Are public clouds right for telecommunications service providers?

"Three, the UK mobile provider, believes it will reduce costs 30% by moving its OSS and BSS to public cloud, going from 95% on-premises to 90% in the public cloud and rationalizing its IT portfolio, cutting about 70% of its applications, in the process."

Omdia[1]

Introduction

Telecommunications service providers (telcos) have been transforming into digital service providers for years. Now with the early stages of 5G and edge deployments coming and as more agile and efficient cloud-native architectures take hold, they must consider whether and when they should build out and manage the infrastructures that comprise these new services as private clouds or take advantage of the public cloud resources available from others. In other words, should they buy or rent? These decisions should typically be based on partnership strategy, economics, operations capacity and capability, architectural choices, time-to-market requirements, or some combination of these, and providers should also consider how they may change over time. This whitepaper highlights the relative advantages of private and public cloud environments among the key decision criteria.

Cloud considerations and choices

Service providers can build their own private cloud platforms, which requires investments in hardware and software infrastructure. This option allows telcos to select their hardware of choice and then deploy and operate workloads, applications, and functions, locating them where they see fit. This choice usually maximizes flexibility and opportunities to customize for special requirements.

Another option is to use public clouds, such as Alibaba, Amazon Web Services (AWS), Microsoft Azure, Google Cloud, or IBM Cloud. The public cloud option offers on-demand infrastructure, which reduces capital investment as well as ongoing operational and life-cycle management. Public cloud can offer a great incubator environment to not only create new applications and services but also bring them to market and scale rapidly. As these innovations mature and stabilize, however, private clouds may offer more efficient delivery and operations. Public clouds can also be used as backup infrastructure for private clouds with low consumption-based costs.

A third option brings public cloud platforms into private clouds using on-premises solutions (hardware and software stacks) offered by public cloud providers. While this approach provides more control over deployment location, the actual infrastructure is preset and service providers have more limitations on flexibility and control of their business and services.

1 *Omdia, "Public Cloud for OSS," Nov. 2020.*

"SaaS and public cloud will make inroads into the market for monetization platforms by growing more than 6.5X from 2019 to 2025 and increase its share to over 14% of the total spend in this segment."

Analysys Mason[2]

How do these clouds compare

Overall, creating a private cloud requires capital investment and the operating expenses associated with labor to run and maintain that infrastructure plus the physical location and environment costs. The advantage comes in control. Service providers can define the cloud infrastructure hardware and software platform, optimize it (i.e., bare metal, containers, automation), and control its life cycle.

Telcos can better support applications with special requirements for low latency, large amounts of storage, specialized compute, and network fabric and protocols. This opens the door to more ways to differentiate service offerings with artificial intelligence algorithms, for example, and avoid excess costs for things like high bandwidth transport or public cloud graphics processing unit (GPU) use. Special public cloud resources like GPUs may also have limited availability, and access during peak demand may be limited by customer priority. In a private cloud environment, telcos can control and address data privacy and protection needs, including data sovereignty, as well as regulatory requirements. In some cases, public cloud providers have the right under their terms of service to collect data they might use in a competitive manner.

Businesses often select private clouds to deliver established applications and services with consistent demand and 24x7 operations more efficiently than they could with public clouds. These "permanent" workloads can drive high consumption costs on public clouds.

Table 1. Decision considerations among private and public clouds and on-premises stack

Consideration	Private cloud	Public cloud	Public cloud on-premises stack
Financial treatment	Capital expenditure	Operational expense (opex)	Operational expense
Capital investment	Required	None (costs were opex)	Commodity
Operational expense	People, environmental, space, maintenance	Consumption-based	People, environmental, space, maintenance
Existing infrastructure	Sunk cost, repurpose	Excess	Excess
Consistent operations	High with control of hardware and software	Variability when using multiple public clouds	Variability when choosing more than one solution
People	High	Low/medium	Medium, depending upon support
Capacity provisioned/ scaling	Provision to peak demand. Slower to scale beyond.	Public cloud provider can absorb surge in traffic	Provision to peak demand. Could burst/ scale into public cloud.

2 *Analysys Mason. "Monetisation platforms: worldwide forecast 2020–2025," Oct. 2020.*

Consideration	Private cloud	Public cloud	Public cloud on-premises stack
Life cycle	Control selection and upgrades of hardware and software	Controlled by cloud provider	Controlled by cloud provider
Demand variability	Predictable, less variable demand; must plan capacity	Variable demand supported by cloud infrastructure	Less variable demand to fit capacity; could burst to public cloud
Cost per traffic volume	Cost per gigabyte can be lower for high volume	Cost efficient for low volume or varied traffic patterns	Cost per gigabyte can be lower at high volume
Application dependencies (e.g., multicast, real-time, network fabric, VxLAN)	Adaptable, customizable to meet complex requirements	Most are not supported. Standard compute, memory, storage, network. May offer telco services.	Most are not supported. Standard compute, memory, storage, network.
Compute intensity	GPU, CPU pinning under control	Extra costs for GPU, no CPU pinning	Extra costs for GPU, no CPU pinning
Storage intensity	Control storage costs and performance	High cost of storage services	High cost of storage services
Data privacy and security	Control	Some providers may collect information	Some providers may collect information
Location and data sovereignty	Control	Country- and provider-dependent	Control
Regulatory compliance	Built to requirements	May not meet all requirements	May not meet all requirements
Portability and interoperability	High within private cloud environment	Lower when public cloud(s) used with private cloud	Lower when public cloud(s) used with private cloud
Intellectual property	Owned, controlled, managed. Can differentiate offers.	Subject to terms of service. Can add cost, limit differentiation.	Subject to terms of service. Can add cost, limit differentiation.

Public clouds, on the other hand, are well-suited for applications that do not have critical real-time or low latency requirements, can perform well on standard CPUs, and have no real network dependencies. Yet, each public clouds is distinct and has different strengths for different types of workloads. So one might better fit network intensive applications, and another better fit apps that require top storage performance. No single provider has the best environment to meet unique needs across various types of workloads. Applications well-suited for public cloud typically require a small storage

footprint so that storage service costs do not outweigh other economic benefits. Those economic benefits stem from outsourcing the cloud hardware and software platform plus its operation and maintenance.

When they choose public cloud, service providers pay the operational expenses of the cloud services consumed instead of the combined capital and operational expenses of building, managing, and maintaining their own cloud infrastructure. The financial advantages can be great when application demand fluctuates significantly because public cloud consumption costs scale up or down with actual usage. Software-as-a-Service (SaaS) models can potentially take those cost advantages even further. Back-office applications will often fit well here. Development, experimentation, and testing can also flourish with flexible, on-demand public cloud resources meeting more unpredictable, fluctuating needs.

On-premises public cloud stacks lower costs as they commoditize private cloud infrastructure. These can provide a strong solution for creating application environments that do not require customization or do not have special requirements for operators who want greater control over their infrastructure than public clouds offer. Their ability to extend easily into public cloud can add flexibility, capacity, and resiliency, but this ability is limited to the specific cloud provider.

Additional considerations

Data is currency in the digital world. Massive amounts of data in cloud environments facilitate information sharing for both personal and professional purposes. However, ownership issues can arise because the data of many organizations coexists on shared public cloud infrastructures, requiring protection for two types of property:

- Property of the cloud, or hardware and software that constitute the business assets of the service provider.
- Property in the cloud, or data and applications from individual cloud users who need assurance that they alone retain control over their own information.

Laws grant protection to both but regulations are complex. Because clouds do not have national boundaries, the diverse protection among countries creates even more complexity. Until governing laws and regulations reach maturity, intellectual property ownership in the cloud (e.g., application logic, valuable data, trade secrets) may not be entirely secure.

Lock-in or lack of portability is often overlooked as public clouds are perceived as public and therefore open. Yet each is built to its own specifications and how applications access services (e.g., storage, network) is unique for each cloud, even when, for example, each offers a service based on open source Kubernetes. The development tools in a given public cloud also follow this model. Acquired skills and training need to focus on the specific public cloud used. These factors can be a strong hindrance to moving applications out of one public cloud to a private cloud or another public cloud.

Availability should be considered as well. Public cloud providers have experienced a number of high profile outages (as long as a full week), even as more businesses depend on public cloud infrastructures. When a service provider commits service-level agreements (SLAs) to its customers, its private cloud helps because the provider can design for the resiliency required and can control the resources to remediate the infrastructure problems encountered. However, public cloud providers have improved their performance against SLAs, and telcos need to assess whether public cloud SLAs can meet their workload needs or be negotiated.

Use cases and potential fit with public cloud

Internet protocol television (IPTV) and live video streaming

Internet protocol (IP) multicast-based applications like IPTV or live video streaming cannot effectively be moved to the public cloud. Public cloud providers and their on-premises solutions do not offer native support of multicast protocols, even if they do document workarounds for such applications. Especially because these applications are the foundation of revenue-producing services, they need reliable, efficient, performant infrastructure to deliver them. That infrastructure is private cloud.

Billing

A billing application can usually move to a public cloud without issue. Of course, it may need to be adapted to the particular cloud provider's environment. It is not likely to be running constantly. Rather, it may run daily after a day closes, processing for customers whose specific billing cycle matches that date. As long as it can access its data sources from the public cloud, its processing is straightforward and not complex. It does not have special infrastructure or network requirements, as well.

Charging

Charging functions and especially real-time converged charging functions require low latency, high availability performance with a large volume of transactions per second. They also require strong integration with network functions to monitor network resource usage, which is best achieved with an on-premises deployment or with high speed connections to the public cloud.

Radio access networks (RAN)

RAN workloads have specific timing and real-time synchronization requirements that point to operating them in a private cloud. The first is that the distributed unit requires a location near the radio units it supports. Additionally, there are strict timing requirements in fronthaul and specific timing protocols for running these disaggregated functions on a cloud infrastructure. Public clouds and public cloud on-premises stacks do not meet these timing needs.

Edge applications and services

On-premises public cloud stacks can serve the objectives of enterprise services teams within the telcos, particularly when capital budgets are tight. These teams need to move quickly to make progress toward revenue targets. Public cloud stacks offer turnkey deployment and also bring the specific cloud provider's developer ecosystem with them. The built-in ecosystem can make it easier and faster to partner in delivering value-add services to targeted vertical markets and capturing new revenue streams.

Innovation incubation

When designing a new service or artificial intelligence and machine learning (AI/ML) model, public cloud can provide agile and cost effective environments for initial testing and training. Required resources change as development and the scale and scope of work evolves. Once a service or model is ready for production deployment, however, moving to a private cloud may better suit its technical, scaling, and pricing requirements.

Redundancy for critical services

Because public cloud costs are based on consumption, they provide a means to lower the cost of redundancy for critical high availability services that providers run in private clouds by activating redundancy resources only when needed. As long as other service needs (e.g., regulatory) can be met, using public clouds can reduce the capital investment in backup infrastructure to provide the required redundancy.

Market announcements and analyses

As discussed, adopting public clouds clearly benefits communications service providers (CSPs) for certain applications. Industry analysts are taking notice. A Nov. 2020 Omdia report found that "In general, telco RFPs [requests for proposals] for new OSS [operations support systems] ask for the ability to run the software entirely in public cloud and for the option to host on private cloud and burst into public cloud as needed."[3] While they note in the same report that "CSPs are still typically choosing to run OSS in their private clouds for cost or familiarity reasons," Omdia believes "the OSS industry will gradually move towards a SaaS approach where the telecom operator no longer needs to worry about where the application is hosted. Instead they focus on the OSS applications themselves, their functionality, and their relationships with the vendors that support them."[3]

Analysys Mason forecasts SaaS and public cloud will make inroads into the market for monetization platforms by growing more than 6.5X from 2019 to 2025 and increase its share to over 14% of the total spend in this segment.[2]

Several operators have publicly shared plans for adopting public cloud. For example, the UK mobile provider Three believes it will reduce costs 30% by moving its OSS and business support services (BSS) to public cloud, going from 95% on-premises to 90% in the public cloud and rationalizing its IT portfolio, cutting about 70% of its applications, in the process.[1] As part of its announced partnership, Microsoft will be AT&T's preferred cloud provider for non-network applications. AT&T plans to move most non-network workloads to the public cloud by 2024 with its "public cloud first" approach.[4] OSS and BSS suppliers like Amdocs and Netcracker have also announced partnerships with public cloud providers. For network functions, Telefonica Germany recently announced virtual 5G core with AWS in an industrial use case proof of concept.

On the other hand, some have moved away from public clouds when volume and complexity of services have grown. Dropbox moved out of public clouds and into its own custom infrastructure when its volumes grew, and its 2018 IPO filing emphasized trust, security, and reliability as reasons for having 90% of users' data on its custom datacenters.[5] Cloud providers, however, have continuously improved their offers to address these concerns.

Conclusions for service provider cloud use: Private, public, and hybrid

For telecommunications service providers there is no straightforward answer to whether they should move applications to a public cloud. They must consider whether potential public cloud partners might be direct competitors, delivering core communication services within their strategic time horizons.

---

3 *Omdia, "Public Cloud for OSS," Nov. 2020.*

4 *AT&T press release, "AT&T and Microsoft Announce Strategic Alliance," July 2019.*

5 *United States Securities and Exchange Commission. "Form S-1 Registration Statement," Feb. 2018.*

As one might expect, cost savings and outsourcing certain challenges cannot always align with an application's primary need to deliver its expected services and capabilities. Clearly there are workloads that do fit and moving them to a public cloud will have financial and, possibly, operational advantages.

When business uncertainties, whether market or execution risks, are high, public clouds offer an efficient way to expand footprint, launch pilot programs, and elastically scale. When production services and applications have established, more consistent demand, private clouds are often able to deliver economic advantages and improved control.

Given a mixed set of workloads, each with its own demand characteristics, CSPs can benefit from a hybrid cloud environment, using both private and public clouds. To implement hybrid cloud effectively, network services and applications should be interoperable and portable horizontally across private and public clouds, and with uniform consistent interfaces for development, operations, and life-cycle management. The infrastructure becomes elastic, fungible, and resilient.

Cloud-native platforms like Red Hat® OpenShift® can help CSPs establish a consistent development, continuous integration and delivery (CI/CD) pipeline, and operational environment with better security for applications and services whether built or purchased. Operators and their vendors can write solutions once and then choose the cloud where they run, simplifying the use of hybrid cloud by running workloads on multiple clouds without modification.

Service providers can also choose to burst capacity beyond that of their private clouds into public ones. They can migrate applications from private to public cloud when appropriate and avoid integration efforts for the new environment. If a better public cloud option or partnership comes along, they can readily move applications over or use multiple public clouds for more resilient applications and services deployments than one cloud can provide. Finally, CSPs have the freedom to pull back applications that do not produce greater value or are well supported in a public cloud environment.

Learn more about Red Hat OpenShift or try Red Hat OpenShift in private or public cloud.

North America
1 888 REDHAT1
www.redhat.com

Europe, Middle East, and Africa
00800 7334 2835
europe@redhat.com

Asia Pacific
+65 6490 4200
apac@redhat.com

Latin America
+54 11 4329 7300
info-latam@redhat.com

facebook.com/redhatinc
@RedHat
linkedin.com/company/red-hat

redhat.com
#F26141_1220