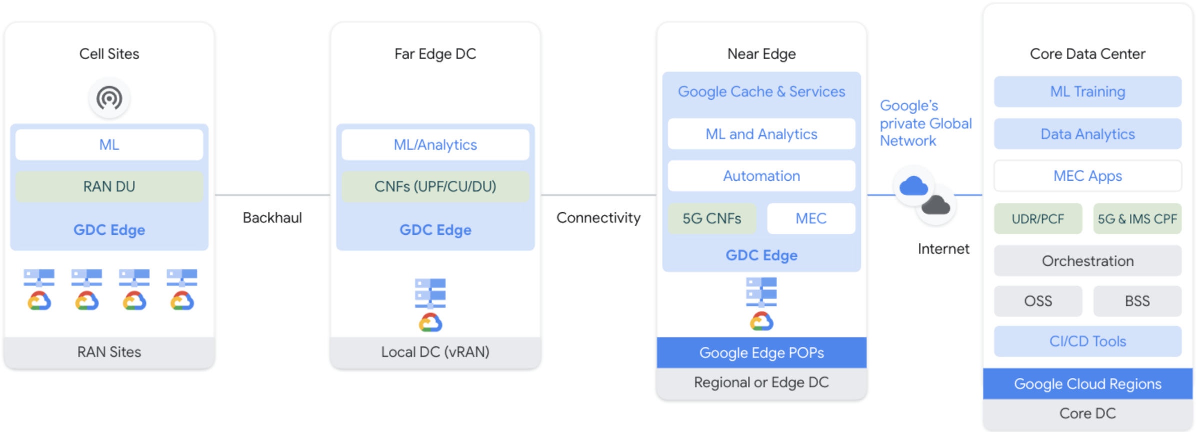

Source: Google Cloud

Google Cloud also offers a network function validation program that simplifies software upgrades and de-risks production environments, as well as network analytics tools to ingest, organize, and process network-related data and drive actionable insights for network optimization.

In June 2022, Google Cloud announced new partner-led private network solutions running the control plane functions on a Google Cloud Region, user plane on GDC Edge at the network edge or on-premises, and vRAN on-premises, which bring together over 200 partner edge applications for industrial IoT and manufacturing, media and entertainment, and retail, as well as horizontal applications for commerce, analytics, security, etc.

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Omdia commissioned research, sponsored by Google Cloud

Conclusion

The journey toward cloud-based networks is moving forward rapidly, and Google Cloud is well positioned to be the partner for CSPs willing to accelerate their transition toward cloud-native core and access networks.

Operators such as Airtel, AT&T, Bell Canada, Deutsche Telekom, Jio, Orange, Telus, and TIM are already working with Google Cloud.

Others that want to build more flexible, agile, efficient, and reliable cloud-based networks should get started. CSPs should first clarify their expectations and define clear requirements. They should engage with cloud specialists to learn the technology and test it, but also conduct their own assessment of the technology and the business case. In that journey, they can lean on partners like Google Cloud, but service providers should also acquire a minimum set of skills and internal capabilities to play an active role in their transformation and make the most of it.

CSPs can find information and resources to learn more by visiting Google Cloud for telecommunications and read more about cloud-native networks.

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Appendix

Author

Rémy Pascal
Principal Analyst, Mobile Infrastructure
customersuccess@omdia.com

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Get in touch

www.omdia.com
askananalyst@omdia.com

Omdia consulting

Omdia is a market-leading data, research, and consulting business focused on helping digital service providers, technology companies, and enterprise decision-makers thrive in the connected digital economy. Through our global base of analysts, we offer expert analysis and strategic insight across the IT, telecoms, and media industries.

We create business advantage for our customers by providing actionable insight to support business planning, product development, and go-to-market initiatives.

Our unique combination of authoritative data, market analysis, and vertical industry expertise is designed to empower decision-making, helping our clients profit from new technologies and capitalize on evolving business models.

Omdia is part of Informa Tech, a B2B information services business serving the technology, media, and telecoms sector. The Informa group is listed on the London Stock Exchange.

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Omdia's consulting team may be able to help your company identify future trends and opportunities.

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.

Copyright notice and disclaimer

The Omdia research, data and information referenced herein (the “Omdia Materials”) are the copyrighted property of Informa Tech and its subsidiaries or affiliates (together “Informa Tech”) or its third-party data providers and represent data, research, opinions, or viewpoints published by Informa Tech, and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice and Informa Tech does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an “as-is” and “as-available” basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa Tech and its affiliates, officers, directors, employees, agents, and third-party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa Tech will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.

© 2022 Omdia. All rights reserved. Unauthorized reproduction prohibited.