

Independent market research and
competitive analysis of next-generation
business and technology solutions for
service providers and vendors

**HEAVY
READING**
**WHITE
PAPER**

Heavy Reading's Cloud-Native 5G Core Operator Survey

*A Heavy Reading white paper produced for
F5, HPE, Mavenir, and Oracle*



AUTHOR: GABRIEL BROWN, PRINCIPAL ANALYST, HEAVY READING

INTRODUCTION AND KEY FINDINGS

This report presents the results of the Heavy Reading **Cloud-Native 5G Core Operator Survey** conducted in December 2020 and January 2021. The focus is on the 3GPP packet core, associated cloud infrastructure platforms, and end-user services. “Cloud native” is not tightly defined in the context of this report, but is generally taken to mean network functions deployed in containers and centrally orchestrated.

Key findings

The deployment of 5G core networks, and the introduction of standalone 5G services, is now underway and gaining momentum. The survey indicates widespread deployment of 5G core in advanced markets over the next two years. A small majority of respondents (57%) say their company will have 5G core deployed and in live service by 2022. This does not represent the global average view, but the view of the advanced operators represented in the survey. Clearly, it is still early in the cycle, and it will take time for 5G core to reach the global mainstream. Nevertheless, 5G core is now underway and set to scale rapidly.

A slim majority of respondents expect 5G core to produce “clear benefits to the customer experience” for consumers (52%) and enterprises (57%). This is not quite a wholehearted endorsement of the service benefits of 5G core because there is some doubt among just under half of respondents. However, this doubt does not appear to be acute: 37% (consumer) and 36% (enterprise) think the experience will “probably” be better even if “it is not yet clear how significant this will be.” The overall picture, therefore, is one of optimism about the impact of 5G core on the service experience, tempered by some caution about what it will mean in practical terms.

A majority (64%) say the primary reason to deploy 5G core is to “offer advanced new services” to customers. This is a logical and expected result. It is perhaps a little surprising that the lead over “to modernize infrastructure” (32%) was not greater. It is, however, worth keeping in mind that the supposed new services that 5G core may enable largely do not yet exist commercially and it is, therefore, reasonable to be cautious. Moreover, there is an argument that operators need to modernize core infrastructure before they can offer improved or new services. The overall result, therefore, is clear and unambiguous: services are the driver for 5G core.

There is cautious optimism about the commercial value of network slicing, but not outright certainty that this service type has a strong business case and will be a commercial success. The largest group of respondents (48%) think network slicing is an “attractive” proposition but are “not certain of the business case.” That only 30% believe there is “a strong business case” and that network slicing is “very attractive” indicates there is uncertainty on the commercial outlook for this service type.

A large majority (86%) expect their company to operate a common 4G/5G at some stage. The largest group of respondents are the 46% that say it is “essential to have a converged 4G/5G core” just ahead of the 40% that say it is “important and desirable, but not essential to start with.” To achieve a common, converged core within the aggressive deployment timeframes identified is challenging and it is possible that respondents are over-confident or have a permissive definition of what constitutes a common core. It may be more appropriate to think of common network functions or groups of network functions

in the first instance. However, this is a complex topic, with many variations and phases. And again, the overall direction is clear: operators want common 4G/5G core networks.

There is an appetite to use new vendors for 5G core. The survey does not indicate across-the-board vendor swaps are likely, but it does show that operators will take the opportunity to evaluate new suppliers. The largest group is the 30% that expect to use “a mix of new and current vendors.” A quarter (24%) say they will “introduce new 5G vendors and retain existing 4G vendors,” while only 14% expect to use “new vendors across the 5G and 4G core.” With no dominant single answer, the survey indicates a wide variety of vendor selection strategies are in play.

A strong 51% say they are likely to use “two or three vendors to assemble a 5G core.” This is a clear lead over the 26% that are “likely to use a single vendor” and the 17% “likely to use multiple vendors to create a best-of-breed 5G core.” The overall picture is that there will be vendor diversity in the 5G core, and that new entrants have opportunities, but also that the transition will not result in a radical industry-wide overhaul of the vendor landscape. Only 27% of respondents say their company is well prepared in terms of internal skills to integrate a multi-vendor best-of-breed core, whereas a majority (54%) say they are prepared for a pre-integrated solution they can “evolve over time.”

The choice between infrastructure as a service (IaaS) and platform as a service (PaaS) is one of the most critical and far-reaching decisions operators must take as they plan 5G cores. There is a preference for an IaaS model (45%) in the survey base versus 32% for PaaS and 23% for a vendor-integrated full stack. This indicates that both IaaS and PaaS models are likely to be used over the near and medium terms. In terms of future vision, PaaS models and the container as a service (CaaS) variant are generally expected to prevail over the longer term and might have been expected to score more highly. This result, therefore, perhaps suggests that cloud infrastructure vision and execution are not yet quite aligned.

Roughly a quarter of respondents believe their company is ready for large-scale operation of cloud-native 5G core. This, of course, means three quarters are not yet ready. Over a two-year view, however, nearly all respondents (90%) think their company, and the industry at large, will make sufficient progress in infrastructure, applications, and internal skills to go live with large-scale operation. The next two years will be critical to 5G core technology and process development.

Operators will be diligent and purposeful in their adoption of 5G core automation and not “gung-ho.” A clear majority (69%) intend to pursue a “balanced” automation strategy as they select 5G core vendors, versus the 27% that will be “aggressive” and “need end-to-end core automation from the outset.” The response also reveals, categorically, that one-week cycle times are expected to be sufficient for core network configuration changes, with 70% of the response. One-week cycle times are much faster than is typical using today’s operating processes and this probably represents operator ambition, rather than near-term expectations in real-world deployments.

Enthusiasm for cloud-native technology is high. However, only 14% of respondents already have a continuous integration (CI)/continuous deployment (CD) pipeline in place for 5G core. Caution in applying CI/CD to nationally critical mobile core infrastructure is understandable because failures can impact large numbers of people, services, and devices. Activity in CI/CD will build in the near and medium terms: in the next

two years, a majority expect to implement a CI/CD approach to 5G core operations, of which 32% will act in the next 12 months and 27% in the next 24 months.

A wide range of options for CI/CD pipelines are under consideration, which may indicate that diverse implementations are most likely, or that the market is currently immature. System integrator-developed pipelines (26%) and in-house developed pipelines (25%) are represented roughly equally, with the vendor-supplied model split between the platform vendor (14%) and 5G core vendor (also 14%). A platform vendor approach is more likely to be horizontal, while a 5G core vendor approach is more likely to be vertical and specific to the core network domain. Continuous testing strategies for 5G core are similarly diverse and relatively immature.

The survey confirms the widespread expectation that software-based user planes will prevail in the 5G core: a majority (57%) expect virtual network functions (VNFs) to predominate, ahead of cloud-native network functions (CNFs) (35%) and physical network functions (PNFs) (just 8%). The move away from PNFs is well understood and expected. The view that VNFs will be more numerous than CNFs in the user plane for the next three years requires explanation. One reason could be that the performance of a CNF user plane from vendors is not yet at the required level for a high throughput, high availability 5G core. For smaller user plane deployments, such as at the edge, there is likely to be a stronger push for a CNF user plane earlier in the cycle because of the greater need for automation where large numbers of locations are involved.

Half (49%) of respondents expect 80% of today's core functions that currently use Diameter signaling will migrate to a 5G core solution within three years. This is an aggressive timeline and would represent a rapid transition if it came to pass. This finding is also a useful proxy for the rate at which operators expect to move to a common 4G/5G core (identified as an important objective for most operators).

Operators prefer the service mesh approach to managing signaling traffic in a 5G core. A solid 62% select this option over an independent deployment unit (21%). This is not surprising because a service mesh is more cloud native and is clearly the more fashionable option. It does, perhaps, appear a little aspirational. In practice, it is likely that both models will be used in the same network because there are scenarios that the service mesh approach does not readily address; for example, to proxy into and out of a Kubernetes environment for non-5G core signaling. Moreover, distributed Service Communication Proxies (SCPs), based on service meshes, may be construed as both independent deployment units and as service meshes.

Operators are keen to develop the 5G edge cloud model by working with public cloud providers. Close to a fifth (22%) say they are already bringing public cloud edge infrastructure inside the mobile network perimeter to support third-party developers. A further 26% say their company “plans to do this within 12 months” and 32% “within 24 months.” Overall, sentiment is positive: the public edge cloud model looks like it is becoming an important—possibly very important—new architecture to deliver services over 5G.

SURVEY DEMOGRAPHICS

The Heavy Reading **Cloud-Native 5G Core Survey** was conducted online in December 2020 and January 2021. The questionnaire was written by Heavy Reading with input from project sponsors F5, HPE, Mavenir, and Oracle. It was promoted to the Light Reading service provider database and received 72 responses from individuals working at operators with mobile network businesses, after spurious, incomplete, and non-operator responses were removed.

Respondents were asked to self-assess their knowledge of their company's mobile core network strategy; those with "no direct knowledge" or "only a little knowledge" of the mobile core were excluded. Of the 72 that remained, 20 (28%) say they have "in detail" knowledge of the mobile core, 20 (28%) say they have knowledge at a "strategy level" and 32 (44%) say they have a "good general knowledge of our core network strategy." The survey data presented in this report, therefore, represents a well-informed respondent base.

All responses are confidential and are only ever presented in aggregate form. Heavy Reading does not share individual names or company names from the survey. The 72 responses represent between 40 and 49 different operators, with the spread due to how national operating companies, from the same operator group, are counted.

Figure 1 shows the responses by operator type. Half (51%) of respondents work at converged operators with a mobile network, followed by one-third (32%) who work at pure-play mobile operators.

Figure 1: What type of telecom service provider do you work for?

	Percent	Count
Mobile operator	32%	23
Converged operator with a mobile network	51%	37
Cable operator with a mobile network	10%	7
Virtual operator (MVNO)	4%	3
Cloud provider/Other	2%	2

n=72

Source: Heavy Reading

Figure 2 below shows the response by geography. The US is the largest market represented in the survey with 42% of the response. This allows Heavy Reading to compare the US response versus Rest of World (RoW) with reasonable confidence. Where the analysis in this report contrasts US and RoW, it is made clear in the text.

Figure 2: In what region is your organization headquartered?

	Percent	Count
US	42%	30
Canada	4%	3
Central/South America (including Mexico & the Caribbean)	11%	8
Western Europe	17%	12
Central/Eastern Europe	6%	4
Asia Pacific (including Australia)	19%	14
Middle East	1%	1

n=72

Source: Heavy Reading

Figure 3 shows responses by job title. Network engineering and planning is the largest group with 43%, followed by R&D and technical strategy roles with 24%.

Figure 3: What is your primary job function?

	Percent	Count
Corporate management	6%	4
R&D and technology strategy	24%	17
Network engineering & planning	43%	31
Network operations	13%	9
IT and cloud	7%	5
Marketing/sales	4%	3
Other	4%	3

n=72

Source: Heavy Reading

Figure 4 shows that responses are dominated by operators with more than \$1bn in annual revenue and almost half (46%) with revenue of more than \$5bn.

Figure 4: What is your company's approximate annual revenue?

	Percent	Count
Less than \$250m	11%	8
\$250-\$499m	13%	9
\$500-\$999m	13%	9
\$1-\$4.99bn	18%	13
\$5bn or more	46%	33

n=72

Source: Heavy Reading

DRIVERS AND TIMELINES FOR 5G CORE


The deployment of 5G core networks, and the introduction of standalone 5G services, is now well underway and gaining momentum. **Figure 5** shows that one-third (32%) of respondents say their company has already deployed a 5G core and is offering standalone services in the wide area network, and that a further 27% will be live by the end of 2021. A small majority of respondents to this survey say their company will have 5G core deployed and in live service by the start of 2022. This is clearly a bullish outlook and, at first glance, an over-positive representation of the market.

To put this in context, at the end of 2020, there were close to a dozen operators live with 5G standalone (most on a limited scale) in their public networks, so this survey's results are clearly not representative of operators in general. There is no way one-third of the 900 mobile operators worldwide have already deployed 5G core, or even that one-third of the 120+ live 5G networks already use a 5G core. It is more appropriate, therefore, to view this data from the point of view of the advanced operators represented in the survey. It is also possible—perhaps even likely—that friendly user trials of 5G core are considered by respondents to be “live” deployments.

This interpretation that the survey reflects the advanced makeup of the respondent base is seen in regional data. In the US, for example, the major mobile operators have launched 5G, albeit in limited ways and generally not on a mass-market commercial scale. In this region, 47% of respondents say they have already deployed 5G core versus a smaller 22% for the RoW.

Overall, the survey confirms a view that operators around the world are now at the start of a broad-based 5G core deployment. It is clearly still early in the cycle, and it will take time to reach the global mainstream, but 5G core is underway. The next two years will see a lot of activity.

Figure 5: When does your company expect to deploy 5G core and offer standalone 5G in the public wide area network?




n=71

Source: Heavy Reading

The next question asks why operators need, or want, 5G core networks. This question is useful because it is possible to offer 5G services using a 4G core in non-standalone mode, as is the case for the vast majority of 5G networks and subscribers today. The result in **Figure 6** shows a clear majority (64%) say the primary reason to deploy 5G core is to “offer advanced new services” to customers. This is a logical and expected result. It is perhaps a little surprising that the lead over “to modernize infrastructure” (with 32%) was not greater. It is worth keeping in mind that the supposed new services that 5G core will enable barely exist commercially and it is, therefore, reasonable to be cautious about this aspect of 5G core. Moreover, there is an argument that operators need to modernize core infrastructure before they can offer advanced new services. Overall, then, the result is clear and unambiguous: services are the driver for 5G core.

Figure 6: What is the PRIMARY reason for your company to deploy a 5G core in your public wide area network?



n=71

Source: Heavy Reading

Building on the theme of 5G services, the survey tests how confident respondents are about the impact that standalone 5G will have on the customer experience. **Figure 7** below shows a slim majority expect to see “clear benefits to the customer experience” made up of 52% for consumer and 57% for enterprise. This is a positive result for 5G core, but it is not quite a wholehearted endorsement because, there is clearly some doubt among just under half of respondents. This concern does not appear to be acute, however, as 37% (consumer) and 36% (enterprise) think the customer experience will “probably” improve even if “it is not yet clear how significant this will be.” The overall picture is one of optimism that 5G core will have a direct, positive impact on the service experience, but tempered by caution about what this will mean in practical terms.

Figure 7: Will 5G core enable better, or new, customer experiences for consumers and enterprises in the next three years?


	Yes, there are clear benefits to the customer experience	Probably, but it is not yet clear how significant they will be	Unsure—in practice, it is not clear	Most end users won't notice the difference
Consumers	52%	37%	7%	4%
Enterprises	57%	36%	7%	0%

n=71-72

Source: Heavy Reading

One of the service types that is expected to define 5G is network slicing. This is a service that requires a 5G core and cannot be delivered easily on a 4G core network. **Figure 8** asks how attractive network slicing is as a commercial proposition. The results suggest optimism, but not outright certainty that network slicing has a strong business case and will be a commercial success. The largest group of respondents is the 48% that think network slicing is an “attractive” proposition but are “not certain” of the business case. That only 30% believe networks slicing is “very attractive” and there is “a strong business case” indicates there is some uncertainty on the commercial outlook among the majority of the respondent base.

Figure 8: How attractive is network slicing as a commercial proposition for your company?



n=71


Source: Heavy Reading

5G core can be introduced alongside 4G core as an “overlay” deployment with various points of integration to support interworking and/or it can be deployed as a new, integrated common 4G/5G core. The question in **Figure 9** below seeks to understand how important it is to operators to have a common 4G/5G core versus a 5G overlay that may be integrated

over time. The largest group of respondents are the 46% that say it is “essential” to have a converged 4G/5G core just ahead of the 40% that say it is “important and desirable, but not essential to start with.” The key finding, therefore, is that a large majority (86%) expects their company to operate a common 4G/5G at some stage.

To achieve a common, fully converged core within the aggressive deployment timeframes identified previously is challenging and it is possible that respondents are over-confident or have a permissive definition of what constitutes a common core. It may be more appropriate to think in terms of specific network functions or groups of network functions. For example, in terms of common policy control across 4G and 5G, or in some scenarios, of a common user plane, and in others, a common subscriber data solution. This is a complex topic, with many variations and phases.

Figure 9: How important is a common 4G/5G core to your company?



n=72

Source: Heavy Reading


5G CORE VENDOR SELECTION

A new generation of core network is an opportunity for operators to re-evaluate their vendor relationships and consider if they should introduce new suppliers. The expected uptake of cloud-native 5G core makes it strategically important to select the right long-term vendor partners and introduces technical decisions that did not apply in the 3G to 4G packet core transition. For vendors, a new generation core network is a once a decade opportunity to secure a footprint in operator accounts that will generate long-term business.

The survey asks if operators will lean toward existing or new suppliers in their vendor selection strategies, as shown in **Figure 10** below. The responses are pretty evenly spread across the categories with no dominant single answer. This shows there is an appetite to use new vendors in some form, but that a wide variety of strategies are in play. The largest group is the 30% that expect to use “a mix of new and current vendors,” while a quarter (24%) say they will “introduce new 5G vendors and keep the current for 4G.” However, only

14% expect to use “new vendors for 5G and 4G.” One way to interpret this is that there is a likelihood of changes in vendor selection at the individual operator level, but that on an industry-wide basis, although there will be change, there may not be radical transformation in the vendor landscape.

Figure 10: What is your company’s primary strategy when selecting a 5G core vendor?



n=71

Source: Heavy Reading

A view of the data that compares respondents working for operators with different annual revenue reveals some interesting differences in vendor selection preferences. **Figure 11** shows that 41% of respondents at companies with less than \$5bn will “use existing vendors for 4G and 5G.” By contrast, those working for operators with more than \$5bn in revenue (47%) will use “a mix of new and current vendors for a common 4G/5G core.” This indicates that larger operators are more adventurous in vendor selection.


Figure 11: What is your company’s primary strategy when selecting a 5G core vendor?

	All respondents	>\$5bn	<\$5bn
Base:	71	32	39
Select new vendors for 5G and 4G	14%	22%	8%
Introduce new vendors for 5G and keep the current for 4G	24%	16%	31%
Use existing vendors for 4G and 5G	30%	16%	41%
A mix of new and current vendors for a common 4G/5G core	30%	47%	15%
Don't know	3%	0%	5%
Sigma	100%	100%	100%

Source: Heavy Reading

A variant on the multi-vendor question is when the survey asks specifically about how many vendors an operator will use in its 5G core, as shown in **Figure 12**. There is a clear favorite, with 51% saying they are “likely to use two or three vendors to assemble a 5G core,” which increases to 70% in the US (typically larger operators) versus 38% in the RoW. A further 26% are “likely to use a single vendor” and 17% are “likely to use multiple vendors to create a best-of-breed 5G core.” The overall picture is, again, that there will be diverse vendor strategies in play in the 5G core, but not a radical overhaul of the market on an industry-wide basis.

Figure 12: Does your company plan to assemble the functions that make up the 5G core from multiple vendors or from a single vendor?




n=72

Source: Heavy Reading

A challenge with multi-vendor core networks is integration of the component subsystems. In a cloud-native deployment, where network functions are abstracted from infrastructure, this is arguably an even greater challenge that requires expertise not only in the core network (interfaces, interoperability, dimensioning, configuration, etc.), but also in the underlying cloud environment. **Figure 13** below shows that most operators believe they are reasonably well prepared to manage this internally. A majority (54%) say they are “prepared for a pre-integrated, pre-validated solution that we can evolve over time,” but only 27% think they are “very well prepared.” This shows a willingness to use vendor support when needed, rather than to seek to do everything in-house immediately. The picture is one of pragmatic, purposeful progress.

Figure 13: How prepared are you internally to integrate a multi-vendor best-of-breed 5G core?




n=72

Source: Heavy Reading

Inspired by hyperscale public cloud companies, operators are widely expected to pursue 5G core network automation. The advantages are clear: lower cost of operations, faster to change configurations, reduced risk of human error, and so on. **Figure 14** below shows that this expectation is correct, but also that operators will be purposeful in their adoption of automation and not gung-ho. The clear majority (69%) intend to pursue a “balanced” automation strategy as they select 5G core vendors, versus the 27% that will be “aggressive” and “need end-to-end core automation from the outset.”

One reason, perhaps *the* major reason, for operators to be cautious about core network automation is the consequences of network failures. In a core network, even small errors can result in a widespread network failure. For critical network infrastructure, this is not only a problem for customers and the operator, but potentially for society at large. As more and more devices and services—and more critical services—run on 5G networks, the impact of outages and brownouts becomes ever more severe. This single reason explains why operators will be cautious in how fast and how far they rely on automated operation of the 5G core network. Operators want more automation in the core, but they also want to retain oversight and control.

Figure 14: How essential is automation for your 5G core vendor choices?




n=72

Source: Heavy Reading

Automation is addressed again in the next question, as shown in **Figure 15** below. Here the questionnaire asks how fast operators expect to be able to make changes to the topology or services running on the 5G core. The response reveals, categorically, that “balanced” 1-week or less cycle times are expected to be sufficient for core network configuration changes, with 70% of the response. In practice, a 1-week cycle time would be much faster than is typical using today’s operating processes and this result seems bullish given the need for purposeful and careful change in the 5G core. Therefore, Heavy Reading interprets this finding, in part, as representing 1) operator aspiration; and 2) that respondents focused on the “balanced” part of the question.

Figure 15: How quickly would your company expect to be able to roll out changes to the topology or services going across your 5G core?



n=72

Source: Heavy Reading

5G CORE INFRASTRUCTURE STRATEGY

Cloud infrastructure choices are critical to how a 5G core is deployed and operated. In a change to the classic appliance model based on PNFs, the move to cloud native has major implications for the core network, most notably in terms of redundancy and high availability. Cloud infrastructure also has direct and numerous implications for performance and manageability; for example, in terms of throughput, automation, security, dimensioning, scale in/out, and so on. Operators and vendors have addressed some, but not all, of these changes in the transition to virtual 4G core, using VNFs running in virtual machines, and should be somewhat prepared for the cloud-native transition in 5G core.


Figure 16 below shows expectations for the type of infrastructure that will run the 5G core functions. The choice between IaaS and PaaS is one of the most critical and far-reaching decisions operators must make as they plan for the 5G core. Respondents favor an IaaS model (45%) in which the operator provides the compute, network, and storage resources in their own data center and the vendor provides the infrastructure and application software. This model is common in network functions virtualization (NFV) and, to Heavy Reading's understanding, is used in most of the first 5G core networks to go live to date. The advantage is that the operator and vendor can optimize the performance of the 5G core in a timely manner, but at the expense of it being vertically quasi-siloed for the initial deployment.

In second place are those (32%) that favor PaaS. In this model, the operator also provides the cloud infrastructure software (OS, runtime, orchestration, etc.) on which it can deploy a vendor's 5G core applications. This places greater responsibility on the operator and requires more advanced internal skills and capabilities. One advantage of this horizontal model is that operators can create a common private telco cloud platform for all network functions, with fewer silos and less vendor lock-in.

Respondents working for larger operators with revenue of more than \$5bn annually are somewhat more likely to select PaaS with a score of 44% versus 41% for IaaS, and 16% for the vendor full stack. Respondents working for operators with revenue of less than \$5bn, score 23% for PaaS, 49% for IaaS, and 28% for the vendor full stack.

The overall picture, according to this survey result, is that both PaaS and IaaS models are likely to be used over the near and medium terms. This accords with Heavy Reading's view prior to the survey, given that both have advantages and challenges. The PaaS model and the CaaS variant are generally expected to prevail over the longer term and might have been expected to score more highly. Perhaps this result suggests that cloud infrastructure vision and execution are not yet quite aligned.

Figure 16: What is your infrastructure/platform strategy for your 5G core?



n=71

Source: Heavy Reading

Three high-level requirements for a cloud-native 5G core are: 1) cloud infrastructure; 2) 5G core applications (a.k.a. CNFs); and 3) internal skills. **Figure 17** below asks respondents when they expect each of these will be ready for large-scale operation in their company's network. Roughly a quarter (28%, 21%, and 23%) think they are ready today and about the same number (24%, 32%, 29%) believe they will be ready in two years (i.e., by the start of 2023). The largest group (39%, 37% and 37%) think each of the domains will be ready for large-scale operation of the 5G core in one year (i.e., early 2022).

These timelines align pretty well with the schedules for 5G core deployments identified earlier (which is that 57% say their companies will have 5G core deployed and live service in 2022). The overall picture is that a majority of operators think they are close to being ready to operate cloud-native 5G core at scale, but that they are not quite there yet. According to the survey, 2021 is set to be a critical year to prepare for a mass-market 5G core from 2022 onward. It is worth noting that US respondents (typically larger operators) are significantly more bullish than their RoW peers on their "readiness" estimates.

Figure 17: When do you anticipate the following aspects of the cloud-native 5G core technology ecosystem to be ready for large-scale operations in your network?

	Ready for large-scale operation	Ready in one year	Ready in two years	More than two years from ready
Cloud infrastructure	28%	39%	24%	9%
5G core CNFs (applications)	21%	37%	32%	10%
Internal skills and operations	23%	37%	29%	11%

n=70-71


Source: Heavy Reading

The 3GPP-defined 5G core is made up of over a dozen discrete network functions, most of which are in the control plane. In the user plane, the user plane function (UPF) is the key 3GPP element. In most networks, UPFs sit alongside various non-3GPP user plane functions, such as firewalls, in the SGi/N6-LAN. In the 4G core, physical appliances carry much of the user traffic globally; however, in new or refreshed deployments, VNFs are now the more common solution.

The question in **Figure 18** below asks if physical, virtual, or cloud-native network functions will be in the majority in the 5G core user plane over the next three years. The results confirm the widespread expectations that software-based user planes will prevail: a majority (57%) expect VNFs to predominate, ahead of CNFs (35%) and PNFs (just 8%).

The move away from PNFs is well understood and expected. These products are likely to be used in some high throughput fixed wireless 5G core applications but are less likely to be used for mobility. The expectation that VNFs will be more numerous than CNF user planes for the next three years requires more explanation because, long-term, the direction of travel is assumed to be toward cloud native. Perhaps one reason relates to the installed base of VNF user plane functions being upgraded to dual-mode 4G/5G; another could be that the performance of CNF user planes from vendors is not yet at the required level for high throughput, high availability 5G core; yet another could be that operators are not yet fully deployed with cloud-native infrastructure. For smaller user plane deployments at the edge, there is likely to be a stronger push for CNFs earlier in the cycle because of the automation benefits and the requirement to rapidly “spin up” UPFs.

Figure 18: In terms of 5G core user plane, which network functions do you expect to be the majority in your company's network in the next 3 years?




n=72

Source: Heavy Reading

The 3GPP Service Based Architecture (SBA) used in the 5G core is based on the principle that interconnected network functions can access common data repositories and are authorized to access each other's services. Each function exposes its services through a Service Based Interface (SBI), which is a REST interface over HTTP/2 transport. This replaces the point-to-point architecture used in the 4G core. Given the preference (already established earlier in the survey) for a common 4G/5G core, it is reasonable to expect that, over time, 4G mobile core functions will transition to the SBA. This is, however, a complicated picture because existing 4G core functions and subscriber databases currently use Diameter. The new 5G core must either replace or integrate with these existing functions to enable services, such as handover between 4G and 5G, and to support common service offerings across the 4G and 5G footprint.

Figure 19 below shows expectations for the timeline to transition Diameter functions to 5G core. Specifically, it asks when 80% of the core functions that currently use Diameter (i.e., 4G core functions) will transition to 5G; the question is, therefore, a proxy for how fast respondents expect operators to move toward a common 4G/5G core. The results show that almost half (49%) expect that this will happen within three years. This would be, by most measures, a rapid transition, but is not unrealistic if the 3G to 4G core transition is used as guide. It also means half the survey base thinks it will take longer than three years, although relatively few (just 11%) think the process will take more than five years. Note also that the way the question is worded is a little ambiguous; this finding does not necessarily mean 80% of Diameter functions will be removed from the core.

Figure 19: What's your target timeline to transition 80% of core functions currently using Diameter to 5G core?



n=72


Source: Heavy Reading

The SBA allows for many-to-many connections between network functions. Signaling between these functions can be implemented within a service mesh or using a discrete independent deployment unit.

Figure 20 below makes it clear that, with a score of 62%, operators prefer the service mesh approach to managing 5G core signaling traffic. This is more cloud native and is clearly the more fashionable option; therefore, this result is not surprising. The 21% using a discrete, centralized “independent deployment unit” should not be discounted because this is an attractive option in some networks. It is also likely that both models will be used in the same network where there are scenarios that the service mesh approach does not readily address; for example, to proxy into and out of a Kubernetes environment for non-5G core signaling.

A further complication is that compressive 3GPP signaling cannot be implemented in an unmodified service mesh (e.g., the open source Istio) because 3GPP specifies custom headers for services such as flow control. It is also the case that an SCP developed as an independent deployment unit could use a modified Istio service mesh “under the hood.” Therefore, when respondents say they will use a service mesh, it could also apply to distributed SCP deployments.

Figure 20: What is your deployment strategy for 5G core signaling control functions in the next 3 years?



n=71


Source: Heavy Reading

5G CORE AND THE PUBLIC CLOUD

Many of the applications that consumers and enterprises use daily are hosted in the public cloud. It is logical to optimize how 5G users access these clouds to improve the experience of existing services and to enable new services that are not practical, or possible, with today's centralized cloud architectures. One option is for the operator to partner with the public cloud provider to create a "5G edge cloud" platform and commercial offering. Today, this typically involves deploying a public cloud hardware and software stack inside the operator network behind the SGi/N6 core network interface. Developers can deploy services on this edge cloud using tools with which they are already familiar. It is anticipated that, over time, developers will re-architect applications so that the appropriate parts of the service run on the edge (e.g., low latency media delivery), while other parts of the application continue to run in the larger centralized data center (e.g., subscriber profiles).

Figure 21 below shows operator views on bringing public cloud edge infrastructure inside the mobile network perimeter. About a fifth (22%) say they are already doing this. At first glance, this looks like a high number, given there are only a handful of public references for this edge cloud model, of which only a few are in live service. However, many operators are known to be involved in these partnership discussions; on this basis and given the advanced nature of the operators in the survey base, this number looks more plausible. That a further 26% say their company "plans to do this within 12 months" and 32% "within 24 months" shows that sentiment toward this public edge cloud model is positive. This looks like it is becoming an important, and possibly a very important, new architecture that will determine how certain service types are delivered over 5G in future.

Figure 21: Will you bring public cloud edge infrastructure inside your mobile network perimeter (e.g., AWS Wavelength, Azure Stack, etc.) to support third-party services and developers?




n=72

Source: Heavy Reading

Another area where the survey seeks insight is the use of the public cloud to run 5G core network functions. **Figure 22** below shows that operators are primarily interested in this model for specific services and network types. The largest group (38%) says their companies are likely to run core functions in the public cloud “for specific services, such as private networks, roaming, automotive, IoT, etc.” within three years. This shows there is momentum and enthusiasm behind the idea. The next largest group (29%) says this is “under evaluation – we don’t have a good view at this stage.”

The 22% of respondents that say they are likely to run 5G core applications “for the main wide area public network” on this type of infrastructure within three years looks high. As a rule, operators require full control over their main core network, and Heavy Reading’s prior research have detected little enthusiasm from major telcos to host the main mobile core network in the public cloud. However, some smaller operators and MVNOs are open to this model and some larger operators are known to be evaluating moving functions, such as policy control, to the public cloud, particularly in cases where they also plan to move business support systems (BSS) from private infrastructure to the public cloud. Another possible explanation is that public cloud providers are proposing to deploy a private cloud (with the same technology stack and economics as a public cloud) inside carrier networks to support network functions use cases; it is plausible that this type of deployment accounts for part of the interest expressed in this result.

Figure 22: Is your company likely to run 5G core applications on public cloud infrastructure within the next 3 years?



n=72

Source: Heavy Reading


CI/CD AND 5G CORE OPERATIONS

One major advantage of moving from PNFs to VNFs and CNFs running on cloud infrastructure is the ability to quickly change the configuration of services and infrastructure. The survey has already established that automation is important to vendor selection and that operators want the ability to rapidly make changes to the 5G core and associated services via software.

CI/CD is a set of operating principles that enable development teams to deliver code changes more frequently and reliably. A CI/CD pipeline is the implementation of this model and is a defining feature of cloud-native deployments. **Figure 23** below shows just 14% of respondents already have a CI/CD pipeline in place for 5G core, but that a majority expect to implement this within the next two years, of which 32% in the next 12 months and 27% in the next 24 months. This maps fairly closely to the 5G core deployment timelines identified earlier, albeit with some lag time.

Caution before applying CI/CD to the mobile core network is understandable and sensible. This is nationally critical infrastructure and failures can impact large numbers of people and devices—if a mobile connection is lost, all services are also lost—and this may explain why 32% say they are live with 5G core (see **Figure 5** earlier in this report) but only 14% say they are live with CI/CD today.

Figure 23: Do you expect to implement a CI/CD pipeline for your cloud-native 5G core?




n=71

Source: Heavy Reading

With the principle of a CI/CD pipeline is established, the survey then asks how operators will implement this model in their 5G core networks. **Figure 24** below shows that a wide range of options are in play and that the landscape of solutions is diverse. Integrator-developed (26%) and in-house (25%) pipeline implementations are represented roughly equally, with the vendor-supplied model split between the platform vendor (14%) and 5G core vendor (also 14%). A platform vendor approach is more likely to be horizontal, while a 5G core vendor approach is more likely to be vertical and specific to the core network domain. This result possibly under-represents the vendor input into the CI/CD process because, in many cases, Heavy Reading is aware that the core network vendor plays a major role.

Figure 24: How do you plan to enable a CI/CD pipeline for your 5G core network?




n=72

Source: Heavy Reading

To be effective a CI/CD pipeline needs a method to test software prior to deployment to minimize errors in production. One approach is “continuous testing,” which involves “shifting left” to create test plans that can identify defects as early as possible in the software delivery process. Applied in a 5G core context, the expectation is that correcting errors before new configurations are deployed improves the overall reliability of core network services. Test plans and methods can be designed by the 5G core vendor in association with the operator or developed independently of the vendor.

Figure 25 below shows that operators are pursuing a range of strategies for 5G core network testing. A quarter of respondents (25%) have already “shifted testing to the left” and have implemented a continuous testing strategy. This represents the advance guard that is closest to a cloud-native operating model. The largest group is the 33% that say they have developed a continuous testing plan “in-house” but have not yet shifted left; this represents operators that are progressive and moving toward a CI/CD pipeline model, but are also cautious about adopting it too soon. A further 26% expect the vendor to bring their own test plan to integrate with the operator CI/CD pipeline; in practice, 5G technology suppliers are very likely to be involved in the 5G core test plan, so this number may again under-represent the true involvement of vendors.

Figure 25: Which of the following best describes your current 5G core testing strategy?



n=72

Source: Heavy Reading

BACKGROUND TO THIS STUDY

Heavy Reading's 2020 **Cloud-Native 5G Core Operator Survey** was conducted in December 2020 and January 2021 and this analysis was written in January and February 2021. The online survey generated 72 responses from individuals working at communications service providers after non-qualified responses were deleted from the survey. Respondents were asked to self-assess their knowledge about 5G core. Those that reported "no direct knowledge" or only "a little knowledge" of their company's 5G core strategy were excluded from the survey and their responses are not considered in this analysis.