Assessing the effects of 5G SA adoption on TCO and network operation

Lei Shi, James Kirby, Michelle Lam, Gorkem Yigit, Bengt Nordstrom

20 April 2026

Contents

Copyright © 2026. The information contained herein is the property of Analysys Mason and is provided on condition that it will not be reproduced, copied, lent or disclosed, directly or indirectly, nor used for any purpose other than that for which it was specifically furnished. This report has been prepared by Analysys Mason and sponsored by Ericsson, subject to Analysys Mason's editorial judgement and discretion. The analyses contained within this report are the sole responsibility of Analysys Mason and do not necessarily reflect the views of Ericsson, or of other contributors to the research

Analysys Mason AB
Sankt Eriksgatan 63B
Floor 9
112 34 Stockholm
Sweden
Tel: +46 8 587 120 00
stockholm_office@analysysmason.com
www.analysysmason.com
Registered in Sweden: No. 556285-9487
Subsidiary of Analysys Mason Limited, 5th Floor,
22 Upper Ground, London, SE1 9PD, UK, No. 5177472

1 Executive summary

5G standalone (SA) can form a meaningful part of the near-term business case for operators, not only through future monetisation, but also through lower network total cost of ownership (TCO) and more scalable operations.

This study compares two trajectories in a simulated developed mobile market: a conservative scenario and an SA-led scale-up scenario. Both scenarios start from the same baseline of an existing cloud-native dual-mode core (DMC); the difference is in how quickly communication service providers (CSPs) choose to scale SA migration and automation. The same study based on non-cloud-native DMC for 4G and SA would lead to a similar directional result, but fewer operational gains in the core domain since SA is better able to exploit the advantages of a cloud-native environment.

- **conservative**: 5G SA remains limited in scale and most traffic continues to rely largely on 4G and 5G non-standalone (NSA)
- **SA-led scale-up**: SA-capable traffic is migrated more actively to the 5G core and automation scales alongside that migration.

This study takes both quantitative and qualitative results into account for both scenarios.

- Quantitative analysis indicates that there are two main areas of benefit
  - SA-led scale-up can reduce cumulative five-year network TCO by around **13%** versus the conservative scenario. The annual saving profile is back-loaded: as SA traffic share and automation maturity increase, the **relative annual TCO difference** between the two scenarios can exceed **20% by year 5**.
    - larger operators tend to show greater upside because they have more traffic in the network and broader service mixes that benefit more from improved capacity planning, traffic profiling and operational efficiency.
  - the radio access network **(RAN) domain** will show the largest gains: improved traffic profiling, better utilisation of radio resources and earlier refarming of selected 4G spectrum bands can reduce avoidable capacity spend by allowing operators to time upgrades more precisely and, in some areas, defer densification.
    - over time, a second group of benefits becomes more material in the **core domain**, where stronger observability, more structured lifecycle control and higher automation coverage reduce manual effort across provisioning, change management and assurance.
- Qualitative operator interviews and survey responses broadly support the quantitative findings. Operators consistently point to device readiness, automation maturity and operating complexity as the main influential factors. They also expect value to emerge in sequence: first through RAN planning and spectrum use, and later through operational industrialisation as SA traffic grows and automation reaches meaningful scale.

The strategic implication for CSPs is clear. SA should be treated as a transformation programme with two parallel aims: reducing avoidable network cost and preparing the network for what comes next.

In the near term, the priorities are to ensure that SA-capable traffic is progressively carried on the 5G core, that planning assumptions evolve with that migration and that operational improvement scales beyond pilots. Over time, the same foundations that unlock cost savings also create readiness for dynamic slicing at scale, more demanding service assurance, AI-assisted operations and future network programmability.

2 Introduction

2.1 Cost savings can form a meaningful part of the 5G SA business case

5G standalone (SA) is widely recognised as the technical foundation for network slicing and more advanced automation. However, initial operator interviews indicate that SA scale-up is progressing more slowly than industry expectations, for three main reasons.[1]

- **Device readiness and customer experience risk:** SA-capable devices are growing in number, but broad activation depends on device enablement and certification processes from original equipment manufacturers (OEMs). Operators are cautious about moving too quickly while 4G and 5G non-standalone (NSA) remain the dominant experience for most customers.
- **Limited demand for dynamic slicing:** Many operators see value in differentiated services but expect near-term needs to be met by a small number of predefined (static) slices rather than per-customer bespoke slices, due to cost-benefit considerations and governance complexity.
- **Upfront cost and operational complexity:** The introduction of cloud-native core and SA requires upfront investment and often increases near-term operational burden (skills, tooling, dual run),[2] even if operators expect efficiencies to emerge as automation is industrialised.

Against this backdrop, many discussions on SA business cases focus on the **timing and shape of future revenue upside** from differentiated slicing-based services, network exposure/API opportunities, positioning-related services and SA-enabled device/service categories, such as RedCap. These monetisation opportunities are emerging, but are likely to materialise unevenly across operators and over a longer time horizon, influenced by factors that are outside the CSP's control, such as device ecosystem timing, customer adoption and regulatory considerations. In parallel, the **cost-side impact** of SA and automation can be assessed more directly and is more closely linked to operator execution. Decisions on the operating model, tooling, orchestration and automation coverage determine how quickly CSPs can industrialise SA operations and create repeatable efficiencies.

This report therefore focuses on the total cost of ownership **(TCO) and operational efficiency case** for accelerating SA and automation. The intent is to complement that longer-term discussion with a decision-grade view on the cost transformation potential and the conditions required to realise it. Importantly, the same capabilities that unlock cost benefits also prepare the network for future requirements such as scaled dynamic slicing, advanced SLA-based enterprise services and AI-assisted operations.

---

[1] Low-band and mid-band SA readiness remain important preconditions for good SA quality, even if they were not the most frequently cited near-term constraint in the interviewed communication service provider (CSP) sample

[2] Running both 4G and 5G cores in parallel

2.2 TCO saving is measured in networks over five years

The study examines the impact of SA adoption on the mobile network cost base over a five-year period, covering the following scope dimensions:

- **Network scope:** mobile network only, covering radio access network (RAN), core, transport/backhaul and operations support systems (OSS)/operations where these costs are network-facing, for example, assurance tooling and labour related to network operations centres (NOCs).
- **Cost scope:** network capex and opex relevant to SA and slicing adoption, including transitional costs where applicable, for example, integration and dual-run implications.
- **Time horizon:** five-year comparison (2026–2030) to capture both near-term overheads and later benefits as automation matures and SA-device adoption and demand for slices/differentiated services grow.

The study does not quantify revenue upside from new SA-enabled services. It also excludes broader commercial or corporate costs, handset economics and one-off spectrum acquisition costs unless explicitly stated.

2.3 Measurements are presented across two scenarios and three operator profiles

The quantitative analysis compares two execution scenarios in a simulated developed mobile market, using three operator profiles (Tiers 1–3) to reflect differences in scale, service mix and resulting operational complexity (not current SA maturity) across consumer mobile broadband, fixed wireless access (FWA), Internet of Things (IoT), wholesale, public protection and disaster relief (PPDR) and private networks. **The tiering is intended to show how SA affects operators with different traffic intensity and cost structure.**

Table 2.3: (a) Simulated market and operator profile assumptions used in the scenario comparison

Market share	Tier 1	Tier 2	Tier 3
Consumer market share	40%	25%	10%
FWA market share	40%	25%	10%
IoT market share	50%	20%	-%
Wholesale market share	40%	-%	-%
PPDR market share	50%	-%	-%
Private network market share	40%	20%	-%

Both scenarios start from the same baseline: a cloud-native dual-mode core (DMC) is already in place, but most 5G traffic is still carried in NSA mode, meaning the 5G core is not yet the primary traffic-carrying core. In both cases, SA-capable device penetration grows at the same pace over time. The difference is how actively the operator converts that readiness into actual SA traffic, and how far it scales automation and slicing-related operations as complexity increases.

- **Conservative**: SA remains limited in scale, with mainstream traffic continuing to rely largely on 4G and 5G NSA, differentiated services are delivered mainly through existing **quality of service (QoS)** policy, **access point name (APN)**-based separation and **virtual private network (VPN)**-based mechanisms, rather than through scaled end-to-end SA slicing and lifecycle management.

- **SA-led scale-up**: SA-capable traffic is migrated more actively to the 5G core, enabling earlier refarming, automation scales with that migration, and slicing evolves from limited template-based use towards broader operational scale, with more structured lifecycle management as service complexity grows.

Table 2.3: (b) Scenario comparison summary

Item	Conservative	SA-led scale-up
Starting point	Same in both: DMC primarily in NSA mode	Same in both: DMC primarily in NSA mode
SA traffic migration	Limited	Active but capped by device penetration (from 10% to 80% in 5 years)
Slicing approach	NSA-era differentiation / limited static slicing	Static SA slicing initially, later dynamic slicing
Refarming	2100MHz only	2100MHz initially, 1800MHz and 800MHz over time
Automation	Limited	Scaled with SA growth

The next section presents the quantitative results, starting with the overall five-year TCO impact of the two scenarios and then examining the main drivers of value across RAN, core and operations. Detailed methodology and modelling assumptions are provided in the appendix.

3 Quantitative results: TCO and operational efficiency

This section presents the model outputs for the cost-side effects of 5G SA migration and automation over a five-year period (2026–2030). Results are avoidable cost over five years, reported as differences between the SA-led scale-up and conservative scenarios.

The model starts from a mobile-only network cost baseline for each operator profile and decomposes costs into RAN, transport/backhaul, core and network-facing OSS and operations. In the baseline, annual spending continues to rise or remain elevated over time because operators must sustain capacity and modernisation.

The value from faster SA migration and automation comes primarily from more targeted investment timing and lower cumulative upgrade need over the five-year period, rather than from an immediate reduction in network change activity, The operational efficiency improvement comes over time as both SA traffic and automation coverage scale. The TCO saving should not be interpreted from a ‘cost down’ perspective in absolute terms.

3.1 TCO savings[3] ramp with SA traffic share and automation maturity

The year-by-year results follow a consistent pattern: the majority of cost savings are only realised when (i) a meaningful share of traffic is carried on the SA core and (ii) automation and industrialised operations are applied at scale. As a result, the annual saving profile is typically back-loaded, the gap between the SA-led scale-up and conservative trajectories starts to widen in later years.

Figure 3.1 shows the annual relative TCO difference between the two scenarios. Two effects explain the shape of the curve. First, realised savings accelerate as SA traffic share increases (capped by SA-capable device adoption and enablement). Second, later-year savings are amplified as CSPs move into more advanced automation across provisioning, change and assurance, reducing the effort required per service/slice and per network incident.

Scale and service mix also matter. Tier-1 and Tier-2 operators typically show greater upside because they have higher traffic volumes and a larger operational workload base to optimise. In the model, Tier 1 and Tier-2 operators benefit more from SA-led scale-up because faster SA migration reduces dependence on LTE anchoring, which in turn enables earlier refarming and therefore greater capacity relief. Operators with a more diverse service mix tend to gain additional upside from improved traffic profiling and slice-based resource control, which increases the opportunity to improve utilisation as slicing scales (both in number of slices and the size of the network they cover).

---

[3] Savings or TCO reduction refers to lower cumulative five-year network cost relative to a comparator scenario, not to an absolute decline in annual network spending over time.

Figure 3.1: Operator network TCO relative savings per year, SA-led scale-up vs conservative, 2026–2030. [Source: Analysys Mason, 2026]

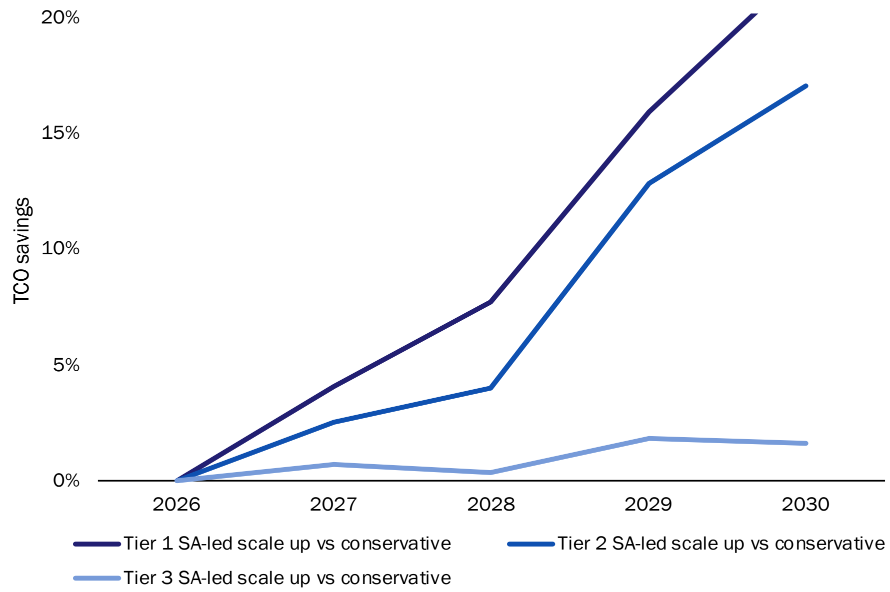

3.2 The main five-year TCO upside comes from earlier RAN savings and later core efficiencies

Figure 3.2 compares the five-year cumulative network TCO in the SA-led scale-up and conservative scenarios for a Tier-1 operator and indicates that the SA-led scale-up path can reduce TCO by around 13% versus the conservative scenario. The waterfall shows that this upside is shared across both capex and opex, with the largest effects typically coming from RAN and transport, and later operational efficiencies increasingly driven by core.

Figure 3.2: Illustrative five-year cumulative network TCO, SA-led scale-up vs conservative (split by domain and capex/opex) [Source: Analysys Mason, 2026]

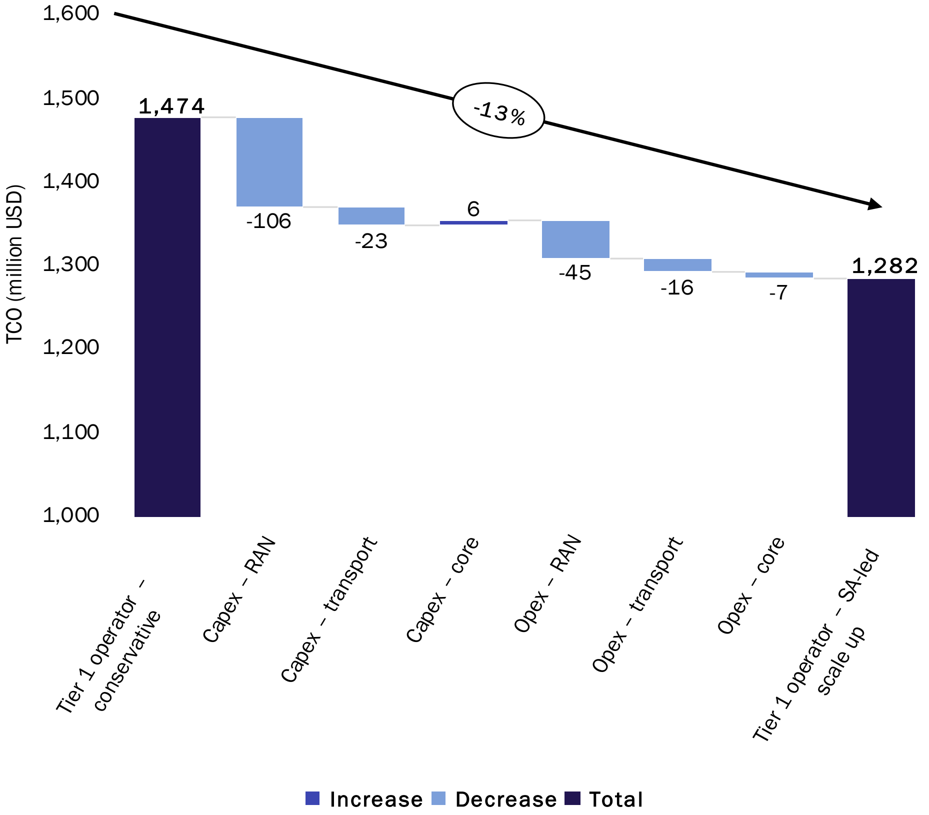

To make the result easier to interpret, Table 3.2 summarises the main efficiency levers that sit underneath the capex and opex split.[4]

Table 3.2: Main five-year TCO efficiency levers in the SA-led scale-up scenario

Cost type	Efficiency lever	How value is created
Capex	More targeted capacity upgrades	Improved traffic profiling and service differentiation enables more precise capacity planning, reducing overbuild and allowing some upgrades to be delayed.
Capex	Earlier 4G spectrum refarming[5]	As reliance on LTE anchoring reduces, portions of 1800MHz and 800MHz can be refarmed from 4G to 5G earlier, improving effective capacity and delaying selected upgrade needs.
Capex	Lower need for densification in selected areas	Better utilisation of existing spectrum and radio resources reduces the need for additional sites in some high-load areas.
Opex	Automated provisioning and change workflows	Automation reduces manual configuration effort and shortens service activation and change cycles.
Opex	Automated assurance and fault handling	Better observability, fault correlation and remediation reduces incident-handling effort and improves operational response times.
Opex	Remote diagnostics and fewer field interventions	Improved visibility and automation reduces truck rolls, speeds up readiness checks and lowers field maintenance effort.

The RAN domain is where the model shows the earliest and most material cumulative impact. This reflects the fact that better traffic profiling, combined with earlier refarming, under SA allows more targeted and improved timing for capacity investment, reducing the need for some later expansion.

This should not be interpreted as an absence of SA-related RAN investment. The SA-led scale-up scenario still requires targeted RAN-enablement costs, including SA-RAN software activation on the existing 5G footprint, selective radio modernisation in refarmed bands, and ongoing mid-band deployment where additional capacity is still required. The cumulative RAN TCO is nevertheless lower than in the conservative scenario because these enablement costs are more than offset over time by earlier refarming, improved traffic profiling and more targeted capacity expansion for the SA-led scale up scenario.

Core economics are structurally different from RAN. The SA-led scale-up scenario requires earlier investment for scaling up SA core operation, as well as additional investment in automation tooling and orchestration to manage increasing operational scale as slicing expands. This is consistent with the experience and expectation from CSPs interviewed during the study that SA scale-up can increase cost and complexity before it reduces it, particularly in cloud-native environments with a higher change cadence. The result is typically higher initial investment and operational overhead.

---

[4] Note that in our cost model, capex includes upfront cost for equipment, initial subscription fees for software, hardware licenses, implementation services; opex includes operational labour cost, recurring software subscription fees, hardware and software support fees.

However, that enablement is what allows later operational efficiencies to emerge through automation and more standardised lifecycle management.

This comparison should not be interpreted as “cloud-native versus non-cloud-native”, since both scenarios start from the same cloud-native DMC baseline. Some operational gains from automation and modernised lifecycle management can already be realised on a cloud-native DMC baseline, even before SA traffic and slicing scale materially. The incremental difference is that faster SA scale-up allows operators to make fuller use of cloud-native operating practices, including stronger observability, more structured lifecycle control and broader automation as SA traffic and service complexity grow. As SA traffic share increases and automation scales, these efficiencies increasingly offset the earlier cost uplift, leading to a more favourable year-5 core TCO than in the conservative scenario.

3.3 Key sensitivities: SA traffic share and automation scale have the biggest influence on outcomes

The five-year cumulative TCO result is most sensitive to two groups of drivers. The first is SA device adoption and related spectrum use, where SA traffic share and the pace of refarming determine how much capacity relief can be realised. The second is core automation scale, where the operating model and the scale of slice lifecycle activity determine how much operational efficiency can be captured. These sensitivities align closely with both the survey and interview evidence: survey responses point to RAN planning and operational efficiency as the main sources of value, while interviewed operators consistently identified device readiness, automation maturity and operating complexity as the main factors that affect their pace and cost for migrating subscribers to SA.

3.3.1 RAN-side drivers: SA traffic share and refarming

The most important constraint on early-year benefits is the share of traffic that can be carried on SA. This matters most in the RAN domain because the RAN cost base is large and several of the main capex levers depend on SA traffic becoming meaningful enough to improve capacity planning, resource utilisation and refarming readiness.

A simple sensitivity illustrates the effect. If SA traffic share reaches only 50% by 2030 rather than the base case of 80%, the modelled five-year cumulative TCO improvement falls from around 13% to around 8%. The effect is seen primarily in RAN, because lower SA traffic share reduces the benefit of SA-RAN spectral efficiency, traffic profiling gains, and delays the point at which operators can confidently optimise capacity planning and refarm spectrum from 4G.

Refarming assumptions are also material because they directly affect the timing and magnitude of RAN capacity investment. In the base case, earlier SA migration reduces dependence on LTE anchoring and supports refarming of the 1800MHz and 800MHz bands. If this refarming is removed in sensitivity testing, the five-year cumulative TCO improvement falls from around 13% to around 11%.

The model does not assume full NSA/EPC switch-off within the horizon; additional simplification and opex benefits would likely arise beyond the model period once legacy layers can be retired more fully.

3.3.2 Core operations drivers: cloud-native operating model and slice scale

Core-domain economics are shaped less by refarming and more by the balance between upfront SA-enablement investment and later operational efficiency as complexity increases. Two assumptions are particularly important: the operating environment for the DMC, and the scale at which slice lifecycle activities need to be managed.

First, although the main comparison is not between cloud-native and non-cloud-native core environments, it is relevant as a sensitivity analysis to test whether the same SA-led trajectory can realise the same operational benefits in a non-cloud-native variant. A sensitivity test using a non-cloud-native DMC variant shows that, while investment for enabling SA is broadly similar, core/OSS operational savings are roughly 40–50% lower because the network cannot realise the same gains in observability, change execution discipline and automated lifecycle management.

Second, slice scale materially affects the operational case for SA-native orchestration. As the slice catalogue grows, manual lifecycle handling becomes increasingly costly unless orchestration and automation scale with it. In the SA-led scenario, a dynamic slicing orchestrator and cloud-native automation practices[6] contain this cost by reducing the manual effort associated with activation, modification and assurance. In the conservative scenario, where differentiation continues to rely on NSA-era mechanisms, lifecycle handling becomes much more labour-intensive as complexity rises.

Sensitivity testing shows that if slice growth is modest (for example from around 5 to around 10 slice instances over five years), core operational savings are materially lower, because the scaling pressure that makes orchestration valuable is weaker. If slice scale is higher (for example from around 5 to around 100[7] slice instances), core operational savings are materially higher because orchestration and automation avoid a large amount of incremental manual lifecycle effort. In other words, the slice-scale sensitivity is driven less by large changes in SA-led operating cost and more by how quickly operational effort rises in the conservative case as service differentiation becomes harder to manage without SA-native lifecycle control. Table 3.5 summarises the directional impact of the main sensitivities on the five-year cumulative TCO result.

---

[6] In the SA-led scale-up scenario, two mechanisms contain this cost pressure. First, investment in a dynamic slicing orchestrator reduces the manual effort required to manage lifecycle changes across a large number of slices (e.g. automating activation, modification, and deactivation workflows that would otherwise require manual intervention per slice). Second, deploying cloud-native automation principles (e.g. Kubernetes-based workload management allows network function instances to scale dynamically, GitOps-based configuration management replaces ad-hoc change execution with version-controlled, auditable and repeatable deployment pipelines and CI/CD practices that allow software updates and policy changes to be tested and rolled out continuously rather than through disruptive maintenance windows.

[7] . This is not a prediction that operators will run 100 slices on NSA; it is a modelling device to quantify the structural difference between an SA-native approach to lifecycle management and NSA-era ‘emulation’ when service differentiation scales.

Table 3.5: Key sensitivity analysis

Sensitivity	Primary domain affected	Base → Variant	Directional impact
Lower SA traffic share by 2030	RAN (main), core (secondary)	80% → 50%	~13% → ~8% 5-year TCO improvement
Refarming of 1800/800	RAN	Both bands refarmed over time in base case → both bands stay as 4G bands	~13% → ~11% 5-year TCO improvement
Non-cloud-native DMC	Core/OSS ops	Cloud-native DMC in base cases → non-cloud-native DMC	core/OSS opex savings ~40–50% lower
Slice scale	Core/OSS ops	base: ~5 to ~50 slice instances low: ~5 to ~10 high: ~5 to ~100	Low: core operational cost savings ~60% lower High: core operational cost savings ~50% higher

4 Qualitative validation and operator insights

This section complements the quantitative modelling by summarising (i) what CSPs reported in interviews and (ii) what respondents indicated in the operator survey. The objective is to test whether the modelling logic aligns with operator experience, and to highlight what barriers operators have experienced or foresee that could limit the pace at which 5G SA and slicing can be scaled.

The qualitative analysis consists of anonymised operator interviews, complemented by a survey focused on expected efficiency impacts from NSA to SA and from slicing-enabled automation. The interviews are weighted towards larger, more mature CSPs. As a result, the insights are most representative of operators already investing in SA readiness and automation rather than those at the earliest stages of 5G deployment.

Survey results should be interpreted as directional expectations. Several responses implicitly combine SA-driven effects with broader improvements from modernisation and cloudification, so the survey is best used to understand where operators expect the largest benefits and how those benefits change over time, rather than as an estimate of savings attributable to SA alone.

4.1 Survey findings: what operators expect to improve, and when

The survey results reinforce two themes that are consistent with the quantitative model: benefits increase over time as SA traffic share grows, and the operational and planning impacts are perceived as at least as important as pure technology features.

4.1.1 Operators expect operational efficiency to improve as SA and automation scale

Respondents expect the operational effort required per unit of traffic to improve materially over time as SA becomes more established and automation scales. The survey indicates a steady reduction in full-time employee (FTE) hours per GB when comparing SA to NSA over the period, with the largest improvements appearing in later years. This supports the modelling outcome where operational efficiencies are back-loaded rather than achieving immediate savings in the first one to two years of SA roll-out. When asked to rank the most impactful operational efficiency drivers by 2030, respondents place the highest emphasis on:

- remote-first operations and reduced field intervention
- automated provisioning
- improved network planning and fault correlation.

Figure 4.1: (a) FTE hours per GB of traffic NSA vs SA, share compared to year 1 NSA [Source: Analysys Mason, 2026]

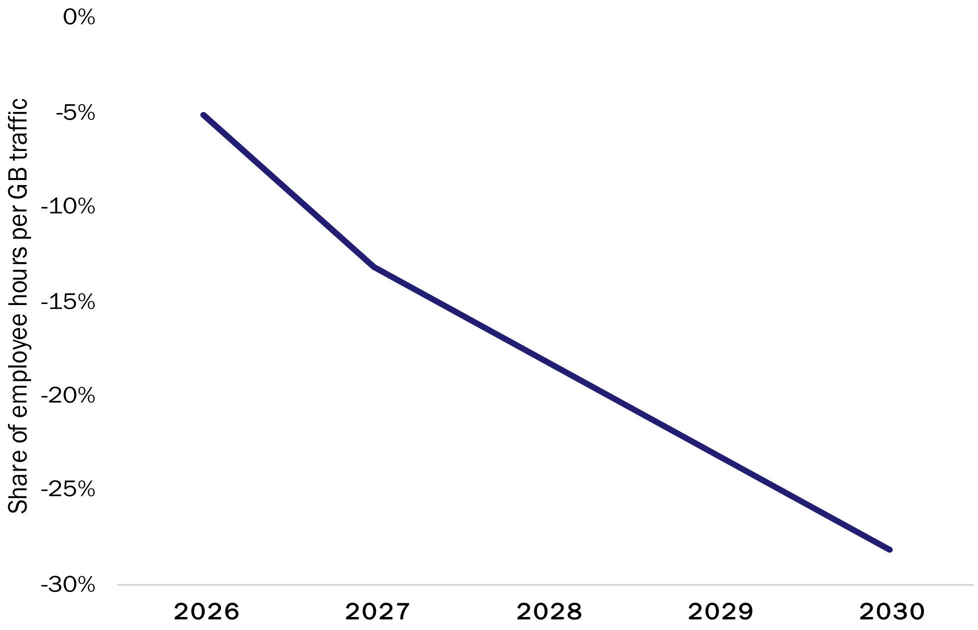

Figure 4.1: (b) What operational efficiencies do you expect to achieve for your 5G SA network compared to your 5G NSA network by 2030?, 0-5 (0 no impact – 5 strongest impact on total annual time or cost required to operate your 5G SA network) [Source: Analysys Mason, 2026][8]

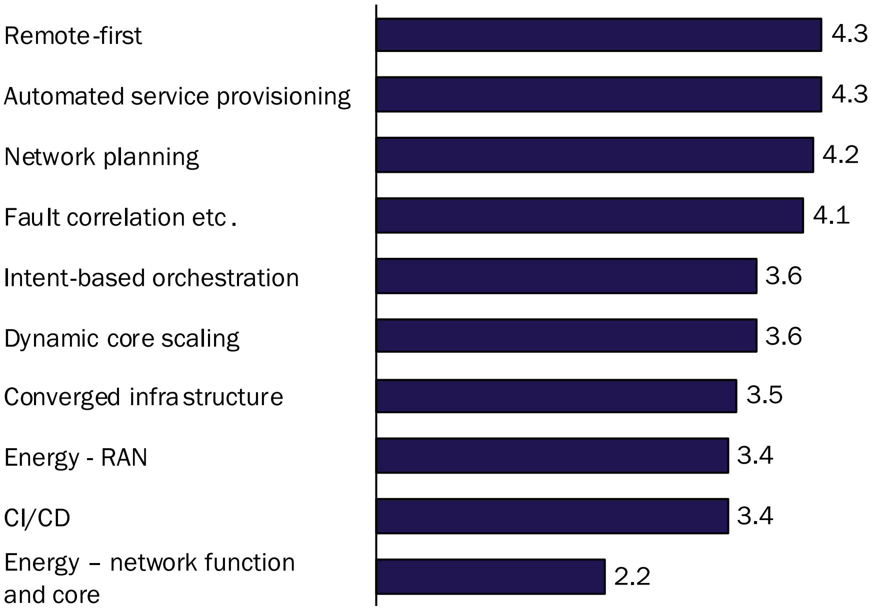

---

[8] Note: Chart labels are shortened for readability. Full survey wording is provided in Annex B.

4.1.2 Operators expect capex efficiency to come mainly from RAN planning and capacity discipline

Respondents also expect capex efficiency to improve as operators gain better control over capacity planning and upgrade timing in an SA-led environment. When asked to rate the sources of capex efficiency by 2030, the highest-scoring responses relate to RAN and planning:

- RAN-related efficiency and improved capacity planning are rated most highly, followed by scaling and spectrum-related levers
- transport slicing and shared-core effects are seen as relevant but secondary.

This aligns with the model's emphasis on earlier, more material, capex leverage in RAN (and to a lesser extent transport), with outcomes dependent on whether operators translate improved traffic profiling and policy control into more targeted capacity upgrade decisions.

Figure 4.1: (c) How does the capex efficiency (spend per GB of new capacity) of your 5G SA network compare to that of your 5G NSA network? How do you expect this to change? [Source: Analysys Mason, 2026]

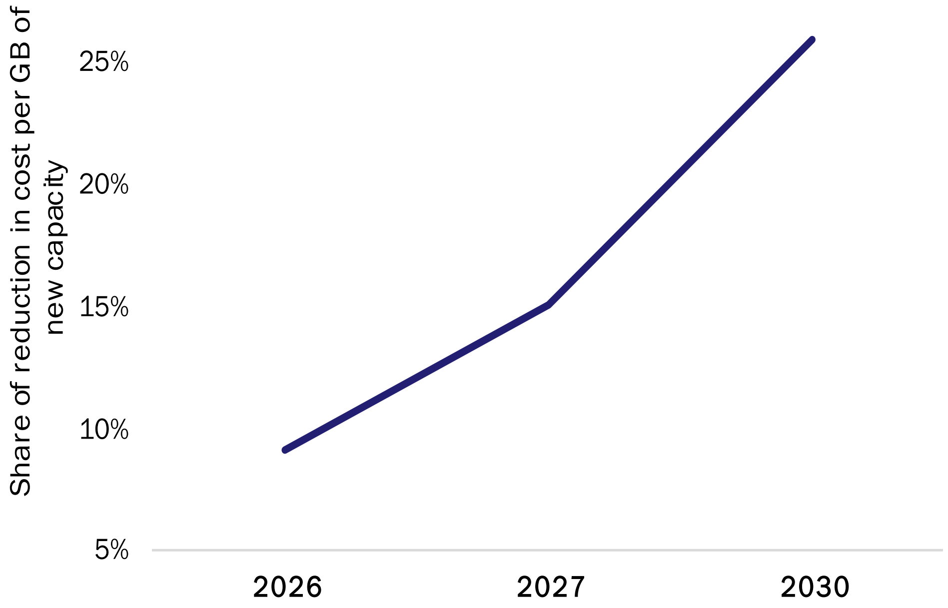

Figure 4.1: (d) What are the main capex efficiencies you are achieving or expect to achieve by 2030 from your 5G SA network compared to 5G-NSA?, 0-5 (0 no impact – 5 strongest impact) [Source: Analysys Mason, 2026][9]

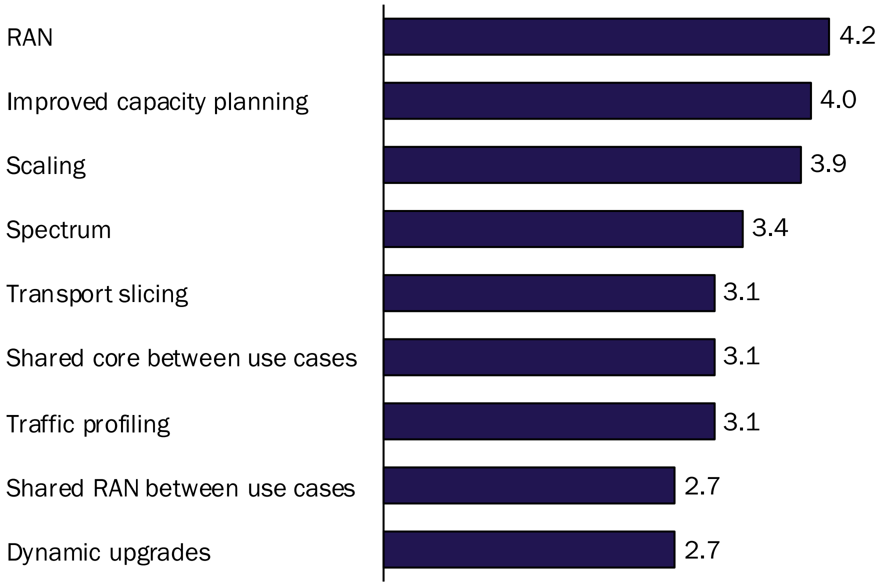

4.1.3 Operators see slicing value first in operations and selected high-load use cases

The slicing-focused survey results point to two practical conclusions. First, respondents expect slicing to scale over time in terms of the number of slice instances and operational constructs, implying that the need for lifecycle discipline and repeatable operations will increase over time. Second, the highest-rated benefits of slicing and slice-based automation are concentrated in operational and planning outcomes:

- network planning and fault correlation are rated highest
- cross-domain QoS and SLA support also rate strongly
- operational mechanics such as automated slice management and dynamic scaling are rated as important, but not the only value driver.

Notably, items such as reusable templates and test slices are rated lower, which may indicate that operators either take these as prerequisites rather than benefits, or have not yet scaled these practices sufficiently to view them as a major efficiency driver.

Figure 4.1: (e) What operational benefits do you expect to achieve from introducing network slicing and slice-based automations for your 5G SA network and how would you rate their expected impact on operational

---

[9] Note: Chart labels are shortened for readability. Full survey wording is provided in Annex B.

efficiency of your 5G SA network by 2030, 0-5 (0 no impact – 5 strongest impact) [Source: Analysys Mason, 2026][10]

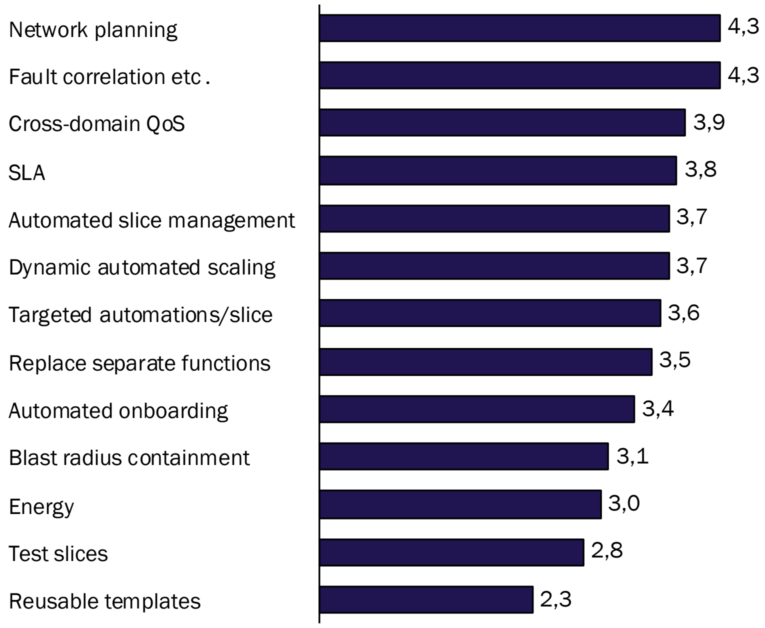

4.2 Interview findings: what operators report in practice

Interviews provide richer context on why SA scale-up is progressing unevenly, where operators see tangible benefits already, and why many benefits remain back-loaded.

4.2.1 SA adoption is driven by capacity, slicing and long-term operational benefits

Operators described SA as strategically necessary even when the short-term business case is challenging. The most common drivers fall into three categories.

- **Capacity, throughput and spectrum efficiency.** Multiple operators view SA as essential to manage future traffic growth as incremental spectrum additions become harder to secure and more expensive to deploy. Several interviewees highlighted that SA reduces reliance on dual connectivity and the ‘anchor tax’ associated with NSA operations. One operator reported markedly higher throughput on SA than NSA in its own environment and linked that improvement to removing LTE anchoring and improved scheduling behaviour in an SA-led configuration.

- **Enabling slicing and differentiated service models.** Operators consistently described SA as the prerequisite for scalable slicing, particularly for dynamic slicing and more demanding enterprise SLAs. Several noted that slicing under NSA conditions is operationally complex due to EPC reliance and dual control planes, and that SA is required for a simpler and more scalable model. Early slicing-related use cases mentioned in interviews include priority services

---

[10] Note: Chart labels are shortened for readability. Full survey wording is provided in Annex B.

for public safety, performance uplift for FWA and selected enterprise cases where deterministic service behaviour is valued.

- **Long-term simplification.** Several interviewees emphasised that the full operational simplification comes only once legacy layers are retired and SA becomes the dominant network mode. Until then, SA can feel like another layer rather than simplification, even if it is strategically the correct direction.

4.2.2 Device readiness, demand and complexity continue to slow SA scale-up

This section outlines the main barriers and constraints operators encounter when scaling up SA deployments. These insights allowed for more accurate calibration of the cost and effort assumptions in the SA-led scale-up scenario, ensuring they realistically account for the challenges operators encounter when transitioning from legacy systems to SA-driven networks.

- **Device readiness and enablement timelines.** Device readiness emerged as the single most consistently cited constraint, covering both the installed base and operator/OEM enablement process. Multiple operators noted that enabling SA at scale is not only a matter of device capability, but also of operator-led enablement and certification processes. Some cited long certification cycles for new network features and device variants.

- **Limited near-term demand for advanced slicing features.** Several interviewees reported that while differentiated services are conceptually attractive, near-term demand is not yet strong enough to justify broad commercial deployment of highly granular dynamic slicing. Many expect the initial slicing approach to remain limited and template-driven, with only a small number of slice instances mapped to broad service categories.

- **Upfront cost, integration load and compute intensity.** Interviewees repeatedly highlighted the cost and complexity of implementing SA and maintaining it in a multi-vendor environment. Some operators described significant ongoing investment in labs and effort to manage integration and regression testing. Others pointed to increased per-subscriber resource needs in SA relative to legacy EPC, driven by additional network functions and policy/charging components in the SA core architecture. This effect is material to early-year economics and supports treating core-domain impacts as a mix of enablement cost and later operational benefit rather than immediate net savings.

- **Operational transition burden.** Operators emphasised that until legacy layers are retired, SA introduces a period of 'layer-on-layer' operational complexity: more systems, more testing and more failure points. Cloud-native operations also introduce a higher operational cadence and a need for new competencies, which can increase workload in the near term, even where the long-term intent is lower run effort.

4.2.3 Early operational gains are real, but driven mainly by automation

Interviews suggest that early improvements are real but often originate from automation and cloud-native practices rather than from 'SA as a feature' in isolation.

- **Day 0/Day1 activation and provisioning** (primarily driven by cloud-native automation): Reported efficiency gains from speeding up readiness checks before networks go live. Several interviewees emphasised that the breakthrough comes when automation is applied end-to-end and scaled through repeatable templates and processes, rather than through isolated scripts or one-off implementations.

- **Day 2 service operations**: Easier to implement change management because of containerisation – Kubernetes enables new instances to be spun up, scaled, configured and deployed with automatic updates and rollbacks for software upgrades. Using AI-driven anomaly detection enables proactive assurance – faster root cause analysis and lower mean-time to repair.

- **RAN and spectrum-related effects.** A few operators observed measurable but limited improvements after refarming and migrating traffic patterns, but most noted that the largest RAN efficiency gains require SA traffic at scale, which is still some years away in many markets.

- **Slicing benefits in constrained scenarios.** Where slices have been applied in high-load or high-priority contexts (for example emergency services or critical enterprise cases), operators report better traffic isolation and more predictable performance under congestion. At the same time, most operators do not yet see slicing translating into broad, monetisable operational savings; the value is seen as real but concentrated in selected use cases.

4.3 Implications for modelling assumptions

Taken together, survey and interview evidence supports four modelling and messaging choices used in this report.

- T**iming** of benefits matters as much as the magnitude. Operational efficiencies are widely expected, but they are not immediate; they scale with SA traffic share and with the maturity of automation and operational industrialisation.

- The **largest early cost-saving levers** are those that affect capacity upgrade timing and planning discipline (often RAN-led), combined with targeted operational improvements that can be automated relatively early. This reinforces the ‘act now to reduce avoidable spend’ storyline, while keeping outcomes grounded in execution reality.

- C**ore economics are nuanced**. Many operators see SA as compute- and integration-intensive in the near term, which supports treating core impacts as a mix of enablement uplift and later operational improvement rather than assuming net savings from day one.

- Finally, the qualitative evidence makes clear that SA is not only a cost story. Operators view **SA as the platform required to scale slicing**, to support more demanding service assurance requirements, and to build towards more autonomous operations using richer telemetry, repeatable lifecycle processes and automation at scale. These future capabilities are a key part of why operators pursue SA, even when early economics can be challenging.

Section 5 builds on these insights to set out a practical roadmap for execution, including priorities for the next 12–24 months and the operational building blocks required to move from pilots to scalable operations.

5 Strategic implications for CSPs

This section translates the quantitative results (Section 3) and the operator evidence (Section 4) into practical implications for CSPs. The core implication is that 5G SA should be treated as a multi-year transformation programme with two parallel aims: reducing avoidable network cost through better upgrade timing and scaled automation, and preparing the network for future capability such as dynamic slicing at scale and AI-assisted operations. Benefits are not instantaneous; they depend on the pace of SA-core traffic migration, the maturity of automation and the share of traffic carried on SA-capable devices.

5.1 What CSPs should do now

Operators typically face two competing pressures: near-term complexity and investment discipline, versus the longer-term competitiveness and capability benefits that SA can unlock. The findings show that cost and capability benefits from SA increase with greater traffic share and scaled operations; delaying progress postpones learning and can lead to unnecessary capacity investment.

The following actions by CSP leaders can be highly effective in the near term:

- **Treat SA as an operating shift, not a feature roll-out.** The decision about whether the organisation intends to make SA the mainstream traffic-carrying mode over time. Without that intent, many benefits remain theoretical because SA stays marginal in the traffic mix.
- **Align expectations on timing and where value shows up first**. Early impact is most often linked to RAN and transport planning decisions (upgrade timing, utilisation discipline, spectrum choices). Operational efficiencies in core and OSS typically become more visible later, once processes are repeatable and automation is scaled. This avoids expecting immediate run-cost reductions from day one.
- **Treat device readiness as a constraint to manage, not a reason to wait.** Early SA benefits are limited not just by how many SA-capable devices are in the market, but by how quickly operators can enable and certify them for broad use on the network. Even with compatible devices available, SA activation at scale can still be slow because testing and OEM certification take time and traffic migration requires careful pacing to maintain customer-experience.
- **Keep slicing practical but avoid choices that make scale harder later**. Early slicing is likely to be limited and template-driven. The priority is not to maximise the number of slices immediately, but to establish lifecycle discipline and an operating approach that can scale when slice complexity increases, rather than relying indefinitely on ‘slice-like’ differentiation mechanisms.

5.2 Execution priorities for the next 12–24 months

The aim over the next 12–24 months is to move from limited SA usage to sustained operational scale, while preventing operational effort from rising in line with complexity. In practice, most operators find that progress tends to concentrate in four areas:

- **Increasing SA traffic on the 5G core in a controlled way.** In practice, this depends not only on coverage, but on device enablement, RAN quality and operational readiness being sufficient to support broader migration without creating avoidable customer experience risk.
- **Making planning and investment decisions reflect the SA trajectory.** As SA traffic grows, the assumptions used for capacity planning and spectrum strategy need to be updated so that

upgrade timing and refarming decisions reflect where traffic is moving, rather than remaining anchored to NSA-era baselines.

- **Scaling operational improvement beyond pilots.** Pilot automations typically involve disjointed workflows in isolated domains. The largest opex reduction requires those workflows to be standardised into repeatable, machine-executed templates and can be scaled and deployed consistently across Day 0 through to Day 2 operations.
- **Maintaining operational stability as cadence increases.** As SA and slicing mature, the volume and frequency of lifecycle activity tend to increase. The practical priority is to ensure day-to-day operations can absorb that cadence without a step-change in incident load or rework.

Technology transformation alone is not sufficient. Operators also need changes in operating models, skills and ways of working if they are to capture the full benefit of higher automation and lifecycle control. Several operators have moved towards more integrated product, engineering and operations models, with clearer end-to-end ownership across development and operations, in order to support higher change cadence and more automated service management.

5.3 Preparing the network for what comes next

Even though this report focuses on avoidable cost, operator evidence is clear that SA is also pursued because it is a prerequisite for future capability. The practical implication is that SA programmes should be designed to create readiness as well as near-term efficiency; otherwise, operators risk ending up with an SA footprint that is technically deployed but not able to support more demanding service and operational requirements.

Looking ahead, four areas are particularly relevant:

- **Dynamic slicing at operational scale**. Moving from static to dynamic slices requires functions that can translate service intent into dynamic resource allocation and support frequent policy-driven changes and automatic network function scaling and slice lifecycle management that is co-ordinated by the slicing orchestrator.
- **Advanced SLA-based services and stronger performance isolation**. Operators highlighted the value of predictable behaviour under congestion for selected use cases (for example, public safety priority services or critical enterprise connectivity). Delivering this requires enforced QoS policies and slice-specific policy profiles, combined with SLA monitoring using more granular telemetry data to manage resource contention across various slice-based services.
- **AI-assisted operations and higher autonomy.** The service-based architecture in 5G SA provides AI systems with more structured and granular observability and performance metrics that are discoverable across network functions – giving AI live topology awareness to execute intelligent decision-making.
- **6G readiness and programmability.** While 6G is not the focus of this report, the direction is clear: more software-driven networks, more frequent change and stronger linkage between service intent and network behaviour. SA migration and operational industrialisation are foundational steps toward that direction.

Overall, the future-readiness case strengthens the cost-side case rather than competing with it. The same elements that unlock avoidable cost over five years (traffic migration to SA, scaled automation and repeatable lifecycle practices) are also the elements that position the network to support more advanced slicing, more demanding service levels and more autonomous operations over time.

Annex A Detailed methodology

A.1 Scenario framework

This study quantifies the cost-side impact of accelerating 5G SA and scaling automation over a five-year horizon. The analysis compares two execution trajectories, conservative versus SA-led scale up, in a simulated developed mobile market with characteristics typical of a mature European country.[11] Three operator profiles are modelled to reflect differences in scale and segment mix (consumer mobile broadband, FWA, IoT, wholesale, PPDR and private networks). 'Tier' is used to indicate scale and service mix only, not current SA maturity.

A.1.1 Shared assumptions across both scenarios

Both trajectories start from the same baseline network state: the CSP has already implemented a cloud-native DMC, but the majority of 5G traffic remains on NSA. In practice this means:

- the 5G radio layer is in place, but 5G traffic is still anchored to LTE and EPC for most users
- ongoing selective 5G mid-band deployment continues where required by capacity demand
- the 5G core is deployed but not yet scaled as the primary traffic-carrying core
- any 'slicing-like' differentiation is primarily implemented through NSA-era mechanisms (for example APN/VPN-based separation and [5G QoS Identifier (5QI)/QoS policy), rather than end-to-end SA slicing with dedicated lifecycle control.

Note that the headline scenario results reflect the incremental impact of migrating SA-capable traffic to the 5G core and industrialising SA operations, not the one-off benefit of cloudifying the core (which is assumed to be already done in both trajectories).

The model distinguishes between baseline benefits associated with a cloud-native DMC operating environment and the incremental benefits that arise as SA traffic, slicing scale and automation maturity increase.

To separate ecosystem timing from CSP execution choices, we apply the same adoption constraints in both trajectories:

- SA-capable device penetration increases over time (from 10% in 2026 to 80% by 2030 for the baseline scenario)
- the service portfolio becomes more diverse over time, reflected in an expanding slice catalogue (from around five slices in 2026 to around fifty[12] by 2030 for the headline scenario).

---

[11] A fictional market is used to ensure comparability and avoid reliance on any single operator's confidential data.

[12] Illustrative assumption used to represent increasing operational scale and lifecycle complexity as slicing adoption broadens; different values were examined in sensitivity analysis.

A.1.2 Key differences between the two scenarios

The scenarios differ in how proactively the operator converts these underlying constraints into realised SA operation at scale.

Conservative scenario

Conservative represents a cautious scenario in which SA exists but remains limited in its role as the primary traffic-carrying mode. The operator continues to rely on NSA for most traffic and uses existing QoS mechanisms to provide differentiated service. Key characteristics are:

- limited migration of SA-capable devices and traffic to the 5G core
- slicing remains mostly 'emulated' via 4G/NSA-era constructs (APN/VPN separation, 5QI/QoS policy), with limited end-to-end slice lifecycle control
- refarming pace remains conservative because NSA anchoring remains dominant for mass traffic
- automation improves incrementally but remains constrained by operating multiple layers at scale.

SA-led scale up scenario

SA-led scale up scenario represents a forward-leaning execution considered credible in practice. It assumes faster migration of SA-capable traffic to the 5G core and earlier industrialisation of SA operations than the conservative scenario, including the tooling required to manage slicing at larger scale. Key characteristics are:

- proactive migration of SA-capable traffic to the 5G core as coverage and device readiness allow
- progression from limited, template-driven slices to operationally scalable slicing, including dedicated slice lifecycle management as complexity increases
- earlier reduction in dependence on LTE anchoring as SA traffic becomes a meaningful share of total traffic which enables faster refarming staged through the period
- increasing automation coverage across provisioning, change and assurance to handle greater operational scale safely.

Maturity **stages used to describe the conservative and SA-led scale up scenario**

For clarity, we describe the SA-led scale up scenario using three maturity stages. These stages are used to structure when specific cost uplifts and savings mechanisms take effect:

- **Stage 0: NSA-led operation on DMC**
  SA is deployed but not scaled as the primary traffic-carrying mode; differentiation is mainly provided via NSA-era mechanisms (APN/VPN and 5QI/QoS). Operational complexity remains 'layer-on-layer'.

- **Stage 1: SA scale-up with limited automation**
  SA-capable traffic is migrated onto the 5G core at increasing scale. SA RAN activations is applied gradually on existing 5G footprint. The operator adopts a limited, template-

driven slice approach for selected service categories. A key implication is that reduced dependence on LTE anchoring makes earlier refarming feasible, which can materially influence RAN capex outcomes.

- **Stage 2: SA at scale with dynamic slicing and extensive automation**
  As slicing complexity increases and slice counts grow; the operator introduces dedicated slice orchestration and more advanced service and assurance automation to manage lifecycle changes at scale. This is where core/OSS operational efficiency benefits become more pronounced in the model, offsetting earlier enablement overheads.

This scenario structure allows results to be interpreted as avoidable cost over five years. A lower cumulative cost in the SA-led scale up scenario reflects cost that can be avoided by earlier execution and industrialisation, rather than an implication that annual network cost must fall over time. The maturity-stage timelines that define conservative and SA-led scale up are illustrated in Figure 6.1 below.

Figure A.1: Scenario maturity stage development, 2026–2030 [Source: Analysys Mason, 2026]

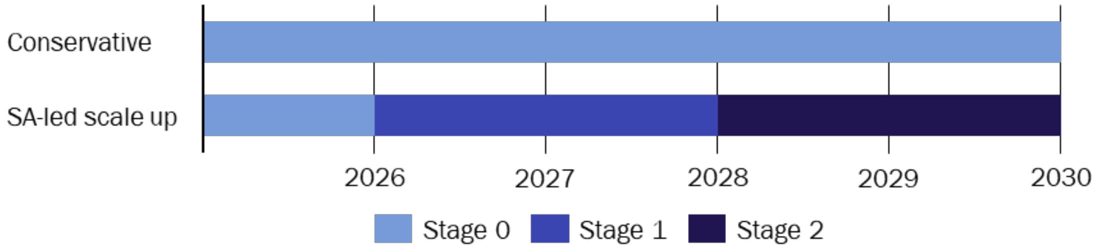

A.2 TCO model overview

The quantitative analysis is based on a bottom-up network TCO model designed to capture where SA migration and slicing-related operational change can shift both capex and opex. The model is scoped to the mobile network cost base and covers the main network domains affected by SA adoption: RAN, core, transport/backhaul and network-facing OSS and operations.

The model is parameterised for each operator profile in the simulated market, reflecting differences in scale and segment mix. This ensures the results capture how SA acceleration affects operators with different traffic loads and operational complexity, not only a single ‘average’ case.

The model separates the economics into two broad types of impact:

- **Capex impacts** are represented primarily through changes in network hardware, software and cloud infrastructure investments. This includes deferral potential where improved utilisation and more targeted upgrades delay capacity expansions, as well as incremental investments required to enable SA readiness and automation at scale.

- **Opex impacts** are represented through changes in operational workload and efficiency in areas such as network and service provisioning, change management, assurance and field operations. These effects are linked explicitly to the pace at which automation coverage and operational industrialisation increase, rather than being assumed to occur immediately upon SA deployment.

Outputs are produced in a form that supports both executive interpretation and technical review. The headline metrics are:

- **Cumulative five-year TCO delta** between scenarios (primary headline metric)
- **Year-5 run-rate delta** (to show the direction as SA traffic share and operational maturity increase)
- Decomposition by **capex versus opex**, and by **domain** (RAN, core, transport, OSS and operations).

A.3 Data sources and validation approach

The model is parameterised using a combination of benchmarks and primary inputs. The cost structures are based off our extensive cost models which we continue to evolve over time, supported by a wide range of vendor and operator feedback. Primary research through interviews is used to validate the practicality of the assumed trajectories and to calibrate the operational logic behind savings mechanisms, including where operators commonly see near-term overheads, where they expect efficiency to materialise and what constraints they encounter when scaling automation and slicing operations.

The modelling approach follows three principles to maintain credibility and comparability. First, the assumed growth in SA-capable devices and traffic is applied consistently across scenarios to avoid attributing ecosystem effects to operator execution. Second, differences between scenarios are defined narrowly and transparently in terms of migration pace and operational industrialisation, so the driver of outcomes is clear. Third, where evidence is limited or operator practice varies substantially, assumptions are treated as ranges and stress-tested through sensitivity analysis rather than presented as a single point estimate.

This methodology is intended to provide a decision-grade view of the cost-side business case for SA acceleration and automation. It does not attempt to quantify revenue uplift from new slicing-enabled services, but it does consider how readiness constraints and operational realities can cap or affect the timing for the realisation of cost benefits.

A.4 Key definitions used in this report

- **5G NSA vs SA:** NSA relies on an LTE anchor and an EPC/dual-core arrangement; SA uses the 5G core and supports a cleaner end-to-end service model for automation and slicing.

- **Static vs dynamic slicing:** static slicing refers to a limited number of predefined slice instances with relatively stable configuration; dynamic slicing refers to more flexible, policy-driven lifecycle management and scaling that can adapt to demand and SLA requirements. We assume NSA supports *slice-like QoS differentiation* (static, coarse) but not end-to-end slice lifecycle/isolation; SA enables true slicing and later dynamic slicing with orchestration.

- **Automation and industrialisation:** automation refers to reducing manual effort in key workflows (provisioning, change, assurance). Industrialisation refers to making these capabilities repeatable and scalable through templates, governance, telemetry consistency and closed-loop operating practices.

Cost reduction vs avoidable cost: where this report uses terms such as 'reduction' or 'savings', these refer to **lower cumulative cost versus a comparator scenario** over five years, not an absolute decline in annual spend over time.

A.5 Key assumptions summary table

Figure A.2: (a) Key assumptions

Key assumption	2026	2030
SA-capable device share	10%	80%
Active slice instances	5	50

Note: Slice count influences orchestration/assurance workload and the extent to which profiling benefits can be realised; values are illustrative.

Figure A.3: (b) Key spectrum assumptions per scenario stage

Stage	Spectrum assumption
Stage 0	NR at 700 (10 MHz), 2.1 (20MHz), 3.5 (100MHz); remaining spectrum on LTE
Stage 1	Refarm 1800 (10MHz) to NR
Stage 2	Refarm 800 (10MHz) to NR

Note: NR carriers can carry NSA and SA traffic; realised SA benefits are constrained in the early years by SA-capable device/traffic share. Mid-band deployment continues in both scenarios where required by capacity demand and is not treated as an SA-specific cost.

Note: RAN enablement assumptions in the SA-led scenario include SA RAN software activation across the existing 5G footprint and 30% radio modernisation in refarmed bands, with the remainder activated by licence. Transitional RAN engineering, optimisation and voice-readiness effort is acknowledged qualitatively but not isolated as a separate major cost line.

Annex B Survey results label definitions

Figure B.1: NBED Figure B.4: Operational efficiency chart labels

Short label in figure	Meaning in the survey / how to read it
Remote-first	More troubleshooting and maintenance handled remotely, reducing site visits and truck rolls
Automated service provisioning	Faster, more automated service activation and configuration
Network planning	Better demand forecasting and shorter planning cycles
Fault correlation etc	Fewer tickets and faster mean-time-to-better fault correlation, isolation and resolution, reducing manual troubleshooting effort
Intent-based orchestration	Customer intent can be translated into policy or declarative workflows in a more automated way
Dynamic core scaling	More efficient scaling of core network functions as demand changes.
Converged infra	Shared infrastructure/functions across use cases instead of separate environments for e.g. IoT, FWA etc.
Energy – RAN	Lower RAN energy consumption through more efficient radio/network operation.
CI/CD	Continuous integration/continuous delivery more automated software release, testing and deployment processes
Energy – network function and core	Lower energy use in network functions and core through more efficient scaling and resource use.

Figure B.2: Capex efficiency chart labels

Short label in figure	Meaning in the survey / how to read it
RAN	Capex efficiency from increased use of NR and more efficient radio resource use.
Improved capacity planning	Better observability and agility improve upgrade timing and reduce overbuild.
Scaling	More efficient core scaling lowers the capex needed to support traffic growth
Spectrum	More efficient use of spectrum reduces the need for incremental capacity investment
Transport slicing	More efficient transport investment through traffic isolation / deterministic forwarding
Shared core between use cases	Reuse of the same core across services such as IoT, FWA, PPDR and private networks
Traffic profiling	Better knowledge of traffic patterns allows more targeted provisioning and lower excess capacity
Shared RAN between use cases	Reuse of the same RAN footprint across more services reduces separate build requirements
Dynamic upgrades	Smaller, more targeted upgrades rather than larger step changes in capacity investment

Figure B.3: Slicing / slice-automation chart labels

Short label in figure	Meaning in the survey / how to read it
Network planning	Better visibility and agility improve planning for slice-based services
Fault correlation etc	Faster remote troubleshooting and better issue isolation by slice
Cross-domain QoS	More coordinated QoS control across core, RAN and transport
SLA	Fewer SLA violations and more reliable delivery of differentiated services
Automated slice management	Lower effort to create, change, scale and retire slices
Dynamic automated scaling	Lower effort to scale services or slice capacity dynamically
Replace separate functions	Use slicing to avoid building separate dedicated networks/functions for each use-case
Automated onboarding	Faster customer/service onboarding with higher intent-to-activation rate
Blast radius containment	Reduced outage/incident impact by containing issues within a slice boundary
Energy	Better traffic information and partitioning support more efficient energy management
Test slices	Faster testing and lower roll-out risk
Reusable templates	Faster deployment and lower effort to build automation for each new slice
Targeted automations/slice	More closed-loop automations and higher automation success rate using slice-specific data

Annex C Glossary

Term	Explanation
5G NSA	5G new radio deployed with LTE anchoring and legacy core dependencies.
5G SA	5G standalone deployment using the 5G core, enabling cleaner end-to-end service control.
5G core (5GC)	The core network architecture for 5G SA, supporting service-based interfaces and policy control.
Network slicing	Partitioning network resources and policies to support differentiated service behaviour and assurance.
Static slicing	A limited set of predefined slice instances with relatively stable configuration and governance.
Dynamic slicing	Policy-driven slice lifecycle management and scaling that adapts to demand and service-level agreement (SLA) requirements.
Orchestration	Automation and coordination of lifecycle actions across network functions and domains.
Automation coverage	The share of targeted workflows executed without manual steps, measured for key processes.
OSS	Operations support systems used for network management, assurance and lifecycle processes.
NOC/SOC	Network operations centre / service operations centre functions responsible for monitoring and incident handling.
CNF	Cloud-native network function, typically deployed as containerised software on cloud infrastructure.
Cloud-native	Operational model and architecture built around container platforms, automation pipelines and frequent change cycles.
DMC	Dual-mode core, supporting both legacy core operation and 5G core capabilities during transition.
LCM	Lifecycle management of network functions and services (deploy, change, scale, retire).
MTTR	Mean time to repair/restore, often split into detection, diagnosis and remediation phases.
CI/CD	Continuous integration/continuous delivery, automated pipelines for testing and releasing changes.
GitOps	Operating model where desired system state is declared and enforced through version-controlled workflows.
Observability	Ability to understand service and network behaviour through metrics, logs and traces.
Closed-loop operations	Automated detection and remediation actions triggered by policies and telemetry signals.
FWA	Fixed wireless access, broadband delivered over mobile radio access.
PPDR	Public protection and disaster relief services often require prioritisation and high assurance.
TCO	Total cost of ownership for the network cost-base over the defined horizon.
Capex / opex	Capital expenditure (investment) and operating expenditure (run costs).