analysys mason

THE EVOLUTION TO CLOUD-NATIVE NFV: EARLY ADOPTION BRINGS BENEFITS WITH A FLEXIBLE APPROACH

NOVEMBER 2017

Caroline Chappell

Ref: 2011421-463

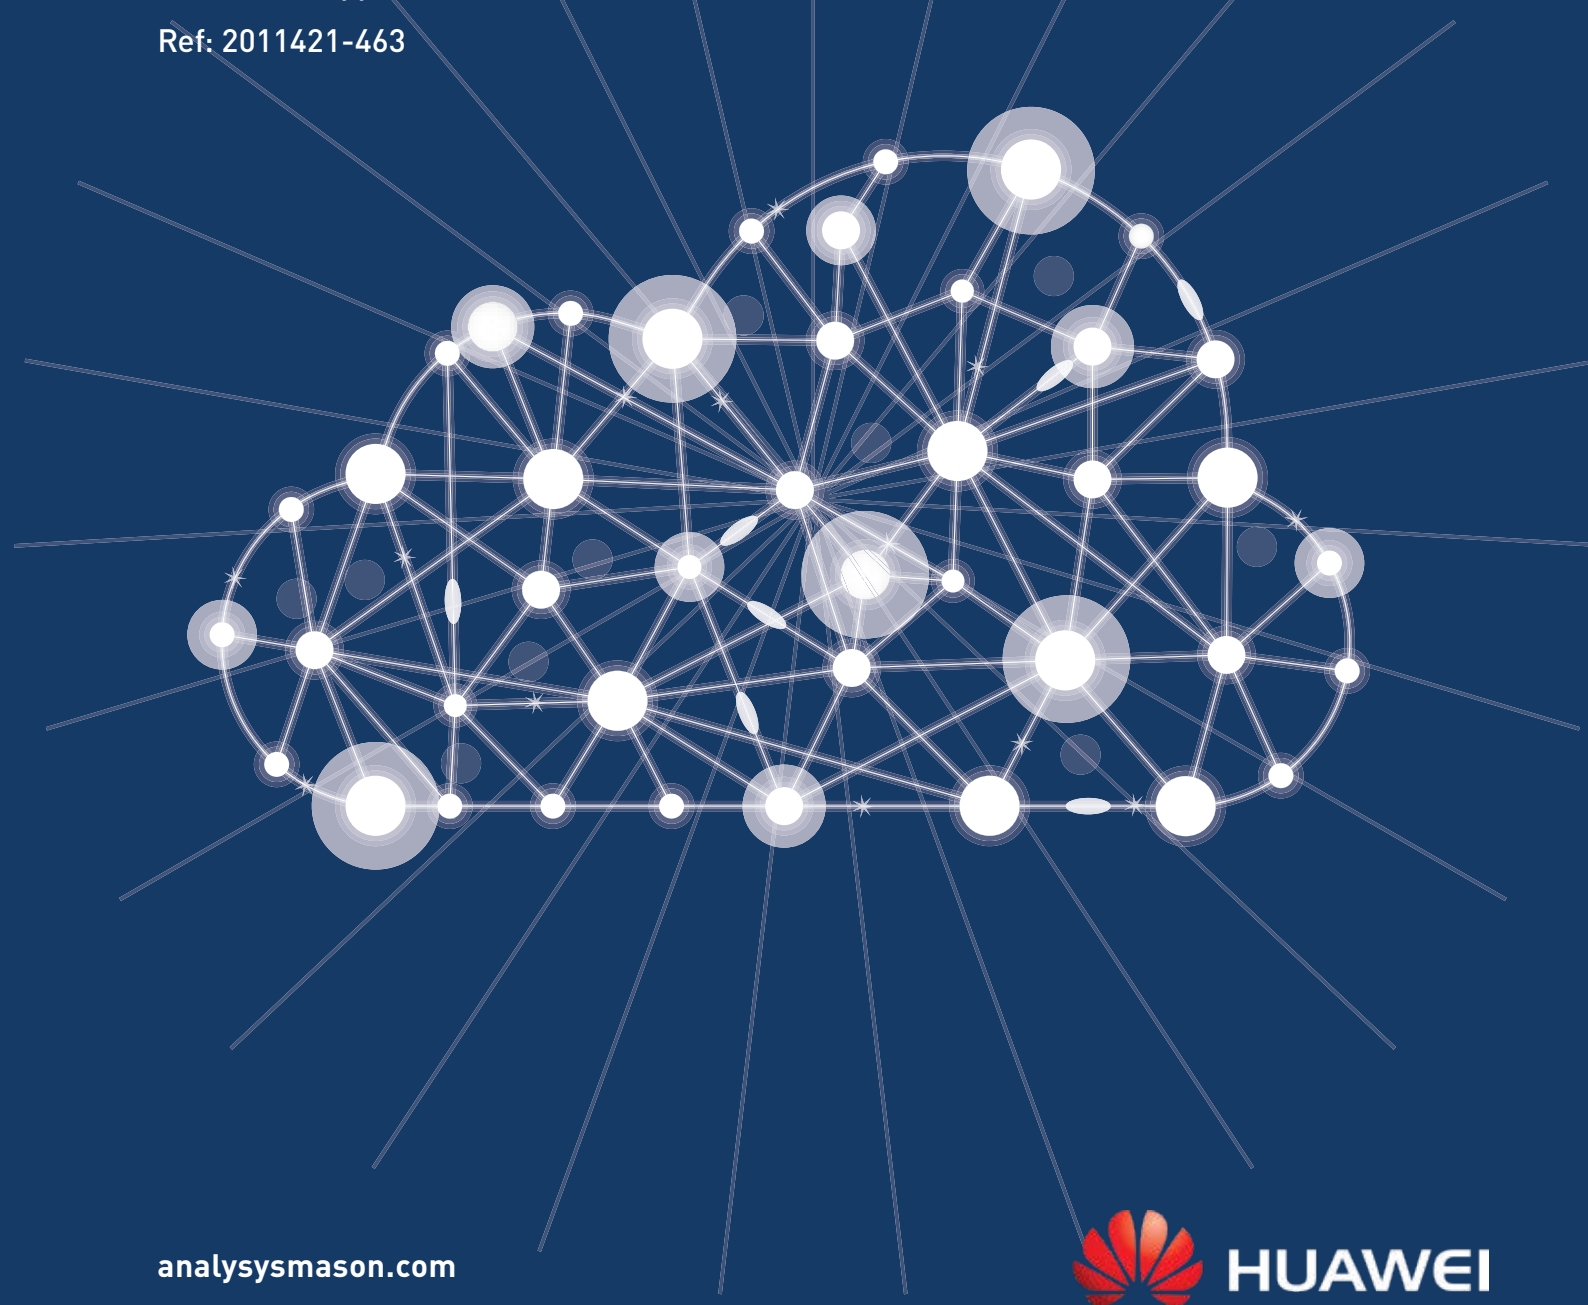

analysysmason.com

HUAWEI

Contents

List of figures

1. Executive summary

Software is the key enabler of the digital era. Webscale companies such as Google, Alibaba, Facebook and Tencent are using cloud-native computing as the fastest and most efficient way of developing and deploying software in the cloud. Other companies also recognise that they must apply cloud-native computing in their businesses or risk being left behind in the digital race. As communications service providers (CSPs) 'softwareise' their networks through network function virtualisation (NFV), they too should adopt cloud-native computing if they wish to achieve the full benefits of this initiative.

In the IT domain, the Cloud Native Computing Foundation (CNCF) has created a definition of cloud-native computing based on a combination of a microservices design pattern, containerisation and orchestration. The CNCF definition is helpful, but it is important to note that the term 'cloud native' covers a spectrum of implementation approaches adapted to varying business requirements across different industry sectors. There is no 'one size fits all' application of cloud-native computing that is right for every telecoms use case, but there are common principles which CSPs can adopt to achieve similar benefits to those enjoyed by the webscale companies.

This paper describes a telecoms-appropriate approach to cloud-native software design, deployment and operation which CSPs and their vendors are developing. It explains how the cloud-native principles of statelessness and horizontal service decomposition can make virtual network functions (VNFs) more resilient and faster to change and upgrade. It suggests that CSPs should explore containerisation in use cases where it makes sense but use other, more mature and secure deployment options for critical use cases. The paper also explains key aspects of VNF lifecycle management automation, such as continuous integration and A/B testing, which CSPs can start to adopt today in preparation for the 5G cloud-native network of the future. The paper recommends that CSPs procure VNFs that have been designed with cloud-native computing in mind, understand the deployment benefits and risks of using containers and virtual machines (VMs) and develop and/or procure the skillsets and tools needed to operate and manage cloud-native VNFs in a highly automated way.

Finally, the paper points out that the extreme level of automation needed to operate cloud-native network functions and network services requires a different organisation and expertise than CSPs have in place today. The way CSPs migrate their organisations will vary depending on their individual circumstances and preferences but best practice approaches are emerging, for example starting small with a 'tiger team' that can build an initial DevOps blueprint or seeding the existing organisation with NFV capable staff organised into a virtual team. CSPs need to be prepared to work with vendor partners who are the 'developers' in the telecoms DevOps domain, and to leverage vendor expertise to accelerate the rate at which they adopt cloud-native computing capabilities. The paper cites examples of CSPs working on joint innovation and cloud-native computing proofs of concepts (PoCs) and commercial deployments with a vendor partner, Huawei.

CSPs which adopt cloud-native computing in a network context early will gain significant competitive advantages, for example through their ability to deliver services at speed, reduce network capex and opex, and address exciting new revenue opportunities in the enterprise sector.

2. What is cloud-native network virtualisation?

2.1 Drivers for the cloud-native network

Every business on the planet needs to become cloud-literate if it is to thrive in the 21st century. This means that companies must adopt cloud-oriented behaviours exemplified by the FANG[1] and similar companies. Facebook, Amazon, Netflix and Google are responsible for the rapid growth of the cloud and/or run their businesses in the cloud, so their names (and those of other webscale giants, such as Tencent and Alibaba) are synonymous with cloud literacy.

The outstanding cloud behaviour displayed by these companies is their mastery of software development and delivery, since software is powering the digital era. The FANG companies have pioneered cloud-native computing as the fastest and most efficient way of developing and deploying software in the cloud. Their cloud-native computing competence reinforces their dominant market positions.

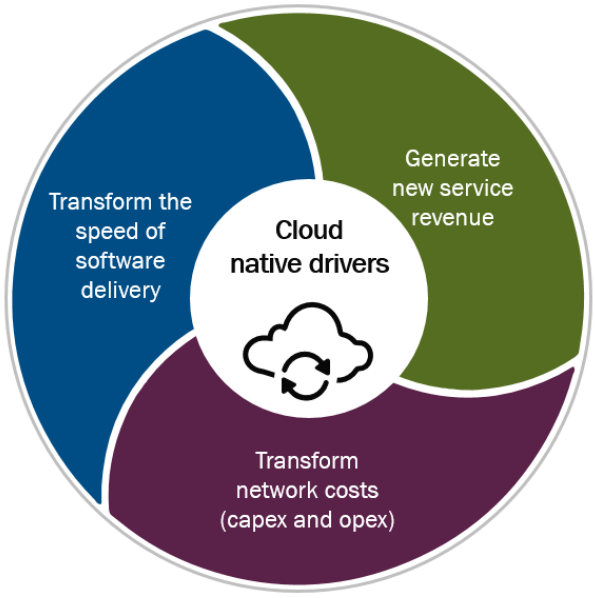

Figure 2.1: Drivers for applying cloud-native computing to the network [Source: Analysys Mason, 2017]

The drivers for CSP adoption of cloud-native computing are clear:

- **Cloud-native computing transforms the speed of service delivery.** Leading CSPs realise that they must adopt cloud-native computing if they want to transform themselves into digital service providers with the ability to deliver new software-based services as fast as their FANG competitors. As CSPs become more software-capable, driving cloud-native computing into their networks, they can differentiate their network services in ways that are not possible in traditional, appliance-based networks. This will help them become more competitive in their existing markets. Advanced CSPs which are already implementing cloud-native network functions have found that they can introduce innovative features faster than their rivals, expand their service portfolios to increase customer stickiness and trial new services at low risk, 'failing fast' if service take-up is low or scaling quickly if the services are successful.

- **Cloud-native computing is key to generating new service revenue**. Cloud-native computing is required to support an increasing range of emerging market opportunities, such as smart homes/meters/cars and other Internet of Things (IoT) services, augmented/virtual reality services and 5G network slices aimed at

[1] Facebook, Amazon, Netflix and Google are cloud pioneers and digital service provider exemplars.

specific vertical sectors (including health, automotive and public safety). Because IoT use cases are so varied in their bandwidth and latency requirements, CSPs want a cost-efficient method of tailoring the network for each use case. They cannot afford to create multiple, separate physical networks, but cloud-native network virtualisation enables them to generate virtual network slices automatically, with specific characteristics suited to particular use cases. They can then apply differentiated pricing to each slice. As this paper will discuss, cloud-native computing enables operators to respond at speed to new business requirements that promise to yield new sources of revenue. It supports deployment automation at scale, so CSPs can launch, customise and remove new services rapidly with little risk.

- **Cloud-native computing transforms the cost of building and operating the network.** Cloud-native computing is designed to be a highly efficient means of deploying and managing software. It achieves cost efficiencies through the use of extreme levels of automation, such as ultra-high levels of resource utilisation, and significantly faster deployment, scaling and recovery of applications than is possible with earlier virtualisation approaches. When these capabilities are applied to networking software, they can have a significant impact on capex and opex.

2.2 The IT cloud community has pioneered cloud-native computing

It is important to understand why FANG companies pioneered a different software-development paradigm for the cloud.

Traditionally, software applications have been designed as sets of tightly-coupled functions. Such applications are unable to leverage the flexibility and resource utilisation potential of the cloud. This is because their monolithic architecture does not enable support the differentiated and highly efficient scaling and restoration of individual application components. Monolithic applications must be handled with care in the cloud environment, because they are not designed to cope with common hardware and cloud availability zone failures. They are typically known as ‘pets’ and require special management and virtualisation support.

Application developers working with the cloud soon realised that traditional software architectures were not optimal for the cloud. New patterns of software development that could leverage the capabilities of the cloud began to appear in 2011, as Figure 2.2 shows, and the first cloud-native orchestration systems were under development a year later. The Cloud Native Computing Foundation (CNCF), a non-profit foundation, was created in 2015 to support the adoption of cloud-native IT environments.

Figure 2.2: The evolution of cloud-native computing in the IT industry [Source: Analysys Mason, 2017]

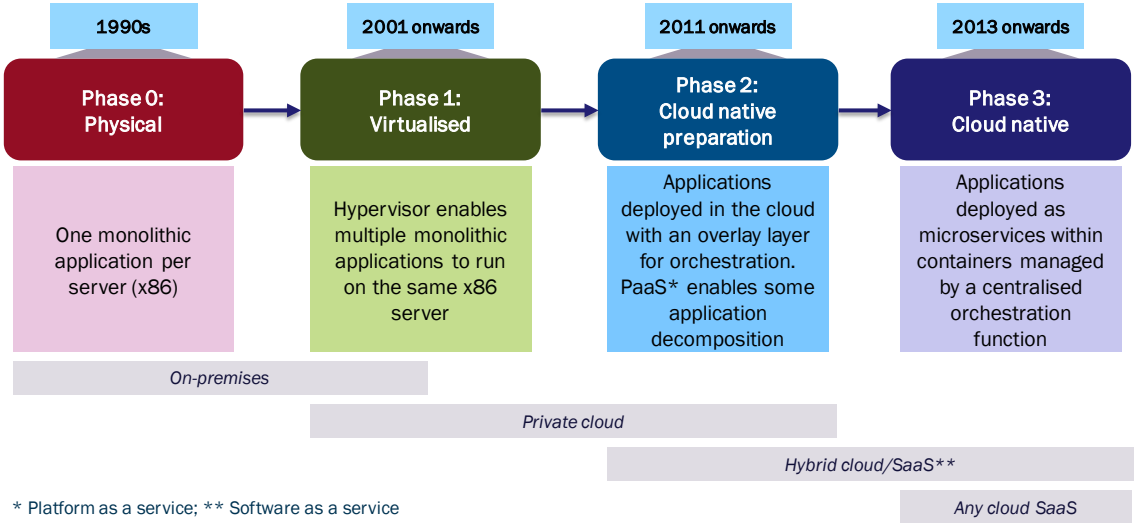

According to the CNCF, all systems and applications characterised as 'cloud native' have three fundamental characteristics. They are:

- **Microservices oriented**. In a cloud-native architecture, applications are developed as 'micro' services which are stateless (that is, each individual application does not have persistent data storage) and loosely coupled. Larger systems can be composed from microservices, since the latter's loosely-coupled status makes them reusable in different application contexts. Microservices can easily and automatically be replicated in the cloud, where individual, stateless instances can be treated as 'cattle': if one instance falls over, another can immediately take its place and further replacements quickly generated. Microservice architectures are therefore highly resilient; they support the rapid release of application enhancements or new features, improving time to market, and they incur lower operational costs.

- **Containerised**. Each microservice (application, process, etc.) is packaged in its own container for isolation and ease of deployment. A container is a lightweight virtualisation mechanism that encapsulates an entire application run-time (that is, the minimal package of dependencies, libraries and configs needed to run the application). Containerisation lets applications run independently and portably across different cloud infrastructures. Figure 2.3 illustrates the difference between containers and virtual machines (VMs).

- **Dynamically orchestrated.** Containers are actively scheduled and managed to optimise resource utilisation. Container orchestration systems are responsible for tasks such as provisioning host resources and assigning (scheduling) containers to hosts, instantiating a set of containers, rescheduling them if they fail, integrating containers through application programming interfaces (APIs) and scaling them.

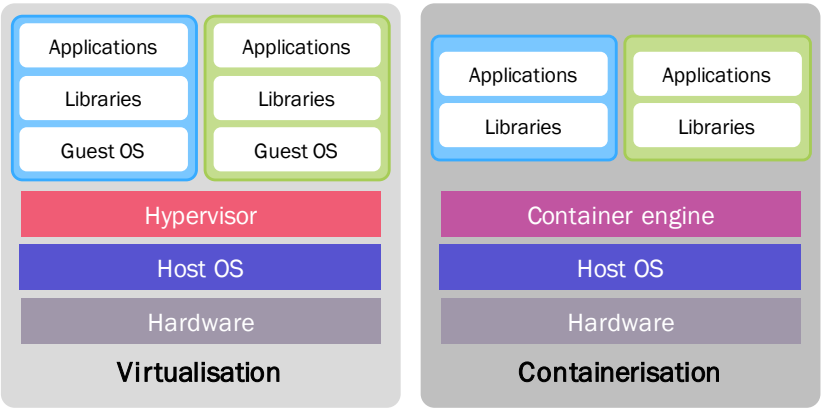

Figure 2.3: How containers differ from virtual machines [Source: Analysys Mason, 2017]

The CNCF definition is helpful, but it is important to note that the term 'cloud native' covers a spectrum of implementation approaches. This is due to the varying business requirements and appetite for risk across different industry sectors, and the fact that there is no 'one size fits all' application of cloud-native computing that is appropriate for every use case. For example, there is no industry standard for the size of a microservice. It is more critical that a 'microservice' is stateless than small. Verizon uses the term 'microservices' to refer to the functions within its enterprise SDN implementation – routeing, WAN acceleration, load balancing, firewall, NAT and intrusion detection. Meanwhile, Amazon defines a microservice as a function that is no bigger than a 'two-pizza' team can run: that is, a team that doesn't need more than two pizzas for a team meal.

2.3 Telco industry progress towards cloud-native NFV

The telecoms industry has specific business constraints and requirements for VNFs which are shaping its approach to cloud-native computing. VNFs and the network services they participate in have specific

requirements and characteristics, for example related to latency and throughput, which do not apply to IT applications. As a result, CSPs should be guided by business need and technology practicalities when evaluating the fundamental characteristics of cloud-native computing.

Cloud-native computing technologies such as containers are evolving and immature, so they may pose unacceptable risks for certain VNFs that need to support mission-critical parts of the network and millions of subscribers. It may not make economic sense, particularly for user-plane VNF components, to modularise them to the level of granularity of 'two-pizza' microservices. Today, there are very few complex VNFs that consist entirely of 'two-pizza' microservices instantiated in containers and orchestrated to support dynamic load handling and scalability in a distributed cloud environment. Such VNFs are typically regarded as experimental by the market and tend to be deployed for small-scale, non-critical use cases at the moment. However, containerisation is a fast-moving technology domain which holds significant promise for 5G, so it is important that CSPs and vendors take such pioneering steps.

There are, however, VNFs that judiciously apply key cloud-native principles, including a microservices-based pattern of functional decomposition and appropriate use of containers, to achieve resiliency and agility benefits. The cloud-native implementation of such VNFs is business-driven, in that it balances the unique needs of individual VNFs with what is technologically possible today using a cloud-native approach to software design, deployment and operation.

These principles include:

- Support for the statelessness of VNF components, and their ability to be loosely coupled
- The decomposition of monolithic VNFs into components of appropriate scope, with the emphasis on creating modular units that make sense from a functional and economic perspective, rather than on the size of those components. Such decomposition supports the independent scaling and lifecycle upgrades of individual modules and the ability to expand a VNF with further modules as needed over time
- Support for containerisation where it can safely be adopted, given the current level of maturity of the technology, alongside VMs
- Automation (orchestration) of VNF components (VNFCs), VNF and network service lifecycle management.

2.4 Cloud native is an urgent goal, but most vendors are moving slowly towards it

When ETSI Network Functions Virtualisation Industry Specification Group (NFV ISG) published its seminal paper on network functions virtualisation[2] in 2012, its founder members were clear that their end goal was to build a network based on cloud-native principles, even though cloud-native concepts were still in their infancy. Progress towards applying cloud-native principles has been slower than advanced CSPs would have liked, because of the difficulties of refactoring traditional, monolithic networking software, built for a non-virtualised, proprietary hardware environment, for the cloud.

Most vendors have taken a step-by-step approach to NFV. In Phase 1, shown in Figure 2.4 below, they merely 'forklifted' their network function software out of proprietary appliances and onto hypervisor-enabled commercial off-the-shelf (COTS) servers, without re-architecting it for the cloud. They assumed that the resulting 'pet' VNFs would be manually operated alongside their physical networking function (PNF) counterparts, and so have not prepared their VNFs for orchestration. Because the VNFs' software architecture is still monolithic, the entire VNF must be replicated for redundancy. The first phase of NFV has proven that network functions can be virtualised on commodity hardware with a similar performance

2 See https://portal.etsi.org/NFV/NFV_White_Paper.pdf

profile to proprietary appliances and this low-risk approach has secured mainstream market support for NFV. However, the initial phase of NFV has not achieved the significant capex and opex savings originally envisaged by the founders of ETSI NFV ISG.

Leading vendors, however, have created network function software and its automated operations (orchestration) in a cloud-friendly way, without passing through the virtualisation phase. Such vendors aim to deliver greater benefits than are achievable through virtualisation alone, while recognising the current limitations and immaturity of cloud-native technologies, such as containerisation, when applied to the network.

As a first step towards cloud-native VNFs, such vendors started to modularise their network functions. They also ensured that their modular components were orchestration-ready according to DevOps principles, which dictate that a component's operational aspects must be defined at design-time, so that an orchestration system automatically understands its execution requirements at run-time. VNF components may run in virtual machines (VMs) where there is a need for a mature, highly reliable and secure deployment environment, or in containers if the requirements for a specific component are less stringent.

In fully cloud-native 5G networks, CSPs would like to see more VNF components running in containers and decomposed to a greater degree, but further research and development is needed to test the validity of these assumptions, as explained in the next section.

Figure 2.4: Progress towards cloud-native VNFs [Source: Analysys Mason, 2017]

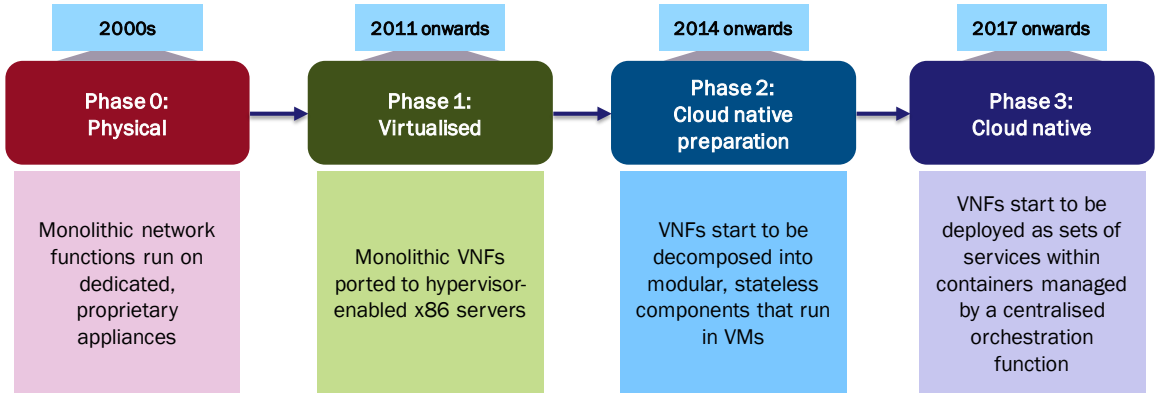

3. Key principles of cloud-native computing in a telco context

3.1 CSPs should evaluate VNFs from three perspectives to ensure they are future-proofed for cloud-native computing

To start their journey towards a 5G cloud-native future, CSPs should:

- Procure VNFs that have been **designed** with cloud-native computing in mind. Such VNFs may not be microservices-based but they will have key features, such as statelessness and control/user plane separation in the case of a virtual Evolved Packet Core (vEPC), which will support migration to a cloud-native end state.

- Evaluate the **deployment** benefits and risks associated with different virtualisation environments – virtual machines (VMs) or containers – at this point in the market. Containers are faster to deploy and recover, but VMs offer stability and mature operations features.

- Engage in the development/procurement of the skillsets, tools and automation needed to **operate and manage** VNFs in a cloud-native way; for example, DevOps methodologies, continuous integration and testing tool chains, and automated feedback loops that support VNF scaling and healing.

3.2 Designing VNFs for the cloud

To be ready for cloud-native computing, a VNF must incorporate two foundational software patterns:

***Stateless design** is a cornerstone of a VNF that is developed natively to run in the cloud*	Statelessness can be achieved by 'vertically' decomposing a VNF and network functionality into three types of component: stateless (process) components, a stateful data layer and, in certain scenarios (for example, control plane VNFs) a load balancing module. This enables CSPs to scale the three types of component independently and manage each appropriately. The stateless components are 'cattle' that can be distributed across blades/servers, and so any single instance of a stateless component (VNFC) can fail without affecting the correct operation of the VNF overall. Since such VNFCs do not hold state, they are, in effect, fault-tolerant copies of one another, providing N-way resilience (active–active) for better resource utilisation.
***Service decomposition** is key to achieving statelessness, and simplifies the architecture of a VNF by decomposing it horizontally into discrete process modules*	Service decomposition implies the horizontal modularisation of VNFs into discrete processing units. 3GPP is working on a service-based architecture (SBA) for the 5G network, characterised by modular services and lightweight interaction protocols. It is hoped that the SBA will provide a standardised, microservices-based architecture for the 5G network, enabling faster service development cycles and upgrades, the dynamic discovery and launch of services on demand in the network, and the ability to compose services held within a service library together, like Lego blocks, in agile and innovative ways. There is no assumption about the granularity of SBA components or whether they should be deployed in containers or VMs, since determining the right level of functional decomposition in the network is a work in progress. CSPs and their vendors are still using proofs of concept (PoCs), trials and commercial deployments to learn about the performance, resilience, speed and other trade-offs that must be made when different VNFs and VNFCs are modularised.

VNFs that are vertically and horizontally decomposed, such as a cloud-native policy and charging rules function (PCRF), for example, can be customised faster to meet new customer requirements, since policy rules can be updated and new ones released more simply in a modular architecture than they can in a monolithic function, where more regression testing is needed. Figure 3.1 illustrates the vertical and horizontal decomposition of a VNF.

Service decomposition in the mobile core network (virtual EPC) is mandating the separation between control and user planes, known as CUPS. CUPS is a prerequisite for a cloud-native, 5G mobile core where services with low latency requirements will need support from large numbers of distributed, virtualised gateways at the edge of the network. In this scenario, it will no longer be practical to maintain a VNF architecture where control and user plane functions are tightly coupled in specific gateways. If they are decoupled, CUPS can

be scaled and upgraded independently and appropriately, control functionality can be centralised, network slicing elegantly supported, and capex and opex reduced.

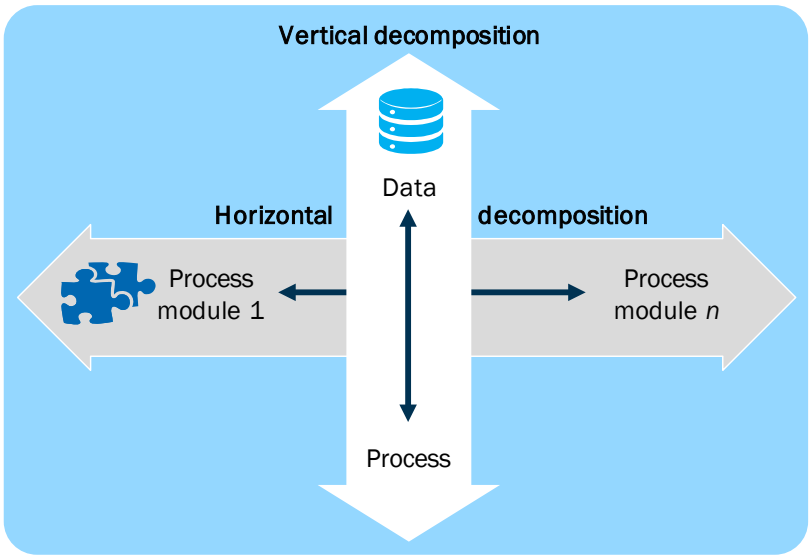

Figure 3.1: Vertical and horizontal decomposition of VNFs [Source: Huawei/Analysys Mason, 2017]

3.3 Cloud-native deployment of VNFs

Containers have been adopted very rapidly as a virtualisation method in the IT world, because they enable much faster software deployment than VMs. As containers are so lightweight, it is considerably quicker to deploy and tear them down, they consume fewer resources (so more of them can share the same resource pool), and because a Docker container encapsulates an entire application run-time, it is portable across a highly distributed infrastructure, including across multiple clouds. Containers are preferable to VMs where high density is required (for example, more than ten instances of an application on the same machine) and where instances are short-lived.

Containers have disadvantages, however, due to potential security risks and the current immaturity of the container ecosystem. The kernel operating system on which all containers in a system depend is not just a potential single point of failure: it represents a single point of entrance for a cyber attack. Isolating a fault within a single container is not easy: the fault can quickly flood across all the other containers sharing the same kernel.

Containerisation is a nascent technology and has not yet been hardened with the NFV-specific developments that have been applied to hypervisors and OpenStack, such as support for Enhanced Platform Awareness features (e.g. single-root input/output virtualisation (SR/IOV) and Data Plane Development Kit (DPDK)), and carrier-grade networking capabilities. CSPs that use containers in a networking context face challenges associated with multi-tenancy support, multi-network plane support, forwarding throughput and limited orchestration capabilities. CSPs should work with open-source container communities to help them understand NFV requirements: ETSI NFV ISG has set up a working group to investigate best-practice container use in an NFV context.

Mobile Edge Computing (MEC) is likely to be an area for early bare-metal container adoption due to the small size and therefore resource constraints of edge data centres, and the fact that edge applications will be those that need the very lowest latency.

However, containerisation is expected to co-exist with VM-based virtualisation in the network for a long time. CSPs are faced with a rich choice of deployment options for VNFs, as Figure 3.2 shows:

- VM-only mode
- Container-only mode
- Hybrid mode: containers run in VMs, which provide isolation and security measures
- Heterogeneous mode: some VNFs run in containers, some in VMs, and others in a mix of both.

CSPs should evaluate which deployment mode(s) are appropriate for them, based on the use case they wish to implement and their business requirements, such as the level of risk they wish to incur versus the cost and agility of deployment.

Figure 3.2: Deployment options in different NFV data centres [Source: Huawei, 2017]

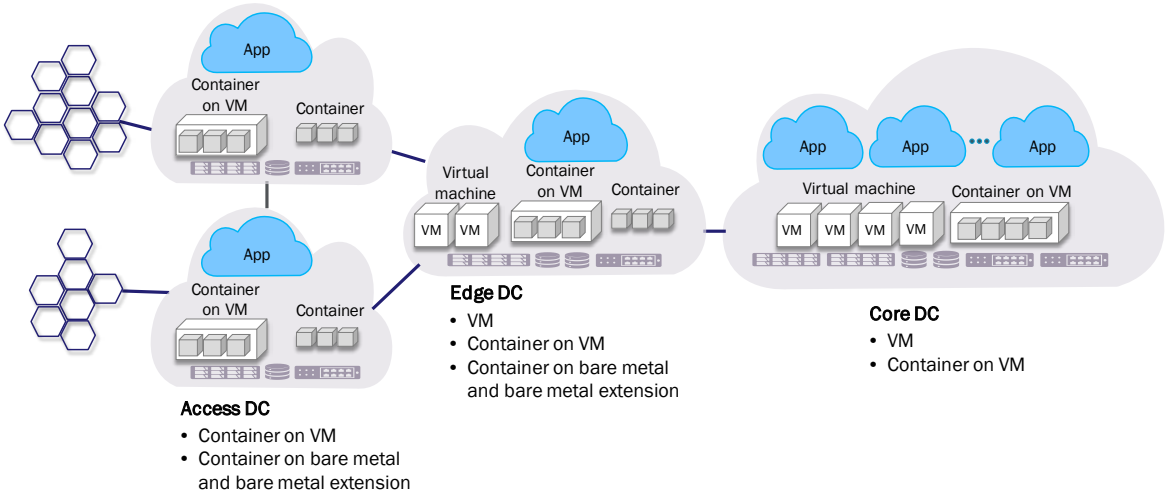

3.4 Cloud-native automation and management of VNFs

Cloud-native applications are inherently built for orchestration. That is, they are designed to be managed automatically as part of a larger (cloud) system with shared resources that can be consumed on demand. Open-source cloud-native orchestration platforms, such as Kubernetes, Mesos and OpenShift, have rich sets of operational tools which developers use to programme the operations and maintenance of their microservices and containers in the cloud. Cloud-native developers naturally use an agile DevOps methodology to design and manage their applications at run-time.

CSPs can adopt a DevOps approach to software development and delivery before they have a fully cloud-native NFV environment. This will ease the transition to cloud-native computing by ensuring that CSPs acquire the right skills and experience ahead of their migration. CSPs will, however, need appropriate orchestration tools and approaches for vertically and horizontally decomposed VNFs which are not yet fully microservices-based and which run in VMs rather than containers. These tools and approaches should provide similar capabilities to those in open-source cloud-native orchestration platforms but they should be adapted and extended for the complex needs of networking components. For example, they will need to:

- **Support continuous integration** (CI) and in-life software upgrades (ISSU) of multiple VNFs from different vendors. Continuous integration of new features and upgrades is a cloud-native capability perfected by the FANG companies: Amazon typically deploys new software every second. CSPs are unlikely to update or upgrade their network functions at this rate for the foreseeable future, but they do need a robust tool chain that supports the automatic onboarding, testing and deployment of enhancements to, and new software versions of, VNF components, without breaking the VNFs and service chains which include those components.

- **Provide generic lifecycle-management** functions across multiple VNFs, again potentially from different vendors. Such orchestration tools must make it easy for CSPs to automate the configuration, instantiation, scaling, healing, maintenance and retirement of VNFs, based on standardised interfaces and templates, as well as feedback from VNFs themselves and their cloud environment.

- **Use big data and analytics** to monitor, correlate and analyse the wealth of VNF and NFV infrastructure performance data to detect service-quality deterioration and drive closed-loop (automated) actions such as the proactive recovery ('recover first') of VNFs and services and cross-domain root-cause analysis. This is an area that is benefiting from the application of machine learning and other AI technologies.

- **Support A/B testing (or split testing)**, an IT DevOps concept which supports the incremental introduction of a new version of a VNF component, thus reducing the risks associated with traditional cutovers between versions and enabling new versions to be deployed faster. As Figure 3.3 shows, traffic can gradually be switched to a new version and its performance continuously monitored so that CSPs can satisfy themselves that it provides the same performance as the version it is replacing. If this is not the case, the new version can be quickly rolled back without any risk of service outage. Once all the traffic has been switched successfully, the old version can be retired.

Figure 3.3: A/B testing accelerates time to market and reduces risk [Source: Huawei, 2017]

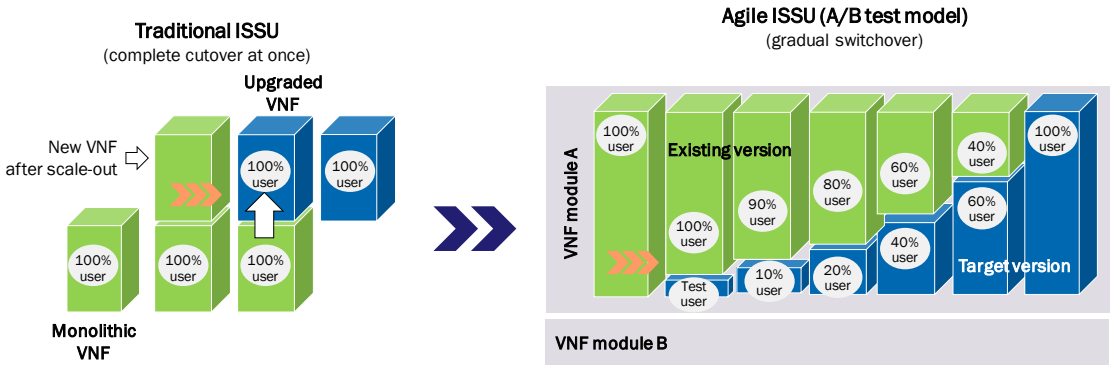

	Traditional ISSU	Agile ISSU (A/B test model)
Service impact	Ongoing service affected	No impact
Operation time	Night (low traffic time)	24 hours
Resources needed	Double	Same in most cases, at most 30% more
Service agility	Difficult	Very good. CI/CD, DevOps, slicing available

The impact of cloud-native automation has been proven in numerous PoCs, field trials and production network deployments. VNFs that have already been onboarded can be launched and configured in minutes; continuous integration/continuous development (CI/CD) tool chains can automate up to 80% of deployment activities and deploy new VNFs, including regression testing, within a few hours. The automation of operations and maintenance (OAM) reduces the risk of service-level agreement (SLA) violations in the first place and accelerates the trouble-to-resolve (TTR) process. Early familiarity with automation and the ability to build it in the right way, through DevOps practices, prepares CSPs to take advantage of full cloud-native computing as soon as they wish to introduce it to the network.

4. Cloud-native network use cases and migration strategy

4.1 When should cloud-native computing be applied?

The decision on when to drive cloud-native computing into the network and for which functions should be based on business factors, such as how strategic the NFV use case is, how agile and scalable it will need to be in future and how long it will continue to be deployed.

New use cases for specific industry verticals and the enterprise sector in general, such as 5G, IoT and network slicing, are of strategic importance. They are the key to the future of the networking business and are expected to generate significant amounts of revenue for CSPs in the coming decades. Their expected longevity justifies investment in building these use cases on cloud-native foundations, and selecting VNFs that address all four cloud-native network principles: stateless service design, service decomposition based on microservices and CUPS, containerised deployment, automated operations using cloud-native orchestration, and a DevOps approach.

When expanding use cases to support mass-market requirements it is essential to manage costs. In this situation CSPs want to bring in virtual versions of network functions they already have in the network to augment capacity and enable new features and capabilities, provide an opportunity to increase operational agility, and to reduce capex and opex. Here, as 5G evolves, it may be appropriate for CSPs to select VNFs that follow a cloud-native software design.

4.2 Applying cloud-native computing: mitigating organisational and operational impacts

The extreme level of automation needed to operate cloud-native network functions requires a different organisation and expertise from that which most CSPs have today. Advanced CSPs are introducing agile DevOps methodologies, the software skills to build automation scripts and new team structures into their organisations as a critical part of their cloud-native journey. The way CSPs implement these capabilities can be very different, however, as what works in one case may not be right for another company. Nevertheless, certain best-practice behaviours are beginning to emerge.

- **Start small.** Some CSPs prefer to create a separate, greenfield organisation, dedicated to implementing the new operational processes and tools needed for cloud-native computing. Initially this is typically a very small ‘tiger team’ which contains the best staff from a CSP’s IT and networking functions. The team builds a DevOps blueprint for cloud-native operations that can then be rolled out across the organisation as part of a larger transformation exercise, as the number of cloud-native VNFs – and the NFV infrastructure on which they run – scales. This transformation exercise typically requires retraining/reskilling and/or the hiring of new operations staff, and CSPs must be prepared to adjust the initial blueprint as the cloud native-capable organisation grows: what works within a single, small team may not necessarily translate to a bigger environment.

- **Seed the existing organisation with NFV-capable staff organised into a virtual team.** Instead of co-locating staff with the right skills in a single, separate team, some CSPs are opting to create virtual teams that span their existing organisational structures. For example, they add people with cloud-native computing skills into teams that already manage the physical mobile core network or services platforms,

enabling them from the inside. This approach has the virtue of not disrupting long-standing reporting structures and responsibilities, and may make it easier to gain greater acceptance for the cloud-native computing approach.

- **Work with vendors as development partners.** The DevOps model needs to be applied differently within a CSP than it does for other enterprises and FANG companies. In a networking context, third-party VNF vendors are the developers and CSPs provide the operations function. The DevOps approach therefore needs to work across multiple company boundaries. CSPs need to take the lead in defining and exposing a coherent operational environment into which VNF vendors can onboard their software. Both CSPs and vendors should work as far as possible with industry-standard automation and modelling tools, orchestrators, APIs and VNF descriptors to reduce the friction between the many NFV ecosystem players involved.

- **Leverage vendor expertise to accelerate cloud-native computing adoption.** Although ETSI NFV ISG envisaged a 'plug and play' environment, where best-of-breed, disaggregated capabilities from different vendors – VNFs, VNF components, NFV infrastructure components, orchestration and management capabilities – can easily be swapped in and out at will, complexity represents a major stumbling block here. Certain highly software-capable CSPs are determined to achieve the original vision of multi-vendor interoperability, but accept that they will move slowly as a result. Meanwhile, CSPs that want to gain results faster may need to implement a full cloud-native NFV stack from a single vendor. Both are valid approaches, and the convergence of cloud-native solutions on key open-source technologies can help to mitigate potential vendor lock-in to a single stack solution, while providing CSPs with rapid access to a leading-edge implementation.

5. Huawei's cloud-native core network solutions

Huawei has incorporated cloud-native principles into its NFV solutions from the start of its development. The company adopted vertical and horizontal decomposition for all core network VNFs. Huawei's solutions that use this architecture are in commercial deployment across the globe. For example:

- A UK operator deployed Huawei's cloud-native CloudPCRF solution in December 2016, cutting the average time for introducing new control policies from 3 months to 2 weeks.

- Huawei and a European galaxy operator have collaborated on multiple cloud-native commercial cases in several OpCo networks
  - The operator's Spanish network has rolled out the world's first standard NB-IoT network based on Huawei's CloudEPC solution and has begun to create innovative business-to-business (B2B) services
  - The Spanish network is also jointly trialling MEC, implementing a CUPS architecture
  - The European operator is also focusing on container-based deployment, working with Huawei to evolve A/B testing and in-service software upgrades (ISSU) for multiple cloud-native core VNFs

- The vendor is engaged with China Mobile and China Telecom on multiple MEC cases, with the latter's Ningbo network an example of a commercial deployment of CUPS.

- A European Tier 1 operator partnered with Huawei to implement the world's first "autonomous" network slicing across an end-to-end 5G network. Huawei's SOC (5G Service Oriented Core solution) was used to support multiple network slices on one physical network infrastructure. Control and user planes were separated and distributed in multiple tiers.

- Huawei is engaged with a UK operator to verify the cloud-native core as the basis for dynamic network slicing and a next-generation data plane.

- The vendor's VNF portfolio is widely deployed in a VM environment. As the business requirements of Tier 1 operators evolve, joint innovation and deployment is helping to prepare Huawei's VNFs such as CloudPCRF, CloudMSE, CloudTAS, CloudEPC, CloudDGW, CloudRGW for commercial deployment in containers as well.

- Huawei has created an open-source Telco DevOps platform and strengthened its tooling support for microservice development and operation. Its CI/CD tool chain aims to simplify multi-vendor integration and lifecycle management, shortening the time from service development to release
  - It worked with a Tier 1 European operator to develop a CI/CD tool which carried out 80% of development-to-integration tests automatically and completed a VNF regression test within 2 hours.

The company also operates six NFV Open Labs, which host over 40 NFV solution providers in a partner ecosystem. Huawei participates in more than 20 NFV-related standards and open-source organisations. It ranks first in terms of the number of project proposals it has submitted to ETSI NFV ISG, and is a platinum member of CNCF, OPNFV, OpenStack and the Linux Foundation. Huawei was also a founder member of Open-O (now merged into the ONAP effort), and is a top developer in the Docker community.

6. Conclusion

The implementation of cloud-native computing in an NFV context needs to be business-driven and pragmatic, taking into account the specific requirements of VNFs. For a variety of business reasons, including cost, it may never be desirable to rearchitect the entire network as a cloud-native software system running exclusively on containers. Further investigation is needed to determine how far VNF components can be modularised, and significant development is needed to ensure that the container ecosystem, including its orchestration, is mature enough to support the requirements of VNFs and network services.

Nevertheless, there are good reasons why CSPs should move on from virtualised monolithic VNFs or should skip this phase altogether if they are at the beginning of their NFV journeys, and procure VNFs that implement the cloud-native principles of statelessness, service decomposition and automated lifecycle management. These reasons include the greater agility and resilience that such VNFs will bring to CSPs' NFV deployments, lower capex and opex costs through cloud-native features such as N-way resilience and orchestration, and the fact that they will prepare CSP organisations for 5G by providing foundational cloud-native capabilities.

CSPs will need to navigate a range of deployment options, and indeed levels of 'cloud nativeness' based on use case, business goals and appetite for risk. Containerised deployment in constrained environments, such as the mobile network edge data centre or customer premises equipment (CPE), may be acceptable but VMs may be needed when the mobile core supports hundreds of millions of subscribers. CSPs will need help and

advice to select an appropriate mix of cloud-native capabilities for a specific use case, and potentially to implement and operate a heterogeneous deployment environment supporting multiple VNFs with different cloud-native deployment profiles.

CSPs which are addressing cloud-native computing in a network context today already benefit from the ability to deliver services faster than the industry average, and they are starting to transform the cost of building and operating the network. They are positioning themselves well to generate new revenue from exciting opportunities opening up in the enterprise market and specific industry verticals. Leading CSPs realise that mastering cloud-native computing early will bring significant competitive advantage, hence the joint innovation and collaboration they are undertaking in a technology domain that is critical to their 5G future.

This report was created with the co-operation of Huawei Technologies Co. Ltd.

About the author

Caroline Chappell (Principal Analyst) is the lead analyst for Analysys Mason's Software-Controlled Networking research programme. Her research focuses on service provider adoption of cloud and the application of cloud technologies to fixed and mobile networks. She is a leading exponent of SDN and NFV and the potential that these technologies have to enhance business agility and enable new revenue opportunities for service providers. Caroline investigates key cloud and network virtualisation challenges, and helps telecoms customers to devise strategies that mitigate the disruptive effects of cloud and support a smooth transition to the era of software-controlled networks. Caroline has over 25 years' experience as a telecoms analyst and consultant.

Published by Analysys Mason Limited • Bush House • North West Wing • Aldwych • London • WC2B 4PJ • UK
Tel: +44 (0)20 7395 9000 • Email: research@analysysmason.com • www.analysysmason.com/research

Registered in England No. 5177472

About Analysys Mason

Analysys Mason is a trusted adviser on telecoms, media and technology. We work with our clients, including communications service providers (CSPs), regulators and end users to:

- design winning strategies that deliver measurable results
- make informed decisions based on market intelligence and analytical rigour
- develop innovative propositions to gain competitive advantage.

We have around 235 staff in 14 offices and are respected worldwide for our exceptional quality of work, independence and flexibility in responding to client needs. For over 30 years, we have been helping clients in more than 110 countries to maximise their opportunities.

Consulting

- We deliver tangible benefits to clients across the telecoms industry:
  - communications and digital service providers, vendors, financial and strategic investors, private equity and infrastructure funds, governments, regulators, broadcasters, and service and content providers.
- Our sector specialists understand the distinct local challenges facing clients, in addition to the wider effects of global forces.
- We are future-focused and help clients understand the challenges and opportunities that new technology brings.

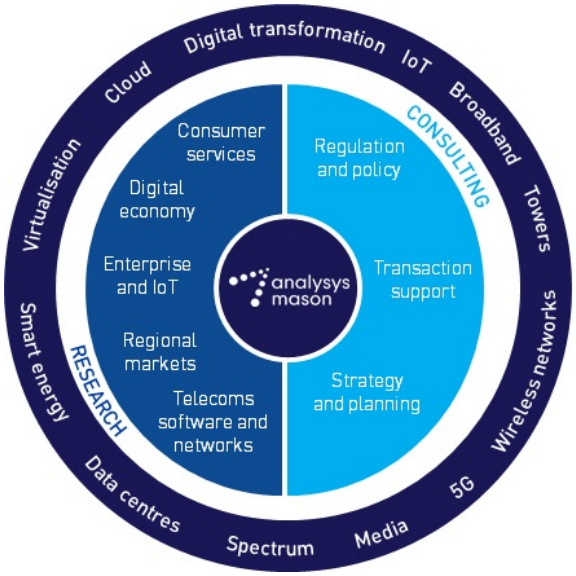

Research

- Our dedicated team of analysts track and forecast the different services accessed by consumers and enterprises.
- We offer detailed insight into the software, infrastructure and technology delivering those services.
- Clients benefit from regular and timely intelligence, and direct access to analysts.

Research from Analysys Mason

We provide dedicated coverage of developments in the telecoms, media and technology (TMT) sectors, through a range of research programmes that focus on different services and regions of the world

The division consists of a specialised team of analysts, who provide dedicated coverage of TMT issues and trends. Our experts understand not only the complexities of the TMT sectors, but the unique challenges of companies, regulators and other stakeholders operating in such a dynamic industry.

Our subscription research programmes cover the following key areas.

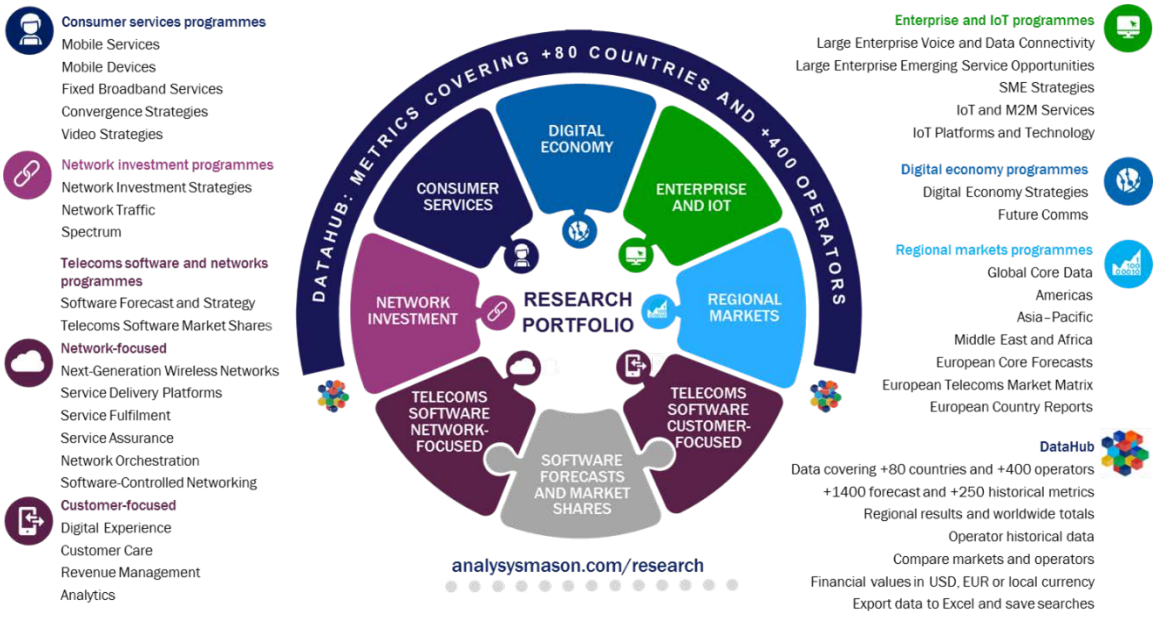

Each subscription programme provides a combination of quantitative deliverables, including access to more than 3 million consumer and industry data points, as well as research articles and reports on emerging trends drawn from our library of research and consulting work.

Our custom research service offers in-depth, tailored analysis that addresses specific issues to meet your exact requirements

Alongside our standardised suite of research programmes, Analysys Mason's Custom Research team undertakes specialised, bespoke research projects for clients. The dedicated team offers tailored investigations and answers complex questions on markets, competitors and services with customised industry intelligence and insights.

For more information about our research services, please visit www.analysysmason.com/research.

Consulting from Analysys Mason

For over 30 years, our consultants have been bringing the benefits of applied intelligence to enable clients around the world to make the most of their opportunities

Our clients in the telecoms, media and technology (TMT) sectors operate in dynamic markets where change is constant. We help shape their understanding of the future so they can thrive in these demanding conditions. To do that, we have developed rigorous methodologies that deliver real results for clients around the world.

Our focus is exclusively on TMT. We advise clients on regulatory matters, help shape spectrum policy and develop spectrum strategy, support multi-billion dollar investments, advise on operational performance and develop new business strategies. Such projects result in a depth of knowledge and a range of expertise that sets us apart.

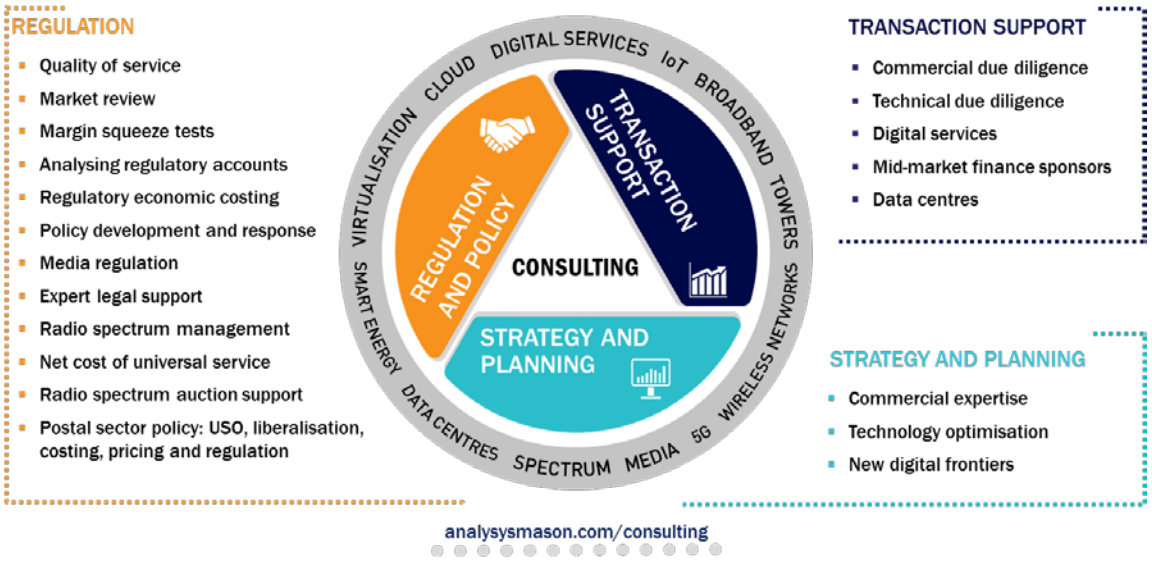

We look beyond the obvious to understand a situation from a client's perspective. Most importantly, we never forget that the point of consultancy is to provide appropriate and practical solutions. We help clients solve their most pressing problems, enabling them to go farther, faster and achieve their commercial objectives.

For more information about our consulting services, please visit www.analysysmason.com/consulting.