

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

**HEAVY  
READING**  

---

**WHITE  
PAPER**

# The Journey to Cloud Native

*A Heavy Reading white paper produced for Juniper, Nokia, and Red Hat*

**JUNIPER**  
NETWORKS

**NOKIA**



**Red Hat**

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

---

## INTRODUCTION

Cloud native is a tough term to pin down. Heavy Reading sees the goal of cloud native as enabling enterprises to build and run highly scalable and flexible applications for deployment in a public, private, or hybrid cloud. This is accomplished through the use of containers, service meshes, microservices, DevOps and continuous integration/continuous deployment (CI/CD) development practices, and declarative APIs. The key benefits of cloud native are listed below:

- Faster time to market (TTM) for new services and applications
- Improved, automated, and comprehensive lifecycle management (LCM)
- 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 architecture and the use of CI/CD
- Lower total cost of ownership (TCO) through the use of containers and microservices, enabling users to deploy only what is needed, rather than entire monolithic network functions

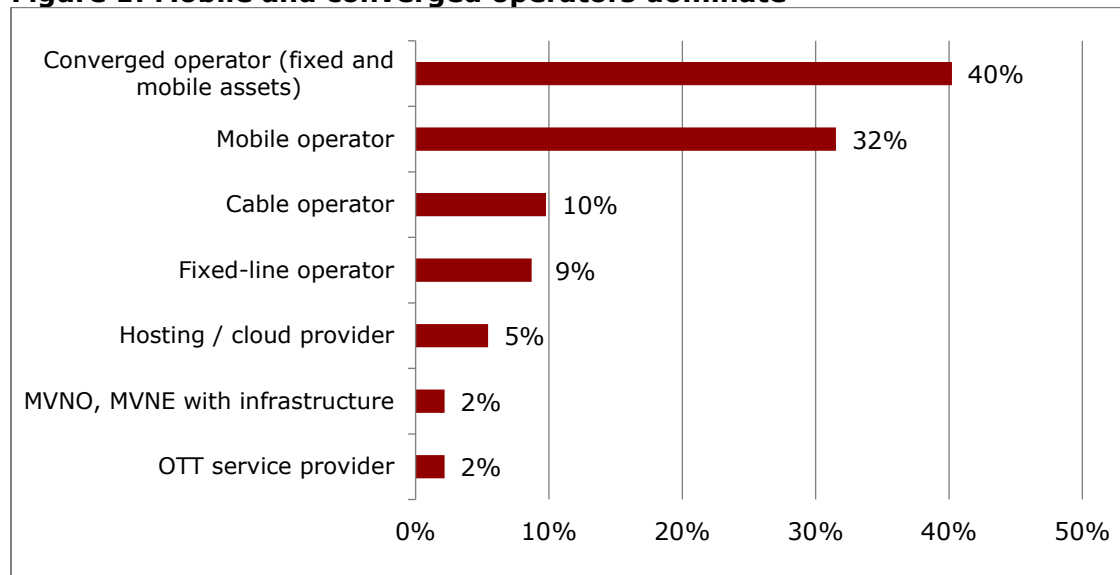
As compelling as these benefits may seem, cloud native represents a fundamental change in the way communications service providers (CSPs) design, deploy, and manage applications and services. Heavy Reading wanted to understand what was motivating the CSPs to adopt cloud native, where they were in the journey, what obstacles they were encountering, and how they were making their choice of cloud native solutions providers and partners.

To answer these questions, Heavy Reading conducted a survey of 92 global CSPs during 3Q21. This report presents the highlights of the survey.

## DEMOGRAPHICS

Mobile and converged operators made up the bulk of Heavy Reading's survey respondent pool, accounting for almost three-quarters of overall responses (see **Figure 1**). An additional 19% came from the fixed-line and cable operator community. The remaining 9% hailed from hosting/cloud providers, mobile virtual network operators (MVNOs)/mobile virtual network enablers (MVNEs) with infrastructure, and over-the-top (OTT) service providers.

**Figure 1: Mobile and converged operators dominate**



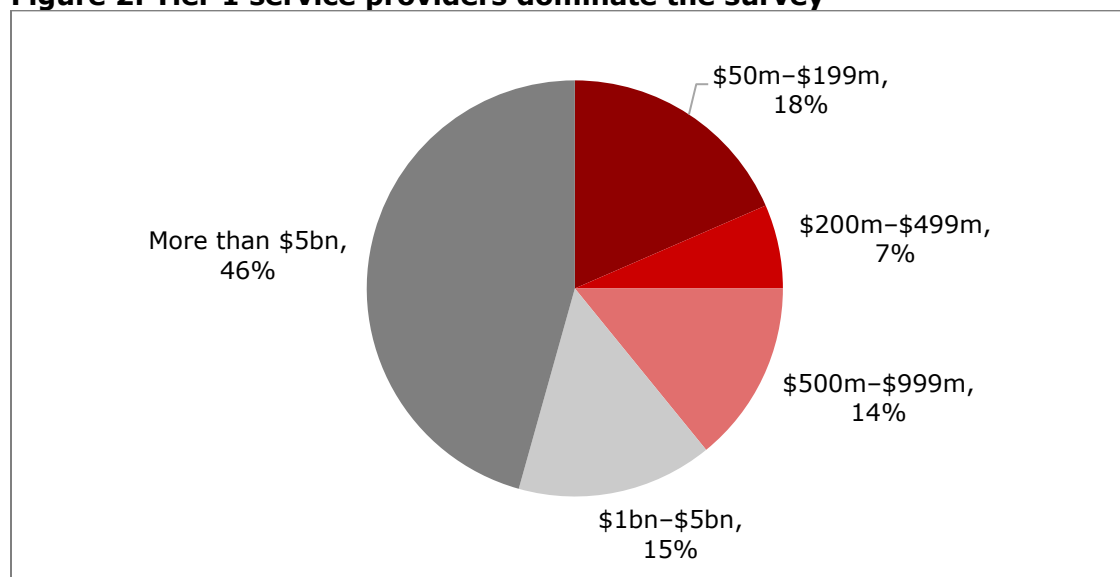
n=92

Q: What type of service provider do you work for?

Source: Heavy Reading

Close to half of respondents represented very large CSPs with annual revenue of over \$5bn (see **Figure 2**). CSPs with revenue of between \$500m and \$5bn made up 29% of the respondent pool, and those with revenue between \$50m and \$499m comprised the remaining 25%. Revenue dictates the capital budget available for funding the transition to 5G, edge computing (EC), and cloud native networking. Heavy Reading research shows that carriers dedicate, on average, 17–18% of revenue to capex.

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



n=92

Q: What is your company's approximate annual revenue (\$USD)?

Source: Heavy Reading

---

## Regional breakdown

Just under half of the survey respondents are from the US. Europe, the Middle East, and Africa account for 20% of respondents. The remainder of North and South America make up 19% of respondents. The remaining 13% of the respondents are from the Asia Pacific region.

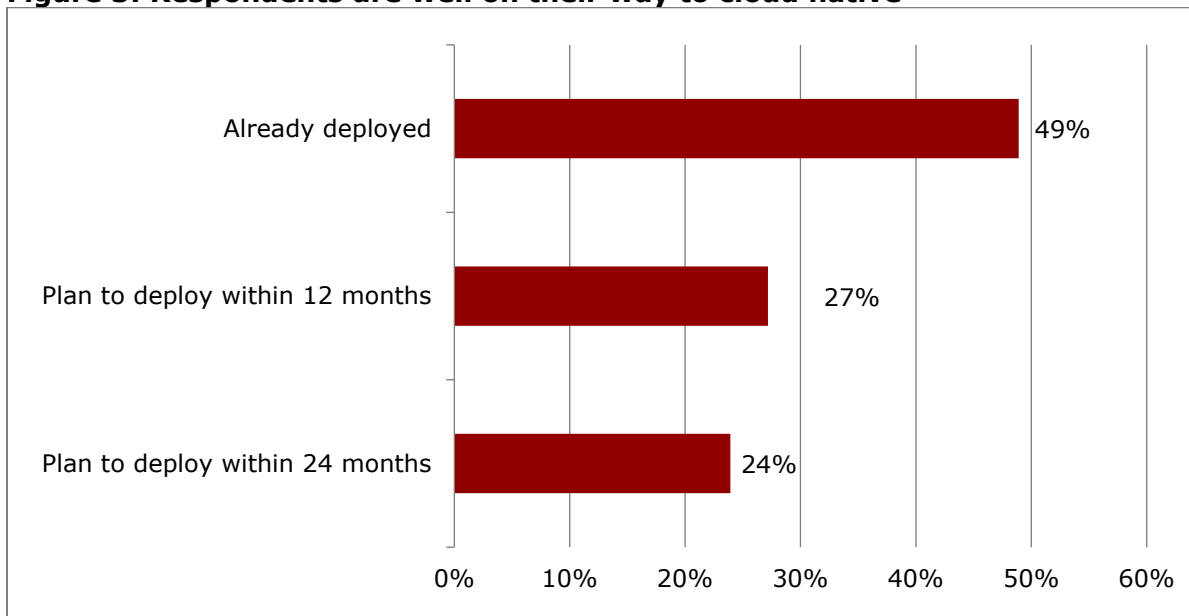
## Job function

As is the case with most Heavy Reading surveys, two-thirds of the respondents (66%) were from technical networking roles: planning and engineering, R&D, and network operations. Almost one-fifth were from management and marketing. The data center, IT, and cloud made up 10% of respondents. The remaining 5% came from finance, security, and software engineering.

## Cloud native plans are already underway

Heavy Reading was interested in understanding the details of cloud native deployments in terms of drivers, challenges, goals, and more, so we restricted the respondent pool to CSPs with plans to deploy cloud native technology within 24 months. They were not hard to find. Half of the respondents have already deployed a cloud native platform, and the remaining half plan to do so within the next 24 months (see **Figure 3**).

**Figure 3: Respondents are well on their way to cloud native**



n=92

Q: Has your organization deployed or does it plan to deploy a cloud native telco or edge cloud platform?

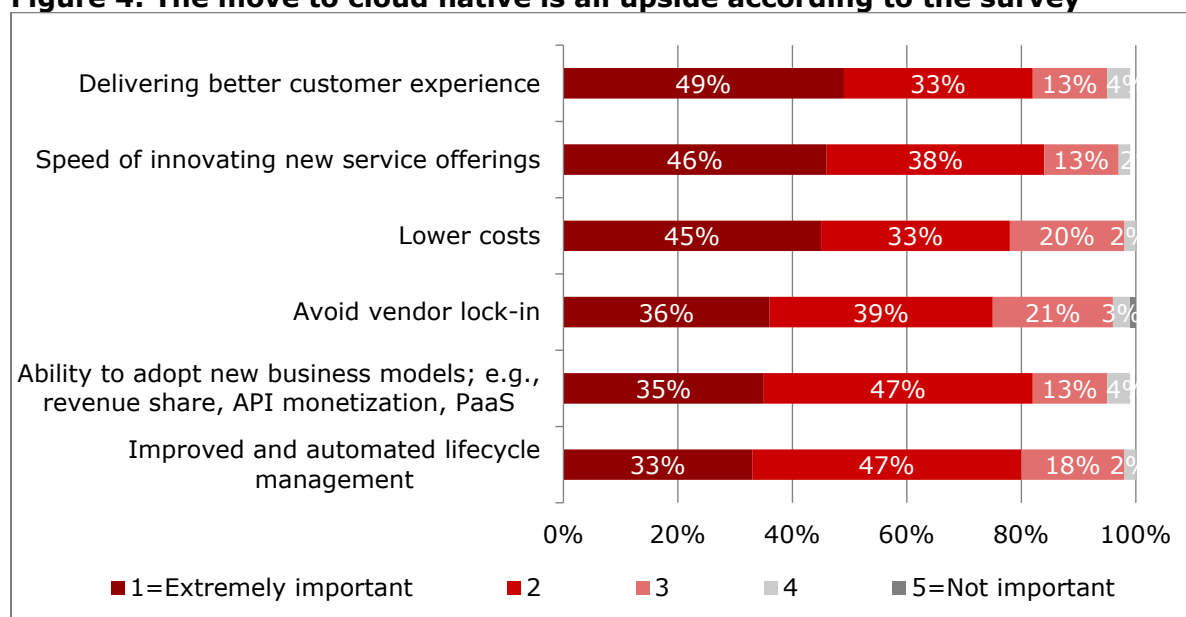
Source: Heavy Reading

## CLOUD NATIVE BENEFITS AND DRIVERS

Heavy Reading survey respondents are expecting a lot of return on investment in cloud native, led by a better customer experience and speed of innovation (see **Figure 4**). All benefits were considered important, as shown by an almost complete absence of option 5 (“not important”) votes and an extremely paltry showing for option 4 votes. “Avoiding vendor lock-in” had the fewest number of combined 1 (“extremely important”) and 2 votes, and even there, they added up to a robust 75% of respondents.

Differences in how these benefits are perceived emerge when looking at the different segments of the survey base. Respondents who have already adopted cloud native technology find all of the benefits to be more important, ranging from 84% combined “very important/important” votes for avoiding vendor lock-in to 94% of respondents for speed of innovation. This is a clear indicator that hands-on experience only serves to underscore the benefits overall of cloud native solutions.

**Figure 4: The move to cloud native is all upside according to the survey**



n=91

Q: How important to your company are these benefits of moving to cloud native?

Source: Heavy Reading

### Services and applications driving cloud native

EC has emerged as the dominant driver of a move to cloud native computing, according to the survey respondents (see **Figure 5**). The microservices architecture of cloud native is key to managing the costs, both capex and opex, of a highly distributed edge storage and compute architecture. A microservices-based architecture, together with automation and artificial intelligence (AI)/machine learning (ML), will enable the rapid rollout of EC with zero-touch provisioning and lights-out maintenance.

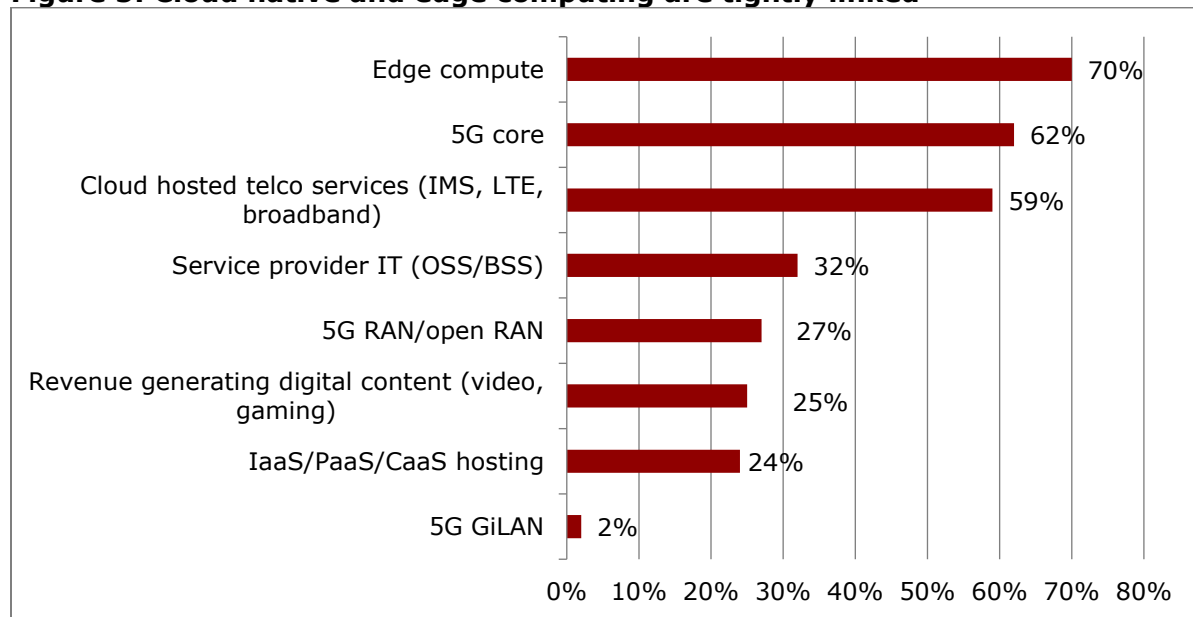
The 5G standalone core is assumed to be cloud native. As CSPs move forward from their current LTE+, non-standalone 5G cores, cloud native will be part of that transition, so it is not surprising that the respondents ranked 5G core second in the Heavy Reading list of cloud native drivers.

Cloud-hosted telco services, which came in third in the list of drivers, require a virtualized environment and will benefit from everything that cloud native brings with it: faster TTM, increased flexibility, improved LCM, and lower TCO.

These first three drivers correlate to leading revenue focus areas for the CSPs (edge and private LTE) or, in the case of 5G core, elements that are part of the evolution to 5G. There is a sharp drop-off in interest for the next group of drivers. Most of these actually depend on the deployment of EC and will increase in interest as that moves into production. In the case of 5G RAN/open RAN, many believe that edge devices will evolve in the near term to include the distributed unit (DU) in a 4G and 5G open RAN architecture. CSPs are currently implementing virtualized RAN to support centralized unit (CU) functions. Getting the DU costs down while incorporating both automation and LCM is a challenge that is pushing out the timeline for cloud native 5G RAN deployment, in terms of volume production implementations, by 12–24 months.

The very low ranking of 5G GiLAN makes Heavy Reading question whether GiLAN will survive the transition to 5G or whether the same service-chained, virtualized functions of the firewall, carrier-grade NAT, DPI, and so on will be completely rearchitected (and renamed) for a 5G microservices-based environment.

**Figure 5: Cloud native and edge computing are tightly linked**



n=92

Q: What are the top three applications or services driving your cloud native plans?

Source: Heavy Reading

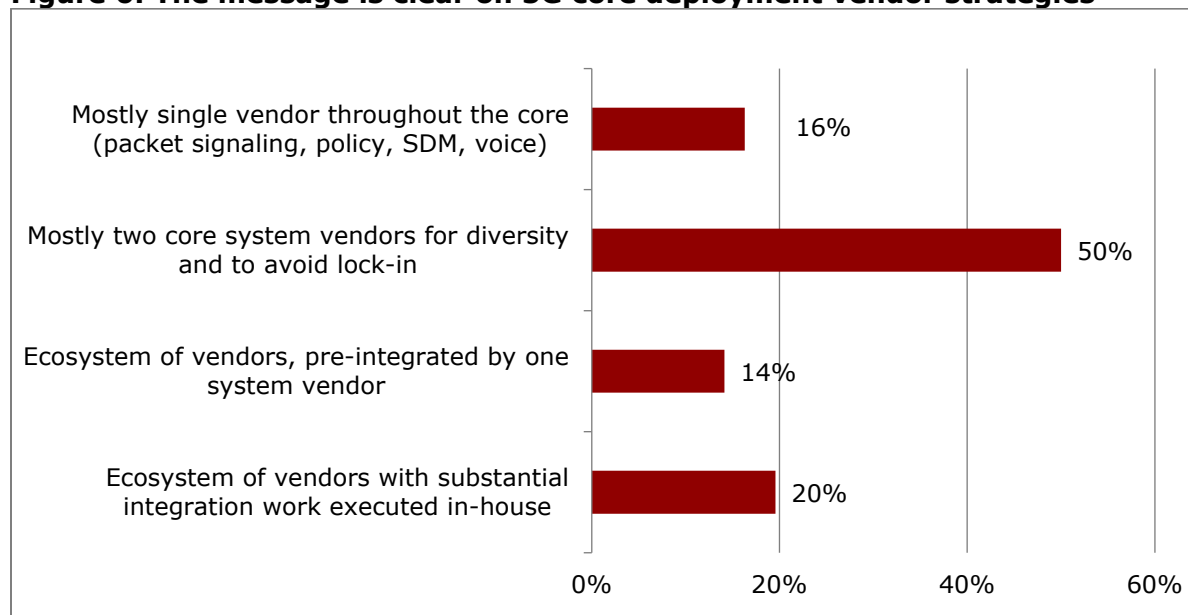
## 5G core deployment strategy

As noted above, the 5G core will be implemented in a cloud native environment. Heavy Reading wanted to understand a little more about how the CSPs would be deploying the 5G core in terms of vendor strategy. It is clear from **Figure 6** that CSPs want to avoid a single-system-vendor strategy, whether that vendor is offering its own fully integrated solution or pre-integrating a complete solution with ecosystem partners. CSPs would rather work with multiple vendors, create their own ecosystem, and do the integration work themselves.

However, the CSPs' preferred choice, by far, is working with two main system vendors for diversity and to avoid vendor lock-in. This is the current vendor strategy in the mobile network, and Heavy Reading's survey respondents are not inclined to change.

The only cut of the data the two-vendor strategy drops to second place is when Heavy Reading focuses on the 26 respondents from the US who have already deployed cloud native solutions; 42% of those respondents selected the ecosystem of vendors with in-house integration versus 38% who preferred the two-system-vendor strategy. Clearly, if there is no viable solution currently available, these large CSPs with significant in-house resources will build the integrated solution themselves. There is no cut of the data where either of the one-vendor strategies moves up in ranking.

**Figure 6: The message is clear on 5G core deployment vendor strategies**



n=92

Q: How have you deployed, or will you deploy, the 5G core to ensure a consistent cloud native architecture and accelerated deployments?

Source: Heavy Reading

---

## CHALLENGES TO IMPLEMENTATION

Heavy Reading asked our survey pool two questions about the challenges they are facing with the implementation of cloud native solutions. The responses led to a couple of overall observations:

- None of the challenges pulled in over half of the survey respondents, demonstrating that there are no overwhelming obstacles to implementation.
- The responses from those that have already implemented cloud native solutions are almost all lower as implementation progresses, the challenges lessen, and there appear to be no “gotchas” (see **Figure 7**).

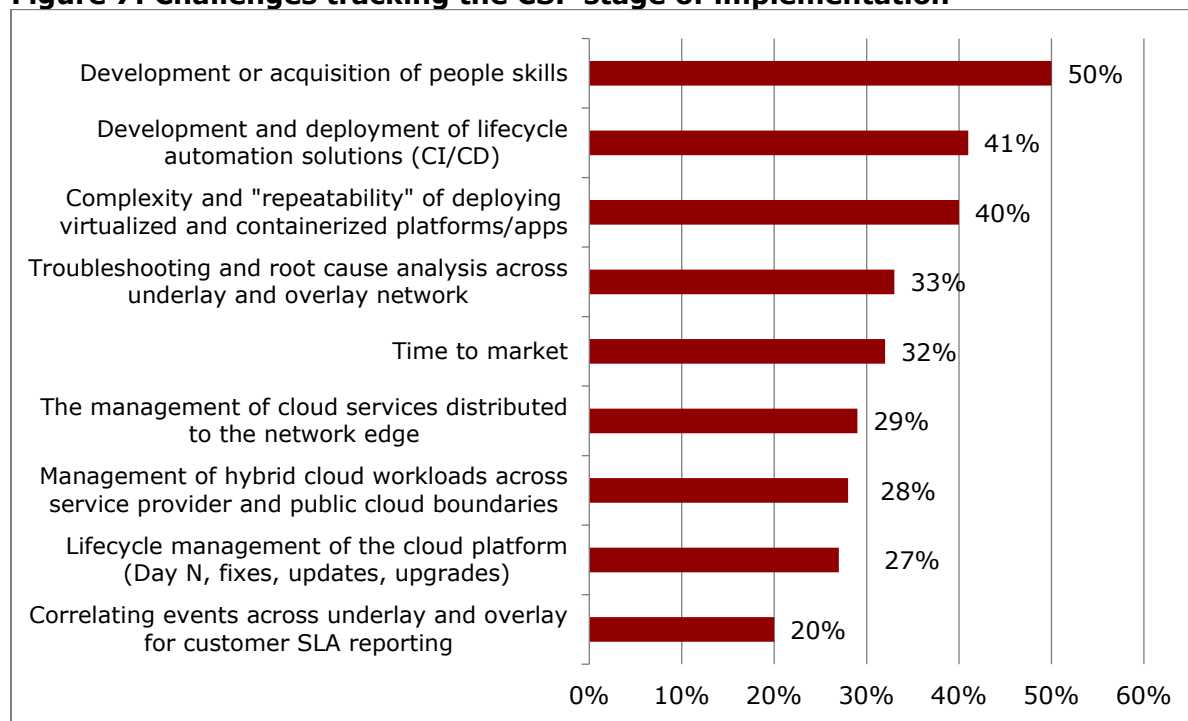
The top response reflects the tech industry overall: we do not have the people needed to support the evolution of technology, in this case, the transition to cloud native. The second most popular response is an elaboration on this point: we do not have people trained in a CI/CD style of development to provide us with cloud native applications.

The third response, “complexity and ‘repeatability’ of deploying virtualized and containerized platforms/apps,” can be paired with the eighth “lifecycle management of the cloud platform (Day N, fixes, updates, upgrades).” One is regarding deployment, and the other is regarding ongoing management. Similarly, the fourth response, “troubleshooting and root cause analysis across underlay and overlay network,” can be paired with the least most popular response, “correlating events across underlay and overlay for customer SLA reporting.” One concerns management across both the underlay and overlay network, and the other is correlating those results for customers’ service level agreement (SLA) reporting. The order of these challenges appears to be just as much associated with where the CSP is in the deployment of cloud native as it is with the inherent complexity of the challenge. The fact that all are clustered within 30 percentage points of each other supports this observation.

The responses reveal additional nuances when looking at specific regions or revenue levels. The most extreme difference that Heavy Reading found was between the response to “management of hybrid cloud workloads across service provider and public cloud boundaries,” with 43% of those with revenue over \$5bn finding this to be a challenge. Only 16% of the respondents with revenue under \$5bn responded the same, demonstrating that smaller CSPs are not currently grappling with multidomain, multicloud issues.



**Figure 7: Challenges tracking the CSP stage of implementation**



n=92

Q: Select the top three challenges your company is facing in deploying a cloud native architecture.

Source: Heavy Reading

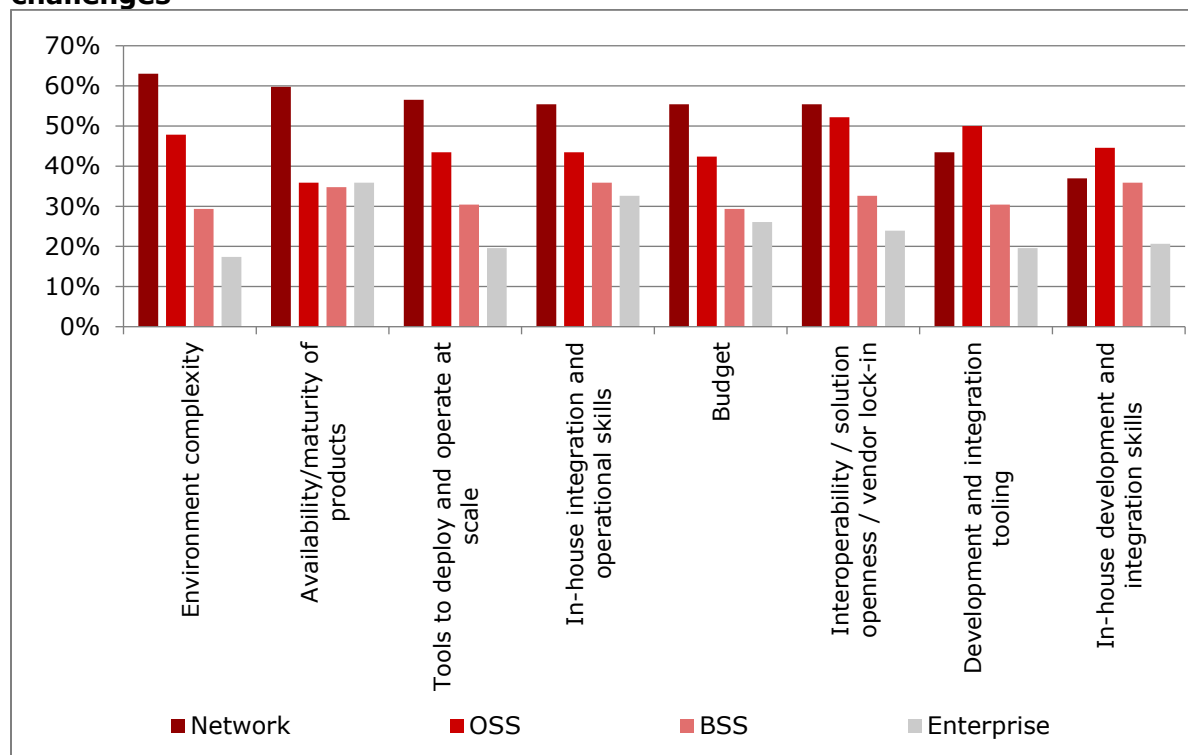
Looking for additional detail on challenges, Heavy Reading queried the survey pool about the areas of deployment where they were experiencing challenges to going cloud native, such as the network, operations support system (OSS), business support system (BSS), or enterprise (see **Figure 8**). Heavy Reading established earlier in the survey that cloud native will be deployed first for network workloads: 59% of respondents have already deployed it and 32% will within 12 months. Respondents plan to transition workloads to the OSS business areas next, then BSS, and last, enterprise, where 32% of respondents will wait 1–3 years to deploy cloud native.

In almost all cases, respondents ranked the challenges in that same order: first network, then OSS, BSS, and enterprise. The only challenges that were considered more severe in an area *other* than the network were “in-house development and integration skills” and “development and integration tooling,” where the OSS space was recognized as a greater challenge than the network. This is not surprising given that most Tier 1 carriers have dozens of OSSs in operation and they do most of the integration work between systems internally. A migration to cloud native brings with it the following changes:

- The move to microservices.
- Standardized access to these microservices via API exposure.
- The integration of multiple operational layers and domains for application management.
- Automation across the application lifecycle through the use of DevOps.

These are profound changes to the application development and management environment of the CSPs and are challenging to tackle with internal resources.

**Figure 8: The network space is seeing the most implementations and the most challenges**



n=92

Q: In which business areas are you experiencing significant challenges to going cloud native? Check all that apply.

Source: Heavy Reading

Looking only at the survey results from respondents who have already deployed cloud native, Heavy Reading notes that there are some differences compared to the rest of the survey base. In the network area, “tools to deploy and operate at scale” is a greater challenge by 11 percentage points compared to respondents planning to deploy cloud native in 1–2 years.

“Budget” in the OSS area plummets between those who have not yet deployed cloud native, with 57% of respondents considering it a challenge, and those who have already deployed cloud native, with only 27% of the respondents finding it to be a concern.

Those who have already deployed cloud native also consider *all of the challenges in the enterprise area to be greater* than the survey base as a whole. This suggests, again, that they have firsthand experience with the challenges of implementing cloud native in the network area and are starting to gain experience with these challenges in the enterprise space.

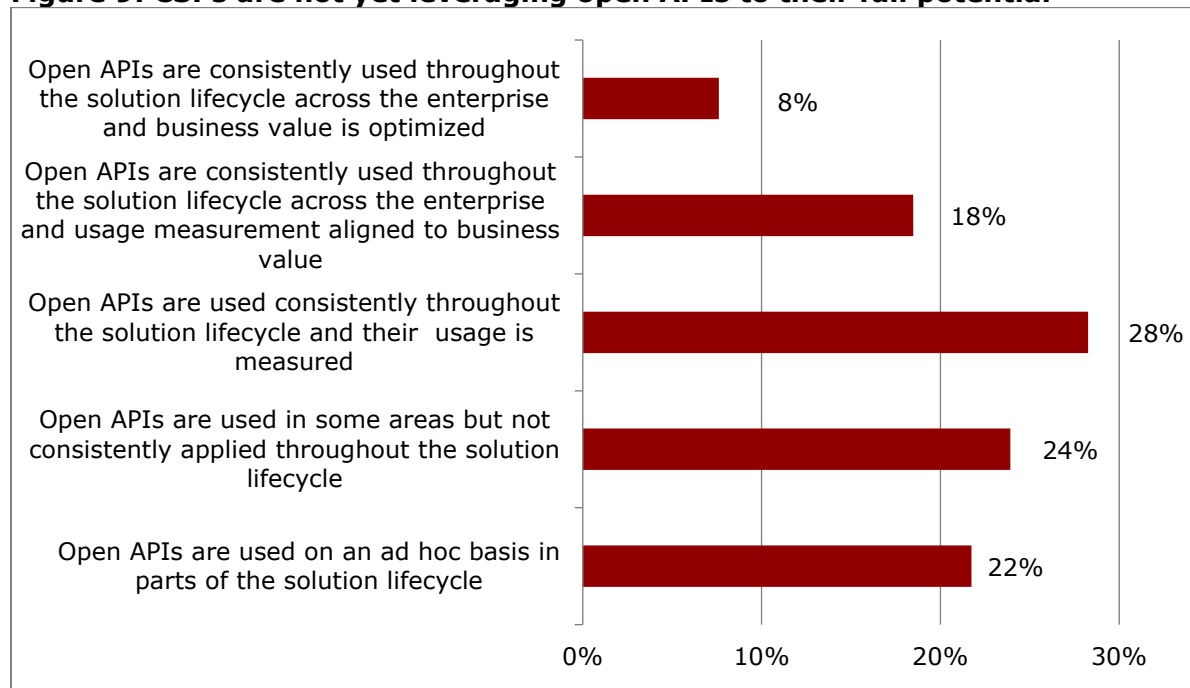
## OPEN APIS KEEP CLOUD NATIVE RUNNING

Cloud native applications are segmented into microservices under the assumption that these microservices can be deployed separately or called by other applications by using well-defined open APIs. These APIs are critical to the seamless operation of cloud native applications, simplifying interoperability and, by enabling and simplifying microservice reuse, accelerating design and CI/CD-based development. APIs can also be used to capture and export telemetry data from cloud native applications and microservices for improved LCM.

Open APIs play a significant role throughout the lifecycle of cloud native applications. To understand the degree to which the survey respondents were leveraging open APIs, Heavy Reading asked them the question shown in **Figure 9**. Taking advantage of new capabilities, the survey pool, as a whole, uses open APIs, with the largest percentage managing to use them consistently throughout the application lifecycle and measure their usage. However, 46% (the bottom two bars) use them only in some areas or on an ad hoc basis. It is encouraging that one-quarter of the respondents are more advanced in their use of open APIs (the top two bars).

The only cut of the survey data that shows a more optimal use of open APIs is that portion of respondents who have already implemented cloud native applications. For that segment, the top two bars increase by 7 percentage points.

**Figure 9: CSPs are not yet leveraging open APIs to their full potential**



n=92

Q: How do you use open APIs in digital projects for sustainable applications and infrastructure?

Source: Heavy Reading

---

## CHOOSING A CLOUD NATIVE SOLUTION

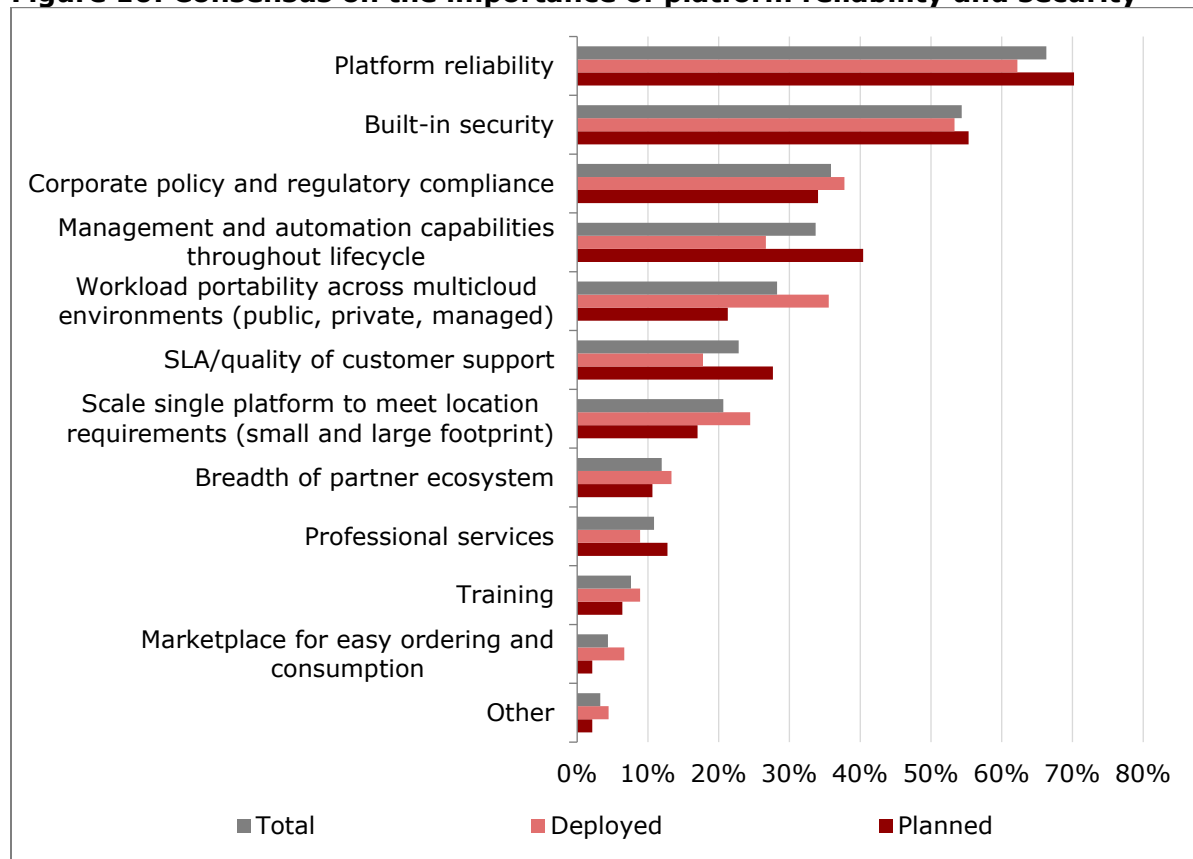
Heavy Reading asked the survey group what they looked for when selecting a cloud native solutions provider (see **Figure 10**). Platform reliability and security were the top picks, regardless of region, company size, or deployment status. They were also the top picks by a wide margin. Note that there is a difference of 30 percentage points between the first and third picks. Everything else is clustered. There is a similar difference between the third pick and the last, or twelfth pick (“other” at 3%), suggesting that there is not widespread agreement on the relative importance of the other criteria listed.

Dividing the results by those who have already deployed cloud native and those planning to do so in the next 24 months, Heavy Reading sees some fairly sharp differences, as shown in **Figure 10**. In “workload portability,” for example, 36% of respondents who have already implemented cloud native cite it as a top criterion, while only 21% of those with cloud native still on the drawing board do the same. Conversely, “management and automation” is important to just 27% of cloud native deployers, but to 40% of those who have not yet deployed cloud native. There are sharp regional differences as well. “Scale single platform to meet region requirements” is a top criterion to 27% of US respondents, but to only 8% of those from Asia Pacific.

“Corporate policy and regulatory compliance” is more important to US respondents, to the large Tier 1 companies (>\$5bn), and to those who have already deployed cloud native. The more experience a CSP has with integrating cloud native technology and cloud-based applications into the enterprise environment, the more familiar they become with the complex and demanding world of cloud regulatory compliance, which is closely tied to security and to moving to a microservices-based cloud native environment. The following is a partial list, only, of cloud regulations that must be followed when transitioning workloads to the cloud:

- Federal Information Security Management Act (FISMA), US
- Health Insurance Portability and Accountability Act of 1996 (HIPAA), US
- Payment Card Industry Data Security Standard (PCI DSS)
- General Data Protection Regulation (GDPR), EU
- Sarbanes-Oxley Act of 2002 (SOX), US
- Federal Risk and Authorization Management Program (FedRAMP), US

**Figure 10: Consensus on the importance of platform reliability and security**



n=92

Q: What are the top three evaluation criteria when choosing cloud native solutions providers?

Source: Heavy Reading

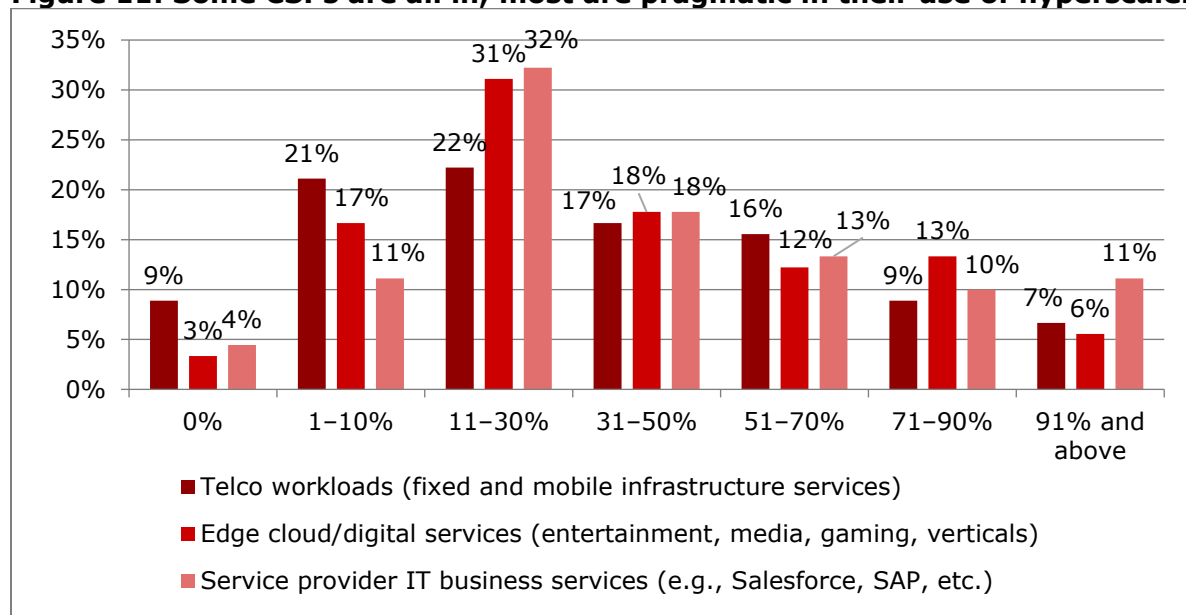
## The role of the hyperscalers

The last three years have seen the explosion of partnerships and alliances between telcos and hyperscalers. The questions are: What percentage of the telcos' workloads would they deploy on hyperscalers' clouds, and would the hyperscalers significantly improve the telcos' ability to create and deploy high margin consumer and enterprise services? Heavy Reading's survey captures the current thinking among the CSPs regarding the first question (see **Figure 11**). The answer to the second question remains to be seen.

CSPs are more than receptive to working with hyperscalers. Heavy Reading had a total of 16% of respondents across workloads claim that they would not work with hyperscalers at all, compared to 24% on the other end of the spectrum, who anticipate handing over 91% or more of their workloads. The percentage of respondents avoiding hyperscalers is almost identical when comparing US to non-US respondents. The percentage of respondents turning 91%+ of their workloads over are, with one exception, in the US.

CSPs appear comfortable working with hyperscalers, using them for up to, in general, 50% of their workloads. They appear, overall, to be more willing to transition their IT business and edge cloud/digital services workloads and less enthusiastic about moving their telco workloads.

**Figure 11: Some CSPs are all in; most are pragmatic in their use of hyperscalers**



n=92

Q: To succeed in your cloud native transformation, what percentage of these workloads will be run within a hyperscaler cloud (e.g., AWS, Google, Azure)?

Source: Heavy Reading

## SIGNPOSTS ALONG THE WAY POINT TO FULL CLOUD NATIVE DEPLOYMENTS

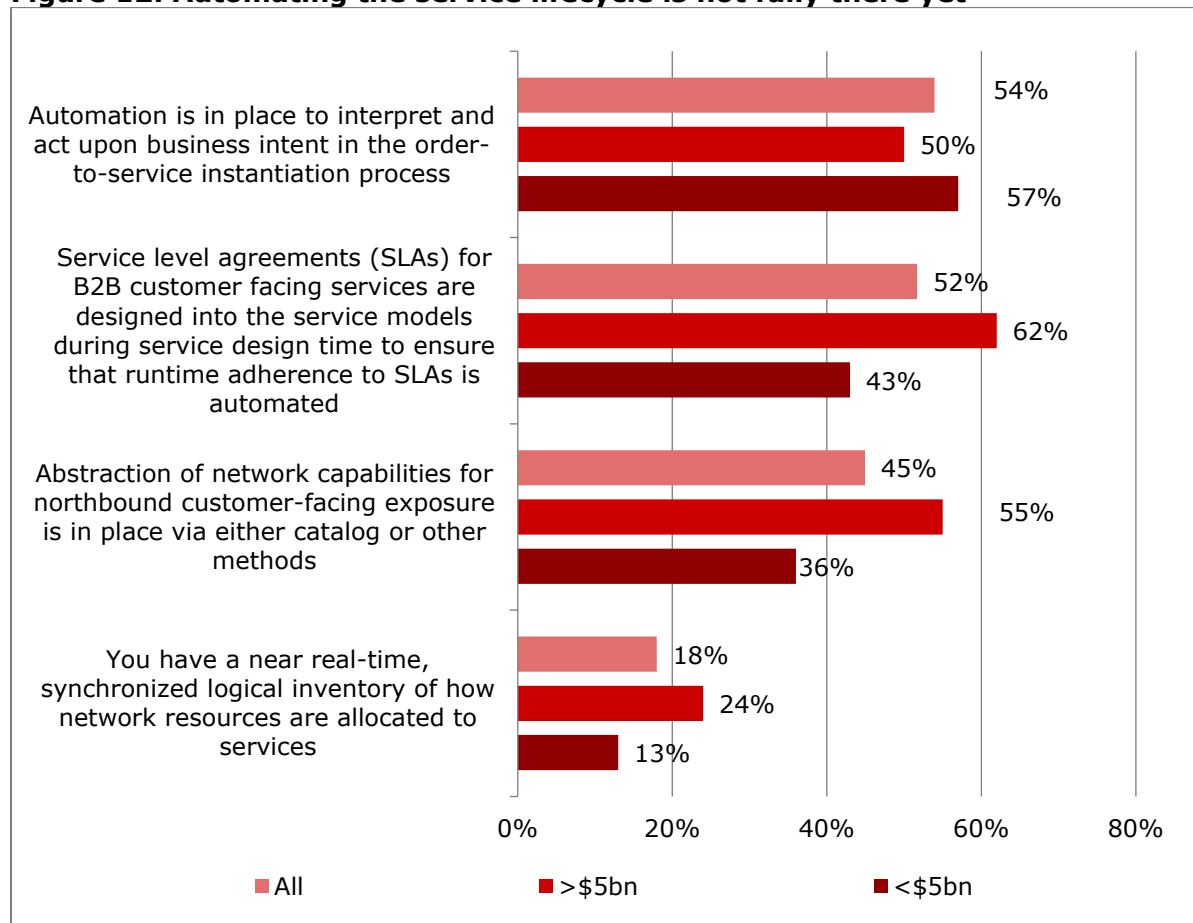
Cloud native adoption is more than the deployment of container-based, microservices-architected applications. Throughout the survey, Heavy Reading sought to capture some of the details that go into a cloud native deployment, such as how open APIs are used, what challenges for which workload CSPs were facing, vendor strategies, solution evaluation strategies, and more. The larger survey consisted of 24 questions. Heavy Reading presents in this report only the highlights of the survey. In the final question in this report, and one of the last in the survey, Heavy Reading wanted to get a feel for where the survey respondents were in their implementation of solutions and technologies that underpin a CSP's cloud native deployment, particularly when it comes to extending the benefits of cloud native solutions to the end customer—benefits that reduce TTM, increase flexibility, and improve LCM.

About half of those surveyed responded that they are able to abstract network capabilities for northbound, customer-facing exposure; that automation is in place to improve, simplify, and accelerate the order-to-service instantiation process; and that runtime SLA data is automatically captured for reporting to B2B customers (see **Figure 12**). Less than 20%, however, reported the ability to do a near real-time, synchronized logical inventory of how network resources are allocated to services. This is a tough task in the cloud native world,

where resources, both physical and virtual, can be spun up or down at will, and the ability to monitor and troubleshoot the environment is dependent on being able to associate which resources are allocated to which services. It becomes particularly challenging in a multicloud environment where visibility into resources outside the CSP's control is limited.

These are very challenging problems to tackle. The responses are statistically consistent throughout the survey base, regardless of region or cloud native deployment status. However, Heavy Reading does see a stronger response to these statements from the larger, >\$5 bn CSPs, as shown in **Figure 12**. Not surprisingly, large Tier 1 companies are significantly further along in automating the service lifecycle, particularly in SLA automation and the abstraction of network capabilities.

**Figure 12: Automating the service lifecycle is not fully there yet**



n=89

Q: How far have you progressed in automating the interworking between the business layer and the network layer with regard to the design, deployment, assurance, and decommissioning of customer-facing services (e.g., the end-to-end [E2E] service lifecycle)? Check all that you have addressed.

Source: Heavy Reading

---

## CONCLUSIONS

Heavy Reading's cloud native survey reveals key insights into CSP strategies. Despite the list of compelling benefits of cloud native in the introduction to this report, CSPs recognize that the most important benefit is improving customer experience. They are moving forward with implementing cloud native, initially driven by EC and private LTE, along with the 5G core. However, looking at some of the details that show mature cloud native adoption, such as automating the E2E service lifecycle or fully leveraging the capabilities of open APIs, it is clear that CSPs have a long road in front of them.

CSPs have decided, however, who they want to work with in their transition to cloud native solutions. They are increasingly comfortable working with hyperscalers and running the cloud native workloads in hyperscaler clouds. But they are adamant about not wanting to work with just one system vendor, thereby leaving themselves vulnerable to vendor lock-in.

The lack of technical resources and the necessary transition to a CI/CD style of application development are obstacles. Nevertheless, they are not deflecting carrier focus from their transition to cloud native. The Heavy Reading survey base has already started this journey. The requirement for cloud native by critical network infrastructure components, such as the 5G core, may have started CSPs on the path, but the benefits, such as faster TTM, improved LCM, and lower TCO, are ensuring that they keep moving down the road to cloud native.