



5G EDGE CLOUD

Author: **Gabriel Brown**, Principal Analyst, Mobile Networks & 5G, **Heavy Reading**

To deliver consistent performance in 5G networks, and low latency ultra-reliable services in particular, it is widely expected that operators will need edge cloud infrastructure. Hosting applications and content closer to the “user” should improve the service experience and enable high performance applications that are impractical, inconsistent, or not possible using only large, centralized cloud infrastructure.

KEY TAKEAWAYS:



There is no single motivation to deploy edge cloud infrastructure that stands out, but rather, several reasons – each with solid support. **To illustrate, “ensuring application performance” scores highest in terms of the primary motivation with 32%**, ahead of “differentiated communications services,” which scores lowest at 19%.



Operators are making progress on their edge cloud deployments. “Developing” scores highest, followed by “early stages” in respondents’ self-assessment of their progress. Less than 20% said they have “not started,” and in all but one case, fewer than 20% claim their deployment as “mature.” This indicates the rollout of edge cloud infrastructure is well underway in advanced operators but far from finished.



At the edge, operators expect to support both containers and VMs now and in the medium term. Consistent with the prior finding on 5G core, **the conclusion is that edge cloud infrastructure must support multi-mode workloads.**




A majority (71%) expect to support less than 100 edge cloud locations in 2020. Looking ahead to end-2023 – i.e., 4 years from now – the picture changes. At this stage, a majority of operators expect to support more than 100 edge cloud locations, but less than 1,000. U.S. respondents selected higher numbers of locations than their peers in RoW, both now and in 2023.



5G Edge Cloud

In an attempt to identify respondents' primary motivation for edge cloud deployment in 5G networks, the first question in this section (Fig 20) allowed only one answer. The result shows there are several reasons, each with solid support and no clear winner. "Ensuring application performance" scores highest at 32%, followed by "vertical industry services" at 28%. These two options in combination account for 60%. The percentage of respondents that are more focused on operators' internal priorities (21% "reduced transport cost" and 19% "differentiated communications services") is lower at 40%, but not by a huge amount.

Fig 20. What is your PRIMARY motivation to move workloads to the edge? (N=143)



- Reduce bandwidth use/cost21%
- Offer differentiated communications services (vs competitors).....19%
- Offer vertical industry services (e.g. in-vehicle scanning for ambulances, advanced real-time analytics for investment banking).....28%
- Ensure application performance.....32%

Given the evenness of response, the conclusion is that the investment case for 5G edge cloud will require operators to pursue multiple motivations simultaneously.

The next question (Fig 21) asked about the stage of development of operators' edge cloud deployments. In all cases, "developing" scores highest, followed by "early stages"; less than 20% have "not started." In all but one case (network virtualization), less than 20% claim their deployment as "mature." The clustering of response around the middle options indicates the rollout of edge cloud is well underway in advanced operators, but far from finished.

U.S. respondents were slightly more likely to select "mature" than their RoW counterparts. However, only in the categories of "containerized network functions" and "containerized workloads" was this significant – to the tune of a 10% higher score. This may indicate U.S. operators are a little more advanced than peers elsewhere.

Asked about the risks of running different types of workloads at the edge (Fig 22), most respondents view edge cloud as "moderately risky" for nearly

all workloads, with security at the edge considered more severe. A third (35%) believe edge security presents "extreme risk." Quality control and cost control are considered somewhat less risky aspects of edge cloud deployment, perhaps because respondents feel these are issues operators are already familiar with when designing and deploying new infrastructure.

One might argue that the inclusion of the word "risk" in the question led respondents to express greater concern about security than they might have otherwise. However, security is highlighted elsewhere in the survey as a primary concern. For example, "security" was identified as the biggest risk of working with external cloud providers to provide enterprise 5G services.

The transition to virtualized and cloud-native telecom networks must consider the infrastructure on which workloads will run. This question (Fig 23) asked specifically about the edge and the mix between containers and VMs in 2020 and 2023. The straightforward analysis of the response is that operators expect to support both VMs and containers now and in the medium term, which leads to the conclusion that edge cloud

Fig 21. Where is your organization in its rollout of edge cloud capabilities? (N=140-144)

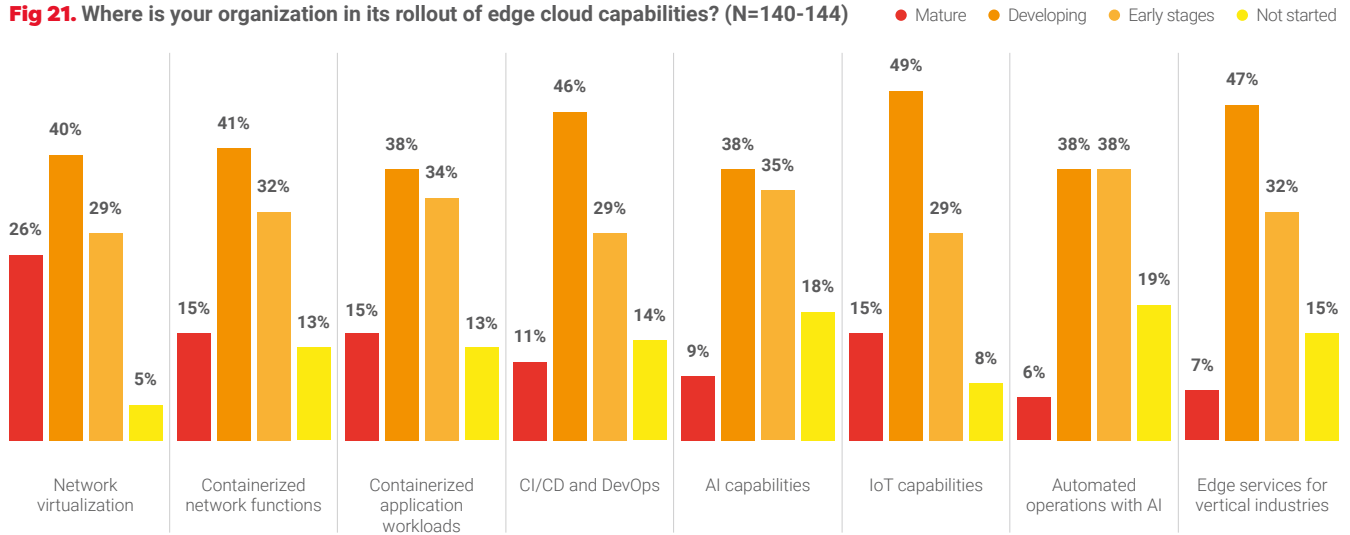


Fig 22. How much risk is involved in running more of the following workload at the edge? (N=143-144)

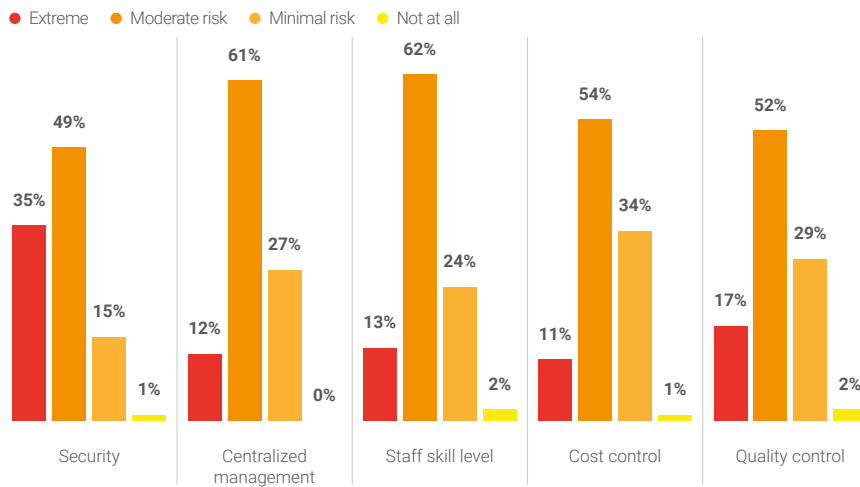
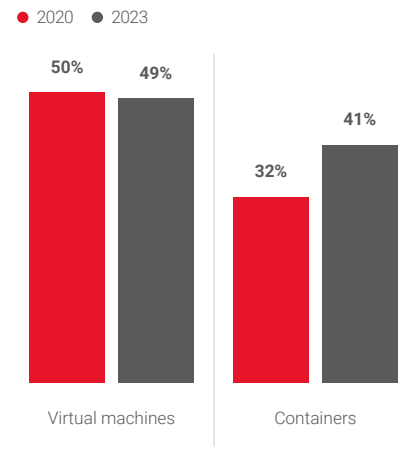


Fig 23. What percentage of your edge workloads will you run in virtual machines or containers? (N=137-144)




infrastructure platforms will be multi-mode. This finding is consistent with the response to the 5G core variation of this question in the section above, where it was established that a majority of operators (82%) would deploy 5G core using both VNFs and CNFs.

One of the most common discussion points in telco edge cloud in the

past few years has been: How many locations will operators deploy? Heavy Reading asked respondents (**Fig 24**) for their views in 2020, and then in 2023. The first finding is clear: a majority (71%) expect to support less than 100 edge cloud locations this year, which indicates larger facilities serving larger numbers of users will prevail in the near term.

Fig 24. How many network edge cloud locations do you support now, and will support in 2023? (N=141-143)



Looking ahead to 2023 – 4 years from now – the picture changes. At this stage, a majority of operators expect to support more than 100 edge cloud locations, with 48% saying between 100 and 999 locations, and 34% more than 1,000. However, this also means a majority expect to have less than 1,000 locations in 2023, given 18% will still have less than 100 locations. Thus, operators likely expect to take an assertive, but also phased and measured, approach to edge cloud deployment in the medium term.

The basic pattern is the same for U.S. and RoW respondents. Note, however, that U.S. respondents selected higher numbers of locations than RoW in both time periods. For example, 35% of U.S. and 18% of RoW respondents expect to have between 1,000 and 9,999 locations in 2023. ■



Edge computing promises to change the way people function and interact with various services. Healthcare services previously only available in hospitals will be delivered in ambulances or in patients’ homes. Sensors will detect faulty machinery in remote sites and prevent problems before they happen.

Smart cars, smart cities, artificial intelligence/machine learning (AI/ML), and the Internet of Things (IoT) are all within reach as it becomes possible to move workloads away from the network’s core out to its edge, where data can be processed and acted upon – practically in real time.

Service providers are increasing edge deployments for multiple reasons, making it necessary to prepare to support different workload types running in virtual machines (VMs) and/or containers. Red Hat offers a consistent, secure open hybrid cloud foundation for digital service providers to build and deploy edge services – a common infrastructure across the compute, storage, and network footprint, with automated provisioning, management, and orchestration to simplify operations. Get ready with Red Hat and our ecosystem of certified partners to respond to the opportunities that edge computing makes possible – even those not yet imagined.




5G SECURITY STANCE

Author: *Jim Hodges, Chief Analyst, Heavy Reading*

Without question, CSPs must turn up their security game in 5G. To accomplish this will require not only stepping up monitoring and general vigilance, but also new strategies to mitigate the distributed threat landscape that 5G will introduce.

KEY TAKEAWAYS:

-  **"Encrypting live data" (56%) and "aggressive policy scanning" (55%) are identified as the top two components of an effective 5G security strategy.** However, "encryption of stored inactive at-rest data" (48%), "consistent infrastructure provisioning," and "conducting vulnerability tests on platforms" (both 47%) are also key considerations. "Automation" is not far behind (42%).
-  Many operators are still in the development or pre-implementation phase in terms of executing governance, risk, and compliance management strategies, according to survey responses. **While 28% have mature, implementable compliance strategies, many more have yet to start to implement** (23%) or are still developing plans, either with external partners (21%) or without external partners (11%). Perhaps even more telling is that only 9% of respondents have 5G security strategies in production.
-  **Survey respondents ranked having "trust in the physical hardware" (51%) as the most important security focus area.** This was followed by "identity and access management" (40%), "isolation and policy enforcement" (38%), and "visibility into trust status and operations" (35%).
-  **Although many survey respondents believe that secure zero-trust deployment and provisioning are of critical importance, more than half (51%) currently either have "limited familiarity" (37%) or "no familiarity" (14%) with zero-trust concepts.** Perhaps even more telling is that as it stands, only 7% of respondents said their company is currently "implementing a zero-trust based security strategy" in commercial deployments.

5G Security Stance

One reason why 5G will necessitate new security strategies is because it will introduce an unparalleled level of data processing, storage, and encryption at the edge to meet extremely tight end-to-end service latency budgets. As a result, respondents (**Fig 35**) expect to adopt a multifaceted strategy that includes focusing on “encrypting live data” (56%), “aggressive policy scanning” (55%), and the “encryption of stored inactive at-rest data” (48%).

In addition, CSPs will continue to focus on the hardware resources of both “physical and virtual platforms” (47%) to ensure adequate resources are available to meet policy enforcement requirements. They also plan to maintain a strong focus on “patching and conducting vulnerability tests” (47%).

Although “automation” ranked sixth (42%), the gap between third and sixth placed priorities is not that significant. In Heavy Reading’s view, this confirms that automation is already considered a key component of an effective 5G security strategy.

5G security hinges not only on a strong strategy, but also on flawless execution. Many survey respondents are still in the development or pre-development implementation phase in terms of governance, risk, and compliance management.

This is concerning since Heavy Reading believes 5G is less well suited to a “build first and secure second” strategy that CSPs have used in previous generations of mobile technology. The next question (**Fig 36**), for example, shows that while 28% have mature compliance that can be implemented, many more have yet to start to implement (23%) or are still

developing plans with external partners (21%) or without external partners (11%).

The remaining 18% represent opposite extremes of the spectrum, with 9% having in-production security and 9% without any plans or development

activity. In fairness, many service providers have yet to roll out 5G, so they have time to complete development plans and implement. However, Heavy Reading interprets these data points as being less progressive than the pace of 5G deployments will necessitate.

Fig 35. As 5G emerges, with more edge activity and smart devices, how do you plan to evolve your security strategy? (select all that apply) (N=141)


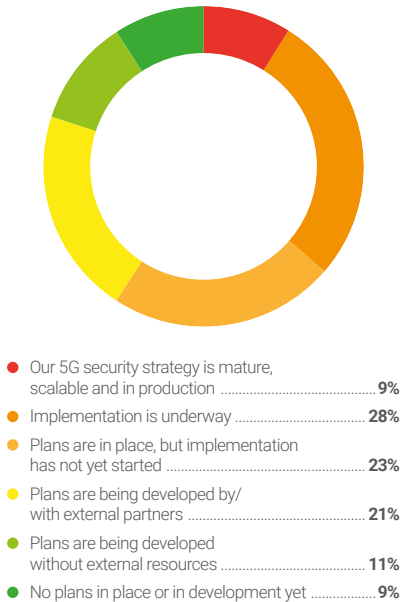


Fig 36. How mature are your plans for governance, risk, and compliance management within your overall 5G security strategy? (N=141)



In addition to governance and compliance, effective 5G security strategies must also consider the infrastructure capabilities on which these policies will run. Essentially, 5G infrastructure must be multi-modal to ensure holistic security coverage.

Based on the “critical” responses, a number of infrastructure areas stand out (Fig 37). Of these, having “trust in the physical hardware” (51%) was considered the greatest concern. This was followed by “identity and access management” (40%), “isolation and policy enforcement” (38%), and “visibility into trust status and operations” (35%). The close rankings of these three areas was not unexpected since 5G networks, particularly the 5G core, are policy-driven and place great emphasis on identity management (both human and non-human) as well as creating specific trust areas to execute sliced-based services.

The 5G infrastructure requirements documented above apply to both

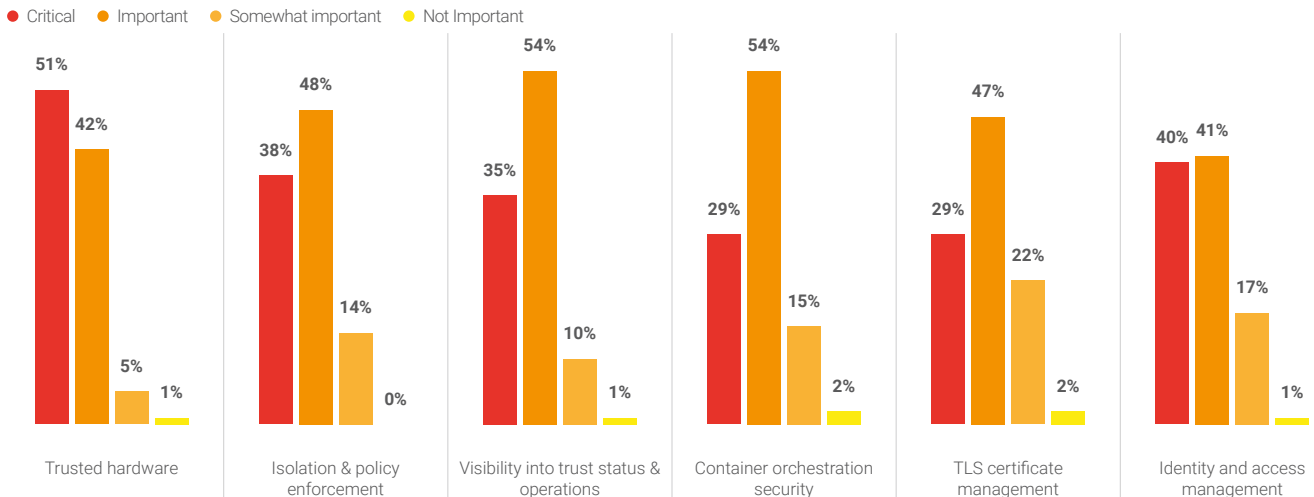
centralized and edge clouds. However, the edge cloud (Fig 38) has unique security requirements related to pushing compliance and risk management to remote devices before they can unleash threat vectors.

Yet, the survey respondents view the technical fundamentals as still very similar in that the foundational focus area continues to be utilizing “trusted hardware” (54%) to ensure policy deployed at the edge is consistent with the “global security posture” (50%).

However, the third- and fourth-ranked attributes – “root of trust for remote devices” (41%) and “integrating edge security best practices with existing security incident procedures” (40%) – do capture that some aspect of the security postures are different at the edge relative to centralized infrastructure.


In order to adequately address edge security requirements, there must also be a strong focus on the management of edge devices. (Fig 39) As with security

Fig 37. For your 5G infrastructure, how important are the following security areas? (N=138-140)



“WITHOUT QUESTION, CSPs MUST TURN UP THEIR SECURITY GAME IN 5G. TO ACCOMPLISH THIS WILL REQUIRE NOT ONLY STEPPING UP MONITORING AND GENERAL VIGILANCE, BUT ALSO NEW STRATEGIES TO MITIGATE THE DISTRIBUTED THREAT LANDSCAPE THAT 5G WILL INTRODUCE.”

Fig 38. How important are the following security capabilities for securing the edge? (N=139-140)



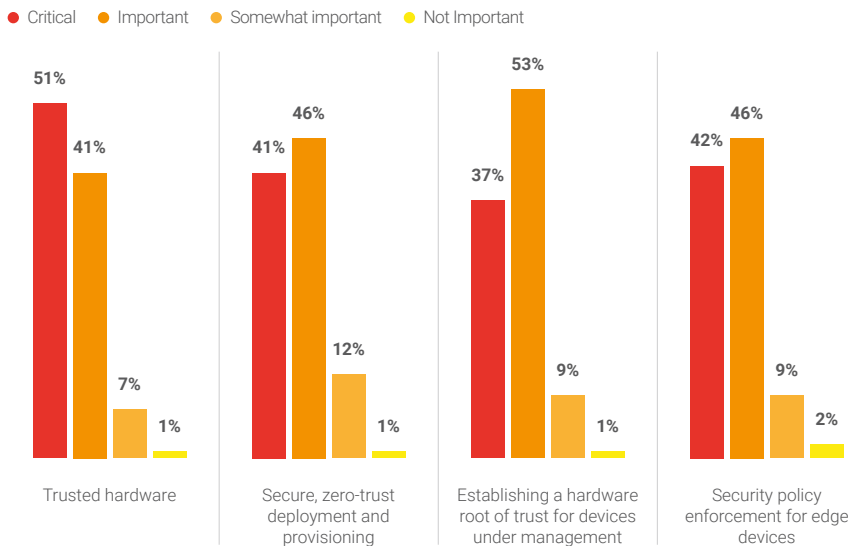
infrastructure, the survey respondents indicated they believe the leading consideration based on “critical” response levels is that the devices must be deployed on “trusted hardware” (51%).

However, a number of other considerations are critical as well. Of these closely ranked capabilities, the ability to support “policy enforcement for edge devices” (42%) ranked second. Not far behind were “secure zero-trust provisioning” (41%) and “establishing a hardware root of trust for devices under management” (37%).

Heavy Reading interprets these findings as confirming that CSPs will initially focus on hardware platforms both for infrastructure and devices. Yet, in the device realm, they are implementing advanced security policy enforcement capabilities that apply zero-trust principles to edge devices.

Support of zero-trust was one area where the U.S and RoW respondents


Fig 39. Please rate the importance of the following security capabilities of edge devices. (N=137-140)



had notable deviation in their security question responses. In this case, while 46% of U.S. respondents felt a “trust no one, authenticate everything” zero-trust provisioning approach was critical, only 34% of RoW respondents felt this way.

As noted above, roughly 4 out of 10 (41%) survey respondents believe that “secure zero-trust deployment and provisioning” is of critical importance. Yet, half (51%) either have “limited familiarity” (37%) or “no familiarity” (14%) with zero-trust concepts (Fig 40).

Fig 40. To what extent do you plan to employ zero trust security concepts for commercial deployments? (N=139)



It can be argued that the majority of CSP employees do not need to understand zero-trust concepts because they typically rely on their security colleagues, which typically represent a very small percentage (usually less than 8%) of the employee base. However, in the 5G era, Heavy Reading believes all employees will need to have a greater understanding of advanced security fundamentals, including zero-trust.

Interestingly, even though U.S. respondents place a higher value on zero-trust, when looking at the familiarity level of zero-trust concepts, the splits are very similar (36% U.S. vs. 38% RoW), which confirms that limited knowledge in this area is a global concern. This is perhaps a reason why only 7% of respondents (9% U.S. and 5% RoW) are currently implementing a zero-trust strategy in commercial deployments. ■



IT security has traditionally been focused on fortifying, maintaining, and policing the data center perimeter – but today that perimeter is dissolving. The way we develop, deploy, integrate, and manage IT is dramatically changing. Public and hybrid clouds are redistributing responsibility for regulatory compliance and security across geographical, sovereign, and vendor borders.

The adoption of containers at scale requires new methods of analyzing, securing, and updating the management of infrastructure and delivery of applications. Mobile apps are spread across a multitude of devices, and more and more infrastructure is moving from hardware to software. This device and infrastructure proliferation is contributing to the complexity of networks as they extend into hostile environments.

The traditional ways of managing security aren't keeping up. Digital transformation demands a change in security programs – security must be continuous, integrated, and flexible in a digital world.

Red Hat wants you to have confidence as you adopt a continuous security strategy. We do that by making open source ready for large-scale production. The goal is to help your business remain competitive, flexible, and adaptable while maintaining security and regulatory compliance. Red Hat knows the landscape and how to innovate within it. We work with you and for you. Our unique subscription model gives digital service providers access to a dedicated team of experts who support our technology 24x7. Visit redhat.com/security to learn more about Red Hat's commitment to protecting customer data and privacy.