

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

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

Edge Computing in Telco Networks: Gaining the Competitive Edge

*A Heavy Reading Market Leader Survey report produced for
F5, Schneider Electric, and Wind River*



Schneider
Electric



AUTHOR: JENNIFER P. CLARK, PRINCIPAL ANALYST, HEAVY READING

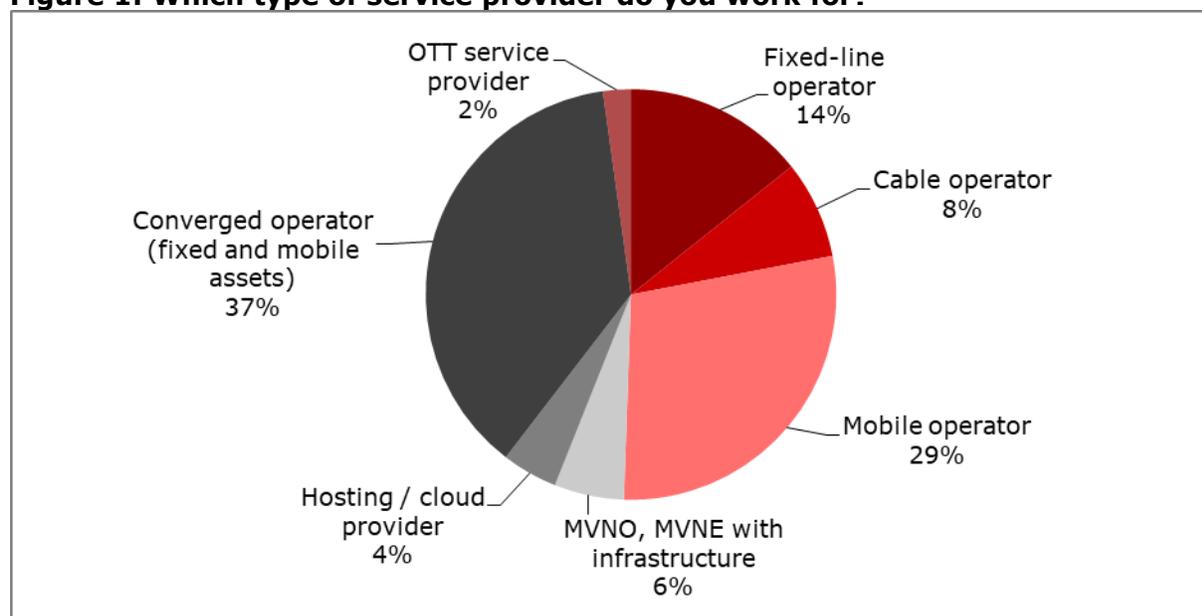
INTRODUCTION

Heavy Reading has been conducting surveys focused on edge computing or with an edge computing component for more than five years. Past surveys show uncertainty regarding what edge computing is and more uncertainty as to how it would or could be leveraged by the telcos. Heavy Reading's current **Edge Computing Survey** demonstrates that service providers (SPs) are deploying edge computing solutions now or are in the planning stages. The direction has been chosen and SPs are (moderately) confident of their plans. However, it is evident from the responses to this survey that although SPs understand the high level drivers of edge deployments, questions remain regarding the details of execution.

DEMOGRAPHICS

Mobile and converged operators made up the bulk of Heavy Reading's survey respondent pool, accounting for two-thirds of overall responses. An additional 22% came from the fixed-line and cable operator community. The remaining 12% hailed from mobile virtual network operators (MVNOs)/mobile virtual network enablers (MVNEs) with infrastructure (6%), hosting and cloud providers (4%), and over-the-top (OTT) providers (2%).

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

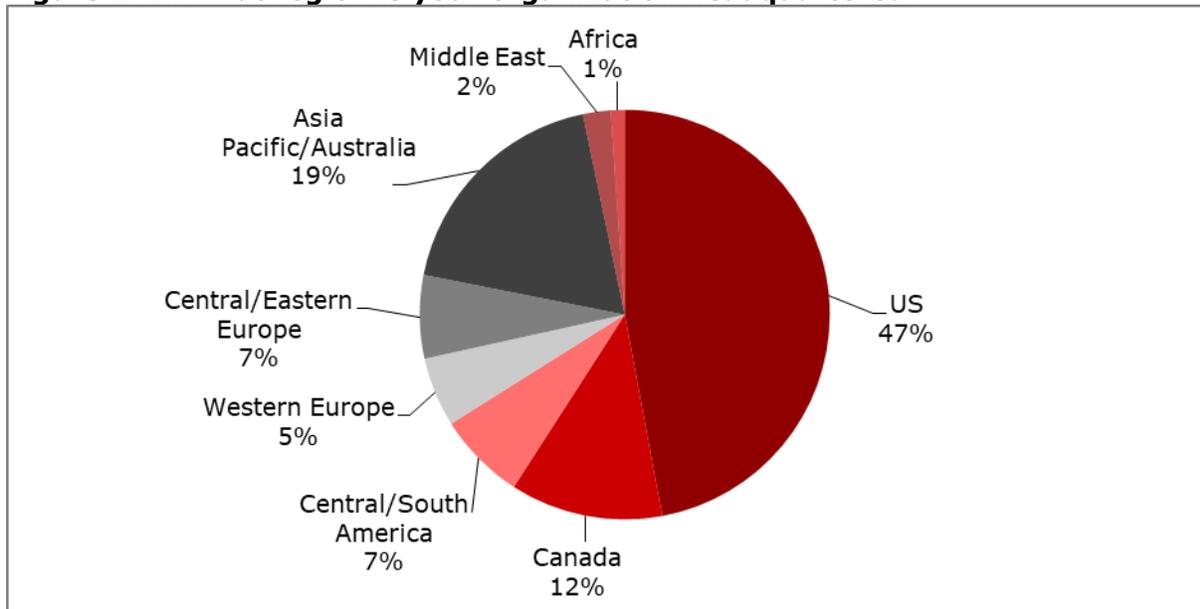


n=91

Source: Heavy Reading

The US accounted for just under half of the respondents, as shown in **Figure 2** below. The Asia Pacific/Australia region provided the second largest block, with close to one-fifth of overall responses. All of Europe, the Middle East, and Africa together accounted for 15%. Canada and Central/South America also had respectable representation at 12% and 7%, respectively.

Figure 2: In what region is your organization headquartered?

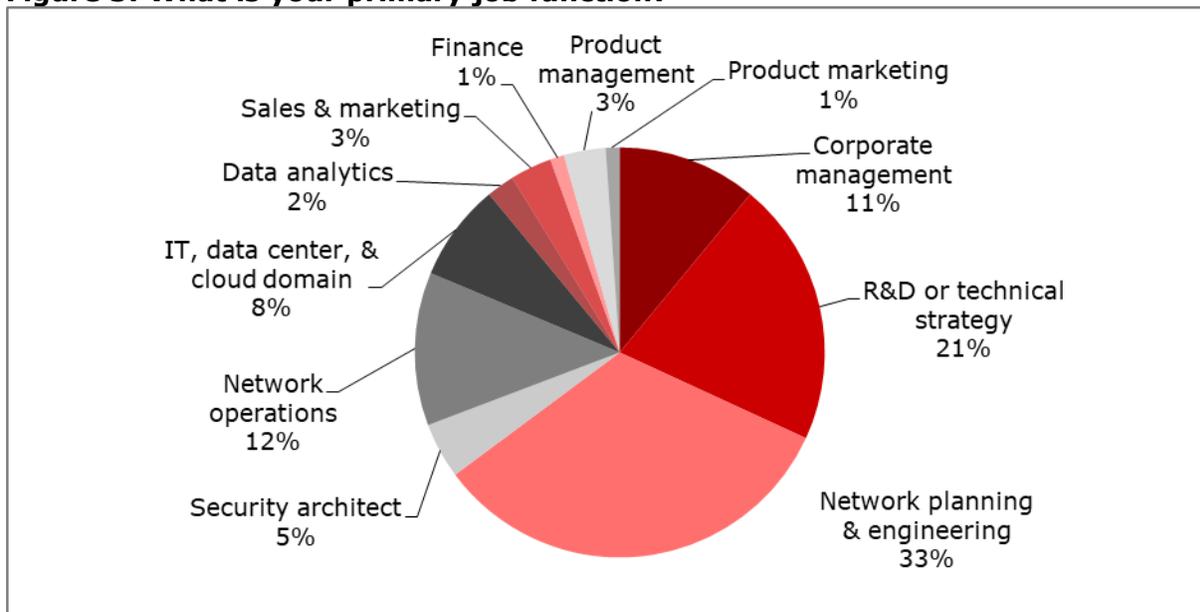


n=91

Source: Heavy Reading

As is the case with most Heavy Reading surveys, two-thirds of the respondents were from technical networking roles: planning and engineering, R&D, and network operations. Almost a fifth were from management, marketing, and finance. The data center is more heavily represented in recent surveys; 10% of respondents were from IT and data analytics in this survey. Security architects, at 5% of respondents, made up the smallest segment.

Figure 3: What is your primary job function?

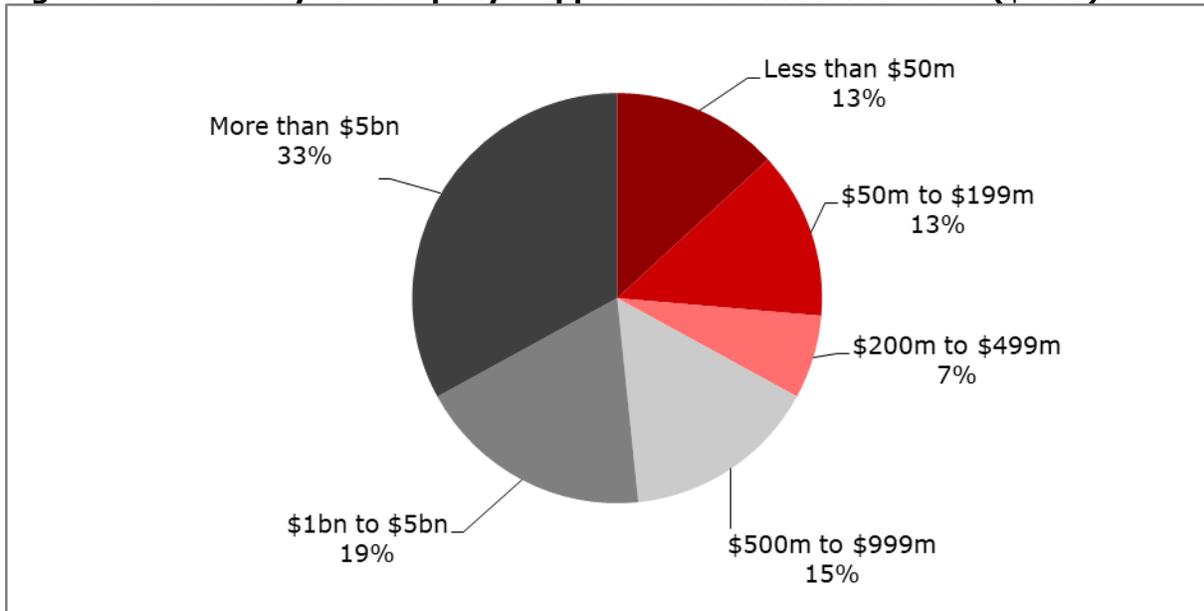


n=91

Source: Heavy Reading

A third of respondents reported their companies have revenues of over \$5bn. Another third have revenues of between half a billion and \$5bn. The remaining third have revenues under half a billion. These smaller SPs include regional business units of larger carriers (e.g., Telefónica), dominant carriers of smaller countries, and OTT vendors, hosting vendors, and MVNOs/MVNEs.

Figure 4: What are your company's approximate annual revenues (\$USD)?



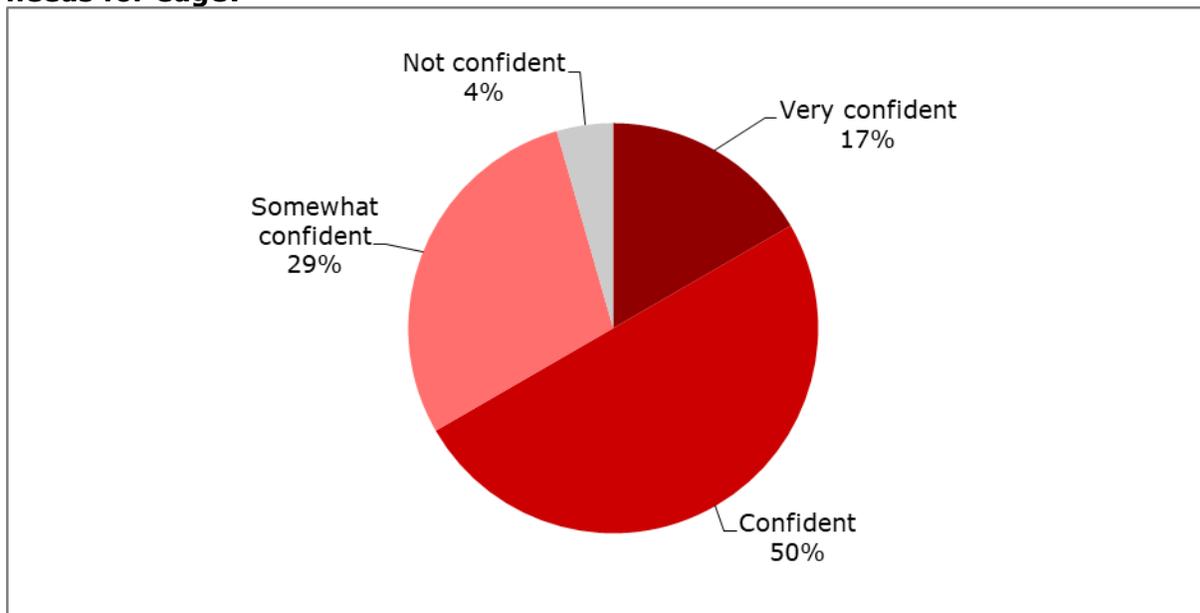
n=91

Source: Heavy Reading

EDGE COMPUTING BIG PICTURE QUESTIONS

Heavy Reading's survey respondents have spent some time planning their edge strategy, and many are sure that they understand where their organization is headed. Half of respondents reported they are "confident" that they have identified their need for edge and an additional 17% are even more assured with their "very confident" response. Nevertheless, as shown in **Figure 5** below, a third are still uncertain as to how their organization will leverage the edge, responding that they are only "somewhat confident" or "not confident" at all.

Figure 5: How confident do you feel that you've accurately identified your network needs for edge?



n=90

Source: Heavy Reading

Heavy Reading weighted the responses to the question for **Figure 6** below to give first choices the most importance and third choices the least importance. The resulting scores read as more of a timeline of implementation than an assessment of the impact of the technology. Internet of Things (IoT) and the connected (not autonomous) car are certainly having an impact today, as seen in the following research and forecast data from Omdia:

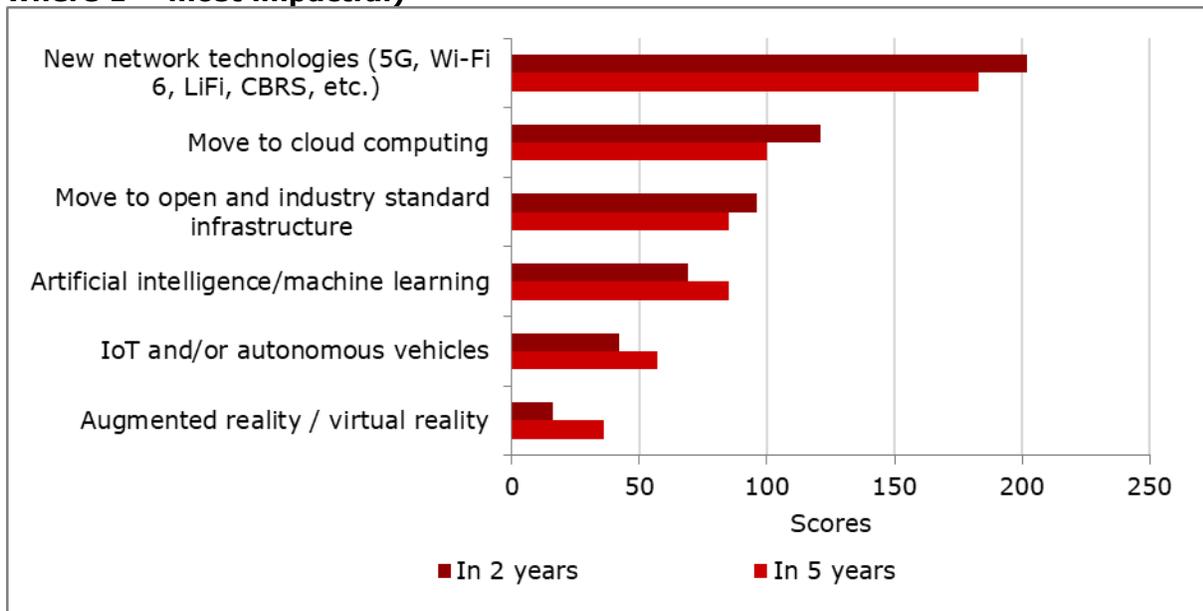
- The number of IoT contracts awarded to carriers grew from under 10 worldwide in 2009 to 1,168 in 1Q20.
- The number of connected cars is expected to grow, worldwide, from 59 million in 2016 to 308 million in 2022.

The responses move, as the carriers themselves tend to move, from the infrastructure, through supporting network technologies, and finally to the application enablement technologies. If a key aspect of edge computing is to support a “data first” network, then the telcos and their partners/vendors face a challenge in shifting their mindset to a more user/data/application-oriented mantra.

The score for “open and industry standard infrastructure” is also quite high at a time when network functions virtualization infrastructure (NFVi) is in flux. The need for application and workload acceleration is encouraging bespoke white box-plus solutions with field-programmable gate arrays (FPGAs) onboard to siphon off compute-intensive processes, relieving the CPU cores.

The transition to private LTE and 5G networks is a daunting task, but it is a task (or set of hundreds of tasks) with which the telcos are very familiar. The broader challenge for the telcos continues to be the struggle to excel in strategic tasks that are aligned more closely with end-user/enterprise data and applications.

Figure 6: Which technologies will have the most impact on the evolution of telco industry in 2 years? In 5 years? (Please rank in order your top three choices, where 1 = most impactful)



n=91

Source: Heavy Reading

What changes five years out? As shown in **Figure 6** above, the only line item that changes from 2022 is artificial intelligence/machine learning (AI/ML), which trades places with open industry standards.

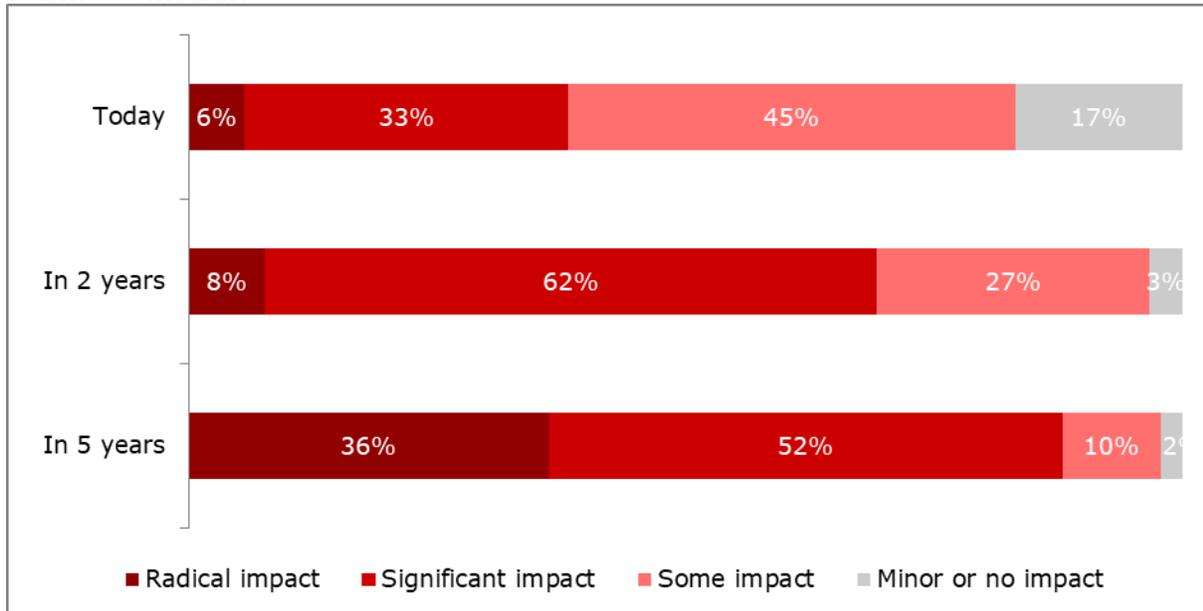
5G is rolling out faster than any previous mobile technology. Omdia forecasts show worldwide 5G subscriptions will be 4x greater than 4G was at the five-year milestone. The number of physical cell sites (i.e., not including multi-sectored antennas) will grow by an order of magnitude as a result of the rollout of 5G. With 5G, carrier ROI is further challenged as unlimited data plans become the norm due to both competition and the reduced cost, compared to 4G, of carrying data.

How will telcos manage their costs—both opex and capex—and support the flood of new cell sites while under increasing competitive pressures? The move to cloud-native, AI/ML, and open standards is key. When it comes to zero-touch provisioning and lights-out maintenance, AI and ML will play a pivotal role.

While they are still at the bottom of the order, IoT and augmented and virtual reality (AR/VR) both garnered higher scores for the “in five years” response (compared to “in two years”). Their implementations will command more respect as their current wins are acknowledged and their future wins attract more recognition.

Virtually the entire survey base believes edge deployments will have an impact on their existing business model in five years. As shown in **Figure 7**, however, the percentage of “radical impact” responses does not change much between today and the next two years.

Figure 7: What level of impact are edge deployments having on your existing business model?



n=91

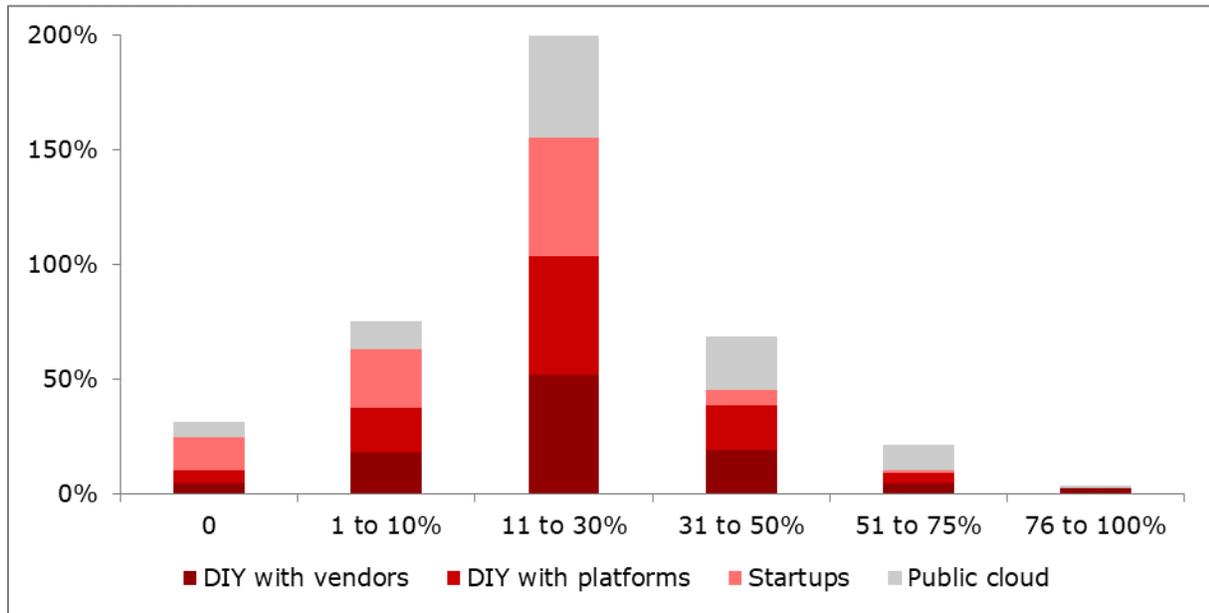
Source: Heavy Reading

Over three-quarters of the survey respondents see edge deployment affecting their business today, and about 70% feel that impact will be “significant” or “radical” in two years. That percentage grows to almost 90% five years out.

This is a very optimistic response, but one that tracks the exuberant rollout of 5G over the same time period. Heavy Reading does not think the response is overly optimistic.

If the results are measured by responses comprising over 10% of edge implementations, the public cloud providers were winners with the greatest percentage of 31–50% of implementations, as depicted in **Figure 8** below. They also constituted the greatest percentage of 51–100% implementations.

Figure 8: What percentage of your edge implementations will be in the following architectures?



n=89

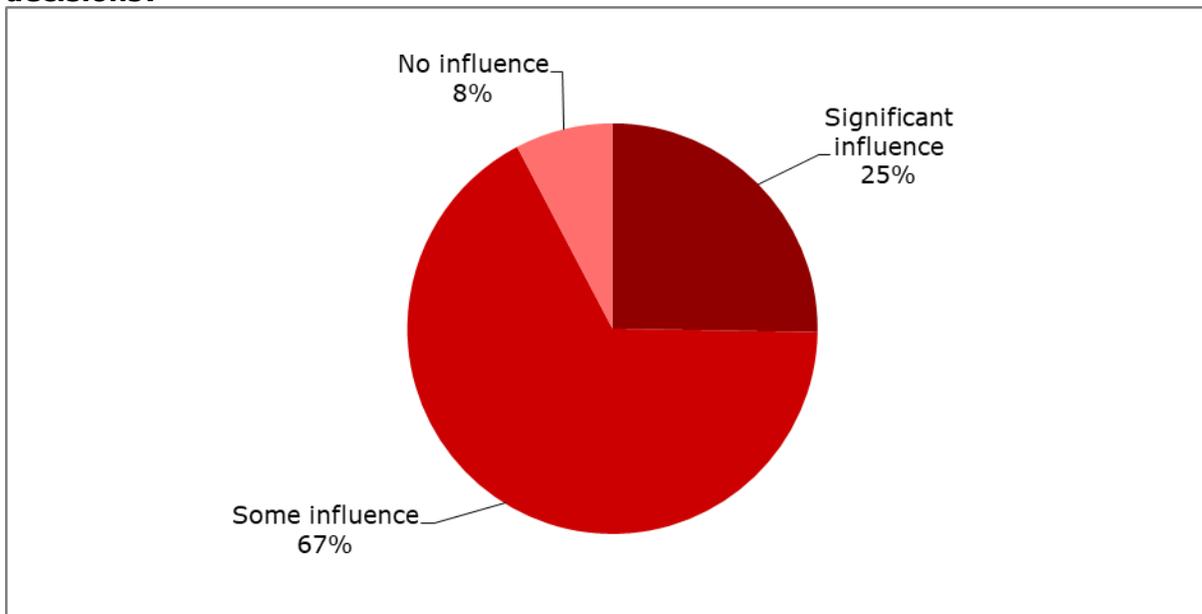
Source: Heavy Reading

Startup companies lagged, with zero responses over 75%. These companies garnered one response of 60% and six in the 31–50% range. The remaining responses were 30% of implementations or below. For a startup, 40% or even 20% of a large implementation can be a boon. More can be a curse, as a startup will struggle to scale development, manufacturing, and service, whether internal or outsourced, and tie 50% or more of its revenue to one or two carrier contracts.

While the carriers are clearly willing to do business with the webscale providers, they will rely mainly on DIY solutions. Heavy Reading offered respondents two flavors of DIY—one driven by standards and the other relying on a vendor assist. But if combined, DIY solutions emerge as the favored solution among the carriers.

As shown in **Figure 9** below, two-thirds of respondents chose the safe response. In other words, new telco/edge cloud platforms will have “some influence,” which translates to “how could they not?” A quarter of respondents admitted that these platforms would have a “significant influence.” A quick cross-tabulation with responses from the previous question shows that these respondents are also expecting to go with a public cloud provider for a high percentage of their sites. “No influence” tracks closely with the DIYers from **Figure 8**.

Figure 9: To what degree will the new telco/edge cloud platforms (i.e., Azure Edge Zones, AWS Wavelength, Google Mobile Edge Cloud) influence your edge purchase decisions?



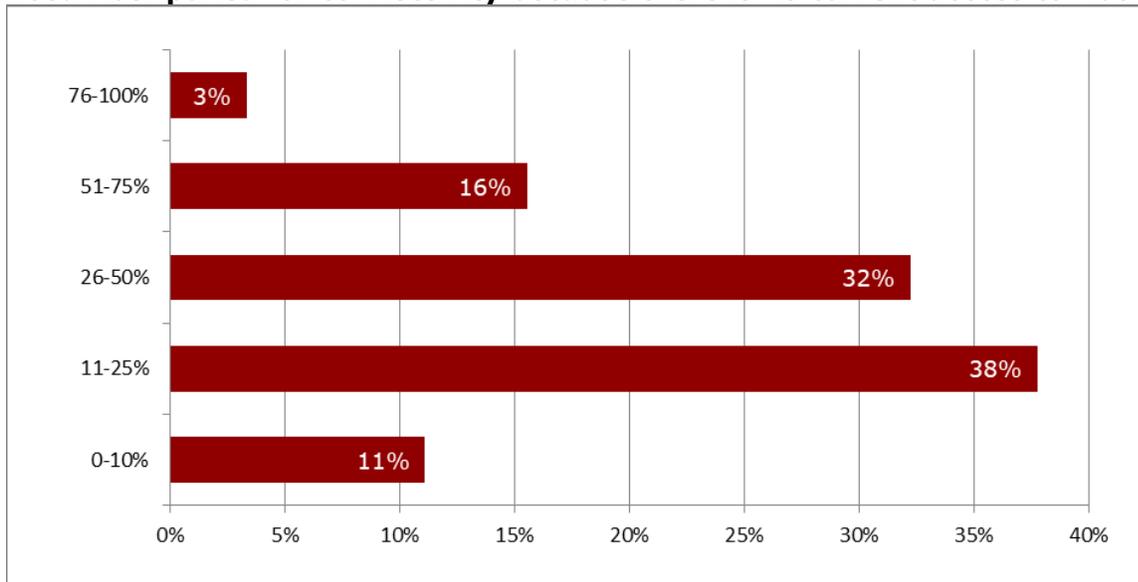
n=91

Source: Heavy Reading

Rolling out edge solutions

Bottom line for **Figure 10** below: Nobody has enough fiber. The 0–10% respondents are a combination of fiber-rich cable operators, Chinese carriers, and respondents who do not build out their own fiber, such as MVNOs and OTT players. Everyone else will be scrambling. The three 76–100% respondents are all carriers with over \$5bn in revenue. The vast majority of responses (86%) fall in between. The fiber needed, for the most part, will not be long lines. They will be fiber spurs and xHaul. The bulk of the fiber windfall is likely to go to dark fiber and xHaul providers, such as cable operators Zayo, Internap, and Windstream.

Figure 10: What percentage of your planned edge computing implementations will need fiber pulled for connectivity because there is no current access to fiber?

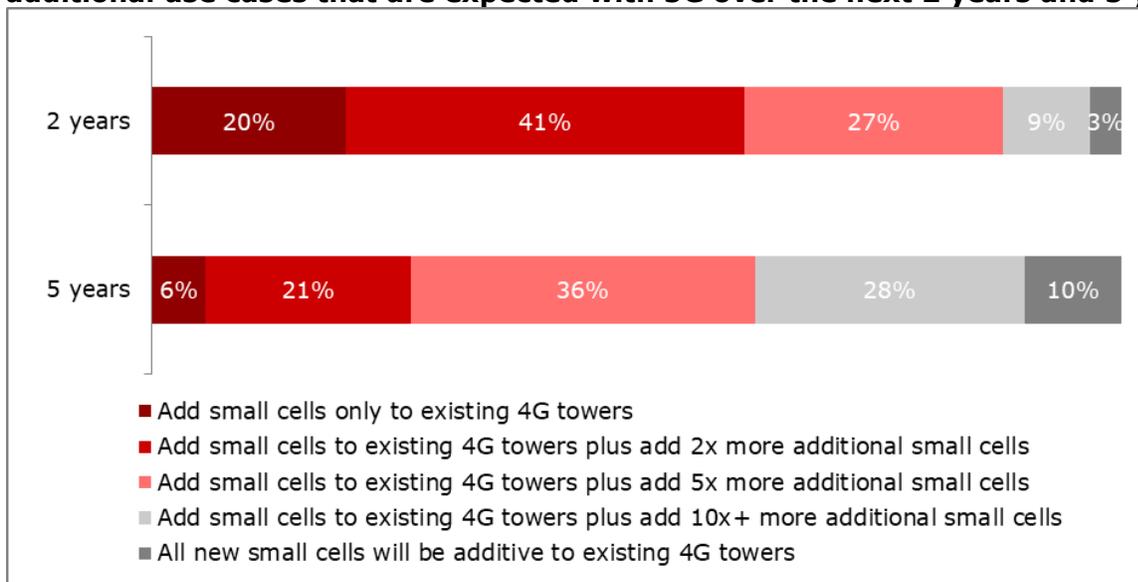


n=90

Source: Heavy Reading

The carriers have a plan. Survey responses shown in **Figure 11** show a glimpse of the methodology that SPs will use to expand their footprint to provide 5G coverage. Over the next 2 years, the majority will double the number of cell sites, with a quarter expecting to grow sites by a factor of 5. Larger carriers (over \$5bn in annual revenue) anticipate aggressive growth of 10x in the next two years.

Figure 11: Which of the following actions will you take in order to support the additional use cases that are expected with 5G over the next 2 years and 5 years?



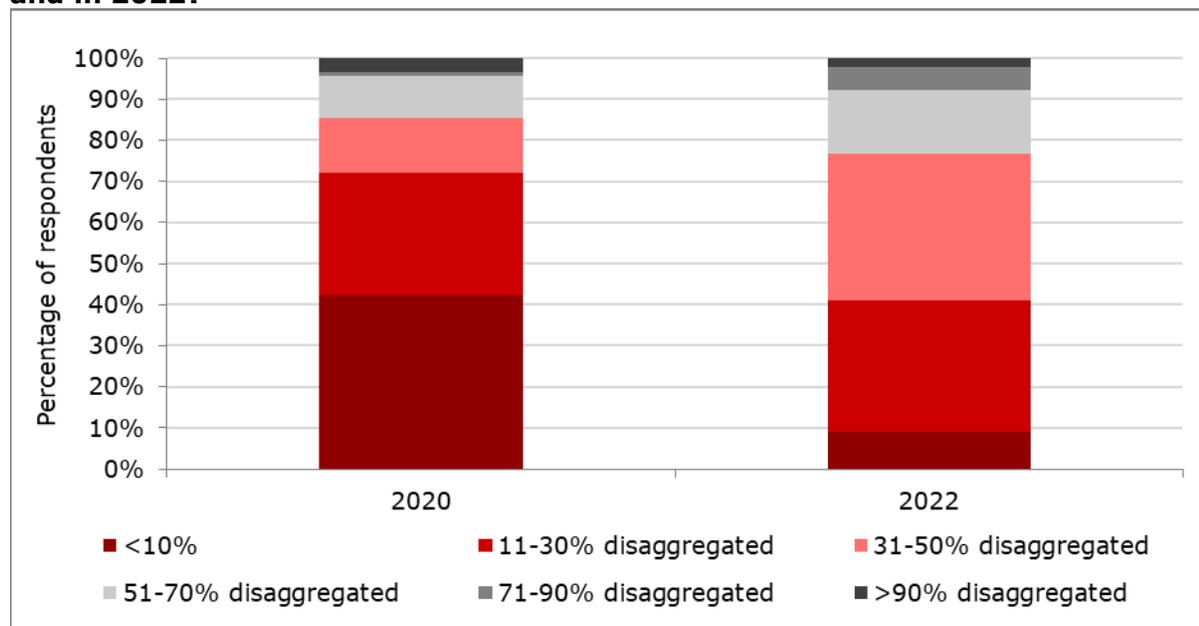
n=90

Source: Heavy Reading

In five years, 75% of respondents expect to grow their number of cell sites between a factor of 5 and 10. This is in keeping with Omdia research that shows global 5G sites (both radio units and distributed sites), excluding China, growing from an estimated 200,000 sites by the end of 2020 to 1,680,000 sites in 2024. China is excluded not only because in-region SPs do not share transport infrastructure numbers, but also because informal estimates suggest China's numbers would more than double the market.

Figure 12 indicates that virtual radio access network (vRAN) is a hot topic for the mobile operators and one that they are moving on today. The use of vRAN is expected to triple in the next two years. This year, over half of respondents will have 10% or more of their 5G network running on a vRAN solution. Two years from now, over half will use vRAN solutions in 30% or more of their 5G network. The percentage using vRAN in over 70% of their 5G deployments will double within the same two-year period. The move to vRAN and a software-driven architecture for 5G will both enable and simplify the move to edge computing.

Figure 12: What percentage of your NEW 5G edge network will run on a disaggregated virtual RAN solution vs. a vendor integrated base station in 2020 and in 2022?

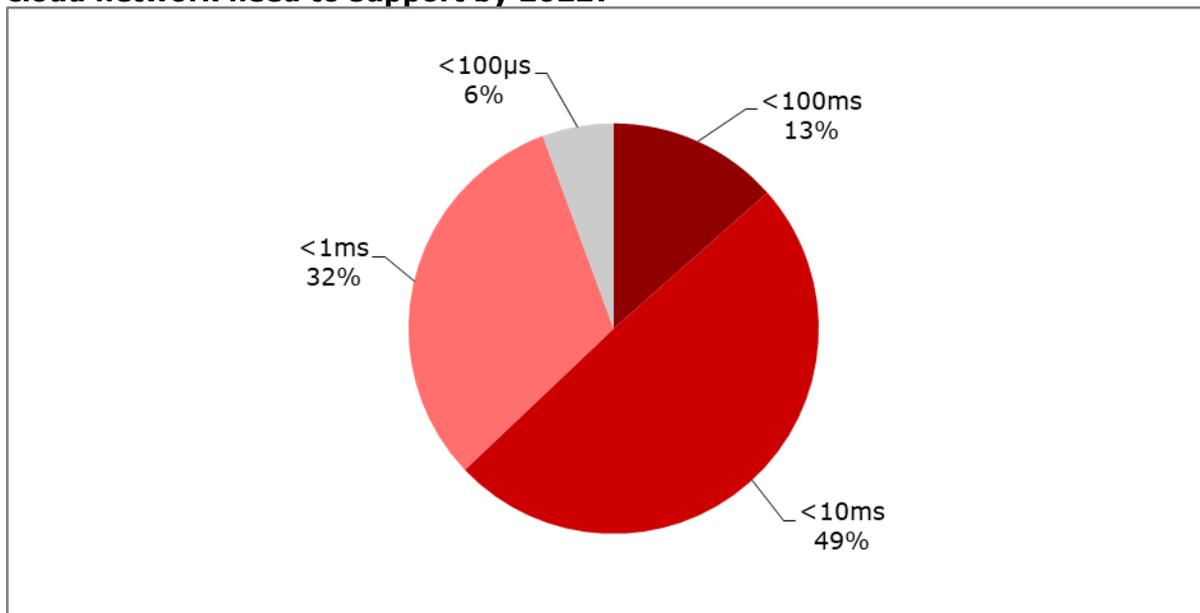


n=90

Source: Heavy Reading

It is fortunate that a larger percentage of respondents (shown in **Figure 13** below) are not looking for sub-100 microsecond response time in the edge-to-core connections. The laws of physics come to play for longer runs, and mobile operators face aggregate network delay (depending on number of hops) and challenges with today's sync technology.

Figure 13: What level of one-way edge-to-core network latency will your edge cloud network need to support by 2022?



n=89

Source: Heavy Reading

However, on the other extreme, not many more respondents will be satisfied with latency approaching 100 milliseconds (ms). As shown in **Figure 13** above, the vast majority, 87%, are pushing for sub-10ms, with 38% pushing for sub-1ms in the next two years. Sub-10ms one-way latency will meet their needs. This is the design attribute that most vendors are building into their solutions, along with the assumption that some low latency applications at the edge may be even more stringent and require sub-1ms on the local short-haul connection.

Edge features and functionality

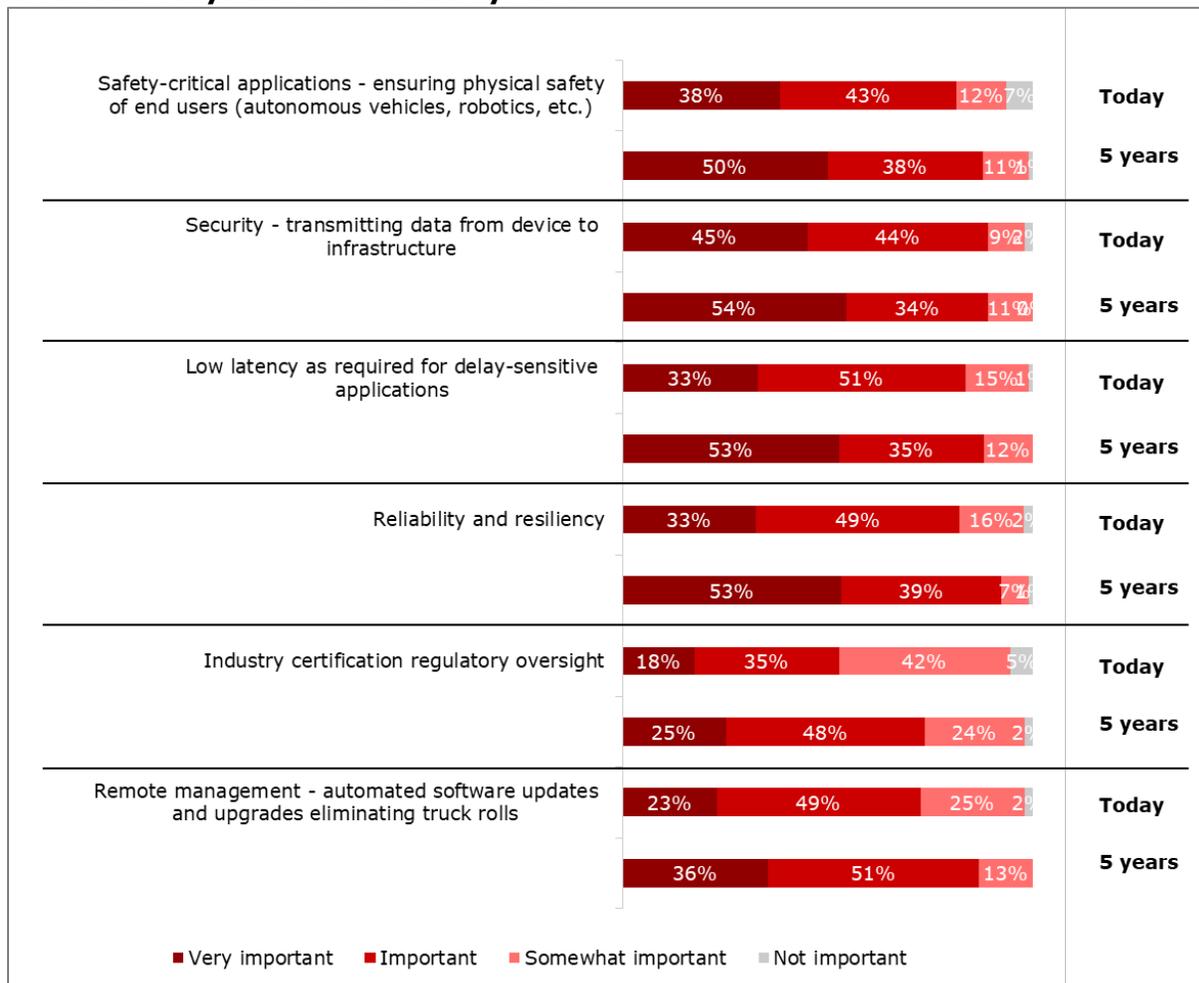
Security—both data en route and security tied to end-user safety—topped out Heavy Reading’s survey, with 89% and 81%, respectively of respondents claiming that they were “very important” or “important” today. However, as shown in **Figure 14** below, “safety-critical applications” also garnered the highest percentage of “not important” votes. This suggests that some respondents did not expect to be dealing with these types of use cases at all. Another possibility is that they did not understand how these use cases were likely to be affected by the network, at least the edge network.

“Low latency,” probably the best publicized aspect of 5G, came in with fewer votes as a “very important” feature but a great deal of support behind it as an “important” feature. “Reliability and resiliency” showed only a shade of difference compared to “low latency” responses.

While “remote management” came in as second to last in terms of “very important” votes, it is second to highest in terms of “important” votes. Carriers cannot afford anything other than a high percentage of zero-touch, lights-out data centers if they are looking at an edge footprint that is 5x–10x that of their current 4G network.

The comparatively modest, yet positive, results for “industry certification,” both today and in five years, sounds like a plea for information. Heavy Reading believes there is an opportunity here for solutions providers to educate the carriers on how edge computing can help them in this area. Carriers are only now becoming aware of how edge computing can assist them their own regulatory issues of lawful intercept and data sovereignty. They need more insights into the way edge can assist them with similar issues tied to other industries. This is a real opportunity for service operators to provide education.

Figure 14: How important are the following intelligent edge features on a 5G network today? In the next five years?



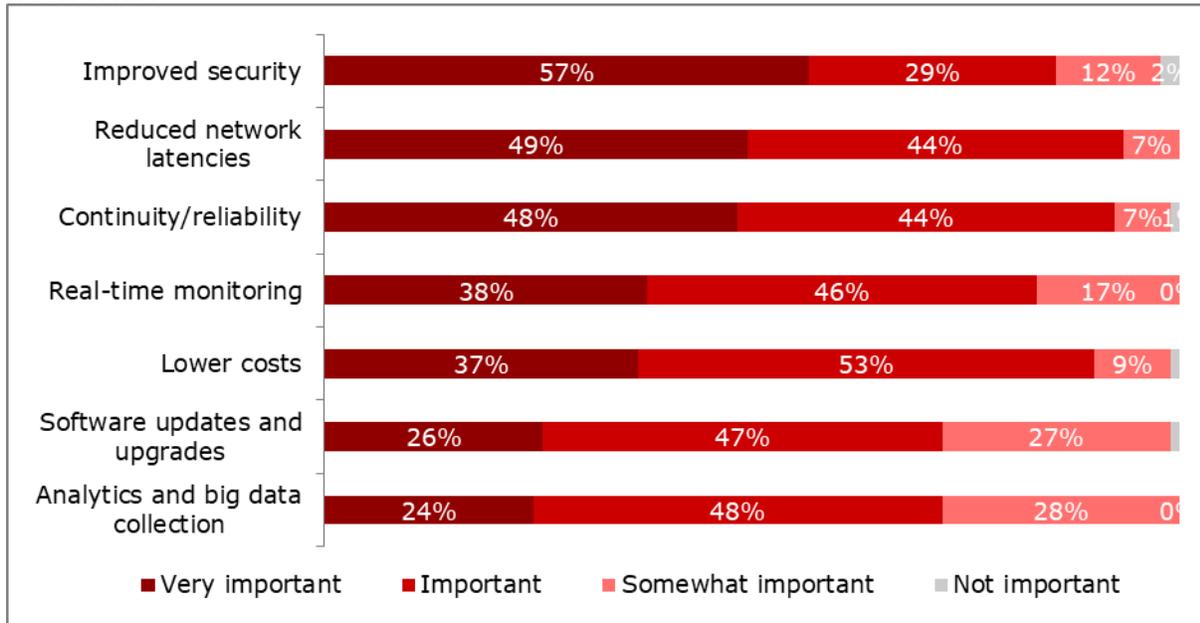
n=91

Source: Heavy Reading

Five years out, everything is important. Not a single issue loses ground in terms of importance. In five years, however, reliability and resiliency top the list, with 92% of respondents claiming that this is “very important” or “important.” Assurance that solutions incorporate carefully thought-out security features appears to be an early hurdle for sales that only becomes more important five years out. Concerns about ongoing reliability and resiliency also pull in issues of service, maintenance, assurance, zero-touch provisioning, ease of configuration, and ease of moves, adds, and changes. These will become the carriers’ driving issues as they scale solutions over the next five years.

Again, improved security tops the list of edge business features in **Figure 15**, this time business features, for the carriers. Improved network latency is also ranked very high by respondents, reflecting the feedback they provided earlier for **Figures 13–14** and the overall importance of low latency capabilities in a 5G network.

Figure 15: How important are the following edge business features to 5G networks?



n=91

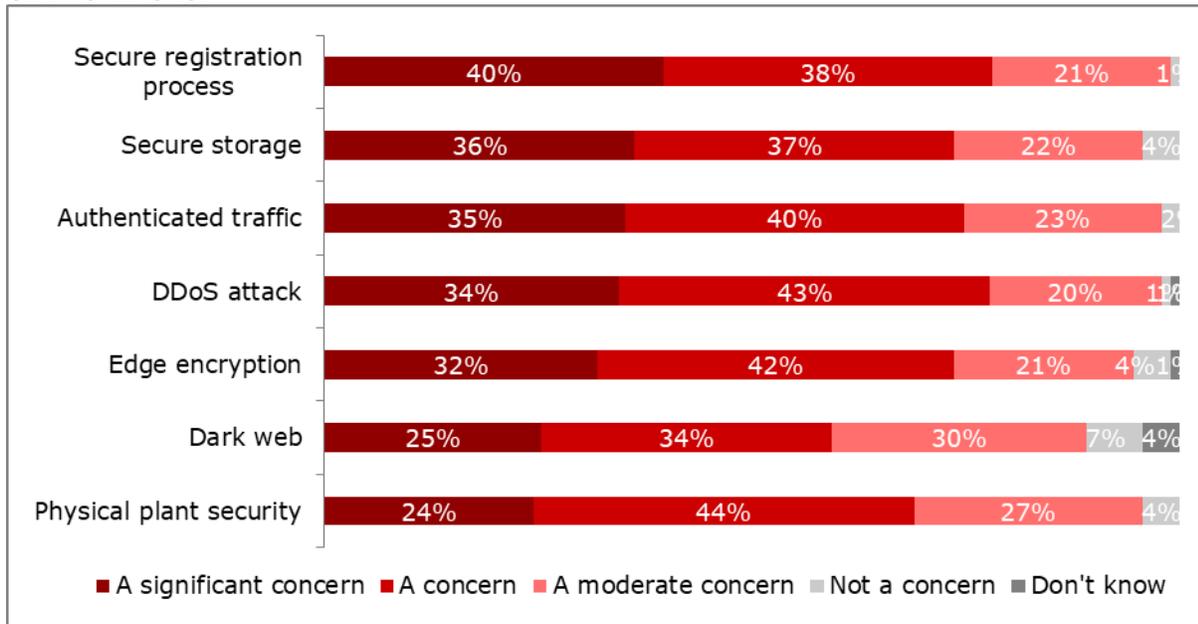
Source: Heavy Reading

Lower cost—with the highest “important” response for this question—is ranked next, rounding out the set of attributes that are top-of-mind with the carriers: security, latency, reliability, and cost. Cost is rarely the top concern among survey respondents, but its impact on the actual buying decision cannot be underestimated. CFOs are very willing to put their thumb on the scale when presented with what they perceive as a bargain.

Are software updates and analytics really of least importance in a list of very important features to the carriers? Heavy Reading believes respondents are viewing them more as tactical concerns, rather than strategic concerns. Analytics affect every other feature on the list. Software updates/upgrades are an element of continuity/reliability and, one could argue, of real-time monitoring.

The most significant aspect of the responses on security concerns in **Figure 16** below is the fact that only “secure registration” is ranked as mostly a “significant concern.” Everything else is ranked as more of a mere “concern” on the anxiety meter, though none dip below to more of a “moderate concern.”

Figure 16: How much of a security concern are the following in the edge environment?



n=91

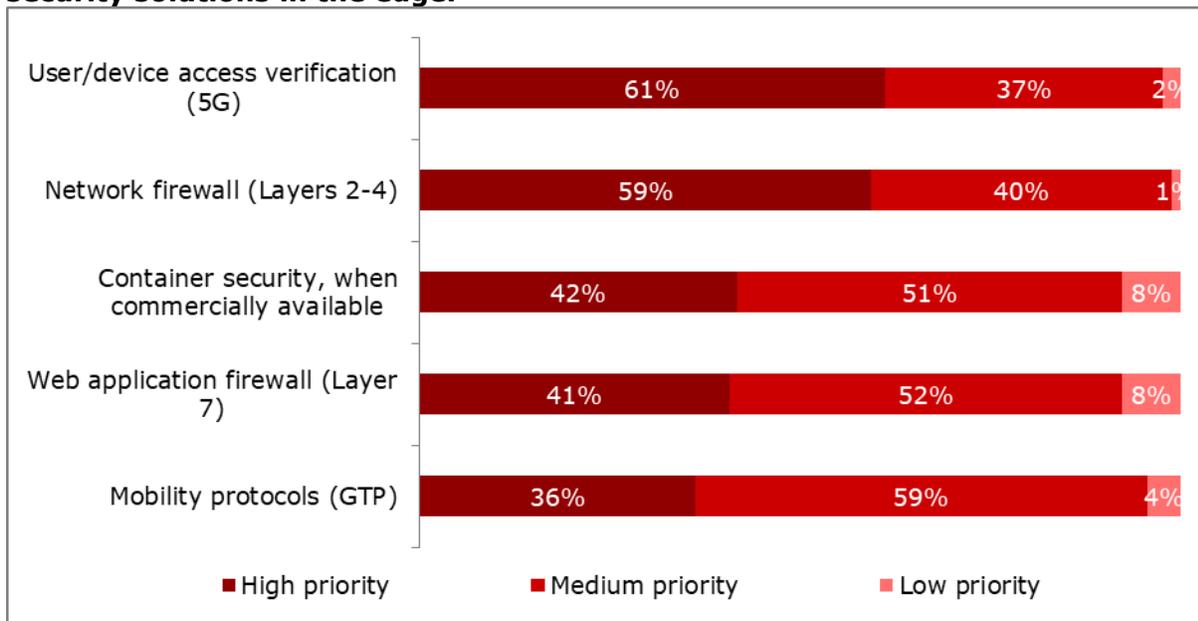
Source: Heavy Reading

“Physical plant security” ranks lowest in “significant concern,” but highest under “concern.” This smacks of a “not my job” response. Carriers often depend on different divisions or separate partners to ensure and provide physical security. To the network personnel, it is a box on a checklist—albeit an important box.

“Dark web” ranks significantly below “DDoS attack” (to which, one could argue, it is related). Heavy Reading believes this is because, while the carrier can guard against, prepare for, and mitigate DDoS attacks, the dark web is more like the weather—difficult to predict, track, or guard against.

Responses to **Figure 17** below are a combination of “go with what you know” and a clear desire to launch a multilayered response to the growing volume of multi-vectored attacks. The response to access verification and network firewalls is almost identical, and these will be the first line of defense and will not lessen in importance. Security is a leading concern when it comes to container networking, which was not designed with a carrier’s idea of security or 5-9s of reliability in mind.

Figure 17: Please rate how you will prioritize implementation of the following security solutions in the edge.



n=91

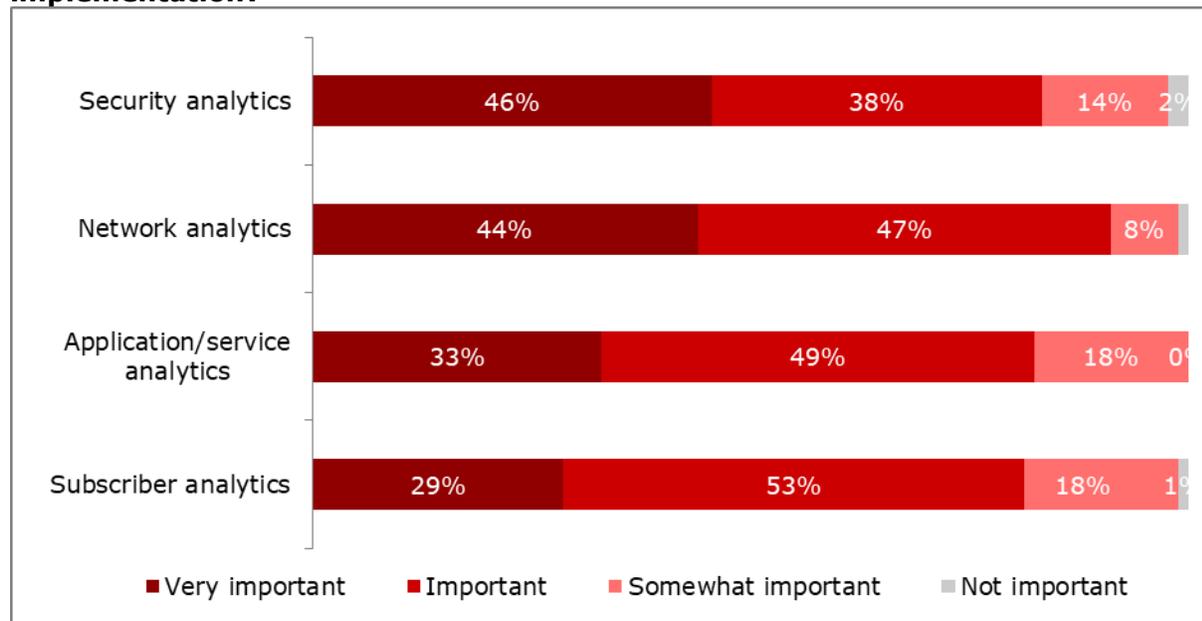
Source: Heavy Reading

Respondents are eager to implement container security when it is available. The high percentage of “low priority” responses is likely due to the fact that most carriers are still in the early days of container deployment. In addition, container security has only recently started to make its way into carrier networks with proofs-of-concept (PoCs). The response to GTP (GPRS tunneling protocol) can be characterized as “of course.”

It is interesting that Layer 7 firewalls show the least amount of exuberance, by comparison only, from the carriers, with 92% stating that it is of high or medium priority. Layer 7 is a key attack vector today. Cloud-based web application firewalls (WAFs) have been deployed by the SPs not only as part of their security portfolio, but also as part of a security as a service offering for enterprise customers. Today, they are being updated to be used with microservices and in a DevOps CI/CD environment.

Heavy Reading believes that the respondent pool looks at analytics as a tactical tool—a subset of, for example, their overall network management or security set of solutions. As shown in **Figure 18** below, security is the only category of analytics evaluated by the respondent pool as more “very important” than “important.” However, it also had a surprisingly high percentage of the less enthusiastic “somewhat important,” as well as the highest (if not statistically relevant) percentage of “not important.” This shows more of a difference of opinion among respondents than any other line item in a security-related question.

Figure 18: How important are the following analytics to your edge implementation?



n=91

Source: Heavy Reading

Network analytics are widely used today and are critical as carriers move to AI/ML-enabled closed-loop management and automation for their expanding networks. Network analytics came in slightly lower than security in “very important” votes, but almost 10 percentage points higher in “important” evaluations, and had the lowest percentage of “somewhat important” selections.

There is no ambiguity among the survey pool regarding the value of network analytics. However, respondents do not seem to be folding application and service analytics into a vision of holistic network management. Such management can evaluate, in real time, all layers of the network, including the applications. It can detect, evaluate, and resolve degradations in performance and service wherever they occur.

The shift from network operations into the customer experience and marketing category of subscriber analytics also portends a move away from the core competency of the survey base. This likely accounts for the respectful but cautious overall ranking of subscriber analytics. With subscriber analytics, SPs are able to move beyond basic parental control and bandwidth upgrade offers to strategic, targeted marketing campaigns. Some of the “other” write-ins touched on subscriber analytics.

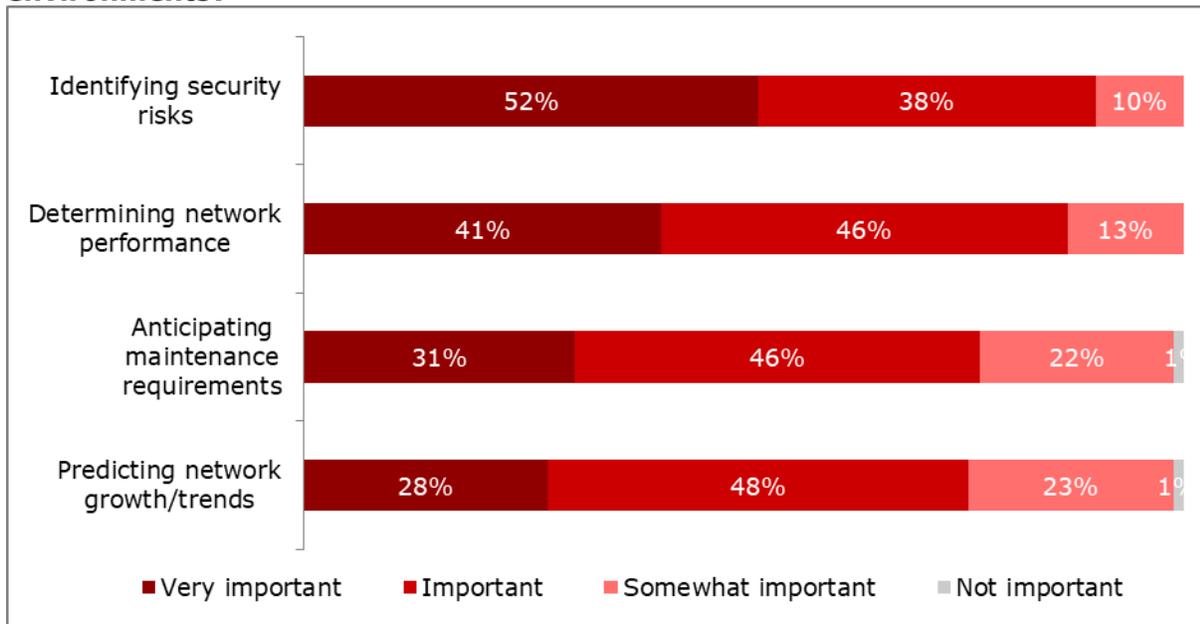
Of note, each of the five “other” write-ins appear to be part of one of Heavy Reading’s specified types of analytics. The write-ins, along with their selected level of importance, included the following:

- “Real-time analytics”: Unspecified level of importance (could be added to network)
- “We have a consumer centric AI security/privacy platform”: Very important (could be added to security or subscriber analytics)

- “Traffic engineering”: Very important (could be added to network)
- “App analytics”: Important (could be added to application/service analytics)
- “User analytics”: Important (could be added to part of subscriber analytics)

Looking at the question of what is important in an edge implementation from a slightly different perspective, the same two key concerns emerge: first security and then the network. When the network is separated into performance, maintenance, and growth/trends, the weight of responses shifts away from “very important” toward “somewhat important.” This shift also indicates a move from the immediate requirements of network performance to the predictive nature of anticipating maintenance needs and forecasting growth trends. The survey shows no disconnect between the need for network analytics and the importance of the network management tasks. The latter rely on analytics to establish a well-understood baseline of performance and extrapolate the trends and anomalies needed to provide predictive maintenance and identify near- and long-term growth trends.

Figure 19: How important are the following attributes in enabling telco edge environments?



n=91

Source: Heavy Reading

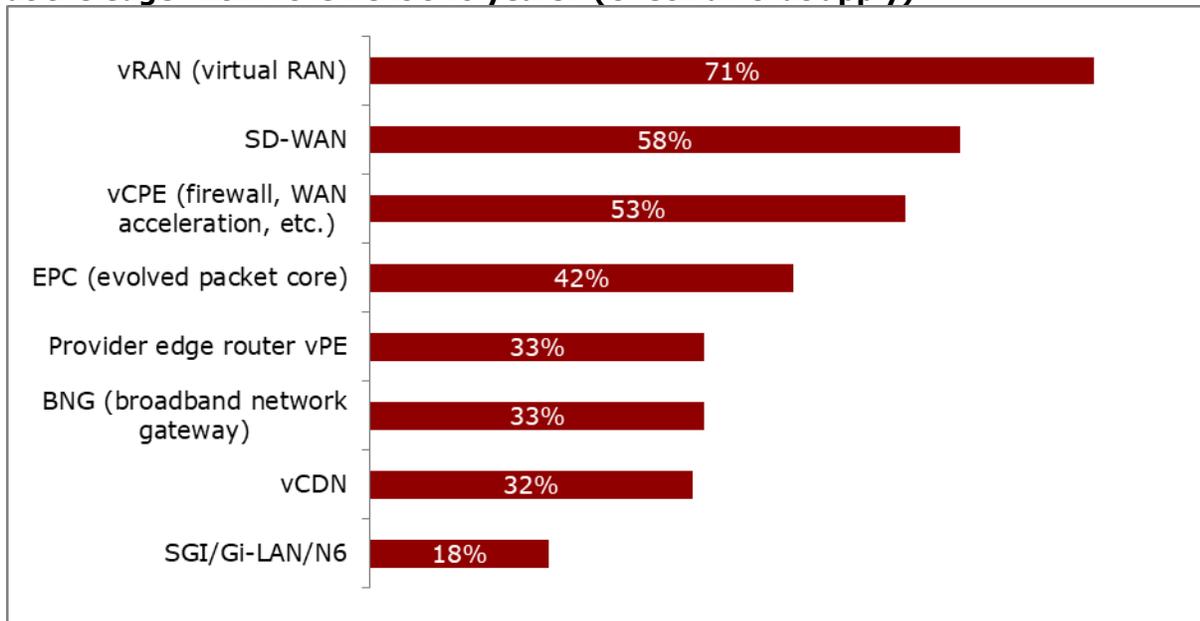
Edge drivers: VNFs, use cases, and vertical industry

Recent Heavy Reading surveys are gratifyingly consistent when it comes to the network functions that respondents intend to virtualize, whether at the core or at the edge of the network. vRAN has led the pack over the past 12 months as SPs plan for and deploy 5G. However, vRAN was a leading virtualized network function for 4G implementation in 2018/2019, as well. Virtualized customer premises equipment (CPE) and software-defined WAN (SD-WAN) continue to be the next two most popular virtual network functions (VNFs), as they are the foundation of the edge network. Evolved packet core (EPC) solutions were early favorites with SPs, as they were largely centralized, residing at the core, and did not involve lengthy, complex and tough to deploy/manage/upgrade service chains.

As shown in **Figure 20**, the virtualized content delivery network (CDN), broadband network gateway (BNG), and virtual provider edge (vPE) are tied with each other in terms of survey pool interest. SPs have tried their luck with CDNs off and on for years. It will be interesting to track whether edge computing and 5G, and the expansion of storage domains they bring, urge carriers to invest more in CDN solutions.

As typical of virtualized solutions, the virtualization of the provider edge router results in a significant improvement in agility and flexibility. However, the provider edge has significant pressures on it to support, in the case of cell site routers, up to 300Gbps (with port speeds of 25Gbps–100Gbps). Provider edge router functionality is also integral to vCPE, vRAN, and network slicing. Heavy Reading expects that virtualized router implementations, particularly at the provider or customer edge, will accelerate.

Figure 20: Which VNFs (virtual network functions) are you planning to implement at the edge within the next two years? (Check all that apply)



n=91

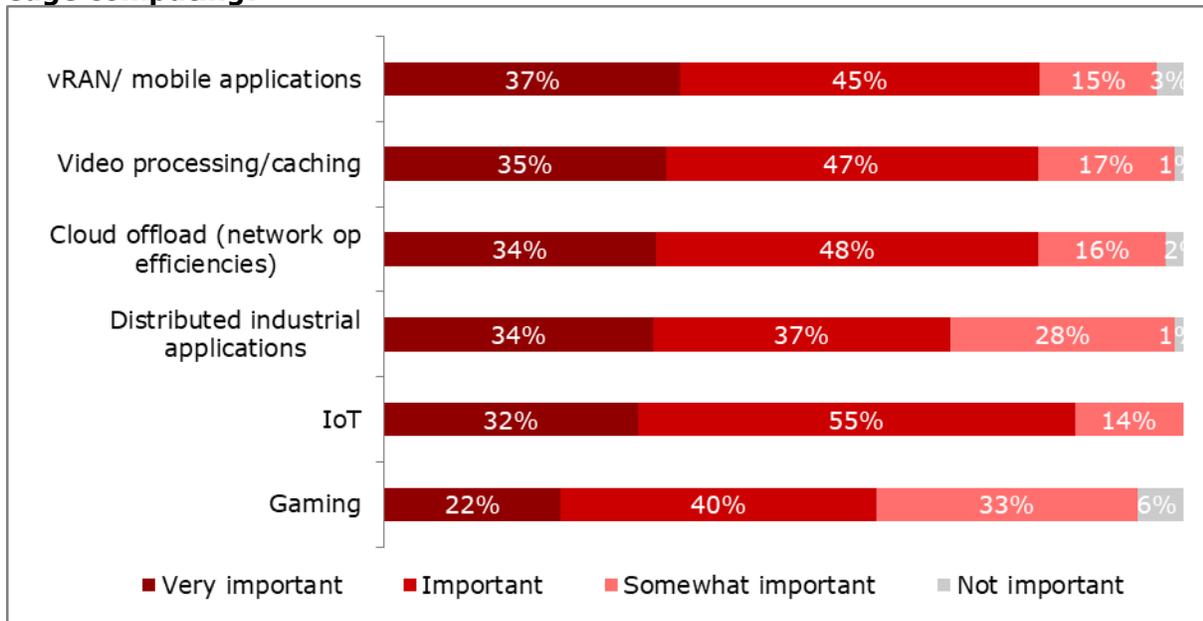
Source: Heavy Reading

While the Gi-LAN ranks lowest in the survey, Heavy Reading believes this is more of an issue with awareness and potentially being too much in the technology weeds for the survey pool. The Gi-LAN certainly benefits the SP, lowering both capex and opex through the use of commercial off-the-shelf (COTS) hardware. It also enhances network flexibility and service agility with simplified network architecture and service lifecycle management and improved management and orchestration capabilities.

Figure 21 below shows that all the use cases are “important,” but none show higher percentages for “very important” (compared to “important”). The widest gap between the scores for “very important” and “important” are in the IoT and gaming use cases. Gaming also has the highest percentage of “somewhat important” and “not important” votes and ranks at the bottom of the use cases—despite the frequency with which it is cited to illustrate the benefits of edge computing.

SPs that Heavy Reading has interviewed confess that the more prosaic use cases are driving edge computing implementations today (such as vRAN, video processing/caching, and operational efficiencies). It is important that SPs acknowledge the impact of these use cases, rather than waiting for high profile, “killer apps” to emerge. “Other” responses included machine learning (“very important”), storage applications (“somewhat important”), and cloud analytics (“somewhat important”).

Figure 21: How important are the following use cases in your implementation of edge computing?

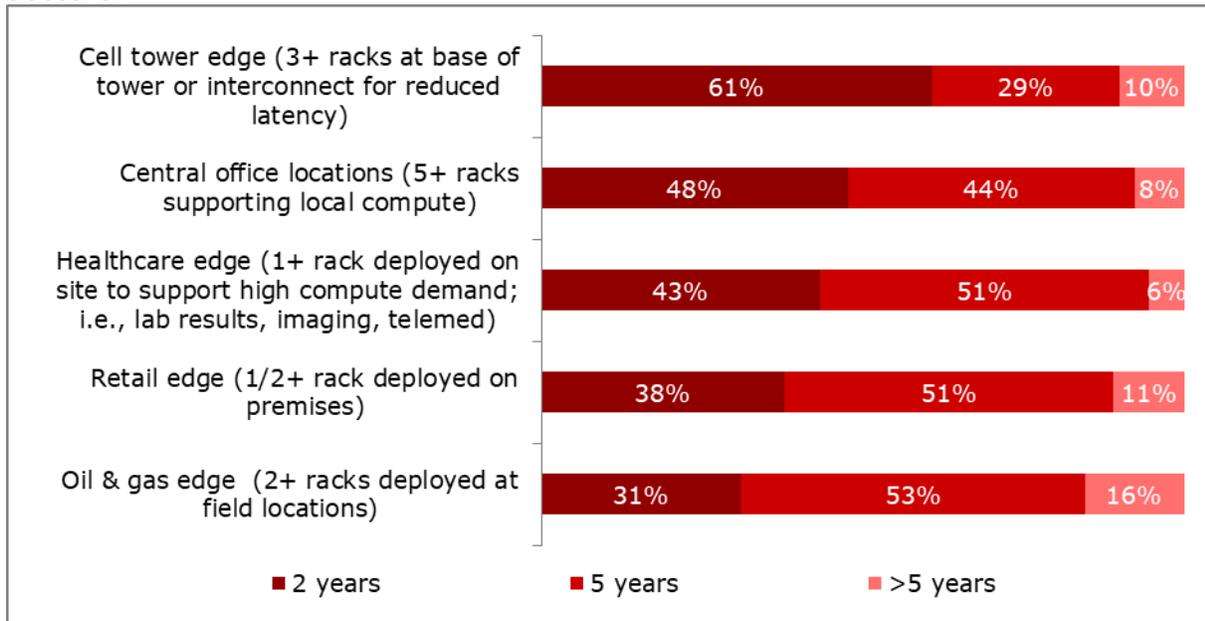


n=88

Source: Heavy Reading

SP responses to **Figure 22** below reveal a logical implementation plan that starts at the RAN edge sites, most likely with the rollout of 5G. They plan to proceed to central offices and then take on an industry that has been significantly affected by COVID-19: telemedicine. Significant deployments will likely be seen in two years, and the bulk of deployments are expected to hit in two to five years. The SPs then expect to tackle two other verticals that have been proactive in planning for edge computing: retail and oil & gas. It is interesting that oil & gas ranks lowest and includes the highest percentage of >5 years, as the oil & gas industry is a leader in the charge to private LTE and 5G networking.

Figure 22: When do you expect mass edge deployments in the following industry sectors?



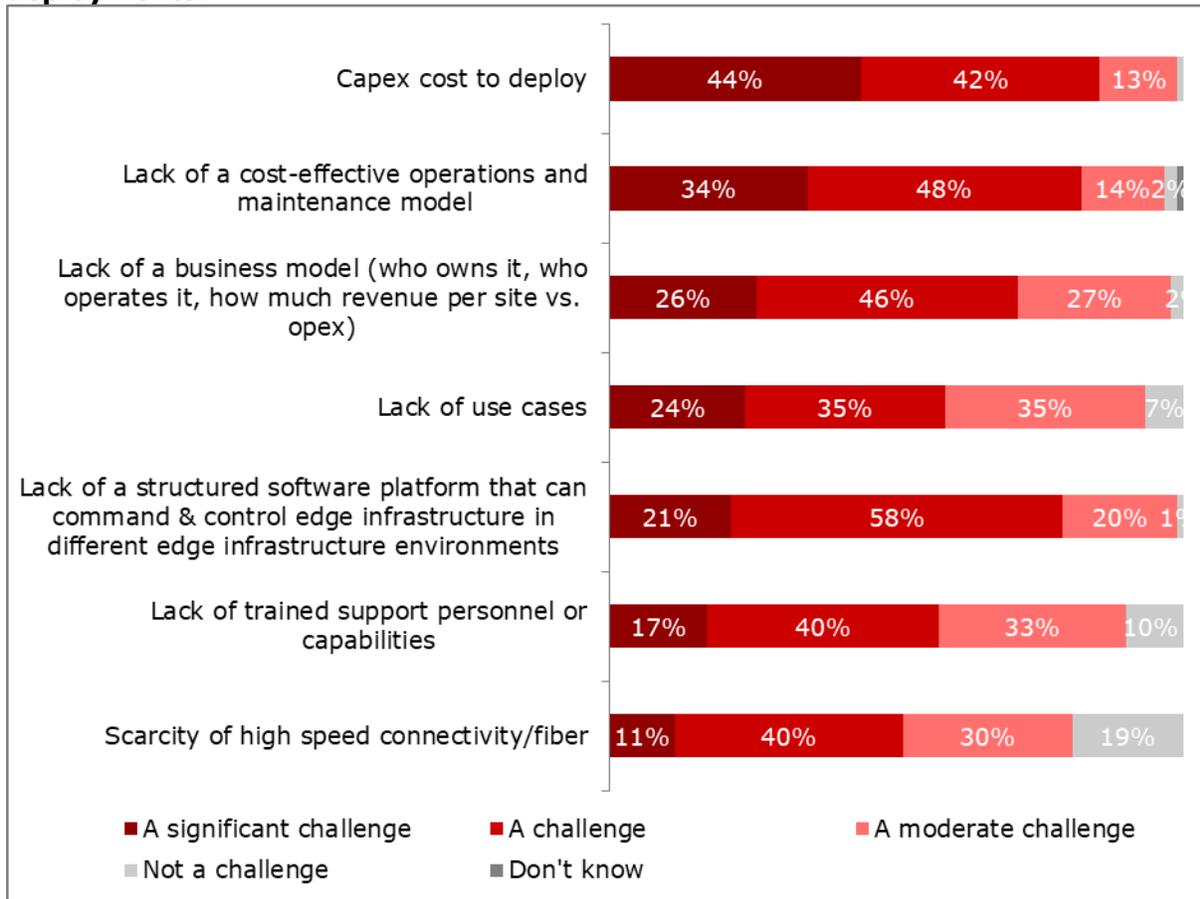
n=90

Source: Heavy Reading

Last question

When looking at the challenges facing the SPs in edge deployments, the first barrier—"cost to deploy"—is the only one cited by more respondents as a "significant challenge" than as simply a "challenge" in **Figure 23** below. All others are weighted heavily by SPs as "challenges." While "lack of use cases" is fourth in a list of only seven challenges, it has a significant impact on the "cost to deploy" response. As mentioned in the analysis for **Figure 21**, vRAN and video processing/caching are important use cases. However, most SPs deploying edge computing today are being driven by general efficiencies in network operations and application performance and by user demand for low latency response time for certain applications.

Figure 23: How much of challenge are the following to full-scale edge deployments?



n=91

Source: Heavy Reading

The second and third most significant challenge can be summarized as “we need a plan”—whether it is for operations and maintenance or a business model that outlines revenue models and organizational structure. Respondents also feel that the software platform for remote command and control is a significant hurdle. This is the first challenge that is largely under the control of the SP’s vendor partners and can be a key differentiator in a sales situation.

The last two—trained personnel and lack of fiber—may delay edge deployments. This is particularly the case in the COVID 19 era, where revenues are declining, headcounts are frozen, and fiber construction schedules have been pushed out.

CONCLUSIONS

Network operators know that their current method of designing and deploying networks will not hold up to near-term demands, let alone future demands. There are multiple waves of change sweeping the network. The COVID-19 pandemic, and the pressures that it has placed on the network, has brought into focus the fact that the world has moved beyond simply connecting people to one in which “always-on” access to digital services is seen as a basic need.

Heavy Reading’s survey results show that SPs believe that edge computing will be a key tool to deliver on the demands of the network now and in the future. Many concerns remain, however. Security definitively heads the list, making enhanced security features and functionality a key differentiator in carrier buying decisions.

Security is followed closely by network performance as a top-of-mind issue. The ante in terms of network performance is continually being raised as devices flood the network, ultra-low latency applications emerge, overall traffic volume doubles every three years, traffic patterns shift with the migration of workloads to the cloud, and work-at-home becomes a permanent fixture.

Cost does not always emerge as a top concern in Heavy Reading’s surveys, partly because the survey base is weighted toward the technical and network operations respondent. It is notable that it emerges as the top concern in the last question of this survey. This result is linked to the fact that many respondents are not sure of the early use cases for edge deployments, combined with the fact that the carrier edge is vast—comprising tens of thousands or even hundreds of thousands of locations. Edge implementation is impossible without a cost efficient solution from both a capex and opex perspective. At this stage of adoption, the use of flashy but still futuristic use cases such as remote surgery and autonomous vehicles is adding to the hype but not helping to establish a pragmatic way forward. The industry has moved past the “fire the imagination” stage of edge computing and into the “how do we actually do this?” stage. SPs are looking to understand the near-term justifications of edge computing, what today’s use cases are, how it will be managed, and how it will be secured.