HEAVY READING

WHITE PAPER

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

Strategies for Connecting the Edge:

2019 Heavy Reading Survey

A Heavy Reading white paper produced for Arista, CoreSite, Fujitsu, and Infinera

ARISTA CORESITE FUJITSU Infinera

AUTHOR: STERLING PERRIN, PRINCIPAL ANALYST, HEAVY READING

EXECUTIVE SUMMARY

For several years, data center trends have pointed to a shift from a few highly centralized mega-facilities to more distributed data centers populating metros. Edge computing is a new evolution of the distributed processing and storage trend that brings high-performance compute, storage, and networking resources closer to users and devices than ever before. The goals of this approach include lowering the cost of data transport, decreasing latency, and increasing locality.

This trend runs directly counter to the historical approach of housing compute and storage in a small number of massive (or hyperscale) data centers placed in often remote locations to minimize the operations costs of running them. Thus, with edge computing, the pendulum is moving from a centralized cloud architecture to a highly distributed one.

Originally thought to be a mobile-only trend, it is now generally understood that edge computing affects all areas of communications – including wireline and wireless applications and connectivity. And, as edge computing redefines the future of data centers, it must also redefine the future of metro network connectivity.

While interest in and hype around edge computing is high, detailed knowledge of the edge is remarkably low. At this nascent stage, the most fundamental questions are being asked: "What exactly is the edge?"; "Where is the edge?"; and "What use cases will drive edge deployments?"

This research project was commissioned to answer these fundamental questions, and many more, with a focus on the evolution of network connectivity to meet edge computing needs. Conducted by Heavy Reading and co-sponsored by Arista, CoreSite, Fujitsu, and Infinera, the survey-based study explores the impacts of edge computing on network connectivity (including data center interconnect [DCI]), capacity demands, technical and architectural requirements, and deployment models.

Using multiple databases within the Informa group, Heavy Reading surveyed key constituents in edge computing in June and July. These constituents included fixed, mobile, and converged network operators, colocation and cloud content providers, and enterprises. (Additional details on the demographics of the survey group are included as an **Appendix** at the end of this report.) The report analysis and key findings are based on the global survey data.

The following sections detail Heavy Reading's key findings from this survey.

Key Findings

Drivers and Use Cases

At this early stage, service providers and enterprises understand the importance of edge computing to their business. In Heavy Reading's survey, 80% of respondents reported that edge computing is at least important for their business, with 20% of the group stating that edge computing is critical and that their business cannot succeed without it. Just 3% of those surveyed believe that edge computing is not important at all.

For telecom operators, the top use cases and business drivers for edge computing are 5G, Internet of Things (IoT), ultra-reliable/ultra-low latency use cases, and high-performance content delivery. Telecom operators overwhelmingly link edge computing with their 5G plans, and we note that 5G will also form the basis of their plans for ultra-reliable low latency communication (URLLC) and IoT applications.

For enterprises surveyed, the top use cases and business drivers are artificial intelligence (AI) applications, ultra-reliable/ultra-low latency use cases, and IoT. Enterprises overwhelmingly link edge computing with AI use cases, whereas telecom operators place AI nearly at the bottom of their edge drivers list. Telecom operators prioritize edge for 5G, but enterprises place 5G on a lower tier.

When asked which applications are most likely to drive their initial edge deployments, both telecom operators and enterprises agreed on industrial or factory automation as the biggest initial application. Industrial automation was selected by 41% of telecom respondents and 55% of enterprises surveyed.

Consistent with the priority placed on edge computing, the expected timelines for adoption are aggressive. All told, 79% of survey respondents expect to have edge compute in their networks in the next 3 years.

For both telecom operators and enterprises surveyed, there was strong agreement on the primary locations for edge computing, namely, multi-tenant data centers (with dense network connectivity), large regional data centers, and corporate/in-house data centers.

Telecom operators and enterprises are also largely in agreement on where edge computing will *not* be housed initially: in street cabinets either at aggregation points or at cell sites. This finding is bad news for those speculating that edge connectivity must be placed at every macro site and small cell in order to deliver on the 5G promise of URLLC and massive machine-type communications (mMTC) applications.

Technology and Architecture

Despite technical immaturity and continuing questions about operations models, network operators clearly expect white box switches and routers to play a major role in their edge computing networks and strategies. According to the survey, 68% of respondents believe that white box switches/routers will be at least important to their overall edge computing plans and strategy, with 21% reporting that white boxes will be critical to success. We also note the packet white box edge scores consistently high across Heavy Reading surveys. The results suggest that, while the challenges are real, the burden remains firmly on the vendor community to overcome those challenges to meet the needs of their customers.

Consistent with overall concerns about data security, Layer 1 encryption/Federal Information Processing Standards (FIPS) certification topped the list of enabling technologies for edge DCI deployments. The edge computing model poses new challenges for security, and concerns are clearly reflected in the survey result. Still, Layer 1 encryption will only address in-flight data and, thus, is only part of a security solution.

Single, multi-rate transponders, IP over DWDM (IPoDWDM), and automation features ranked second, third, and fourth, respectively, as enabling technology priorities for edge DCI deployments. These prioritized functions are all aimed at saving capex and opex. IPoDWDM, for example, saves capex by eliminating transponder shelves and placing long reach optics directly in the routers themselves. Automation is a major theme across networking overall and saves on operations by reducing/eliminating manual interactions and speeding up processes. Single multi-rate transponders, meanwhile, hold the potential to save on both opex and capex by reducing card inventory and simplifying ordering.

The emergence of 400 Gbit/s pluggables, being standardized by the Optical Internet Forum (OIF) as 400ZR, have garnered significant industry attention in 2019, but the 400ZR promise has not yet resonated with most telecom operators – at least for edge DCI applications. According to the survey, just 9% of telecom operators knowledgeable about 400ZR expect to deploy modules for edge DCI, while 31% expect no 400ZR to be deployed. At 60%, the majority of respondents said their organizations are still deciding what to do. One plausible explanation for the finding is that 400 Gbit/s is simply bandwidth overkill for the edge bandwidth requirements. Beyond capacity, another potential pitfall for 400ZR is the short reach, point-to-point implementation proposed. Telecom operators with ROADM-based metros will struggle to find a fit in their networks.

Colocation Strategies

In Heavy Reading's survey, 81% of respondents reported that network-dense colocation was at least important for their edge strategy, with 17% of the group reporting that colocation is critical and that their edge strategy cannot succeed without it. The finding is significant because it indicates that choosing an edge location is not simply a question of "where." Equally significant (or perhaps more significant) is the question of "who's there." Inter-connectivity was a primary factor in peering and content delivery networks (CDNs), and results suggest that edge computing may not be so different in this regard.

Primary colocation benefits differ for telecom operators and enterprises. For telecom operators, the primary colocation benefit is the ability to access the edge where it is needed, followed by the ability to connect to major network service providers and the ability to connect to major content/cloud providers. For enterprises, however, the primary benefit is the time-to-market advantage compared to building their own edge locations. The ability to connect to major network service providers is also important to enterprises (ranked second).

Security concerns tops the list of inhibitors to colocation at the edge, despite the money and efforts that data center operators invest to ensure their facilities and data are secure. In Heavy Reading's survey, 56% of respondents cited security as a primary inhibitor, and it was the top choice among both telecom operators and enterprises. In addition to security, lack of control over infrastructure and applications is a major inhibitor (selected by 47% of the group), followed by high costs of colocation (selected by 43%).

WHAT IS THE EDGE?

The 2018 *State of the Edge* report, produced by several edge companies, provides a succinct definition of edge computing as placing high-performance compute, storage, and network resources as close as possible to the end users and devices. The goals of this approach include lowering the cost of data transport, decreasing latency, and increasing locality. This trend runs directly counter to the historical approach of housing compute and storage in a small number of massive (or hyperscale) data centers placed in often remote locations to minimize the operations costs of running them. Thus, with edge computing, the pendulum is moving from a centralized cloud architecture to a highly distributed one.

In thinking about the edge, there are a couple of points to keep in mind:

- The "edge" is a location, not a thing.
- The specific edge location will vary depending on the use cases/applications as well as edge owner (enterprise, telecom operator, cloud operator, etc.). There will not be a single edge, but rather, multiple edges operating throughout the network.
- Edge computing and multi-access edge compute (MEC) are not synonymous terms. Edge computing is a broad industry concept while MEC, which started out as "mobile edge compute," is a specific ETSI-defined standardized edge computing architecture. While ETSI is a contributor to edge computing, it is not the only one. Other groups active in edge compute include the Open Networking Foundation (ONF), Telecom Infra Project (TIP), IEEE, and 3GPP, among many others. Throughout this paper, we discuss edge computing in the broader sense.

At this early stage, service providers and enterprises understand the importance of edge computing to their business. In Heavy Reading's survey, 80% of respondents reported that edge computing is at least important for their business, with 20% of the group stating that edge computing is critical and that their business cannot succeed without it (**Figure 1**). Just 3% of those surveyed believe that edge computing is not important at all.

Figure 1: Importance of Edge Computing for Your Business

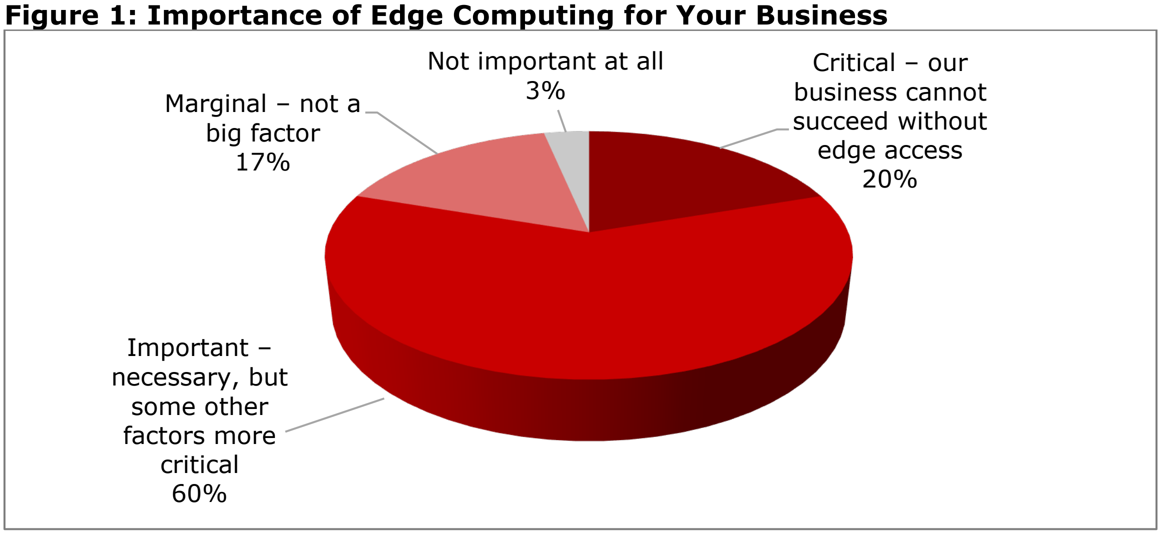

N=91
Source: Heavy Reading

DRIVERS AND USE CASES

As shown in the data above, the need for edge computing is widespread across telecom operators and enterprises, but their views differ on the use cases and applications that are driving that need. For telecom operators, the top use cases/drivers are 5G (selected by 67% of telecom respondents), IoT (selected by 63%), ultra-reliable, ultra-low latency applications (selected by 55%), and high-performance content delivery (selected by 47%). For enterprises, the top use cases/drivers are AI applications (selected by 61% of enterprise respondents), ultra-reliable, ultra-low latency applications (selected by 57%), and IoT (selected by 48%). See **Figure 2** for the full results broken out by telecom and enterprise respondents.

Figure 2: Leading Use Cases/Business Drivers for Edge Compute (Telecom vs. Enterprise)

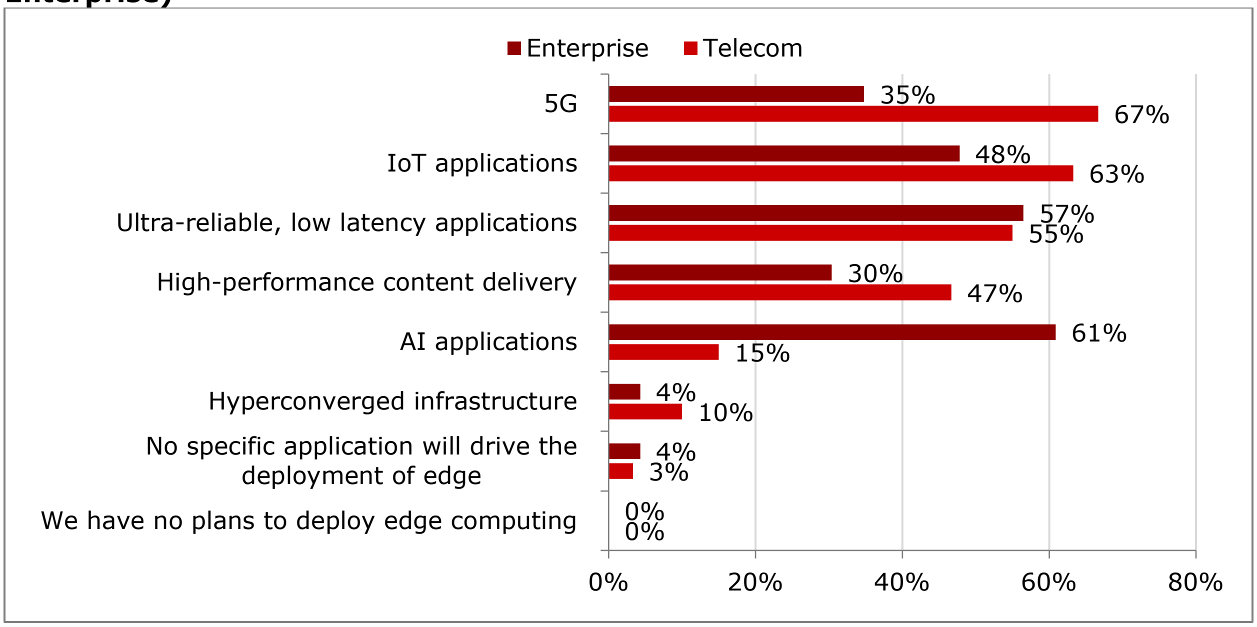

N=60 telecom, 23 enterprise
Source: Heavy Reading

In comparing the two groups, there are two significant differences in drivers. First, telecom operators overwhelmingly link edge computing with their 5G plans. 5G will also form the basis of their plans for URLLC and many IoT applications. Second, enterprises overwhelmingly link edge computing with AI applications, whereas telecom operators place AI nearly at the bottom of their drivers list.

When asked which applications are most likely to drive their initial edge deployments, both telecom operators and enterprises agreed on industrial or factory automation as the biggest initial driver. Industrial automation was selected by 41% of telecom respondents and 55% of enterprises (see **Figure 3** below).

But beyond factory automation, telecom and enterprise views differ. Among telecom operators, augmented reality/virtual reality (AR/VR) is also expected to be a primary initial driver (tied with industrial automation and selected by 41% of telecom respondents), closely followed by safe cities (selected by 39%).

Among enterprises, smart cities with enhanced/localized experiences (selected by 41% of enterprise respondents) is also expected to drive initial edge deployments, though to a lesser extent than industrial automation. Beyond these two applications, all others ranked as secondary or ancillary for enterprise respondents.

Figure 3: Applications Most Likely to Drive Initial Edge Deployments

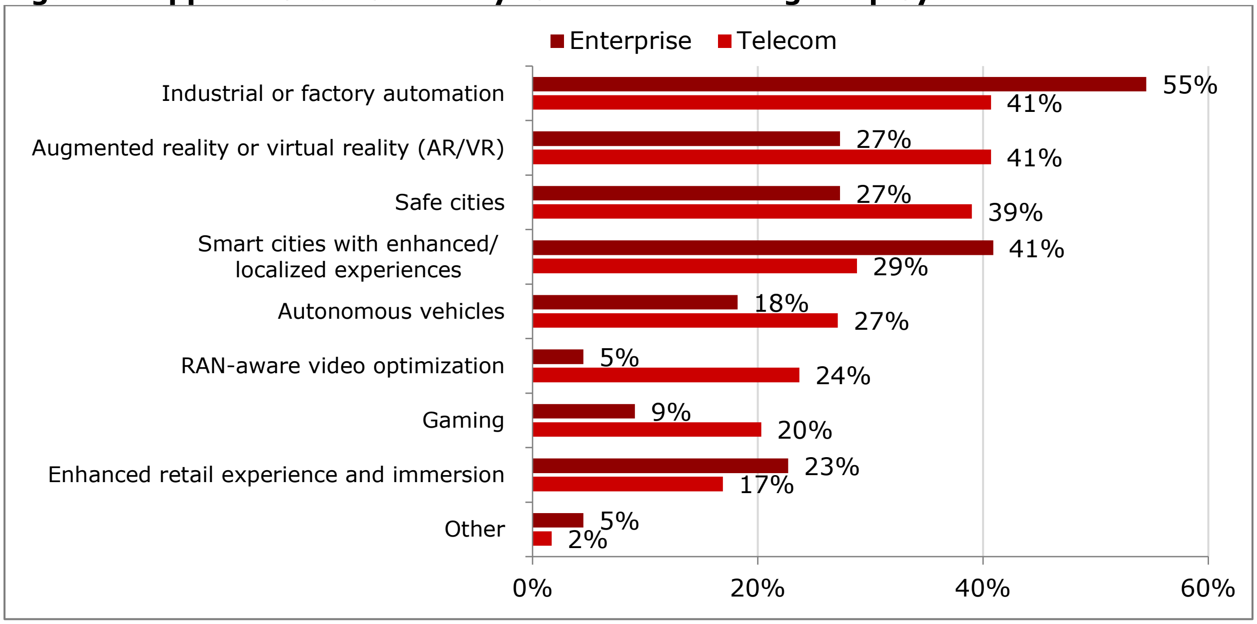

N=59 telecom, 22 enterprise

Source: Heavy Reading

Consistent with the priority placed on edge computing, the expected timelines for adoption are aggressive. While just 17% report using edge computing today, more than one-third (36%) expect to have edge computing deployed by the end of 2020. An additional 26% expect to have edge computing in their networks by the end of 2022 (see **Figure 4** below). All told, 79% of respondents expect to have edge computing in their networks in the next 3 years.

Figure 4: Expected Timelines for Edge Computing

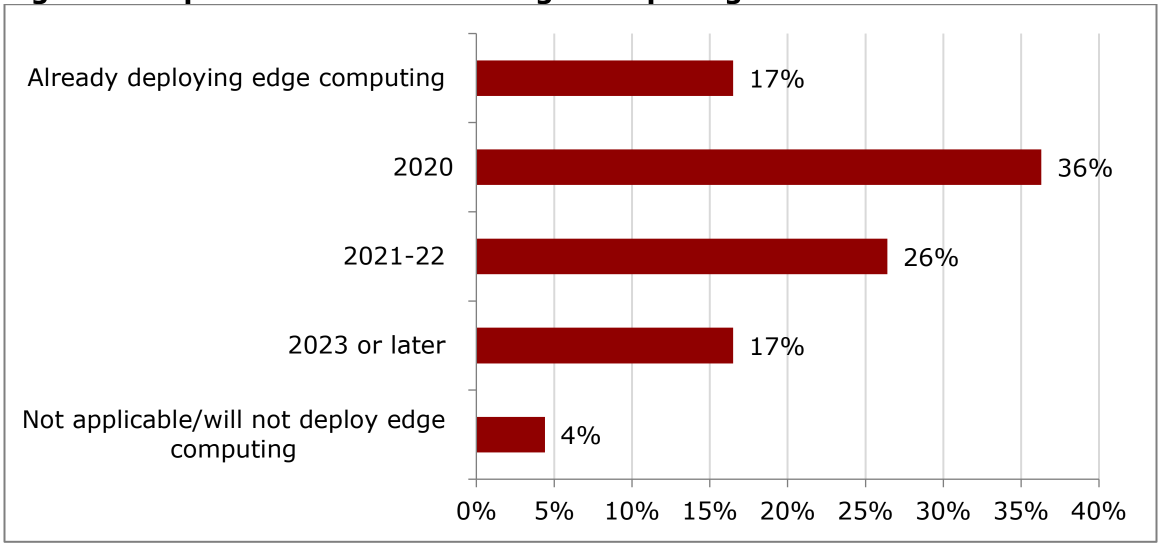

N=91
Source: Heavy Reading

In addition to defining the edge, "where is the edge?" is another industry question of intense discussion and debate. To get a better sense of industry views, Heavy Reading asked survey respondents to select the locations for their initial edge computing deployments.

For the group as a whole, there was strong agreement on the primary locations for edge computing, namely, multi-tenant data centers (with dense network connectivity), large regional data centers, and corporate/in-house data centers (see **Figure 5**).

Figure 5: Locations for Initial Edge Computing Deployments

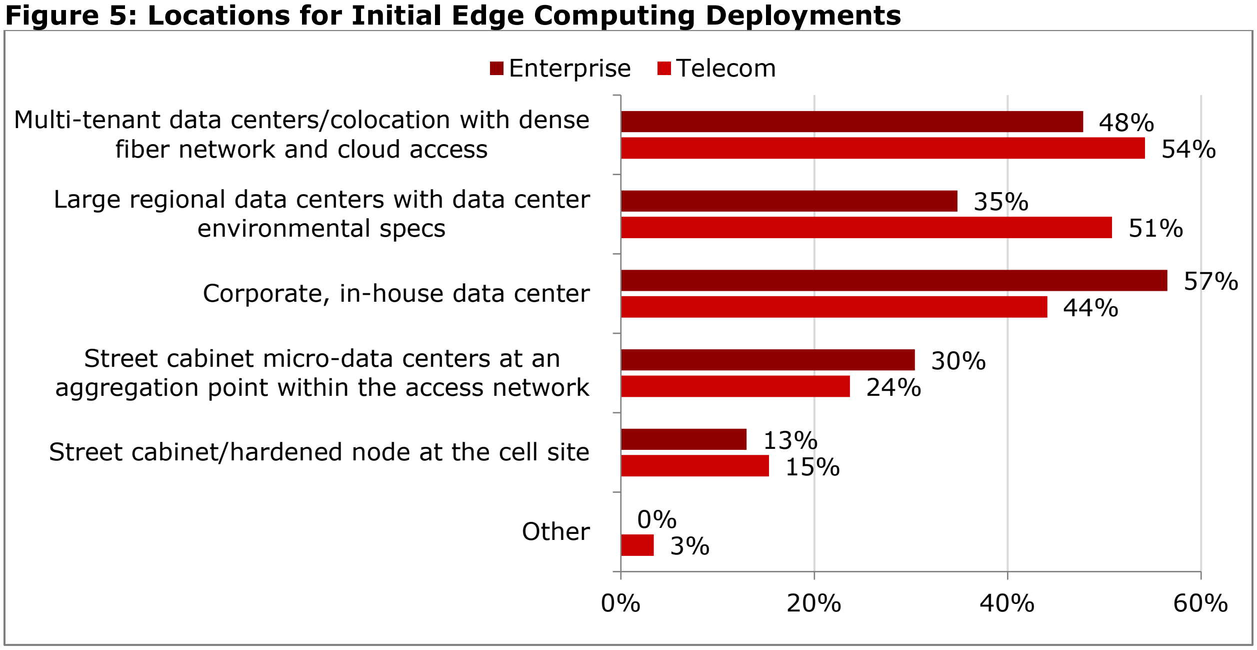

N=59 telecom, 23 enterprise
Source: Heavy Reading

However, there are some differences in how telecom operators and enterprises prioritize these three locations. For telecom operators, multi-tenant data centers are expected to be the top location, followed by large regional data centers and then in-house locations. The majority of enterprises, however, selected corporate/in-house locations, followed by multi-tenant data centers and then regional data centers.

Interestingly, telecom operators and enterprises are also largely in agreement on where edge computing will **not** be housed initially: in street cabinets either at aggregation points or at cell sites. This finding is bad news for those speculating that edge connectivity must be placed at every macro site and small cell in order to deliver on the 5G promise of URLLC and mMTC applications. Service providers and vendors closest to the market have been cautious for some time about how far the edge will spread, and the survey results indicate this cautious viewpoint is becoming more widespread. Still, we note that the question asked only about initial deployments and, therefore, should not be used to extrapolate on longer-term trends – such as those 3-5 years out.

TECHNOLOGY AND ARCHITECTURE

This section addresses enabling technologies for edge connectivity. Heavy Reading asked respondents to select the most important network equipment features and functions for edge DCI over the next 3 years. From a list of features/functions provided, **Figure 6** below shows the results in ranked order from highest to lowest priority, along with overall scores for each. Items ranked first were given a higher weight in the scoring.

Consistent with overall concerns about data security, Layer 1 encryption/FIPS certification topped the list, followed by single multi-rate transponders, IPoDWDM, and automation features. With the exception of security, the prioritized functions are aimed at saving capex and opex. IPoDWDM, for example, saves capex by eliminating transponder shelves and placing long reach optics directly in the routers themselves. Automation is a major theme across networking overall and saves on operations by reducing/eliminating manual interactions and speeding up processes.

Single multi-rate transponders, meanwhile, hold the potential to save on both opex and capex by reducing card inventory and simplifying ordering. However, Heavy Reading cautions that relative pricing was not included in the survey question and that cost premiums on single-rate transponders could erode some of the potential benefits.

Finally, telecom operators and enterprise respondents were largely in agreement on function priorities, particularly regarding Layer 1 encryption, single-rate transponder preferences, and automation. The enterprise respondents, however, placed a higher priority on network management systems and a lower priority on IPoDWDM.

Figure 6: Most Important Network Equipment Features/Functions for Edge DCI

Item	Overall Rank	Score
Layer 1 encryption/FIPS certification	1	401
Single multi-rate (200-600G) transponders (access/metro/LH)	2	394
Integrated optics on routers (IPoDWDM)	3	377
Automation features like zero-touch provisioning, streaming telemetry, etc.	4	376
Network management system, tools, and applications	5	369
CD/CDC ROADM-based flex grid line system architecture	6	299
75 GHz and higher spaced fixed grid line system	7	260
L-band capacity (using L-band transponders and line systems)	8	242

N=90
Score is a weighted calculation, where items ranked first are given a higher weight
Source: Heavy Reading

White box switches and routers are expected to play a major role in edge computing. According to the survey, 68% of respondents believe that white box switches/routers will be at least important to their overall edge computing plans and strategy, with 21% reporting that white boxes will be critical to success (**Figure 7**). An additional 29% see a niche role for white boxes at the edge, while just 3% of respondents see white boxes as inconsequential.

Figure 7: Importance of White Box Switches/Routers for Edge Computing

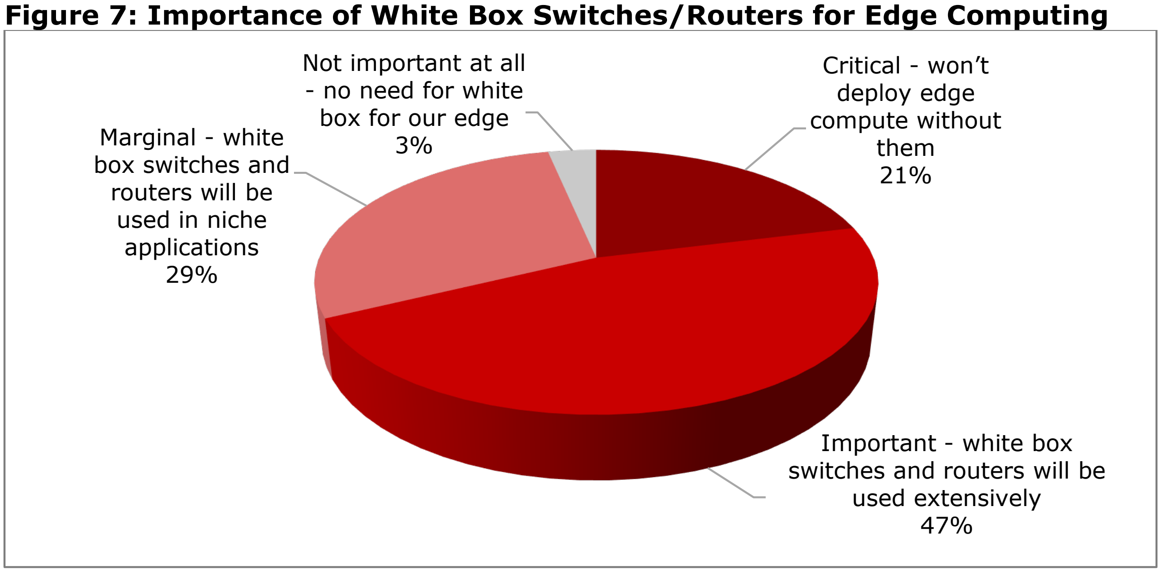

N=90
Source: Heavy Reading

Heavy Reading notes a disconnect in the market between optimistic service provider and enterprise plans (as expressed in **Figure 7**) and the current market reality that few white boxes have been commercially deployed – in edge and elsewhere. The survey results reflect optimism for a white box future, but in order to come to fruition, performance, pricing, and operations thresholds must be met. If performance, pricing, or operations fall short, the packet white box trend will be delayed. How white box hardware and software vendors address these challenges will determine the trend's level of success.

As discussed earlier, integrated optics on routers, or IPoDWDM, is an important enabling technology for edge DCI applications – at least for telecom operators. We asked respondents to estimate how extensively IPoDWDM will be used in their edge networks. **Figure 8** shows results for the 59 telecom operator respondents only.

Figure 8: Expected IPoDWDM Connectivity for Edge Networks (Telecom Only)

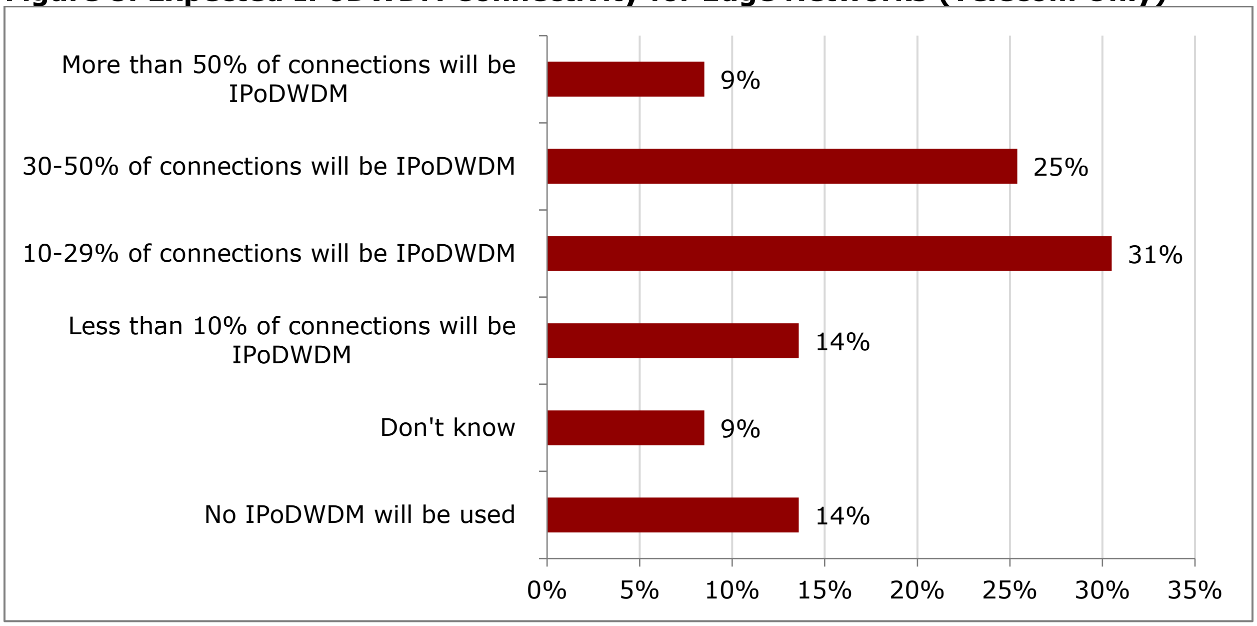

N=59
Source: Heavy Reading

At nearly one-third (31%), a plurality of telecom respondents expect that 10%-29% of their edge connections will be IPoDWDM, while an additional quarter (25%) of them expect 30%-50% of edge connections to use IPoDWDM. Thus, over half of telecom operators surveyed expect IPoDWDM to account for 10%-50% of their edge DCI connections. While few respondents expect more than 50% of edge connections to use IPoDWDM, few also expect no IPoDWDM to be used at the edge (just 14% said no integrated transponders will be used).

The emergence of 400 Gbit/s pluggables standardized by the OIF as 400ZR promises to jump-start the IPoDWDM market by eliminating one of the biggest barriers: the faceplate trade-off. Historically, long reach optics have been much larger than short reach, client optics (essentially double), thus forcing customers to sacrifice up to half of system's switching/routing capacity when integrating long reach optics. With 400ZR, long reach and client optics will, for the first time, use the same module form factors, requiring no capacity trade-off.

However, the 400ZR promise has not yet resonated with most telecom operators. According to the survey, just 9% of telecom operators knowledgeable about 400ZR expect to deploy modules for edge DCI, while 31% expect no 400ZR to be deployed. At 60%, the majority of respondents said their organizations are still deciding what to do (see **Figure 9**).

Figure 9: Expectations to Deploy OIF 400ZR in Edge Transport/DCI Networks (Telecom Only)

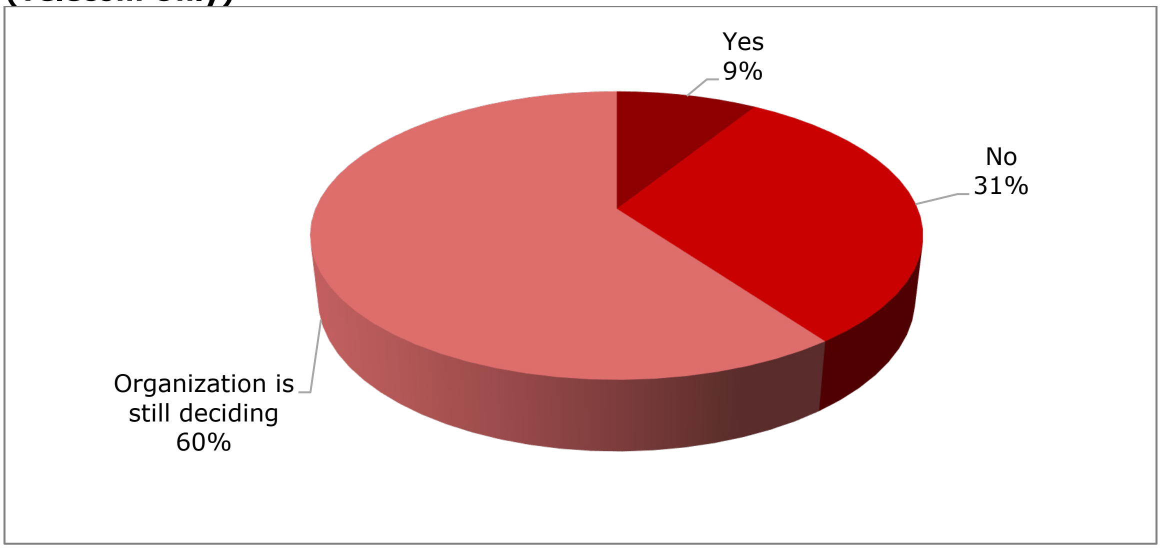

N=45
Only respondents with knowledge of 400ZR included
Source: Heavy Reading

Since the OIF 400ZR implementation agreement is not yet complete, some indecision at this early stage is expected. However, there are likely other factors contributing to the results. One question is: How much capacity will be needed at the edge? While 400 Gbit/s will be needed in large, central locations (as initially proposed by Microsoft in helping define the spec), it may be overkill for the emerging distributed edge.

A second factor is reach. The 80 km reach and point-to-point topology in the 400ZR spec may not be suitable for telecom operator networks, with ROADM-based networks and reach requirements in the 100s and 1,000s of kilometers. It remains to be seen how requirements that came from the hyperscale world will port over to traditional telecom.

As a final point, Heavy Reading notes that the adoption of 400ZR and IPoDWDM is partly (though not completely) intertwined, as 400ZR pluggables will be an important enabler for broader adoption of integrated optics. However, if 400 Gbit/s capacity is simply overkill for the edge, IPoDWDM becomes less appealing. The survey results suggest that – for the edge – this is likely the case, at least over the next 3 years.

COLOCATION STRATEGIES

Survey results highlight the significance of multi-tenant colocation for edge computing – among telecom operators and enterprises. In this section, we explore colocation trends for edge computing in more detail. In Heavy Reading's survey, 81% of respondents reported that network-dense colocation was at least important for their edge strategy, with 17% of the group reporting that colocation is critical and that their edge strategy cannot succeed without it (see **Figure 10**). Just 2% of respondents do not consider network-dense colocation important at all.

The finding is significant because it indicates that choosing an edge location is not simply a question of "where." Equally significant (or perhaps more significant) is the question of "who's there." Inter-connectivity was a primary factor in peering and CDNs, and results suggest that edge computing may not be so different in this regard.

Figure 10: Importance of Network-Dense Colocation for Edge Strategy

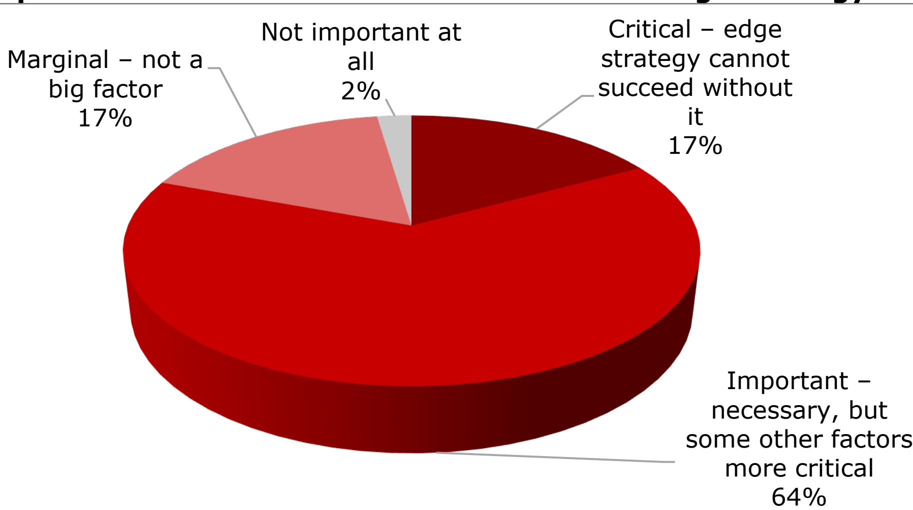

N=88
Source: Heavy Reading

Primary colocation benefits differ for telecom operators and enterprises. For telecom operators, the primary colocation benefit is the ability to access the edge where it is needed, followed by the ability to connect to major network service providers and the ability to connect to major content/cloud providers (**Figure 11** below). For enterprises, however, the primary benefit is the time-to-market advantage compared to building their own edge locations. The ability to connect to major network service providers is also important to enterprises (ranked second), followed by the ability to access the edge where needed (ranked third). Connecting to cloud providers is a bit less important for enterprises surveyed, compared to telecom operators.

Significantly, expected cost savings was listed as the least compelling colocation benefit for both telecom operators and enterprises. Heavy Reading does not believe that the finding indicates that colocation cost is unimportant. Rather, we think that it reflects the reality that colocation is very expensive. High colocation costs are a frequent complaint in Heavy

Reading service provider interviews and are raised again in our survey question about colocation inhibitors (below).

Figure 11: Primary Benefits of Colocation for Edge Strategy

Item	Telecom Rank	Enterprise Rank
Ability to access the edge where it is needed	1	3
Ability to connect to major network service providers	2	2
Ability to connect to major content/cloud providers	3	4
Time to market vs. build our own edge	4	1
Cost savings vs. build our own edge	5	5

N=55 telecom, 22 enterprise
Source: Heavy Reading

Security concerns tops the list of inhibitors to colocation at the edge, despite the money and efforts that data center operators invest to ensure their facilities and data are secure. In our survey, 56% of respondents cited security as a primary inhibitor, and it was the top choice among both telecom operators and enterprises. In addition to security, lack of control over infrastructure and applications is a major inhibitor (selected by 47% of the group), followed by high costs of colocation (selected by 43%; see **Figure 12**).

Telecom operators and enterprises surveyed are largely in agreement on the main inhibitors. However, telecom respondents ranked lack of control higher compared to enterprises and enterprises ranked high costs higher compared to telecom.

Figure 12: Primary Inhibitors of Colocation for Edge Strategy

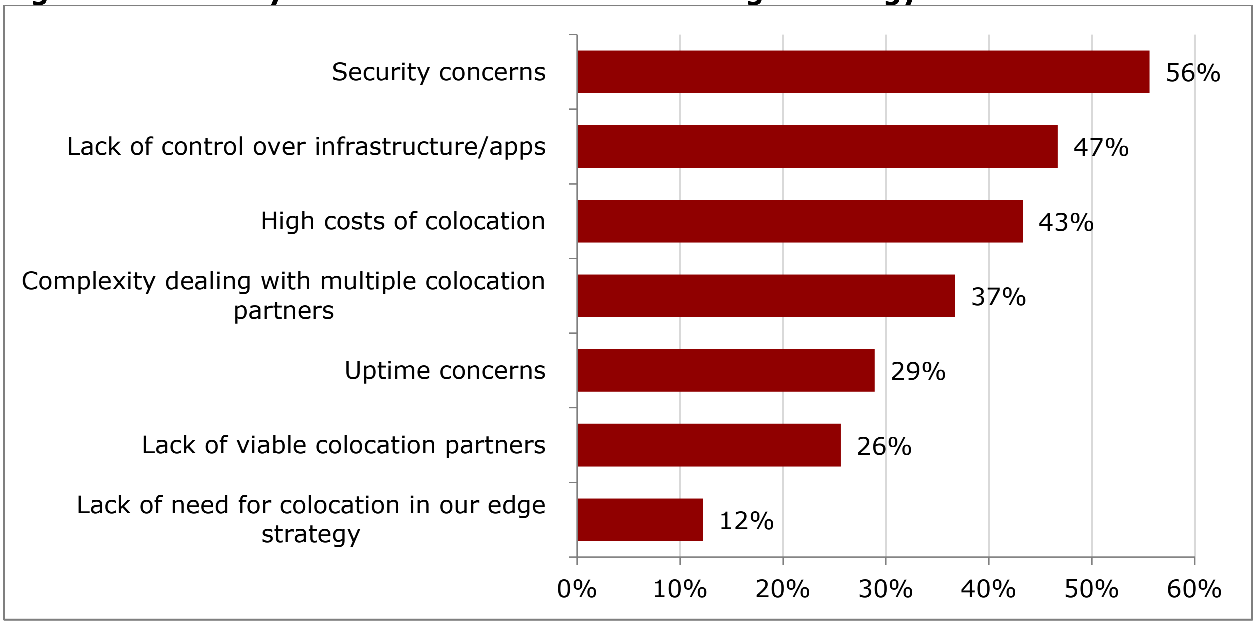

N=90
Source: Heavy Reading

APPENDIX: SURVEY DEMOGRAPHICS

This Heavy Reading report is based on a web-based survey of network operators, enterprises, and cloud providers worldwide conducted in June and July 2019. Respondents were drawn from two sources:

- Telecom operators were sourced from the Light Reading readership database.
- Enterprises were sourced from Informa's Network Computing readership database.

After reviewing responses, 91 were deemed qualified participants and were counted in the results. To qualify, respondents had to work for a verifiable service provider or enterprise. The full survey demographics are detailed below.

Figure 1A shows the type of company respondents work for.

Figure 1A: Respondents by Company Type

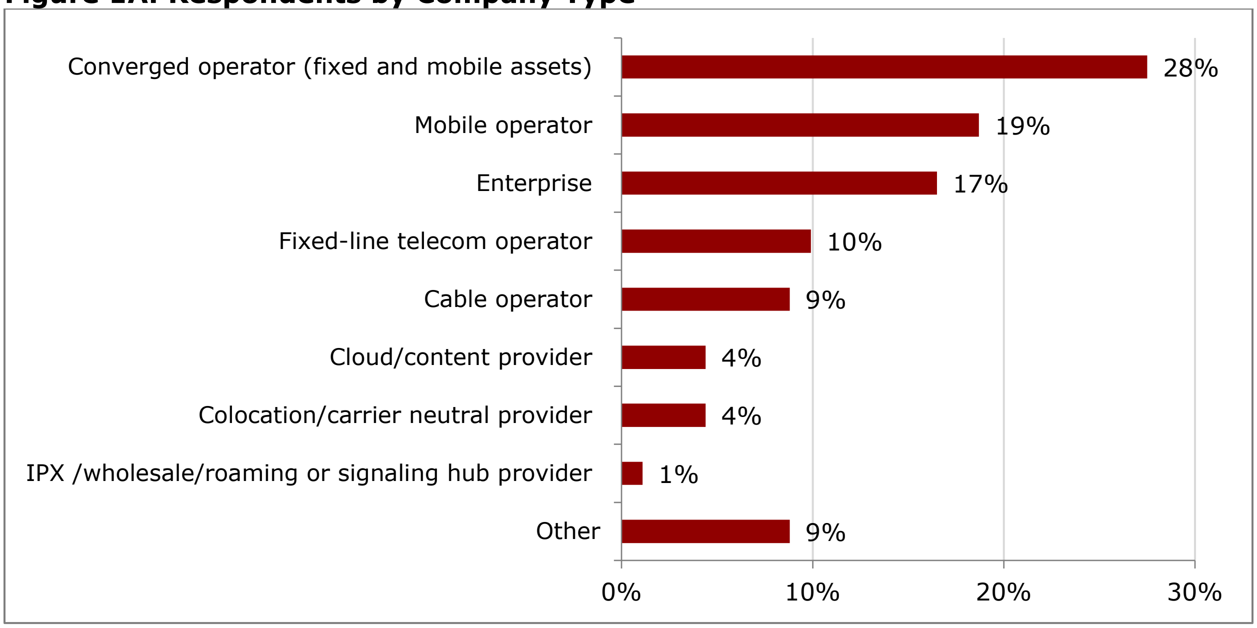

N=91
Source: Heavy Reading

Figure 2A shows respondents' involvement in network planning and/or purchasing of network equipment in the organization.

Figure 2A: Involvement in Network Planning and/or Purchasing

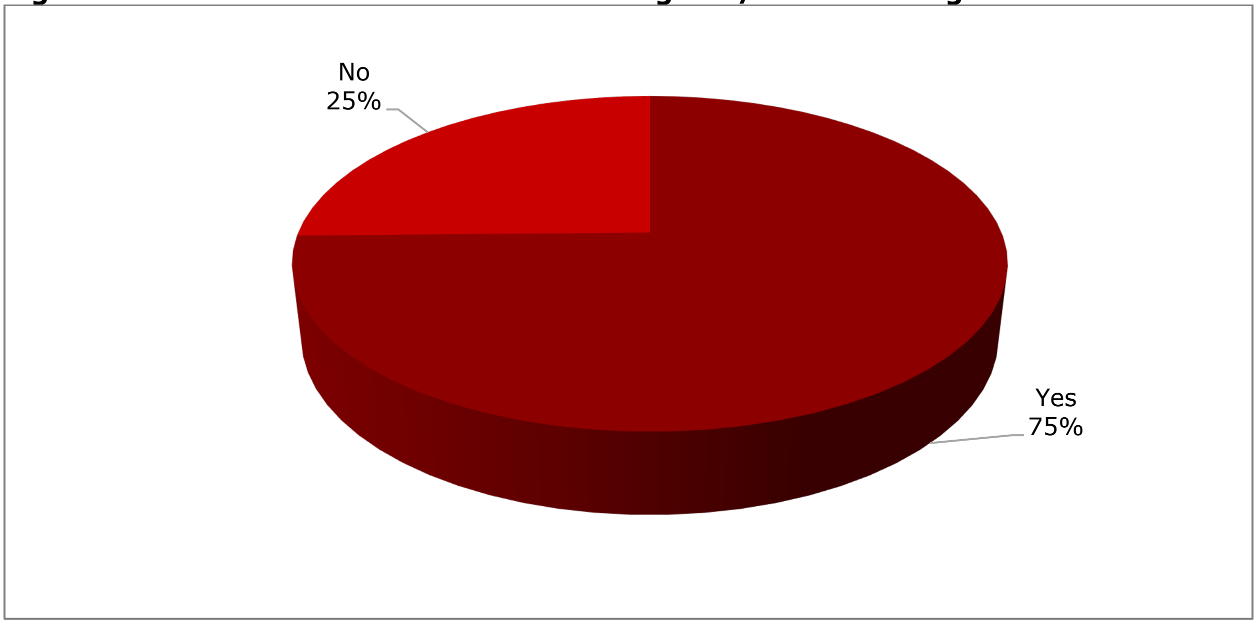

N=91

Source: Heavy Reading

Figure 3A shows respondents' direct involvement in their organization's edge strategy and/or purchasing for the edge.

Figure 3A: Direct Involvement in Edge Strategy and/or Purchasing

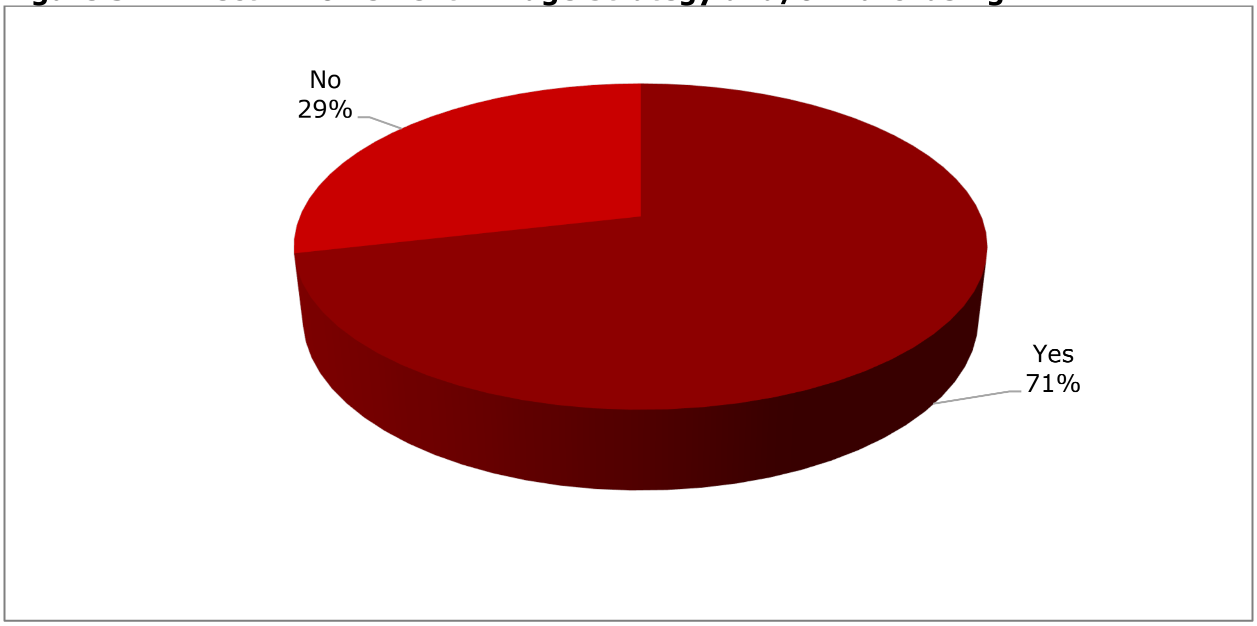

N=91

Source: Heavy Reading

Figure 4A shows survey respondents broken down by geographic region.

Figure 4A: Respondents by Geographic Region

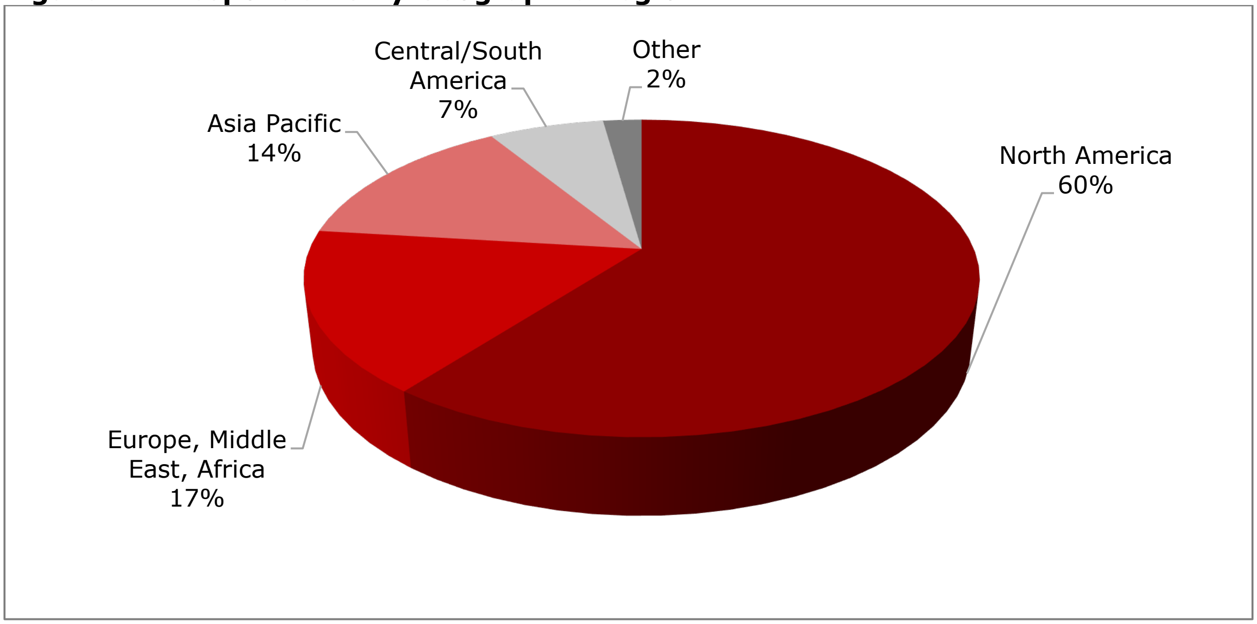

N=91
Source: Heavy Reading

Figure 5A shows survey respondents broken out by company annual revenue.

Figure 5A: Respondent Breakout by Company Annual Revenue

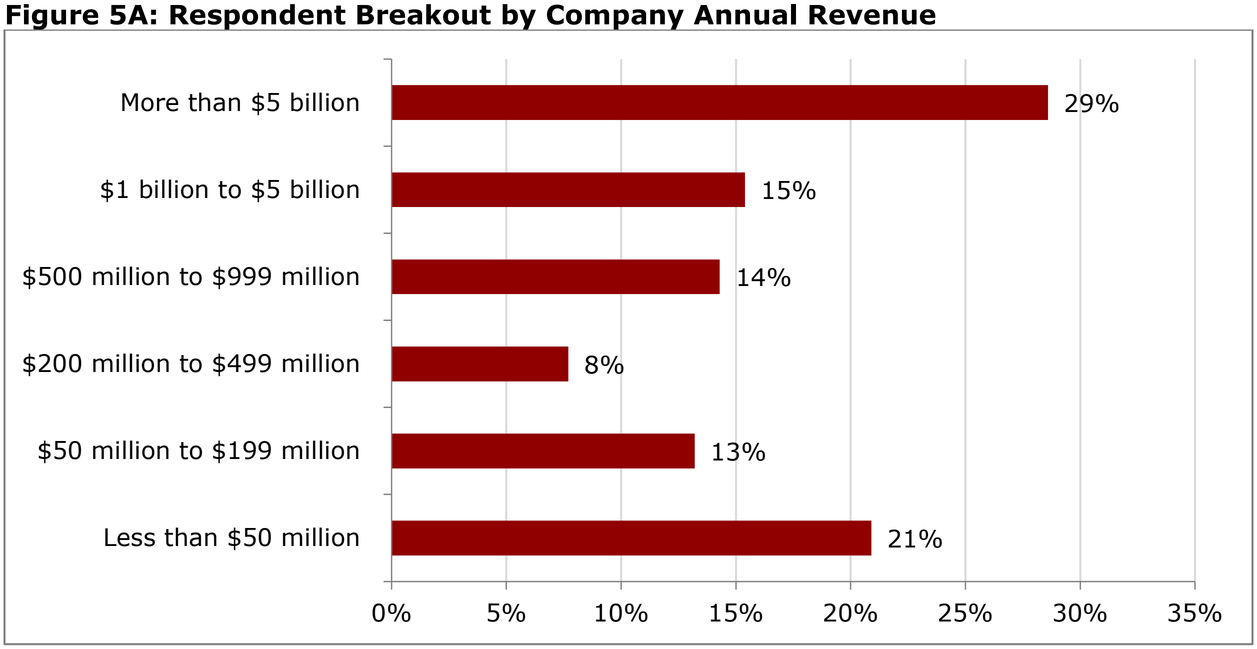

N=91
Source: Heavy Reading

Figure 6A shows survey respondents broken out by job function.

Figure 6A: Respondents by Job Function

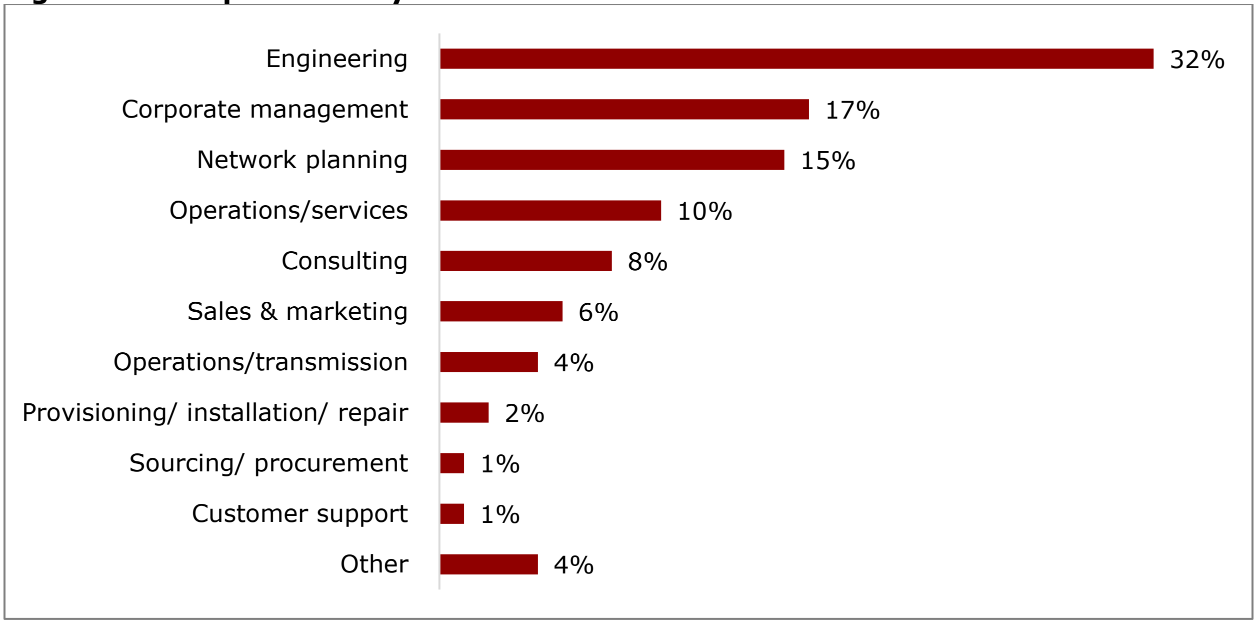

N=91
Source: Heavy Reading