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

HEAVY READING WHITE PAPER

Harnessing Edge Computing for 2022

A Heavy Reading Market Leader Survey Report Produced for Accedian, Kontron, and Red Hat

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

INTRODUCTION

The edge has become the epicenter of strategic focus in the information and communications technology (ICT) sector, opening up new market opportunities across the ecosystem. To understand the industry, technology, and application evolutions driving edge computing, Heavy Reading conducted a survey with 82 CSPs that have launched edge computing solutions or are planning to do so within 24 months. This survey also aimed to examine critical issues, such as the timeline of adoption and the anticipated ROI for communications service providers (CSPs).

Heavy Reading wanted to understand how CSPs are selecting the platforms and systems that power a highly distributed edge. What are the architectures that can support the compute, storage, and latency requirements of enterprise customers in a differentiated and competitive manner without stressing the CSP's budget? What are the management tools, security implementations, performance assurance and automation solutions, and partnering strategies that enable an agile and flexible CSP and can be replicated across the vast edge frontier?

This report presents the highlights of the Heavy Reading survey, revealing insights on the following questions:

- Which vertical industries are using multi-access edge computing (MEC) today?
- What is the primary goal of MEC implementations?
- How are CSPs working with hyperscalers to deploy MEC and minimize their own near-term risk?

SURVEY DEMOGRAPHICS

Mobile and converged operators made up the bulk of Heavy Reading's survey respondent pool, accounting for more than three-quarters of overall responses (see **Figure 1**). An additional 11% came from the fixed-line and cable operator community. The remaining 13% hailed from hosting/cloud providers, mobile virtual network operators (MVNOs)/mobile virtual network enablers (MVNEs) with infrastructure, and over-the-top (OTT) service providers.

Figure 1: Mobile and converged mobile/fixed operators dominate

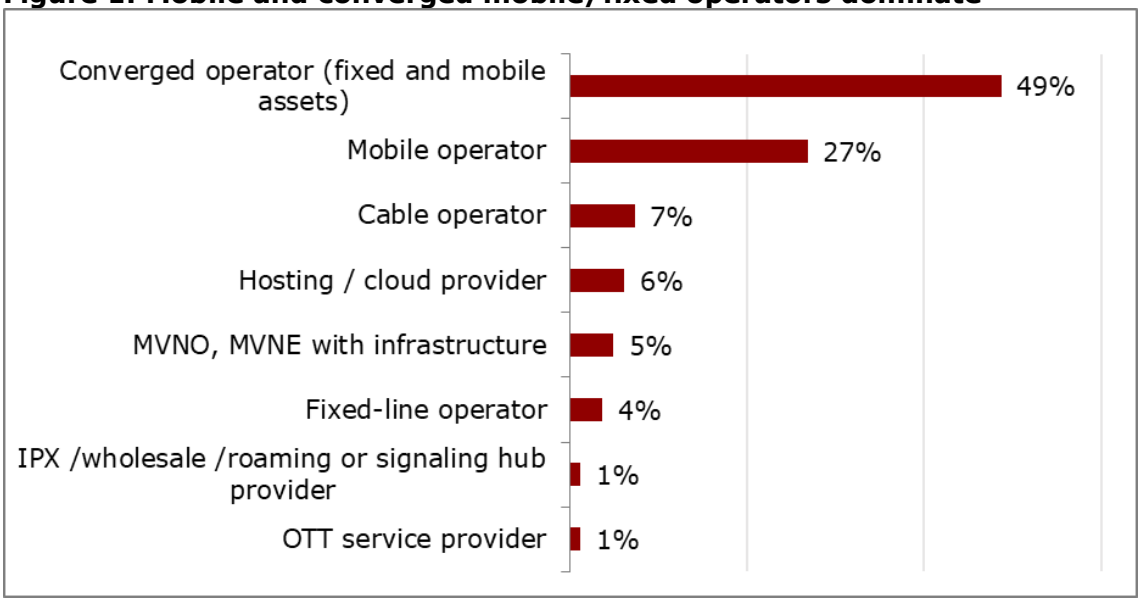

Q: What type of service provider do you work for? (n=82)
Source: Heavy Reading

Among the survey respondents, 40% represented very large CSPs with annual revenue of more than $5bn (see **Figure 2**). CSPs with revenue of between $500m and $5bn made up a third of the respondent pool, and those with revenue of between $50m and $499m made up the remaining 25%. Revenue dictates the capital budget available for funding the transition to 5G, edge computing, and cloud native networking. Heavy Reading research shows that carriers over the past decade have dedicated, on average, 17–18% of revenue to capex.

Figure 2: Tier 1 service providers dominate the survey

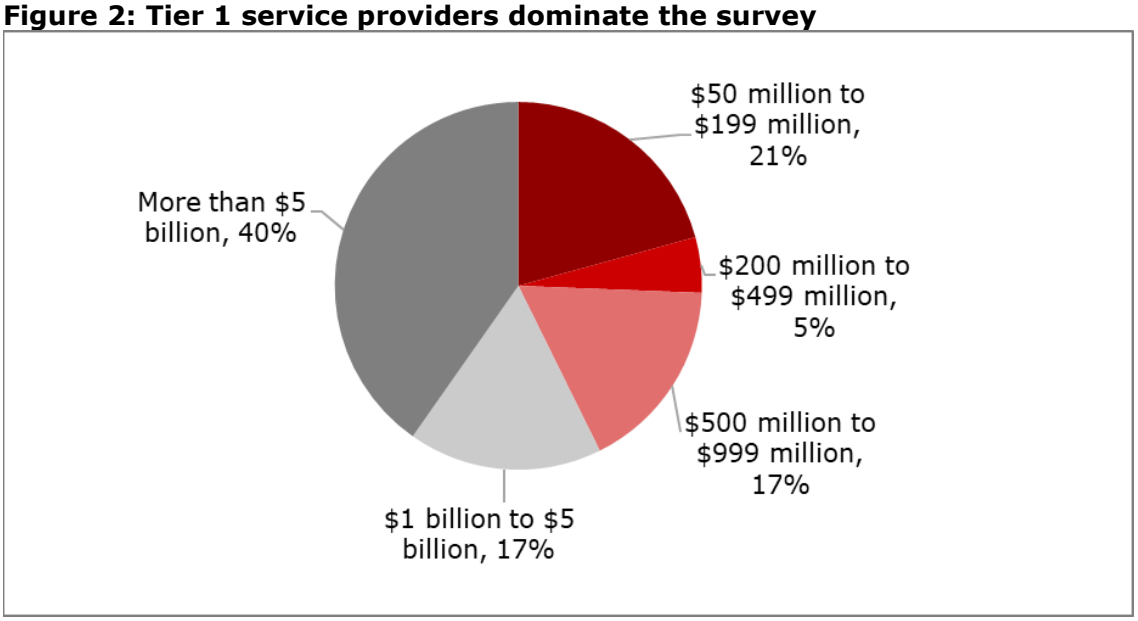

Q: What is your company's approximate annual revenue (US$)? (n=82)
Source: Heavy Reading

Regional breakdown

Just under half of the survey respondents were from the US. Eastern and Western Europe, together with the Middle East, accounted for 27% of respondents. Canada, Central America, and South America made up 16% of respondents. The remaining 10% of the respondents were from the Asia Pacific region.

Job function

As is the case with most Heavy Reading surveys, the majority of respondents—63%—were from technical networking roles: planning and engineering, R&D, product management, and network operations. One-quarter were from management and marketing. The data center, IT, and cloud made up 7% of respondents.

EDGE COMPUTING DEPLOYMENT IS WELL UNDERWAY

Heavy Reading was interested in understanding the drivers, challenges, and goals specifically for edge computing deployments, so the respondent pool was restricted to CSPs that had already deployed edge computing or have plans to do so within 24 months. They were not hard to find. The majority of respondents (39%) have already deployed edge computing, a third plan to do so in the next 12 months, and the remaining 29% plan to do so within the next 24 months (see **Figure 3**).

Figure 3: CSPs are committed to the edge

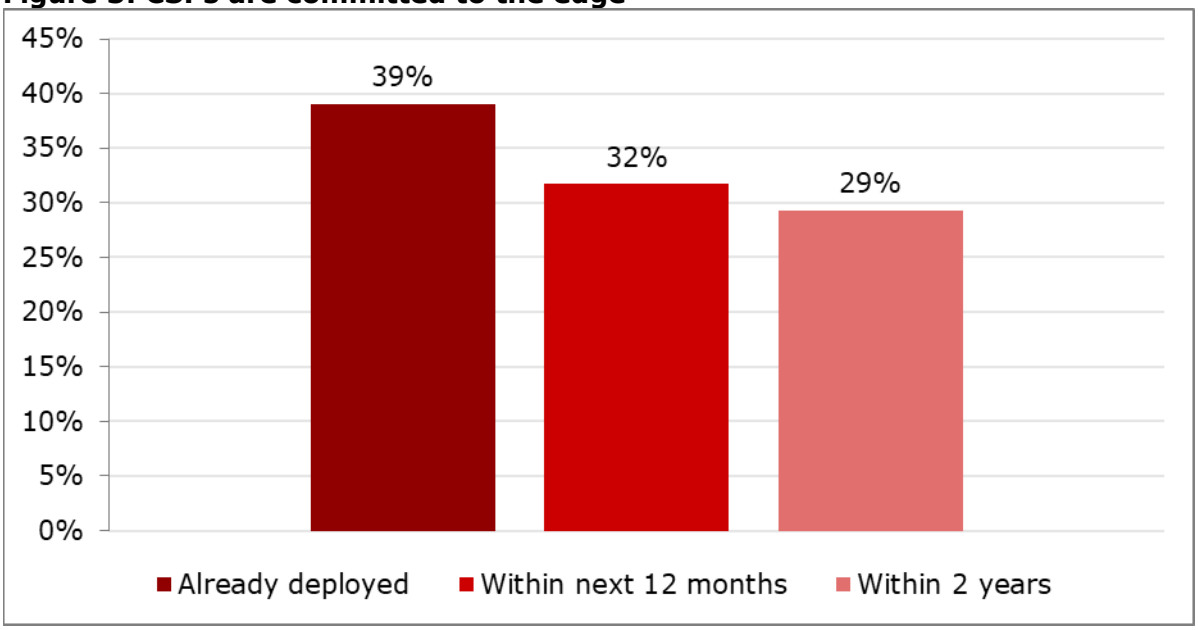

Q: When will your company deploy edge computing solutions in a production environment? (n=82)
Source: Heavy Reading

Benefits driving CSPs to the edge

The most significant benefit of edge computing, according to the survey respondents, is its ability to maintain or improve customer quality of experience (QoE) (see **Figure 4**). In early implementations, this often pertains to support for applications that require ultra-low latency.

Figure 4: Perceived benefits of edge start with customer experience

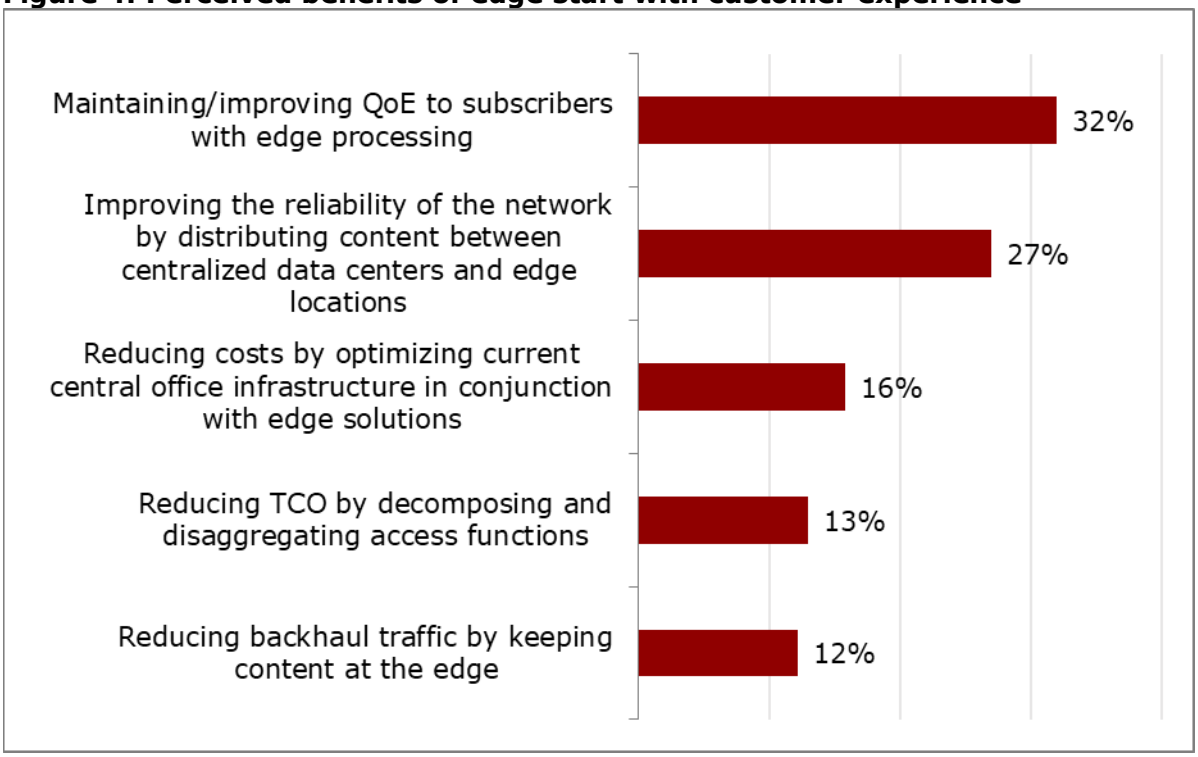

Q: What is the most important benefit of edge computing to your company? (n=82)
Source: Heavy Reading

Every vertical industry has applications that can benefit from the low latency capabilities of edge computing (see **Figure 5**). However, the benefits of edge computing extend beyond support for ultra-low latency. For example, remote healthcare enabled by edge computing provides the ability to process data locally and provide clinical decision support to a wide variety of healthcare workers. Edge computing can also help with Health Insurance Portability and Accountability Act of 1996 (HIPAA) patient privacy compliance. In addition, edge computing platforms that secure patient information locally in a distributed manner are less vulnerable to inadvertent or malicious breaches.

The second most popular response, "improving the reliability of the network," also leans more toward improving customer experience than toward optimizing the network—a benefit that is represented in the bottom three responses.

In fact, if only looking at the responses from very large carriers (annual revenue >$5bn) or at responses from CSPs that have already implemented MEC, there is a 15-percentage-point gulf in "improving the reliability" between these parts of the survey population and the remainder of the respondents. Thus, this is, by far, the most compelling benefit for large and/or experienced carriers.

Figure 5: Ultra-low latency use cases are found across vertical industries

Smart factory/ industrial automation	Healthcare	Entertainment	Transportation	Manufacturing	Energy
Industrial control	Remote diagnostics	Immersive entertainment	Driver assistance apps	Motion control	Smart energy
Machine-to-machine	Remote surgery	Online gaming	Enhanced safety	AR and VR applications	Smart grid
Robot control	Emergency response		Autonomous driving	Remote control	
Process control			Traffic management		

Source: Heavy Reading

Low latency applications are still perceived as the leading use case for MEC, regardless of the respondents' region or deployment status (see **Figure 6**). However, when Heavy Reading divides the survey pool between the large Tier 1 companies with revenue of over $5bn and everyone else, some significant differences regarding dominant use cases emerge. Satisfying growing enterprise demand for private 5G is the dominant use case for large CSPs, followed by improving overall network performance and operations with distributed 5G core functionality.

The smaller carriers, on the other hand, are placing more emphasis on the radio access network (RAN) side of the network as a driving use case. These carriers are looking for edge solutions that can combine distributed unit (DU) and potentially centralized unit (CU) functionality with MEC storage and compute capabilities. The challenge for these implementations will be the per-unit cost, which is already a barrier in 5G DU deployment. This includes environmental hardening, where needed.

Respondents from the cable operators (most of which are under the $5bn threshold) are responsible for the bump seen in smart city importance for smaller carriers. All of the cable company respondents identified the smart city as a leading use case, which is a trend that Heavy Reading sees being carried out in the market, with cable companies collaborating with local governments to accelerate the deployment of MEC.

Blockchain has failed to capture the attention of the CSPs and comes in almost last in the use case list. Artificial intelligence (AI)/machine learning (ML), on the other hand, are critical to the lifecycle management of carrier networks, particularly in a 5G environment. However, the carriers do not see these technologies as driving standalone use cases for MEC, viewing them more as requirements in deploying MEC. Once fully immersed in the approaching "metaverse" dominated by virtual reality (VR)- and augmented reality (AR)-enabled applications, AI/ML will be a critical attribute of emerging applications, but not as a use case itself.

Figure 6: Low latency still a leading use case of edge

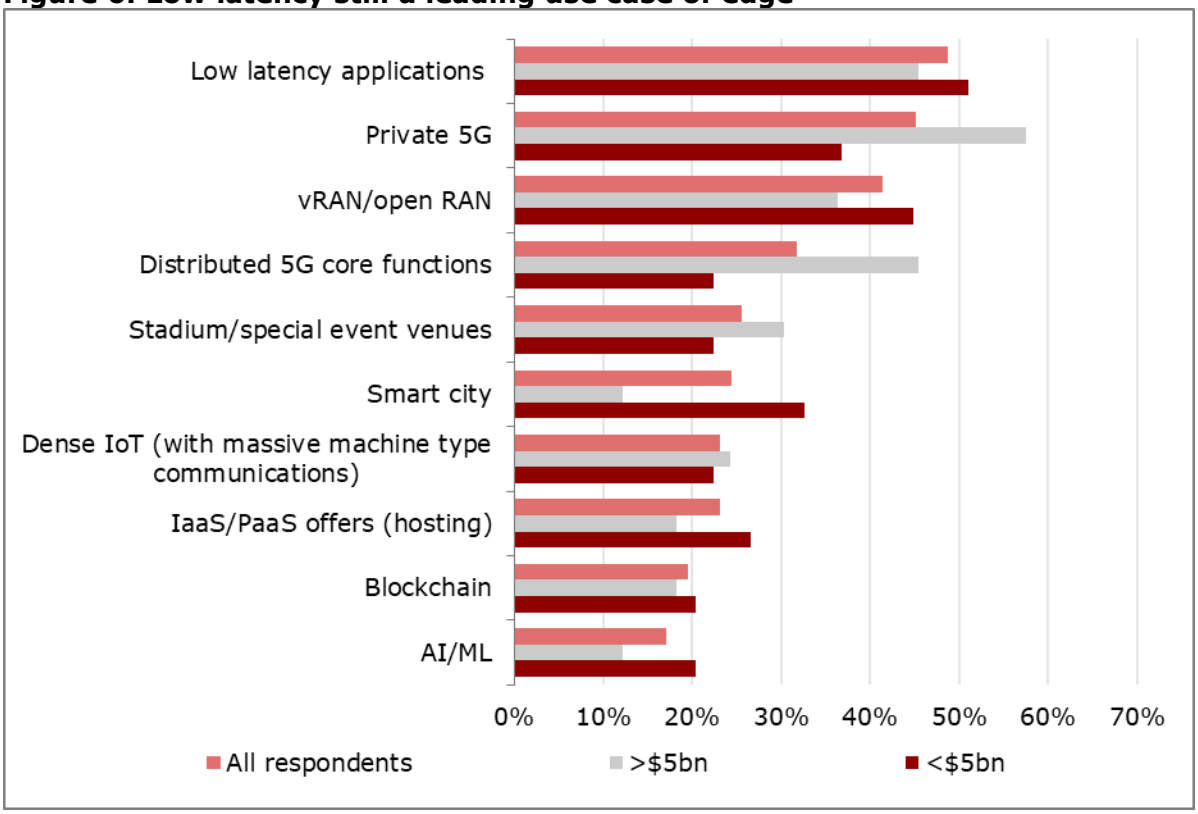

Q: What are your top three edge computing use cases? (n=82)
Source: Heavy Reading

Roadblocks and implementation considerations

The survey respondents appear, at first glance, to agree that a combination of in-house staff, integrators, and system vendors will be needed to deploy MEC (see **Figure 7**). These numbers do not change significantly when Heavy Reading sorts the results by revenue or deployment status. However, when looking only at US-based CSPs compared to all other regions, there is a marked difference. Among the US-based respondents, 38% are confident that they can deploy MEC with in-house staff, compared to 7% of respondents from all other regions. In fact, none of the European respondents (Eastern or Western) expect to deploy MEC exclusively with in-house staff. This can be attributed to many factors, including the highly competitive European market that keeps CSPs very lean in terms of staff, the regulatory environment, and the complexities of operating a network across international borders.

Figure 7: It takes a village to deploy MEC

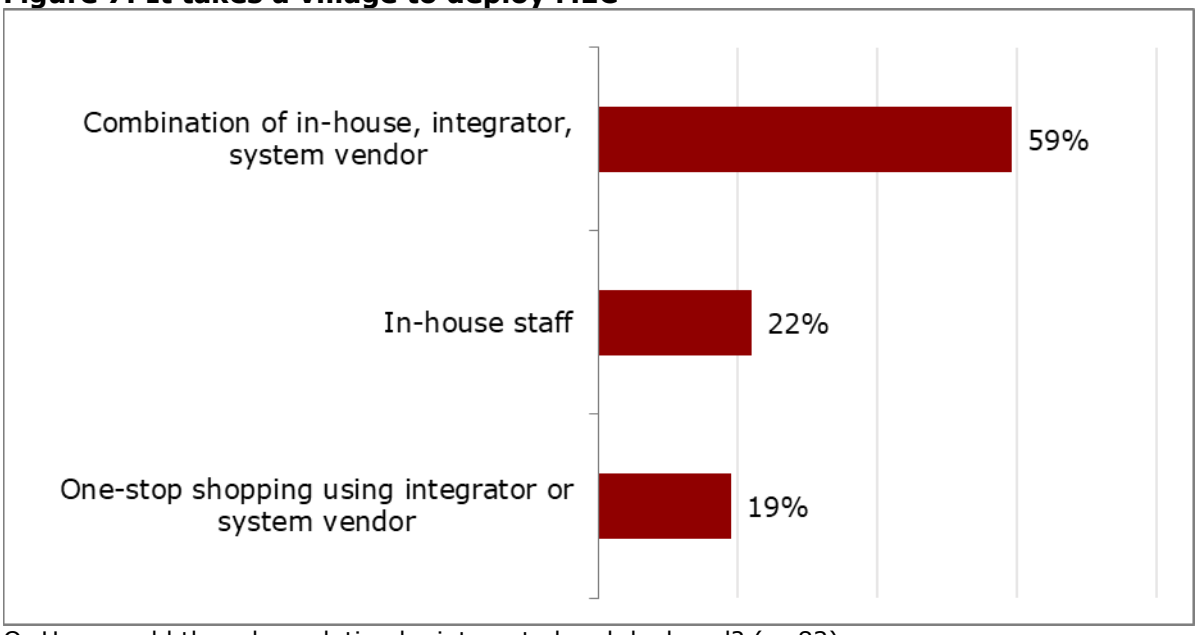

Q: How would the edge solution be integrated and deployed? (n=82)
Source: Heavy Reading

The Heavy Reading survey respondents anticipate a variety of challenges in deploying MEC platforms (see **Figure 8**). None of the suggested challenges were discounted; all pulled in votes from at least one-third of the respondents. None, in this overall view, were selected by more than two-thirds of the respondents.

Systems integration was identified as the most significant challenge. Those respondents that have already deployed MEC evidently learned from experience that it is an even greater challenge than they anticipated; 75% chose it, compared to 60% among the planned deployment population. Conversely, this same population of respondents that have already deployed MEC found space restrictions to be less of an issue (41% compared to 54% for currently planned MEC).

Carriers with revenue of over $5bn saw power as much less of an issue than the remainder of the respondents (39% compared to 57%). This can be attributed to a number of factors. Tier 1 carriers have the market heft to negotiate an attractive electricity rate. In the US, that translates to one-half or less compared to the average consumer rate of 10.5 cents per kilowatt-hour. In addition, power concerns may be outside of the sphere of influence of the respondents, particularly among the larger carriers, where formal green initiatives and carbon-neutral goals are planned and tracked by a separate, dedicated organization.

While large carriers may see power as less of a barrier than their smaller counterparts, they see a lack of environmental hardening as more of an issue (45% compared to 31% for smaller carriers). Carriers looking to deploy storage and compute capabilities at the cell tower, on the roofs of high rise buildings, or in residential telco cabinets need hardened and secure enclosures that are cost-effective, particularly when this environment can demand a scale of thousands of edge points of presence (PoPs).

Figure 8: Challenges reflect the complex landscape of MEC deployments

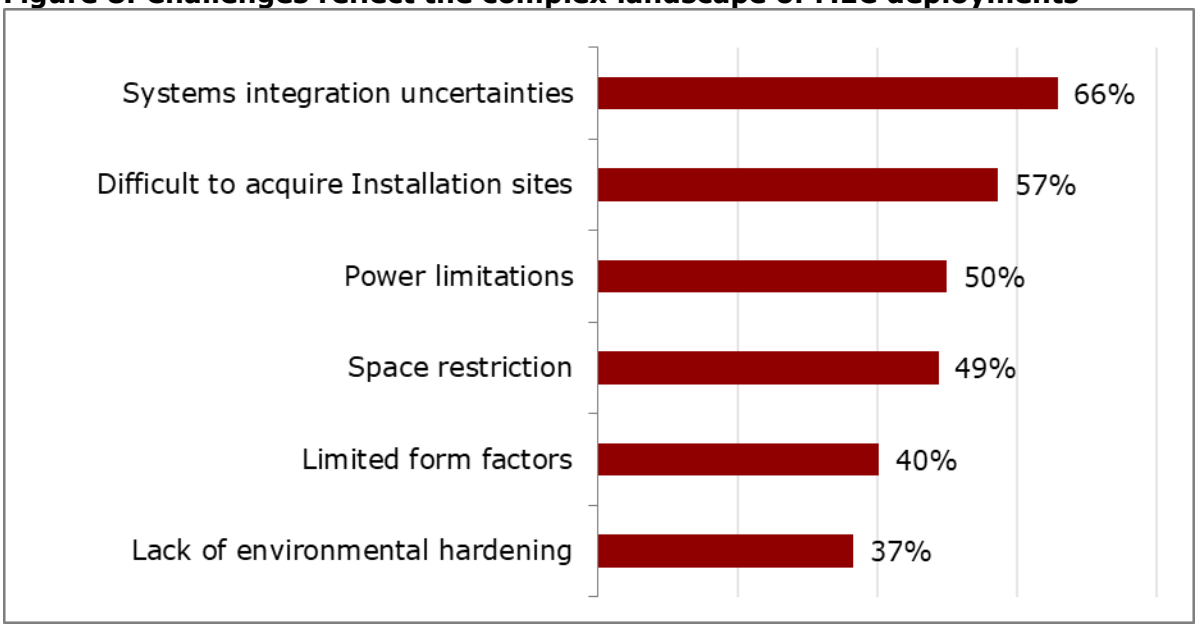

Q: What challenges have you seen, or expect to see, in your edge platform deployments? Select three. (n=82)
Source: Heavy Reading

Heavy Reading asked the survey respondents what they were looking for in an edge platform provider (see **Figure 9**). Product characteristics led by product reliability claimed top marks, while vendor characteristics, such as training and partner ecosystem, clustered at the bottom half of responses. There is only a 29-percentage-point spread between the most popular and least popular responses; all the attributes are important, but none overwhelmingly so. In addition, there is, for the most part, not a significant difference when filtering the data by revenue or deployment status, with some notable exceptions.

Built-in security ranks second in the survey overall but first among carriers that have already deployed MEC. When looking at those versus carriers still in the planning stage, there is a very large swing of 31 percentage points (53% vs. 22%). By comparison, all other platform provider requirements show a swing of, at most, 11 points. Heavy Reading sees a similar but smaller swing for built-in security between the >$5bn carriers, 45% of which voted for security, and the <$5bn carriers, with 27% voting for security.

Edge computing is vulnerable to a number of security risks that attack the platform itself, such as node counterfeiting, node replication, hardware trojan injection, and physical tampering with the edge node. There is also the risk of distributed denial of service (DDoS) attacks or corruption of the routing information (whether malicious or due to human error). CSPs that have deployed MEC have learned from experience that security cannot be an afterthought and bolted onto the solution after the service goes live. It must be integrated into the platform and overall solution, starting with the design phase. Heavy Reading hears this feedback from carriers and vendors every time edge computing is discussed at our conferences, digital symposia, and webinars. It is a critical attribute for CSPs deploying edge computing and vendors designing edge solutions.

The only other significant difference seen when sorting the data by demographic filters is in the professional services response. The very large carriers showed little interest in this (3%), while the <$5bn carriers were much more interested in it (20%).

Figure 9: Top requirements from an edge platform provider to make sure it works and has built-in security

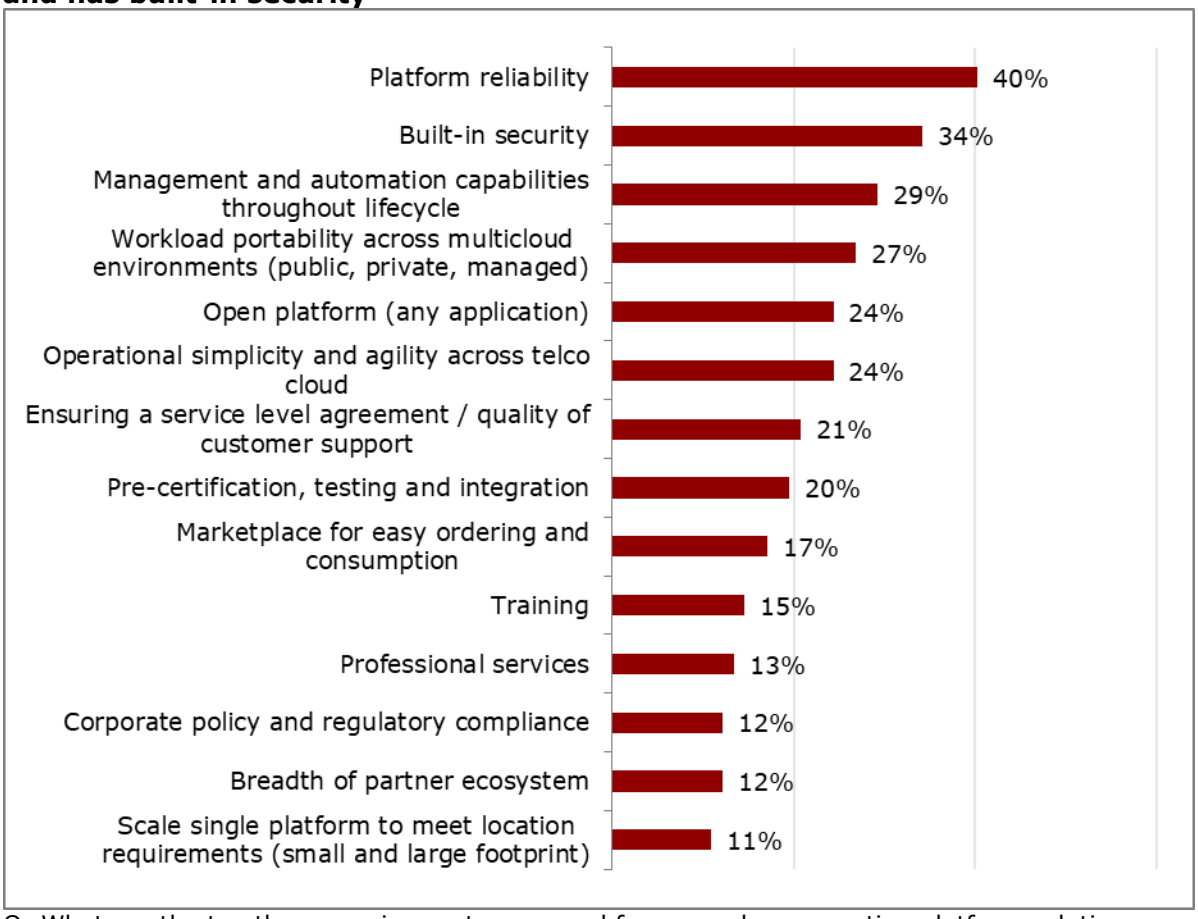

Q: What are the top three requirements you need from an edge computing platform solutions provider? Select three. (n=82)
Source: Heavy Reading

Assuring edge performance

The majority of Heavy Reading's survey base represents mobile and converged operators with both 4G and 5G networks—the so-called brownfield operators. A perpetual priority among these operators is ensuring generational interoperability between networks. It is, therefore, not surprising that 4G/5G network interoperability ranks as the key challenge (see **Figure 10**). CSP concerns about edge computing security, discussed above, make "security at the edge" an obvious choice for second in importance. All remaining challenges have the potential to materially impact MEC performance: managing virtual service layers, containers and microservices, network slices, and so on. How do the CSPs prioritize these challenges? It is interesting to note that the responses from CSPs that are *planning* to deploy edge computing are fairly homogenized, with only a 30-percentage-point difference between their biggest and their smallest challenges. When looking at the responses from

CSPs that have already deployed edge computing, on the other hand, the spread is 56 percentage points. The order of the challenges does not change significantly, but the challenges at the top of the chart show more concern among the deployed CSPs, while the challenges at the bottom show far less. Why is this the case? Real-world experience with the edge. Assuring service performance at the edge is not easy. Controlling and managing edge performance, gaining insight from key performance indicators (KPIs), and correlating virtual and physical layers for automation all have to be considered and built in at the service design phase, not as an afterthought—just like security must be an upfront criterion.

Figure 10: The experience of CSPs that have deployed edge shows in their ranking of performance challenges

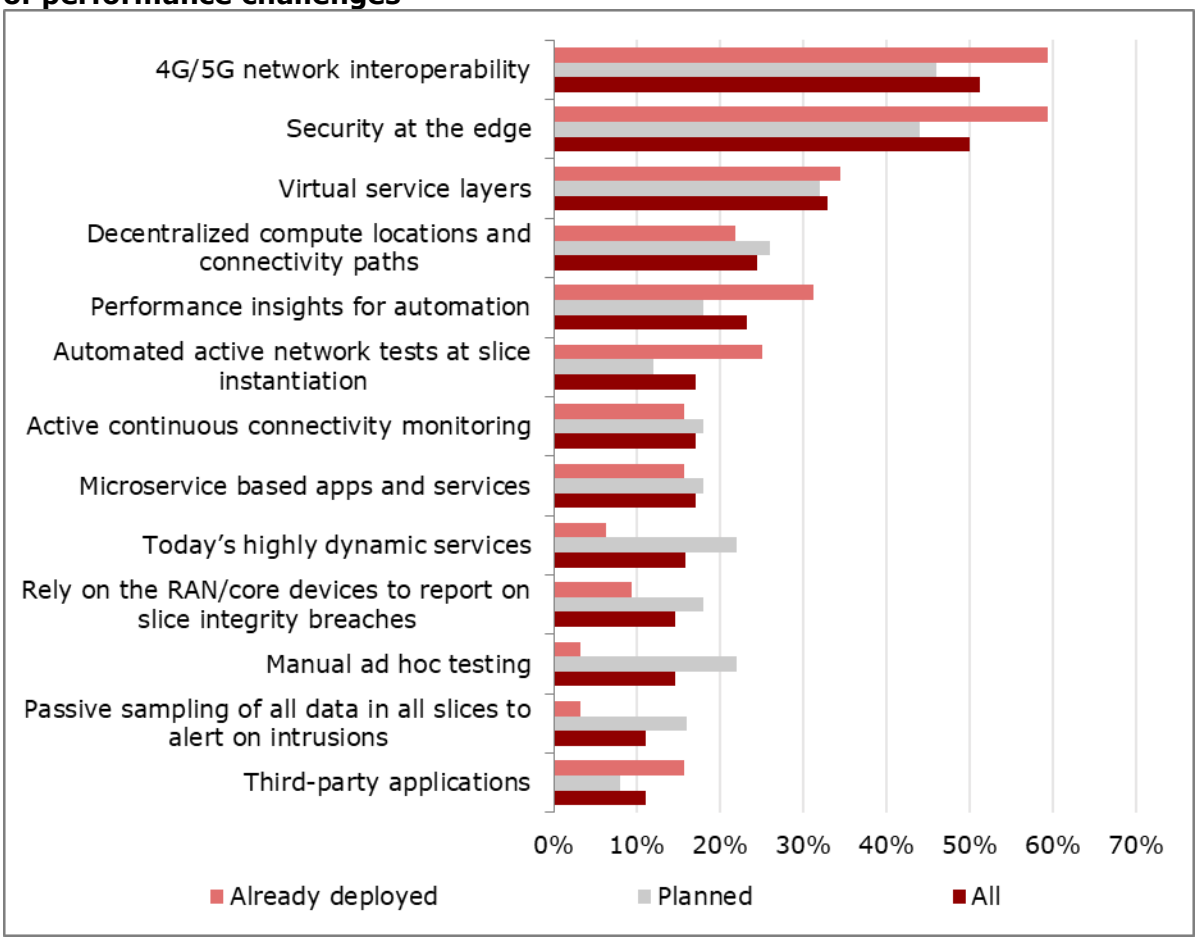

Q: What are the biggest challenges in assuring the performance of MEC-based services at the edge? Select the top three. (n=82)
Source: Heavy Reading

When Heavy Reading asked the survey respondents what performance monitoring and assurance capabilities they would need to address these challenges, their responses showed a significant difference between those that have already deployed edge computing and those still in the planning stages (see **Figure 11**). For the respondents with edge deployments still one to two years out, each capability accounted for not less than 20% and no more than 38% of the respondents. For those who have already deployed, the responses have a much wider range from 16% to 56%. Furthermore, respondents who have already deployed the edge place far greater importance on the first two capabilities of network

security and network performance quality of service (QoS) KPIs. All remaining capabilities are of less interest—less than or equal to respondents in the planning stage.

The message here is not "implement security monitoring and network performance QoS KPIs and forget about everything else"; it is one of prioritization. CSPs are in the early stages of implementing edge computing. Their in-the-trenches experience, so far, shows that the early focus must be on security monitoring and network performance. Everything else will follow as the service and vendor edge assurance products mature. It is interesting to note that network performance is number five among the planned deployments but rises to number two among those that have already deployed the edge. Similarly, end-to-end device monitoring drops from number one to number four or five once deployment has started. These results suggest that having a QoE testing app for phones may seem good before deployment, but once deployed, continuous network monitoring becomes more important.

Figure 11: Network security and network performance rise to the top for assurance priorities

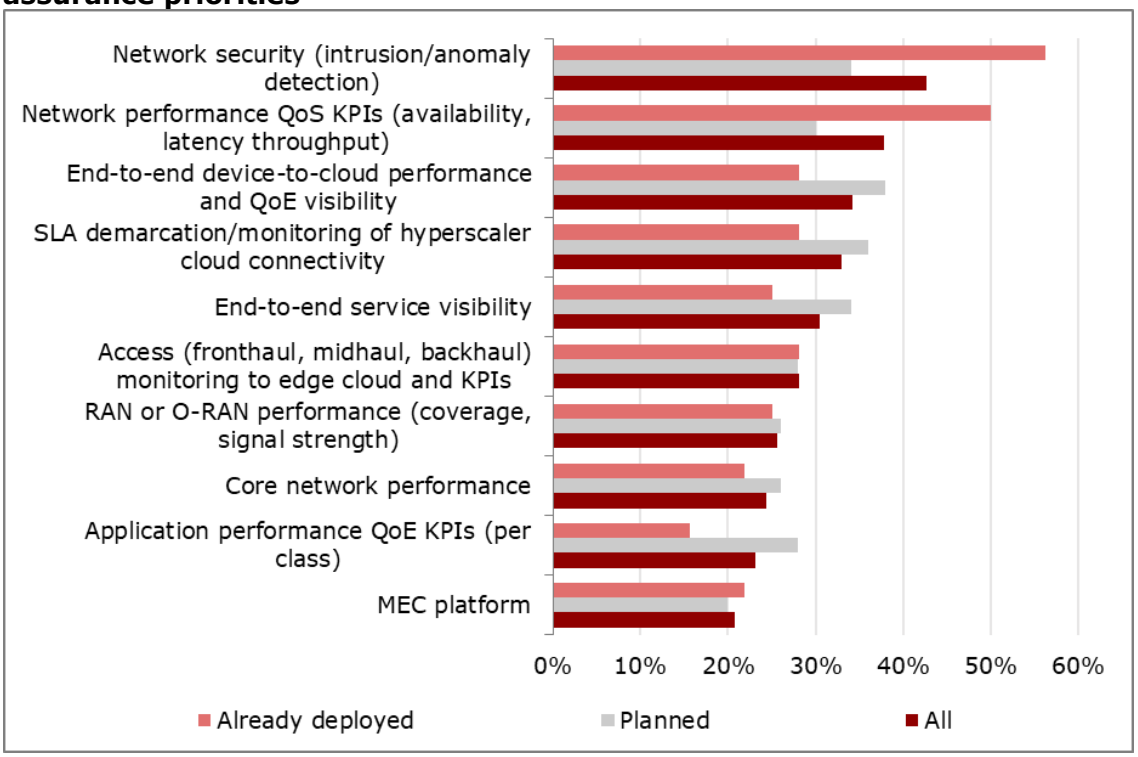

Q: What performance monitoring and assurance capabilities will your organization need for edge computing infrastructure and services? (n=82)
Source: Heavy Reading

The unbridled and continued growth of the network measured by any metric, be that traffic, devices, or cell sites, has encouraged network operators to step away from their multi-monitor network displays and start to rely more on automation for monitoring and proactive maintenance. In addition, the pandemic turned many central offices into lights-out, zero-touch facilities virtually overnight. As a result, CSPs that have been reluctant to relinquish ongoing network control to automated tools are now starting to embrace them (see **Figure 12**).

When Heavy Reading asked survey respondents who would be consuming their metrics data, "internal closed-loop automation or ML/AI tools" was the emphatic number one response. All the humanoid consumers—operations, partners, the end customer—were clustered together, 15 to 20 percentage points behind. This response is even more pronounced among very large carriers (>$5bn), with 48% choosing automation tools as the primary consumer of the data. The need for high quality network performance data and KPIs are essential for success and include accuracy, granularity, and timeliness of KPI data—all analytics that generate actionable insight—when being used directly for closed-loop automation or AI/ML tools.

Figure 12: Assurance that metric data will feed directly into closed-loop automation tools

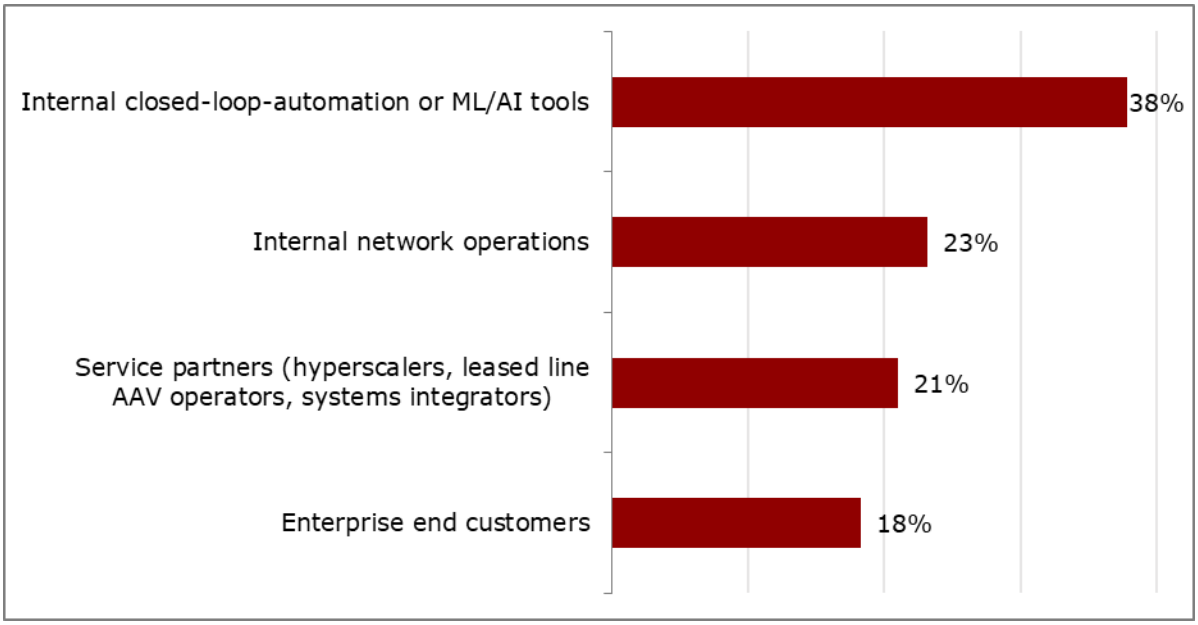

Q: For service and network performance, who will be the primary consumers of your metrics data and KPIs? (n=82)
Source: Heavy Reading

The role of hyperscalers

Hyperscalers have introduced dedicated edge products and embedded their software stack into operator infrastructure, including Internet of Things (IoT) devices and network gateways. They have introduced products dedicated to the telco market, such as Wavelength from Amazon Web Services (AWS), Azure Edge Zones from Microsoft, and Anthos for Telecom from Google Cloud. According to Heavy Reading's survey results, their efforts have paid off, as CSPs have unquestionably decided to partner with hyperscalers in their MEC services.

When looking at reasons to partner with hyperscalers, it is interesting to examine the least popular responses first (see **Figure 13**). They provide notable insights: the CSPs are overwhelmingly planning to work with hyperscalers, and they do not expect to save money by doing so.

So why *are* CSPs partnering with them? The most important reason is ease of deployment. It is even more dominant with the largest carriers—with 67% selecting it compared to 49% for smaller carriers. The same enthusiasm comes from carriers that have already deployed MEC; 66% selected ease of deployment as their top reason compared to 50% of those that have not yet deployed MEC. According to the survey, the larger and further along the carrier is on the journey to MEC, the more likely it is to partner with hyperscalers. The second most popular response—accelerate service time-to-market—is also much more important to large carriers and those that have already deployed MEC, increasing between 12 and 15 percentage points above responses from smaller carriers and those still in the planning stage. The only other response showing a statistically relevant difference between different segments of the survey base is in the importance of out-of-region reach, which is, not surprisingly, much more important to smaller carriers and to those that have not yet deployed MEC (15 and 19 percentage points, respectively).

Enhanced cloud-based security is the third most important reason for partnering with hyperscalers, according to Heavy Reading's survey base, with little variation due to CSP size, region, or deployment status.

Figure 13: Hyperscalers make it easy

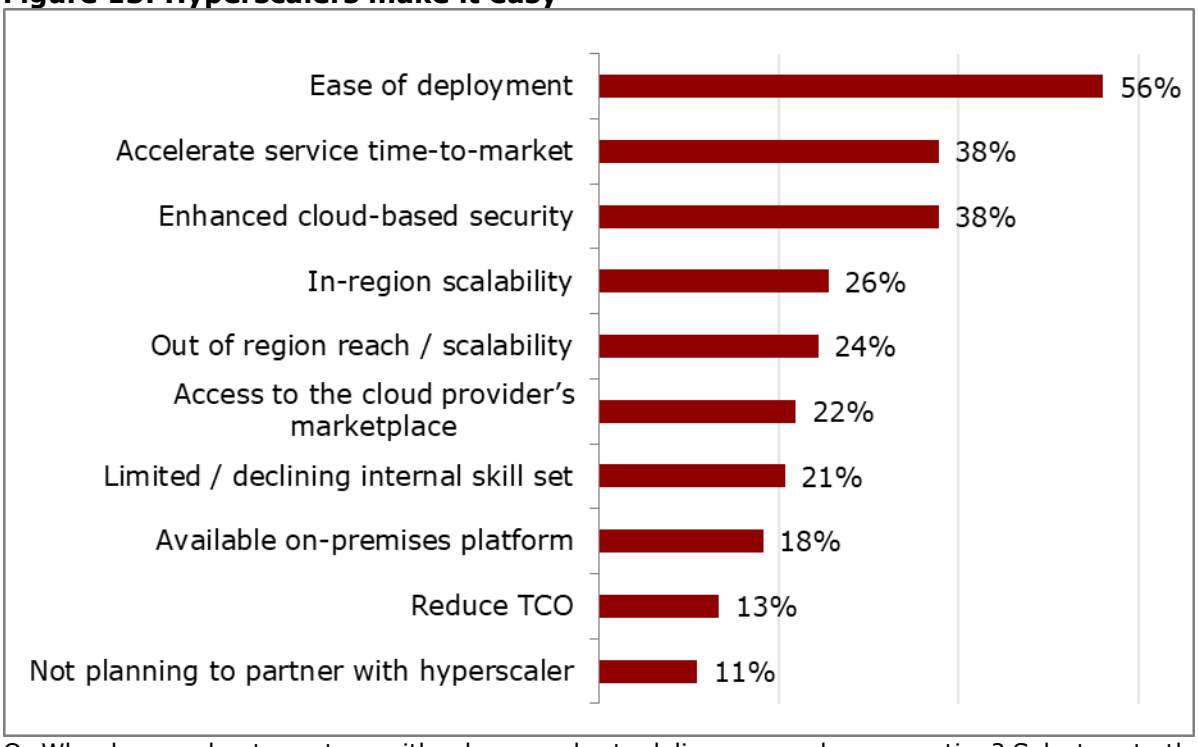

Q: Why do you plan to partner with a hyperscaler to deliver your edge computing? Select up to three. (n=82)
Source: Heavy Reading

THE VIEW FROM THE EDGE

Heavy Reading's survey results show that carriers have committed to edge computing and are progressing rapidly with implementations. Improving the customer experience has risen to the top in terms of goals, rather than improving network operations or lowering costs. The key use cases that have emerged during these early implementations are private 5G and the low hanging fruit of applications demanding very low latency. Overall, 5G is an important driver of the edge, with MEC in the RAN and the distributed core as favored use cases, along with private 5G. Cable, fixed, and smaller operators are also finding leverage with edge computing in their forays into stadiums and smart city implementations.

The deployment of edge computing brings with it issues of scale and complexity. Implementation requirements of space, power, site selection, and so on are all acknowledged as challenges but far from insurmountable with the right partners. CSPs are most concerned with overall network performance and security. In fact, those companies that have already deployed the edge have a heightened concern about these issues. They are looking for help from their traditional vendor and integrator partners, from their network monitoring and assurance tools, and from the hyperscalers.

The overall message that has emerged from Heavy Reading's most recent edge computing survey is the pivot to improved customer experience as the key goal and the acknowledgment that CSPs must clear new paths to achieve this goal. They must do so by leaning into automation, particularly in overall lifecycle management, building in comprehensive security protections and performance control from the design phase forward, and partnering with hyperscalers for ease of deployment, accelerated time-to-market, and enhanced cloud-based security.