analysys mason

Perspective

Evaluating private versus public cloud models for CSPs' cloud-native mobile core deployments

June 2024

Ameer Gaili

Contents

List of figures

This perspective was commissioned by 5GDNA. Usage is subject to our disclaimer and copyright notice. Analysys Mason does not endorse any of the vendor's products or services.

1. Executive summary

Communication service providers (CSPs) have an important opportunity with 5G standalone (SA) to give their 5G investments a boost and to drive new revenue growth. CSPs can use 5G SA to establish a truly cloud-native mobile core to enhance network performance, enable new consumer and enterprise use cases, improve customer service and reduce total cost of ownership (TCO). CSPs are committed to this path and have strong plans to adopt cloud-native and containerised mobile cores (across 5G SA, 5G non-standalone (NSA), legacy (4G/LTE), private 5G, and IP multimedia subsystems (IMS)) to position themselves for future revenue growth and innovation in areas such as AI.

However, CSPs face the challenge of choosing between private or public cloud as the ideal deployment model for their cloud-native cores. The public cloud has emerged as a potential option for CSPs in recent years, but its suitability for deployment first requires a thorough evaluation of all technical and strategic factors, as well as a deep analysis of TCO, not just an analysis of capex.

Analysys Mason conducted a survey in early 2024, in collaboration with the 5G Deterministic Networking Alliance (5GDNA), which included responses from 60 Tier-1 (accounting for 70% of respondents), Tier-2 and Tier-3 CSPs worldwide. Based on this survey and other Analysys Mason research, we found that private cloud deployments can satisfy more technological and commercial criteria than public cloud deployments, and private cloud models can also provide CSPs with the optimal TCO in the medium to long term.

Figure 1.1: Key considerations for CSPs when choosing between private or public cloud for a cloud-native mobile core deployment model

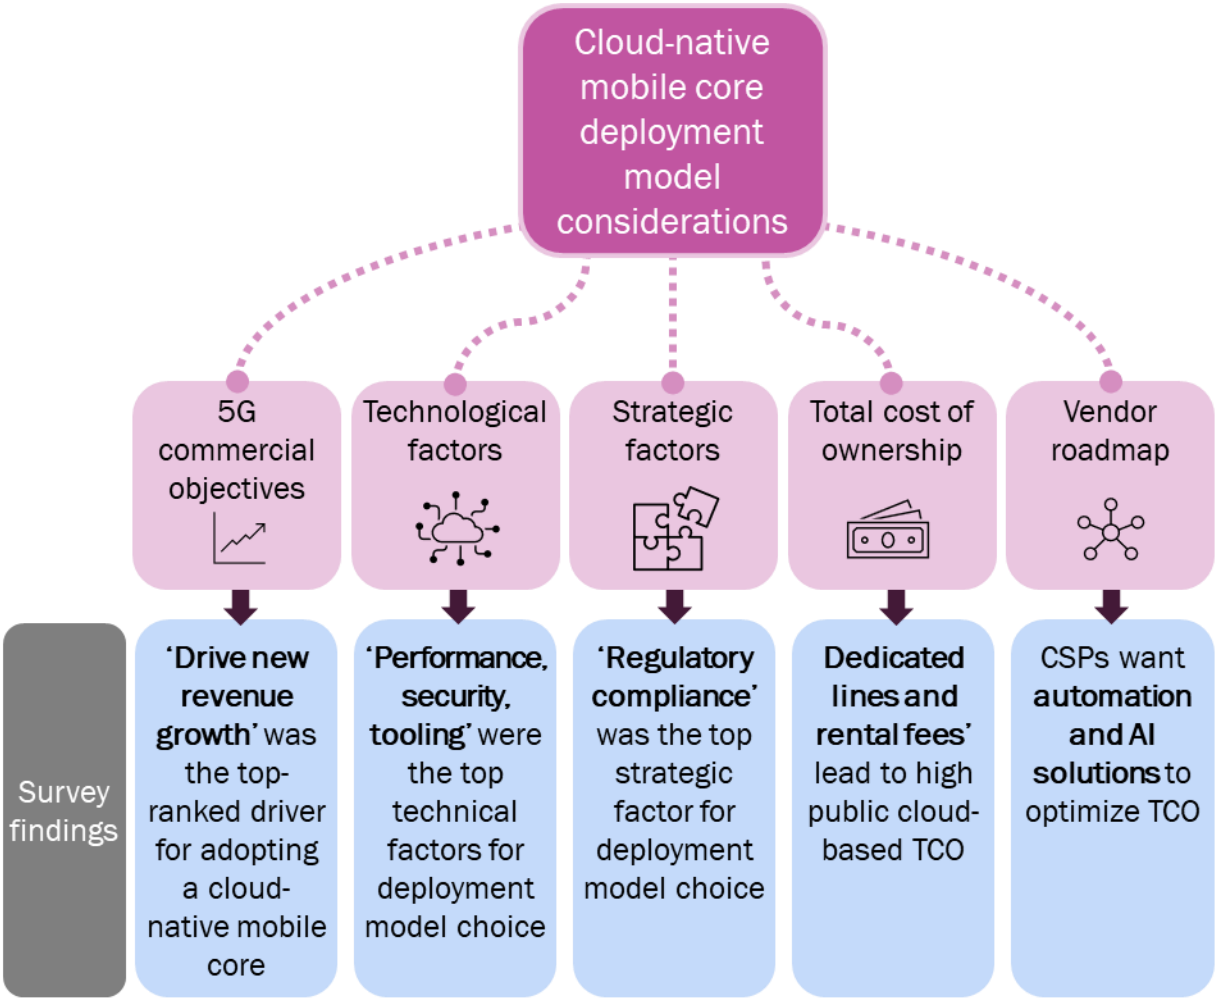

Source: Analysys Mason

Our main survey findings are as follows.

- **83% of Tier-1 CSPs expect to deploy 5G SA within the next 2 years.** CSPs are accelerating their adoption of 5G SA, driven by a need to generate new revenue growth, improve customer service and increase operational efficiencies/cost savings. These important goals indicate a pressing need to quickly identify the ideal deployment strategy.

- **CSPs consider containerisation and cloud-native technologies to be the central foundations of 5G SA core.** CSPs' transition to a cloud-native core starts with containerisation, which is necessary for modular, flexible network infrastructure. CSPs aim to move away from physical network functions (PNFs) to loud-native network functions (CNFs). We expect that CSPs' CNFs will account for 60% of their network estate 3 three years, on average.

- **Private clouds remain the preferred solution for 98% of CSPs' mobile core deployments**. CSPs stated that they believe that private cloud offers better security, superior network performance and compliance with data sovereignty requirements compared with public cloud. While CSPs appreciate the on-demand scalability aspects of the public cloud, many CSPs that are trialling public clouds have failed to achieve desired network performance levels.

- **While public cloud deployments offer superior short-term benefits, private cloud TCO can be lower after 2 to 3 years of operations.** The main factors that contribute to high public cloud opex in the medium to long term include ongoing rental fees when deploying a hybrid model, and dedicated line charges and traffic costs when deploying a centralised public cloud model.

- **Private cloud TCO can be further optimised through automation and AI.** Automation reduces operational costs and enhances network performance, while AI applications offer predictive maintenance and improved network management. By investing in these technologies, CSPs can optimise their private cloud infrastructure, ensuring that it meets the commercial objectives of 5G and lays the foundation for future advancements (including 6G), therefore maximising capex.

2. CSPs are committed to adopting cloud-native mobile cores to fulfil their commercial ambitions for 5G

2.1 5G standalone demands a shift from network function virtualisation to a cloud-native architecture

Network function virtualisation (NFV) was introduced 10 years ago and was CSPs' first use of cloud-based infrastructure strategies. This involved abstracting network functions from dedicated hardware to virtualised environments, which marked the first step towards a more-flexible and efficient network architecture. Despite their significant investments in NFV, CSPs faced, and continue to face, many challenges in realising the promised opex and capex benefits of the transition from PNFs to virtualised network functions (VNFs).

The advent of 5G necessitates a more-radical transformation to achieve opex and capex savings. Designed from the ground up as a cloud-based network, 5G introduces a horizontal cloud platform architecture that leverages microservices and containerisation. This new design paradigm enhances network agility, enables cloud-native automation and positions CSPs to optimise their TCO while unlocking new revenue opportunities.

The mobile core and, in particular, 5G SA, represents an ideal opportunity for CSPs to introduce cloud-native technologies. CSPs can use 5G SA technologies and practices as the catalyst to accelerate their cloud-native infrastructure and network function adoption. Many CSPs plan to embrace cloud-native technologies across their portfolio of mobile cores, including legacy (4G/LTE), 5G NSA, private 5G core and IMSs. The mobile core is poised to command a substantial portion of CSPs ' investments in network cloud infrastructure in the short term and will play a critical role in shaping the future of telecommunications infrastructure as CSPs begin to virtualise the RAN.

2.2 CSPs are set to accelerate their adoption of cloud-native network functions, driven by 5G SA roll-outs

Analysys Mason research[1] shows that as of March 2024, 64% of Tier-1 CSPs have commercially launched cloud-native network 5G SA cores. While most of these deployments are fairly small scale and can even be classified as live commercial network trials in a single region, they represent an ambition and commitment by

[1] For more information, see Analysys Mason's *Cloud Transformation Benchmark 2023: results and key findings*.

CSPs towards embracing CNFs. Within 2 years, 83% of Tier-1 CSPs are expected to have deployed 5G SA CNFs, which further indicates their strong commitment to change. Our survey results indicate that China and developed Asia–Pacific (DVAP) are the regions that are furthest ahead in terms of their implementation timelines, while Central and Eastern Europe (CEE) and emerging Asia–Pacific (EMAP) will be the slowest regions to adopt CNFs.

Within 2 years, 71% of Tier-1 CSPs worldwide will have also deployed private 5G cores as CNFs, highlighting an ambition to better serve enterprise customers to meet the latter's increasingly stringent low-latency demands and to support a variety of enterprise applications that can lead to new revenue-generating opportunities. While legacy and 5G NSA CNF implementation lags behind 5G SA and private 5G, many CSPs are converging multiple core domains on a single cloud-native platform, taking cloud-native implementation a step further by having a single horizontal platform to further increase efficiencies. However, other CSPs are taking their time to define and implement their 5G SA core strategies before extending this to other mobile core domains.

> By 2028, we intend to be an entirely cloud-native organisation. We are still in the early stages of roll-out, but our main emphasis is on 5G SA. This core will be a converged core for all mobile core domains.
>
> *Networks VP, Tier-1 CSP in Western Europe*

Figure 2.1: CSPs' timeline for adopting cloud-native network functions by mobile core domain, worldwide, 2024

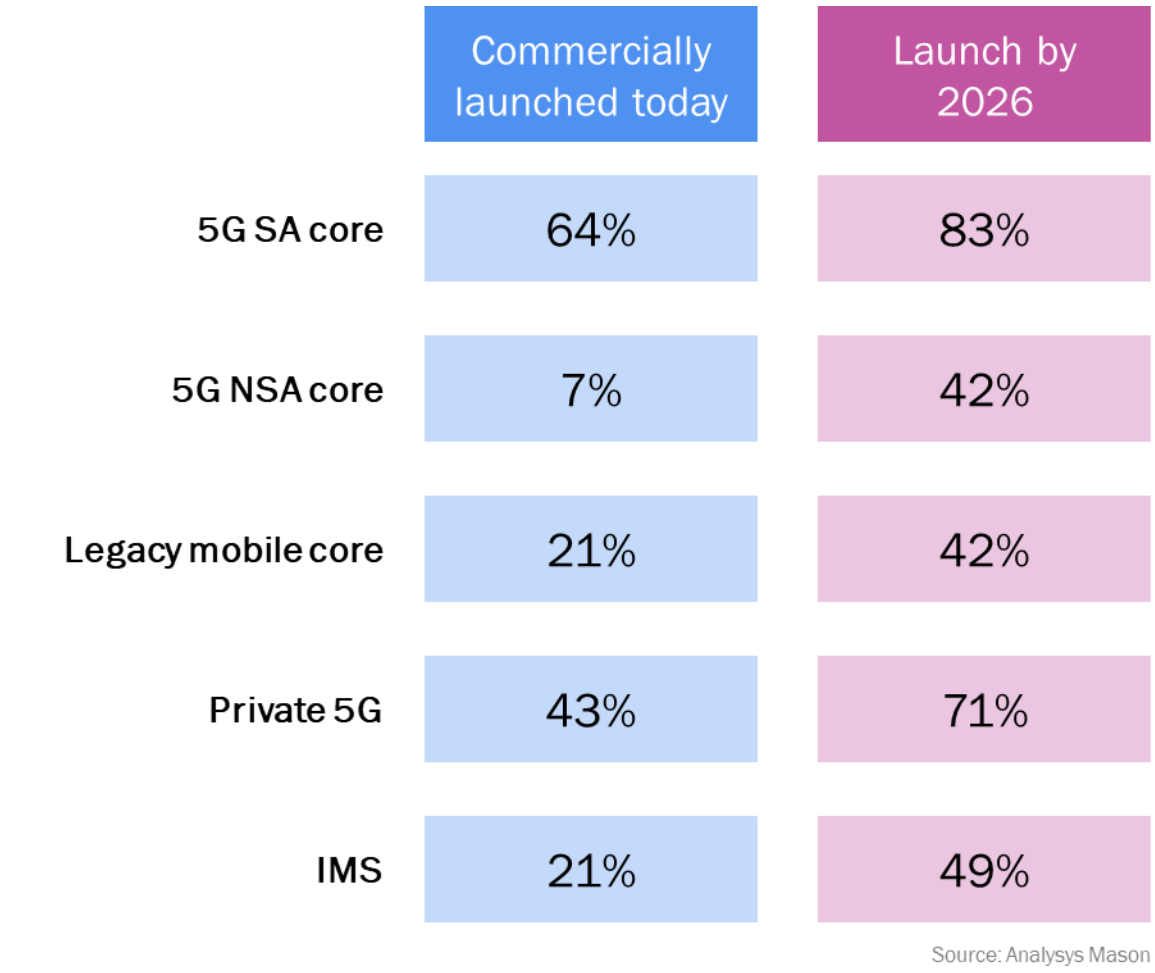

2.3 CSPs can only achieve their commercial objectives through the adoption of cloud-native cores; this is a progressive strategy that starts with containerisation

CSPs are mainly focused on commercial objectives as the drivers for introducing cloud-native mobile cores, and they understand the clear link between the commercial objectives and the adoption of cloud-native cores. Analysys Mason's survey of 60 Tier-1, Tier-2 and Tier-3 CSPs world conducted in early 2024 showed that CSPs' top two motivations for introducing a cloud-native core are to drive revenue growth and to improve customer experience (see Figure 2.2). These sentiments are shared by CSPs regardless of size and across all regions. These main motivations were followed by CSPs' desire to increase efficiency/opex savings, which indicates that they want to drive new revenue growth that is backed by more-efficient operations.

Figure 2.2: CSPs' top motivations for implementing a cloud-native/containerised mobile core, worldwide, 2024

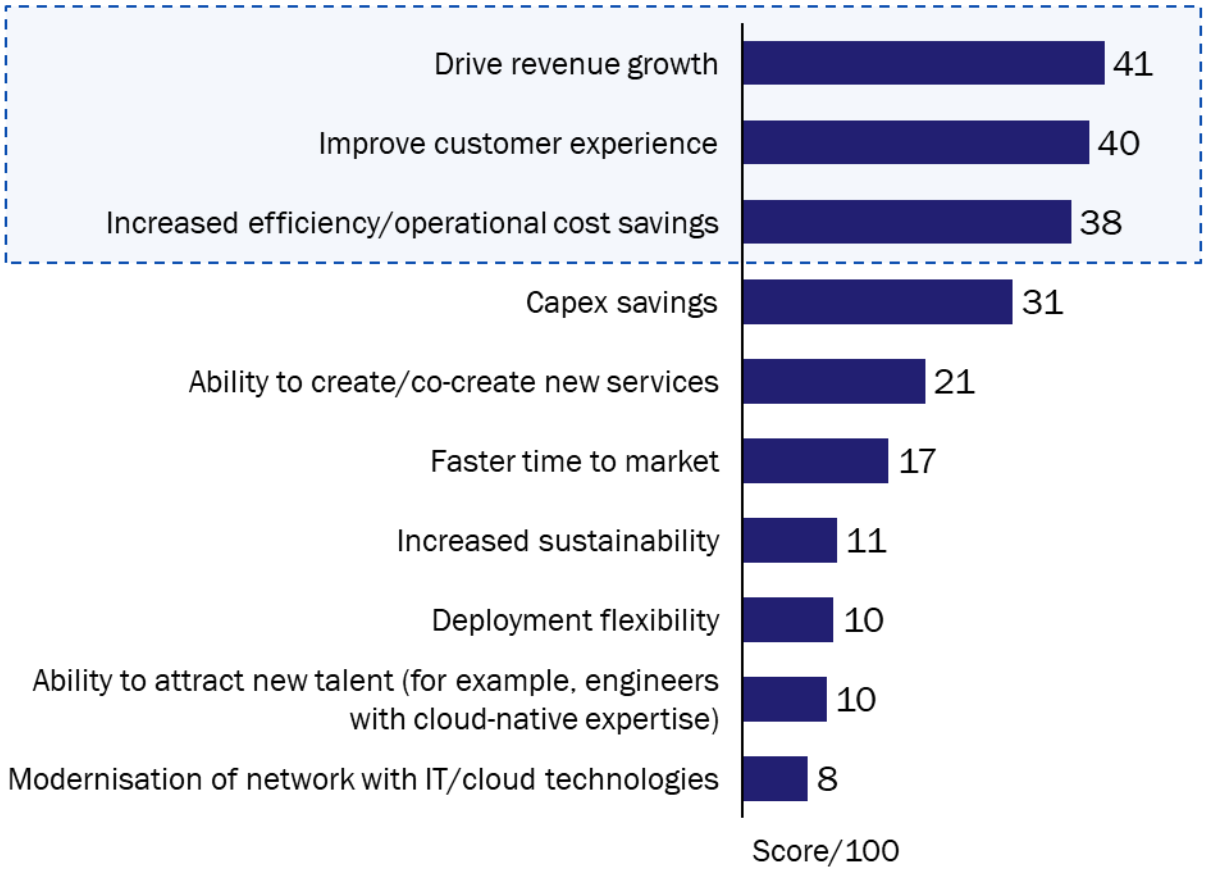

Explanation of scores: CSPs were asked to rank their top three motivations. Top-ranked motivations received the highest weighting.

Source: Analysys Mason

Cloud-native mobile cores can support CSPs' ambitions in the following ways.

- **Driving new revenue growth**. Cloud-native architectures allow CSPs to deploy digital services significantly faster than traditional networks because their modular design supports rapid deployment, iterations and updates. This agility allows CSPs to respond quickly to market demands and customer needs, seizing revenue-generating opportunities and strategically positioning them for emerging enterprise and consumer use cases.

- **Improving customer experience.** By adopting cloud-native cores, CSPs can enhance network reliability and performance. The inherent resilience of cloud-native architectures, and in particular, Kubernetes, with capabilities such as automated healing and load balancing, ensures that network services remain robust and continuous. The ability to provide more digital services more quickly will also contribute to improving customer experience.

- **Increasing efficiency/operational cost savings.** Cloud-native automation is a new paradigm for CSPs and one that will enable significant operational and cost savings. This technology minimises manual network management, setting the stage for autonomous networks that operate based on business intentions. With networks becoming self-sustaining, the need for human intervention is drastically reduced. This shift allows CSPs to reallocate their workforce and resources from routine maintenance to pursuing more-lucrative

revenue-generating opportunities. The following cloud-native properties are essential for achieving this goal: containerisation, a microservices architecture, observability and a horizontal platform.

CSPs are clear on their desires for a cloud-native core but to achieve the benefits of this approach, CSPs should have the right implementation strategy, with containerisation being a key technology element that must be adopted as a first priority. The above-mentioned commercial and TCO benefits rely on containerisation because it provides the modular building blocks of a cloud-native core that is essential for the network's flexibility, and container orchestration platforms such as Kubernetes are the foundation of highly automated networks.

With this in mind, CSPs have very strong ambitions to accelerate their adoption of containerisation. Today, PNFs still constitute a significant portion of CSPs' mobile core environments (see Figure 2.3). However, in 3 years' time, PNFs will become the minority network function environment because CNFs are set to represent 60% of a CSP's network estate, on average. This strong shift towards CNFs highlights CSPs' pressing need to keep pace with the broader cloud-native transformation trends across all other industries, and to be able to meet the quickly evolving enterprise and consumer demand for more-advanced network use cases.

Figure 2.3: Average level of mobile core proportions by environment, worldwide, 2024[2]

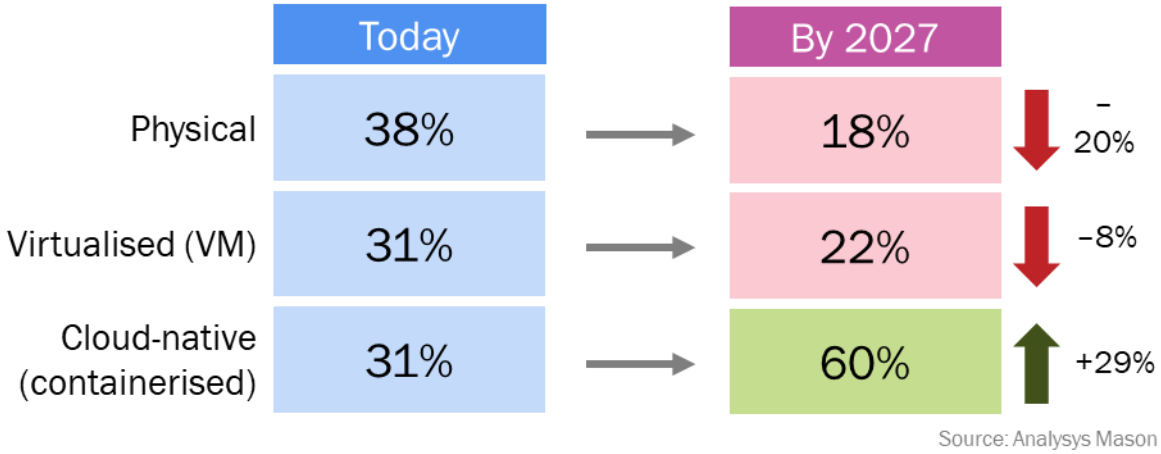

In summary, CSPs have strong cloud-native ambitions and a robust understanding of the benefits of transitioning to a cloud-native and containerised environment. However, as CSPs deploy cloud-native mobile cores, they face a critical decision regarding the choice of infrastructure for their cloud platform: whether to opt for private telco cloud or public cloud environments.

[2] Question: What proportion of your mobile core is cloud-native today, and in 3 years' time, versus based on virtualised or physical infrastructure?; *n* = 60.

3. CSPs believe that private cloud is the ideal option for unlocking the full benefits of cloud-native mobile cores

3.1 Private clouds continue to be the preferred solution for mobile cores for 98% of CSPs

Analysys Mason's research indicates a strong preference among CSPs for private cloud infrastructures to run their mobile cores. Today, 98% of CSPs select private clouds for their mobile core deployments, with just 2% currently adopting a public cloud model.

Public cloud as a potential mobile core deployment model is a recent development that has gained the attention of some CSPs. However, actual commercial deployments remain very limited, with most being small-scale, live network trials. Public cloud adoption will primarily be driven by CSPs in North America and Western Europe as part of a mix of deployment model strategies that will be highly dependent on the market and on the use case. Public cloud providers (PCPs) in these two regions also have the greatest availability of in-country data centres. Even in these markets, the public cloud remains a niche deployment model. After conducting extensive evaluations of their deployment options, many CSPs that have launched a 5G SA core have chosen private clouds.

In the following section, we discuss the perceived advantages of private cloud over public cloud, and which the drivers of CSPs' preference of private clouds for their mobile core deployments.

Figure 3.1: Proportion of private cloud and public cloud mobile core deployments today, worldwide, 2024

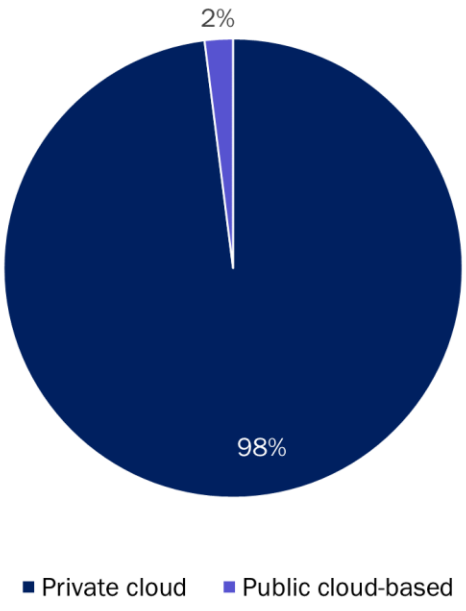

Source: Analysys Mason

3.2 Private clouds offer major advantages over public clouds in three key technical areas

A private cloud deployment model for deploying a mobile core can have three key benefits over the public cloud: better security, superior network performance and compliance with data sovereignty requirements. A potential fourth benefit, greater TCO savings, is discussed in section 4.

- **Security** is the top reason for not adopting the public cloud, according to surveyed CSPs. CSPs are concerned about data security breaches and the need to meet stringent requirements mandated by their governments (see Figure 3.2). Security is a particular concern in the cloud-native era because employing containerisation, microservices architectures and Kubernetes greatly increases the attack surface of a network and increases the number of potential vulnerabilities.

  CSPs believe that private clouds offer superior security for mobile core environments compared to public clouds, primarily due to their dedicated resources and customisable nature. Unlike public clouds, where infrastructure is shared among multiple tenants, private clouds provide CSPs with exclusive control over their environments. This exclusivity significantly reduces the risk of unauthorised access and data breaches, which are more prevalent in shared settings. The fact that each public cloud provider has suffered at least one major security breach heightened concerns among many CSPs.

  In private clouds, CSPs can also tailor security protocols to fit their specific needs, implementing advanced encryption, strict access controls and bespoke firewall protections. CSPs want the same fine-grained level of customisation to be available in the public cloud before they place confidence in this environment.

- **Network preference.** Figure 3.2 also shows that CSPs have major network performance concerns when deploying a mobile core on the public cloud. These include fears that a public cloud environment may compromise the latency, carrier-grade reliability, consistency and predictability of a mobile core network. A significant number of surveyed CSPs that have already trialled public cloud-based mobile cores said that they were unable to achieve the same guaranteed performance and reliability as they experienced when running the mobile core in a private cloud. In addition, public cloud service-level agreement (SLA) guarantees are not yet on par with private cloud's 'five nines' (99.999%) reliability guarantees (5.26 minutes of downtime per year), further contributing to CSPs' hesitation in fully adopting public cloud solutions for their critical network functions.

  Data residency and sovereignty issues are a growing concern for CSPs due to geopolitical threats and increased privacy legislation. A private cloud is formed from dedicated infrastructure and as such, it enables data to be stored and managed within a specified jurisdiction, allowing organisations to easily comply with local data protection laws and avoid the legal and ethical complexities associated with cross-border data transfers. PCPs do not have in-country data centres in every geography and therefore may be unable to comply with these laws, which makes public cloud solutions unsuitable for adoption by CSPs that must comply with these regulations. The management and control of encryption keys pose a further sovereignty issue for CSPs. PCPs typically manage root encryption keys on behalf of all their customers, a situation which many CSPs find unacceptable.

Figure 3.2: CSPs' top reasons for not adopting or considering adopting a public cloud-based deployment model, worldwide, 2024

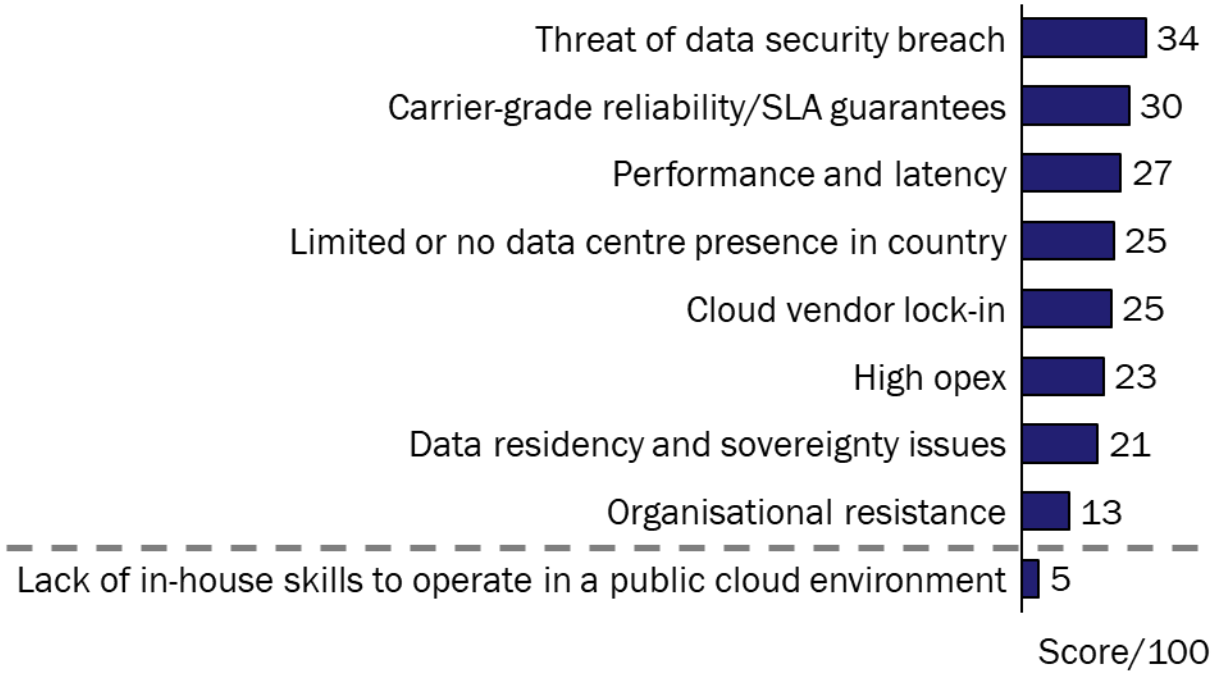

Explanation of scores: CSPs were asked to rank their top three reasons. The top-ranked reason received the highest weighting.

Source: Analysys Mason

3.3 Public cloud will remain a niche strategy for mobile core deployments

At first glance, the public cloud presents some enticing advantages, such as on-demand scalability and access to PCPs' well-respected IT and artificial intelligence/machine learning (AI/ML) tools.

On-demand scalability allows CSPs to dynamically adjust their resources to meet fluctuating network demands. In a public cloud environment, CSPs can quickly scale up their infrastructure during peak usage times without the need for significant upfront investments in physical hardware. This capability is particularly beneficial for handling unexpected surges in network traffic, which can vary due to special events or the launch of new services.

PCP cloud-native tools and services are mature and comprehensive. These providers are well-versed in operating cloud-native environments, and they have a portfolio of AI tooling that is useful for the migration of IT workloads for modernisation and cost-efficiency purposes but represent a discontinuity for network operations.

Despite these benefits, public cloud mobile core deployments will remain a niche strategy due to the technical challenges outlined in the previous section, and strategic challenges such as a lack of in-country data centres outside of North America and Western Europe. As a result, the public cloud may not become the primary model for many Tier-1 CSPs' large scale, mission-critical networks in the near future.

However, there are some specific use cases that are more suitable and compelling to CSPs, for example, improving network resiliency and service continuity by having a backup or hot standby in case the primary

network goes down or experiences technical issues. In addition, the public cloud model offers a platform within which to quickly and cost-effectively experiment with new services, especially for new B2B/enterprise services. Smaller CSPs and MVNOs may also find the public cloud model suitable for their operations, given its infrastructure flexibility and lower upfront costs.

4. CSPs should consider all costs when comparing the TCO between private and public clouds

4.1 The private cloud model's ease of operations and lower reliance on specific expertise are its main advantages over public clouds for achieving TCO savings

Many CSPs are attracted by the short-term potential TCO savings of a public cloud mobile core deployment. However, these CSPs are generally overly optimistic about these benefits and have no direct experience of using it or have limited familiarity with the detail of public cloud deployments. In our survey, those CSPs currently trialling and evaluating the public cloud expressed concerns about lacking the necessary in-house expertise to fully achieve the benefits of this model. These concerns include a lack of operational expertise in using and maintaining PCP tools and protocols; the challenge of keeping pace with PCPs' fast release cadences for software updates that have to be aligned with network functions and the complexities of managing the co-existence of legacy infrastructure and a public cloud environment. This indicates that CSPs need to make significant upfront investments in the requisite skillsets and automation capabilities to efficiently operate their networks in the public cloud and/or engage with external partners to fill these skillsets and operational gaps. All of these factors can make a considerable impact on the TCO and CSPs should take these factors into account when evaluating the adoption of public clouds.

Over the last decade, CSPs have made significant investments in technology and skillsets for their private clouds. Many are reluctant to write off these investments. Instead, they prefer to evolve their existing private cloud infrastructure using cloud-native technologies. This approach allows them to capitalise on operational familiarity and existing skillsets, facilitating a smoother, less-radical transition to cloud-native networks. Another significant TCO advantage of private cloud over public cloud is the integration with legacy infrastructure, which prevents or minimises disruptions in network operations across various geographies, markets and network domains during this transition. Furthermore, existing SLA guarantees, established and trusted vendor relationships that have been built up over years in existing private telco clouds provide CSPs with greater comfort in terms of stability, reliability and manageability.

4.2 Network and data transfer costs may result in a significantly higher TCO for public clouds compared to private telco cloud

When CSPs consider deploying their mobile cores on either a private or public cloud, they must weigh up a variety of other key capex and opex factors, in addition to assessing which deployment model best suits their technical requirements and internal capabilities, as outlined in the previous section. To help CSPs in making an informed decision, the critical TCO elements are detailed below.

Figure 4.1: Capex and opex criteria for mobile core deployments

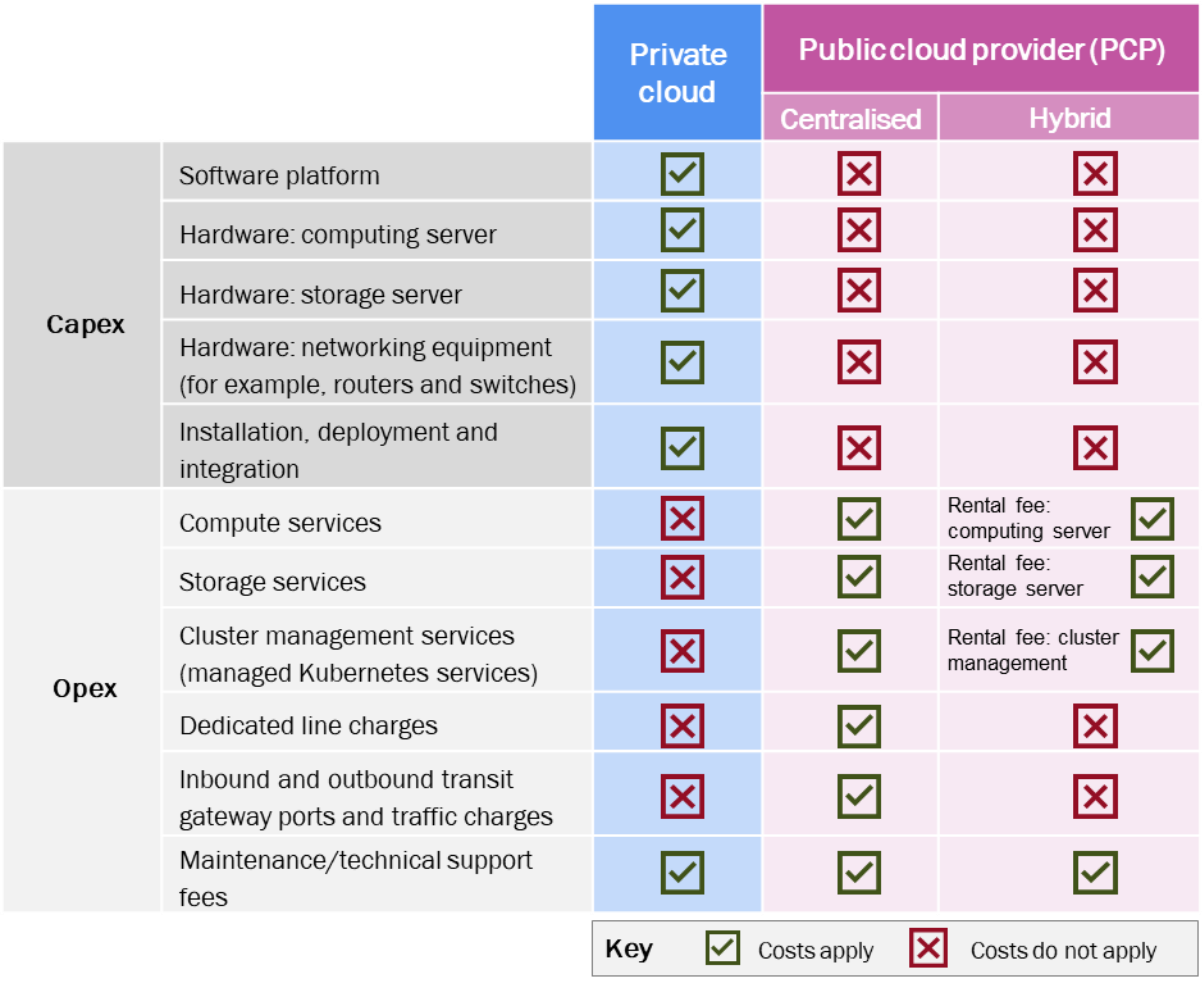

Source: Analysys Mason

Figure 4.1 compares the cost items involved in deploying a mobile core using three different models: private cloud, centralised public cloud, and hybrid public cloud. In the centralised public cloud model, the entire mobile core is hosted on the public cloud. This requires the CSP to pay for compute, storage and cluster management services, as well as dedicated line charges for connecting its network to the public cloud, and traffic charges for data ingress and egress (data moving in and out of the cloud).

In contrast, the hybrid public cloud model involves the PCP supplying the compute, storage and networking hardware, which is installed at a CSP data centre and paid for on an opex basis. This setup eliminates dedicated line charges and traffic costs because only the management plane connects to the public cloud.

Private cloud deployments initially incur a higher TCO than both types of public cloud model due to the upfront capex requirements. These expenses include hardware for computing, storage and networking, as well as cloud software platform costs and expenses related to installation, deployment and integration. However, in our interviews, several CSPs expressed concerns that over the medium to long term, the TCO for public cloud models would likely surpass that of private clouds after 2 to 3 years of operation due to the significant opex involved. The main drivers and cost factors for this potentially higher TCO vary depending on the public cloud model.

For the **centralised public cloud model**, compute, storage, networking and managed Kubernetes are consumed as-a-service under an opex model. Although these costs are initially lower than the upfront capex required for building a private cloud environment, the overall TCO for this model may become more expensive over the long term for large-scale networks such as the 5G core. While PCPs typically offer discounts of up to 40% on compute, storage and egress for reserved instances over several years, the medium- to long-term costs still remain significant.

Dedicated lines and traffic costs also significantly contribute to the public cloud opex and should be included in the TCO evaluation. Dedicated lines, or direct cloud connections, are required to provide stable, high-speed connectivity to the public cloud from CSPs' data centres. This connectivity is necessary because, despite the core running in the public cloud, it still needs to integrate with the CSP's broader network, including legacy (4G) infrastructure and the RAN. These dedicated lines are essential for handling the substantial data volumes inherent in mobile network operations. Establishing these connections involves significant initial expenses, and the ongoing costs are determined by bandwidth capacity, which can go up over time.

Traffic costs are another major source of high opex associated with this model. Most PCPs charge for data egress – the data sent from the cloud to the internet or to other data centres. CSPs need to move large amounts of data constantly, leading to far higher egress costs than most enterprises. As a result, identifying data flow, and working out how the data can be managed and utilised within a single cloud region (many PCPs do not charge ingress fees) to minimise these costs, will require significant planning and restructuring so that CSPs can ensure these costs do not become excessive. Even when considering the on-demand scalability aspect of public clouds, traffic charges remain substantial due to the constant high flow of data, both network data and network observability and monitoring data – and if CSPs are aiming for more automation, then they will need as close to real-time data as possible, which will incur significant costs.

In the **hybrid public cloud deployment model**, the main costs are the rental fees for compute, storage and networking hardware. This model involves ongoing rental payments for the hardware provided by the PCP and installed at the CSP's data centre in a similar way to private cloud model. It saves on dedicated line and traffic fees and provides a more-predictable and manageable cost model compared to the variable costs associated with the centralised model. However, these costs may still result in a higher TCO than the private cloud in the long run, due to the opex-based procurement model, which has accumulating ongoing costs, similar to that of centralised public clouds explained above, but without the benefits of cloud economics. This model lacks the elastic capacity and on-demand scalability benefits of centralised public clouds, and the rental equipment typically comes in a limited set of specifications and modularity, which may lead to overprovisioning of hardware resources.

To summarise, while private clouds can come with significant capex, the main infrastructure-related opex criteria that a CSP will need to consider is the ongoing maintenance and technical support fees. In the medium to long term, CSPs should be aware that the TCO for public cloud mobile core deployments is likely to exceed that of private clouds due to the significant dedicated line and traffic charges for a centralised public cloud deployment and the annual hardware rental fees for a hybrid deployment.

4.3 CSPs can improve private telco cloud TCO by introducing cloud-native automation and AI

While the private cloud model is not perfect, it has the potential to provide the optimum TCO for CSPs and to maximise the return on network investments. There are several ways to optimise a private cloud to ensure that it delivers high performance, carrier-grade reliability, increased security and meets the commercial 5G objectives of driving new revenue growth and improving customer experience. On top of acquiring the right cloud-native

technologies, CSPs need to implement true cloud-native automation and lay the groundwork for integrating AI into the network (see details below).

- **Implement cloud-native automation.** Integrate automation solutions that provide full visibility and control over network operations. Automation can significantly reduce opex and enhance operational efficiency by streamlining network management processes. It also has a direct impact on customer service by improving network reliability and performance and allows for greater agility and reduced time to market for new services. The ideal automation method involves using a DevOps and CI/CD based approach.

- **Integrate AI.** Use the same foundations that were built to put cloud-native automation in place and establish the necessary data architectures to facilitate AI applications, including large language models (LLMs). In addition, identify the optimal model training methods and resources required. AI can offer immediate benefits, including more-efficient network operations and maintenance.

CSPs currently face significant challenges in implementing these steps due to the steep learning curve associated with operating cloud-native mobile cores. It is therefore crucial for CSPs to have access to tools that enable seamless and agile network operation while maintaining their legacy infrastructure. Partnering with a provider that can deliver these tools, along with comprehensive support and a clear roadmap for future innovations in automation and AI, is essential. Operations and maintenance (O&M) represent the most-challenging and sensitive aspects for CSPs, playing a critical role in achieving sustainable revenue growth while minimising TCO. Below are details of some of the key areas of improvement for O&M that can be gained by adopting a cloud-native private cloud mobile core deployment.

- **Cross-layer fault locating and demarcation.** To enable efficient fault-finding in the network, it is essential that all parts of the network are instrumented, and that data is cleaned and channelled into a data lake. This means an observability platform is required to provide the full view of the network via a single dashboard. Seamless interoperability across cloud stack layers, enabled by automation, is crucial for effective cross-correlation throughout the network. This capability allows for rapid fault identification, with full-stack vendors well-positioned to provide these capabilities due to their native integrations. CSPs can also benefit from AI-driven diagnostics to locate faults more efficiently using chat-based AI applications, which some advanced vendors are able to provide.

- **Predictive maintenance.** CSPs can develop and implement preventative measures for their mobile core operations based on historical data and trends. When training LLMs, the CSP must use its own data. While automated decision-making may not be ready for live network use, predictive maintenance can be a useful tool to help network operations staff to manage their workloads and decide on the best course of action.

Looking ahead to the future, CSPs have the opportunity to lay the foundation for 6G with future-proofed architecture and investments that can be reutilised and maximised throughout the 6G era, which will be a fully cloud-native era marked by AI. With a private cloud and cloud-native mobile core, CSPs will have the ability to design their own roadmaps because they will own their infrastructure. In the AI era, compute capacity will be at a premium. CSPs can mitigate the risk of overprovisioning their network infrastructure by using any spare capacity to train AI models or by renting it out to others for training purposes. This shows that TCO can be lower for private cloud deployments in the medium and long term when comparing direct mobile core operations with public cloud deployments, and it shows that private cloud can also become a valuable asset for future network monetisation.

5. Conclusions

CSPs are actively pursuing the adoption of cloud-native mobile cores to fully realise the potential of 5G technology. The shift from NFV to a cloud-native architecture is crucial for achieving operational and capital expense savings. With 5G SA, CSPs have a significant opportunity to introduce cloud-native technologies, enhancing network agility and enabling cloud-native automation. Many Tier-1 CSPs have already launched cloud-native 5G SA cores, demonstrating their commitment to this transformation, with the rest of the market likely to follow suit in the coming years. As CSPs converge multiple core domains on a single cloud-native platform, supported by key architectural properties such as containerisation, microservices architecture, observability and a horizontal platform, they are positioning themselves to better serve consumer and enterprise customers and support a variety of new applications to drive new revenue growth.

Private clouds remain the preferred deployment model for most CSPs, while the public cloud is expected to remain a niche deployment model suited for certain use cases and markets. Most CSPs in our survey believe that private clouds can offer superior security, network performance and compliance with data sovereignty requirements, all of which are critical for mobile core environments. While public clouds provide benefits such as on-demand scalability and advanced tooling, concerns about performance reliability and data residency have limited their adoption. CSPs must carefully evaluate their options, considering all technological and strategic elements, while factoring in organisational limitations to determine the best approach for their cloud-native mobile cores.

Although public clouds may appear to offer short-term savings, the long-term costs associated with dedicated lines, traffic charges and ongoing rental fees can lead to higher TCO than private cloud in the span of just a couple of years. CSPs can optimise their private cloud investments through cloud-native automation and AI integration, further reducing TCO and enhancing network performance. By introducing these elements, CSPs can ensure that their cloud-native infrastructures are future-proof, supporting their commercial objectives and paving the way for 6G advancements with minimal new capex.

6. About the author

Ameer Gaili (Analyst) is a member of the Cloud research practice, and mainly contributes to the Cloud Infrastructure Strategies and Edge and Media Platforms research programmes. Prior to joining Analysys Mason, Ameer was a strategy consultant at a boutique management consultancy. Ameer holds an MEng in chemical engineering from the University of Manchester.