Produced by

HEAVY READING

Light Reading

in Partnership with

5G americas

atis

5G Network Strategies

OPERATOR SURVEY 2022

Jennifer Clark
Principal Analyst, Edge Computing
& Cloud Infrastructure
Heavy Reading

5G and Edge Computing

Edge computing is a focal point of innovation for operators and their partners. The survey results show that multi-access edge compute (MEC) deployments are accelerating and that operators need to build roadmaps that emphasize support for new applications and services that do not overleverage partnerships with the hyperscalers.

Key takeaways

- Early edge deployments are largely focused on overall improved network costs and performance, not on creating new revenue streams and enabling new applications.
- The current MEC market favors the manufacturing and healthcare verticals. However, respondents predict a rapid uptick over the next 24 months in the deployment of edge services to the retail, media and entertainment, and transportation & logistics verticals.
- The deployment of MEC combined with a move to containers are very complex transitions and somewhat outside the current skill set of the operators. This is making partnering with the hyperscalers an increasingly attractive (if hazardous) shortcut.

The drivers of edge computing deployment can be divided by those that improve performance and operational costs and those that enable new applications and revenue streams. The respondent pool, regardless of company size (over or under $5bn in annual revenue) or location (US vs. Rest of World [RoW] based) selected “reduce bandwidth use/cost” as the top driver for edge computing (Figure 16). Related drivers of “improve application performance in general” and “improved resilience” were much more important to smaller operators than to Heavy Reading’s over $5bn segment. All three drivers translate to money in the pocket of the operators—the first by the more efficient use of bandwidth, the other two by providing consistency in performance—thereby lowering pay-outs for service-level agreement (SLA) breaches. This perspective has been reinforced by one-on-one conversations with the operators, which shared that their early MEC rollouts were aimed not so much at reducing latency for some demanding applications, but at establishing consistent delay and jitter for all applications. From this perspective, early MEC implementations are acting as a pressure valve. They are helping to absorb sudden demands on the network and, as one Tier 1 operator commented, they are giving their enterprise customers some “extra runway” to handle ever increasing network loads.

The largest operators showed much greater interest in drivers that enabled new revenue streams (as opposed to lower costs). “Differentiate services vs. competitors” is far more important to the largest operators (65% vs. 28% for respondents with annual revenue of under $5bn). Heavy Reading posits that the “competitors” these Tier 1 operators are responding to are not only other telcos, but also hyperscalers (the ultimate “frenemies” as both partners and competitors).

“Meet data sovereignty and regulatory requirements” claims the “also ran” spot in this survey. However, it can provide low hanging fruit for the operators, particularly in some heavily regulated healthcare, defense, and government applications.

Fig16: What are your top motivators for moving workloads to the edge? (Select all that apply.)

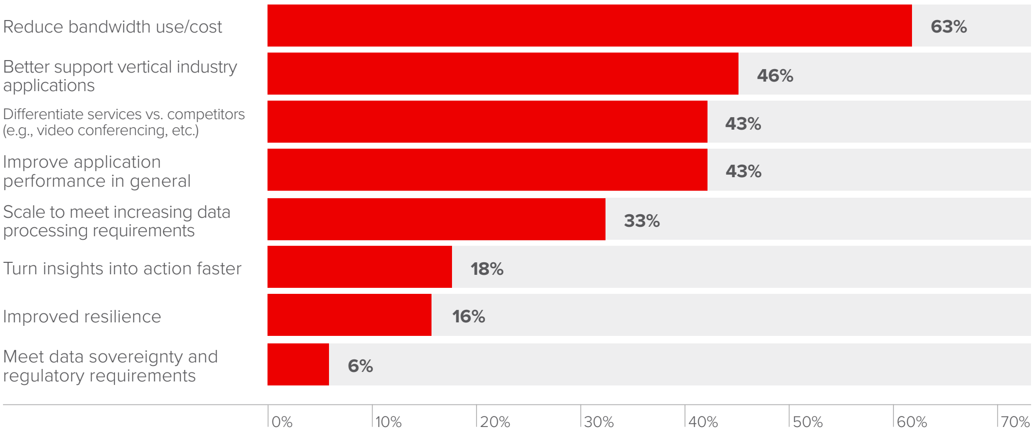

Source: Heavy Reading

n=87

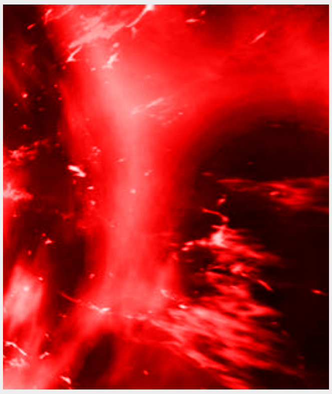

According to Figure 17, healthcare, automotive, financial services, and manufacturing respondents have the greatest percentages of current implementations. “Automotive” clearly does not refer to fully autonomous vehicles on public highways. Rather, it is a catchall for a variety of auto-centric applications: insurance, Internet of Things (IoT) cameras and devices, onboard entertainment, and campus autonomous vehicles, with some manufacturing and fully autonomous driving trials and proofs of concepts (PoCs) thrown in. The media & entertainment vertical shows the largest growth spurt over the next 24 months, followed by retail and transportation & logistics.

In terms of regional differences, areas outside the US are less interested in the financial vertical—33% with “no near-term plans”—but show a heightened focus on deploying MEC services within the next 24 month in both the healthcare (59%) and transportation & logistics (71%) verticals. These verticals can be heavily subsidized outside of the US and frequently bubble to the top in Heavy Reading’s edge surveys, particularly when the spotlight is on Europe.

> “Survey results show that multi-access edge compute (MEC) deployments are accelerating and that operators need to build roadmaps that emphasize support for new applications and services”

Fig17: When will your organization start to offer 5G edge services for the following industries/use cases?

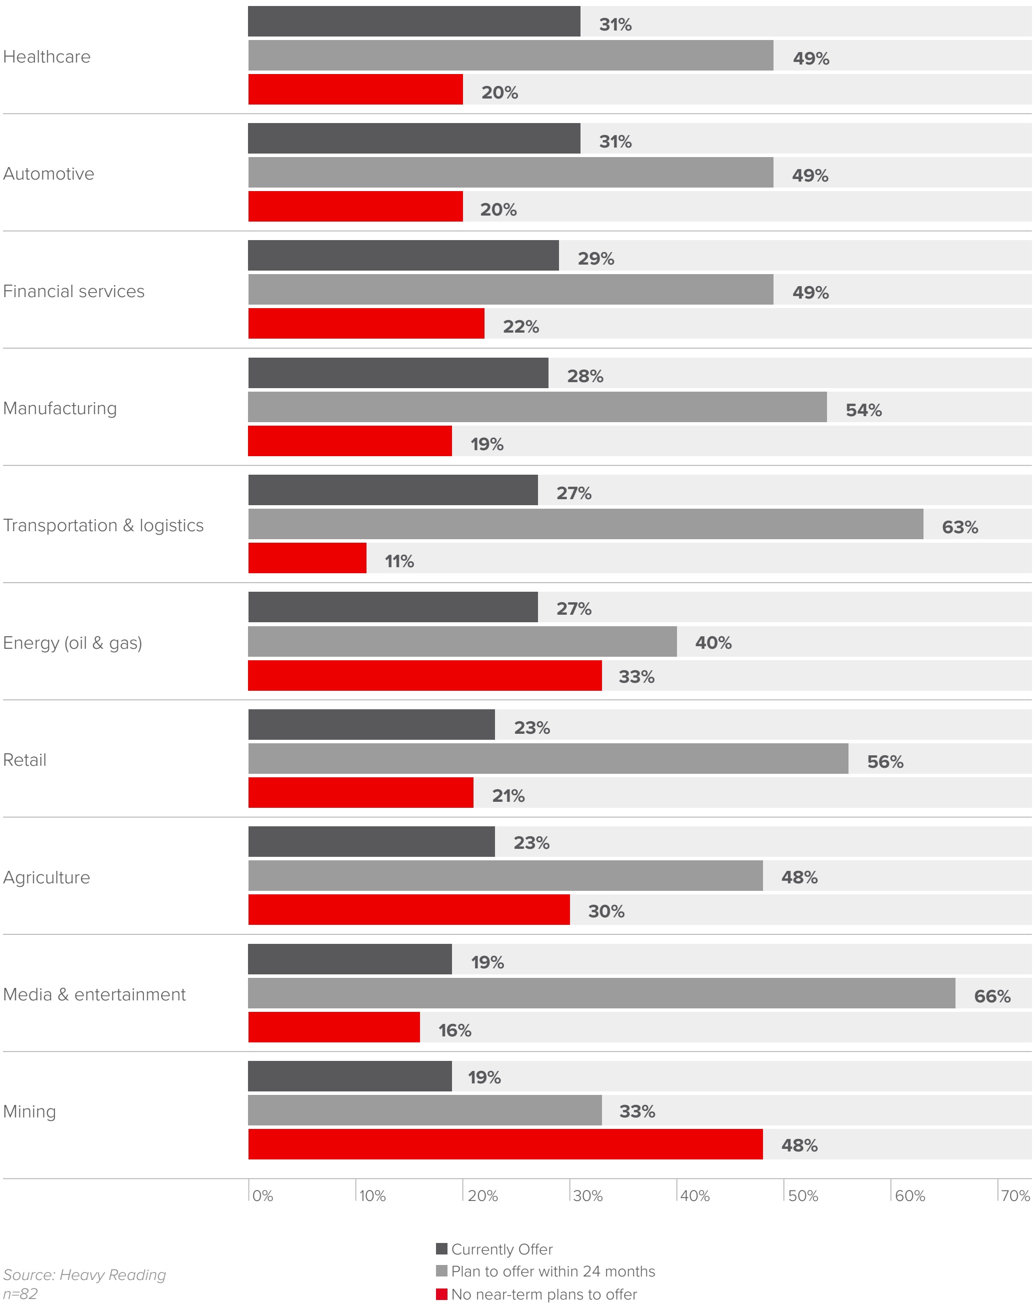

Source: Heavy Reading
n=82

While mining shows the greatest percentage (almost half) of respondents with “no near-term plans” to deploy MEC, these responses are heavily weighted toward smaller operators (61% vs. 25% for >$5bn operators) and operators outside of the US (66% RoW vs. 21% US). However, the demographic spread for the related vertical of energy is different. Heavy Reading believes this is because US Tier 1 operators, in particular, have been targeting mining operations for private 5G and view MEC as a key enabler of private 5G.

Edge services can be offered via a variety of deployment models with varying degrees of ownership, control, and geographic reach. As shown in Figure 18, a “hybrid public/private telco cloud infrastructure,” giving CPSs the most attractive combination of both control and reach, garners a third of survey responses overall. However, when looking only at the largest operators (>$5bn in revenue), that percentage grows to almost 50%. The next two deployment models fall precipitously to around 15%. The remaining choices for these Tier 1 operators are in the single digits—as they are with the overall survey respondents.

Fig18: **How will your organization primarily offer edge services?**

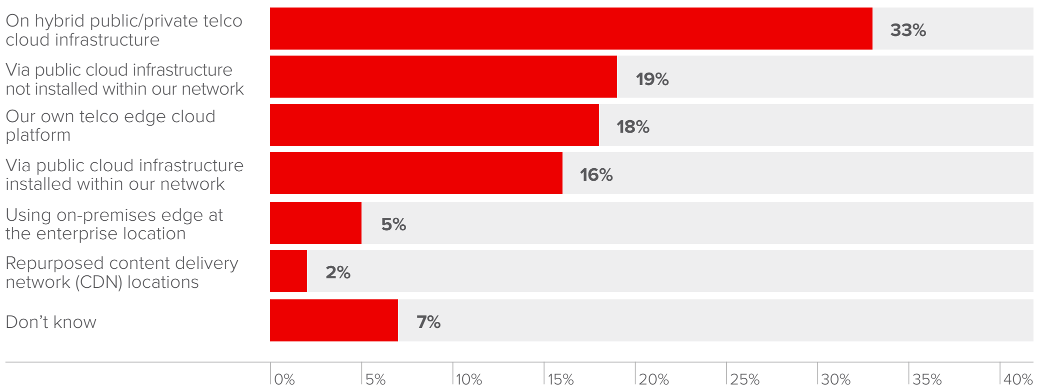

Source: Heavy Reading *n=84*

Operators with annual revenue of under $5bn divide their responses almost equally between the first four responses, with little variation between the US and RoW. The reluctance to deal with hyperscalers that was seen from some regions outside the US in earlier surveys has dissipated as operators have been seduced by the ease with which they are able to get an edge service up and running in their own region and beyond.

Heavy Reading was surprised to see any operator claim that the primary edge deployment model would be on-premises edge at the enterprise location. The four respondents that made this selection are all mobile or converged operators from both the US and RoW. Three of the four have revenue of over $5bn. The most logical explanation is that these edge deployments are in a private 5G network or were driven by a specific IoT enterprise app.

The two respondents that decided on the practical strategy of leveraging repurposed content delivery network (CDN) locations were a cable operator and a mobile virtual network operator/enabler (MVNO/MVNE).

The use of containers enables enterprises to build and run highly scalable and flexible applications for deployment at the edge and/or in a public, private, or hybrid cloud. Key benefits of containers/microservices include the following:

- Lower total cost of ownership (TCO), enabling users to deploy only what is needed, rather than entire monolithic network functions.
- Faster time to market (TTM) for new services and applications.
- Ability to decouple the application from the infrastructure, simplify application development, and enable applications to run in a highly distributed fashion.
- Increased cadence of small and regular updates to applications enabled by the microservices/containerized architecture and the use of CI/CD.

As compelling as these benefits may seem, the transition to containers and containerized network functions (CNFs) from virtual machines (VMs) and virtual network functions (VNFs) is much more complex than the transition from appliance-based network functions to network functions virtualization (NFV) and VMs. Containers and cloud native represent a fundamental change in the way operators design, deploy, and manage applications and services. With this in mind, it is not surprising that almost half of respondents claim that less than 25% of their edge cloud workloads are containerized today (Figure 19). It is encouraging that over 50% of respondents expect 51% or more of their workloads to be containerized by 2025. Heavy Reading notes some acceleration of container adoption predicted by the very large operators and US operators, but the numbers are statistically similar to the survey base overall.

When examining the factors that are limiting edge deployment (Figure 20), the most interesting detail to note from the data is that there is only a 6-percentage-point difference between the number one limiting factor, "cost and complexity of infrastructure," and the fourth factor, "integration/compatibility between ecosystem components." And each of the four factors was identified as a barrier by around half of the respondents. These top four factors combine to form a significant barrier to deployment. It is this barrier that is encouraging operators to partner with the hyperscalers for their MEC deployments: because it is easier. They are not expecting to reduce TCO, as has been pointed out in other recent MEC surveys.

Limited customer demand, pulling in a third of respondents, is concerning, but the percentages are predictably lower (hovering around 20%) for larger operators and US operators. ■

Fig19: What percentage of your edge cloud workloads is containerized now or will be by 2025?

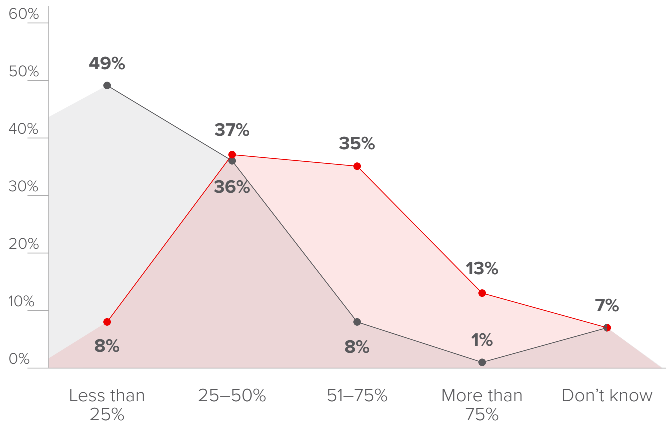

Source: Heavy Reading *n=84*

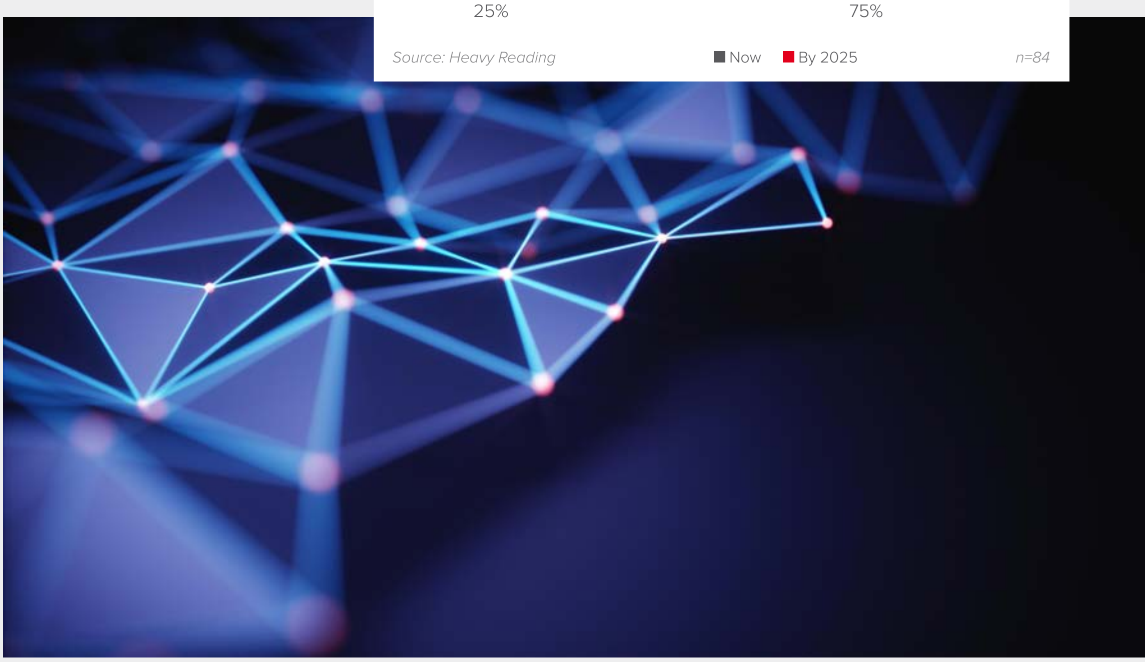

Fig20: What is most limiting your 5G and edge cloud deployment? (Select three)

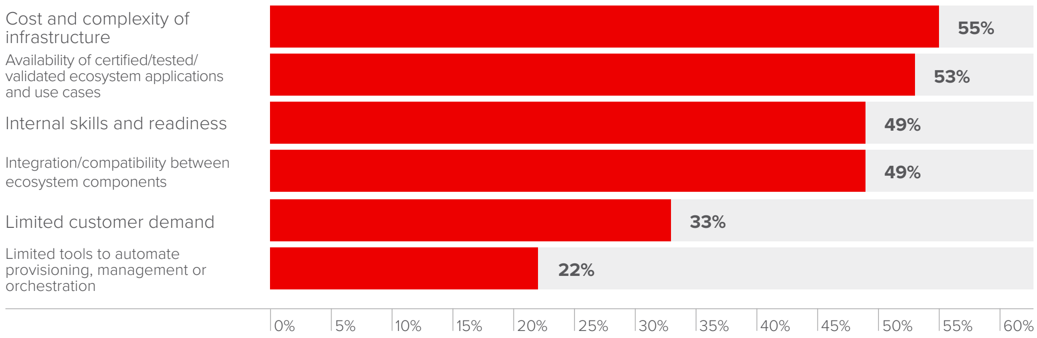

Source: Heavy Reading *n=82*

Executive Summary

Digital service providers (DSPs) who use the powerful combination of 5G with edge computing offer better user experiences and support bandwidth-hungry apps through a more flexible, agile, and resilient network. Using cloud-native solutions for their radio access networks (RANs) allows them to quickly scale to dynamically meet changing demand. With multi-access edge computing (MEC), service providers can deliver innovative, latency-sensitive services and applications for enterprise customers, and capture new revenues.

DSP organizations are looking for a unified, horizontal platform—from the core to the edge—with a consistent deployment and operations experience. Red Hat provides the flexibility to develop and deploy with speed and ease in any cloud you choose. Together with our ecosystem partners, we help customers make the most of 5G opportunities with pre-integrated offers to build out services with confidence and without fear of lock-in. For the foreseeable future, digital service providers will have many different kinds of workloads in multicloud environments. Our telco-grade hybrid cloud solutions provide a consistent, predictable foundation that lets service providers move between clouds as strategic, business, or technology needs evolve. Partnering with Red Hat, DSPs can provide the experience that customers expect, and address cost, resilience and regulatory requirements.

Red Hat

Jim Hodges
Research Director, Cloud,
and Security
Heavy Reading

5G Security

Many mobile operators are now well into the execution phase of commercializing their 5G core and RAN networks. One of the key considerations in this execution process is the definition and implementation of the unique 5G security capabilities vital to support cloud native services in centralized and edge cloud configurations.

Key takeaways

- Mobile operators, US operators especially, are confident that they have in place effective security strategies that will enable them to support and secure 5G services.
- These strategies rely on several foundational capabilities to secure 5G infrastructure and devices. Leading the way here are “trusted hardware” (46%) and “identity and access management” (43%) for infrastructure security. Trusted hardware is also considered a critical component for securing edge infrastructure (48%) and even device endpoints (46%).
- The survey respondents believe that as edge deployments scale, they will need to evolve their security strategies and focus on “policy compliance scanning” (66%), “implementing zero-trust principles” (59%), and “encryption of data at rest” (50%).

While operators’ 5G security strategies will vary due to specific regional market and threat requirements, as Figure 31 documents, globally some common capabilities will play a strategic role in the execution.

Of these, the two that stand out based on “critical” inputs are the well-established security fundamentals of “trusted hardware” (46%) and “identity and access management” (43%).

Other newer capabilities bundled in the 30–33% range, such as “isolation & policy enforcement” (33%), “container orchestration security” (31%), “continuous image security scan & vulnerability analysis” (31%), “automated security remediations” (31%), and “visibility into trust status & operations” (30%), also resonated with the survey respondents. These results confirm that orchestration, policy, and automation are also key components.

Fig31: How important are the following capabilities for securing 5G infrastructure at your company?

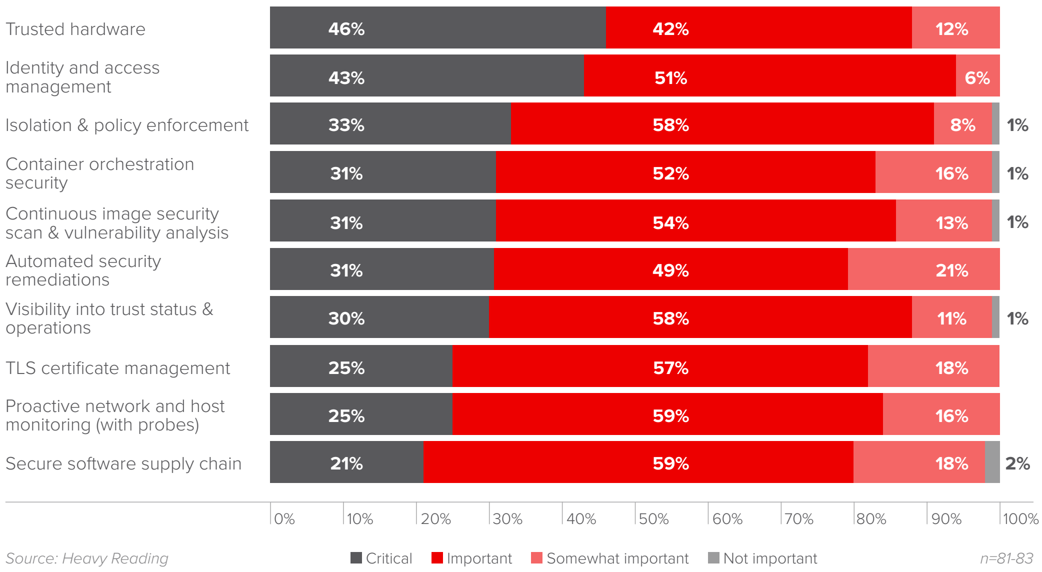

Source: Heavy Reading *n=81-83*

Fig32: How important are the following capabilities for securing edge infrastructure at your company?

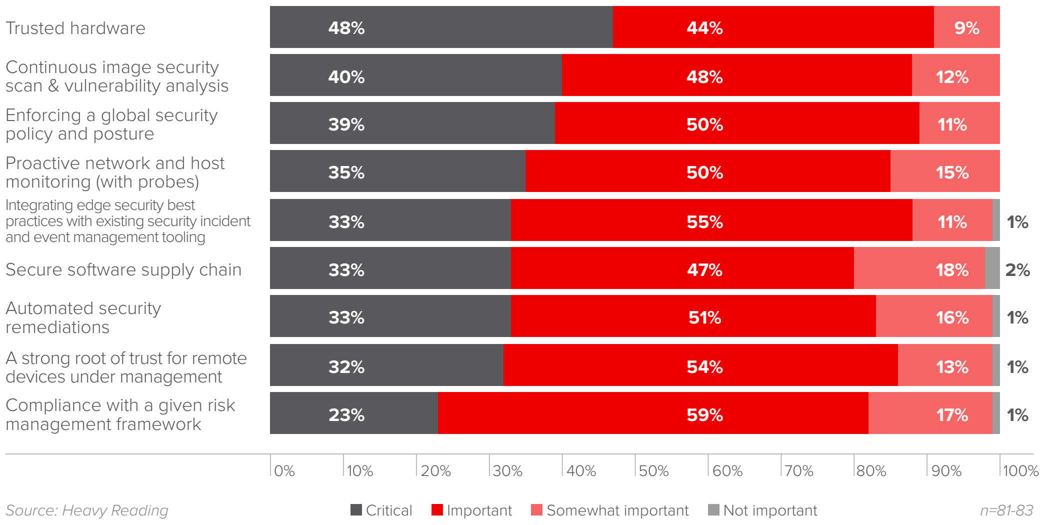

Source: Heavy Reading *n=81-83*

One consideration that a mobile operator must address in its security strategy is prioritizing the specific security capabilities that will be most important for securing edge infrastructure. As Figure 32 shows, based on “critical” inputs, the top three capabilities are not that different from the capabilities documented directly above.

This includes “trusted hardware” (48%) attaining the highest ranking. Other capabilities of note include the second-place scoring of “continuous image security scan & vulnerability analysis” (40%), which moved up from its fourth-place ranking in Figure 31, and “enforcing a global security policy and posture” (39%).

Consistent with the previous figure inputs, several capabilities fell into a narrow data range (32–35%). “Proactive network and host monitoring (with probes)” (35%) leads the pack, followed closely by “integrating edge security best practices,” “secure software supply chain,” and “automated security remediations” (all 33%).

Based on these two inputs, it is clear that 5G security will rely on a number of important capabilities both in the core and at the edge. Leading the way are trusted hardware and capabilities that enable the ability to support continuous image scanning and adaptive and automated policy control.

> “Many mobile operators are now well into the execution phase of commercializing their 5G core and RAN networks.”

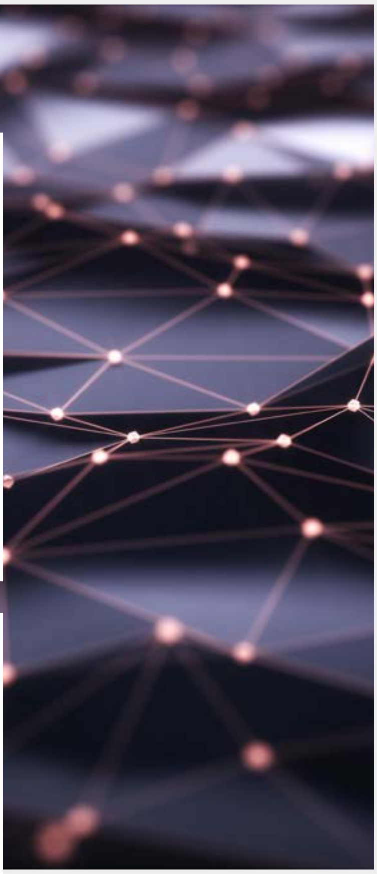

Fig33: How important are the following capabilities for securing endpoints, such as smartphones and IoT devices?

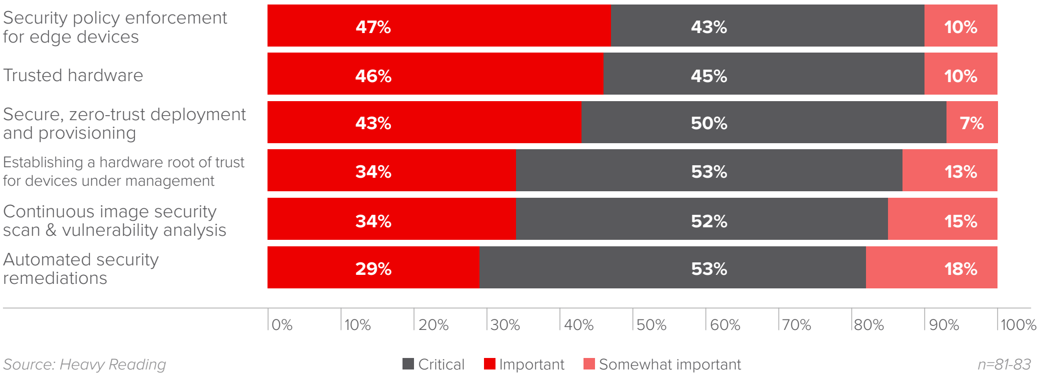

Source: Heavy Reading *n=81-83*

Fig34: How will your company evolve its security strategy, as 5G emerges with more edge activity? (Select all that apply)

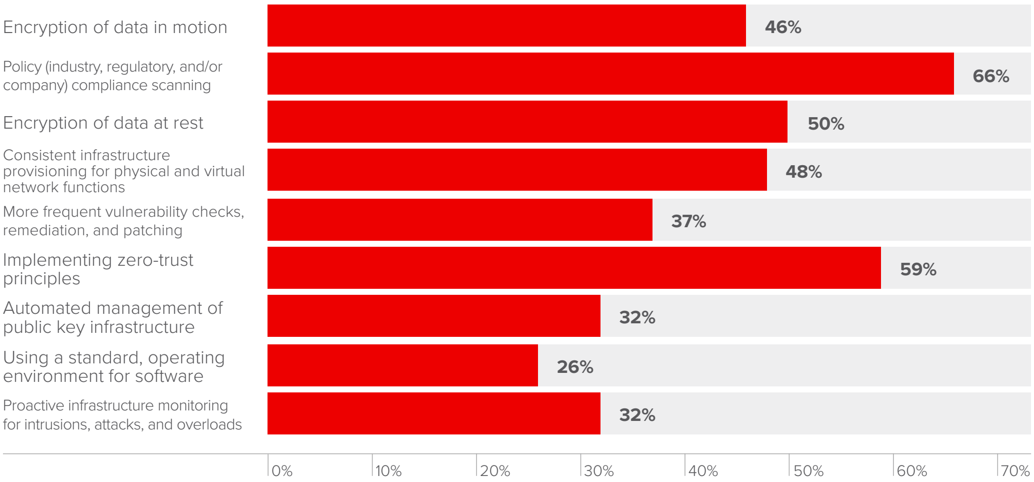

Source: Heavy Reading *n=82*

All generations of mobile technology have introduced new types of mobile devices with new form factors, enhanced performance metrics, and storage capabilities. 5G will not disappoint in this regard with the ability to support ultra-low latency cloud native services. However, with these new capabilities comes a greater risk that these new devices can be “weaponized” to support cyber-attacks. This applies to not only smartphones, but also IoT devices that will be crucial to fuel machine-to-machine (M2M) service innovation.

To address these new device-centric threats, mobile operators will also need to support new capabilities to secure these devices. As Figure 33 illustrates and as seen with infrastructure, based on “critical” inputs, policy-related capability, specifically “security policy enforcement for edge devices” (47%), is a top consideration.

“Trusted hardware” (46%), consistent with previous input, scored highly, attaining a second-place ranking. In third place with a score of 43% is “secure, zero-trust deployment and provisioning,” which, in Heavy Reading’s opinion, reinforces the importance of adopting zero-trust provisioning for cloud networks.

The next question in the survey provided additional granular security insights into the impact of the evolution and adoption of 5G services at the edge. As Figure 34 illustrates, there are several top-of-mind concerns. Of these, “policy compliance scanning” attained the highest ranking (66%), followed closely by “implementing zero-trust principles” (59%) and then “encryption of data at rest” (50%).

Other capabilities of note include “consistent infrastructure provisioning for physical and virtual network functions” (48%) and “encryption of data in motion” (46%). Heavy Reading views this data as reinforcing the importance of zero-trust principles and the value of policy-based scanning, as well as confirming the value of data encryption at the edge.

The final question of this section assessed overall operator 5G-related security confidence and strategy readiness. As Figure 35 shows, despite the challenges, 78% of the respondents believe that their "5G security strategy supports the new requirements associated with disaggregated and distributed network infrastructure" and will even support "the ability to secure 5G network slices" (78%).

They also are confident that their 5G strategy will support all-important zero-trust principles (69%) and will incorporate a "self-learning autonomous system leveraging [machine learning/ artificial intelligence] ML/AI" (67%) and even "services running in public clouds" (67%).

The one area where respondents' confidence is lower relates to the maturity of their security strategy and ability to scale and be production-ready (59%). Unlike the other questions in this section in which the data trends between the US and RoW respondents were quite similar, a greater number of US respondents agree that their 5G security strategies were mature, scalable, and in production compared to their RoW counterparts (US =76%; RoW = 48%). One factor likely influencing the variance is that 88% of US respondents agree that they have the internal resources and skill sets to secure their 5G network compared to only 51% of RoW respondents. As a result, a lower percentage of RoW operators agree they are ready to implement zero-trust principles (US = 79%; RoW = 63%) and 5G network slices (US = 88%; RoW = 71%). ■

Fig35: Do you agree or disagree with the following statements?

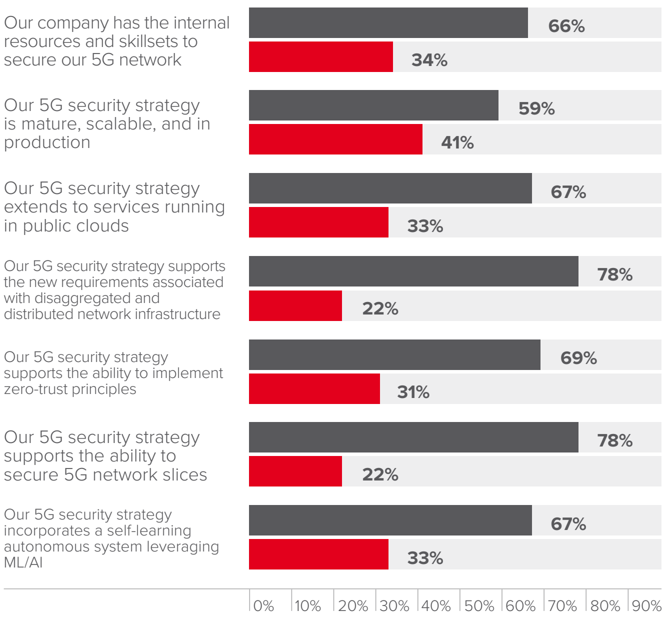

Source: Heavy Reading

■ Agree ■ Disagree

> "Mobile operators, US operators especially, are confident that they have in place effective security strategies that will enable them to support and secure 5G services."

Executive Summary

You are only as secure as your weakest link. As application environments evolve, security teams are increasingly challenged to keep up with the changing risks, compliance requirements, tools, and architectural changes introduced by these innovations. Traditional perimeter-based network security is no longer effective on its own. Security should be implemented within each layer of the application and infrastructure stack. Automation is a critical part of scaling how the organization addresses security and compliance monitoring.

Red Hat wants to help you have confidence as you adopt a continuous security strategy to maintain security and regulatory compliance, while helping your business remain competitive, flexible, and adaptable. Red Hat provides telco-grade technologies to build, manage, and automate hybrid clouds more securely as part of a layered, defense-in-depth security strategy, and our broad partner ecosystem extends these capabilities even further. You can take advantage of the capabilities at each layer in your environment, including operating systems, container platforms, automation tools, Software-as-a-Service (SaaS) assets, and cloud services. Visit redhat.com/security to learn more about Red Hat's commitment to protecting your environments and the data and privacy of your customers.