tmforum

An essential
CSPs
PUBLIC
CLOUD:
but not singular solution for

Author: Tim McElligott, *Senior Analyst*
Editor: Dawn Bushaus, *Managing Editor*

December 2019

Contents

The big picture

When US communications service providers (CSPs) <u>CenturyLink</u> and <u>Verizon</u> sold off their data centers in 2016 and 2017, many in the telecoms sector felt the blast of air from the burst balloon of telco cloud opportunity. But by the time <u>AT&T</u> and <u>Telefónica</u> followed suit in 2019, the sentiment had shifted from lost opportunity to a realization that there may be a better way for CSPs to participate in the cloud and platform economy: by partnering with public cloud providers.

Technologically, culturally and in terms of business models, the public cloud is still the quandary for telcos. It is no less complex than traditional networking and is perhaps even more complex. It is also risky.

> In the minds of many CSPs, the public cloud still contains too many unknowns regarding security, performance, true cost of operation, interoperability, and control of both customer data and network connectivity. Most of all, it does not adhere to traditional telco processes, practices and standards.

Historically, any one of these concerns would be show-stoppers for CSPs that prefer a more principled engineering environment. This explains why, according to our research, most CSPs have deployed less than 5% percent of their operations software in the public cloud, at a time when investment in new operational and business support systems (OSS/BSS) should be on the rise in preparation for 5G.

More operators are deploying BSS in the public cloud, however, and this summer <u>AT&T and Microsoft announced a blockbuster $2 billion, multi-year deal</u> to move all AT&T's non-network workloads to the Microsoft Azure public cloud. AT&T's simultaneous proclamation that it is becoming a "public cloud first" company is likely to convince other operators to embrace public cloud, not only for telco operations, but also to target enterprise customers.

Two-way street

The onus in these burgeoning relationships is not all on CSPs. Cloud providers must work closely with operators to address end customers' needs. They must also recognize that they are working with partners who have invested considerably in network functions virtualization – investments telcos are unlikely to abandon. Just as CSPs need to adapt to cloud environments, cloud providers must learn how to engage with CSPs operating multi-generational networks.

This report looks at CSPs' strategies to move corporate IT, OSS/BSS and network applications to the cloud, with particular focus on OSS/BSS. It also examines the strategies that leading public cloud providers are using to address the telco market. Both sides are coming to the realization that the future lies not in public cloud alone but in a combination of cloud architectures.

Read the report to understand:

- What the drivers are for CSPs to embrace cloud architectures, and conversely why CSPs are wary of cloud

- Which types of workloads CSPs are moving to the cloud first

- The difference between cloud-native, public cloud, hybrid cloud and multicloud

- How Amazon Web Services, Google and Microsoft are partnering with CSPs and how they manage multicloud environments

- How OSS/BSS suppliers such as MYCOM OSI and Optiva are delivering cloud-native applications

- Why the AT&T-Microsoft deal is important and how other leading CSPs such as Dish Network, Rakuten, Telefonica, Three UK, TIM, Truphone, Vodafone and Verizon are partnering with cloud providers

Section 1

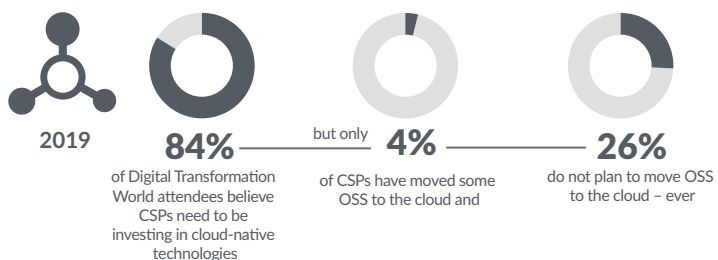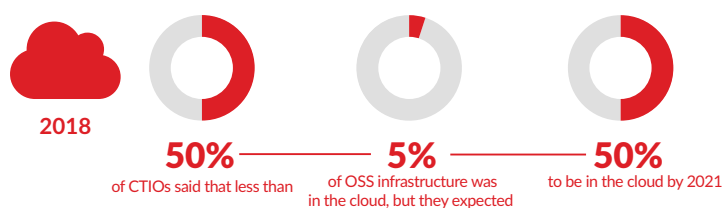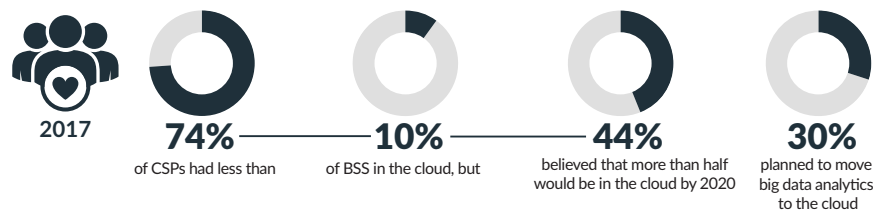# Have CSPs truly changed their minds about cloud?

It isn't as if communications service providers (CSPs) are against the concept of cloud. As noted in the introduction to this report, most large operators built their own data centers a decade ago only to sell them off when they failed at being cloud providers. But most CSPs are hesitant to put their own workloads into a public cloud environment, particularly operational and business support systems (OSS/BSS). This begs an important question as CSPs once again look to target enterprises, this time through partnerships with public cloud providers: If the cloud is not good enough for telco operations, why would enterprises put their workloads there? CSPs will have to come up with an answer.

Expectations for CSPs to move BSS to the cloud have always been higher than for moving OSS. In 2017, TM Forum published a report called *Cloud BSS: The migration begins*, which found that operators were beginning to test the cloud waters by moving some customer management and revenue management applications. But there was no strategic objective to move the entire BSS stack.

At the time, three quarters of CSP respondents reported having less than 10% of their BSS in the cloud (see graphic). But operators had begun adopting the software as a service (SaaS) model for some BSS applications, partnering with companies such as Salesforce and SAS.

This is likely why nearly half of respondents to the survey predicted that more than 50% of all BSS would be in the cloud by 2020. We are surveying CSPs now for a *Future BSS* report to be published in January, but early indications are that operators are nowhere near achieving that prediction.

Cloud OSS/BSS then and now

2017

74%
of CSPs had less than

10%
of BSS in the cloud, but

44%
believed that more than half would be in the cloud by 2020

30%
planned to move big data analytics to the cloud

2018

50%
of CTIOs said that less than

5%
of OSS infrastructure was in the cloud, but they expected

50%
to be in the cloud by 2021

2019

84%
of Digital Transformation World attendees believe CSPs need to be investing in cloud-native technologies

but only **4%**
of CSPs have moved some OSS to the cloud and

26%
do not plan to move OSS to the cloud – ever

TM Forum, 2019

One finding in the BSS survey has proven true, however: Big data analytics are workloads suitable for cloud. In 2017, nearly a third of CSP respondents said they were already deploying or planned to deploy analytics in the cloud. The growing adoption of artificial intelligence (AI) and machine learning is reinforcing this trend. A 2019 survey of Digital Transformation World attendees revealed that a full 84% of respondents believe CSPs need to be investing in cloud-native technologies to support this transformation.

Read this report to learn more about the role for cloud in data analytics:
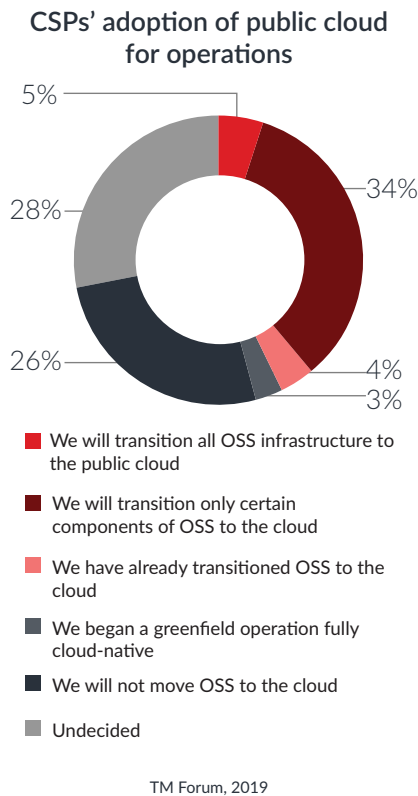
HOW TO LEVERAGE DATA ANALYTICS

Uncertainty about OSS

There is much less certainty about whether OSS workloads are suitable for cloud, despite previously high expectations. In 2018, TM Forum conducted a survey of CTIOs which revealed that half of CSPs had moved less than 5% of OSS infrastructure to the cloud, but they expected this percentage to increase to 50% within three years.

While their prediction has two more years to come true, operators do not appear to be on a trajectory to realize it. Only 4% of the 99 CSP respondents surveyed for our November 2019 report *Future OSS: Towards an open digital architecture* said their companies have already

transitioned some OSS applications to a public cloud (mostly service assurance applications), and only 5% are planning a full OSS stack transformation (see graphic). A third said they will move select applications, but a quarter do not plan to move OSS to the public cloud – ever.

CSPs' adoption of public cloud for operations

- We will transition all OSS infrastructure to the public cloud
- We will transition only certain components of OSS to the cloud
- We have already transitioned OSS to the cloud
- We began a greenfield operation fully cloud-native
- We will not move OSS to the cloud
- Undecided

TM Forum, 2019

Read this report to learn more about the future of OSS:

FUTURE OSS: TOWARDS AN OPEN DIGITAL ARCHITECTURE

FUTURE OSS
Towards an open digital architecture

Other research reinforces TM Forum's findings. A 2019 report from Ovum notes that only 20% of core telecoms systems are virtualized, and only a fraction of them are in the public cloud. Ovum's analysis suggests that CSPs are missing a big opportunity to evolve but still predicts that telcos will be among the fastest-growing adopters of public cloud this year and next. In fact, Ovum has predicted that the number of CSPs planning to move workloads such as revenue management and network management to the public cloud will double by the end of the year. At press time, their final numbers had not been published.

Pros & cons

The CSPs we surveyed and interviewed for this report were able to cite more reasons for deploying workloads in public clouds than against. However, one negative reason could be enough to deter adoption if an operator considers it an insurmountable problem.

The word clouds on pages 6 and 7 show the pros and cons. Larger words signify a greater number of people choosing it as a benefit or challenge. Curiously, some reasons appear in arguments for and against, such as reliability and availability. This indicates that while some operators are convinced of cloud's reliability, others are decidedly not. Additional analysis of the pros and cons follows the word clouds.

Cloud benefits

Time to market   Agility   Risk control   Reliability/availability

Operational performance   Process management

Escalation   Cloud BSS integration   Modern architecture   Greenfield

Automation   Regulatory concerns   Security   Reuse

Flexibility   Trial open source   Enable innovation   Distributed architecture

CapEx/OpEx reduction   Maintenance   Hardware failure

Easier upgrade   Service awareness   Monetization

Adaptive integration   Scalability

TCO   On-demand   Responsiveness   Lack expertise   DevOps   Data privacy

Resource optimization   Predictive maintenance   Standardization

TM Forum, 2019

Scalability – the ability to scale applications economically is important to respondents, but even more critical is ability to scale on demand. This lies at the heart of CSPs' efforts to automate operations and optimize networks for real-time response and reporting.

Total cost of ownership – Rakuten CTO Tareq Amin claims that operating a cloud environment costs 45% less than running a traditional network (see page 20). Time will tell if this is true, but the assumption is that moving applications to the cloud lowers operating expenses. The amount of savings likely will depend on the level of managed services provided and the complexity and size of the network.

Time to market – despite years of process automation and improvement, CSPs still are unable to create services as quickly as internet-based companies. They believe cloud will help level the playing field by speeding service design and testing.

Automation – enabling automation is the top attribute CSPs are looking for in next-generation OSS and BSS solutions. If the early deployments of BSS in the public cloud succeed, CSPs may be more likely to move OSS to the cloud.

Predictive maintenance – CSPs would love to reduce the time it takes to perform upgrades and maintenance, and they believe cloud deployments, particularly managed services, could help by allowing them to perform maintenance tasks without taking components out of service.

Lack expertise – CSPs realize they don't have the necessary software skills for cloud and are ready to let the experts handle the unfamiliar language, components, technologies and culture of cloud networking.

Cloud drawbacks

GDPR compliance

Data security

Customers' concerns Unstable

Proprietary platform   Not scalable Investment costs

Inflexible   Lacks global availability   No quality assurance

Physical security

Customer data protection

Performance   Unreliable   Complex management   Costly

Prefer private   Too risky   Lack of control

Can't customize

TM Forum, 2019

Physical security – this concern is mostly about unauthorized access to servers and systems in the cloud. It has been a difficult barrier for cloud providers to overcome, but they are taking steps to improve security. Google, for example, has adopted stringent identity and access management restrictions. More detail is available in the company's security whitepaper.

Investment costs – respondents cite somewhat ancillary costs in this concern, referring to the cost of changing providers and moving data that has been stored with one cloud provider to another, sometimes called egress. Also, transactional costs associated with continually accessing data can add up, and CSPs worry about the cost of managing analytics.

Lack of control – similar to physical security concerns, CSPs are not comfortable giving up operational control during times of outage, performance degradation, feature rollout, etc.

Unreliable – a 2014 report from the International Working Group on Cloud Computing Resiliency (IWGCR) about global cloud availability showed an average of 99.1% availability. Although it's an outdated figure, it is still often cited as the norm for cloud-based systems and is a far cry from the 'five nines' reliability CSPs are used to.

GDPR & data security – when it comes to customers' privacy and data security, CSPs have a good reputation to protect. They also are heavily regulated and do not feel comfortable trusting data stored in the cloud. This is especially true in countries with strict data residency rules such as the EU's General Data Protection Regulation (GDPR).

Need for change

In the end, it may not be that CSPs have changed their minds about cloud, but rather that the weight of the negative arguments is lessening. The need for change is undeniably urgent, and operators are realizing they cannot transform on their own.

Adoption of public cloud by prominent global operators like AT&T (see page 16) is convincing others to give public cloud a chance. Their minds are slightly more at ease about going in the direction they know is inevitable if they are ever to realize their goals of cost efficiency, agility and being a leading player in the enterprise.

In the next sections, we'll look at what it means to be cloud native and why some CSP leaders are betting on the technology.

Section 2

Go native or go home

Moving a handful of applications to the cloud is one thing, but operating as a cloud-native service provider is an entirely different objective. Eventually, all communications service providers (CSPs) will become cloud-native operators, but for now they are relying on their suppliers to pave the way with cloud-based network functions and support applications. In some cases, they are giving vendors an ultimatum: Go cloud native or risk losing our business.

For now, most suppliers are simply "lifting and shifting" workloads to the cloud, meaning they are porting traditional software systems to a hosted cloud environment. Some vendors are hosting their applications with public cloud providers like Amazon Web Services, Google or Microsoft, while others host them in their own private clouds or in CSPs' private clouds. But only a few suppliers are delivering applications designed from the outset to operate in a cloud environment.

Vendors' reluctance

Our *Future OSS* survey reveals that suppliers have a long way to go to deliver cloud-native apps. Only 7% of supplier respondents said they plan to move their entire OSS stack to the cloud, and nearly a third have no plans at all to offer cloud-based solutions (see graphic opposite).

The lack of support for cloud-based solutions is less about indecision than it is about reluctance. As IBM's Bill Lambertson, Director of Cloud for Telco, Media and Entertainment at IBM explains:

> " Vendors are not in a hurry. They made a big investment in these CapEx technologies."

Suppliers' plans to move OSS solutions to the cloud

7%

13%

30%

20%

30%
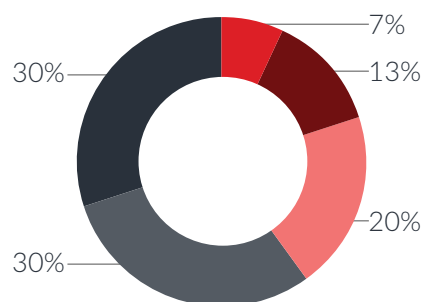
■ We plan to migrate all OSS solutions to the cloud

■ We plan to migrate certain components of OSS solutions to the cloud

■ We have already transitioned OSS solutions to the cloud

■ Undecided

■ We do not plan to transition OSS solutions to the cloud

TM Forum, 2019

What is cloud native?

Cloud native is a term used to describe container-based environments. When applications are "containerized", they are bundled with all the related configuration files, libraries and dependencies required to run efficiently and reliably in different computing environments.

Containers isolate an application and its dependencies into a self-contained unit that can run anywhere. Hardware and operating systems are virtualized, which means the same operating system is shared with other hosted applications. But this does not make them cloud native. To be cloud native, solutions must be broken into functional blocks and run as microservices within containers on elastic infrastructure using Agile DevOps processes and continuous delivery workflows.

In a traditional IT environment, operations teams manage the allocation of infrastructure resources to applications manually. In a cloud-native environment, applications are deployed on infrastructure that abstracts the underlying compute, storage and networking functions. Cloud-native platforms handle tasks such as scheduling of workloads, scaling and load balancing, debugging, and authorization through software tools that manage containers.

Reducing costs

More broadly, being cloud native relates to achieving the promised cost savings that result from running workloads in a cloud, says Eric Troup, Microsoft's CTO, Worldwide Communications and Media Industries. In a cloud-native environment, all resources necessary to run any workload or application are software defined and load balanced and can expand dynamically to support workloads as they scale up or down.

> " Cloud-native software must be developed in a way that lets the cloud work the way it was intended to work, so that it can take advantage of the technology and economics the cloud offers," Troup explains.

The lift-and-shift approach is usually meant to be a temporary fix and cannot take full advantage of a cloud-native architecture. CSPs eventually will phase out these applications, or they will have to be redesigned to be cloud native. In the interim, Troup says, they equate to driving a 50-passenger bus with a single passenger in it, which is inefficient.

"You can transport that one individual at a much lower cost if you can instantly and dynamically downsize the bus and then size it appropriately if the number of passengers increases," he explains.

It's an evolution

For most CSPs and suppliers, the lift-and-shift approach is a stepping stone. It is as much a part of the planning process as re-architecting solutions to be cloud native. And some applications

may never move to the cloud. Making the determination depends on several factors, including:

- Age or remaining shelf-life of a system

- Functional issues such as meeting latency requirements

- Rules and regulations that prohibit storing an application's data in a cloud environment

- Cost and other technical constraints of redesigning

Quick wins

Nevertheless, the lift-and-shift approach can deliver quick benefits, according to George Glass, TM Forum's VP of Architecture & APIs and former Chief Systems Architect at BT. It can help operators tackle "technical debt" by removing dependence on legacy infrastructure, for example.

> " The lift-and-shift approach does not leverage all public cloud's advantages," Glass explains. "On the other hand, completely rewriting an application and transforming it to be cloud-native typically takes a long time and requires significant investment."

He adds: "CSPs need a middle ground that allows for gradual migration in a controlled manner to a cloud-native architecture. The TM Forum Open Digital Architecture (see page 10) provides a framework for this middle ground as well as a full migration path forward in the public cloud."

Read Glass' two-part series on why CSPs should embrace cloud:

Cloud agnostic?

There are differing opinions about whether being cloud native means that a company can put a workload into any public cloud or multiple public clouds without modification. A similar question plagues CSPs relying on systems integrators to run their back-office environments. Suppliers must choose which integrator to team up with and which middleware platform to write application program interfaces (APIs) for, thereby limiting their ability to address the full market.

The difference is that containers improve the onboarding process by reducing the need for integration as proprietary components are secured inside the container. MYCOM OSI announced its cloud-native Experience Assurance and Analytics (EAA) solution in early 2018, leveraging the Red Hat OpenShift Container Platform. This cut the time it takes to deploy and operate the system by over 75%. MYCOM OSI then partnered with AWS to launch its public-cloud based Assurance Cloud SaaS offering in early 2019.

The process of porting applications to different cloud environments is not effortless, however. Each environment has its own way of operating, which means that a workload optimized for a Microsoft Azure environment may not be optimized for Google Cloud. This typically leads to a hybrid cloud deployment because suppliers do not want to build for two different public cloud platforms when they could use their own private cloud for redundancy with less development work.

Using the Open Digital Architecture to transition to cloud

Despite a steady upswing in public cloud deployments, it isn't a sure bet for any company, as is evidenced by analyst firm IDC's finding that a full 81% of companies "repatriated" in 2018, which means they moved their public cloud deployments back into private or hybrid cloud environments. For CSPs, making the leap to public cloud will be an evolutionary process, and the TM Forum Open Digital Architecture (ODA) can help them get there.

TM Forum Open Digital Architecture: a work in progress

People

Organizations

Things

Engagement management

DECOUPLING & INTEGRATION

Party management

DECOUPLING & INTEGRATION

Core commerce management

DECOUPLING & INTEGRATION

Production

DECOUPLING & INTEGRATION

Intelligence management

TM Forum, 2018

"CSPs will embrace public cloud because it's much less expensive, more scalable and easier to manage," says George Glass, TM Forum's VP of architecture & APIs. "The big question for most operators is how to embrace cloud-native applications and public cloud."

As noted, CSPs basically have two options: lift and shift existing applications to the cloud or re-architect them. The former short-changes operators on the benefits of cloud, but the latter takes too much time and is costly. The ODA provides an evolutionary path between the two.

Part of the Open Digital Framework, (see page 34) ODA is fundamentally designed as a component-based architecture, with the business services of a component exposed as a set of Open APIs. The APIs can be, and typically are, further decomposed into a set of services and microservices. The advantage of using microservices is that they can
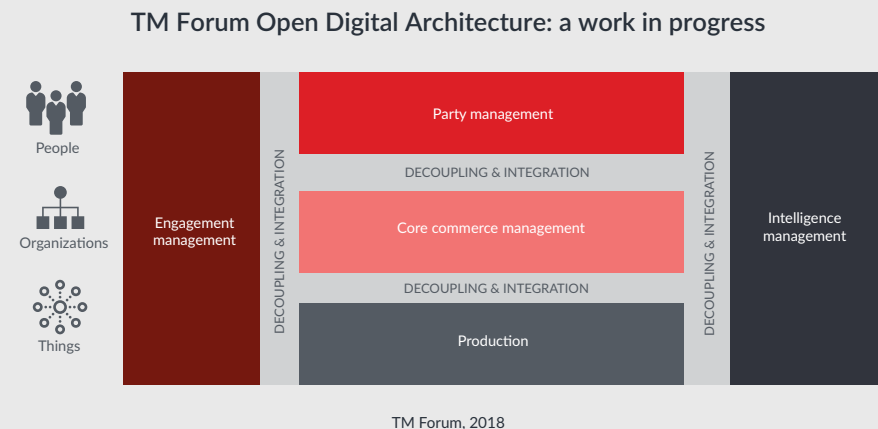
be managed on scalable infrastructure using Agile development practices.

"CSPs and other organizations can use ODA to set an architectural vision and plan relevant roadmaps to implement it," Glass explains. "The transformation guides that are part of the ODA enable organizations to work backwards from their target architecture."

For example, companies can use optimization, re-engineering or abstraction to manage migration from their current hosting infrastructure, which is usually a mix of dedicated hardware, virtualized infrastructure, private cloud deployments and some public cloud deployments, to a truly cloud-based environment.

"Controlling the migration allows CSPs to change the architecture of their IT estate gradually, moving relevant components to the cloud only when the business is ready," Glass says, adding that ODA

accepts that not everything has to be cloud-native and that simply cloud-enabling some legacy components may be appropriate to manage costs. In addition, re-engineering or abstraction allows for the decomposition of a component into services and microservices enabling a company to move to a cloud-native deployment.

"This supports an Open Digital Framework concept called 'pace layering', where components that support innovation and differentiation can be identified and abstracted from legacy infrastructure using Open APIs," Glass explains. "They are then rewritten using DevSecOps techniques to provide continuous integration and continuous development (CI/CD). This helps CSPs deliver new services rapidly and maximize return on investment."

If you'd like to learn more about the ODA project, contact George Glass.

What about NFV?

With all the excitement about cloud-native architectures, it's easy to forget about the fervor that surrounded the 2012 ETSI-sponsored whitepaper that 13 of the world's largest telcos authored extolling the virtues of network functions virtualization (NFV). So far, NFV, which was supposed to be transformative on its own and act as a catalyst for 5G, has not delivered.

"Virtual network functions [VNFs] have been around a while but haven't really taken off," Adrian Comley, General Manager of Dynamic Network Services at BT said during a panel discussion at SDN NFV World Congress in October 2019. "There are no global deployments of service chains of VNFs at scale anywhere."

But the investment CSPs have already made in NFV could stall investments in cloud, because operators have a limited amount to spend and are unlikely to abandon investments in VNFs when they are still counting on it to support 5G. Plus, there isn't a lot of incentive make VNFs cloud native if they are effective in supporting services.

Comley was part of a panel discussion about moving from NFV to cloud-native architectures which also included representatives from Orange Spain, Verizon, HPE and Intel. Vickie Lonker, Verizon's VP of Product Management and Development, said she struggles with cloud-native technology because of its limitations.

"We are squarely grounded in a place where NFV infrastructure is where we are deploying network functions," she said. "The things I care about, like LANs and WANs and interfaces are not 'native' in cloud native. Until those concepts become something we can translate between, it will be tough to adopt cloud as a greenfield for us. It will be a hybrid world we will have to manage for quite a long time."

Lonker added that it is difficult to build cloud-native applications that also work in a network because the architectures are very different.

The goal should be to move away from traditional VNFs, which were developed to disaggregate hardware from software, and break them into microservices so that core components from them can discover and build on each other, Comley added. But, he admitted, "we are not there yet."

Watch the SDN NFV World Congress panel discussion:

Edge needs cloud

One area of the network that most panelists agreed could benefit from a cloud-native architecture is the mobile edge. Most multi-access edge computing (MEC) functions by default will be cloud native, so they are a good place to start, according to Lonker.

"They will be different types of workloads than we deploy in current universal CPE [customer premises equipment] applications today," she said.

Multi-access edge computing was conceived as a cloud solution in part because it is a distributed environment that has closer proximity to end users which helps when latency is an issue. For example, Verizon's 5G Edge service combines its 5G Ultra-Wideband network with AWS' cloud services in order to partner with enterprise customers to develop large-scale, latency-sensitive applications at the edge.

Domenico Convertino, VP of Product Management at HPE, agreed, saying that cloud-native MEC is deployable today.

> " Whatever cloud-native services are running at the edge have a completely different set of economics than those running in an NFV environment," he explained. "So, it is technology and business that is driving this innovation."

In the next section, we'll compare the leading cloud providers, focusing on their hybrid and multicloud strategies.

Section 3
Who are potential cloud partners for CSPs?

When it comes to partnering with cloud platform providers, communications service providers (CSPs) have several options depending on whether they want to pursue public, private, hybrid or multicloud deployments. In this section we compare several cloud providers' strategies, with an emphasis on public cloud providers.

Enterprises are partnering with public cloud providers because the economics are too hard to ignore, particularly when technology changes are far outpacing their internal IT teams' abilities to manage them. CSPs have been more hesitant, however, because they have concerns about reliability, security and a giving up control. While they are still unlikely to put network workloads in a public cloud, they are beginning to move BSS workloads, especially revenue management, billing and charging, and customer management to the cloud (see Section 4).

Public cloud providers have been working to harden their facilities and improve reliability, and they are expanding their strategies to include hybrid cloud management of workloads across different cloud and network environments. Like CSPs, hyperscale cloud providers have been reluctant to share ownership of the customer, which is essentially what happens in a hybrid cloud, but this portion of their businesses is growing.

Market leaders

A recent report about the public cloud market from Synergy Research Group shows that during Q3 2019 the top four public clouds – Amazon Web Services (AWS), Alibaba Cloud, Google Cloud and Microsoft Azure – accounted for 72% of the worldwide market for infrastructure and platform as a service, up from 57% at the beginning of 2016.

Amazon held its 40% share of the nearly $90 billion business, while the other three providers showed steady gains. IBM stands as the leader in hosted and managed private cloud services, according to the research firm.
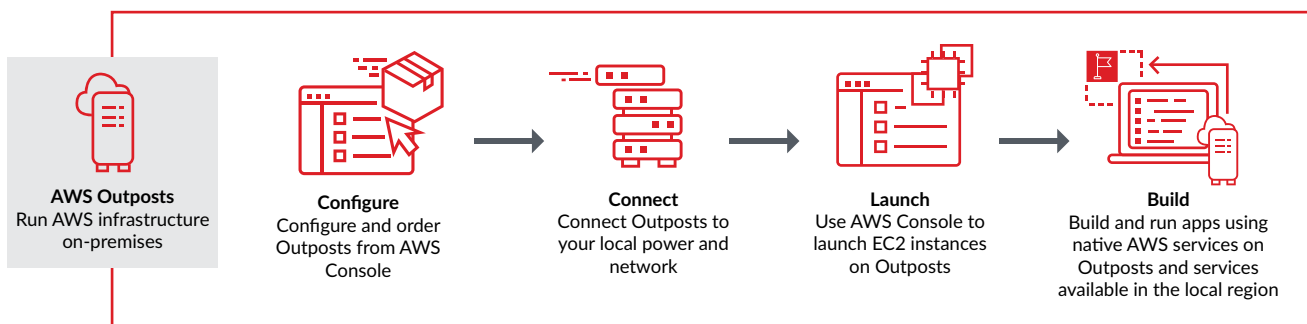
The table below lists the major cloud providers partnering with CSPs. Further analysis of AWS, Google Cloud, Microsoft Azure and IBM Cloud follows, with a focus on their hybrid or multicloud strategies.

Cloud providers at a glance

Cloud provider	Primary focus	Recent developments
aws	Public cloud	Offers AWS Outposts, a managed service through which users get AWS-configured hardware and software delivered to their on-premise data center or co-location space to run applications in a cloud-native manner, without having to run it in an AWS data center
Google Cloud	Public cloud	Google Cloud's new cloud management platform, Anthos, acts as a management layer that runs on top of the Google Cloud Platform and Google Kubernetes Engine; it allows users to manage their entire hybrid, multicloud environment from a central system
Azure	Public, hybrid and multicloud	Microsoft Arc is a multicloud management layer that extends Azure capabilities to other cloud platforms including AWS and Google
Alibaba Cloud	Public and hybrid cloud	Alibaba bought a major stake in ZTESoft, including OSS/BSS assets, creating the subsidiary now known as Whale Cloud; Whale Cloud provides cloud, big data and AI-enabled solutions to global telecom operators, governments and enterprises in more than 80 countries and regions
IBM Cloud	Private and hybrid cloud	Since acquiring Red Hat, IBM has broadened its reach in cloud computing; together, the companies can substantially move cloud and platform adoption forward, particularly at the intersection of telecom and the enterprise as their paths merge
ORACLE Cloud	Hybrid cloud	Although new and still in development, Oracle's Generation 2 Cloud platform features an Autonomous Database, which automates patching, upgrades and tuning; platform services include application development, business analytics, data management, integration, security, AI and blockchain; one of the first apps created for telecom is the cloud-native Billing and Revenue Management (BRM) solution

TM Forum, 2019

How AWS Outposts works

AWS Outposts
Run AWS infrastructure on-premises

Configure
Configure and order Outposts from AWS Console

Connect
Connect Outposts to your local power and network

Launch
Use AWS Console to launch EC2 instances on Outposts

Build
Build and run apps using native AWS services on Outposts and services available in the local region

TM Forum, 2019 (source: AWS)

Amazon Web Services

Many CSPs are considering partnerships with AWS because of its market share. AWS' answer to operating in and managing a hybrid cloud model is AWS Outposts, which the company announced at its 2018 annual event. Outposts is a managed service through which users get AWS configured hardware and software delivered to their on-premise data center or co-location space to run applications in a cloud-native manner without having to run it in an AWS center. AWS operates and manages the service. The graphic above shows how it works.

This approach isn't really about multicloud, rather it focuses on AWS services running in the AWS cloud or on the customer's premises, using VMware's software-defined data center platform. So far, however, it is still only a concept. AWS said it would officially launch OutPosts in late 2019, but at press time no announcement had been made.

It is unclear whether AWS will support multiple clouds as Microsoft and IBM do, but in an interview published at the beginning of December, AWS CEO Andy Jassy said

that companies using a multicloud strategy typically don't divide the business evenly across providers.

> "
>
> If you do, you have to standardize on the lowest common denominator [in terms of feature set], and these platforms are in pretty radically different spots right now in terms of capabilities and the ecosystem and maturity," Jassy said. "You also end up taking your development teams and asking them to not just make the big shift from on-premises to the cloud, but then force them to be fluent across multiple providers, which development teams hate."

Google Cloud

Google introduced its new cloud management platform, Anthos, in April 2019. Anthos acts as a management layer that runs on top of the Google Cloud Platform and Google Kubernetes Engine (GKE), allowing users to manage their entire hybrid, multicloud environment from a central system.

Anthos supports users running applications and workloads on-premises or in public and hybrid cloud environments, including AWS and Azure. Like Amazon, Google Cloud is also a VMware partner and leverages that partnership for hybrid on-premises deployments.

In November, Google enhanced its hybrid cloud approach with the acquisition of CloudSimple, a startup that simplifies the migration of VMware workloads running on-premises to public cloud environments without the need to re-architect applications. This sounds a lot like lifting and shifting, but the companies say that users can still leverage benefits of cloud like performance, elasticity and integration with other cloud applications.

Dispelling cloud myths

BSS provider Optiva is partnering with Google for its cloud-based BSS. At Digital Transformation World 2019 representatives from the two companies dispelled some of the misperceptions, or myths, around public cloud. The biggest myths include:
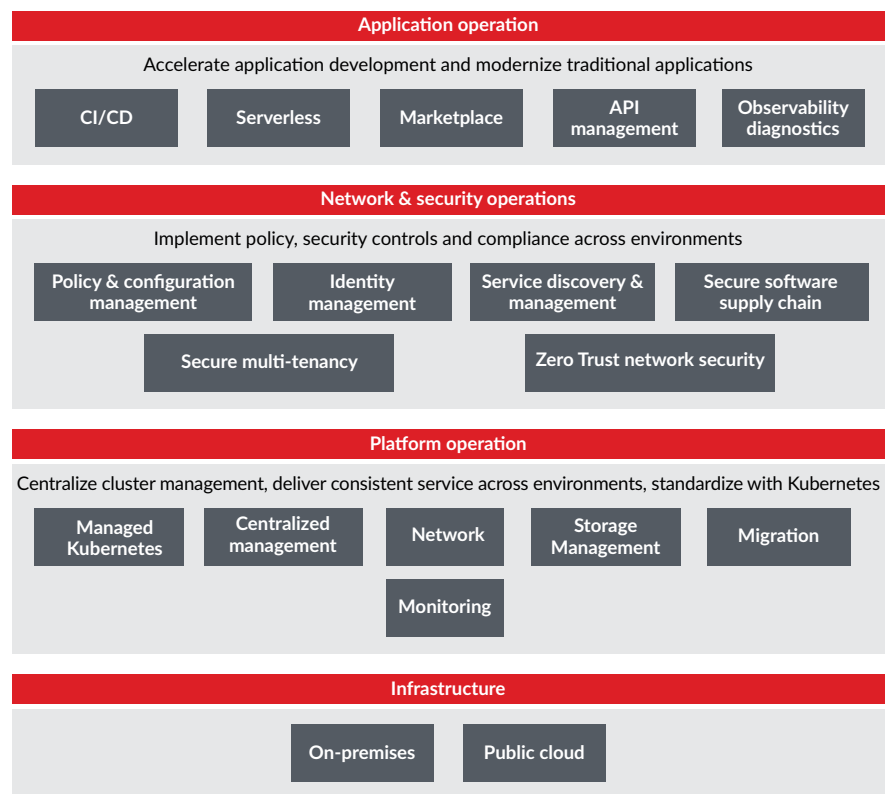
1. Public cloud is not secure

2. Private cloud is the same as public cloud

3. All public clouds are the same

4. Migrating to the cloud is the same as being cloud-native

5. Public cloud is more expensive

Since partnering with Google Cloud in February 2018, Optiva claims it has been able to save a Tier 1 Canadian CSP up to 70% by moving revenue management and charging into Google Cloud. A quick poll conducted during the session revealed that given the new information presented, 64% of the audience found public cloud to be more secure than on-premises solutions, a double-digit improvement from the beginning of the session.

Google Cloud Anthos Platform

Application operation
Accelerate application development and modernize traditional applications

| CI/CD | Serverless | Marketplace | API management | Observability diagnostics |

Network & security operations
Implement policy, security controls and compliance across environments

| Policy & configuration management | Identity management | Service discovery & management | Secure software supply chain |

| Secure multi-tenancy | Zero Trust network security |

Platform operation
Centralize cluster management, deliver consistent service across environments, standardize with Kubernetes

| Managed Kubernetes | Centralized management | Network | Storage Management | Migration |

| Monitoring |

Infrastructure

| On-premises | Public cloud |

TM Forum, 2019 (source: Google)

Google also announced the acquisition of cloud storage company, Elastifile, in July and in June said it would acquire data analytics company Looker for $2.6 billion. At the time of the Elastifile announcement, Deepak Mohan, Research Director at IDC, said that increasingly applications require on-premises level of performance for latency and consistency alongside the scalability benefits of the cloud and that the acquisition of Elastifile will better enable Google Cloud customers to meet this mix of needs as they deploy such workloads to the Google Cloud Platform.

Elastifile was already a Google Cloud partner for Anthos and works to augment GKE On-Prem for deployments in Google Cloud, on-premises and in hybrid cloud configurations. The Looker acquisition continues Google's strategic commitment to multi-cloud by empowering companies to create a cohesive layer built on any cloud database, including Amazon Redshift, Azure SQL, Snowflake, Oracle, Microsoft SQL Server or Teradata, as well as in other public clouds and on-premise data centers. In early November, Google Cloud amped up its hybrid and multicloud sales effort by authorizing 50 partners to start reselling Anthos in early 2020.

Microsoft Azure

Microsoft has been offering tools for hybrid cloud deployments since 2016 when it launched Azure Stack, a portfolio of products that extends Microsoft Azure cloud services and capabilities to a customer's cloud environment of choice by enabling hybrid and edge computing applications to be deployed and run consistently across location boundaries.

In November Microsoft introduced Microsoft Arc, a multicloud management layer that extends Azure capabilities to other cloud platforms including AWS and Google. Arc give CSPs what they have long wanted in their own networks: a single view of all applications and workloads running on any and all platforms. This allows Microsoft to offer public, hybrid and now multicloud services to CSPs while also helping them deliver the same to their large enterprise customers.

Advances in the Azure Arc Resource Manager, Azure's deployment and management service, give users a central, unified approach to manage Windows and Linux Servers, Kubernetes clusters and Azure data services at scale across on-premises, multicloud and edge. At its recent annual Microsoft Ignite event, CEO Satya Nadella said the company believes the hybrid model is a permanent, not temporary, state.

Microsoft Azure Arc enables a hybrid model
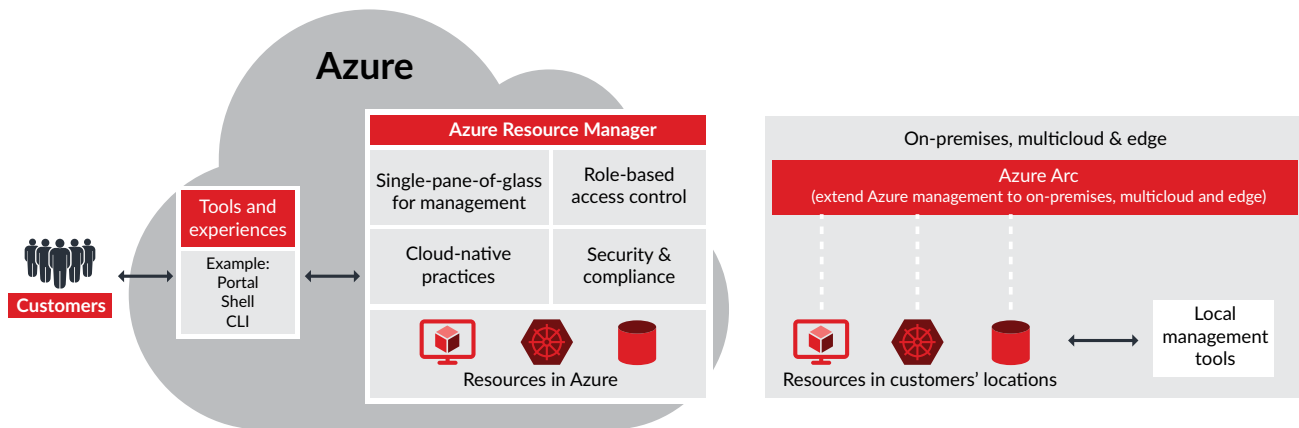
TM Forum, 2019 (source: Microsoft)

Indeed, Eric Troup, Microsoft's CTO, Worldwide Communications and Media Industries, says the discussion among CSPs around hybrid cloud has moved very quickly, causing the distinction between public and private cloud to fade, while at the same time increasing their level of confidence in the overall cloud solution. Azure Arc expands the control plane for policy management across multiple cloud platforms including private data centers and CSPs are more comfortable with that, he says.

Developing the ability to work across environments is a challenge CSPs understand well. The moves by the top three hyperscale cloud providers to support multicloud environments is likely to give operators' more confidence about moving applications to the cloud since it gives them a choice about where to put their most sensitive workloads.

IBM Cloud

IBM has a different approach than the big three public cloud providers in that its primary role often is to act as

an enabler, helping to accelerate clients' journeys to the public cloud, whether it be with AWS, Google Cloud, Microsoft Azure or any combination of them. IBM provides multicloud management as well as skills needed for transformation to cloud-based solutions. In addition, IBM can help CSPs deploy artificial intelligence (AI) and analytics, which gives the company an advantage over some of the competition.

IBM compartmentalizes its cloud efforts in the telecom space around different workloads: IT, OSS/BSS, network and business-to-business. According to Bill Lambertson, IBM's Director of Cloud for Telco, Media and Entertainment, CSPs are making significant progress in using public cloud for corporate IT applications (for example, human resources).

"As long as there is a public cloud property within the data sovereignty area, it makes sense to leverage it whenever you can," he says.

Lambertson also sees a lot of interest from CSPs in moving the BSS stack to public cloud. He explains:

> "They know they need to move to a fully cloud-native environment, but we are seeing signs that there will be a coexistence of virtual machines and cloud-native functions for at least another three to five years."

IBM is helping operators establish their cloud-native architectures using microservices, Kubernetes and Red Hat's Open Shift platform. Using Kubernetes, CSPs can securely deploy, run and manage their data and applications on whatever cloud or combination of clouds they choose to use. All of IBM Cloud, including IBM Watson and IBM Blockchain Platform, runs on Kubernetes, enabling massive scale and workload diversity.

In the next section, we'll look at how some leading CSPs are partnering with cloud providers.

Section 4

Trailblazers begin adopting cloud

While many communications service providers (CSPs) are moving slowly to embrace cloud, some like AT&T, China Unicom, Dish Network, Telefónica, Three UK, TIM, Vodafone and Japanese startup Rakuten are leading the way. Indeed, AT&T's announcement in July that it is becoming a 'public cloud first' company is a huge vote of confidence in the technology, which could persuade other operators to follow suit.

The table on page 17 shows some early partnerships between CSPs and cloud providers. It includes public cloud platforms as well as software vendors whose applications are hosted in a cloud environment. AT&T's $2 billion deal with Microsoft is perhaps the most notable, because the telco has said it will move all non-network applications and workloads, including operational and business support systems (OSS/BSS), to the Microsoft Azure platform by 2024.

As part of the multiyear agreement, Microsoft will be AT&T's preferred cloud provider, and the arrangement includes joint development of AI tools and applications to advance 5G. It also includes giving AT&T's 250,000 employees access to Microsoft 365 for productivity tools, and Microsoft will help AT&T design, test and build edge computing capabilities.

Public cloud first

In a joint statement with Microsoft, AT&T said it is becoming a "public cloud first" company: "That initiative will allow AT&T to focus on core network capabilities, accelerate innovation for its customers, and empower its workforce while optimizing costs."

In reality, AT&T and other CSPs will adopt cloud strategies that rely on a combination of public, private, hybrid and multicloud platforms. In fact, AT&T announced another significant cloud alliance with IBM for its AT&T Business Solutions unit at nearly the same time as the Microsoft deal (see panel on page 18), and the telco already has a relationship with AWS, primarily in its marketplace.

Microsoft CEO Satya Nadella with AT&T CEO John Donovan

CSPs' partnerships with cloud providers

CSP	Cloud partner	Type of cloud	Strategy
AT&T	Azure	Public	As part of a $2 billion deal, the companies will collaborate on cloud, AI and 5G. Microsoft will be AT&T's preferred cloud provider, and AT&T will move all its non-network applications and workloads to the Azure platform by 2024, including OSS and BSS
AT&T Business	IBM Cloud	Hybrid	AT&T Business Solutions is moving internal software applications to IBM Cloud and is using Red Hat's open source platform to manage workloads and applications
China unicom 中国联通	Alibaba Cloud	Public	China Unicom partnered with Alibaba Cloud to launch a new access solution in the African market; the combined public cloud and connectivity services help Chinese and other Asian enterprises who want to do business in Africa
dish	To be determined	To be determined	Dish Network has committed to launching a 5G broadband network by the summer of 2023 that could serve 70% of the US population; Founder and Chairman Charlie Ergen says operations will be cloud native and therefore cost 25% less than other 5G network buildouts
Rakuten	Netcracker An NEC Company / NOKIA	Public	Rakuten is developing a greenfield end-to-end, cloud-native architecture for a new 5G network; this includes working with Nokia for the core network and Netcracker for cloud-based OSS/BSS; Rakuten was expected to launch services in October 2019 but now is aiming for March 2020
Telefónica	Google Cloud	Public	Telefónica has expanded a local agreement it had with Google in Spain into a global agreement that includes offering Google Cloud services alongside other public cloud services and managing the workloads and relationships across cloud providers
Three.co.uk	mycomOSI / NOKIA	Private (Core) Public (OSS)	As part of its 5G core network migration, Three UK is using MYCOM OSI's public cloud-based Experience Assurance and Analytics suite, which runs on the AWS cloud platform, to provide service assurance of its Nokia core network, which includes new VNFs and existing PNFs
TIM	Google Cloud	Public	TIM is aiming to become a key player in cloud and edge computing by partnering with Google; the CSP will expand its service portfolio for the Italian enterprise market by offering public, private and hybrid cloud services
TRUPHONE	Optiva	Public	Truphone is deploying the Optiva Online Charging System, which runs on the Google Cloud platform; the company is also using AWS and IBM Cloud as part of other deals
verizon	aws	Public/hybrid	Verizon is deploying AWS Wavelength as part of its 5G edge computing offering to give developers the ability to deploy applications that require ultra-low latency; Verizon 5G Edge provides mobile edge computing and high-volume connectivity between users, devices and applications, while AWS Wavelength supports the parts of an application that require ultra-low latency to the edge of the network while connecting back to the full range of cloud services running in AWS
vodafone	ERICSSON / Optiva	Private (EPC) Public (BSS moving to public)	Vodafone Idea in India is deploying Ericsson's Cloud Packet Core on a private cloud and the Optiva Online Charging Engine; the telco is trialing the Optiva solution on a Kubernetes-enabled private cloud with an option to move to public cloud (Google Cloud)
vodafone	amdocs	Hybrid	Vodafone Germany has deployed Amdocs' end-to-end digital cloud-native and open platform, which runs on the X platform and includes customer care, commerce and a dynamic catalog
vodafone business	IBM Cloud	Hybrid	The companies formed a $550 million venture that provides optimization and management services from IBM Multicloud Manager as part of Vodafone's Business Multicloud Platform to large multinational corporations using existing public cloud offerings from AWS, Google Cloud, Microsoft Azure and Alibaba Cloud; it also leverages the private IBM cloud

TM Forum, 2019

Staying competitive

AT&T's cloud deals are important because they signal a willingness to partner to save money and increase agility. This is the only way CSPs can remain competitive.

"Every carrier on the planet has the same issue and is heading down the same path," says Eric Troup, Microsoft's CTO, Worldwide Communications and Media Industries. "They need to cut costs, be much more agile than they currently are and focus their investment where the revenue really is.

> 66
>
> AT&T bit the bullet early and made a huge commitment," Troup adds. "Other carriers are taking longer to make that decision, but the economics are so fundamental that everybody will be driven in this direction over the next five years."

TM Forum's George Glass, VP of Architecture & APIs agrees: "This is a big step for AT&T, one that all telcos eventually must take. Embracing public cloud is not a question of 'If?', but 'When'?

"As CSPs build trust with public cloud providers, they will move applications to the public cloud, eliminating the need to build and run data centers," Glass adds. "They will embrace public cloud because it's much less expensive, more scalable and easier to manage."

AT&T & IBM target enterprises

One day before AT&T announced its partnership with Microsoft, AT&T Business Solutions unveiled a cloud initiative with IBM aimed at innovating and collaborating in the enterprise market. The deal is a strategic alliance primarily focusing on private and hybrid cloud solutions.

IBM will be the primary developer and cloud provider for AT&T Business Solutions' operational applications and will help manage the AT&T Communications IT infrastructure, on- and off-premises and across different clouds – private and public. AT&T Business Solutions will also utilize Red Hat's open source platform to manage workloads and applications.

The end goal is to better serve enterprise customers. For example, AT&T and IBM will collaborate on multi-access edge computing services to help enterprise clients capitalize on 5G network speeds and internet-connected devices and sensors at the edge of the network. It will also allow for the rapid transport of data to and from multiple clouds and billions of edge devices with increased reliability and security, addressing any latency concerns in the process.

"There is a learning curve that operators need to go through on their own before they realize that being in the enterprise space for voice and connectivity doesn't mean they can automatically have the skills and knowledge to be in the enterprise space for cloud and IT," says Bill Lambertson, Global Director of Cloud for Telecom, Media and Entertainment at IBM. "Keeping up with the changes in technology, from IaaS [infrastructure as a service] to PaaS [platform as a service] to Kubernetes to work coming out of the Cloud Foundry and serverless technology, makes it challenging to maintain the right skill sets. Many have recognized they need a partner to help them on their cloud journey."

Watch AT&T CEO John Donovan and IBM CEO Ginni Rometty discuss their companies' cloud partnership:

Vodafone tries it all

At Vodafone, all cloud options are on the table, as is evidenced by the company's widespread use of cloud platforms across its operating companies. In India, for example, Vodafone Idea is deploying Ericsson's Cloud Packet Core and Optiva's

Online Charging Engine, which is hosted on Google Cloud (see panel on page 19). Vodafone Germany has deployed Amdocs' cloud-native customer care and commerce solutions, and in late November Vodafone Egypt became the country's first cloud service provider to reach VMware Cloud verified status.

Vodafone Business has teamed up with IBM as part of a $550 million deal to form a venture that uses public, private, hybrid and multicloud platforms to deliver new digital services to its enterprise customers. The partnership falls just short of being a legal joint venture, according to Santi Ribas, Head of Cloud & Security Offers and Innovation at Vodafone Business. He explains:

"We wanted to avoid a loose partnership and make sure everything worked together."

Toward this end, the team works in a joint office where co-leaders from each company, including Ribas, develop joint strategies, make decisions about roadmaps and investment, and provide the structure necessary to support operations and services.

The venture delivers workload optimization and management services from IBM Multicloud Manager as part of Vodafone's Business Multicloud Platform to large multinational corporations using existing public cloud offerings from AWS, Google Cloud, Microsoft Azure and Alibaba Cloud. It also leverages the private IBM cloud.

While the public hyperscale cloud providers do a great job of deploying and managing availability of workloads and providing security, Ribas contends that together Vodafone and IBM take managing the lifecycle of public cloud workloads to a new level. Through IBM tools and professional services, the venture provides a single management platform that:

- Provides a cost assessment of workloads prior to deployment

- Manages workloads across multiple clouds

- Conducts planning within the structure of customers' budgeting processes

Partnering for cloud BSS

Three years ago, Optiva pledged to invest $100 million to re-architect its product line to be cloud native and operate in public cloud environments. Now Vodafone Idea and UK-based Truphone, a mobile virtual network operator (MVNO), are using the company's Online Charging Engine, which runs on the Google Cloud platform, to reduce costs and add new capabilities.

"[Public cloud] is a game-changing technology that has not yet been adopted widely enough by operators, and the hesitancy does not work to their advantage," says Shay Assaraf, Chief Marketing Officer at Optiva. "If cloud gives you the ability to run containers, do it. If it allows you to do continuous deployment and integration, do it. Don't bring your proprietary software to the cloud. You won't see the benefits."

For Vodafone Idea, reducing costs is critical because of the volatility in the Indian market. The company, which resulted from a merger between Vodafone India and Idea Cellular, is the third-largest mobile operator in India, behind startup Reliance Jio and Bharti Airtel.

Vodafone Idea is using the Optiva Charging Engine in a private

Vodafone cloud using Kubernetes. As part of a two-phased strategy, the operator will move the BSS application onto the public cloud platform pending a successful pilot program. The company also signed a multimillion-dollar deal with IBM to manage hybrid and multicloud environments, and to apply analytics and AI-based security in its IT environment.

Truphone is also using the Optiva Charging Engine to launch an MVNO service. Truphone is a connectivity and eSIM provider that operates its own network infrastructure with points of presence in London, New York, Amsterdam, Sydney, California and Hong Kong and partners with other operators to create a global network.

The operator is taking an all-options approach to cloud deployment. Besides partnering with Optiva, Truphone is working with Soracom, a global provider of smart connectivity for the internet of things, whose remote SIM provisioning platform is on the AWS cloud. In addition, it has announced a partnership with Movius, a provider of cloud-based secure mobile communications software operating on the IBM Cloud.

- Manages and optimizes network and service availability

- Provide an analytics layer comprised of AI and machine leaning to leverage data from all cloud and network domains

- Manages assets against customers' security policies

Ribas says the goal of the joint venture is to prepare customers for the digital world and help them discover what they can do with the cloud. The team will co-develop new services that combine

Vodafone's position in IoT, 5G and edge computing with IBM's multi-cloud and professional services. Vodafone also uses IBM's Multicloud Management platform to set up storefronts for enterprises across public cloud providers.

Watch Vodafone and IBM executives discuss their partnership:

Partnering for cloud OSS

CSPs Three UK and Globe Telecom in the Philippines are partnering with MYCOM OSI for cloud-based service assurance that runs in the AWS public cloud. MYCOM OSI launched its Assurance Cloud platform in February 2019 after announcing a formal partnership with AWS a month earlier, but the company has been working since 2017 to make its service assurance solution cloud native by using containers and microservices.

Three UK selected the company as a partner in 2017 for a transformation project to develop a 5G distributed cloud network. The operator also partnered with Nokia, which managed the project and is providing cloud management and many of the core network elements. The network is now serving 9 million customers, nearly all of whom are using data services and about 65% of whom are also using voice.

Three UK's transformation goals were to migrate to public cloud, increase agility, improve customer experience and double Net Promoter Score (NPS) by 2021. To do this, the company requires a service assurance solution capable of delivering a combined view of the network and services, plus end-to-end management and closed loop automation.

Benefits of cloud native

By making its service assurance application cloud native, MYCOM OSI has been able to reduce the time it takes for CSPs to deploy it from days to an hour. Other benefits include on-demand scaling and feature activation, and better interoperability with network cloud deployments. In addition, the partnership with AWS enables the supplier to leverage the cloud platform's analytics tools, which end users can use to build and manage data lakes.

Globe Telecom began working with MYCOM OSI in 2016 when the company was looking for a converged network performance management solution to replace multiple siloed systems and manage network and service quality across all network domains, technologies and equipment vendors. Now the company is running MYCOM OSI's cloud native-assurance application in the AWS cloud.

The move has delivered tangible benefits in a very short time period, according to Niles Howsare, OSS Advisor, Globe Telecom. The biggest overall benefit is the cost efficiency of

managing network performance during a time of rapid growth in terms of both users and traffic, he says.

Globe Telecom can now deliver a unified data set that supports proactive resource assurance across networks through a single view. The operator also can support virtualized network technology in real time and quickly scale on demand to expand application processing power.

Rakuten rolls its own

Japanese ecommerce giant Rakuten is developing its own cloud-native 5G network and back-office stack with help from Nokia and Netcracker. Rakuten CTO Tareq Amin says his company can build a 5G network for 30% of the cost of a traditional one, with a total cost of ownership that is 45% lower.

Rakuten had planned to launch 5G services in October, but the company has pushed its start date to spring of 2020. According to Amin, the company was able to build a virtual radio access network (vRAN) in eight months, and it took less time to implement Netcracker's digital BSS, which will reduce the time it takes to onboard new customers from 3 hours to 10 minutes. The graphic below shows the company's architecture.

Rakuten's architecture
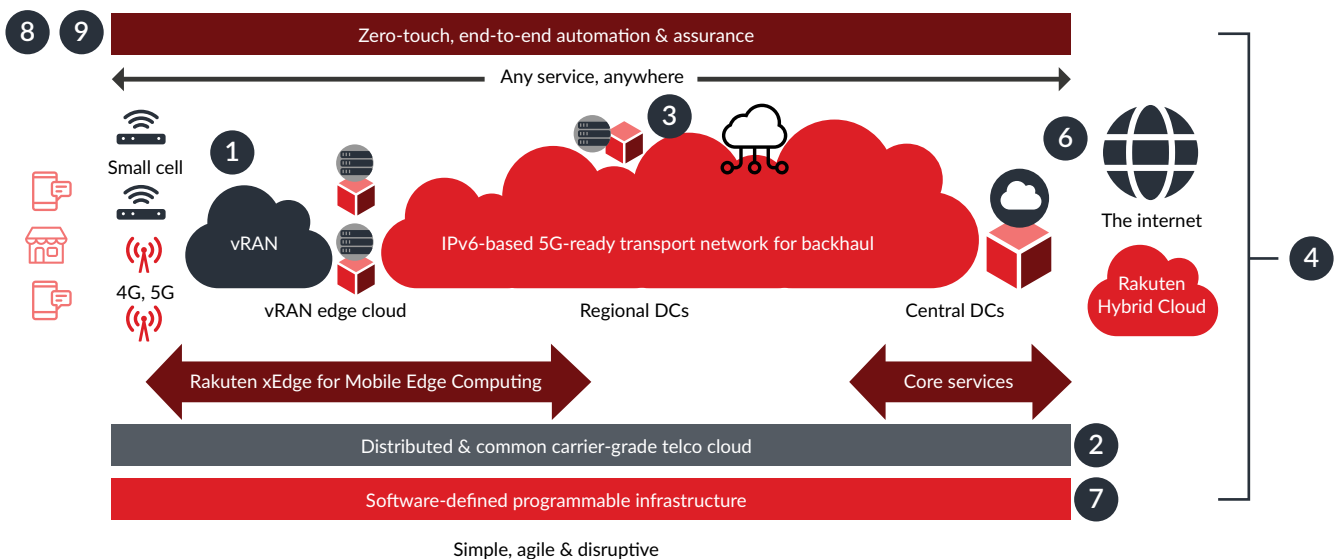
TM Forum, 2019 (source: Rakuten)

Amin says that using virtual network functions and cloud-native BSS also will significantly reduce the time it takes to perform upgrades and maintenance because Rakuten will be able to apply fixes and perform upgrades and maintenance without taking services down.

As an example, he points to Netflix, which performs thousands of software upgrades a day without maintenance windows. CSPs on the other hand must schedule network maintenance and upgrades during low-peak traffic times to minimize impact.

Amin's bold plans garnered him recognition as FierceWireless' most powerful person in wireless technology in 2019. More importantly, Rakuten has drawn the attention of the telecoms industry to the potential benefits of cloud-native technology and virtualization.

However, it remains to be seen whether the company will be successful. In addition to postponing its rollout, Rakuten reported a significantly higher operating loss for its third quarter as a result of increased investment in the mobile network and lower revenue targets due to the delay.

Watch an interview with Rakuten's Amin filmed at Digital Transformation World 2019:

TAREQ AMIN
CTO
Rakuten

Dish aims high

Dish Networks plans to use public cloud in its gutsy push to become the next nationwide 5G mobile operator in the US. While some observers have called the idea a fantasy, others say it could be a masterstroke.

Dish Founder and Chairman Charlie Ergen first floated the idea in May 2018. The company has since committed to launching a 5G broadband network by June 2023 that could serve 70% of the US population. Dish is set to acquire Sprint's prepaid businesses and customers including Boost Mobile and Virgin Mobile for $5 billion following completion of the merger of Sprint and T-Mobile, but it will still have to build a nationwide network and clear many regulatory hurdles.

Given the tight time frame, the only way to achieve this is through partnerships, virtualization and cloud networking. Dish is looking to build a fully cloud-native operation and is

looking at public and private cloud options for its 5G core, vRAN, mobile edge, local and market data centers, and OSS/BSS.

Like Rakuten's Amin, Ergen believes that going cloud-native will help the company get to market much more quickly, and he believes that building a cloud-native network will cost 25% less than a traditional network.

If Dish is successful, "it will provide a really good example of where to look for inspiration for what it looks like to move to a fully cloud-native operating model that supports both container-based and VNF-based virtualization technologies," comments IBM's Lambertson.

Understandably, comparisons to Rakuten have been made in terms of Dish being a risk-taking innovator. Time will tell if the two companies can succeed in disrupting the 5G market. One thing is clear, however: CSPs and their partners need to be specific when discussing what it means to be cloud native and consider options beyond public cloud. While Rakuten claims to be building a cloud-native network, the reality is that it will not be 100% cloud native. It also is not clear whether Rakuten and Dish will be able to stick with the more economical public-cloud option or revert to a more hybrid or private cloud environment.

Cloud at the edge

In December, Verizon and AWS announced a partnership to deploy Verizon's 5G Edge computing service in the cloud. Verizon will use AWS Wavelength to provide developers the ability to deploy applications that require ultra-low latency to mobile devices using 5G. In an interview with CNBC, Verizon CEO Hans Vestberg said that the service is live in Chicago with companies such as game publisher Bethesda Softworks and the National Football League using it.

The solution is aimed at enabling latency-sensitive use cases like machine learning inference at the edge, autonomous industrial equipment, smart cars and cities, internet of things, and augmented and virtual reality. Verizon 5G Edge provides mobile edge computing and high-volume connectivity between users, devices, and applications, while AWS Wavelength supports the parts

of an application that require ultra-low latency to the edge of the network while connecting back to the full range of cloud services running in AWS.

Vestberg said the companies began collaborating 18 months ago as Verizon was disaggregating the hardware and software in its mobile network and AWS was developing its cloud service for the edge. The deployment is a combination of public and hybrid cloud with latency-sensitive application components working at the edge and the rest of the app running in AWS' public cloud.

"We have capabilities in the 5G network like extremely low latency, enormous throughput, and we can do slices of the network," Vestberg said. "In order to do those together you need a partner and the first partner we chose was Amazon. They were the most advanced and we are the most advanced, so the compatibility

between our companies was great and our engineers have done terrific work during the last 18 months."

While this deal is initially positive for both parties, it will provide some good insights into how the dynamics between CSPs and cloud partners evolve.

Watch the CNBC interview with Vestberg:

The next section offers guidance to help CSPs begin transitioning to cloud.

Section 5

Make it happen – Strategies for embracing cloud

Communications service providers (CSPs) largely have resisted partnering with public cloud providers because of questions about reliability and security, but they should take the lead of companies like AT&T, Vodafone and Verizon and embrace it for at least some workloads. Not only do the economics and the complex demands of consumers and enterprises demand it, but the cloud is ready. Public cloud has overcome most objections around quality and security, and operators can use other cloud options such as private, hybrid and multicloud to complement public cloud deployments. Following are steps CSPs can take now:

Evaluate public cloud

A full 84% of attendees at Digital Transformation World 2019 said CSPs need to be investing in cloud-native technologies. However, operators will not reap the full benefit without putting more workloads into the public cloud where the total cost of ownership and operation are lowest. They should begin evaluating individual workloads now to determine, even hypothetically, which are suitable for cloud. Whenever possible, they should choose public cloud because the economics are undeniable.

Build partnerships

CSPs will need to build many types of partnerships to target enterprises. Network-facing partnerships with cloud providers (AT&T and Microsoft, for example – see page 16) should focus on proper management of workloads, including scale, performance and security. CSPs can leverage customer-facing partnerships within industry verticals to co-create solutions leveraging the network. Finally, ecosystem partnerships are needed to manage business and technology relationships between CSPs, application developers, platform providers and cloud providers. These should cover interconnection (APIs), revenue sharing, revenue, and quality assurance, data sharing and compliance are required.

The relationship between CSPs and cloud providers at times will be competitive and conflicting, but all parties need to recognize when 'coopetition' is in the best interest of the customer. Parties must look for ways to bring joint solutions to market or form a joint venture as Vodafone has done with IBM (see page 19).

Trust but verify

Operators know how to engineer, monitor and manage networks, and the cloud is a network. However, cloud providers have developed cross-platform management tools to provide managed services, so this will be an area of competition. CSPs should not take their well-trained eyes off the ball. They should exploit their operational and business support system (OSS/BSS) tools to provide managed services.

In cases where a cloud provider is delivering multicloud service management, CSPs must continue to monitor performance to assure quality and identify areas where they could add value for customers and to ensure that the cloud provider is delivering on their commitments by maintaining visibility into service and network performance.

Adopt an open architecture

The "lift-and-shift" method of migrating workloads to the cloud does not equate to being cloud native, nor does it leverage most of the potential benefits of moving to the cloud. More importantly, it provides no roadmap for doing it the right way, which is migrating to a cloud-native architecture and into the public cloud.

TM Forum's Open Digital Framework, which includes the Open Digital Architecture and Open APIs (see page 10), provides a way for CSPs to adopt cloud at their own pace and an opportunity to move much faster. The architecture provides transformation techniques such as optimization, re-engineering or abstraction to select the desired path to cloud-based deployment of IT applications, allowing CSPs to set a vision, create a roadmap and execute on the plan.

Require cloud native

Operators can go only so far in their cloud-native quests because they depend on their suppliers to deliver cloud-native solutions. Not all vendors are as progressive as the ones highlighted in this report. Many are deliberately moving slowly to embrace cloud-native architectures because they want to preserve the status quo as long as possible. CSPs should insist on cloud-native solutions in all requests for proposal, because even if some workloads will not go into the public cloud now, there may come a time when operators want to change their strategies.

Use what works

Unlike whales, birds and butterflies, which push themselves to incredible limits to get from one place to another, people in the telecoms industry consistently underestimate the pain and challenges of migration. It's a tough road with only one golden rule: If there is a safe path that others have followed successfully, use it.

CSPs should look to pioneers that have successfully moved BSS and OSS workloads into the cloud, and they should be willing to share their migration experiences with others by contributing to standards organizations like TM Forum and open source projects. This will make the path less risky for everyone.

Create new lines of business

5G will make it much easier for CSPs to enter new markets and create new lines of business. Operators should leverage the public cloud in the same way that greenfield competitors like Rakuten and Dish Networks intend to – to get to market more quickly at much lower cost. CSPs can introduce public cloud-based BSS and OSS to support new opportunities in the internet of things, managed networks and multi-access edge computing.

Additional features & resources

SaaS has arrived for Telco OSS

Dirk Michel
SVP SaaS Business Operations at MYCOM OSI

Recent years have seen a rapid growth in the adoption of a 'cloud first' strategy by leading global CSPs, driven by benefits relating to speed, scalability, elasticity, Total Cost of Ownership and functional flexibility, which in turn lays the foundation for a step-change in business agility. With 5G roll-outs underway across the globe, the focus is now on the <u>cloudification of OSS applications</u>, including Network and Service Assurance, which are business critical to operating as well as delivering the promise of 5G: orchestrating carrier-grade digital experience guarantees across speed, latency and connection density, underpinned by AI/ML and closed loop automation.

But can Network and Service Assurance – which has traditionally been located on-premise close to CSPs networks – really be delivered from the public cloud as Software-as-a-Service (SaaS)? As Globe Telecom presented at Digital Transformation Asia 2019, large-scale Service Assurance can indeed be migrated to SaaS, that not only achieves significant cloud-based benefits but is also faster than on-premise systems, despite the lack of physical co-location with network infrastructure. Globe Telecom may be a leader in SaaS adoption but several other Tier 1 CSPs are also implementing Service Assurance in the public cloud including the UK's <u>Three</u> and Vodafone.

Learn more about how Globe Telecom is using Service Assurance SaaS:

At the forefront of Service Assurance in the public cloud

MYCOM OSI's <u>Assurance Cloud™</u>, deployed at Globe Telecom, Three and Vodafone, is the telecom industry's first carrier-grade service assurance SaaS offering. It delivers business agility by eliminating months of on-premise deployment time and effort, avoiding system infrastructure management risk and costs and enabling on-demand expansion and feature activation. With a predictable subscription service model, the Assurance Cloud™ is 5G-ready and future-proof, enabling CSPs to grow and accelerate their journey to Digital Service Providers (DSPs), aided by cloud-native AI/ML at scale to realize the vision of highly automated network and service operations.
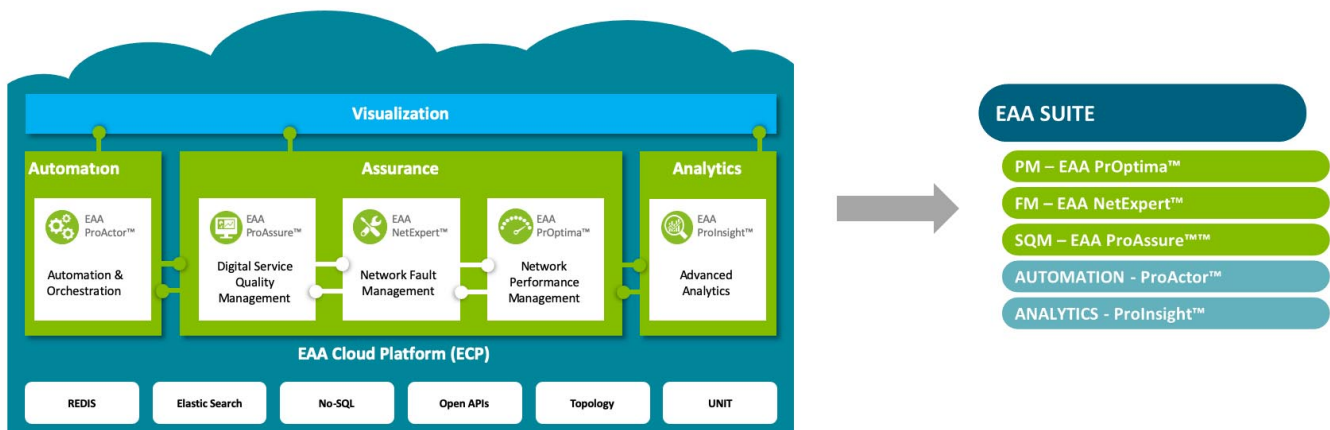
Powered by Amazon Web Services (AWS), The Assurance Cloud™ leverages AWS's advanced security, scalability, availability, platform orchestration, AI/ML services and global footprint to deliver MYCOM OSI's Experience Assurance and Analytics™ (EAA) suite of applications and associated catalogs of data acquisition agents and use case-based SmartPacks™ with ever-increasing levels of automation. Announced at Mobile World Congress in February 2019, The Assurance Cloud™ today already assures some of the largest, fastest and most complex CSP networks in the world, including some of the world's first 5G networks.

Completely eliminating the cost and complexity of infrastructure management

MYCOM OSI controls the cloud infrastructure and middleware on which the Assurance Cloud™ operates, avoiding much of the time and complexity involved in deploying the EAA applications on different 'flavors' of on-premise infrastructure and IT environments. Compared to the on-premise approach, it enables CSPs to take full advantage of the Assurance Cloud™'s deep integration with AWS and helps eliminate large up-front private infrastructure deployment costs, time, complexity and effort.
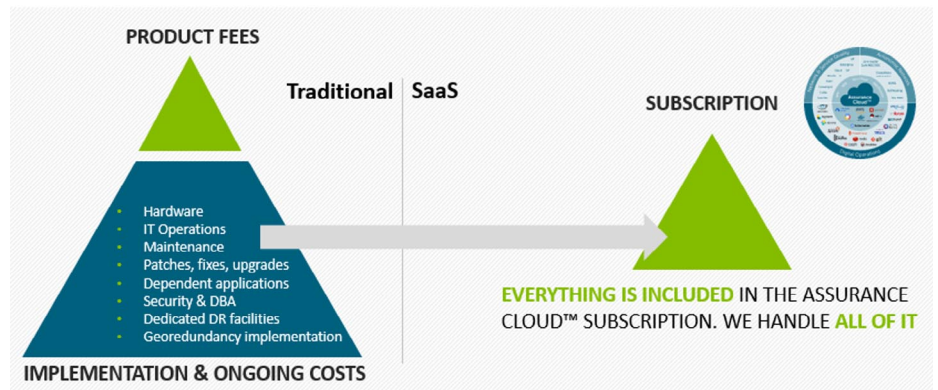

The Assurance Cloud™ is highly automated and deployable, ready to ingest data in under one hour. Time-to-value is typically compressed to under 3 months.

A subscription business model for Telco Service Assurance

MYCOM OSI has created a flexible SaaS offering incorporating a simple, predictable subscription model, tailored to the changing needs of telecommunications service providers. It combines flexibility with predictability, with customers able to "right-size" the initial subscription package and updating the package parameters as and when their requirements change. Changes to the subscription fee are possible on-demand, if for example the CSP intends to expand the user group to

more departments and teams, but crucially this remains under the full control of the CSP.
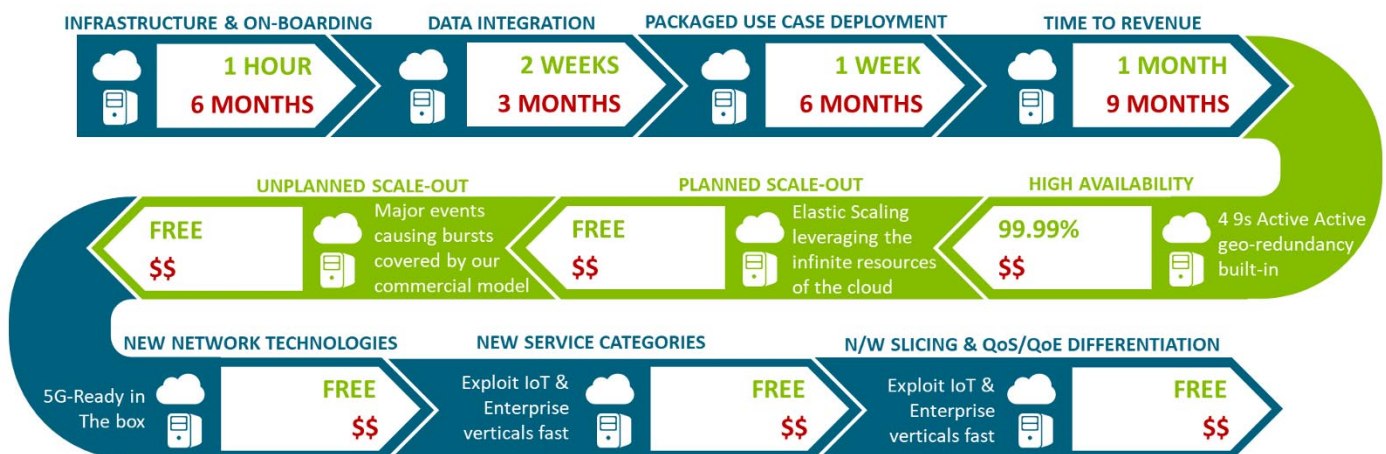
Delivering advanced assurance functionality and faster innovation via the public cloud

AI/ML is a core aspect of assurance-driven closed loop automation, which becomes an essential - not an optional - capability in the 5G era. But integrating AI/ML into cloud-native network and service applications involves leveraging rapidly evolving algorithmic systems such as Amazon SageMaker, as well as adapting to significant fluctuations in specialized compute resource requirements.

MYCOM OSI's Assurance Cloud™ enables on-demand scalability and advanced assurance-driven automation and predictive analytics by leveraging the cloud-native AI/ML services of the AWS public cloud and by incorporating open standards and APIs such as the TM Forum Open APIs and Open Source MANO (OSM). CSPs can additionally leverage integration with peripheral OSS applications to automate and optimize network and service operations across a range of scenarios. As network and service orchestration becomes simultaneously more complex and more business-critical in the 5G era, this assurance-driven automation requirement becomes essential.

Several Tier 1 customers have transformed operations by adopting the Assurance Cloud™ in the following ways:

- **Adopting AI/ML driven closed loop automation:** Advanced analytics (AI/ML) made possible by the vast compute resources of AWS, enabling intelligent closed loop automation, supporting stated targets of upwards of 80 – 95% automation

- **Adopting AI/ML driven predictive analytics:** Training of advanced ML models and creation of predictive inferences at high accuracy underpins a new era in preventative and proactive operational processes for CSPs

- **Adopting a new approach to innovating and risk:** A business model that eliminates barriers to CSP growth through 'all-inclusive' capabilities that seamlessly cater to 5G networks with no additional cost or Service Assurance deployment

- **Unifying their overall assurance estate for cross-domain use cases** through the unification of assurance functions and data into the Assurance Cloud™
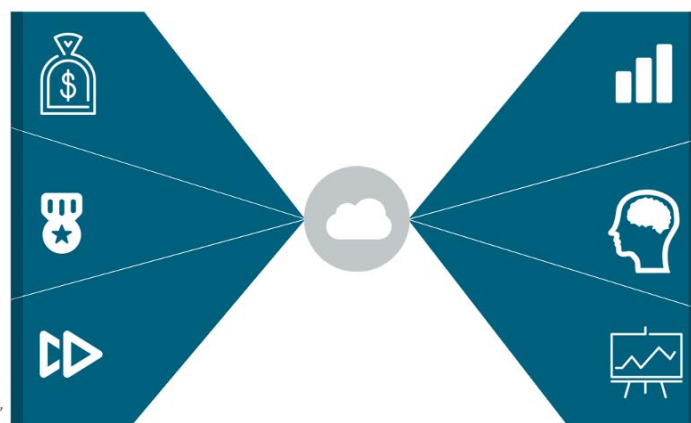
Digital Agility

Assure your digital assests and fast-track your evolution to Digital, giving you the first mover advantage.

Carrier Grade

Secure, geo-redundant, always-on, durable, real-time Assurance for your digital services and infrastructure

On-Demand

Remove bottlenecks to onboarding new users, new networks, new technologies, new Opcos and domains.

5G Ready

Deploy 5G and network slices faster and more reliably while tackling complexity and increase RoI.

Autonomic Networks

Start the Journey towards NFV MANO assurance technology that helps you approach Zero-Touch Operations.

Predictable

Our commercial model follows a flat-fee, future-proofed subscription with the flexibility of a SaaS model

The Experience Assurance & Analytics™ suite of applications and solutions

The Assurance Cloud™ provides the full Experience Assurance & Analytics™ (EAA) suite of applications from the public cloud. EAA provides the functional building blocks for CSPs to move towards ever-increasing automation of their NOC/SOC operations, collapsing the traditional data silos of Service Quality Management (SQM), Performance Management (PM) and Fault Management (FM), supported by a platform of common resources and data models that enable end-to-end visibility and control across all network domains. Based on modern containerized cloud-native technology, it evolves to meet the demands of CSPs today and in the future. EAA:

- Eliminates the silos that constrain CSPs' visibility of the customer experience, with cross-domain comparable KPIs based on a common platform

- Leverages orchestrators for CNF, VNF and PNF technologies to pro-actively automate increasingly complex manual tasks in real time, enabling the resolution of many issues before they impact the customers, reducing MTTR and enabling SLA compliance

- Enables data-driven decision-making, by deploying AI/ML to uncover hidden issues and trends, and implement predictive maintenance and capacity planning, unlocking value from a wealth of network, product and context data

- Provides carrier-grade assurance capabilities out-of-the-box for emerging technologies including 5G, with a catalog of 200+ pre-integrated use case-based SmartPacks™

- Provides a rapid release cycle supported by DevOps and CI/CD, responsive to the changing needs of CSPs

- Supports leading open technologies including OSM, ONAP, OpenDayLight and many more, through active, non-exclusive involvement in open source projects and standards bodies

Subscribers to the Assurance Cloud™ have access to the entire 200+ SmartPacks™ catalog on-demand and can activate or de-activate any given solution at any time, within the parameters of their selected subscription package. This represents a significant improvement in the way CSPs can experiment with services, because the instantiation of assurance of those services is frictionless, eliminating a bottleneck in service innovation without compromising on carrier-grade performance.

The cloud is essential to Telco Service Assurance in the 5G era

Network and Service Assurance will continue to play an increasingly important role in the move to advanced 5G monetization. Network slicing promises to unlock new revenue streams across a range of enterprise verticals, and it is not difficult to view the full 'stand alone' 5G network of the future being split into thousands of end-to-end slices, each catering to a unique set of performance requirements.

Network and Service Assurance is critical to efficiently orchestrating the overall network in this future scenario

and there are three key characteristics that mean the public cloud, if not completely essential, has significant advantages over on-premise deployments. First, enabling 'tenants' of individual network slices to operate and assure their network and services will require the on-demand instantiation of assurance capability, especially given the high turnaround of tenants across thousands of slices. Second, for the CSP itself, the prospect of orchestrating speed (eMBB), latency (uRLLC) and connection density (mMTC) guarantees across thousands of concurrent networks slices comes with significant uncertainty, across compute resource capacity requirements as well as the business scenarios it will cater to. Third, the efficient automation of network slices will involve deep integration with advanced AI/ML services and external orchestration systems.

Moving into the 5G era, it will therefore be essential for Network and Service Assurance systems to

support on-demand access to new features, massive scalability and advanced automation capabilities. The public cloud is uniquely placed to deliver on this requirement. Public cloud providers such as Amazon are progressing further into distributed cloud architectures that support service delivery and assurance from the edge, for low latency and compute intensive scenarios, with services such as Outpost and Wavelength. Looking into the medium-term horizon, innovations such as these mean that MYCOM OSI's Assurance Cloud™ is uniquely placed to lead the way in driving public cloud adoption in the category of telco Service Assurance.
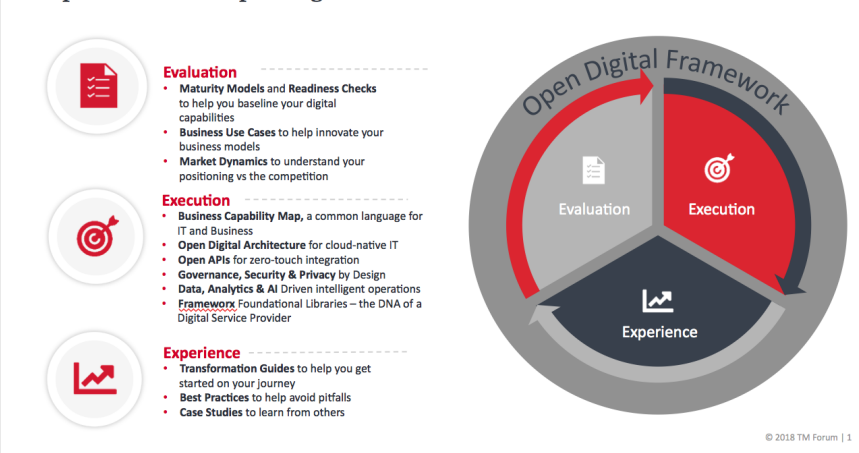
TM Forum Open Digital Framework

Delivering the tools to go from concept to cash in just 18 days

The TM Forum Open Digital Framework is an interactive, continuously evolving collection of tools, knowledge and standards that give communications service providers (CSPs) an end-to-end migration path from legacy systems to modular, cloud-native IT components. Simply put, it is a blueprint for service providers to deliver intelligent operations fit for the 5G era.

A prototype version of the framework is available now for TM Forum members to explore. It is being developed through the TM Forum Collaboration Program and Catalyst Program, and builds on the success of the Forum's established Open APIs and the Frameworx suite of standards. Specifically, it includes:

- **Open Digital Architecture (ODA)** – an enterprise architecture blueprint, common language and key design principles for modular, cloud-based, open digital platforms that can be orchestrated using AI

- **Open APIs** – 50+ standardized REST-based APIs to facilitate zero-touch integration and zero-touch partnering

- **Data & AI standards** – an industry-agreed data model,

Components of the Open Digital Framework

Evaluation
- **Maturity Models and Readiness Checks** to help you baseline your digital capabilities
- **Business Use Cases** to help innovate your business models
- **Market Dynamics** to understand your positioning vs the competition

Execution
- **Business Capability Map**, a common language for IT and Business
- **Open Digital Architecture** for cloud-native IT
- **Open APIs** for zero-touch integration
- **Governance, Security & Privacy** by Design
- **Data, Analytics & AI** Driven intelligent operations
- **Frameworx** Foundational Libraries – the DNA of a Digital Service Provider

Experience
- **Transformation Guides** to help you get started on your journey
- **Best Practices** to help avoid pitfalls
- **Case Studies** to learn from others

© 2018 TM Forum | 1

together with standards maximizing the potential of AI to enhance customer experience and increase operational efficiency
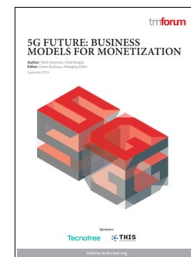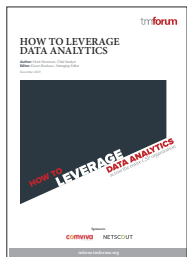
- **Reference implementations** – a framework for assembling and validating ODA components in the Forum's Open Digital Lab, fostering the creation of a services marketplace

- **Practical guidance** – guides and videos showing how the Open Digital Framework can be used to transform the core business and enable new business growth

- **Foundational libraries** – normalized models providing a common language for business processes and information that

simplifies and de-risks transformation projects

The goal of the Open Digital Framework is to help service providers increase agility and drastically reduce the development cycle for products and services from 18 months to 18 days. Much of the collaborative work that is part of the framework is already available, but it helps to organize it and make it more accessible. The framework is a work in progress and will improve through crowdsourcing.

If you would like to learn more about the project or how to get involved in the TM Forum Collaboration Community, please contact Andy Tiller.

TM Forum research reports

HOW TO LEVERAGE DATA ANALYTICS

ENTERPRISE Orchestration SERVICE

FUTURE OSS: TOWARDS AN OPEN DIGITAL ARCHITECTURE

DIGITAL TRANSFORMATION TRACKER ASIA: SEIZING NEW OPPORTUNITIES

5G FUTURE: TARGETING THE ENTERPRISE

5G FUTURE: BUSINESS MODELS FOR MONETIZATION

SECURITY IMPERATIVES FOR DIGITAL TRANSFORMATION

CTIO OUTLOOK 2020: TRANSFORMING OPERATIONS TO MONETIZE 5G

TIME TO KILL THE RFP? REINVENTING IT PROCUREMENT FOR THE 2020s: Volume 1

TAKING ON THE GREENFIELD OPERATORS WITH DIGITAL SUB-BRANDS

FUSING OUTSIDE PLANT DATA AND AI to maximize CapEx

BLOCKCHAIN: WHERE'S THE VALUE FOR TELECOMS?

NETWORK AS A SERVICE: ADDRESSING THE FULL ENTERPRISE OPPORTUNITY

DYNAMIC INVENTORY The foundation of intelligent automation

5G: EVOLUTION AND REVOLUTION

Meet the Research & Media team

Report Author:
Tim McElligott
Senior Analyst
tmcelligott@tmforum.org

Chief Analyst:
Mark Newman
Chief Analyst
mnewman@tmforum.org

Customer Success &
Operations Manager:
Ali Groves
agroves@tmforum.org

Global Account Director:
Carine Vandevelde
cvandevelde@tmforum.org

Report Editor:
Dawn Bushaus
Managing Editor
dbushaus@tmforum.org

Editor, Digital Content:
Arti Mehta
amehta@tmforum.org

Commercial Manager,
Research & Media:
Tim Edwards
tedwards@tmforum.org

Digital Marketing
Manager:
Anna Kurmanbaeva
akurmanbaeva@tmforum.org

For more about TM Forum's Open Digital Architecture Project, please contact **George Glass**