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

HEAVY READING
WHITE PAPER

5G Edge Cloud Infrastructure: For vRAN, Industrial, and Automotive Applications

A Heavy Reading white paper produced for Wind River

AUTHOR: GABRIEL BROWN, PRINCIPAL ANALYST, 5G, HEAVY READING

5G EDGE CLOUD AND THE NEW NETWORK ARCHITECTURE

Operators worldwide are deploying 5G at pace, with new devices coming to market all the time and 5G penetration on track to exceed 25% by 2023 in advanced markets. And the 5G era is only just getting started. The mobile industry is still in the early phases of a long-term transition that will see integrated 4G/5G networks and service propositions evolve as new capabilities and network architectures are introduced.

One of the most anticipated changes in the 5G network is the deployment of edge cloud infrastructure capable of supporting network functions, media content, and end user applications. This new nexus in the mobile network architecture enables operators to offer higher performance services – particularly low latency services – that are impossible or impractical to deliver with 4G and cannot be effectively delivered from large, centralized data centers.

Definitions vary, but in operator networks, the "infrastructure edge cloud" is generally agreed to span the terrain between large centralized data centers, through "near edge" regional and metro-area locations, and into the "far edge" access network. This paper focuses primarily on edge clouds deployed in the radio access network (RAN) by mobile operators and then extends the discussion to on-premises edge clouds for industrial applications and to roadside units (RSUs) for automotive services.

The paper argues that edge cloud is integral to the 5G system design and to operators' ability to offer high performance services to advanced enterprise customers. It draws on a Heavy Reading survey of professionals working for communications service providers, automotive companies, industrial companies, and general enterprises to identify the lead applications, technology solutions, and challenges to deployment.

The survey questions were developed by Heavy Reading in association with Wind River. The questionnaire was fielded to respondents in January 2020. It generated a total of 141 responses from network operators, 55 from automotive companies, and 46 from industrial and general enterprises. Technical and engineering roles predominate in the responses; however, corporate management and customer-facing roles are also well represented.

Key Findings and Executive Summary

Among the key findings from the survey are the following:

38% of operator respondents identified the edge cloud as "critical" to their 5G strategy and 47% as "important." When asked if network functions or end user applications would drive the case for deployment, "network functions" (39%) came out just ahead of "end user applications" (36%). A quarter (25%) selected "both at the same time."

The software infrastructure on which edge applications will be deployed is critical. The largest number of operator respondents (34%) selected a combination of OpenStack and containers/Kubernetes as their preferred platform. This would enable them to continue with their stable, tested virtualized infrastructure and applications and migrate to "cloud native" over time. The smallest number (15%) selected "direct to cloud native" and a quarter each opted for "proprietary virtualization" (26%) and "OpenStack" (25%).

Half of respondents expect to have virtualized RAN (vRAN) in commercial service within the next 2 years, of which 29% expect this to be "at scale" and 21% at "small-scale" deployment. This is an important endorsement because to make edge cloud platforms economical, operators need workloads that will justify the investment in facilities, hardware, and software infrastructure, as well as the ongoing operational expense. Some argue that vRAN should be the anchor tenant for this "far edge" cloud because the RAN is high volume and widely distributed and is, therefore, a large addressable market.

Operators are bullish on vRAN, but not naive. The major challenge for vRAN is "systems integration" (58% of the response), followed by the "operational complexity" of multi-vendor vRAN (46%). "Edge cloud server performance" comes in third (36%).

In terms of performance, a majority think hardware accelerators are "critical" or "important" to run vRAN baseband software on x86 servers. There is a clear difference between the 43% that believe accelerators are critical for high capacity sites and the 27% that think accelerators are critical in general. This will be an interesting area to track in the coming years as general-purpose silicon becomes available that is better suited to Layer 1 RAN processing.

A majority (60%) believe it will be practical to virtualize a 5G baseband (a distributed unit [DU]) function for commercial deployment within 2 years. Heavy Reading believes feature and performance parity with vendor-integrated 5G base stations will be challenging in this time period and tends toward the view that virtual a 5G baseband will be achieved, but with some performance compromises. Asked if and how vendor-integrated base stations might have advantages over vRAN, the largest group of respondents (59%) said the "ability to meet deterministic low latency performance requirements" is a "major advantage."

Enterprise respondents do not believe LTE or Wi-Fi will meet the performance requirements of industrial use cases with scores of just 12% and 2%, respectively. Clearly, the perception is that 5G is superior for industrial applications. Just over a quarter (28%) appear sold on 5G already, saying "performance advantages will force its use" where wireless is needed in industrial networks. The majority (58%), however, say that the claimed, theoretical advantages of 5G still "need to be proven in practice." They are favorable toward the technology, but are not yet convinced.

Among respondents in the automotive sector, when asked which cellular vehicle to-everything (C-V2X) applications will lead deployment, vehicle-to-network is the clear leader with 56%. In second place, with 36%, is vehicle-to-infrastructure. Vehicle-to-vehicle (V2V) scores almost as highly at 35%. Vehicle-to-pedestrian has high potential to improve road safety, but requires technology and policy development to be reliable; it is no surprise that it scores lowest at 9%.

Asked if they believe 5G is needed for Level 4 (L4) autonomous driving, the results are interesting: a majority (64%) do not think 5G is required. However, 27% say a "network connection is needed in practice" and 22% say L4 autonomous driving would "benefit from network assist even if it is not necessary." These views suggest an important role for network connectivity and are aligned with Heavy Reading's view.

Intriguingly, a third of respondents (34%) plan to integrate satellite into their 5G automotive offering. However, just as many have "no plans" (35%) and "don't know" (31%). Heavy Reading notes that Elon Musk, head of Tesla, the world's most innovative car manufacturer, recently tweeted about 5G and satellite communications.

EXPECTATIONS FOR MASS-MARKET 5G SERVICE

To begin the survey, Heavy Reading asked operator respondents when they expect 5G to be the standard, mainstream mobile service for both enterprises and consumers. In other words, when will 5G become like 4G is today? When will 5G be simply what you get when you renew or switch mobile service plans? Although there are a few 5G smartphones in use today connected to 4G networks, in the vast majority of cases, this also involves customers buying new devices.

Figure 1: When do you expect 5G to be the de facto, mainstream solution for consumer and enterprise?

	2021 year-end or before	2022 year-end	2023 year-end	2024 or later
Consumer	40%	33%	16%	11%
Enterprise	36%	37%	16%	11%

N=141
Source: Heavy Reading

As shown in **Figure 1**, a majority of operator respondents expect that within 3 years (i.e., by the start of 2023), 5G will have become the mainstream mobile service. A significant number (36% enterprise; 40% consumer) think this will happen sooner (i.e., by the start of 2022). Even if this response is over-optimistic – since it reflects the views of professionals in this sector that are working in advanced economies – it suggests the mobile industry is set to go through a period of rapid, and deep, change.

5G EDGE CLOUD INFRASTRUCTURE

Among operator respondents, 38% identified the edge cloud as "critical" to their 5G strategy and 47% as "important." This is a strong endorsement. Heavy Reading then asked if network functions or end user applications would drive the case for deployment. By a narrow margin, "network functions" (39%) came out ahead of "end user applications" (36%). A quarter of respondents (25%) selected "both at the same time," as shown in **Figure 2**. The picture painted by this result is that most individual operators will lead with one or the other strategy, but the industry as a whole will move on both strategies at once. For vendors, this means support for both is important.

Figure 2: Which application types will be deployed first on your edge cloud infrastructure?

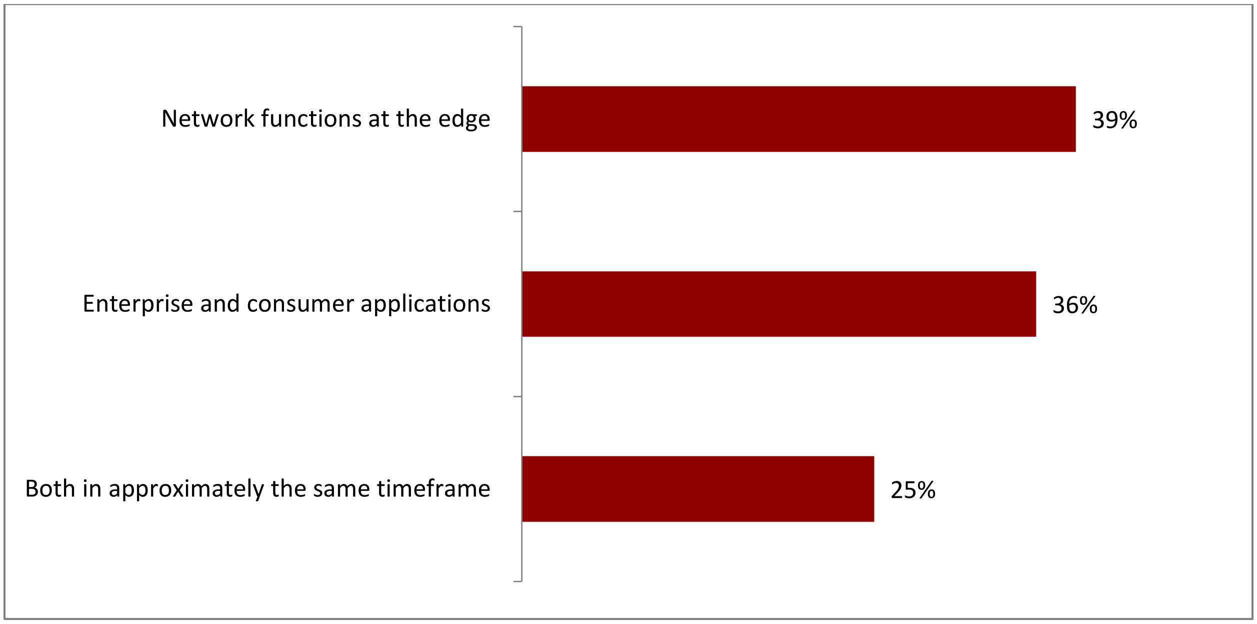

N=141

Source: Heavy Reading

When Heavy Reading delved more deeply into which specific application types are attractive at the edge, the same pattern emerged. **Figure 3** shows the mobile core, a set of network functions, ranked highest (42% "extremely attractive"). This ranking is consistent with a view that the distributed core is an enabler of edge services and, in some respects, a prerequisite. In second place, "gaming and virtual and augmented reality" (VR/AR; 37% "extremely attractive") are popular end user applications characterized by the need to reliably support high bit rate video, very fast response times, and low error rates. Clearly, these are good candidate applications for the edge.

Figure 3: How commercially attractive are the following edge cloud applications to your company within the next 2 years?

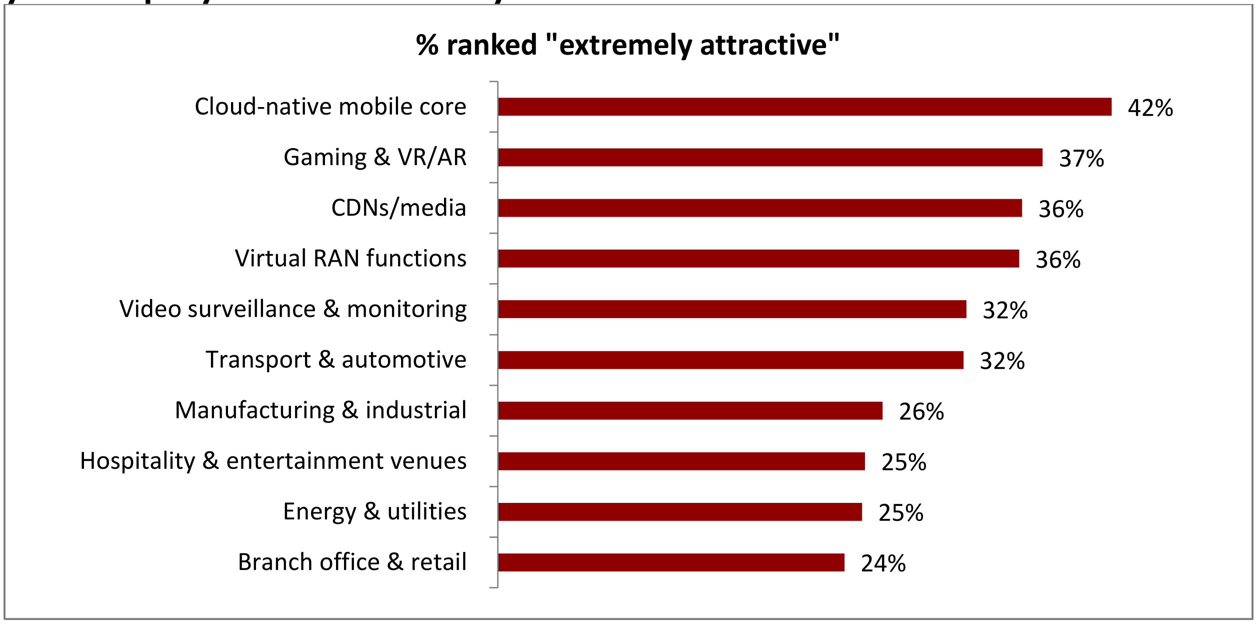

N=140
Source: Heavy Reading

The software infrastructure on which edge applications will be deployed is critical. Classic cloud infrastructure software is designed for large centralized data centers and can be challenging to adapt to the edge because of the different operating models, the larger number of locations, and the smaller hardware footprints at the edge.

Broadly speaking, operators have three options for edge cloud software infrastructure: an OpenStack distribution, proprietary virtualization/cloud platforms, or so-called "cloud-native" platforms (e.g., bare metal with containers and Kubernetes orchestration). The survey asked which of these high level options operator respondents might deploy.

The response, shown in **Figure 4**, does not clarify the picture enormously, but it does reflect that this is a complicated decision with many moving parts. The assumption is that operators want to be cloud native; however, not only is this not tightly defined, it also does not consider the readiness of network functions and operating processes. For example, virtualized network functions (VNFs) can run in virtual machines (VMs), but not containers. This requires the evolution of the VNF to a cloud-native network function (CNF).

The smallest number (15%) selected "cloud native" only, probably because this is an emerging technology and telecom operators have less experience with. A quarter each opted for "proprietary virtualization" (26%) and "OpenStack" (25%). The largest number (34%) selected "a combination" of OpenStack and Kubernetes, which would enable them to continue with their stable, tested virtualized infrastructure and applications, migrating to cloud native over time.

Figure 4: Thinking about your 5G edge cloud requirements, what type of software infrastructure will you deploy?

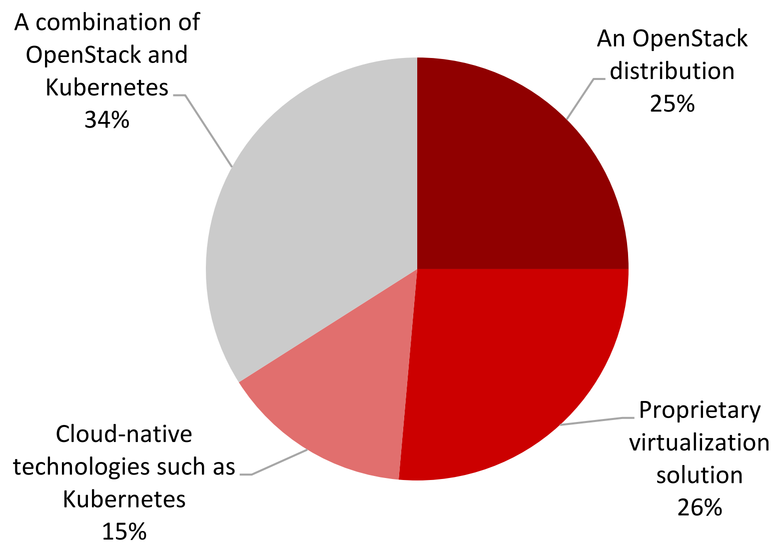

N=144
Source: Heavy Reading

Edge cloud locations need to be connected via underlay and overlay networks, controlled by software-defined networking (SDN). The survey response, shown in **Figure 5**, indicates that multiple SDN strategies are likely. A solid 27% opted for the "proprietary solution" and 34% for "vendor supported open source" – both options are in line with expectations. The largest group, the 39% that selected "open source networking solution," needs some explanation. Previous Heavy Reading surveys have determined that operators generally prefer vendor supported SDN solutions, typically based on open source. One possible explanation for why this may be different in this case could be that edge cloud has characteristics that make open source more attractive and/or that edge cloud infrastructure software comes with networking pre-integrated.

Figure 5: Thinking about your SDN requirements for geo-distributed cloud, what type of networking solution will you deploy?

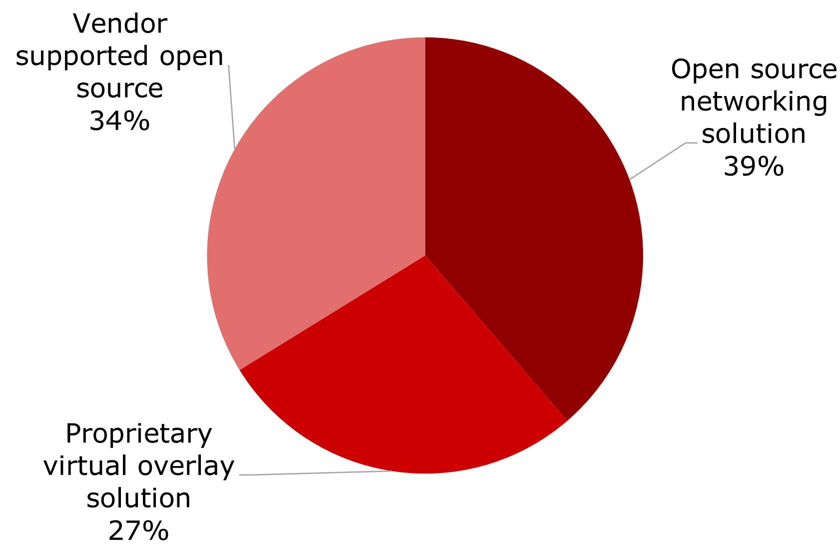

N=142
Source: Heavy Reading

vRAN DEPLOYMENT AND EDGE CLOUD

A converged edge cloud infrastructure that can support a combination of vRAN, core, media, and end user multi-access edge computing (MEC) applications would create a new nexus in the mobile network architecture and enable new modes of services delivery.

To make cloud platforms economic in the RAN, which is the closest part of the network to the application execution environment, operators need workloads that will justify the investment in facilities, hardware, and software infrastructure. Some argue that vRAN should be the anchor tenant for this "far edge" cloud because the RAN is already high volume and widely distributed; therefore, vRAN is one of the largest addressable markets for infrastructure edge cloud. Along this line of thinking, multi-purpose edge cloud deployment can "ride" the vRAN buildout.

The questionnaire tested how interested operators are in vRAN and how likely they are to trial and deploy this technology in their live networks. Based on this survey, expectations are running high. **Figure 6** shows that half of respondents expect to have vRAN in commercial service within the next 2 years, of which 29% expect this to be "at scale" and 21% a "small-scale" deployment.

Figure 6: What is your current thinking on deployment of virtual RAN in your network over the next 2 years?

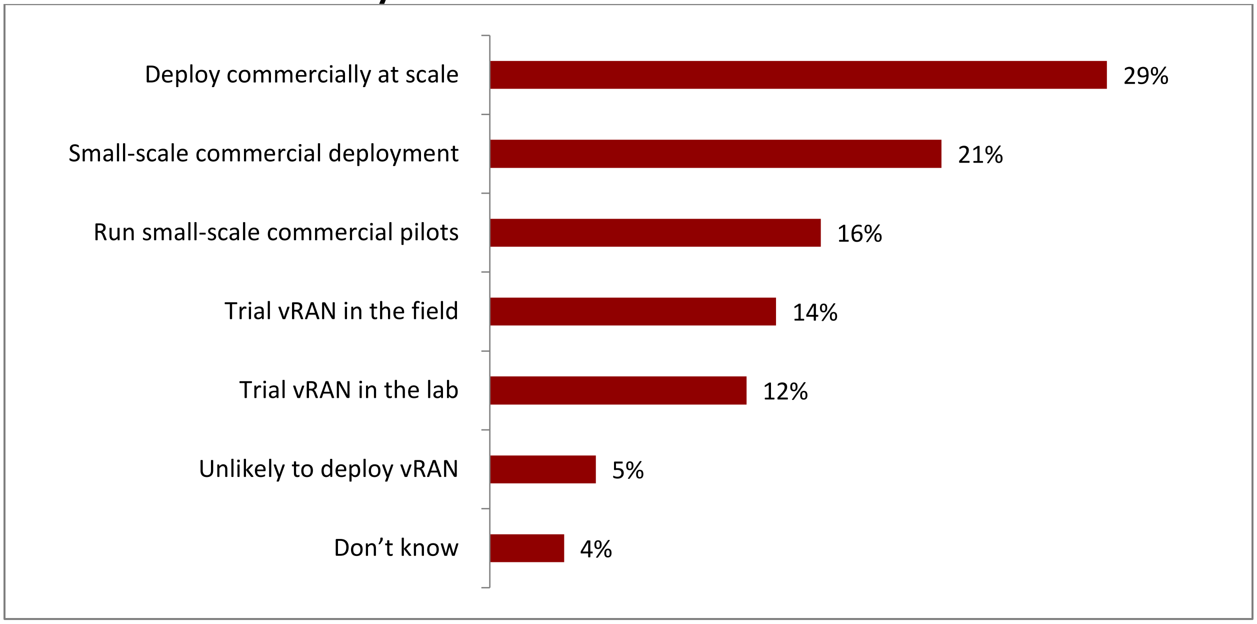

N=141

Source: Heavy Reading

This could be seen as a bullish response given that the first vRAN deployments have only recently gone commercially live, with the largest deployment consisting of just a few hundred sites. However, in Japan, Rakuten Mobile will launch 4G service in April 2020 with more than 4,000 macro vRAN sites deployed on edge cloud facilities. It intends to launch 5G on the same edge infrastructure later in 2020. Therefore, within 2 years (i.e., by the end of 2021), it is plausible that other operators will be able to scale vRAN. Keep in mind that "at scale" is subjective and does not mean the entire network is converted to vRAN. Note also that vRAN is sometimes (legitimately) understood to include partially virtualized 5G vRAN deployments where only the central unit, which processes the Layer 3 radio control plane, is virtualized.

On the other side of the scale, 71% do not expect to deploy "commercially at scale" over the near term. Thus, it could also be argued that this result shows that operators are positive about vRAN, but are not over-committed to unrealistic deployment schedules.

Deploying edge cloud for vRAN is not straightforward at this time; it involves new technology with demanding performance requirements. In the next question, shown in **Figure 7**, respondents were asked to select their top three challenges. (The question generated 356 responses by 140 survey takers for an average of 2.5 responses per person.)

The lead challenge, "systems integration" (58%), speaks to the difficulties associated with setting up a vRAN using hardware and software from multiple vendors, and then testing and verifying performance. "Operational complexity" of multi-vendor vRAN comes second with 46%. Both results are in line with expectations. Interestingly, "edge cloud server performance" comes in third (36%); this point is picked up again in the following question.

Challenges like operations and business support system (OSS/BSS) integration and local market factors rank significantly lower; this is perhaps because most respondents do not yet have direct experience running a vRAN. During Heavy Reading's discussions with operators and vendors that do have live vRAN deployments, these are both identified as important challenges that come to the fore when systems are actually deployed.

Figure 7: What are the biggest challenges to deploying distributed edge cloud for virtual RAN (vRAN)? (Select up to three)

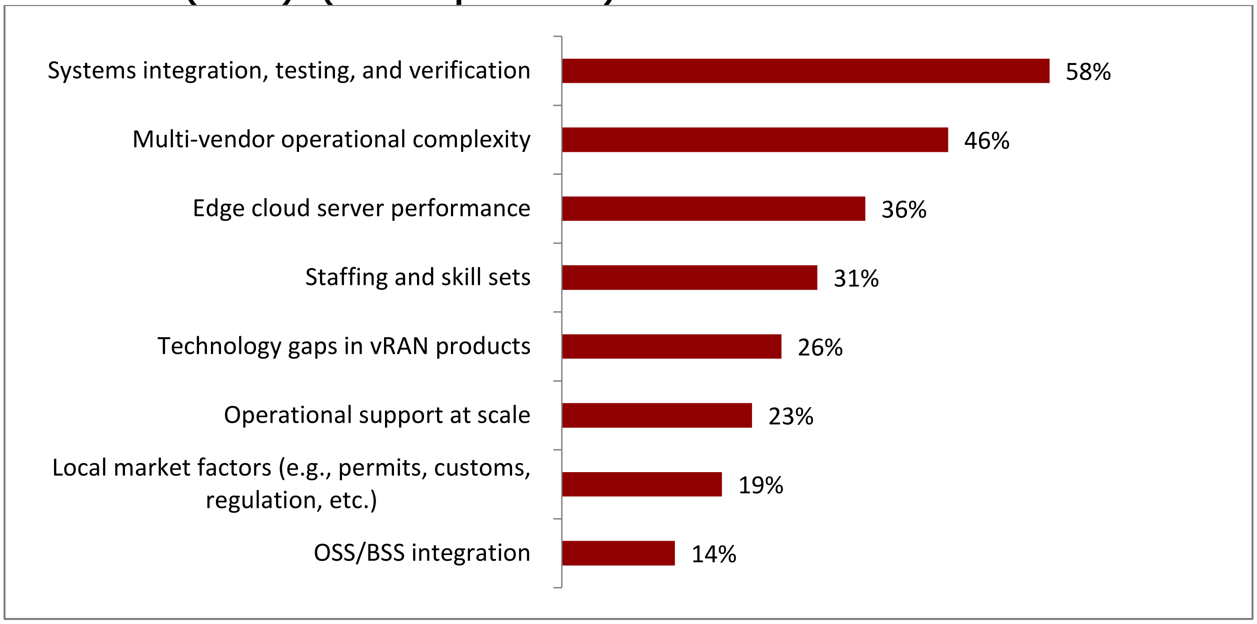

N=140
Source: Heavy Reading

Picking up the point on edge cloud hardware performance, the next question asked about x86 servers and the importance, or otherwise, of accelerators. Hardware accelerators are used to improve Layer 1 performance in vRAN because x86 processors do not currently match dedicated baseband processing hardware. This is particularly acute where higher order multiple input, multiple output (MIMO) and massive MIMO configurations are used (as they increasingly are). Using accelerators, which are typically field-programmable gate arrays (FPGAs) integrated on the board or in network interface cards (NICs), is common in wireline network functions virtualization (NFV). Thus, it stands to reason that the technology would be useful in radio processing, which is, generally speaking, more demanding. The survey results bear out this statement. As shown in **Figure 8**, a large majority think hardware accelerators are "critical" or "important."

Heavy Reading split the question between two categories:

- "High capacity sites," which process more traffic and users and tend to use carrier aggregation, higher order 4x4 MIMO, 8T8R, and massive MIMO.
- "All sites," which would include suburban and rural and may be single carrier sites using 2x2 MIMO at the lower end.

There is a clear difference between the 43% that believe accelerators are critical for high capacity sites and the 27% that think accelerators are critical in general. This will be an interesting area to track in the coming years as new general-purpose silicon becomes available that is better suited to Layer 1 RAN processing. This includes x86- as well as graphics processing unit (GPU)- and ARM-based processor designs. Conceivably, in the future, general-purpose silicon will ship with embedded software libraries that mean hardware accelerators will not be needed.

Figure 8: How important are hardware accelerators in x86 servers for vRAN applications?

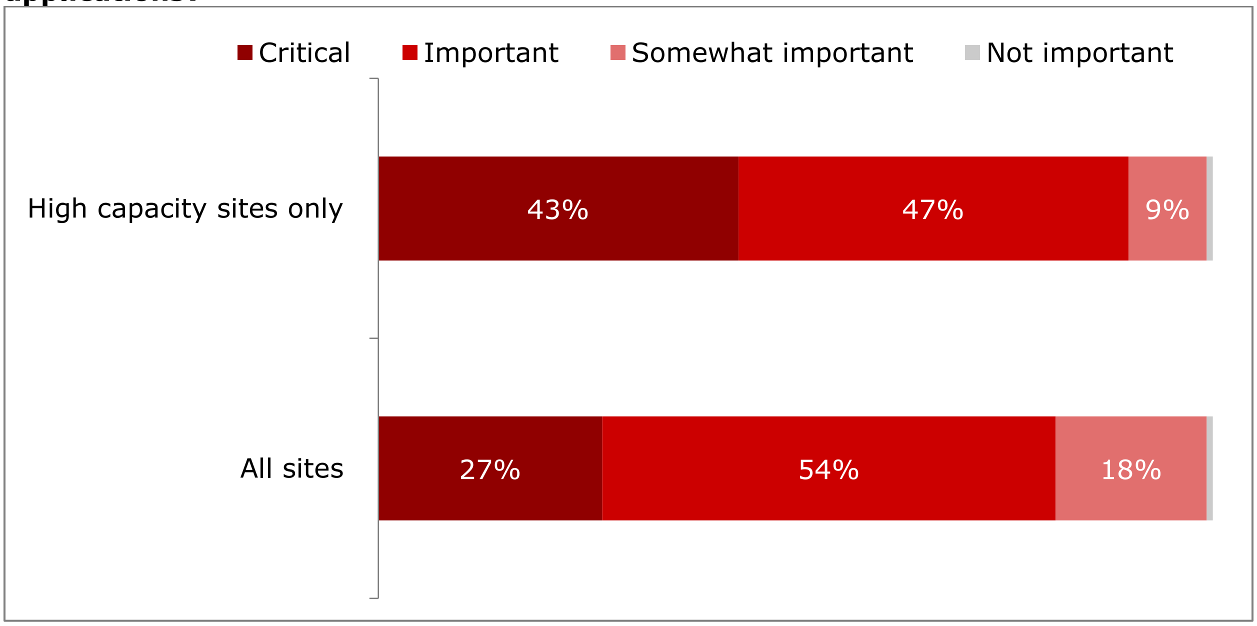

N=138
Source: Heavy Reading

The vRAN systems on the market today support LTE, typically single carrier with 2x2 MIMO, with the more advanced products able to support 4x4 MIMO. The state-of-the-art in classic, vendor-integrated RAN is significantly more advanced, and it is imperative for vRAN that this gap is closed significantly and rapidly. Open vRAN vendors do not need exact feature parity with integrated systems, but they do need to be competitive in the next couple of years. One example of this is support for 5G vRAN.

Asked if it is practical to virtualize a 5G baseband, a.k.a. DU, function for commercial deployment within 2 years, a majority (60%) said yes. Of these respondents, 31% said "yes – with feature and performance parity" and 29% said "yes – but with feature or performance impacts." This positive view on the part of the survey base is shown in **Figure 9**. Heavy Reading believes feature and performance parity will be challenging in this time period and tends toward the view that virtual 5G baseband will be achieved, but with compromises relative to vendor-integrated state-of-the-art.

Figure 9: Do you think it is practical to virtualize a 5G baseband DU function for commercial deployment within 2 years?

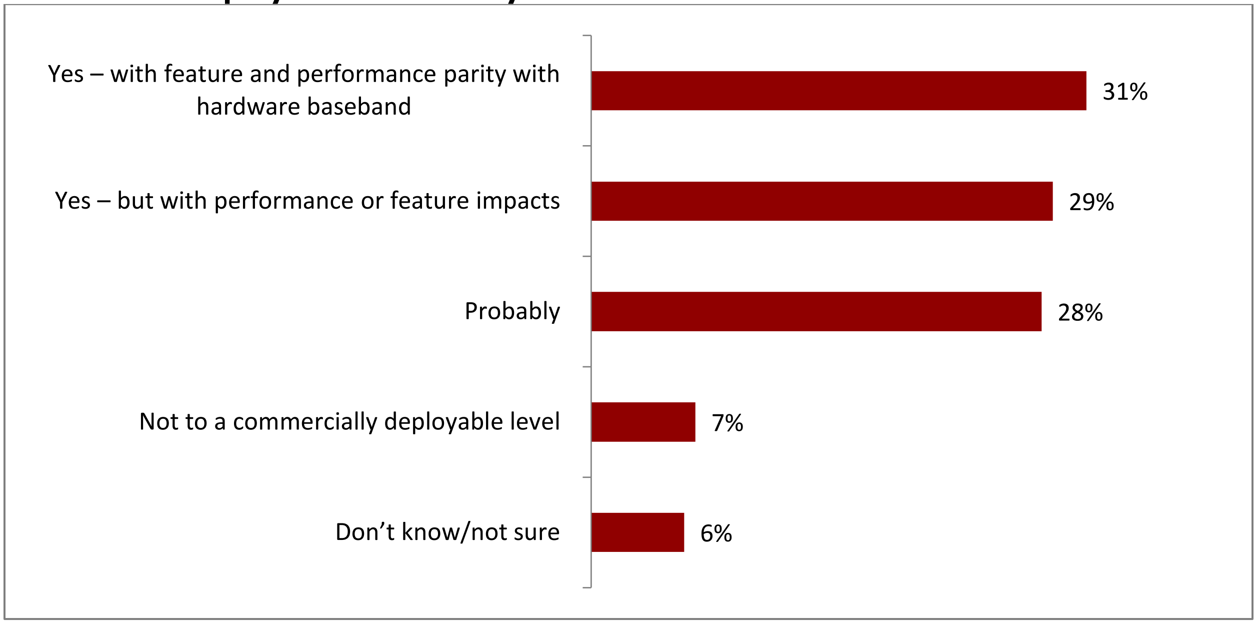

N=139
Source: Heavy Reading

Open vRAN does not negate the value of vendor-integrated, state-of-the-art RAN. The RAN market is large and growing, and operators have many diverse requirements; there is clearly an important role for both. Asked if and how vendor-integrated base stations might have advantages over vRAN, respondents led with the "ability to meet deterministic low latency performance requirements," with 59% considering this a "major advantage," as shown in **Figure 10**.

Figure 10: How advantageous are vendor-integrated 5G base stations compared to a disaggregated virtual RAN solution using COTS hardware and open source?

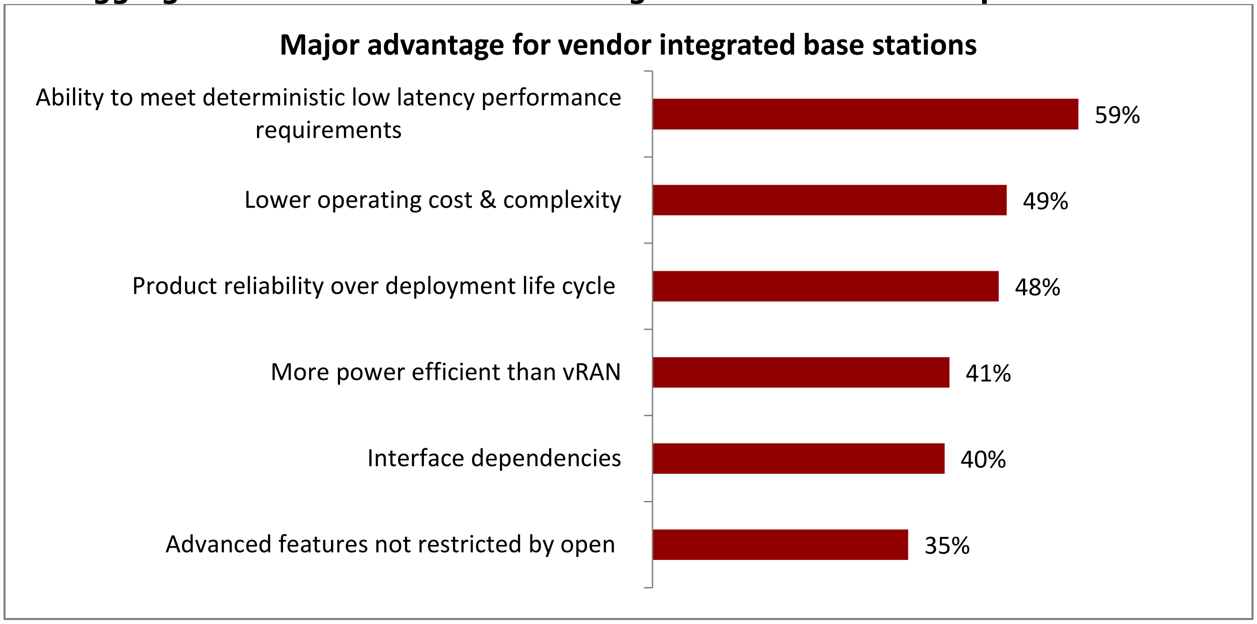

N=136
Source: Heavy Reading

ENTERPRISE AND INDUSTRIAL 5G EDGE CLOUD

The 5G system is designed, in part, for enterprise customers. The intent is to apply 5G technical capabilities across diverse industries. This applies, in particular, where there is a need for ultra-reliable low latency communications (URLLC), which cannot be supported over LTE. Examples are cyberphysical systems with deterministic performance requirements, such as industrial robotics or remote control of machinery. There are also other demanding applications that could be served with existing wireless technologies but may be challenged at scale. Examples include automated guided vehicles (AGVs), UHD video, machine vision, VR/AR data goggles, and so on.

In many cases, to meet URLLC performance targets, the 5G mobile network will have to be deployed near or at the same physical location as the application itself; in other words, on an edge cloud infrastructure. Moreover, because many of these applications generate commercially sensitive information, for example, related to production figures or intellectual property, enterprises may not want data to leave their premises on the public network. In these cases, an edge cloud with local breakout is again a requirement.

This section of the paper discusses enterprise and industrial users' perspectives on 5G by analyzing the responses to three critical questions. **Figure 11** shows how respondents see 5G for industrial use cases relative to LTE and Wi-Fi. Clearly, the perception is that 5G is superior for industrial applications, with only 12% and 2%, respectively, saying that LTE and Wi-Fi will be good enough in future. There is then, on this evidence, a strong appetite for newer, better wireless technologies. And just over a quarter (28%) appear sold on 5G already, saying "performance advantages will force its use" where wireless is needed in industrial networks. The majority (58%), however, say that the claimed, theoretical advantages of 5G still "need to be proven in practice." They are favorable toward the technology but are not yet convinced. This accords with the proof-of-concept (PoC) activity and trials large industrial companies have been running as they seek to integrate 5G into their production environments.

Figure 11: Which of the following statements best characterizes your view on 5G for industrial use cases, relative to LTE and Wi-Fi?

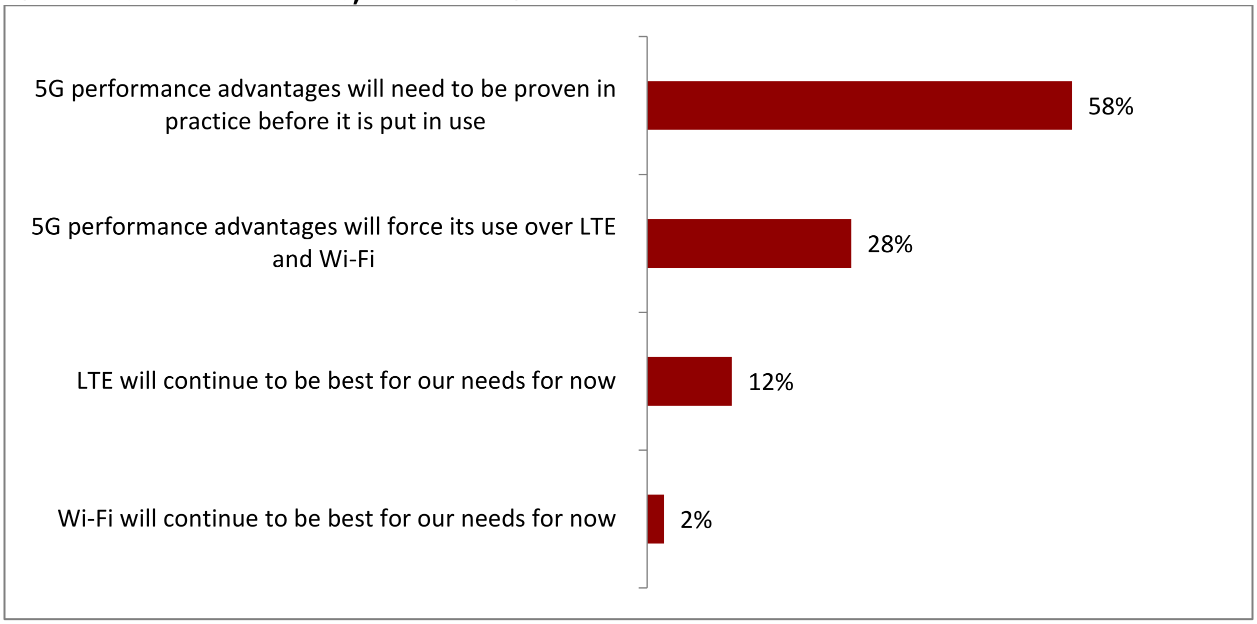

N=43
Source: Heavy Reading

To determine the view that 5G needs to be deployed on-premises to meet demanding performance targets, Heavy Reading asked respondents how likely they are to deploy a private 5G network over the next 3 years (**Figure 12**). The 3-year timeframe was selected because 3GPP Release 16 is due for standardization in mid-2020 and contains many important enhancements for industrial and vertical users. The start of 2023 will be enough time for them to be integrated into vendor solutions, but is nevertheless ambitious.

A quarter of respondents (24%) said it is "very likely" they will deploy a private 5G network; this fits with the cohort identified in a prior question that are convinced about the performance advantages of 5G. A further 38%, the largest respondent group, are more cautious and believe it "likely" they "could execute a deployment." Presumably, they are waiting for proof the technology delivers its promises before they commit. Overall, this is a positive response for private 5G networks. If anything, it might be a little too bullish.

Figure 12: Within the next 3 years, how likely is your company to deploy a private 5G network on-premises to support industrial applications?

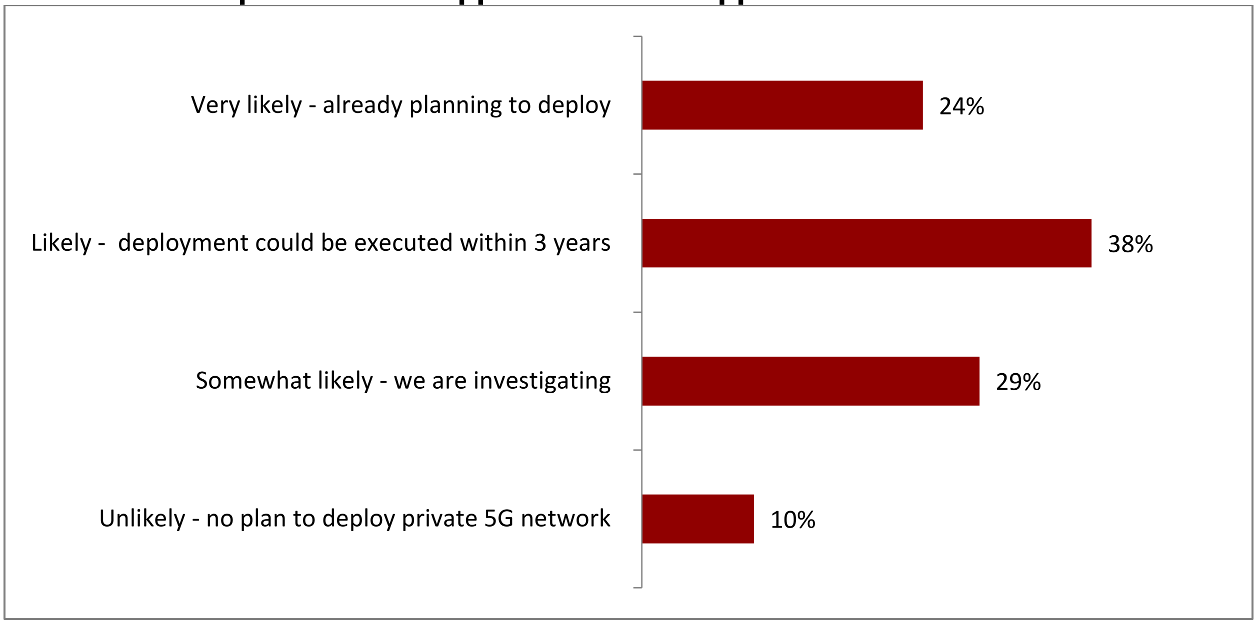

N=43
Source: Heavy Reading

Analytics and monitoring of edge cloud infrastructure are critical to day-to-day operations. High performance industrial applications need high performance resilient infrastructure. When asked about the importance of various types of embedded edge cloud monitoring, **Figure 13** shows that respondents believe embedded monitoring is more important to infrastructure than it is to the applications themselves or to deriving business insights. The most important use of embedded monitoring is for "security" of the edge cloud infrastructure, which is identified as "extremely important" by a majority (56%) of respondents. "Performance" comes second (44%).

The relatively lower scores for "business insight" (33% "extremely important") and "application analytics" (19% "extremely important") imply that respondents will have alternative, dedicated tools at the application and business insight layers. These are still important, however, and the response does not preclude that edge infrastructure monitoring could feed data into those tools and contribute to business insight. But it does imply that the focus for edge vendors and edge implementers should be based on secure performance above all else.

Figure 13: How important is having integrated analytics and monitoring of the edge cloud infrastructure in your industrial 5G network for the following areas?

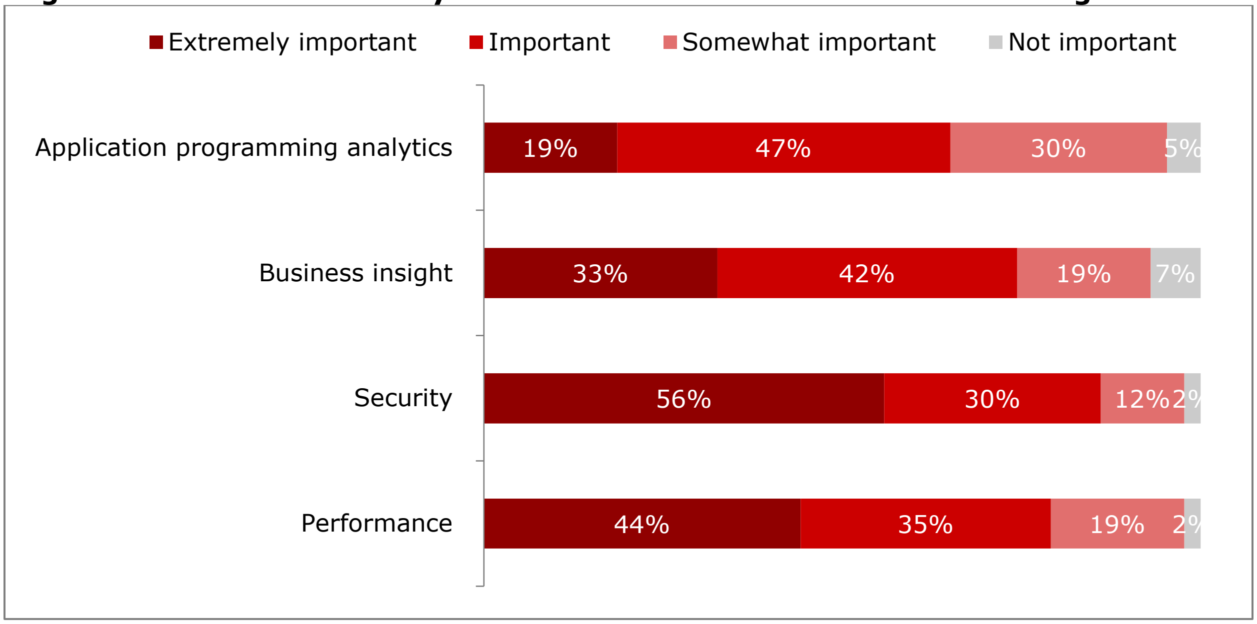

N=43
Source: Heavy Reading

AUTOMOTIVE SECTOR AND 5G

The automotive sector is going through a period of rapid and deep change, with the connected vehicle trend (telematics, navigation, infotainment, etc.) occurring at the same time as advances in driver assistance/automation and mass-scale electrification. Mobile networks, in the broadest sense, are important enablers for each of these. The 3GPP has a dedicated 5G development track for automotive and elsewhere. There is also activity in the IEEE on the Dedicated Short-Range Communication (DSRC) standards group, as well as in the satellite sector.

The 3GPP C-V2X initiatives have emerged as the leading solutions for connected vehicles. This work started in LTE and is now gaining pace in 5G. A key part of the C-V2X concept is the RSU, which is a highly distributed edge cloud node from a hardware and cloud infrastructure software perspective. In time, as deployment gains pace, RSUs are likely to see significantly denser deployments than current RANs, making it an important market for "far edge" infrastructure platforms. There are also several potential synergies with mobile access; for example, Heavy Reading expects to see RSUs with integrated small cell access.

To gain insight into how the automotive sector views 5G, Heavy Reading surveyed respondents working in the sector. The questionnaire garnered a total of 56 respondents, of which over half (62%) were based in the U.S.

In terms of which C-V2X applications will lead deployment, **Figure 14** shows vehicle-to-network is the clear leader with 56% of the response. This refers to vehicles communicating with cloud services for information like real-time traffic routing and is the simplest V2X service. In second place, with 36%, is vehicle-to-infrastructure, which refers to communications with RSUs for services like traffic signal timing and priority. V2V scores almost as highly at 35% and refers to direct communications without necessarily needing an RSU and is targeted at collision avoidance and related safety applications. In 5G C-V2X, V2V communications depend on what are known as wireless "side links," which are being specificed in 3GPP Release 16 and developed further in Release 17. They may be present in commercial products from 2022 onward. Vehicle-to-pedestrian has high potential to improve road safety, but requires significant technology and policy development to be reliable. It is no surprise that it scores lowest at 9%.

Figure 14: Which cellular V2X applications will be deployed first? (Select two)

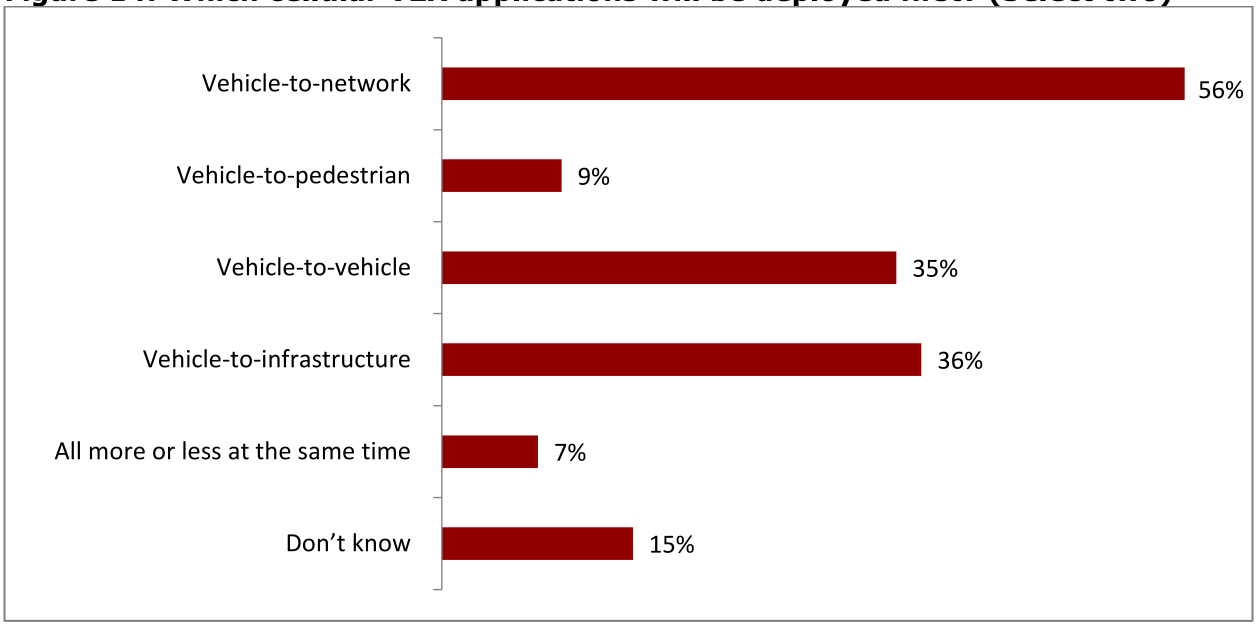

N=55

Source: Heavy Reading

One of the biggest sources of hyperbole in the early years of 5G development was equating autonomous vehicles with 5G. The intent was to show how 5G could help advance connected cars and autonomous driving. However, these statements were often extrapolated out of context and taken to mean that autonomous driving could not happen without 5G. This is incorrect because autonomous is developing without 5G and because it is, and always was, unrealistic to assume that 5G coverage would exist everywhere an autonomous vehicle might travel.

When Heaving Reading asked respondents in the automotive sector, as shown in **Figure 15**, if they believe 5G is needed for L4 autonomous driving, the results are interesting. The largest group (36%) says it is a "requirement." This might indicate a large percentage of the respondent base has bought the hype, although this is mitigated by the qualifier in the answer that a high performance network "can help coordinate vehicles." The other interpretation is that a large majority (64%) do not think 5G is required.

Another interesting issue is the concept of compute offload from vehicles to edge cloud infrastructure. Today's autonomous driving systems require large amounts of expensive compute to be deployed in the vehicle. If some of this could be offloaded to the edge cloud network infrastructure, the cost per vehicle cold be reduced, so the argument goes. This is interesting and a point for further study. Working against it is the view that on-vehicle compute will become far more efficient as algorithms are refined and embedded in dedicated silicon. There is also the question of whether vehicles can self-drive when a high bandwidth, low latency, ultra-reliable network connection is not present.

Perhaps more closely aligned with Heavy Reading's view are the 27% that say a "network connection is needed in practice" and the 22% that say L4 autonomous driving would "benefit from network assist even if it is not necessary." These are different, but not a million miles apart, and they suggest an important role for network connectivity.

Figure 15: Do you believe 5G is needed for Level 4 (L4) autonomous vehicle operation?

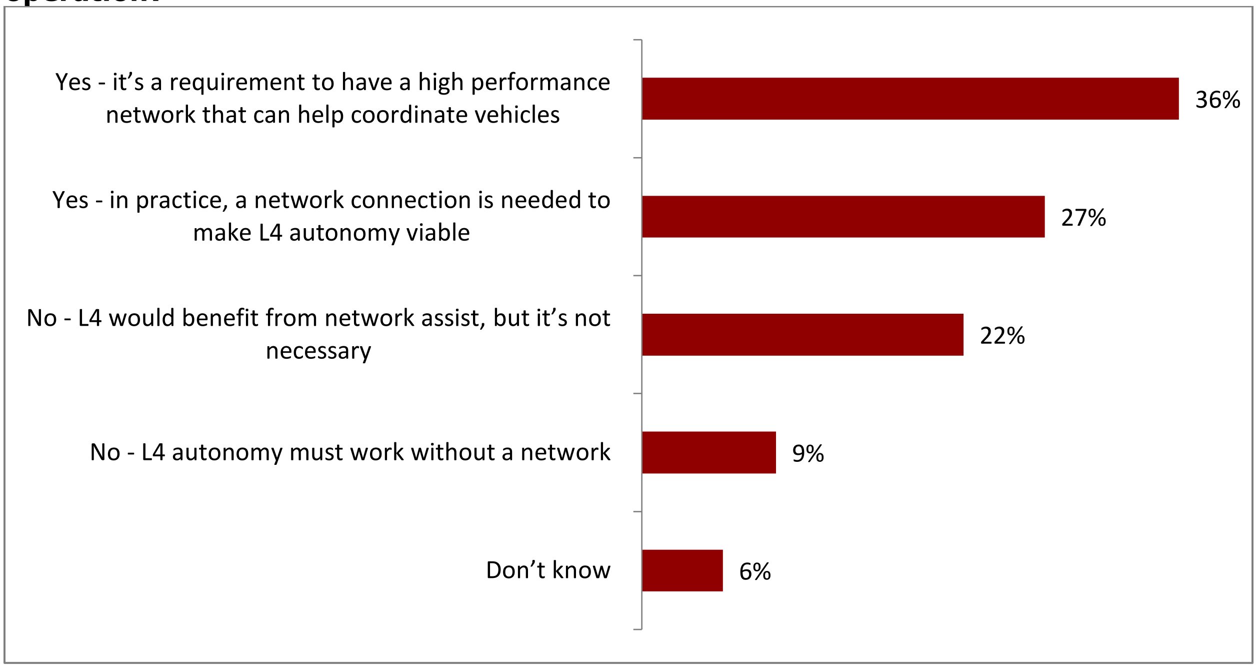

N=55
Source: Heavy Reading

As noted above, mobile network coverage is generally good, but it is far from ubiquitous. In geographically large markets, in particular, it is hard to cover everywhere with terrestrial systems. There are now several non-terrestrial systems in development for wide-area broadband access, with low orbit satellite systems showing promise. At face value, a two-way satellite could be a good solution for vehicle connectivity as an adjunct to cellular and for use cases where low latency is not vital.

Figure 16 shows a third of respondents (34%) think this is a good idea and plan to integrate satellite into their 5G automotive offering. What is also interesting is how broad the response to this question is, with equal shares for "no plans" (35%) and "don't know" (31%). The "no" camp might be explained by factors like the additional cost and business relationships satellite would require, by the fact that satellite for two-way mobile vehicle communications is less mature than 5G for the mass market, and by the lack of support for low latency services. The "don't knows" speak for themselves.

It may be worth noting that Tesla's Elon Musk recently weighed in on 5G on Twitter by claiming the mobile industry is getting "a bit too greedy in its spectrum grab." No doubt he was referring to the U.S. C-Band debate and thinking about his own satellite investments in the first instance. However, it does not take a leap of the imagination to see a link to connected cars.

Figure 16: Are you planning to integrate satellite connectivity into your 5G automotive offering?

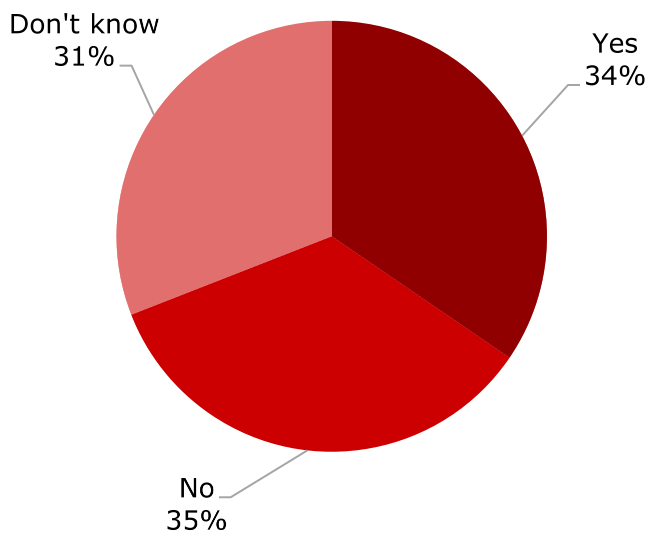

N=53
Source: Heavy Reading

ABOUT WIND RIVER

A global leader in delivering software for intelligent connected systems, Wind River offers a comprehensive, edge-to-cloud software portfolio designed to address the challenges and opportunities communications service providers face when evolving and modernizing their networks.

Wind River is focused on accelerating the massive innovation and disruption at the network edge by being active collaborators and contributors to the StarlingX project. As a commercial deployment of StarlingX, the Wind River Titanium Cloud platform will be key to enabling new business opportunities and innovative applications across multiple market segments.

Visit www.windriver.com for more information.