

SECRETS OF DATA MANAGEMENT

FOR THE 5G CORE

WINTER 2020

1

CLOUD NATIVE FOR THE 5G CORE – WHAT’S DRIVING THE MARKET?

5

2

ESSENTIAL COMPONENTS FOR A SUCCESSFUL DATA MANAGEMENT STRATEGY

13

3

HOW AND WHEN TO USE HYBRID-CLOUD/MULTI-CLOUD DEPLOYMENT

23

4

WHY OPERATORS SHOULD CONSIDER A MULTIVENDOR 5G CORE

29

5

SUMMARY

35

6

APPENDIX
CONTINUOUS INTEGRATION & DEPLOYMENT (CI/CD) FOR MOBILE OPERATORS

37

WHY READ THIS BOOK?

To say that managing data in the 5G core is critical for the future of operators is no exaggeration. Get it right, and you can achieve 30% savings in OPEX, simplify the network and monetize new 5G use cases.

Get it wrong, and – at best – you face a tsunami of unmanageable data. At worst, you lose ground to Amazon, Google and Microsoft, as OTTs take control of operator data and monetize it at your expense.

Secrets of Data Management provides a coherent approach to managing operator data. With it, you can successfully chart a course through this complex terrain, rapidly monetize 5G and get answers to critical questions such as:

- How can operators successfully partner with OTTs?
- What does cloud native really mean for operators?
- How do public, private and multi clouds fit within 5G data management?

- What are the pros and cons of a multivendor approach to the 5G core?
- How have operators successfully embraced Continuous Integration & Deployment (CI/CD)?

Secrets of Data Management crystallizes the insights that Enea Openwave has gained through 15 years of managing subscriber data in some of the largest operator networks in North America, Europe and Asia. It also includes unique insights by world-leading industry analysts from Heavy Reading and Strategy Analytics.

1

CLOUD NATIVE FOR THE 5G CORE - WHAT'S DRIVING THE MARKET?

By Sue Rudd

STRATEGY ANALYTICS

[Back to contents](#)

When a headline in late 2019 declared, “NFV Is Dead – The Cloud Killed It” the industry went crazy, widely reposting the story declaring that NFV was dead. In reality they said nothing of the sort; it was simply a headline designed to increase readership (and it worked!).

In reality, NFV was the first stage of a paradigm shift to network virtualization. With 5G Service Based Architectures (SBA), however, service providers are getting ready to move to the next stage with true cloud-native software, decomposed service functions implemented as microservices and “as a service” capabilities for networking.

Cloud technologies are fundamentally reshaping the way the industry thinks about many aspects of communications. In this chapter we summarize the background to this discussion, examine the so-called hyperscalers and their agenda, and provide a short tutorial on “cloud native”.

“ Operators that fully embrace a cloud-native 5G service-based architecture (SBA) gain massive benefits. ”

Why is 5G driving service providers to go cloud native?

With 3GPP Release 16, cloud native becomes a requirement of the 5G Standalone (SA) core, and operators worldwide have slowly begun to adopt this new way of thinking as they anticipate new use cases and associated revenues.

Enterprise IT operations and cloud hyperscalers already have a proven robust model for highly automated, cost effective scalability within the data center (DC) by replacing virtual machines (VMs) with cloud-native software for DC compute, storage and I/O. Now service providers are starting to do the same – not only in the DC but across the network.

Next-generation networks must operate as a telco cloud in order to:

- Meet the instantaneous loads that bursty 5G traffic imposes
- Scale telco cloud service instances massively and instantaneously
- Dynamically manage distributed compute, connectivity bandwidth and data stores
- Handle greater complexity
- Automate operations for hundreds of simultaneous services and network slices that demand response times in milliseconds

Specific benefits of cloud-native platforms for mobile operators

Operators that fully embrace a cloud-native 5G service-based architecture (SBA) gain massive benefits. To realize these, operators need to fully adopt cloud-native principles, including: independence of applications from underlying infrastructure and technologies; stateless processing; and decomposition of applications into modular, independently upgradable microservices that scale with distribution of load across lightweight containers. These principles demand new processes – and new thinking – and are discussed further in these pages.

Specifically, the benefits include:

- **Resiliency**, as any processor or storage resource can be replaced without loss of service
- Six 9s **availability** with virtualization and minimal redundancy – n+k (not 2*n)

- **Reduced TCO** with a simplified and modular architecture
- **Simplified interoperability** testing with well-specified interfaces for specific function requests and responses
- **Accelerated time to market** for new service functions/service flows
- **Inherent security**, as service and network functions are authorized to communicate only with other authorized service functions

In fact, when fully implemented, 5G SBA makes possible a highly automated network operations environment that can onboard and upgrade new service software seamlessly without service interruption.

Cloud-native software can efficiently integrate multiple domains, such as public, private and hybrid clouds. Cloud-native deployment also leverages new accelerated processes based on Continuous Integration/Continuous Deployment (CI/CD), in which software blocks are developed and delivered in parallel with integration and testing using automated processes.

¹ *Strategy Analytics estimates that cloud native virtualization will eventually reduce overall TCO for 5G by 28-35%

WHAT IS 5G CLOUD NATIVE?

Separating cloud ready from cloud native

The term cloud native has become overhyped. In particular, the terms “cloudified” and “cloud ready” are sometimes (deliberately!) interchanged with cloud native. Cloud-ready typically means that legacy proprietary software has been converted from structured monolithic code running on specialized hardware to a VM that can run on commercial off the shelf (COTS) hardware allowing limited virtualization of nodes with limited savings. What then is “cloud native”?



How to spot cloud-native software

A true cloud-native approach decomposes software for Service Functions or Network Functions into **microservices**, each of which can be instantiated dynamically as required. Microservices are functional software modules decomposed appropriately for replication and reuse. They are loosely coupled and independently deployable to ensure testability, maintainability and scalability. 5G SBA defines service and network functions at a level appropriate for microservice decomposition, with open interfaces for multivendor interoperability and APIs for third-party software.

Cloud-native apps are often described as **stateless**. This simply means that the state information is stored externally and not embedded as part of the app. By externalizing a copy of state information, the app remains “state aware,” but if a node or software instance fails, it can be immediately re-instantiated and have the service flow restored based on live session-state information. This makes network failures transparent to the user, minimizes the chance of signaling floods following a brief outage and enables high-reliability six 9s services to run on less-reliable three 9s COTS hardware.

The first cloud-native use case: Private 5G networks

Enterprise IT applications on private 5G networks have become the first cloud-native revenue generator as droves of enterprises turn to online cloud services operating across multiple DCs, both on premise and in the hybrid cloud.

Microsoft Azure has successfully positioned itself as the leading cloud provider to enterprises and is now moving to support private LTE and 5G networks for enterprises, including on-premise WiFi and hosted cloud platforms “as a service”. Azure’s acquisitions in 2020 beefed up its range of service platforms and added new services at the edge to draw in enterprises to its cloud platform.

Azure is not alone. Amazon Web Services (AWS) is actively building out 5G edge infrastructure for enterprise customers with Outpost while, in parallel, Facebook’s connectivity division is driving disruptive innovation of its own via its [Telecoms Infrastructure Program \(TIP\)](#).


“Data center hyperscalers got to the cloud first, and will not be content to simply partner with operators.”

What is the hyperscalers’ agenda?

The problem operators face, of course, is that data center hyperscalers got to the cloud first, and will not be content to simply partner with operators.

These OTT behemoths – including Microsoft Azure, Google, Amazon Web Services and Facebook – have an insatiable appetite to “land and expand,” effectively scaling into any new area and taking over, backed by massive funding. Of course, operators have had to contend with OTTs for years, so in some ways this is an old battle. The difference now is a slew of new weaponry they must contend with.

There is even talk of OTT hyperscalers hosting telecom services and becoming unregulated telecom players. Telcos are, therefore, under massive pressure to match both the agility and cost structure of highly automated cloud hyperscalers as they deploy and begin to scale their new 5G networks.



Techcos becoming telcos

WHAT IS 5G CLOUD NATIVE?

The UDM and 5G UDSF²

Special high-performance dynamic data stores are required to store this externalized state data. In 5G, this is the unstructured data storage function (UDSF).


To support cloud networking for subscribers and UE authentication in 5G in parallel, the traditional home subscriber server (HSS) is split between the authentication server function (AUSF), the unified data manager (UDM) and the user data repository (UDR). To meet performance and data integrity requirements, the databases associated with these functions have to be capable of being fully distributed, scaling elastically and rebalancing their information stores dynamically.

² UDM, UDR, UDSF, and Network Data Layer are discussed in greater depth in Chapter 2

Network or common data layer

As the diversity of dynamic and static data types grows, an increasing range of applications, signaling and service flows need access to this stored information. As SDN/NFV rolled out it became essential to provide access via a common or network data layer (NDL). This layer provides a common mechanism for access to diverse legacy data stores and subscriber data. Fast-changing session/state data, subscriptions, policy and configuration data, and subscriber profiles are accessible via the NDL.

According to a recent survey of mobile operators, more than 50 percent of operators plan to move to a common network data layer across their network functions as they roll out 5G.



Advantages mobile operators have over web hyperscalers

Although these cloud providers could theoretically end up hosting a portion of 5G edge infrastructure and become formidable 5G opposition, many enterprise customers dislike that hyperscalers attempt to isolate them and ‘nail up’ customers to their walled gardens. Most would prefer a neutral third party, such as a telecoms provider, to offer a comparable but hybrid cloud that optimizes access seamlessly across AWS, Salesforce, MS Azure, etc.

While it is true that the hyperscalers got to the cloud first, mobile operators have their own arsenal of weapons in this continuing battle. But they need to leverage these advantages by:

- Embracing open ecosystems to accelerate industry-wide innovation at reduced cost through open application programming interfaces (APIs) and open standards

- Delivering best-of-breed solutions that provide the latest and greatest functionality in a timely manner
- Leveraging on-demand software defined networking (SDN)-controlled bandwidth and access to compute and database resources to allow customers to pay only for what they need when they need it
- Promoting the hybrid-cloud approach that the overwhelming majority of enterprises want

Google, Microsoft, Azure, AWS and even Facebook Connect have the potential to become full-service cloud, data and voice service providers. More likely, however, cloud hyperscalers will, once more, try to go over the top and turn operators into commodity bit pipes as they assiduously avoid regulators. In the long run, if operators can leverage the low latency and dynamic service enablement of 5G, the cloud hyperscalers will allow them to be part of the value chain. Hopefully, a diverse ecosystem of hyperscalers, application-specific cloud providers, carriers, data center services and platform suppliers will evolve creatively over time.

“Many enterprise customers dislike that hyperscalers attempt to isolate them and ‘nail up’ customers to their walled gardens.”




Sue Rudd
STRATEGY ANALYTICS

Sue has over 30 years' experience working with Mobile Telecoms and Internet service providers. At Strategy Analytics she focuses on matching new technology to business opportunities for 5G Network Slicing and Service Based Architecture (SBA), SDN/NFV, Multi-Access Edge Compute (MEC), Video Delivery Optimization and Telco Cloud as well as OpenRAN/ vRAN, HetNets, Small Cells and Wi-Fi interoperability.

WHAT IS 5G CLOUD NATIVE?

More than 50% of operators plan to move to a common network data layer for 5G



Source: Technology Innovation Council

Cloud native containerization for efficiency, flexibility and lowered costs

The 5G SA core has been designed to be truly cloud native based on containerized microservices for cloud-native service and network functions, while legacy VMs can also be containerized and continue to operate in parallel for pre-5G services under 5G non-standalone (NSA).

Containers are executable units of software that package together application code, libraries and dependencies in common ways so they can be run anywhere. The difference between a container and a VM is that VMs need a complete operating system (OS) installed to support the application, whereas a container packages the application software with any application-specific OS and leverages it on the host's OS kernel.

Microservices are rapidly replacing traditional VMs based on conventional monolithic software that frequently correspond to pre-5G network node.



2

ESSENTIAL COMPONENTS FOR A SUCCESSFUL DATA MANAGEMENT STRATEGY

The service-oriented, cloud-native nature of 5G networks requires a new approach to managing operator data. In this chapter we examine the data to be managed and some of the tools needed to do this successfully, in particular the network data layer and the unified data manager.

[Back to contents](#)

What data needs to be managed?

Operator data includes:

- **Subscription data:** Stores the type of services the subscriber or a device is allowed to use on the network, as well data pertaining to subscribers and their identity
- **Policy data:** Sets the priorities, rules and constraints about how services can be accessed and used
- **Session data:** Context data that is created on the network as subscribers or devices connect and use different services
- **Application data:** Data that is specific to an application. For example, the Network Repository Function might store data about the services that are registered in the network, etc
- **Configuration data:** Optional data required to configure network functions

In particular there are compelling use cases that depend on effective use of mobile identity, such as knowing who is allowed to connect to the network, how they are connected and provisioned, etc.



If you get this right, data can lead to knowledge, which creates multiple advantages, including network efficiencies; a seamless experience for users across devices and sessions; enterprise mobility; intelligence on how your network is performing; and the ability to deliver identity-based personalized services.




Get it wrong, and you risk missing the boat in this rapidly expanding market. Worst of all, potential signaling storms can occur as functions start/restart but cannot find their data quickly.

So how do you get it right?

A coherent approach to managing operator data

3GPP 5G architecture is open and service based, and as such, it should be embraced by operators to maximize agility and avoid vendor lock in. Services share data and can be combined based on open registration capability. Sadly, operators haven't fully accepted this, so end up recreating the same monolithic functions that also remain in data silos, just with different interfaces.


The new architecture is an opportunity to review processes and best practice. In particular, from the experience of Tier-1 operator deployments worldwide, two functions are critical to this process. The first is the network data layer, which provides user data repository (UDR) and unstructured data storage function (UDSF) functionality. The second is unified data management (UDM). Both were introduced in Chapter 1, the following sections clarify the value each brings to 5G architecture.



“Signaling storms can occur as functions start/restart but cannot find their data quickly.”

How does a data layer help?

In a 5G cloud-native solution, applications must be allowed to be stateless and separated from the processing of their data. Not separating application logic from application data causes widely acknowledged problems associated with vendor lock in, where large vendors seek to control applications (and the network) by keeping data embedded within their application, often in a proprietary format. The good news for operators is that in a cloud-native environment combined with REST API-based interfaces, it is easier to separate data from applications in a way that simply wasn't possible before by using the network data layer (NDL), also referred to as a common or shared data layer as the common backend for all stateless applications.



Commercial and technical advantages of implementing a data layer


The data layer:

- 1. Provides the operator with the means to truly own the data schema**
Be it structured (subscriber data) or unstructured (session data), a data layer puts the data back into the hands of the operator.
- 2. Enables efficient automation, scalability and resiliency**
Subscription data and subscriber state are externalized and stored in the NDL. Network applications operate statelessly (write once, read anywhere). Any network application instance can service any subscriber – typically saving 30% of server resources.
- 3. Creates network simplification by reducing coupling of IT systems and eliminating data duplication**
A network data layer decouples application network functions from their data, thus avoiding fragmentation or islands of data. It enables new network applications to be added with faster time to market.
- 4. Enhances reliability and protection from outages**
Signaling storms are often the cause of catastrophic network outages. Because the data layer is de-centralized, session state and subscriptions can be synchronized, safeguarded using overload protection, and made available across the network. Six 9s reliability over three 9s commodity hardware has been proven in the field.
- 5. Enables the edge to be subscriber aware**
The data layer supports multitier deployment scenarios and can facilitate the dynamic provisioning of subscriber data close to the edge in an efficient manner, such as by intelligently replicating only the data that is needed at the edge.

- 6. Enables sharing of data across network slices**
The data layer is slice aware and enables partial and full replication across slices, as well as data isolation within a slice with specific access control.
- 7. Allows responsible sharing of data**
The NDL provides an interface that makes subscriber, session, application and policy data accessible in near real-time, while also sharing it with authorized, onboarded applications.

See [The 6 Benefits of a Network Data Layer](#) for a complete explanation of the benefits of deploying a data layer.


What are the Most Important Factors for Choosing a Network Data layer?



Source: Technology Innovation Council

As shown previously, more than 50 percent of mobile operators in a survey stated their firm intention to implement a data layer. The reasons varied; however, the most common reason was to enable a unified view of data – see accompanying figure.

ADVANTAGES OF A NETWORK DATA LAYER



Concerns over implementing a data layer

In discussing the data layer, concerns may arise, including:


- 1. “I don’t have the resources to change all my data stores.”**

Introducing an NDL does not mean changes to data stores. It means having an API that can deal with legacy data stores and help migrate to a new world of cloud native.
- 2. “I don’t want to get locked in to a hyperscaler.”**

Operators gain openness and APIs, but there is a fear of handing over the keys to AWS, Google, Microsoft or Facebook. However, the operator retains ownership of the data schema and is able to make changes as required. Additionally, regulations such as GDPR tend to favor operators, since hyperscalers have little appetite for such responsibility.
- 3. “The NDL is not feasible for my network situation.”**

Operators have concerns over latency, scale and data synchronization, but these all improve with a NDL. For example, a decentralized data store is capable of millisecond response time with standard interfaces. The alternative is fragmented data islands with variable response times; such a system can only move as fast as the slowest island.
- 4. “Is the NDL commercially rolled out and battle tested (since no one wants to be first)?”**

NDLs with zero downtime on COTS hardware are operationally ready and commercially rolled out in multiple Tier-1 operators.



One common NDL delivers all use cases

How does unified data management help?

The UDM works in tandem with the NDL, performing user authentication for network functions. The UDM is abstracted from hardware so it can be scaled and updated without the need for manual configurations. It can be stateless or stateful, depending on the network architecture. As discussed in Chapter 1, the UDM is the equivalent of the home subscriber server (HSS) in 4G networks. The UDM communicates with different network layers and functions to share user data across the network.

Like the 4G HSS, the UDM stores customer profile and authentication information, provides keys for encryption of the information, and supplies it to other functions that manage mobile network sessions and connections. It is responsible for all user identification handling, access authorization and subscription management in 5G. Importantly, the architecture of the UDM is software defined, cloud native, and often stateless; storage is abstracted to the user data repository (NDL), and the function is run by microservices.

Best practice in selecting a UDM

Ensure your UDM:

1. Is cloud native

The 5G core has been designed for cloud deployment, with all the tools and features of a telco-grade cloud. The UDM must make use of cloud concepts like stateless design to scale effectively and to achieve high availability. A cloud-native UDM can scale to manage the data from billions of connected 5G devices. Furthermore, a cloud-native UDM can be managed with the same tools as the rest of the 5G core.

2. Scales core to edge

The most advanced and flexible UDM architecture uses a stateless front end and a data store back-end data layer or NDL. The UDM can scale to be at the edge with the other 5G control plane functions and use the core data layer capability to transfer and share data updates.

3. Interworks with existing 4G HSS

Ensure it can interwork with an existing 4G HSS. Replacing an existing 4G HSS has not proven to be easy, since service providers have usually added business logic to their HSS. As the UDM essentially replaces the HSS in 5G, it is another candidate for adding non-standard functionality and new business logic.

In addition to the above principles a leading UDM solution will offer the following:

- Flexible support for 5G use cases and interworking with 4G
- Ability to scale efficiently and handle demanding network sizes and use cases
- Powerful rules engine for introduction of customized rules and interfaces
- Rigorous implementation of the latest standards and full 3GPP compliance

3

HOW AND WHEN TO USE HYBRID-CLOUD/ MULTI-CLOUD DEPLOYMENT

Subscriber data management has kick-started migration from silos of data for individual applications to a network data layer with single provisioning, where the same data is available for all applications as they need it. But why would an operator want to expand this to a hybrid cloud, particularly a multi-cloud – or a hybrid heterogeneous cloud with multiple vendors like AWS, Azure, or Google?


[Back to contents](#)

Why consider hybrid-cloud solutions?

- **Flexibility in deployment**
Public cloud resources can provide flexibility in periods during which there is significant provisioning work to be done but no scope for reducing georedundancy. Such work could make use of temporary compute and storage, were the application able to run in connected clusters, both on private and public clouds.
- **Fast partner access to data**
An operator might have a need to expose subscriber data to external applications. This could be done within the operator's own data center; however, the operator might have reasonable concerns about providing access and the privacy and resilience of its system. These concerns can be met by providing an external cloud IaaS availability zone where only *relevant* data is replicated.
- **Enabling 5G private enterprise networks**
There are other uses for the remote partial data set to be synchronized with the operator's main data store. For example, private 5G networks for enterprises, which are among the early use-cases for 5G, may well be deployed such that much of their own infrastructure is in a private (or even a public) cloud. For example, where these are being deployed as a slice on the network, it makes sense for the relevant information to be made available "locally" but to eventually be synchronized in the network to minimize the real-time communication, while also enabling a seamless transition from local enterprise to truly mobile.

- **Partnership for mobile edge computing (MEC) deployments**
Hyperscalers are clearly developing their mobile edge computing strategies. As such, operators need to ensure that the discussion includes them and their requirements. It is clear that there are different approaches being taken, with some operators considering exclusive partnerships. However, we anticipate that most operators will not want to restrict themselves but open the door to numerous partnership opportunities.

The multiple cloud approach, meanwhile, provides flexibility that can lead to better value through more competition, as well as the ability to work with more partners and enterprises. But doing so raises the question of who to partner with.



5G operator solutions could require a hybrid cloud mix


Who should I partner with?

- For many operators, an established or incumbent infrastructure provider may appear an attractive proposition, but there are factors to take into account, including:
- **Established infrastructure as a service (IaaS) vendors**
Most players in this field are now providing hybrid on-premise/public cloud variants. They may have attractive offers with particular software (really becoming more SaaS), but that might not enable the full flexibility that operators require, nor would it generally enable deployment across the widest available range of providers. Operators will typically want to have their own IaaS strategy that involves multiple cloud vendors to provide flexibility and value.
 - **Aggregators/API providers**
These provide good flexibility and resource management. They typically do not provide software directly but have partnership programs that identify software that works well with their hybrid infrastructure. Typically, the management across infrastructures is consistent and avoids management overhead of using multiple IaaS approaches. They are also increasingly able to provide hybrid public/private deployments with on-premise and public cloud orchestration.

“ Operators will typically want to have their own IaaS strategy that involves multiple cloud vendors. ”

- **Software vendors**
The right software will deliver highly flexible solutions using standard open interfaces, which can be deployed on many different platforms while retaining operator control over data. They can help with deployment of a multi-cloud approach and may provide flexible licensing in terms of where the software is deployed, but will not typically provide the infrastructure itself.
- **Integrators/large solution vendors**
Where these are offering appropriate software, they may be able to provide an infrastructure that can neatly integrate externally with suitable public clouds, sometimes also working with aggregators to provide a good overall solution. Operators should ensure that the software platform being offered is able to meet the use cases required.

Whichever partner is chosen, there are ways of approaching hybrid deployments that will help.



Partner solutions require common operator controls

First, consider the platforms that you want to support, and initially deploy small scale for specific use cases or customer segments like Internet of Things (IoT) or cloud gaming. Consider the costs and how to keep them updated - which data centers/availability zones offer you the best value? With this in place, when you need to provide additional capacity, it is simply a matter of turning on more nodes. It gives you the opportunity to rehearse scale-out before you are forced to.

Additionally, it is crucial to consider the facilities that your partners use. What is the value to them of rapid access to data? Being within the same accessibility zones in the same IaaS vendor can be a significant benefit from a latency perspective, as well as typically reducing costs for network access.

Concerns and challenges of a hybrid-cloud development

There are genuine concerns that operators considering hybrid-cloud deployments will have. Based on our experience, adopting the following principles will mitigate these issues and provide the additional advantage of flexible network right sizing.

Avoid forming a working relationship with only one IaaS

You should not be tied to one specific vendor. Ultimately, lock in limits best value and leads to other issues, such as lack of ownership of the data-model.

Ensure you retain the ability to update data models

Over time, what information should be stored in operator subscriber records is likely to change. An open design schema owned by the operator ensures that you are not at the mercy of expensive customization services operated by an incumbent vendor.

Ensure you retain control of your data

Operators need to ensure that infrastructure providers have the ability to ensure data stays within the geographical jurisdiction demanded by regulators and that access rights are controlled.

In addition to the above, there are specific technical principles to adhere to, including:

- **Keep your options flexible by ensuring your data is already synchronized**

When deploying small-scale, the best flexibility is provided when the data is always kept updated (or very nearly so) within an accessibility zone. This means keeping a minimal deployment running and synchronizing against your data. Then if it is necessary, you can pass some processing to that area so it does not need to be duplicated across the internet before it can be used.

“
Rehearse scale-out before you are forced to.
”

- **Data replication across multiple cloud deployments**

We mentioned above the possibilities of replicas of part of the data being provided in a public cloud for partners to use. In that case, consideration could be given to how the data should be replicated. Is there a benefit to reducing the amount of traffic across a public internet link by only replicating the full traffic to the small-scale replica and then using that as a source for the part of the data identified?

- **Resource orchestration in a multi cloud deployment**

Managing multiple infrastructure vendors could certainly be challenging, not least because their interfaces vary significantly. Operators may learn to perform administrative tasks on multiple platforms as they all, for example, provide varying mechanisms to adjust sizing. The risk is that it becomes so much of a challenge to perform when necessary that you leave a system wrong-sized when you want to scale back in and end up with substantial extra costs. An operator may benefit from using a consistent approach, such as that provided by an aggregator.

Heterogeneous clouds and the network data layer

There are significant advantages for operators to deploy a data layer, which can utilize a heterogeneous cloud overlay, including:

- Flexibility in cost and sizing. If more capacity is required than is available in the private cloud, a heterogeneous deployment provides rapid relief. This reduces platform costs and (subject to regulatory considerations) allows the use of whichever service provides the best value.
- Rapid access for partner APIs to co-located data, while minimizing the risk to the core network. Rapid access for APIs (or network slices) could be located in multiple infrastructure providers and even in partners' private clouds.
- A model that can be developed for replication of information and data to the edge.

Operators gain many benefits from hybrid-cloud solutions across multiple hyperscalers. Even if deployment across multiple hyperscalers is not anticipated for some time, it is critical to make decisions early in order to enable such deployments later.

“
It is critical to make decisions early in order to enable such deployments later.
”

4

WHY OPERATORS SHOULD CONSIDER A MULTIVENDOR 5G CORE

By Gabriel Brown



[Back to contents](#)

To develop their 5G networks and service offers, operators need to modernize the core network and migrate to standalone (SA) operation using 5GC.

This chapter, by Heavy Reading's Gabriel Brown, discusses the transition to cloud-native 5GC deployment. The focus is on the importance of open, multivendor 5GC, with specific reference to data management functions that can be deployed as part of a best-of-breed (multivendor) core network. In particular, it highlights the importance of an independent network data layer (NDL).

Multivendor 5G core


The mobile core is made up of several key domains. The packet core, which manages sessions, mobility, policy and subscriber authentication, is the heart of the system (this remains the case in 5G). It is surrounded by the IP multimedia subsystem (IMS), SGi-LAN functions (network address translation, firewall, etc) and various other data management, IP networking, application, signaling, and security functions. Classically, a mobile operator will use a few different vendors split roughly along the lines of these major functional areas. These same domains will be present in 5G networks and are likely to be similarly multivendor in nature.

In terms of the 5GC specifically (aka the 5G packet core), there are several options for multivendor deployment. Being cloud-native and based on standardized functions that communicate over standardized APIs makes it simpler and lowers risk to deploy a multivendor system. This enables operators to select functions - or more likely, closely related groups of functions - that are best-of-breed and can be deployed and managed on common cloud infrastructure using common orchestration and management tools.

How many operators are considering multivendor?

Heavy Reading research indicates that multivendor 5GC is attractive to operators. The figure shows the results of a question in Heavy Reading's 5G Network & Services Strategy Operator Survey. In this survey on multivendor 5GC, the largest number of respondents (48%) say they are "likely to use two or three vendors to assemble a 5G core" and the next largest (26%) say they are likely to use "multiple vendors to create a best-of-breed core." This makes a combined 74% of operators that intend to pursue a multivendor 5GC strategy.

Thinking about your 5G core network, do you plan to assemble the functions that make up the service-based architecture (SBA) 5G core from multiple vendors or from a single vendor? (n=150)



- Likely to use two or three vendors to assemble a 5G core (48%)
- Likely to use multiple vendors to create a best-of-breed 5G core (26%)
- Don't know / too early to say (15%)
- Likely to use a single vendor (11%)

Source: Heavy Reading's 5G Network & Services Strategy Operator Survey, Q1 2020

The survey results give a useful sense of the extent to which operators are keen to pursue multivendor 5GC design. However, it is also useful to keep in mind that multivendor comes with challenges around system design, interoperability, operations, troubleshooting, vendor management, and so forth. More granular vendor selections for true best-of-breed may be better suited to larger operators with greater in-house capability. Smaller operators with more limited integration resources may prefer multivendor strategies where they select pre-integrated groups of functions.

“74% of operators intend to pursue a multivendor 5GC strategy.”

Options for multivendor data management in the 5G Core

One way to build a multivendor 5GC is to select data management functions from specialist providers. There are several different scenarios for best-of-breed 5G data management vendor selection being discussed and evaluated in the market, including:

- **Fully independent NDL:**
Incorporates a next-gen database layer capable of supporting structured data formats (such as a UDR, the master database) and unstructured data (such as the UDSF to store dynamic state). This model can be adopted by any size of operator and is being most actively considered by Tier 1 operators/operator groups as part of a strategic data management capability. Frontend functions, such as the UDM, AUSF and EIR, can be selected from an alternative vendor. There should be no lock in between the database layer and the frontend applications.
- **NDL with frontend functions:**
In this scenario, the operator selects the database layer (such as UDR+UDSF) and the key frontend data management functions (such as UDM+AUSF+5G EIR) from a single vendor. This is attractive to larger and midsize operators that want the flexibility and optionality of multivendor 5GC but without having to integrate everything in-house. However, it remains critical that the interface between the frontend applications and the NDL remains open because this means the operator retains the option to swap frontend vendors at a later stage.
- **Pre-integrated 5G data management:**
Taking pre-integration a stage further, operators may seek to add other related functions to a combined UDR/UDSF/UDM solution, including for example, the policy control function (PCF) or network exposure function (NEF). This would not account for the entire control plane — specialist access and mobility function (AMF) and SMF could be sourced independently — and the user plane would not be included. However, it would account for around one-third of the 5GC functions from a single vendor. Again, it remains critical that there is an open interface to the NDL.

Integration between 4G and 5G Core

Today's 5G networks are deployed in NSA mode. The majority of subscribers in these networks are currently LTE subscribers, and 5G subscribers frequently roam from areas of 5G coverage to LTE. Therefore, the 5GC must interwork with the 4G core and, perhaps, the 4G RAN. There are several different options for this interworking, ranging from a fully integrated common core to physically and logically distinct 4G and 5G core networks.

Heavy Reading believes operators will seek to migrate to cloud platform environments that enable them to support 4G and 5G core networks on common infrastructure. In time, some operators may also migrate to integrated 4G/5G core functions where appropriate. These transitions will occur over an extended period and, in practice, the market is likely to be characterized by diverse 5GC implementation and migration choices. Not all technology decisions are taken on day one, which emphasizes why it is important that operators maintain the flexibility to introduce new capabilities and vendors throughout the 5G cycle. This is another reason to pursue an open deployment model and to avoid a closed NDL interface.

Taken from [“Multi-Vendor 5G Core: Best-in-Breed Subscriber Data”](#) by Gabriel Brown.



Gabriel Brown
PRINCIPAL ANALYST –
MOBILE NETWORKS & 5G

Gabriel Brown leads mobile network research for Heavy Reading. Starting from a system architecture perspective, his coverage area includes RAN, core, and service-layer platforms. Key research topics include 5G, LTE Advanced, virtual RAN, software-based mobile core, and the application of cloud technologies to mobile networking.

5

SUMMARY

This book has focused on one of the most fundamental aspects of an operator's business and suggested a coherent approach to managing the most valuable asset an operator has: subscriber data. In 5G, the concept of a subscriber now expands to 24/7/365 always-connected intelligent devices, enterprise private networks and consumers.

[Back to contents](#)

Indeed, 5G subscribers could eventually number in the billions, therefore, managing this data (subscriber, profile, application and live session data) should be a focal point for any CXO's strategic thinking for network services architecture. Simply turning a blind eye to this data or taking a fragmented approach is not commercially viable.

In summary:

- 5G data management is a strategic, imperative business issue. Mobile operators have advantages over hyperscalers, but a coherent approach to data management is critical. Managing data so that it is available at point of need, accurate, synchronized and private, is foundational for a robust data strategy.
- Hyperscaler hybrid-cloud environments can enable innovation and scale but should also be considered from the perspective of data management, orchestration, APIs and long-term operation. There are important lessons to consider when deploying hybrid cloud environments and choosing partners.
- Taking steps to CI/CD is both necessary and achievable and will drive technical and business agility.
- Open standards and interoperability give operators real choices. Multivendor strategies require thought but in the longer term will bring numerous advantages, while the cloud native data layer gives operators vendor choices like never before.

The implementation and rollout of 5G service-based architecture is not merely an occasion to change radio access. This really is a once in a "G"eneration opportunity to examine operational processes, data management and the important integration of the core and the edge. Critically, it provides a solid foundation for new use cases to monetize 5G effectively and rapidly.

“Mobile operators have advantages over hyperscalers, but a coherent approach to data management is critical.”

6

APPENDIX

CONTINUOUS INTEGRATION & DEPLOYMENT (CI/CD) FOR MOBILE OPERATORS

[Back to contents](#)

In this Appendix, we widen our remit beyond pure data management and summarize the recent move to CI/CD working in our industry and how mobile operators can benefit.

The bywords of the large internet players in terms of software and application software rollout is a nightly rollout of software changes known as continuous integration and deployment. Translating this to a 5G telecom environment is problematic and requires a great deal of consideration.


The 5G telecom environment is different: infrastructure may come from multiple vendors that may be on different pipelines and the first thing to note here is that continuous integration and continuous deployment cannot mean continuous chaos for operations teams. Large internet players have the advantage of controlling the deployment, the application APIs and the supporting infrastructure. In the telecom environment, this is more of a coalition.

Steps for mobile operators to get to CI/CD


1. Set clear objectives for the new system

In effect, examine the reasons for CI/CD in your environment. After doing so, you may conclude, “If it ain’t broke don’t fix it!”. The goal for process changes must always be to improve time to deployment for new features and patches. The traditional operator model of deployment is diligent but can take 6-12 months for major version software rollout in a multisite architecture. The tools and processes have to mature to an open integrated environment. A timeframe reduction to weeks for software rollout instead of months should be a realistic target.

VISION



REALITY



Continuous Integration & Deployment

2. Define the scope

The next step is to set the scope of where the new process will be applied and examine the gap between vision and reality. The perfect Dev Ops environment (illustrated on the left) is great for a single vendor controlled environment, but the reality for production is more like the right-hand side: a jigsaw of stakeholders, interfaces, timing and tasks – all with the basic watch words of test, test, test to build a trusted, repeatable system within a live operations environment.

3. Select your CI/CD tools

In taking the first steps to CI/CD, we recommend agreeing to the tools and installing the pipeline operation (lab, pre-production) across multivendors early on. Consider the best in open common toolsets developing rapidly – Jenkins, Ansible, Pytest/xtesting and Allure to name a few. These enable a better understanding in virtualized deployment models with containers and orchestration (and their supporting layers of Kubernetes, docker, Grafana, Prometheus). The tools and functions for the pipeline have to be open so that multiple vendors can be on-boarded and there is no lock in.

4. Integrate CI/CD into your mainstream processes

The next collaborative step is the harmonization and integration of the process (quality gates, testing, etc). Fundamentally, there must be versioned interfaces, a schema model and an active interoperability program outside of the operator’s environment. In respect to the quality gates and process, this is an inflexion point to consider change to operational processes. Some of the major benefits of moving to virtualization – not just in common hardware and infrastructure but also in changes to process and operations – have been seen here. At each stage of the pipeline, there must be a verifiable metric. It may seem like common sense but coordinating this in a multivendor environment to create automated processes is a challenge.



“Continuous integration and continuous deployment cannot mean continuous chaos for operations teams.”

CHECKLIST FOR ESTABLISHING AN EFFECTIVE CI/CD MODEL SUITABLE FOR 5G ROLL OUT

There are a number of key success factors in establishing a collaborative CI/CD model suited to a 5G rollout. The following is a quick, non-exhaustive checklist specifically from the perspective of subscriber data management/UDR.

- ✓ Does the vendor follow an Agile DevOps process model?
- ✓ Versioning – versioned access & admin interfaces, versioned data schema
- ✓ Standards support (specific versions of 3GPP specifications, lists non-supported items, etc)
- ✓ Automated testing at each stage of the lifecycle
- ✓ Security screening (for example, only authorized software can be deployed at the end of the pipeline, vulnerability scanning, adoption of operator defined processes, etc)
- ✓ Solution interdependency (Interfaces, versions, IAAS, NFVI/CAAS versions)
- ✓ Scalable sizing, pre-defined automated scale
- ✓ Clear tested resiliency & recovery plans (because things will break!)
- ✓ Audit trail for schema, data and functional interfaces
- ✓ Clear verifiable interoperability
- ✓ Clear key performance indicators
- ✓ End-to-end test environment for solution (mirroring production environment & configuration)
- ✓ End-to-end common pipeline environment

The vendor processes must also align to make this work in a telecom environment. The role of CI/CD could be assigned to a system integrator, but verifiable oversight is needed, as this is effectively the core of the business and the keys to drive it should not be handed to an integrator.

As highlighted earlier the scope/roadmap of the transition to CI/CD should be set clearly; not all processes are suited to all aspects. For example, the introduction of canary testing into a data model/schema for a 5G UDR system should be considered carefully, because by nature, data that is synchronized systemwide for consistency and a canary test of data could fail. The result of that combination could be set changes in the overall system that are difficult to roll back.

In summary, a move to a CI/CD model is worth the effort; however, because the environment is a mix of vendors, infrastructure, processes and pre-existing systems, it needs to be scoped, phased and accompanied by verifiable testing at its core before it can be operationalized.

VENDOR



Vendor oversight needed even with a system integrator

ENEA CLOUD NATIVE DATA MANAGEMENT FOR THE 5G CORE

5G core networks require a new approach to the management of data, thanks to their service-oriented, cloud-native nature. Enea's complete 5G Data Management portfolio stores and manages data across all 5G core and edge functions, supporting multivendor 4G/5G interworking.

Our cloud-native suite spans the common network data layer (NDL), scaling the control plane with critical 3GPP functions including UDM, UDR, UDSF, AUSF, PCF and EIR.

Tier 1 operators in North America and Europe have selected Enea 5G data management based on features such as:

- Clear separation between network data layer and applications, avoiding vendor lock in
- Open orchestration & automation, enabling choices independent from the 5G core vendor
- Platform agnostic, with support for any PaaS, private cloud or public cloud
- Software architecture purpose built and optimized for 5G

FIND OUT MORE

Stratum Network Data Layer

Stratum is a 5G network data layer that is open, cloud native and simplified.

Stratum solves the problem of vendor lock-in by collapsing all your vendor data silos into one common network data layer.

Enea Unified Data Manager

The Enea Unified Data Manager provides UDM functions in 5G networks and supports interworking with HSS in 4G networks. It supports all relevant interworking scenarios and manages all subscriber and device data. It can also perform the 5G-EIR function and be colocated with the AUSF function.

ENEA
Openwave Division

For further information on the above or advice on any of the topics covered in this book, please contact us.

Email info@owmobility.com

owmobility.com