

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

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

# Telcos and Hyperscalers: Together on the Edge

*A Heavy Reading Market Leader Survey Report produced for  
Ericsson, Google Cloud, Intel, and VIAVI*

**ERICSSON** 



**Google Cloud**

**intel**®

**VIAVI**

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

---

## Introduction

Hyperscalers have embedded their software stacks and integrated dedicated edge products into operator infrastructure, including Internet of Things (IoT) devices and network gateways. They have also introduced products dedicated to the telco market, such as Google Distributed Cloud Edge from Google Cloud. Hyperscaler efforts are paying off as communications service providers (CSPs) pivot to partner with hyperscalers across the edge value chain. With this survey, Heavy Reading aims to understand how the market for multi-access edge computing (MEC) services is not only progressing but also changing to reflect the advantages and challenges of these CSP/hyperscale cloud provider (HCP) alliances at the edge.

Heavy Reading sought to answer questions such as:

- What is driving CSPs to partner with the HCPs: improved time-to-market, ease of deployment, service reach, or something else?
- What challenges remain to be addressed through these CSP/HCP partnerships?
- How confident are CSPs that they can meet their performance, reliability, and recoverability goals with HCP-enabled MEC services?

The edge has, paradoxically, become the epicenter of strategic focus in the information and communications technology (ICT) sector, opening new market opportunities across the ecosystem. To understand how CSP and HCP alliances are building a new MEC value chain, Heavy Reading conducted a survey of 101 global CSPs that either 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 changing perceptions of the HCPs and their alliances with the CSPs.

This report presents the highlights of the Heavy Reading survey, sharing insights on


- The changing momentum of MEC deployments.
- Where CSPs are looking to leverage the HCPs versus where they are looking to build internal expertise.
- How CSP attitudes have changed about the benefits and risks of working with the hyperscalers on MEC.

---

## Survey demographics

Mobile and converged operators made up the majority of the survey respondent pool, thus accounting for more than half of overall responses (see **Figure 1**). Cloud and hosting providers are growing as a service provider segment in this survey, accounting for 16% of respondents. The fixed-line and cable operator community make up 20%. Mobile virtual network operators (MVNOs) and mobile virtual network enablers (MVNEs) with infrastructure contributed 7% of respondents. The remaining 5% of respondents included over-the-top (OTT) service providers and IPX/wholesale/roaming or signaling hub providers.

**Figure 1: Mobile and converged mobile/fixed operators dominate**




Note: Numbers in figures throughout this report may not total 100 due to rounding.

Q: What type of service provider do you work for? (n=101)

Source: Heavy Reading, 2023

Among the survey respondents, 51% represented large CSPs with annual revenue of more than \$1bn (see **Figure 2**). CSPs with revenue of between \$500m and \$999m made up 19% of the respondent pool, and those with annual revenue under \$500m made up the remaining 31%. Revenue dictates the capital budget available for funding the transition to 5G, edge computing, and cloud native networking. Heavy Reading research shows that over the past decade, carriers have dedicated, on average, 17–18% of their revenue to capex, a percentage they are working to lower with a transition to commercial off-the-shelf (COTS) servers and virtualized and containerized network functions.

**Figure 2: Tier 1 service providers dominate the survey**




Q: What is your company's approximate annual revenue (USD)? (n=101)  
Source: Heavy Reading, 2023

### Regional breakdown

Some 40% of the survey respondents were from the US. Eastern and Western Europe, together with the Middle East, accounted for 25% of respondents. Canada, Central America, and South America made up 11% of respondents. The remaining 25% of the respondents were from the Asia Pacific region.

**Figure 3: Geographic breakdown of respondents**



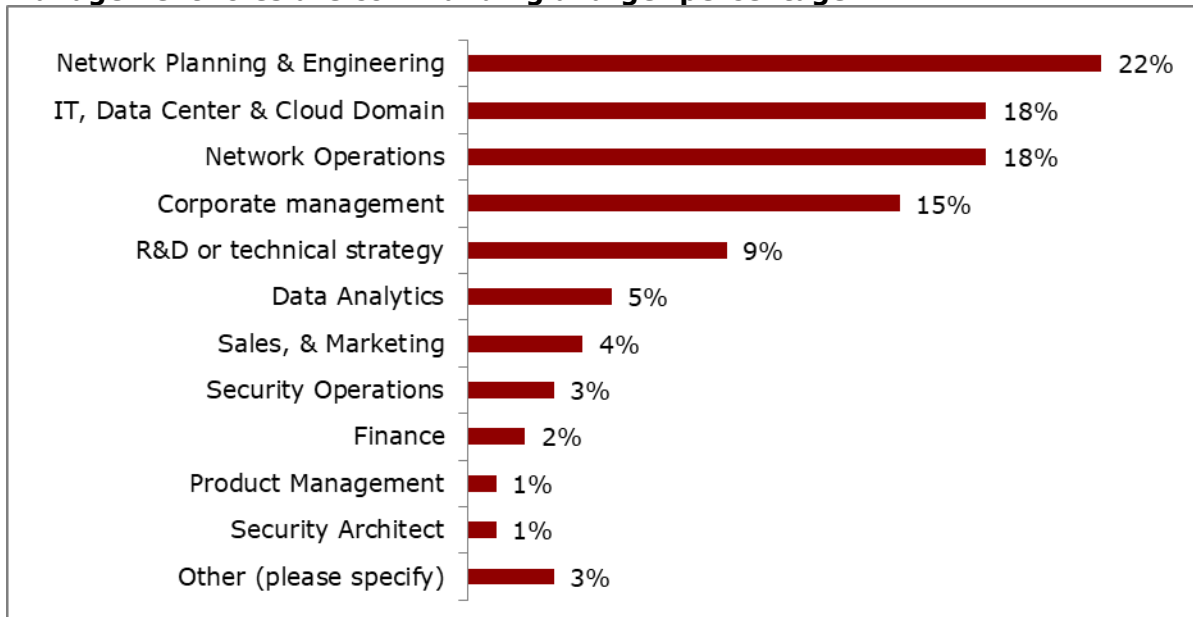
Q: In what region is your organization headquartered? (n=101)  
Source: Heavy Reading, 2023

---

## Job function

Most Heavy Reading surveys have shown over half of respondents in technical networking roles: planning and engineering, R&D, and network operations. In this Edge Computing Survey, the technical networking role percentage falls just under this bar, at 49%. Meanwhile, IT, data center & cloud domain job functions are taking a larger role—in this case, 18%; newer roles, such as data analytics (5%), are gaining more visibility. Management, finance, and marketing account for 22%. Security operations and security architect contribute 4% to the mix, leaving 3% “other.”

**Figure 4: Technical networking roles still dominate the survey, but IT and management roles are commanding a larger percentage**



Q: What is your primary job function? (n=101)

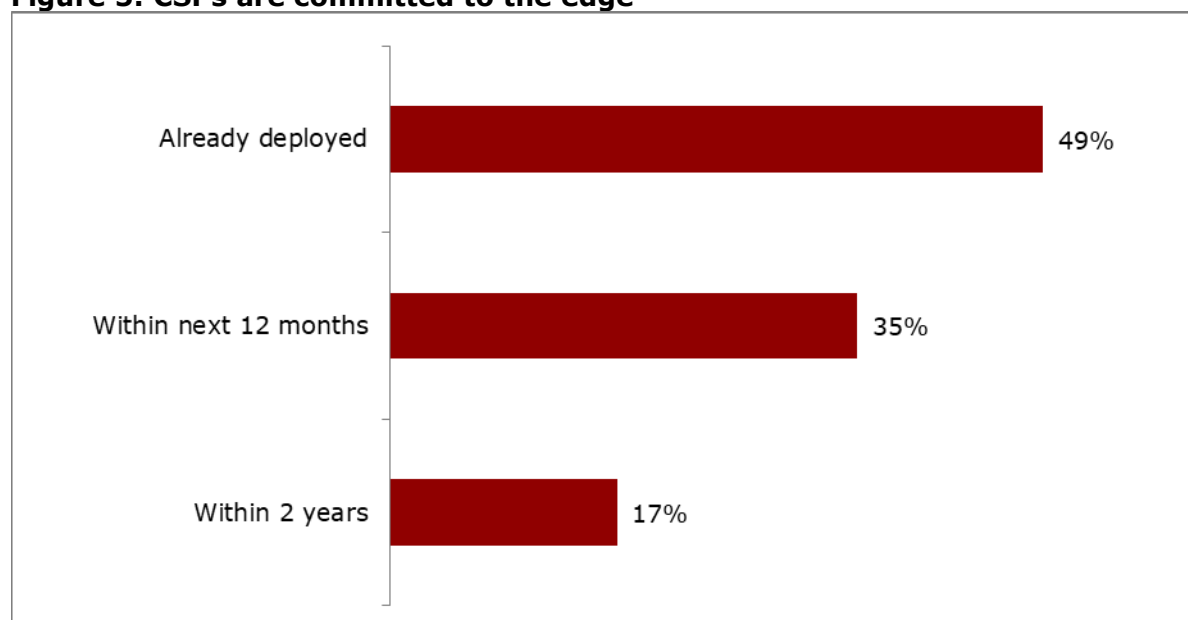
Source: Heavy Reading, 2023

---

## Edge computing production implementations are gaining momentum

All the survey respondents have either implemented edge computing in a production environment or will be doing so within the next 24 months (see **Figure 5**). Heavy Reading excluded any respondents whose plans for edge computing are further in the future. The survey results show that just under half of respondents (49%) already have production deployments; over a third will be implementing within the next 12 months, and 17% plan to deploy within two years. It is gratifying to note that Heavy Reading asked the survey base this same question just under a year ago. At that time, the results showed that 39% had already deployed edge computing, with 32% planning to deploy in a year and 29% planning to deploy within two years. The current survey results demonstrate that not only are CSPs and HCPs implementing edge computing solutions, but they are keeping pace with their deployment plans.

**Figure 5: CSPs are committed to the edge**



Q: When will your company deploy edge computing solutions in a production environment? (n=101)  
Source: Heavy Reading, 2023

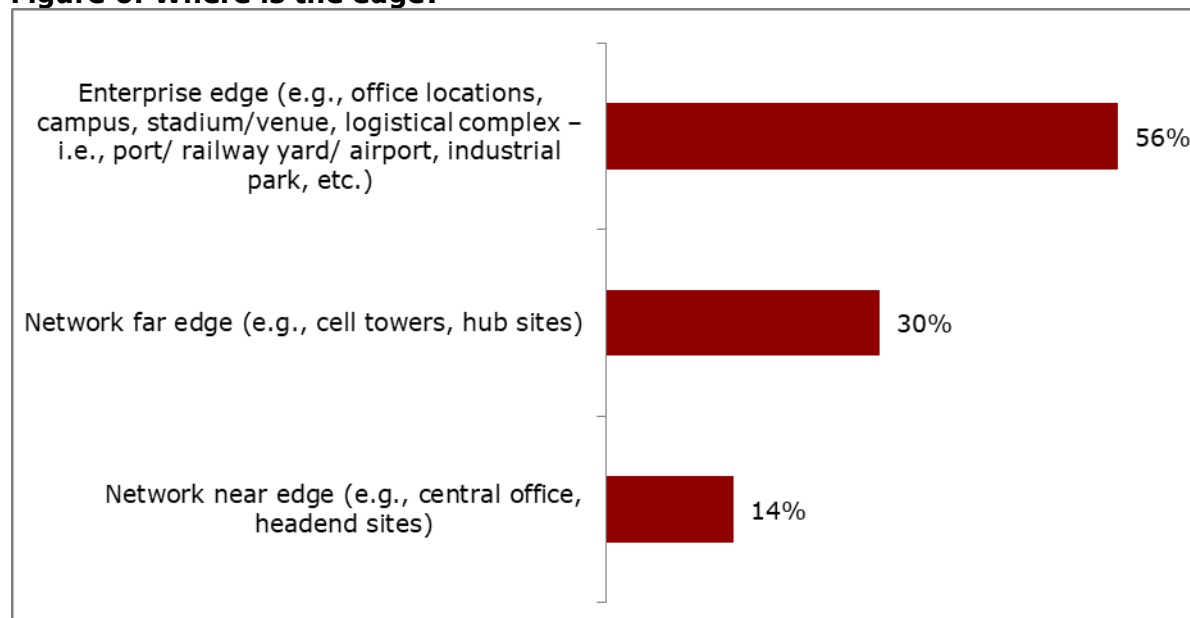
This Heavy Reading survey also showed that next year, respondents will grow their deployments, as measured by the number of sites, by 10–25%, on average. For those planning to deploy over the next 12 months, 25% are looking at 50–100% growth in the number of sites, and 6% are expecting 100–200% growth in the number of sites.

## Where on the network continuum are edge deployments taking place?

Edge computing means different things to different people. When Heavy Reading asked survey respondents where they felt edge platforms should be deployed, the majority (56%) pointed to the enterprise edge as the most appropriate location for deployment. This suggests that those surveyed are largely focused on reducing latency, improving overall application performance, and/or supporting massive IoT implementation. Meanwhile, 30% pointed to the network far edge as their preferred location, and the remaining 14% suggested the network near edge would meet their needs for edge deployments.

These percentages show a shift in how service providers think about their edge deployments. When edge computing implementations first began to emerge three years ago, the bulk of CSP implementations focused on operational efficiencies. The latency targets were about 20ms round-trip delay, satisfactory for most use cases at the time but not sufficient for enabling augmented reality/virtual reality (AR/VR)-enabled applications for either the consumer or enterprise. Instead, the inter-related use cases were intended to provide a relief valve for network pressure during times of traffic stress; keep more traffic local to reduce traffic hair-pinning; and reduce traffic volume at the network core. These use cases can be justified with a straightforward ROI analysis. On the other hand, as deployments reach further toward the edge of the network, implementation costs increase along with the number of edge nodes, and the spread creates end-to-end service management challenges. Far edge and enterprise edge implementations must be justified by new revenue streams. It is both interesting and encouraging, therefore, that survey respondents see the enterprise edge assuming dominance in edge deployments.

**Figure 6: Where is the edge?**



Q: Where does a managed edge platform make the most sense? (n=99)

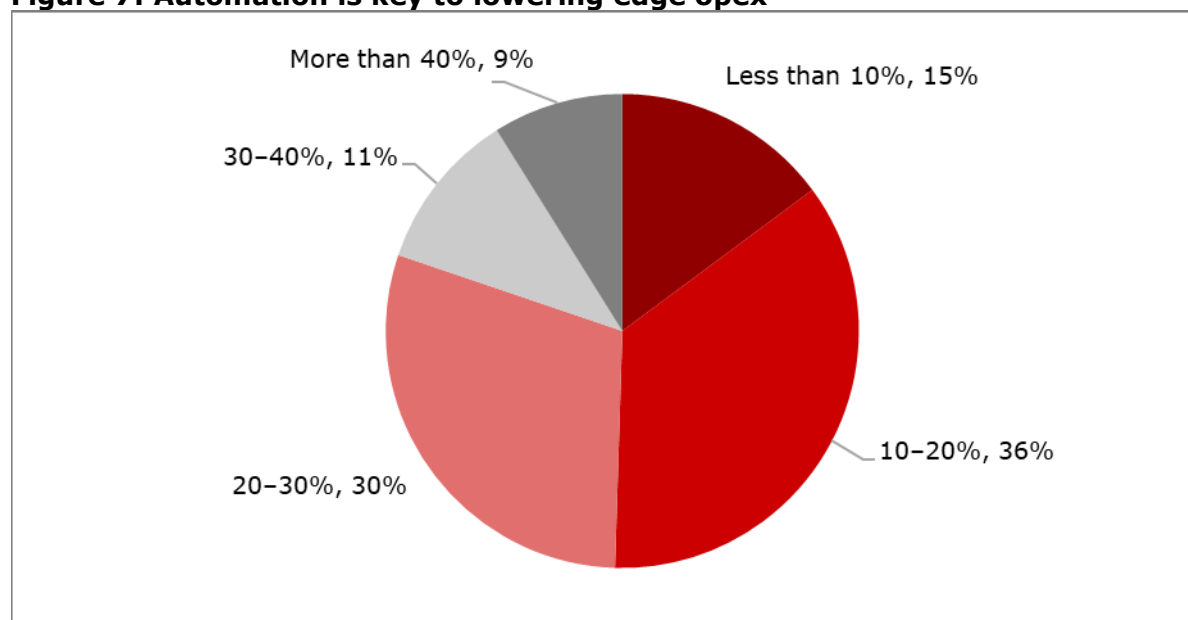
Source: Heavy Reading, 2023

---

## Leveraging automation

This survey makes clear that CSPs rely on automation to lower the overall costs of managing 5G networks for edge use cases (see **Figure 7**). Every stage of the network lifecycle has already benefited from automation, and its impact will continue to grow with the adoption of predictive artificial intelligence (AI), the use of continuous integration/continuous delivery (CI/CD) development and software testing techniques, the use of cloud native computing, the adoption of containerized network functions (CNFs), and more. Growth in the number of edge computing locations cannot be supported without the ability to automate deployment, update software, and scale locations up and down on demand without operator intervention.

**Figure 7: Automation is key to lowering edge opex**



Q: What percent will telecom network automation help you reduce the operational cost of managing your 5G networks for edge use cases? (n=101)

Source: Heavy Reading, 2023

## Deploying edge computing in a hybrid cloud environment

The survey asked respondents what motivated them to move to a hybrid telco/edge cloud platform for their edge implementation. The responses, for the most part, are variations on a theme, "It is faster and easier." The top response, chosen by half of the respondents, was "flexibility to scale deployments quickly and as needed." This is the top value proposition regardless of region, edge deployment status, or CSP revenue tier.

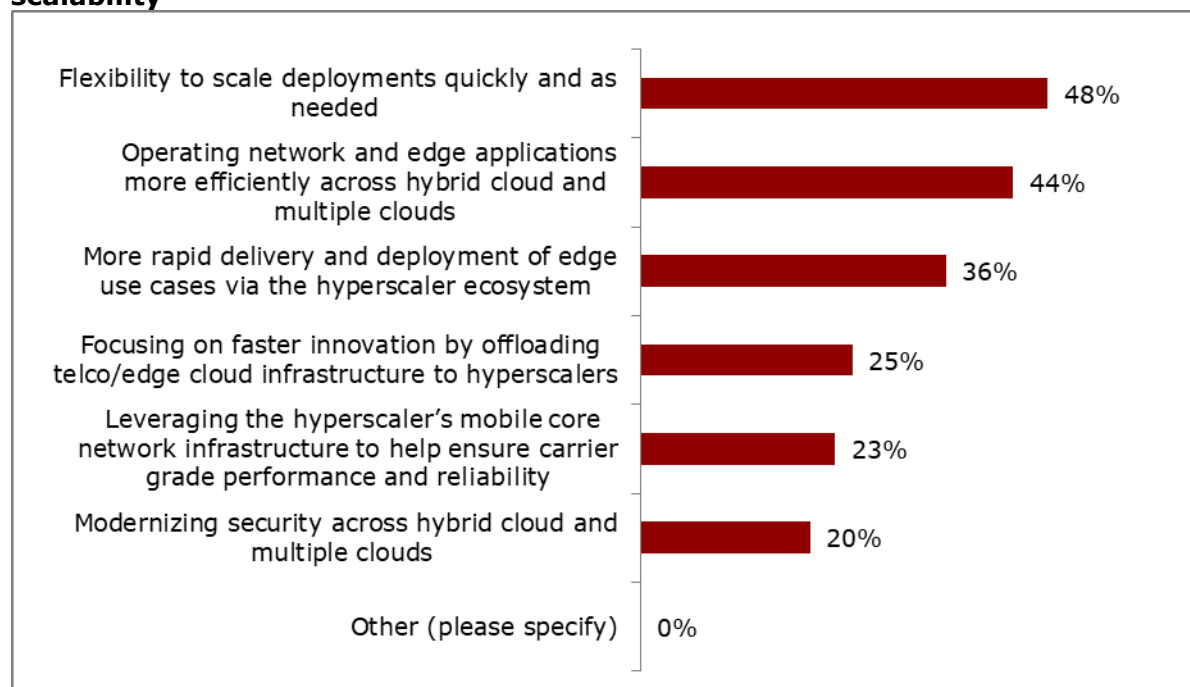
However, there is statistically relevant variation in some of the other reasons when looking only at survey respondents from CSPs with more than \$5bn in annual revenue. For these largest of carriers, the second most popular response, "Operating network and edge applications more efficiently across hybrid cloud and multiple clouds," falls by 9 percentage points to 35%, indicating that end-to-end management is not as much of a benefit as are speed and flexibility of deployment. Even more notably, the score for "Modernizing security across hybrid cloud and multiple clouds" (already at the bottom) falls an additional 10



percentage points to 10% for the more than \$5bn crowd, indicating that this is not so much a value proposition as it is a concern.

Interestingly enough, “Leveraging the hyperscaler’s mobile core network infrastructure to help ensure carrier-grade performance and reliability” does not vary in the percentage of responses from the largest carriers, nor does “... offloading telco/cloud infrastructure.” The hyperscalers have clearly come a long way in overcoming uncertainty among all CSPs, including the largest, about the HCPs’ ability to deliver carrier-grade performance and be entrusted with their telco/cloud infrastructure.

**Figure 8: Value propositions for hybrid cloud emphasize speed, agility, and scalability**



Q: What are the top value propositions that would cause you to consider a telco/edge cloud platform? (Select two) (n=101)

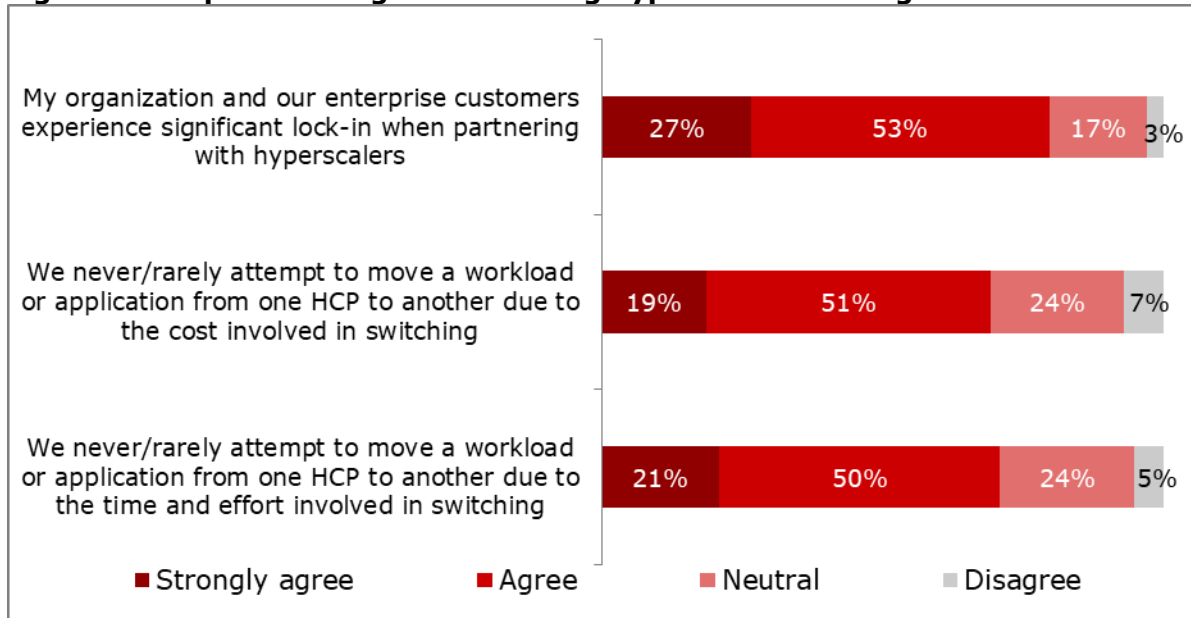
Source: Heavy Reading, 2023

There is, however, one aspect of working with the HCPs that causes concern among the CSPs. That concern is lock-in and the difficulty of moving workloads from one HCP to another. A significant majority of respondents, between 70% and 80%, agree or strongly agree with all three statements in **Figure 9**.

Looking at only the largest CSPs, the results reflect even more potential disgruntlement: 38% *strongly agree* with the last statement—that they experience significant lock-in—compared with 22% for all other revenue segments.

As a result, once CSPs move an application or workload to a specific HCP, the difficulty of moving the workload motivates the CSP to leave it there. Although CSPs partner with multiple HCPs, they rarely deploy the same workload across multiple clouds. In other words, the cloud providers are used for specific workloads. The term “multicloud” might as well be called “one workload, one cloud.”

**Figure 9: Respondents agree: switching hyperscalers is tough**



Q: To what extent do you agree with the following statements? (n=98)  
 Source: Heavy Reading, 2023

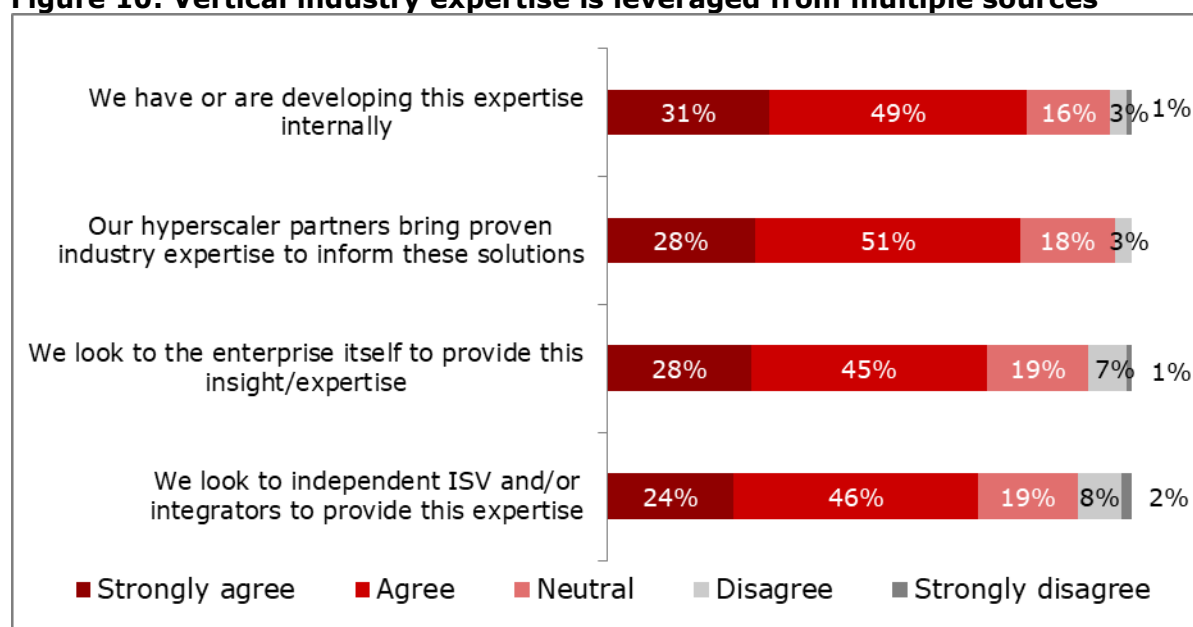
### Building or borrowing vertical industry expertise?

Until now, the market for edge services has been characterized by solutions that address specific use cases for specific enterprises. Those solutions have been highly customized and hard to replicate, even within the same vertical. Are CSPs willing to tackle the challenge of amassing enterprise expertise, or are they looking to their edge ecosystem to offload this resource-intensive task (see **Figure 10**)? At first glance, it appears that the CSPs are neutral about where this expertise originates, whether from within their own organizations or from independent software vendor (ISV) partners, cloud partners, or the enterprise itself. Between 70% and 80% of respondents agree or strongly agree with each choice. A 10 percentage point difference is very tight. However, developing expertise internally does edge out the other choices as the most popular response.

When examining only the largest CSPs (over \$5bn in annual revenue), the responses are more highly differentiated. Looking to the enterprise itself for expertise decreases as a strategy as looking to the hyperscalers and the CSP internal organization increases. For the hyperscaler response, the combined number of “strongly agree” and “agree” changes by only 1 percentage point, but approximately 10 percentage points shift from “agree” to “strongly agree.” This is not the case with building expertise internally, where the “strongly agree” percentage increases by 9 percentage points for the largest carriers while the “agree” percentage remains virtually unchanged. The takeaways are clear:

- All methods are viable for all CSPs.
- But smaller carriers are more likely to look to their ecosystem partners for expertise.
- And Tier 1 carriers are focused on building expertise internally.

**Figure 10: Vertical industry expertise is leveraged from multiple sources**



Q: To what extent do you agree with the following statements on looking for industry expertise when you deploy an edge solution? (n=100)

Source: Heavy Reading, 2023


The survey contained a follow-up question regarding the use of AI, querying respondents regarding primary owners of training data, AI algorithms, and model training for network AI applications. The response choices were the CSPs, the hyperscalers, and third-party AI solutions providers. A full 70% of Tier 1 companies favored themselves, with a response of “CSPs.” Only 10% favored hyperscalers. For all smaller-revenue telcos, the “CSP” response plummeted to 29%, and the hyperscaler response was 37%. This reinforces Heavy Reading’s observation that the smaller the CSP, the more likely it is to turn to the hyperscalers for help with issues foreign to their core expertise, whereas Tier 1 providers are building expertise in-house.

## CSPs see many drivers for partnering with the hyperscalers

The hyperscalers bring a wide range of benefits to the edge. How these benefits are valued varies with the size of the CSP (see **Figure 11**). The largest CSPs are much more interested in DevOps, AI/ML, and time-to-market than their smaller counterparts. Everything below these top three drivers is of more interest to smaller carriers, particularly cost control/total cost of ownership (TCO), the number one driver for the smaller telcos. “Talent acquisition” shows the most dramatic difference, with 36% of the smaller CSPs citing it as a significant driver, placing it third among drivers, whereas only 13% of larger CSPs, 23 percentage points fewer, see it as a driver, placing it last among drivers.

CSPs already have an uphill climb to acquire talent in areas such as CI/CD software development, data sciences, and AI, where they are competing with attractive (to the general workforce) organizations such as the hyperscalers. For smaller CSPs, this hill can be more like a wall, compelling them to look for partners to fill out their portfolios of skills.

**Figure 11: Drivers for partnering with hyperscalers vary depending on size of CSP**



Q: What considerations drive your preference for using hyperscalers? (Select top three) (n=101)

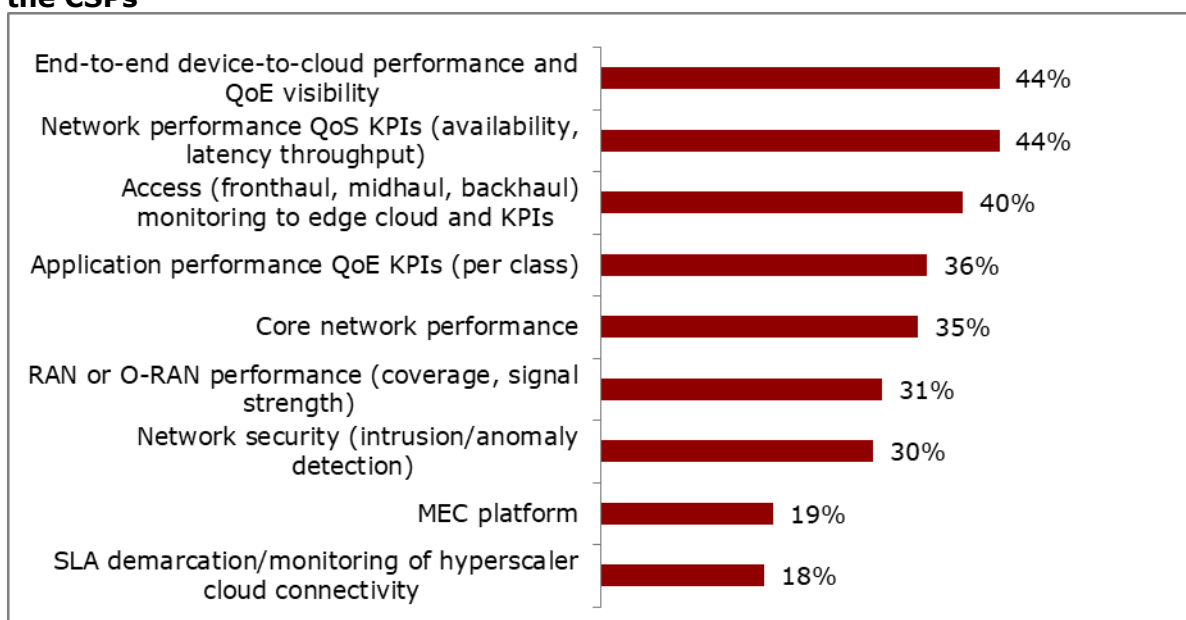
Source: Heavy Reading, 2023

Edge computing infrastructure and service management in a hybrid cloud environment can be difficult and complex. The CSPs are starting out internally with multiple (in some cases, dozens of) operations support systems (OSS), inventory systems, and data repositories, which may or may not affect the edge computing service. End-to-end service management across the public or private cloud then requires insights into device and application performance that is outside the control of the CSP. Managing MEC services demands tight coordination between the enterprise and hyperscaler to establish not only the key performance indicators (KPIs), but also how these KPIs will be enforced, measured, and reported.

When Heavy Reading asked survey respondents, “What performance monitoring/assurance capability will your organization need for edge computing infrastructure and services?” the top response for most demographic segments was, “End-to-end device-to-cloud

performance and QoE [quality of experience] visibility.” It is this holistic, bird’s-eye view of the end-to-end service that is most difficult to capture, and most of the other capabilities listed are prerequisites for that. All the other capabilities vary in terms of survey responses according to region, type of CSP, and edge deployment status. Each capability is important to specific subsets of the survey. In general, however, survey results show that, after end-to-end service management, the CSPs are concerned with the overall network, then the access network, then core, then the radio access network (RAN). Applications performance and QoE rank near the top for most demographic segments. MEC platform and service-level agreement (SLA) rank near the bottom because not all CSPs are managing their own MEC platforms, and even fewer offer SLAs for MEC services, nor are they being offered SLAs from their HCP partners.

**Figure 12: All performance monitoring & assurance capabilities look attractive to the CSPs**



Q: What performance monitoring/assurance capability will your organization need for edge computing infrastructure and services? (Select top three) (n=101)

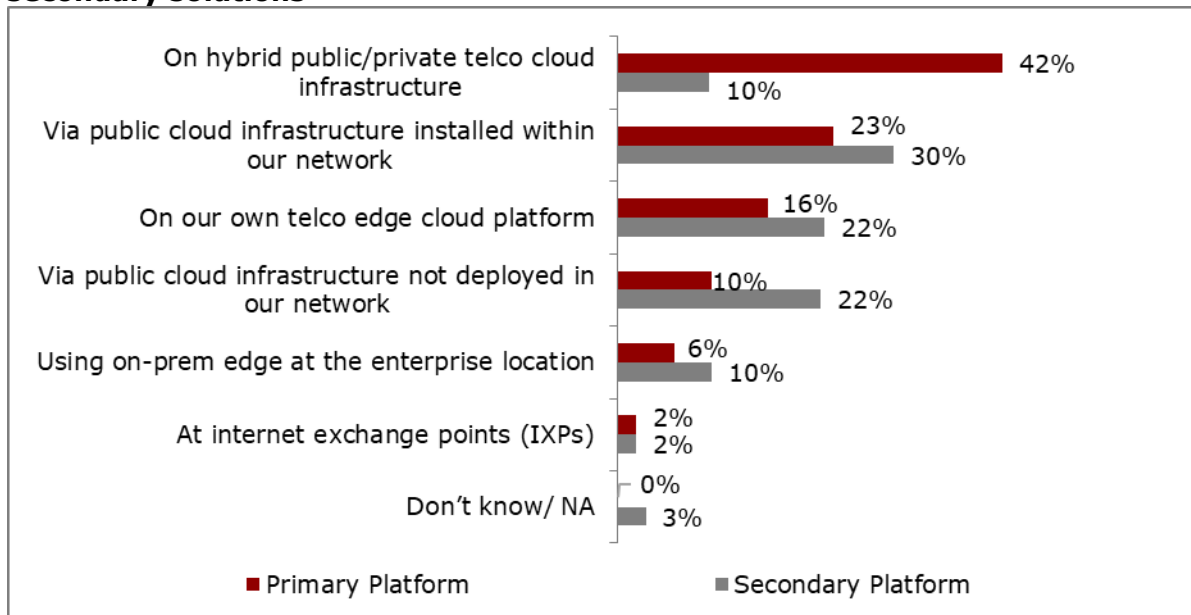
Source: Heavy Reading, 2023

### The MEC platform landscape is changing

As was mentioned in the previous section, CSPs have ranked MEC platform management far down the list of needed monitoring capabilities, and it is clear why in the following exhibit (**Figure 13**). The primary edge platform that CSPs rely on for edge services uses hybrid public/private telco cloud infrastructure. The performance of these platforms has been solid enough that they are not seen as negatively impacting overall edge service performance and efficiency. The world is likely to see this type of platform continue to grow in terms of adoption.

Heavy Reading was more surprised to see that 16% of CSPs cited their own telco edge cloud platform as their primary platform, and 22% cited it as a preferred secondary platform. It is likely these percentages will decrease as CSPs become intolerant of the time, cost, and complexity of deploying and managing edge platforms at scale. Heavy Reading expects them to turn more to the hyperscalers for platform solutions. Similarly, although internet exchange points (IXPs) pulled in only 2% of survey responses for both primary and secondary platforms, this platform solution is expected to grow as CSPs realize they can leverage the availability of IXP locations with edge solutions that are more cost-effective than many of the other platform locations.

**Figure 13: CSPs largely agree on their primary edge platform but use multiple secondary solutions**




Q: What is the primary type of platform your organization is offering with edge services? (n=98)  
 Source: Heavy Reading, 2023

### The big picture of CSP deployments in the HCP cloud

In the final survey question, Heavy Reading stepped back from the edge to query respondents about what elements of their network use cases they were looking to deploy on HCP cloud infrastructure, as well as when (see **Figure 14**). The magnitude of the responses was somewhat surprising.

More than 75% of respondents said that they will deploy all these use cases, at scale, on the HCP cloud within the next two years. That makes a certain amount of sense in the case of IT-adjacent use cases such as OSS and business support systems (BSS). It is more surprising for cloud RAN and the packet core, particularly given the CSPs' continuing concerns about HCP lock-in. However, that appears to be balanced by hopes for a lower TCO, a more agile DevOps environment, and a more rapid feature introduction. That element of shared risk may be appealing to the risk-averse CSPs. Does Heavy Reading believe this will unfold in the way that respondents predict? Answer: CSP respondents may have indicated their intent, but the rollout schedule will likely stretch over a more extended period.

**Figure 14: Multiple CSP use cases are in dead heat as drivers of HCP cloud deployments**



Q: What is the timeframe on deploying HCP cloud infrastructure at scale? (n=98)

Source: Heavy Reading, 2023

## The view from the edge

One of Heavy Reading's key goals in conducting this survey was to investigate how the CSPs and HCPs are working together on the edge to get an idea of the trajectory these alliances are taking. From a MEC service perspective, it is clear that CSPs are increasingly comfortable relying on the HCPs in a variety of areas, including edge platforms, vertical industry expertise, DevOps, AI/ML and automation software, increased scalability, and accelerated service delivery. Earlier concerns about performance and reliability have largely subsided, but the issue of HCP lock-in remains. Heavy Reading believes these "frenemies" are maturing into trusted partners. The next two years might not see the degree of CSP movement of critical workloads to the public or hybrid cloud, as is suggested in **Figure 14**. But HCPs will certainly exert influence and assist with solutions for all aspects of the CSP MEC value chain.