

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
AMD, F5, Oracle, and Red Hat*

AMD 



ORACLE
Communications



Red Hat

AUTHOR: GABRIEL BROWN, SENIOR 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 autumn 2022. 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; however, it is generally taken to mean network functions (NFs) deployed in containers and centrally orchestrated. Published in January 2023, this is the second Heavy Reading survey on 5G core, with the first report published in March 2021.

At the start of the survey, to help guide them, respondents were presented with the following definition:

"This questionnaire asks for views on cloud native 5G core and the migration to standalone 5G architecture. The focus is on 3GPP 5G core elements, associated cloud infrastructure platforms, and end-user services enabled by 5G core."

The 5G core is a vital part of the 5G system and essential to enable standalone (SA) 5G operation (i.e., 5G without a dependency on LTE). Without a 5G core, 5G radios can only be deployed in non-standalone (NSA) mode attached to a 4G core. NSA offers benefits in terms of capacity and end-user data rates, but otherwise largely restricts the evolution of 5G to the 4G system architecture and service model.

The 5G core introduces the service-based architecture specified by 3GPP. Although the 3GPP does not define an implementation model, the specifications are written such that the 5G core will likely be deployed on software-defined infrastructure. In practice, this means virtual machine (VM) or cloud native infrastructure.

Key findings

A majority (64%) of respondents expect 5G core to produce "clear benefits to the customer experience," an increase of 10% over the 2021 survey. A third of respondents (33%) believe 5G core will "probably" enable new or better services, "but it is not yet clear how significant they will be." This result is, overall, a solid endorsement of the services-driven thesis for 5G core investment, with some equivocation in the survey base about how much impact it will have on the customer. One interesting skew in the data is that of the 22 respondents in R&D and technology strategy roles, a smaller 41% identified "clear benefits" to the customer experience compared to 76% of the 17 respondents in corporate management or sales roles.

Respondents are cautiously optimistic about network slicing for enterprise services. Of all respondents, 49% say they "see value in slicing and are incorporating it into our enterprise wide-area offerings." However, network slicing is an iconic 5G service capability, and it merits attention that half the respondents are not highly enthusiastic. Close to a third of respondents (31%) believe slicing will be part of their enterprise offer but will play "a small role" relative to private networks with on-premises equipment. Given that the private network model is likely to be offered in combination with network slices over time, this 31% is also a reasonably positive signal for network slicing for enterprise services, if not unequivocally enthusiastic.

A combined 66% of respondents say their company will introduce a new core for 5G versus 35% that “will evolve the current 4G core to 5G users/traffic.” There is common agreement that a combined 4G/5G core running on a cloud platform is the target architecture; however, there is some variation in how long operators expect this migration to take, with response split almost evenly between “maintain the current 4G core for the medium term and launch a new 5G core” (33%) and “launch a new 5G core and then migrate 4G users to the new platform in the short-term” (28%).

5G core will not generate a major change in how operators select and manage vendors. Most operators expect to directly manage the relationship with their suppliers, and very few (only 6%) expect to work with a single-vendor partner. This is not an invaluable finding, given that proposals to disrupt the mobile core market are regularly floated (for example, 5G core-as-a-cloud service), so it is useful to check sentiment toward new sourcing models. **There is, however, an appetite to use smaller, innovative vendors.** Of all respondents, 42% “have several scenarios where smaller vendors will be useful,” which is a strong signal to these companies that there is a market for their product.

Telco cloud strategies remain under development, although there is a clear and strong preference for private cloud over public cloud. Of all respondents, 37% will deploy 5G core using an infrastructure as a service (IaaS) model versus 23% that favor a platform as a service (PaaS) model for the telco cloud, and 20% will use vendor full stack. There is little appetite to run 5G core fully in the public cloud (just 4%), but some appetite to run a hybrid private and public cloud infrastructure (16%). A horizontal PaaS is attractive to operators because it enables a cloud operating model with common continuous integration/continuous development (CI/CD) tools; however, it is a challenge for telecom operators because of, for example, the maturity of telco cloud stacks, their operational capabilities (e.g., skills), and their existing resiliency and monitoring models. For the mobile core network, the challenge is amplified because high availability is so crucial. In the mobile core, operators understandably have a very low appetite for risk.

About a quarter of respondents (23% and 23%) say their organization already uses CI/CD/continuous test (CT) successfully in live deployment, and 20% are also using “shift left” methods for ongoing service management. The 34 respondents that work at operators with more than \$5bn in annual revenue were more likely to be in commercial service with CI/CD/CT than smaller operators. Telecom operators have been relatively slow to adopt modern cloud operating processes. P An advanced cohort of this size may perhaps indicate that a threshold has been crossed and that a larger number of operators may follow with CI/CD in the near future.

Environmentally sustainable networks and energy efficiency are of great importance. **In the 5G core, the leading initiative to reduce energy consumption is to “move as many functions as possible to a common infrastructure platform” (52%).** From a core network perspective, this is probably the most consequential move operators can make. However, there are many additional moves to improve the power efficiency of the cloud platform itself, including improving CPU efficiencies (45%), investigating hardware offload/accelerators (39%), using cleaner/renewable energy (37%), and minimizing cooling (34%).

There is enthusiasm about the benefits of network analytics and network capability exposure via APIs. Operators will initially use the Network Exposure Function (NEF) for internal applications or to work with carefully curated partners (65%). However, the scores for “creating a developer ecosystem” (58%) and “going beyond basic NEF functionality” (48%) are also strong and indicate operators are optimistic that 5G network APIs could, in the future, enable them to become open technology platforms on which other parties can innovate.

About half of the respondents (51%) expect their engineering team to drive the Network Data Analytics Function (NWDAF) procurement process, and close to a third (31%) expect the customer care team to lead the process. Lower scores for business operations (13%) and data science/analytic groups (5%) to drive the NWDAF procurement do not mean these groups are inconsequential. Although the engineering teams will run the main process, there will be significant input from related teams within the operator, and it is virtually certain that their requirements will be incorporated into the process.

The survey indicates that most service providers expect to start offering 5G SA roaming “within 12 to 24 months” (41%). More than a third also indicated that they expect this offer “within 12 months” (37%). Taking into account the complexity of roaming, the timing of the 5G SA rollout, and the need to trial before a launch, a 12- to 24-month timeframe seems a more realistic goal.

SURVEY DEMOGRAPHICS

The Heavy Reading **Cloud Native 5G Core Survey** was conducted online in September 2022. The questionnaire was written by Heavy Reading with input from project sponsors AMD, F5, Oracle, and Red Hat. It was promoted to the Light Reading service provider database and received 83 responses from individuals working at operators with mobile network businesses.

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 83 that remained, 39 (47%) say they have “in detail” knowledge of the mobile core, and 44 (53%) say they have knowledge at a “strategy level.” 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 83 responses represent 51 different operators.

Figure 1 shows the responses by operator type. Of all respondents, 42% work at pure-play mobile operators and 35% at converged operators with a mobile network, for a combined 77% of the total. The remainder is from an assortment of operator types.

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

	Percent	Count
Mobile operator	42%	55
Converged operator with a mobile network	35%	29
Cable operator with a mobile network	12%	10
Virtual operator (MVNO)	5%	4
Cloud provider/Other	6%	5

n=83

Source: Heavy Reading

Figure 2 shows the response by geography. The US is the largest market represented, with 51% of the response. This allows Heavy Reading to compare the US to the Rest of World (RoW) with reasonable confidence. Where demographic filters are used in the analysis, it is made clear in the text of the report.

Figure 2: In what region is your organization headquartered?

	Percent	Count
US	51%	42
Canada	10%	8
Central/South America (including Mexico and the Caribbean)	12%	10
Western Europe	6%	5
Central/Eastern Europe	2%	2
Asia Pacific (including Australia)	12%	10
Middle East & Africa	7%	6

n=83

Source: Heavy Reading

Figure 3 shows responses by job title. Network engineering and planning is the largest group, with 31%, followed by R&D and technical strategy roles, with 27%.

Figure 3: What is your primary job function?

	Percent	Count
Corporate management	15%	12
R&D and technology strategy	27%	22
Network engineering and planning	31%	26
Network operations	7%	6
IT and cloud	15%	12
Marketing/sales	5%	5

n=83

Source: Heavy Reading

Figure 4 shows that responses are led by operators with more than \$5bn in annual revenue (41%), followed by 19% with between \$1bn and \$4.99bn. These responses represent national-scale operators, predominately in developed markets. However, there is also a good number of smaller-scale operators in the respondent base.

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

	Percent	Count
Less than \$250m	12%	10
\$250–\$499m	17%	14
\$500–\$999m	11%	9
\$1–\$4.99bn	19%	16
\$5bn or more	41%	34

n=83

Source: Heavy Reading

Figure 5 shows the spread of respondents according to the number of mobile subscribers. This is a useful reference because it helps understand how 5G core strategies may vary among very large national operators (e.g., in the US or India), multinational group operators (e.g., in Europe), and smaller operators (e.g., MVNOs run by fixed broadband or enterprise service providers, or rural operators serving smaller communities).

Figure 5: How many mobile subscribers does your organization have in its national market (including MVNO customers)?

	Percent	Count
Less than 1 million	10%	8
1–5 million	23%	19
6–9 million	6%	5
10–25 million	17%	14
26–50 million	13%	11
More than 50 million	30%	25
Don't know	1%	1

n=83

Source: Heavy Reading

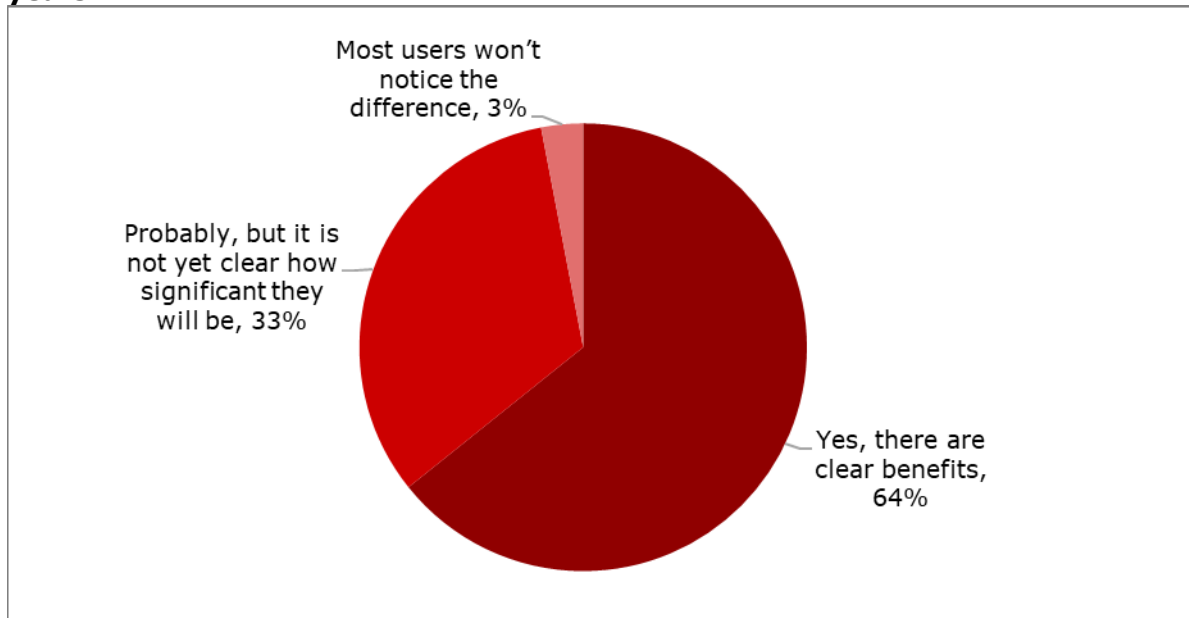
5G CORE SERVICE STRATEGIES

In this section, operators are asked about their views on 5G services and high level deployment strategies. To support new services is, in principle, a major reason to deploy 5G core. The questionnaire tests how optimistic operators are on this front.

The question in **Figure 6** asks if operator respondents think 5G core will enable new or better customer experiences in the next three years. A majority (64%) say “yes, there are clear benefits,” which is slightly higher than the equivalent question in the prior survey (about 18 months previously). A third of respondents (33%) believe 5G core will “probably” enable new or better services, “but it is not yet clear how significant they will be.” This is, overall, a solid endorsement of the services-driven thesis for 5G core investment, with some equivocation in the survey base about how significant that will be to the customer.

One interesting skew in the data is that of the 22 respondents in R&D and technology strategy roles, a smaller 41% identified “clear benefits” to the customer experience, while 45% selected the “probably” option. This compares to 76% of the 17 respondents in corporate management or sales roles. This 35% delta is worth noting, but it is unclear how significant it is. Are R&D people close enough to customers to know what impacts their service experience? Has management been taken in by the incessant hype around new 5G services?

Figure 6: Will 5G core enable new or better customer experiences in the next three years?



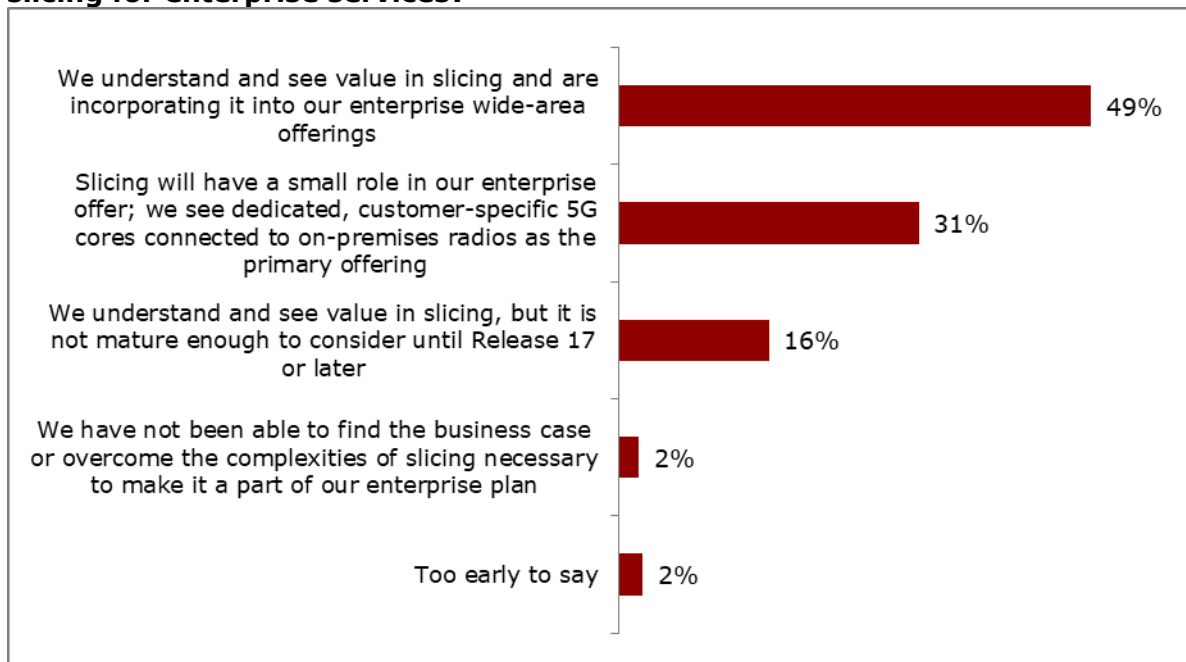
n=83

Source: Heavy Reading

One new service that is only possible with a 5G core and an SA network is network slicing. **Figure 7** asks how respondents view network slicing for enterprise services. The result is cautiously optimistic, with the largest group (49%) saying they “understand and see value in slicing and are incorporating it into our enterprise wide-area offerings.” That these respondents say they “understand” network slicing is of particular importance because, arguably, a lot of the hype around slicing was generated and consumed by people without a good understanding of the technology.

Network slicing is, however, an iconic 5G service capability; in that regard, the finding that half the respondents are not highly enthusiastic merits attention. Close to a third of respondents (31%) believe slicing will be part of their enterprise offer but will play “a small role” relative to private networks with on-premises radios and dedicated enterprise core networks. This is not a direct alternative to network slicing because, over time, the private network model with on-premises equipment is likely to be offered in combination with network slices. Therefore, in that sense, this 31% is also a reasonably positive signal for network slicing for enterprise services, even if it is not unequivocally supportive.

Figure 7: Which statement best reflects your organization’s view on network slicing for enterprise services?



n=83

Source: Heavy Reading

5G CORE DEPLOYMENT STRATEGIES

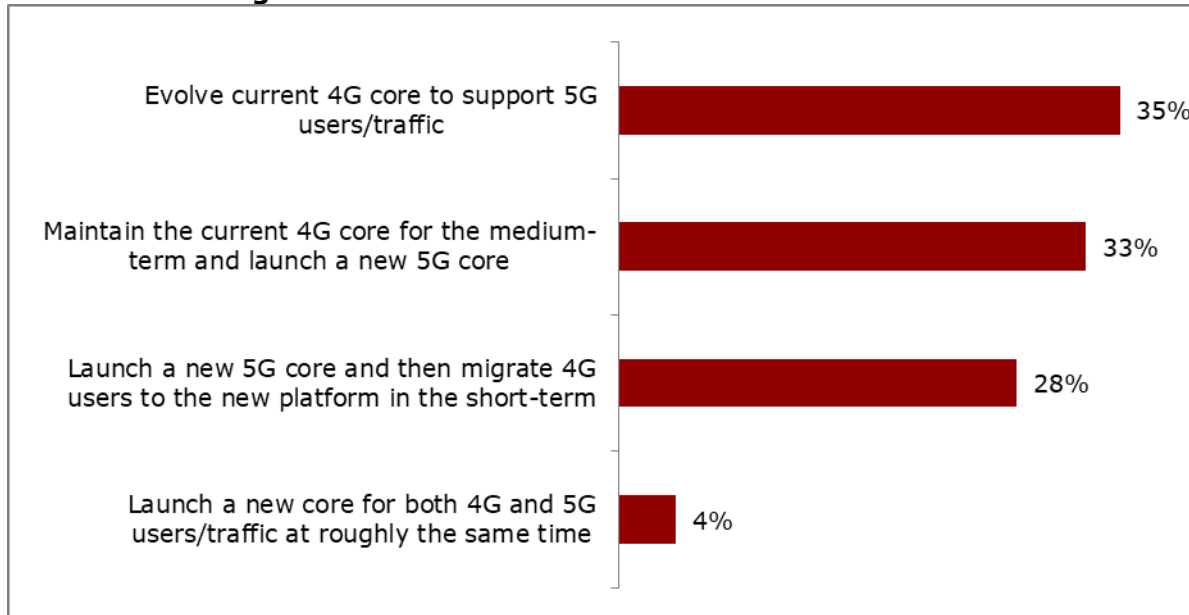
Operators have several options for how they deploy a 5G core. The previous survey (2021) established that a large majority of operators wanted, and intended, to use a common core for 4G and 5G. The 2022 survey (**Figure 8**) asked about phasing in a 5G core introduction and the journey to a common 4G/5G core. Overall, a combined 66% of respondents say their company will introduce a new core for 5G, versus the 35% that say they will “evolve current 4G core to 5G users/traffic.”

Of the 66% “new 5G core” vote, only 5% expect to launch a new common core for both 4G and 5G at the same time, which reflects the risk to service associated with such a large change. The rest are split almost evenly between “maintain the current 4G core for the medium term and launch a new 5G core” (33%) and “launch a new 5G core and then migrate 4G users to the new platform in the short-term” (28%). Short and medium term are not defined in the question; therefore, this should only be seen as an approximation of the time period over which operators will move to a common 4G/5G core. The key finding is there will be a new 5G core platform and then a phased introduction of common core functionality.

The 35% that will “evolve the current 4G core” is higher than Heavy Reading had anticipated because although there are operators known to be pursuing this strategy, received wisdom is that most are deploying a new core for 5G. This result may reflect how parts of the core can be updated for 5G, even when new 5G core functions are introduced. An operator could, for example, maintain its 4G subscriber data systems for 5G core, or it could reuse 4G components such as policy servers and MMEs for 5G core for a period.

Taking a broader view of “core,” it is widely anticipated that most operators will reuse their existing 4G IMS core to support voice and messaging in the 5G network, which could also conceivably play into the 35% response.

Figure 8: Which of the following best represents your organization’s strategy for core network migration from 4G to 5G?



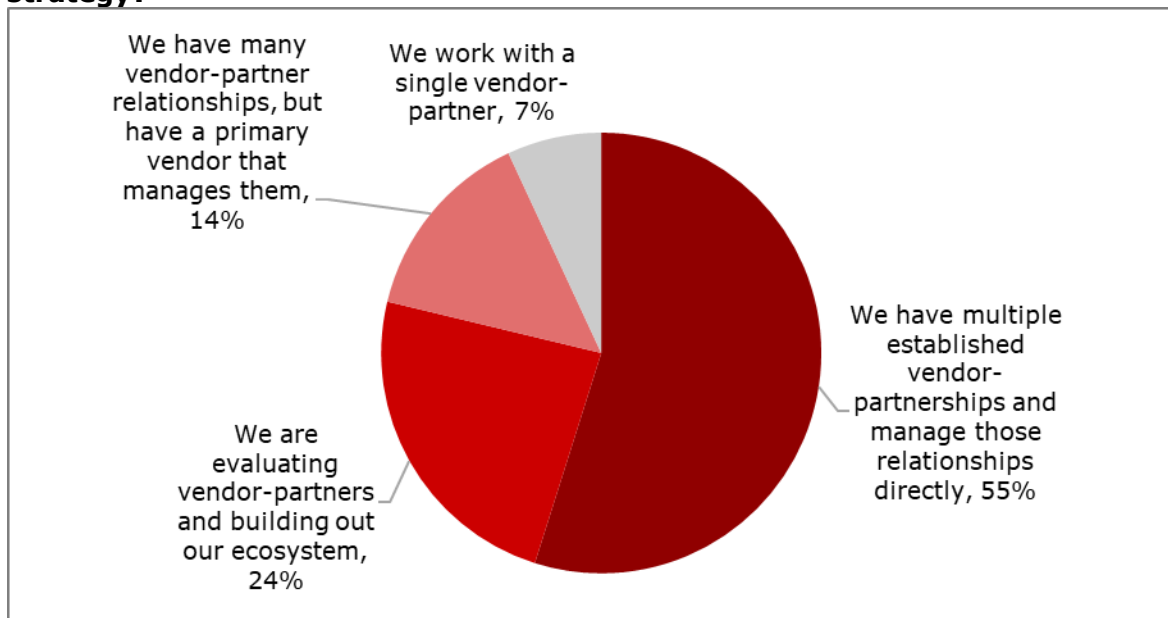
n=83

Source: Heavy Reading

Figure 9 sheds some light on how operators are working with vendors to deploy and operate their 5G core networks. The primary result is that most respondents (55%) “have multiple established vendor-partnerships and manage those relationships directly.” Among the operators with more than \$5bn in annual revenue, this jumps to 74%. In other words, most operators, especially the larger ones, expect to continue business as usual. Based on this result, it does not look like 5G core will generate a major change in how operators select and manage vendors. This is not an invaluable finding, however. Given that proposals to disrupt the mobile core market are regularly floated—5G core as a service in the cloud, for example—it is useful to check sentiment toward new sourcing models.

Not all operators have made their vendor decisions. Of all respondents, 24% are in the process of “evaluating vendor partners,” which is not an insignificant number. This rises to 33% when the 16 responses from the Middle East, African, South American, and Central European regions are isolated. Given the relatively small number of respondents in these regions, we are cautious about overinterpreting this data; however, it does perhaps point to different timelines in different parts of the world.

Figure 9: How are you working with vendor partners in your 5G core network strategy?



n=83

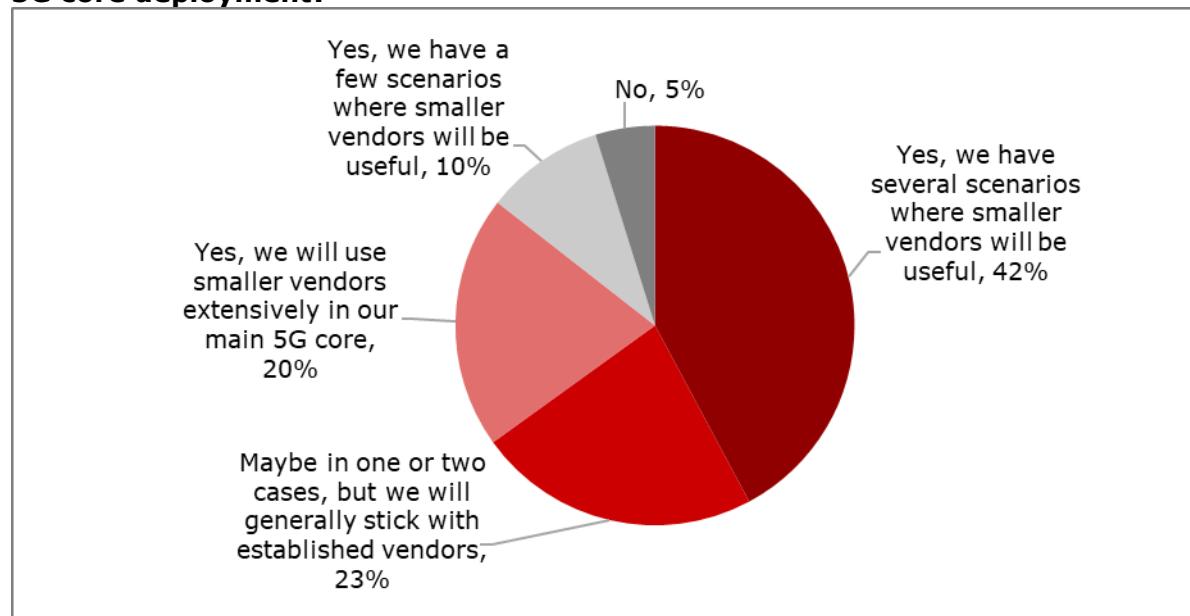
Source: Heavy Reading

Continuing the vendor selection theme, the questionnaire asked about the likelihood operators would use smaller, innovative vendors in the 5G core. Given 5G core is software-centric, one might expect to see more vendors being used. However, the core is so critical that operators are reluctant to take risks. **Figure 10** shows that 20% will “use smaller vendors extensively in their main 5G core,” which means 80% will not. Therefore, the overall analysis is that there is limited appetite to move away from established suppliers for a major part of the core.

Nevertheless, 20% is still a decent number, given how concentrated the mobile packet core market is today. And certainly, there is an appetite to use smaller, innovative vendors. The largest group of respondents (42%) say they “have several scenarios where smaller vendors will be useful,” which is a strong signal to these companies that there is a market for their product. Although the survey does not identify which use cases these might be, IoT, automotive, enterprise, private networks, and edge, among others, might be included.

A combined 33% are less enthusiastic but do have opportunities for innovative suppliers. The 10% that say “yes, we have a few scenarios where smaller vendors will be useful” are open to working with smaller companies for specific reasons on a limited basis. Almost a quarter (23%) say “maybe in one or two cases,” indicating they would consider it but are not really committed to this path. The 5% “no” speaks for itself.

Figure 10: Does your organization expect to use smaller, innovative vendors in its 5G core deployment?



n=83

Source: Heavy Reading

One way some operators appear to be deploying 5G core is to start with the basic functions (e.g., AMF, SMF, and UPF) and reuse subscriber data and perhaps policy functions from the 4G core. This allows for a faster, simpler introduction of 5G SA, with functions such as signalling NFs, network APIs, and 5G-native data management being introduced later.

In **Figure 11**, the survey attempted to investigate the potential phasing of 5G core deployment over different time periods. This is a detailed question that requires close knowledge of the 5G core strategy to answer accurately; therefore, there is potential for "noise" in the response. Nevertheless, the results are interesting.

It is no surprise that subscriber data NFs, which include, among others, UDM, UDR, and AUSF, are rated highly, with almost half of the respondents (49%) expecting their introduction within the Phase 1 period "at launch or close to launch." In the Phase 2 option, 40% believe they will be introduced "within 2 years."

Policy and charging NF introduction seem to be fairly evenly split between the Phase 1 and Phase 2 timeframes, with 39% and 42%, respectively. Signalling NFs were more likely to be implemented within the Phase 2 period (45%), although 36% expect them to be introduced at Phase 1. Within the signalling group, there are several NFs that have matured greatly in functionality within 3GPP Releases 16 and 17; therefore, it is logical that the deployment time scales move to Phase 2 (within two years).

Network exposure NFs receive more votes for Phase 2 introduction (41%) and fewer for Phase 1 (25%). Although it is expected that this area of NFs is more likely to be a second phase for most service providers, there has been much interest and debate around the advantages of the NEF for 5G monetization, which could explain why some operators would choose to introduce this more quickly.

Figure 11: After your initial deployment of 5G core (AMF/SMF/UPF), when will your organization introduce the following NFs?

Responses	Phase 1 – at launch or close to launch	Phase 2 – within 2 years from 5G core launch	Phase 3 – more than 2 years from 5G core launch
Network exposure NFs (e.g., NEF, third-party AFs)	25%	41%	25%
Policy and charging (e.g., PCF, CHF)	39%	42%	14%
Signalling NFs (e.g., NSSF, NRF, SCP)	36%	45%	13%
Subscriber data NFs (e.g., UDM, UDR, AUSF)	49%	40%	6%

n=83

Source: Heavy Reading

5G CORE INFRASTRUCTURE STRATEGY

It is now a common expectation that 5G core should be a cloud native deployment. However, the transition to full cloud native infrastructure is a challenge that requires deep technical expertise. This section gathers information on operators' cloud infrastructure strategies, their perceptions on the readiness of the cloud native ecosystem for 5G core workloads, the role of open source infrastructure software, and the maturity of CI/CD/CT tooling.


The decision on the type of cloud infrastructure on which to deploy 5G core is critical. **Figure 12** reveals that several models are in play. Many respondents (37%) favor IaaS. In this model, compute, network, and storage are provided by, and housed in, an operator data center, while vendors provide the infrastructure software and 5G core functions (i.e., workloads). In the previous Heavy Reading survey, IaaS also led, although there has been a small decline of 8% since then. IaaS is attractive because the 5G core vendor typically takes on most of the integration risk and will be able to offer support and maintenance terms close to those offered for vendor full-stack deployments.

IaaS is generally seen as a stepping-stone to a PaaS model in which the operator runs a telco cloud that acts as a horizontal platform for multiple workloads. PaaS increases vendor flexibility and, perhaps of most importance, enables a cloud operating model with common tools and CI/CD pipelines. This is a challenge for telecom operators in general because of, for example, the maturity of telco cloud stacks, their operational capabilities (e.g., skills), and their existing resiliency and monitoring models. For the mobile core network, the challenge is amplified because high availability is so crucial. Operators understandably are very cautious about the risk of service disruption in the core domain.

There was hardly any interest (just 4%) in the respondent base for deploying the 5G mobile core fully in the public cloud. This is not surprising, as there are many aspects of public clouds that need to mature (technical and business) before they reliably run 5G core workloads for national-scale operators and before operators are comfortable with this model. There is some appetite (16%) for a hybrid model where the 5G core is partially

hosted on-premises and partially in the public cloud, but again, this is limited. Heavy Reading is aware that there is some interest from operators to use a public cloud technology stack deployed as a private cloud on-premises; however, the survey did not test that model (and this would fall into the PaaS category in any case). The conclusion is that operators want to directly control the cloud infrastructure that runs their core network.

Figure 12: What is your infrastructure/platform deployment strategy for your 5G core?



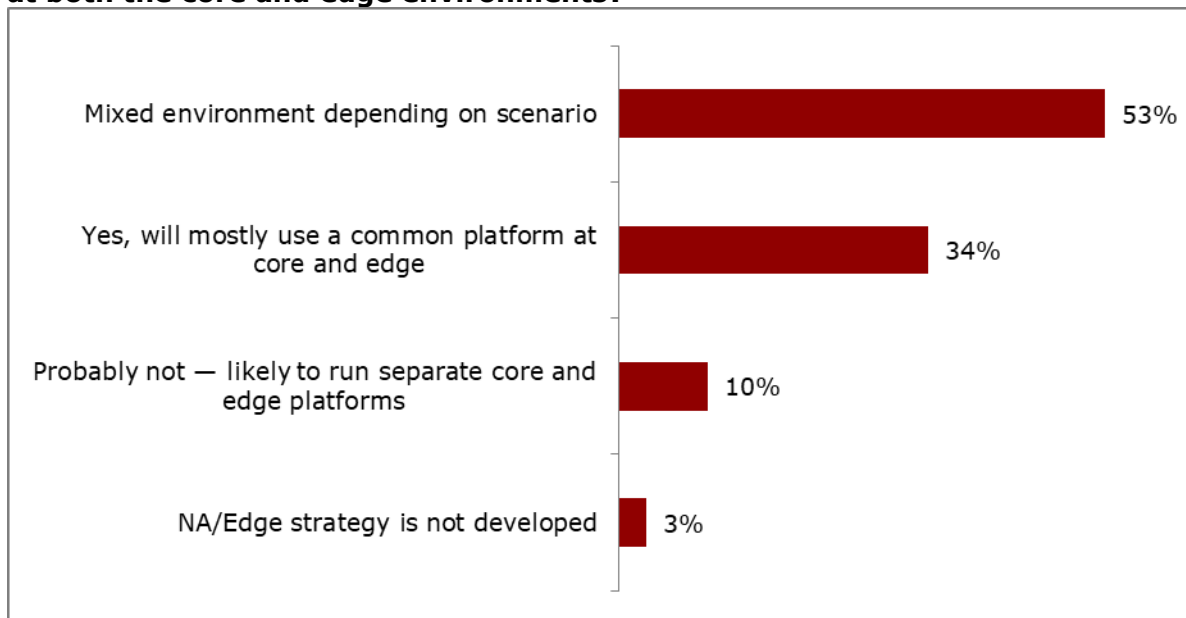
n=83

Source: Heavy Reading

The 5G architecture makes it easier for operators to distribute network functions and services to edge locations. This might be useful, for example, for services that require predictable low latency, for resiliency, or to minimize transport costs. In **Figure 13**, the survey asked if respondents expect to operate a common cloud infrastructure platform at both the core and edge environments. About half (53%) believe it will be “a mixed environment depending on scenario,” while about a third (34%) believe the reverse and “will mostly use a common platform at core and edge.”

This is an interesting response because, at face value, there are advantages to consolidating on a single cloud infrastructure platform—certainly at the infrastructure software layer—because it reduces overhead related to interoperability, testing, and maintenance and simplifies operations. Therefore, we might have expected a larger number opting for a common platform. However, given the question does not specify software infrastructure, respondents could have taken it to also mean hardware, and here there is a clearer case for the diversity of switching and server infrastructure at the edge. A second factor is that several high profile operators are publicly known to use a different cloud infrastructure stack for virtualized radio access networks at the edge than in their centralized data centers, where many core functions will reside. This will be an important area to watch in the coming years.

Figure 13: Will your organization operate a common cloud infrastructure platform at both the core and edge environments?



n=83

Source: Heavy Reading

Open source software is supported by community development and peer reviews. It should be easily accessible and open to modification, which can have multiple benefits in an agile development environment such as a 5G cloud infrastructure. The question in **Figure 14** seeks to understand the rationale behind an organization’s decision to use open source in the infrastructure software layer.

While all the offered answers are strong reasons to use open source within a telco cloud infrastructure, **Figure 14** shows weighted average scores for the respondent rankings with service agility (first) closely followed by strong security (second) rank as the most important. The ability to influence future features/capabilities (third) and easier application modernization and application integration (fourth) rank lower, but only by a small amount.

Figure 14: How important are the following in your organization’s decision to use open source within the cloud infrastructure that will support your 5G core?

Responses	Weighted Average Score	Rank
Increased service agility	265	1
Strong security	262	2
Ability to influence future features/capabilities	250	3
Better for application modernization	234	4
Easier application integration	234	4

n=83

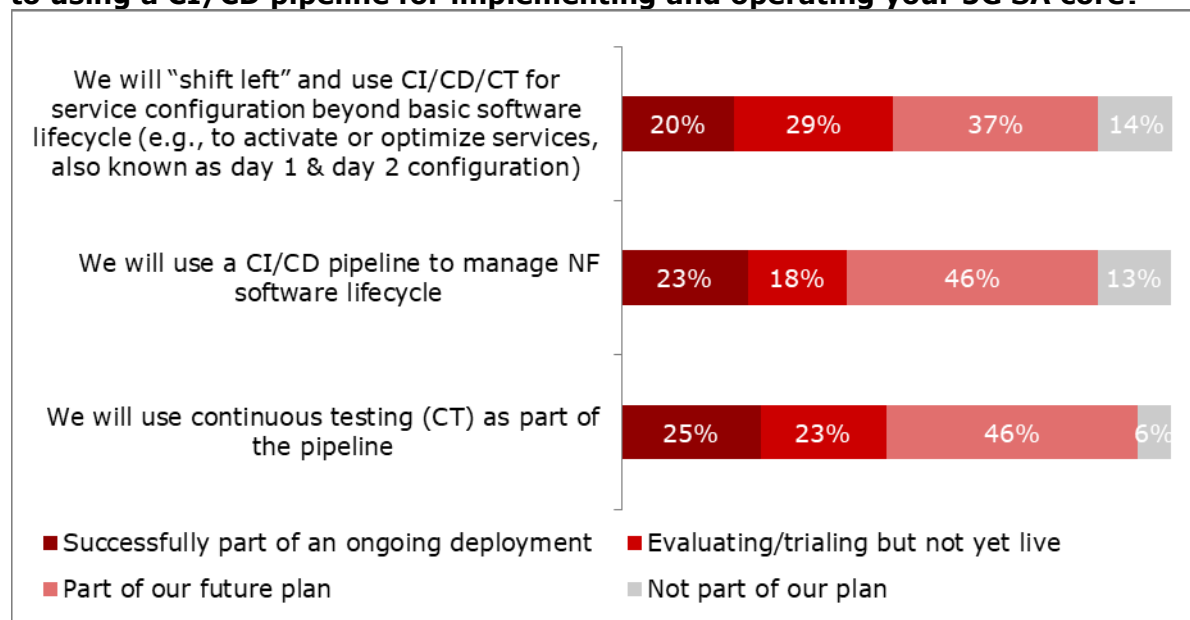
Source: Heavy Reading

Cloud Native Network Functions (CNFs) allow rapid changes to configuration and services via software revisions. Software lifecycle approaches, such as CI/CD, enable greater automation of these changes on shorter time scales with less risk of manual error. The CT and shift-left approach moves testing earlier into the development lifecycle (hence, “shift left”). Introducing these approaches into the 5G core software lifecycle, in principle, improves quality and agility through deployment pipelines.

Figure 15 shows what respondents believe to be the status of CI/CD/CT pipelines for deployment and management of their 5G SA core. The major finding is that about a quarter (23% and 25%) say their organization already uses CI/CD/CT successfully in live deployment, and 20% are also using it for “service configuration beyond the basic software lifecycle.” The 34 respondents that work at operators with more than \$5bn in annual revenue were more likely to be in commercial service with 38%, 38%, and 32%, respectively.

For various reasons, operators have been relatively slow to adopt modern cloud operating processes. Perhaps an advanced cohort of this size may indicate that a threshold has been crossed in CI/CD and that a larger number of operators will adopt these practices in the near future?

Figure 15: What best represents your organization’s current situation with regard to using a CI/CD pipeline for implementing and operating your 5G SA core?



n=83

Source: Heavy Reading

Containerization technologies, such as Kubernetes, originated in enterprise and public cloud environments. As 5G core implementation has begun, limitations in certain areas of containerized technology have emerged (e.g., in telco protocol support, networking, visibility, and network management).


In **Figure 16**, the survey asks, “what are the top containerization implementation issues your organization has encountered transitioning to 5G cloud native core?” Respondents were asked to select the top three and, on average, selected 2.4 responses each.

As expected, there is a large spread across the response. Security is the top answer, selected by 49% of respondents. The security challenge scales from something as simple as the Kubernetes security configuration not being default at the time of initial deployment to more fundamental issues that require deep expertise in telco cloud and 5G core security.

Visibility across containerized workloads within the same node is a well-known industry challenge, typically requiring additional solutions (e.g., container networking interfaces [CNIs], agents, and sidecars). Therefore, it is no surprise that “visibility, operate and manage/troubleshoot” is a very close second (48%).

The spread of answers across the other areas, including “reliance on a single vendor” (39%), “separate clusters” (37%), and “circumventing Kubernetes networking” (35%), illustrates that Kubernetes and container technologies are still reaching maturity. However, it is promising to see that “Kubernetes is not yet mature enough for telco protocols” gains the lowest number of votes, indicating that progress has been made.

Figure 16: What are the top containerization implementation issues your organization has encountered in transitioning to a 5G cloud native core? (Select top three)



n=83

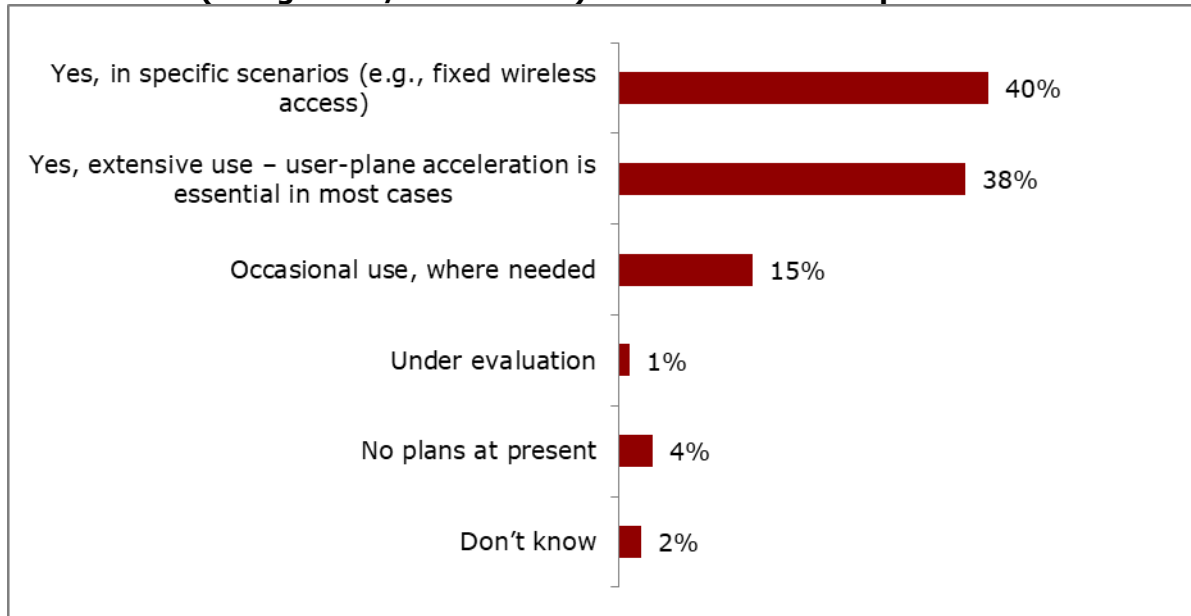
Source: Heavy Reading

Hardware acceleration for certain cloud workloads (e.g., for artificial intelligence/machine learning [AI/ML] and video processing) has been developing over several years. In the 5G core, the rise of high capacity and low latency workloads is driving interest in using data processing units (DPUs) or SmartNICs to optimize processing. This can directly reduce the number of CPU cores required and overall power consumption but introduces friction to workload portability. There is a very active discussion about the need for hardware acceleration in the 5G core, framed around the merits of software agility versus hardware performance optimization.

Figure 17 asks specifically about hardware acceleration in the data plane. “Yes, for specific scenarios (e.g., fixed wireless access)” scores highest with 40%. This is a logical result because fixed-access services generate much greater throughput per connection than mobile services, and the case for hardware acceleration is therefore stronger.

The 38% score for “yes extensive use—user-plane acceleration is essential in most cases” is interesting. On the one hand, it shows a majority do not have this view. On the other hand, it is a strong score and indicates a widespread expectation that accelerators will be important in 5G core deployments to handle the projected increase in user-plane traffic from 5G connections. The overall analysis, therefore, is that this survey does not help resolve the debate between software and hardware user-plane performance optimization but merely continues it.

Figure 17: Does your organization expect to make extensive use of hardware acceleration (using DPUs/ SmartNICs) in its 5G core user plane?



n=83

Source: Heavy Reading

ENVIRONMENT AND SUSTAINABILITY

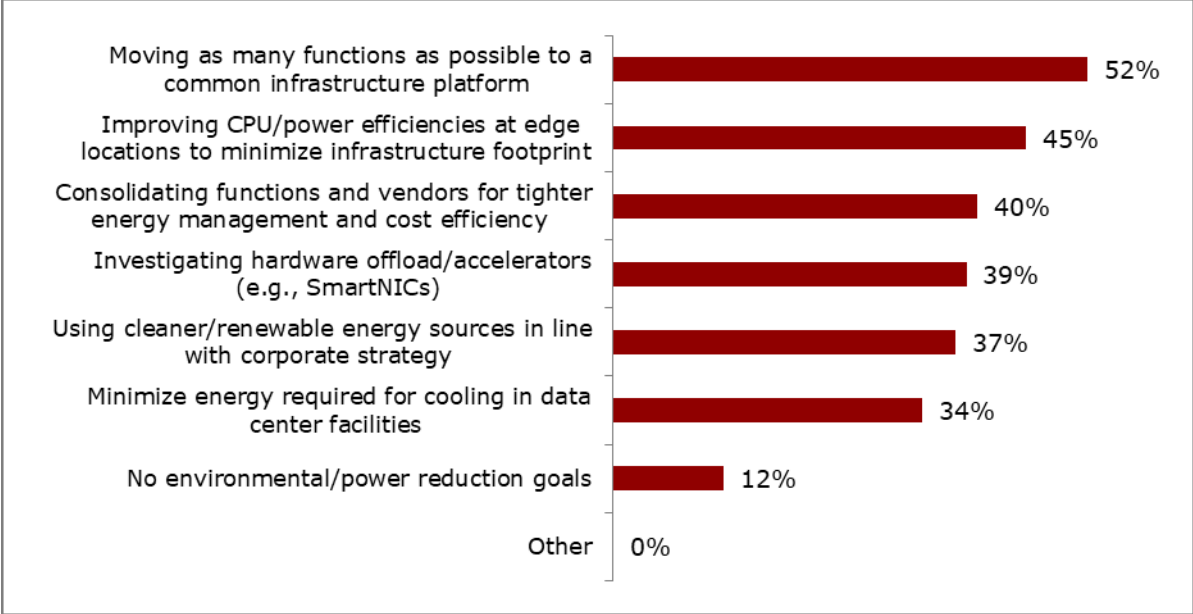
Environmentally sustainable networks are of great importance, especially as they relate to energy efficiency. One question in the survey addresses methodologies and approaches that will help operators meet their environmental targets in the 5G core network.

As background context, it is well known that most operators are looking at power reduction initiatives, as well as “tech cleanups,” to remove legacy equipment that is notoriously power-hungry from their networks. Several recent announcements have also indicated substantial power savings by moving and consolidating workloads from multiple disparate appliances to a common telco cloud platform.

Figure 18 looks at a number of energy-saving options for the 5G core and asks what actions can help operators meet their environmental goals. Respondents were asked to select a top three and selected an average of 2.6 answers each. The top answer—“moving as many functions as possible to a common infrastructure platform” (52%)—mirrors the approach being taken by many operators today. From a core network perspective, consolidating on a common cloud platform is probably the most consequential move operators can make to reduce energy consumption.

There are, however, many additional ways to improve the power efficiency of the cloud platform itself. This explains the answers to the other energy efficacy options offered, from improving CPU/power efficiencies (45%), consolidating functions (40%), investigating hardware offload/accelerators (39%), using cleaner/renewable energy (37%), and minimizing cooling (34%). This spread in the response indicates operators will combine multiple approaches in their power reduction strategy.

Figure 18: What changes will your organization make to its 5G core infrastructure to meet environmental goals and reduce power consumption? (Select top three)



n=83
Source: Heavy Reading

5G ANALYTICS AND NETWORK EXPOSURE

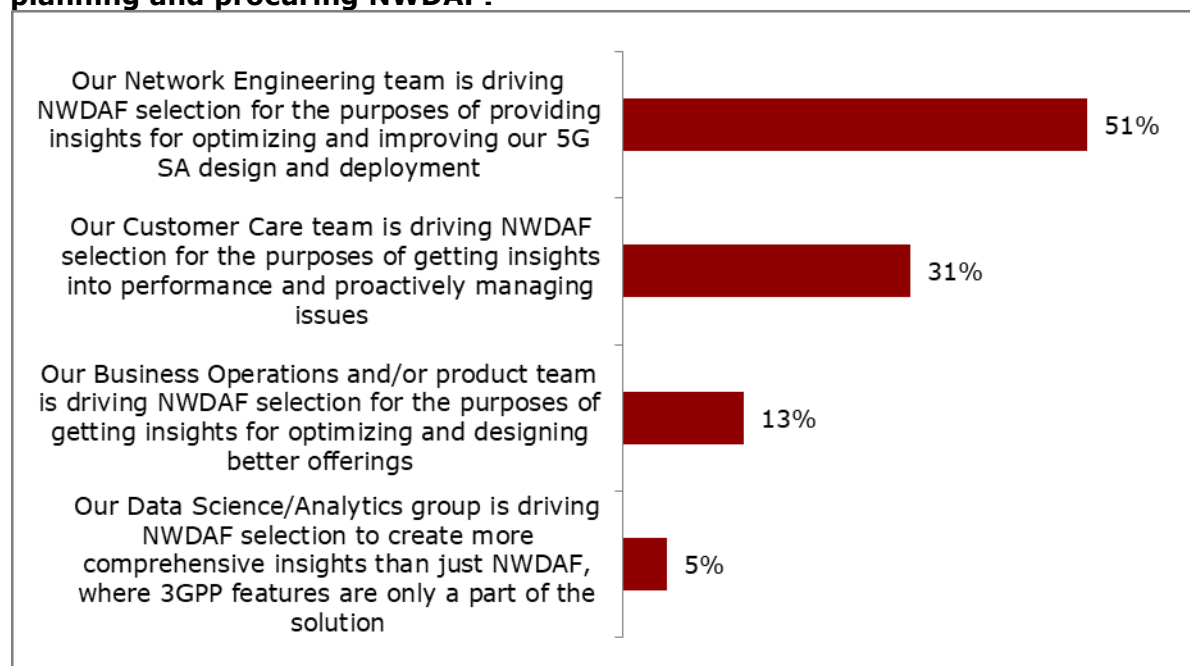
The 5G NWDAF provides a solution for streamlining how analytics data from the core network is produced and consumed and how action is taken on the data. It is a new function in the mobile core architecture. It is fair to say that demand for the NWDAF has largely been secondary to the main rollout of the 5G SA core; however, with maturing standards and the growing importance of analytics, it is now well supported by vendors.

Figure 19 illustrates how operators plan to procure NWDAFs. About half the respondents (51%) expect their engineering team to drive the selection. There is a close coupling between the NWDAF and other 5G core NFs and NetOps tools, which require engineering team insights and knowledge. This explains this vote.

The next highest category is the 31% that says the customer care team will drive the procurement process to get better insights into performance and customer experience. Analytics can have a big impact on reducing churn and ensuring customer satisfaction; therefore, it is reasonable that this constituency has an influence on purchasing analytics solutions.

The lower scores for business operations (13%) and data science/analytics groups (5%) to drive the NWDAF procurement do not mean these groups are inconsequential. Although the engineering teams will run the main process, there will be significant input from related teams within the operator, and it is virtually certain that requirements from these two constituencies will be incorporated into the process.

Figure 19: Which statement best describes your organization's process for planning and procuring NWDAF?



n=83

Source: Heavy Reading

As service providers migrate their core networks to 5G, there are opportunities to upgrade the SGi/N6 LAN domain. Sitting outside of the main 3GPP standardization path, the SGi/N6 LAN is important to how customers connect to the internet and into enterprise and cloud environments. It is critical to network security and important to enable additional service functionality. The move to distributed edge use cases has major implications for the N6 interface and implies that functions in the N6 LAN—especially security—must also migrate to the edge.

Figure 20 addresses the upgrade path from SGi-LAN to N6, asking respondents to rank which service analytics are of the highest priority in the transition. “DDoS protection and firewall alerts” was rated highest, followed by “volumetric traffic managements across apps and mobile network.” In the 5G era, the focus on security is greater than ever, and it is not surprising that both DDoS protection and firewall alerts are the highest priority in the N6 LAN. And in the era of unlimited plans, data rollovers, and zero-rated apps, “volumetric traffic management across apps and mobile networks” has clear importance, accounting for its second-place rank.

Figure 20: Which SGi-LAN service analytics are of the highest priority to your upgrade path as you transition to 5G SA with N6 LAN?

Responses	Weighted Average Score	Rank
DDoS protection and firewall alerts	250	1
Volumetric traffic management across apps and mobile network	206	2
Subscriber change/tracking	192	3
Video optimization metrics	182	4

n=83

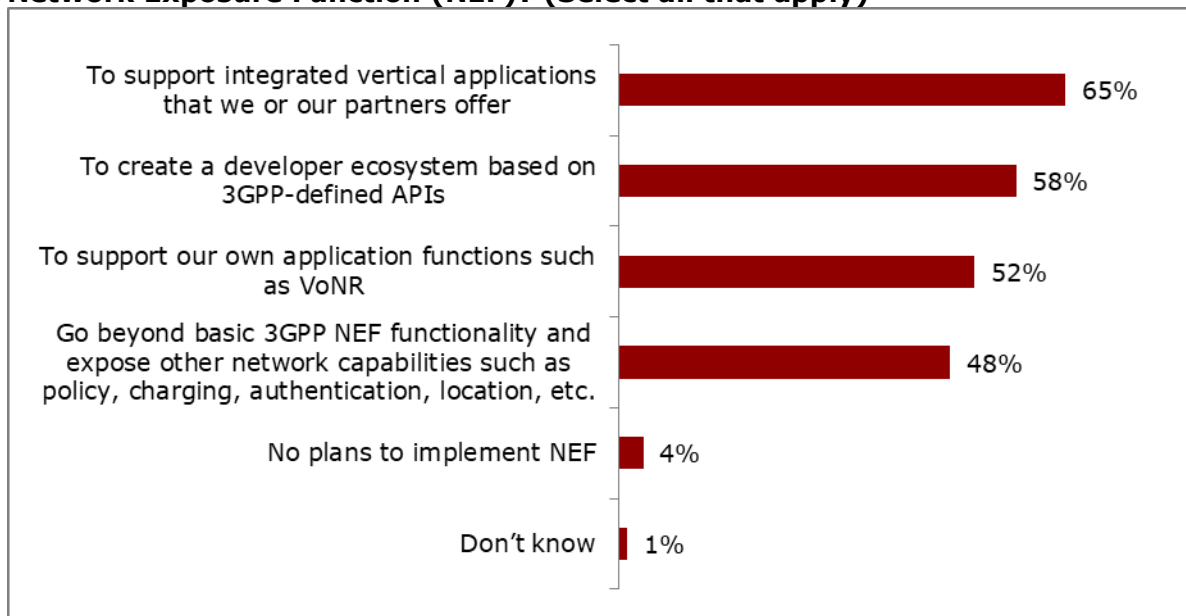
Source: Heavy Reading

5G SA will help service providers be more agile in how they introduce and enable new services due in part to capabilities such as network APIs. The NEF facilitates exposing network capabilities securely to internal services, partners, and customers. This is a step toward a robust and developer-friendly ecosystem that can support vertical industries, developers, and internal application development.

Figure 21 addresses operator views on how NEF may be used. The question asks respondents to “select all that apply” for an average of 2.3 responses per person. It is reassuring to see that only a very slim minority of choices have “no plans” (4%) or “don’t know” (1%), indicating that respondents believe NEF offers value.

“Vertical applications that we or our partners offer” (65%) was the top answer and indicates operators will initially be more comfortable working internally and with carefully curated partners on network capability exposure. However, the scores for “creating a developer ecosystem” (58%) and “going beyond basic NEF functionality” (48%) are also strong and indicate operators are optimistic that 5G networks may become more open technology platforms on which other parties can innovate and profit.

Figure 21: Which of the following reflects your organization’s view for using the Network Exposure Function (NEF)? (Select all that apply)



n=83

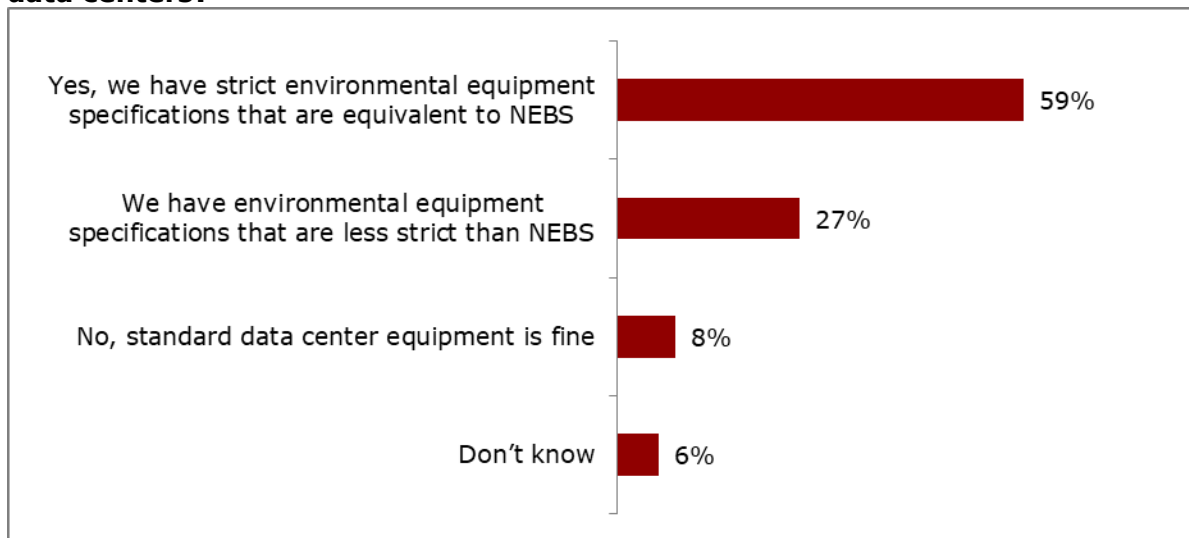
Source: Heavy Reading

EDGE DEPLOYMENT FOR 5G CORE

The 5G architecture allows operators to consider edge deployment scenarios to support specific use cases. Low latency use cases, for example, often require the deployment of a UPF toward the edge of the network to support more stringent SLAs and/or to support local breakout to a private network, a cloud, or the internet.

The survey confirmed that operators have specific environmental requirements for core network equipment deployed at edge locations. Most telecommunication equipment adheres to the US-specified network equipment-building system (NEBS) or similar guidelines, which document how to run equipment in various levels of controlled environments. When asked if these guidelines would also be applied at the edge (see **Figure 22**), a majority (59%) have equipment guidelines equivalent to NEBS at the edge, and this rises to 71% in the US versus 46% for RoW. Organizations using a form of “environmental equipment specifications that are less strict than NEBS” (27%) is second, “no, standard data center equipment is fine” (8%) is third, and finally, only 6% “don’t know.”

Figure 22: Does your organization have a strict set of safety and environmental guidelines/Network Equipment-Building Practice (NEBS) requirements for edge data centers?

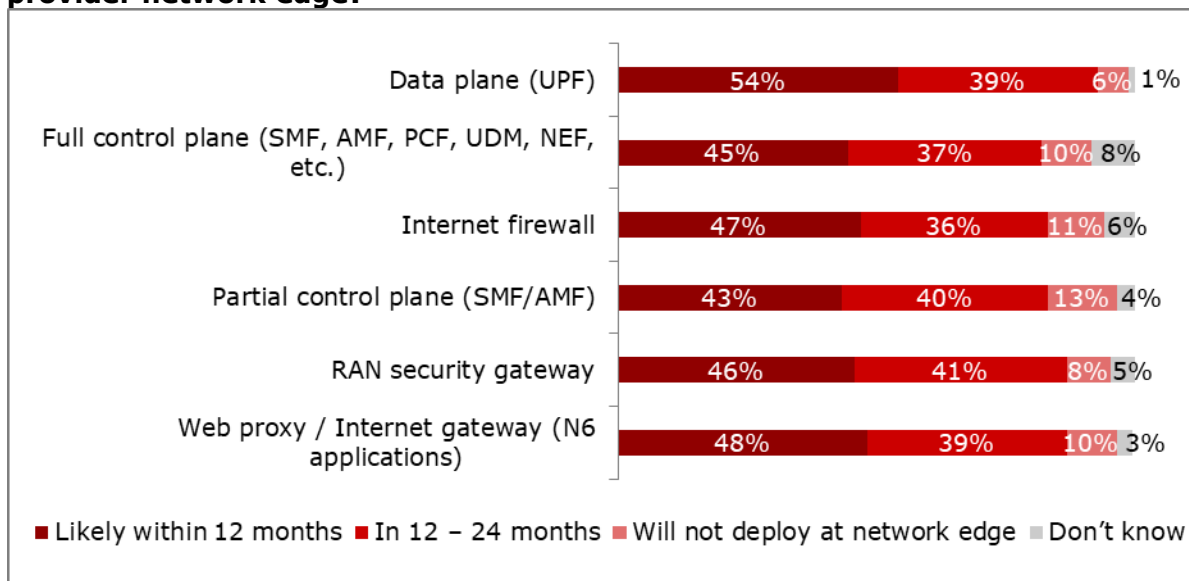


n=83

Source: Heavy Reading

According to the survey, most operators will deploy 5G core functions at the edge within two years. **Figure 23** shows that about half will do so in the next 12 months and 35–40% in the next 24 months. There is not much difference in the timelines for the functions listed in the question, perhaps indicating that when operators deploy an edge user plane, they will move to a relatively complete core deployment at the same time.

Figure 23: When will your organization deploy 5G core functions at the service provider network edge?



n=83

Source: Heavy Reading

5G SA ROAMING AND SECURITY

This section of the survey asks a broad security strategy question and seeks opinions on 5G SA roaming and the associated security of roaming interconnects.

In **Figure 24**, security capabilities are ranked in terms of importance to an organization's 5G core network strategy. The top-ranked answer is "a security and governance framework which includes consistent and compliant configuration management," followed by "security implemented as part of an overall DevSecOps strategy across your organization," in second place. All the answers are ranked closely on the weighted average score, indicating that, as expected, all these capabilities are considered important. It is perhaps interesting, therefore, that the high rank for security in the context of configuration management infers operators are focused on the need to build in 5G core security at the start of the deployment and will place more emphasis on the application lifecycle later in the operating phase.

Figure 24: How important are the following security capabilities to your organization's 5G core network strategy?

Responses	Weighted Average Score	Rank
A security and governance framework that includes consistent and compliant configuration management	220	1
Security implemented as part of an overall DevSecOps strategy across your organization	213	2
A platform with integrated and layered security	201	3
Continuous security checks throughout the entire application lifecycle	196	4

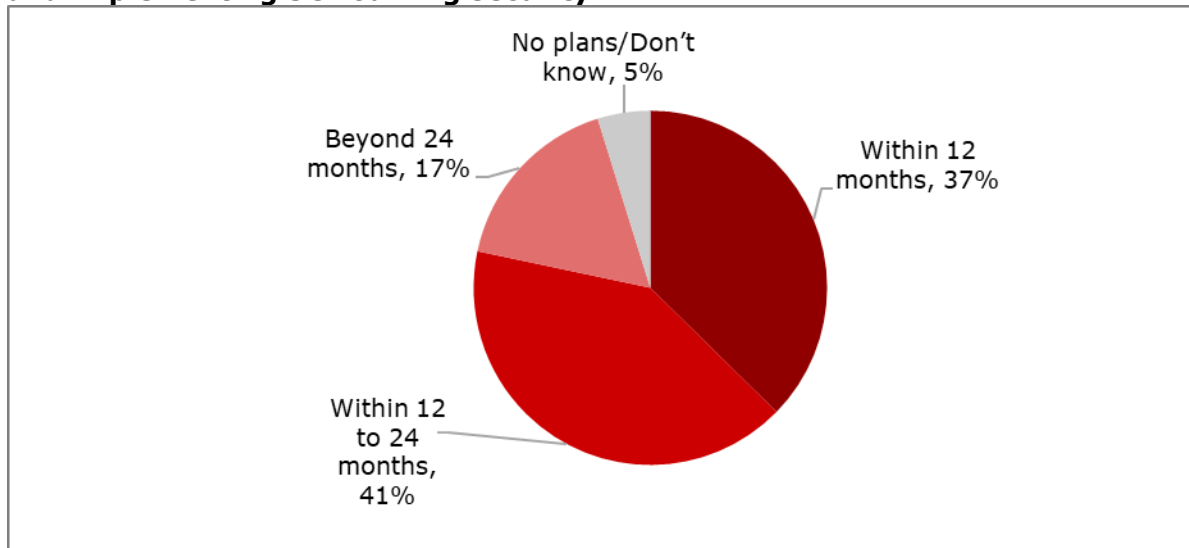
n=83

Source: Heavy Reading

5G SA roaming will be important to many 5G users, from smartphone customers through to enterprise IoT devices that cross national borders. It is dependent on SA services being live and on roaming arrangements (technical and commercial) with partners in other national markets. This is a complex process with various implementation choices and security considerations to resolve. To Heavy Reading's knowledge, at the time of writing (January 2023), there are no live SA roaming arrangements; however, there are a small number of trials under way, with a target of live services in the second half of 2023 if the trials go well.

Respondents were asked (see **Figure 25**), "when does your organization expect to start offering 5G SA roaming and implementing 5G roaming security?". The respondents split between "within 12 months" (37%) and "within 12 to 24 months" (41%). Considering the complexity, timelines for 5G SA rollout, and the need for operator roaming trials before a launch, the 12 to 24 months seems to be a more realistic timeframe for advanced and fast-follower operators. Indeed, filtering only the 26 respondents in network planning roles, the number expecting to launch within 12 months drops to 23%, while the number that anticipates beyond 24 months rises from 17% to 31%.

Figure 25: When does your organization expect to start offering 5G SA roaming and implementing 5G roaming security?

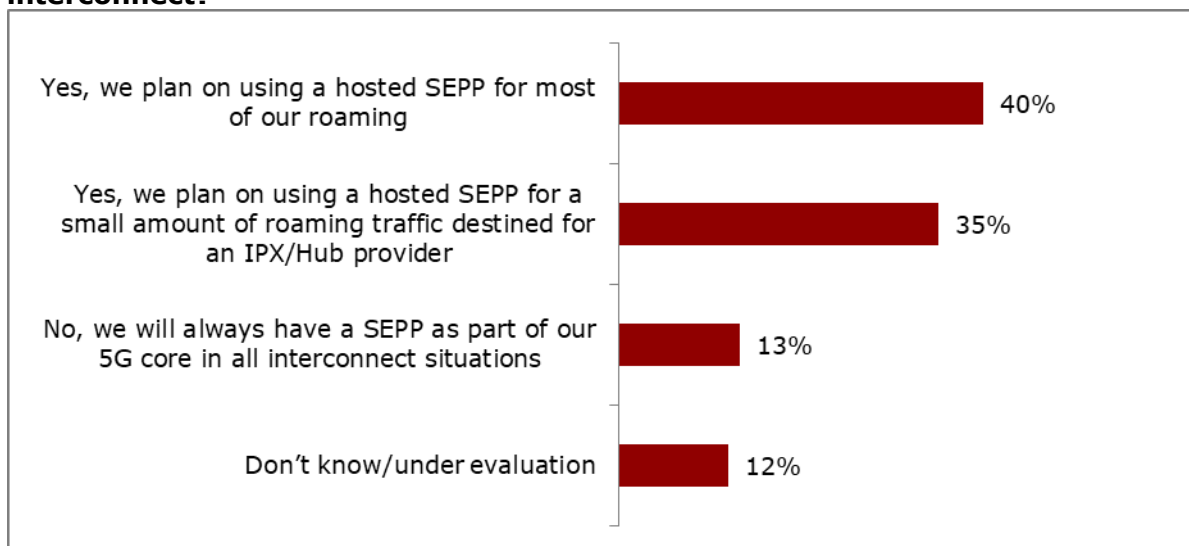


n=83

Source: Heavy Reading

Figure 26 asks respondents about selecting the Security Edge Protection Proxy (SEPP) and whether their organization would consider having a third party run it on a hosted basis. The top answer to this question is “yes, we plan on using hosted SEPP for most of our roaming” (40%), followed by “yes, we plan on using a hosted SEPP for a small amount of roaming traffic destined for an IPX/Hub provider” (35%). Surprisingly, given the maturity of 5G SA roaming, both “no, we will always have a SEPP ...” (13%) and “don’t know/under evaluation” (12%) came out very low.

Figure 26: Would your organization consider having a third party (e.g., an IPX/hub provider) run your SEPP (Security Edge Protection Proxy) for roaming interconnect?



n=83

Source: Heavy Reading

As mentioned in **Figure 25**, many operators are planning to launch 5G SA roaming and are studying both hosted and SEPP-SEPP interconnect. At this stage, it might be too early to get a full, representative opinion on whether SEPP is likely to be hosted or run on the operator's 5G core platform. However, considering the complexity and previous roaming generations, IPX providers commonly hosted 4G Diameter Edge Agent (DEA), the 4G equivalent to SEPP. This could explain why a combined 75% (40% and 35%) of respondents have answered that SEPP will be hosted.

BACKGROUND TO THIS STUDY

Heavy Reading's 2022 **Cloud Native 5G Core Operator Survey** was conducted in September 2022. This analysis was written in November and December 2022 and published in January 2023. The online survey generated 83 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.