

Ioannou, Nikos; Kokkinis, Dimitris; Katsianis, Dimitris; Varoutas, Dimitris

Conference Paper

5G RAN Slicing for Industry 5G Verticals: A Comparative Techno-Economic Analysis of Private Network Deployment Versus Network Slicing

ITS 33rd European Conference 2025: "Digital innovation and transformation in uncertain times", Edinburgh, UK, 29th June – 1st July 2025

Provided in Cooperation with:

International Telecommunications Society (ITS)

Suggested Citation: Ioannou, Nikos; Kokkinis, Dimitris; Katsianis, Dimitris; Varoutas, Dimitris (2025) : 5G RAN Slicing for Industry 5G Verticals: A Comparative Techno-Economic Analysis of Private Network Deployment Versus Network Slicing, ITS 33rd European Conference 2025: "Digital innovation and transformation in uncertain times", Edinburgh, UK, 29th June – 1st July 2025, International Telecommunications Society (ITS), Calgary

This Version is available at:

<https://hdl.handle.net/10419/331278>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

5G RAN Slicing for Industry 5G Verticals: A Comparative Techno-Economic Analysis of Private Network Deployment Versus Network Slicing

Nikos Ioannou¹, Dimitris Kokkinis¹, Dimitris Katsianis², and Dimitris Varoutas^{1*}

Abstract

The strategic decision to build a private 5G network or buy a 5G network slice from a telecom operator presents a critical techno-economic dilemma for industries requiring advanced automation solutions that rely on connectivity infrastructure. This study addresses this decision by conducting a comparative techno-economic analysis specifically targeting three distinct 5G industry vertical scenarios: a seaport (outdoor), an automotive factory (indoor), and a power plant campus (hybrid indoor and outdoor). The modeling of the scenarios is based on data from real-world deployments and industry benchmarks and is used by the techno-economic bottom-up model to estimate the economic impact of 5G RAN slicing and 5G private network deployments per scenario, emphasizing metrics such as Capital Expenditure (CAPEX), Operational Expenditure (OPEX) and Net Present Value (NPV) of costs. The results show that RAN slicing can reduce CAPEX by 33–68% and OPEX by 12–51% compared to private 5G deployments, primarily through macro-site reuse and shared resources but at the expense of customization and the highest isolation and security levels, which can greatly reduce the potential cost savings. This study identifies which industries benefit from the customization and control offered by owning infrastructure and which industries can scale and deploy more efficiently via slicing arrangements, offering organizations actionable insights to align their strategic connectivity decisions with financial resources, operational priorities, and technical needs.

1 Introduction

The rollout of 5G has opened a new chapter in industrial automation by combining gigabit data rates, millisecond-level latency, and massive device density within a single wireless framework. However, industry firms face a pivotal choice when adopting 5G-enabled solutions: deploying a stand-alone private network that they own and manage end-to-end or purchasing an operator-managed network slice of a public Radio Access Network (RAN). The economic and operational stakes of this decision are particularly high for mission-critical verticals that depend on continuous deterministic connectivity.

Recent studies have explored the performance benefits of private networks [1] and slicing [2], but few have quantified the full life cycle cost implications in realistic industrial settings. Moreover, the techno-economic calculus changes noticeably with the physical characteristics of the site, such as an open-air seaport with moving cranes and vessels, an indoor automotive factory dominated by metal machinery and autonomous robots, or a hybrid power-plant campus that mixes control rooms, turbine halls, and high-voltage yards. These contrasting

¹ Dept. of Informatics and Telecommunications, School of Science, National and Kapodistrian University of Athens, Athens, Greece

² Dept. of Digital Industry Technologies, School of Science, National and Kapodistrian University of Athens, Psachna, Greece

* Corresponding author: D.Varoutas@di.uoa.gr

environments drive different radio planning assumptions, performance requirements, spectrum strategies, and service-level agreements (SLAs), ultimately shaping both CAPEX and OPEX.

Building a private 5G network enables the tight tailoring of radio parameters, security policies, and quality-of-service (QoS) guarantees to the use-case at hand. Ownership also allows enterprises to maintain sensitive telemetry and intellectual property on-premises, in alignment with strict safety or regulatory mandates. However, private deployments entail significant up-front outlays for spectrum licenses, gNodeB equipment, and mobile core functions, as well as a long-term commitment to in-house expertise for maintenance and upgrades [3].

Conversely, purchasing a 5G slice transfers most of the infrastructure burden to the mobile network operator. Slices can be provisioned quickly, scaled elastically, and paid for as operating expenses, which appeals to firms with limited capital budgets or uncertain traffic forecasts. However, reliance on a third party may constrain customization, expose production data to external networks, and lock the enterprise into recurring fees that could surpass ownership costs over time [4].

Given these trade-offs, a rigorous techno-economic evaluation is required. This study develops a bottom-up Discounted Cash Flow (DCF) model for three representative European sites, a container seaport (outdoor), an automotive factory (indoor), and a gas-fired power-plant campus (hybrid), and populates it with benchmark data from commercial pilots and vendor quotes. Monte Carlo simulations capture the uncertainty in traffic growth, spectrum pricing, energy costs, and equipment lifecycles. Key indicators include OPEX and CAPEX and NPV of costs under each deployment option.

Our contribution is three-fold:

1. We provide a side-by-side cost comparison of private 5G and RAN-sliced solutions in heterogeneous industrial environments.
2. We identify the technical drivers that swing the economics toward ownership or slicing.
3. We translate the modelling results into actionable guidelines for operations managers and policymakers seeking to foster efficient 5G adoption in Europe.

The remainder of this paper is structured as follows: Section 2 describes the three case study environments and their connectivity requirements. Section 3 outlines the techno-economic modeling framework. Section 4 presents the comparative results for CAPEX, OPEX and NPV cost analyses and discusses the findings. Finally, Section 6 summarizes the study and identifies potential directions for future research.

2 Industry Vertical Scenarios

Table 1 illustrates the Key Performance Indicators (KPIs) that shape the network design in each of the three industrial settings analyzed in this study. The metrics were derived from vendor white papers, operator pilot results, and ITU-R M.2410 baseline requirements [5], and subsequently validated against publicly available deployment reports from reference sites (e.g., Piraeus and Hamburg ports, Stuttgart automotive plants, and combined-cycle power stations in northern Europe).

Table 1: Connectivity requirements for the three case-study environments.

Metric	Seaport Terminal	Automotive Factory	Power-plant Campus
Typical footprint	10 km ² mixed quays & roads	0.5–1 km ² metal-dense halls & yards	2–4 km ² outdoor yard plus reinforced buildings
Latency budget	≤ 20 ms	< 1 ms	5–10 ms
Connection Density	1 up to 5×10^4 devices / km ²	0.1 up to 1×10^6 devices / km ²	0.5 up to 1×10^5 devices / km ²
Reliability target	99.999%	99.9999%	≥ 99.9999%
Resource Isolation	Full Sharing	Partial Isolation	Full Isolation
Mobility	High: 30–50 km/h	Moderate: 5–20 km/h	Low: 0–10 km/h

2.1 Seaport Container Terminal

Modern megaports orchestrate continuous container flows between quay cranes, automated guided vehicles (AGVs), stacking yards, and local transport links. Wide area coverage (up to 60 km²) must accommodate moving assets and shifting line-of-sight conditions caused by stacked containers and metal superstructures. Time-critical applications include remote crane control and traffic coordination, both of which demand sub-20 ms round-trip latency and slice isolation to guarantee safety [6]. Meanwhile, the high-bandwidth eMBB capacity supports HD closed-circuit television (CCTV) and augmented-reality (AR) guidance for maintenance crews. Device mobility drives the need for macro cell overlays combined with beamforming small cells on crane booms.

Table 2: Seaport Terminal Device Distribution Model.

Device	Estimated Quantity	Primary 5G Service
Smart Container & Asset Trackers	90,000	mMTC
HD CCTV & Surveillance Drones	200	eMBB
AR Headsets for Maintenance	50	eMBB
AGVs	100	URLLC
Remote-Controlled Cranes	20	URLLC
Connected Traffic Lights	10	URLLC

2.2 Automotive Assembly Factory

Car-manufacturing plants pack hundreds of robots, collaborative robots (cobots), and programmable logic controllers (PLCs) into metal-rich halls. Closed-loop motion control and real-time quality inspection require sub-millisecond user-plane latency [7]. The connection density can exceed one million endpoints per square kilometer when wireless sensors, tool heads, and AR headsets are counted. Electromagnetic reflections and Faraday-cage effects necessitate dense indoor small-cell grids and precise time-synchronization [8]. Reliability targets reach 99.9999% for production lines, where unexpected downtime translates directly into lost output.

Table 3: Power-plant Campus Device Distribution Model.

Device	Estimated Quantity	Primary 5G Service
Wireless Production Tools & Trackers	250,000	mMTC
Cobots	500	URLLC
Motion Control Systems	1,000	URLLC
AGVs for Parts Delivery	150	URLLC
Quality Control Cameras	100	eMBB

2.3 Power-Plant Campus

Gas-fired and coal-fired power stations extend over several square kilometers and combine control rooms, turbine halls, cooling towers, and high-voltage switchyards. Mission-critical SCADA traffic and protection relays tolerate latencies up to 10 ms but demand availability above 99.9999% [9]. Massive machine-type communications (mMTC) support thousands of sensors that monitor vibrations, temperatures, and emissions. Although mobility is limited, RF propagation must contend with reinforced concrete, high-EMI equipment, and outdoor weather extremes, motivating a hybrid topology of ruggedized indoor small cells and outdoor macro sites.

Table 4: Seaport Terminal Device Distribution Model.

Device	Estimated Quantity	Primary 5G Service
Predictive Maintenance Sensors	100,000	mMTC
Perimeter Security Sensors	2,000	mMTC
SCADA & Grid Control Systems	500	URLLC
Differential Protection Relays	100	URLLC
AR Headsets for Remote Support	20	eMBB
Inspection Drones & Cameras	15	eMBB

3 Techno-economic Modeling

The analysis follows a bottom-up techno-economic modeling framework over a 10-year horizon and is illustrated in Figure 1. The approach comprises four sequential phases: (1) Scenario Design, (2) Coverage and Capacity Modeling, (3) Network Dimensioning, and (4) Cost Analysis and Allocation. Each phase builds on the previous one to estimate the required network resources and the associated CAPEX and OPEX under the two deployment options (private network vs. network slicing). The techno-economic model is based on an extended version of [10].

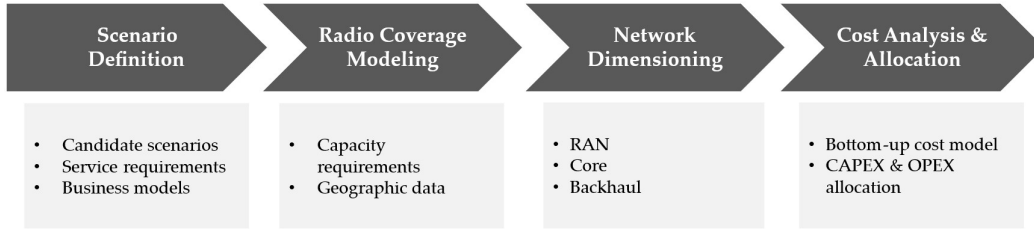


Figure 1: Four-phase methodology of the study.

3.1 Scenario Design

Three representative industrial sites are defined with distinct physical footprints and service requirements, as defined in Section 2.

Each scenario supports three 5G service classes and respective slices (eMBB, mMTC, URLLC) to capture broadband, massive IoT, and ultra-reliable low latency use cases. For planning purposes, we assume each eMBB device requires on average 100 Mb/s. Stringent latency and reliability targets influence the placement of edge computing (MEC) and redundancy levels as discussed later.

Within each scenario, we compare two business models: Private 5G and 5G Slicing by Mobile Network Operator (MNO). In the Private model the enterprise builds and operates its own network. This entails 100% greenfield deployment of 3-sector macrocell sites over the area, full allocation of dedicated spectrum (we assume exclusive use of the local 3.5 GHz band), a dedicated 5G core, and all CAPEX/OPEX borne by the enterprise. In the 5G Slicing model the mobile operator provides service using existing infrastructure: approximately 90% of the required macro sites and 50% of fiber backhaul are reused from the operator's network (only ~10% new macrocells and all new small cells are added). The slicing model shares spectrum and core: spectrum is provisioned by the operator (up to 3.5 GHz) as needed and the operator's shared 5G core is used.

In both models, indoor small-cell deployment is assumed to be fully enterprise-funded (100% greenfield) for coverage. Cost sharing differs in Private, all costs are assigned to the enterprise, while in Slicing, indoor-cell costs (and any dedicated slice equipment) remain with the enterprise, and all other costs are allocated by network usage and isolation level. These assumptions affect how shared resources and Virtual Network Functions (VNFs) are partitioned (see Phase 4).

Table 5: Alternative business models and deployments.

Attribute	Private 5G	5G slicing by MNO
Ownership & Operation	Fully owned and operated dedicated network	Operated by commercial MNO.
Infrastructure Deployment	100% Greenfield	~90% Brownfield (reuse of existing MNO macro-cell infrastructure); 100% Greenfield for small and indoor cells.
Spectrum Utilization	All spectrum used in the area in 3.5GHz band	Based on capacity requirements in the area in all 5G bands up to 3.5GHz.
Network Coverage	3 Sector Macro-cell base stations and small cells for outdoor & indoor	3 Sector Macro-cell base stations and small cells for outdoor & indoor.
Core Network	Dedicated core network	Shared core network.
Cost Allocation	100% own-funded (no cost-sharing)	100% of indoor deployment, Rest of costs allocated proportionally to network resources usage based on level of isolation.
Infrastructure Reuse	0%	High reuse (about 90% of existing commercial infrastructure), including 50% of fiber backhaul.

3.2 Coverage and Capacity Modelling

In this phase we compute the radio coverage and capacity needs from the scenario inputs. We model outdoor coverage with 3-sector macro-cell sites (radius $R_{macro} = 1$ km) and outdoor small cells ($R_{small,in} = 0.3$ km), plus indoor small cells ($R_{small,in} = 0.1$ km) for indoor areas. Let A_{out} and A_{in} be the outdoor and indoor areas to be covered. Under the assumption of standard hexagonal coverage per cell (or sector), the number of sites is determined by area:

$$N_{macro} = \left\lceil \frac{A_{out}}{\frac{3 * \sqrt{3}}{2} * 3 * \left(\frac{3}{2} * R_{macro}\right)^2} \right\rceil \quad (1)$$

$$N_{small,out} = \left\lceil \frac{A_{out}}{\frac{3 * \sqrt{3}}{2} * R_{small,out}^2} \right\rceil \quad (2)$$

$$N_{small,in} = \left\lceil \frac{A_{in}}{\frac{3 * \sqrt{3}}{2} * R_{small,in}^2} \right\rceil \quad (3)$$

These expressions ensure full-area coverage by macrocells and small cells. Next, we assess capacity-needed deployment which is driven mainly by eMBB devices assumed to require up to 100Mbps connection speed as

$$R_{tot} = N_{eMBB} * 100 \frac{Mb}{s} \quad (4)$$

where N_{eMBB} is the number of broadband devices.

A single small cell can support

$$C_{small} = B * h \quad (5)$$

Where B is the assumed maximum spectral efficiency $h = 9$ bits/s/Hz and bandwidth B according to each scenario. Thus, the number of small cells required by capacity is

$$N_{small,cap} = \left\lceil \frac{R_{tot}}{C_{small}} \right\rceil \quad (6)$$

The final small-cell count is taken as the maximum of the coverage- and capacity-based estimates:

$$N_{small} = \max(N_{small,cov}, N_{small,cap}) \quad (7)$$

Ensuring that both coverage and demand are satisfied.

3.3 Network Dimensioning and Resource Allocation

Given the number of macro and small-cell sites, we enumerate all required network elements (RAN, backhaul, core, MEC, and slice-related) by applying engineering rules. Table 6 summarizes the key dimensioning rules and quantities. For RAN infrastructure, the Private model uses all required macro sites as new sites, whereas in the Slicing model only 10% macro-sites are newly built (the rest are reused). All required outdoor and indoor small cells are deployed new in both models. We then apply the following rules to compute equipment counts and civil works.

For Core Network and MECs, the network requires distributed edge computing and centralized core data centers to host the 5G core functions and slice management entities. We assume a multi-tier core. One MEC is placed per M macro sites based on the latency required in each scenario. We set $M=20$ for very low latency (≤ 1 ms) scenarios and $M=50$ for higher latency (up to 20 ms).

Slicing is modeled by VNFs and respective servers in DC and MECs. VNF sizing rules of one user-plane VNF instance about 1250 eMBB users, or 100,000 IoT devices, or 1,000 URLLC connections are respective Control Plane VNF capacity of 20,000 devices are based on [11][12]. The redundancy is provisioned according to the level of reliability:

- 99.999% leading to x2 redundancy only for user plane VNFs.
- 99.9999% leading to x2 redundancy for all VNFs.
- 99.99999% leading to x2 redundancy for all VNFs and respective hardware.

Table 6: Key Network Dimensioning Rules in the Techno-Economic Model.

Asset / Network Element	Dimensioning Rule
Macro site (3-sector), new site	10% of sites MNO 5G Slicing and 100% of sites of 5G private network
Macro site, reuse of existing	90% of sites MNO 5G Slicing and 0% of sites of 5G private network
Small-Cell outdoor site, new site	100% of sites
Small-Cell indoor, new site	100% of sites
Civil works & Macro site acquisition, new	1 per new macro site (full construction)
Civil works, existing Macro site upgrade	1 per reused macro site (reinforcement/upfit)
Civil works Small cell, new	1 per new small-cell site (full construction)
Civil works Indoor cell, new	1 per new indoor cell (full construction)
Equipment Macro Cell	3 per macro site (three sector antennas/RRHs)
Equipment Small Cell	1 per small cell
Equipment Small Cell	1 per indoor cell
Installation Macro Cell	1 per sector (installation per macro cell)
Installation Small Cell	1 per sector (installation per small cell)
Installation Indoor Cell	1 per sector (installation per indoor cell)
Baseband unit (BBU)	1 per site (baseband processing unit)
Edge Server Macro Cell	1 per site (site server for local functions)
Edge Server Small Cell	1 per site (site server for local functions)
Edge Server Indoor Cell	1 per site (site server for local functions)
Connection to power grid	1 per site (outdoor only)
Connection to fiber network	1 per site (outdoor only)
Fiber Backhaul	Calculated per km using spanning tree model
MEC node	1 per 30 sites (MEC location density)
MEC installation	1 per MEC site (setup per edge facility)
MEC racks & servers	4 racks per MEC, 5 servers per rack (20 servers per MEC)
Core data center (central DC)	1 per 500 sites (number of central DCs)
DC racks, servers, switches	max 10 racks, 50 servers, 20 switches per DC (for core VNFs)
Network Management & Orchestration	1 per network (one NM&O system for slicing)
Slice license	1 per slice type (three slice licenses total)
User-plane VNF (eMBB)	1 per 1250 eMBB devices
User-plane VNF (mMTC)	1 per 100,000 IoT devices
User-plane VNF (URLLC)	1 per 1,000 URLLC devices
Control_Plane VNF (eMBB)	1 per 20,000 devices
Control_Plane VNF (mMTC)	1 per 20,000 devices
Control_Plane VNF (URLLC)	1 per 20,000 devices

3.4 Cost Modelling and Allocation

Using the dimensioned assets, we compute total CAPEX and OPEX with a bottom-up cost model. Each asset type is assigned a unit purchase cost and an annual Operational and Maintenance (O&M) cost (sourced from industry and regulatory benchmarks [13][14]) and an economic lifetime. OPEX includes energy (€0.20/kWh [15]) and maintenance. We also assume 2.5% annual inflation on expenditures and a 10% overall administrative overhead on total OPEX. All future cash flows (costs in our analysis) are discounted at 8% per annum to compute NPV over the 10-year study period.

Table 7: Cost assumptions per Asset / Network Element.

Asset	CAPEX	OPEX	Lifetime	CAPEX trend	Cost Category
Spectrum (MHz/pop)	0.05 €		15		Spectrum
Site acquisition Macro Cell site	6,000 €	3,000 €	25	-1.00%	Site/Tower
Site acquisition Small Cell site	2,000 €	800 €	25	-1.00%	Site/Tower
Civil works Macro Cell New site	35,000 €	700 €	25	-1.00%	Site/Tower
Civil works Small Cell New site	10,000 €	200 €	25	-1.00%	Site/Tower
Civil works Indoor Cell New site	1,000 €	20 €	25	-1.00%	Site/Tower
Civil works Macro Cell Existing site	3,000 €	60 €	25	-1.00%	Site/Tower
Equipment Macro Cell	12,000 €	4,104 €	10	-4.00%	RAN
Equipment Small Cell	5,000 €	775.60 €	10	-4.00%	RAN
Equipment Indoor Cell	5,000 €	600.40 €	10	-4.00%	RAN
Installation Macro cell	7,000 €	140 €	10	-1.00%	RAN
Installation Small cell	2,000 €	40 €	10	-2.00%	RAN
Installation Indoor cell	1,000 €	20 €	10	-2.00%	RAN
Server Macro Cell	2,000 €	100 €	8	-3.00%	RAN
Server Small Cell	1,500 €	75 €	8	-3.00%	RAN
Server Indoor Cell	1,500 €	75 €	8	-3.00%	RAN
Baseband Unit	1,000 €	50 €	8	-4.00%	RAN
Fiber Backhaul	10,000 €	200 €	35	-1.00%	Backhaul
Fiber Indoor	1,000 €	100 €	25	-1.00%	Backhaul
Connection to fiber network	1,000 €	20 €	20	-1.00%	Backhaul
MEC infrastructure	10,000 €	2,240 €	15	-1.00%	Core
Connection to power grid	1,000 €	20 €	20	-1.00%	Site/Tower
MEC infrastructure	12,000 €	2,240 €	15	-1.00%	Core
MEC installation	10,000 €	500 €	15	-3.00%	Core
MEC racks	10,000 €	200 €	12	-3.00%	Core
MEC servers	3,000 €	500.40 €	7	-4.00%	Core
Management & Orchestration	500,000 €	100,000 €	5	-4.00%	Core
Slice Licence	0 €	100,000 €	1	-1.00%	Core
VNF Licence	0 €	100 €	1	-1.00%	Core
DC infrastructure	20,000 €	400 €	15	-1.00%	Core
DC Racks	10,000 €	200 €	12	-3.00%	Core
DC Servers	10,000 €	850.40 €	7	-4.00%	Core
DC Switches	2,000 €	275.20 €	7	-3.00%	Core

3.5 Cost Allocation within Scenarios

Finally, costs are allocated between the vertical and the MNO as follows. In the private network case, 100% of CAPEX and OPEX are assigned to the enterprise. In the 5G slicing case, by convention the enterprise fully funds indoor small cells and any dedicated slicing resources while the remaining shared-network costs are split by usage (e.g. cell throughput) and slice isolation requirements, as summarized in Table 8. Concretely, in the fully shared isolation level the cost of all network (except indoor) is allocated based on capacity usage by the devices. In

the partial isolation the cost allocation is similar to shared isolation except for slicing assets/equipment which are shared based on the ratio of User Plane VNFs (dedicated) to Total VNFs. In the fully isolated scenario, the cost of slicing assets/equipment is fully allocated to the vertical.

Table 8: Cost allocation rules based on network isolation level.

Isolation level	Cost Allocation Rule
Shared	Cost of network (except indoor) is allocated based on capacity usage.
Partially	Slicing assets/equipment are allocated based on the ratio of User Plane VNFs to Total VNFs.
Full	Slicing assets/equipment are fully allocated.

4 Results and Discussion

4.1 Cost Comparison

The findings indicate that the total costs (CAPEX and OPEX), are lower under the network slicing model provided by operators compared to private 5G, across all three scenarios.

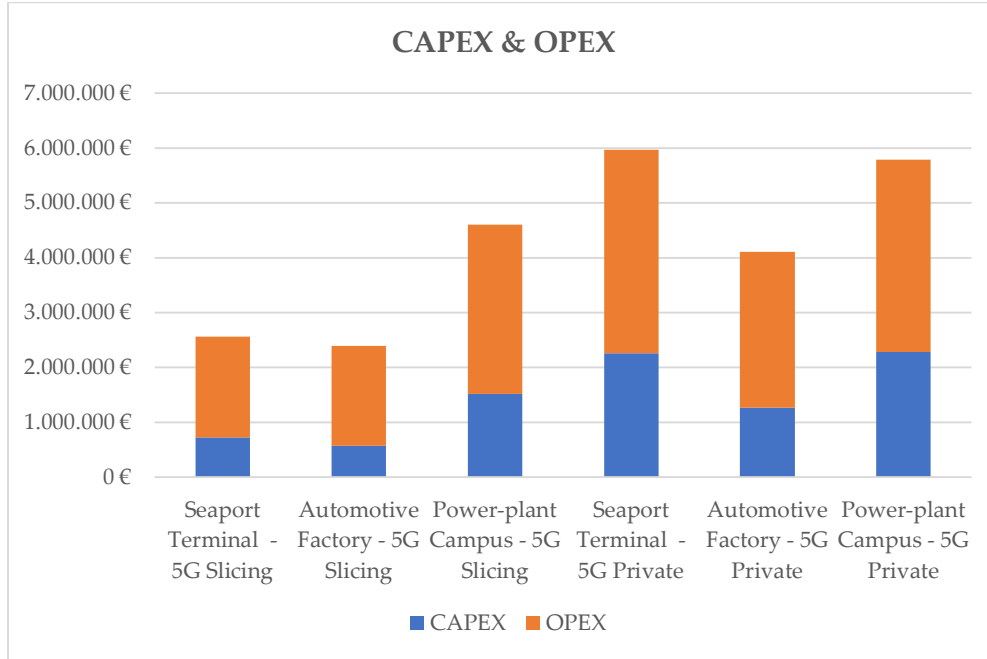


Figure 2: CAPEX & OPEX comparison between all cases and scenarios.

For the seaport terminal, for example, the private 5G deployment requires roughly €2.257M in CAPEX and €3.714M in ten-year OPEX (total NPV ~€4.207M), whereas the slicing model yields only €0.724M CAPEX and €1.838M OPEX (NPV ~€1.745M). This corresponds to a 68% reduction in CAPEX and 51% reduction in OPEX when using slicing (NPV falls by 59%). Likewise, in the automotive factory scenario, CAPEX drops from €1.265M (private) to €0.569M (slicing) and OPEX from €2.840M to €1.825M (reductions of 55% and 36%, NPV down 43%). The power-plant campus, which mixes indoor and outdoor zones, sees smaller but still

significant savings: CAPEX €2.281M vs €1.519M and OPEX €3.507M vs €3.090M (slicing gives –33% CAPEX and –12% OPEX, NPV –25%). In all cases, the private 5G option incurs roughly 2–3× higher up-front capital and ongoing O&M than the slicing alternative. In short, slicing consistently dominates on cost: CAPEX and OPEX are substantially lower in every scenario.

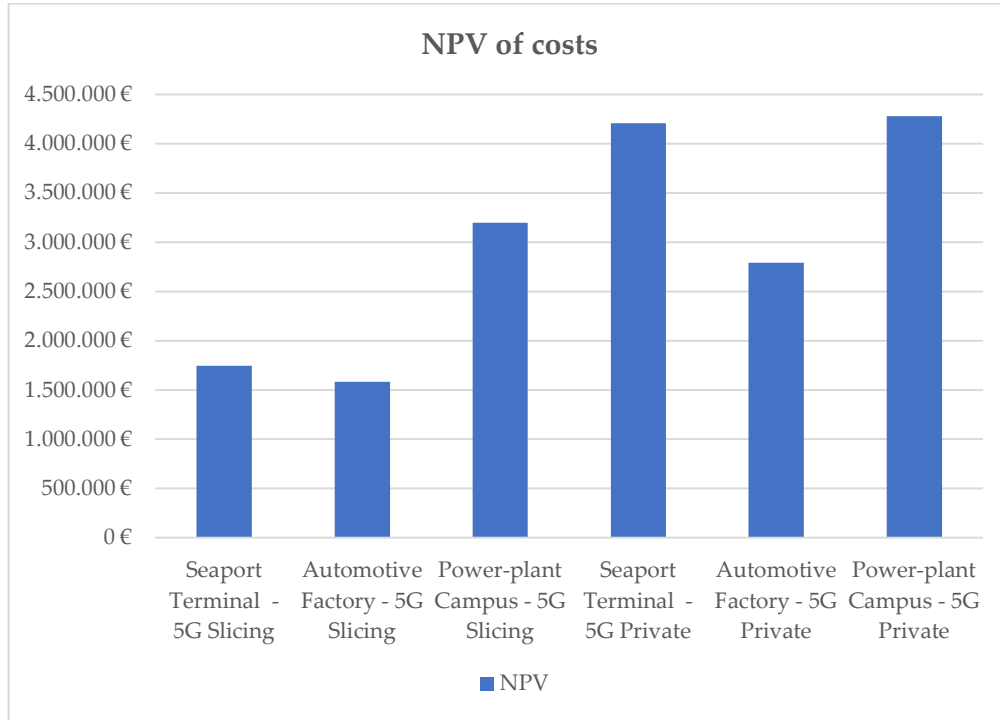


Figure 3: NPV of costs comparison between all cases and scenarios.

The breakdown of costs reveals why. Licensed spectrum costs remain negligible (<0.5% of total) in every scenario and model. The main differences lie in RAN, core, and site costs. Under slicing, most macro - cell sites and backhaul are reused, so site/tower costs fall sharply. For example, in the seaport case the combined site and RAN share of total cost drops from ~41.3% to ~23.6% (17.3% RAN + 6.3% site in slicing). On the other hand, core-network functions (and shared backhaul) become a larger fraction: e.g. seaport core costs rise from ~47.1% (private) to ~67.4% (slicing) of total. In the automotive and power-plant cases similar patterns emerge in slicing mode, core operations dominate (~63-77% of cost) while enterprise - funded RAN/site equipment shrinks to ~15-30%. Thus, the economic advantage of slicing is driven by reusing existing infrastructure (reducing RAN and site CAPEX) at the expense of relying on the operator's core (raising its cost share). Across scenarios, slicing cut total CAPEX by 68%, 55%, and 33% (seaport, factory, power) and OPEX by 51%, 36%, and 12%.

4.2 Technical Drivers Shaping the Business Case

Several technical and economic factors underlie the results. First, infrastructure reuse and spectrum allocation differ. The slicing model assumes the MNO can reuse about 90% of macro sites. This reuse dramatically lowers enterprise CAPEX and off loads costs to future discounted OPEX in NPV terms. In contrast, a private network is built greenfield with exclusive 3.5 GHz spectrum. The need for dedicated spectrum licenses and full control means private deployments incur large up-front costs in base stations and core network functions which greatly reduce investment viability.

Scale, network management and device scale play a key role. Our sensitivity analysis (tornado charts) shows that in private networks the enterprise's O&M cost and coverage area have the biggest impact on total cost, along with the number of mMTC (massive IoT) devices. In the slicing model, slice-licensing fees (capacity charges) and again mMTC device count dominate, while network management costs (e.g. orchestration, SLAs) remain significant. In practice, a larger IoT deployment will raise costs in both models, but especially increases fees paid to the operator in slicing. Likewise, if an enterprise already has in-house expertise to manage a private network, that may favor ownership; but lacking that, it may prefer the managed-slice option even if it pays recurring fees.

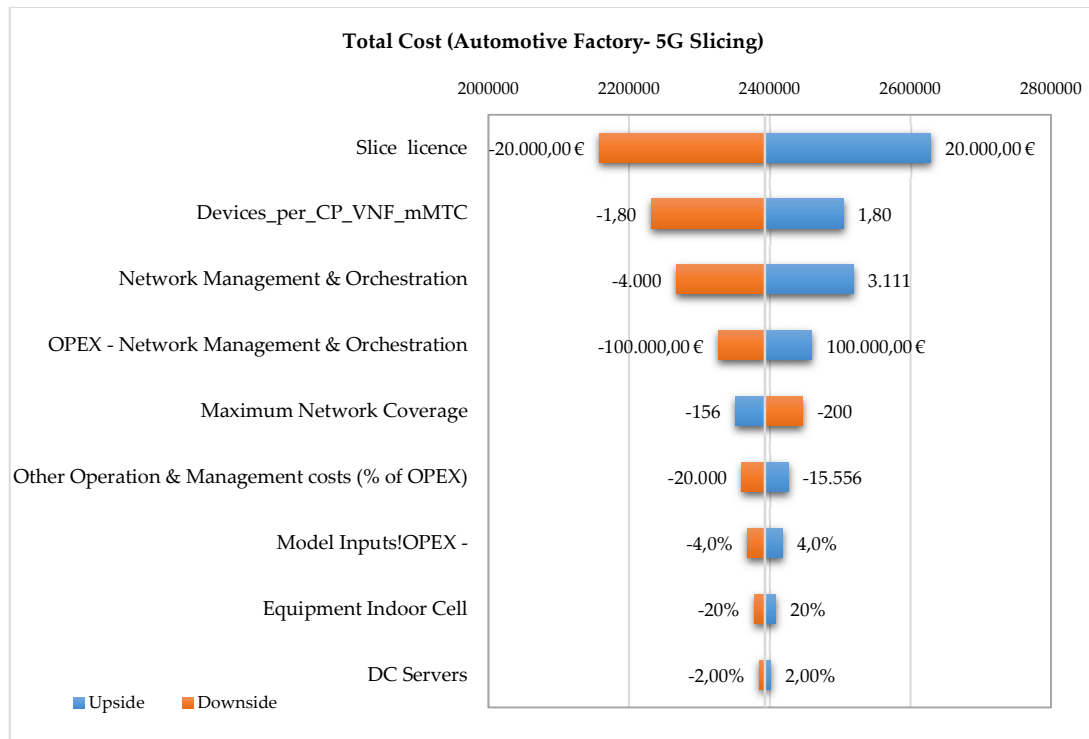


Figure 4: Sensitivity analysis of TCO for Automotive Factory 5G slicing case.

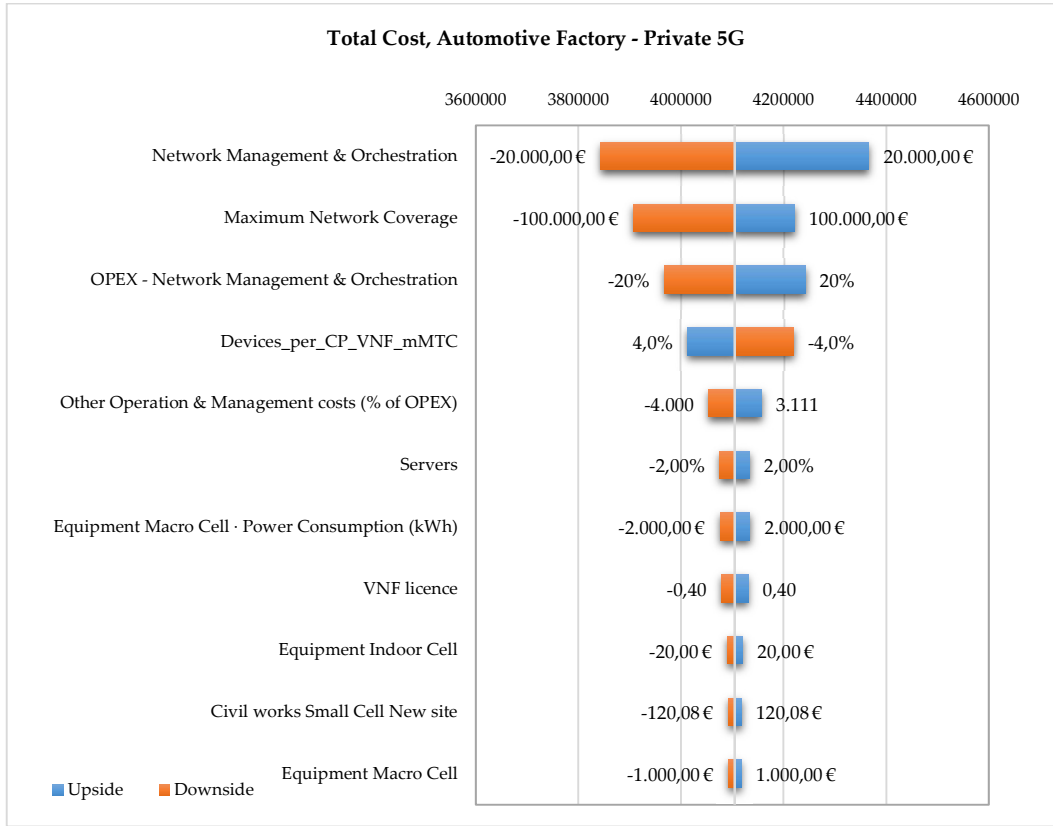


Figure 5: Sensitivity analysis of TCO for Automotive Factory private 5G case.

Finally, private networks allow full control over isolation and redundancy (e.g. fully dedicated slices and MEC), which can make meeting the highest reliability targets easier. The power plant scenario also assumes high resource isolation to protect critical traffic. In an operator slice, sharing with other tenants could risk contention unless strict SLAs are enforced. Thus, mission-critical verticals with absolute uptime mandates may lean toward private ownership despite higher cost, while less critical deployments can accept the slicing trade-offs. Especially in the case of power-plant where the reduction in total costs is less than 25% and given a reasonable profit margin for the MNO could lead even less cost saving for the power-plant. This indicates the importance of isolation and security considerations for the total costs.

4.3 Strategic Implications and Operational Guidelines

For operations managers the analysis highlights a clear cost–performance tradeoff. If the enterprise’s primary goal is to minimize CAPEX/OPEX and it can tolerate operator-driven SLAs, a 5G slice yields 30–70% lower costs. However, if the plant requires the lowest possible latency, highest reliability, or has rapidly growing device counts, owning a private 5G network may be justified despite higher expense. Managers should thus quantify their critical KPIs (latency, reliability, device volume) and compare them with the modeled scenarios. Importantly, network management overhead is a major driver: enterprises with limited telecom expertise might favor slicing to outsource O&M, whereas those with strong IT and operation teams can capitalize on private-network flexibility.

For MNOs, the results underscore the business case for offering industrial slices. Operators can repurpose existing infrastructure (macrocells and backhaul) and sell dedicated slice capacity at

a premium, gaining nearly all the core and spectrum revenue. Because core functions account for over two-thirds of total cost in slicing scenarios, operators should price slices to cover both amortized core investments and expected heavy loads from enterprise users especially in high mMTC and URLLC classes. The sensitivity findings suggest that operators could attract more enterprise demand by bundling network management/orchestration services (since these costs greatly affect customers) and by offering scalable pricing as device counts grow. Additionally, operators should invest in local MEC and flexible spectrum allocation to meet the strict latency and reliability needs of factory customers. The model shows a large profit opportunity in slicing (given its lower CAPEX), but also indicates that operators must maintain high SLAs and flexible offerings to win over industry clients.

The techno-economic insights can also inform regulatory support for industrial 5G. Policymakers should ensure affordable access and availability of licensed spectrum for private networks (e.g. local or shared licenses), since spectrum costs, while small in our model, are a factor for private deployment. Moreover, regulations that encourage infrastructure sharing (e.g. co-location incentives or neutral-host frameworks) will magnify the cost savings of slicing, as our model shows. Finally, given the variability across sectors, policies should not favor one model exclusively. Rather, by fostering competition among MNOs for slice offerings and facilitating private investment, policymakers can help the viability of both deployment modes.

5 Conclusion

This techno-economic comparison shows that leveraging 5G RAN slicing on a shared infrastructure can substantially lower the total network costs compared with deploying independent private systems. In the seaport and automotive factory case studies, slicing achieved large CAPEX/OPEX reductions and much faster break-even times, reflecting the high device counts and moderate reliability requirements of these environments. In contrast, the hybrid power-plant campus, with its extreme uptime and traffic-isolation needs, sees only a modest NPV gain (on the order of 20–30%) under a slicing model. This smaller savings gap arises because mission-critical sites require full control. In practice, operators at power facilities insist that control systems must be completely separated from external networks to maintain safety. Industries with less stringent latency/reliability targets (e.g., ports and most manufacturing) can often profitably outsource connectivity via slicing, whereas ultra-mission-critical sites may justifiably invest in full private networks despite the higher cost.

The results also underscore the key architectural trade-offs. Hardware reuse and virtualization greatly favor slices with lower per-unit spectrum and processing costs but cannot fully substitute the enhanced isolation and redundancy of dedicated RAN and core hardware. Private deployments provide complete isolation (and thus higher security/reliability), whereas slices, by definition, share infrastructure. For example, factory URLLC applications (<1 ms budget) and power-plant control loops benefit from localized MEC and distributed 5G core nodes to meet latency and six-nines reliability. Conversely, less critical slices (e.g., campus broadband or non-real-time IoT) can ride on centralized cores without on-site data centers.

The heterogeneous results imply that no connectivity model universally dominates, so policymakers and business strategists should have alternative choices available. Regulatory bodies should ensure spectrum availability for private enterprise use (through local or shared licensing schemes) while encouraging infrastructure sharing. Neutral-host frameworks,

incentives for co-locating on existing towers, and multi-operator RAN agreements can magnify the cost advantage of slicing. From the operator's standpoint, this analysis highlights a clear business case for industrial slice carriers to repurpose their existing sites and sell high-value virtual networks. Finally, enterprises must align their decisions with their skills and objectives. Firms with strong in-house IT/OT teams and mission-critical mandates may accept higher CAPEX for on-premises control, whereas others prefer to offload management and pay-as-they-go fees.

The quantitative conclusions of this study are based on modeling assumptions and input estimates that introduce uncertainty. Actual costs can significantly differ based on geography, vendor pricing, and technology generation. Moreover, our analysis focused on three European-style scenarios at a moderate scale, but the results could differ in megacities or developing economies. Other operational issues, such as the learning curve for private network teams or regulatory delays in spectrum assignment, are also not considered in the modelling. These limitations suggest that precise savings figures should be interpreted in context.

Future research could test various other cost allocation and estimation methodologies, such as long-run incremental cost (LRIC), to more accurately estimate annual costs for the operator and available pricing thresholds (up to the point when private 5G is preferred). Finally, as 6G architectures emerge, their effect on ownership models must be assessed because the respective technological advances might enable fully autonomous enterprise networks or radically new sharing schemes (e.g., dynamic spectrum sharing across terrestrial and satellite domains). Future analysis should explore how such trends and associated regulatory changes will shift the cost-benefit trade-offs of private vs. shared networks.

References

- [1] S. Roblot et al., "Techno-economic analyses for vertical use cases in the 5G domain," Jun. 2019, Accessed: Jan. 05, 2023. [Online]. Available: <http://arxiv.org/abs/1906.09746>
- [2] S. E. Elayoubi, S. Ben Jemaa, Z. Altman, and A. Galindo-Serrano, "5G RAN Slicing for verticals: Enablers and challenges," *IEEE Communications Magazine*, vol. 57, no. 1, pp. 28–34, Jan. 2019, doi: 10.1109/MCOM.2018.1701319.
- [3] A. Rostami, "Private 5G networks for vertical industries: Deployment and operation models," in *IEEE 5G World Forum, 5GWF 2019 - Conference Proceedings*, Institute of Electrical and Electronics Engineers Inc., Sep. 2019, pp. 433–439. doi: 10.1109/5GWF.2019.8911687.
- [4] M. Datar, E. Altman and H. L. Cadre, "Strategic Resource Pricing and Allocation in a 5G Network Slicing Stackelberg Game," in *IEEE Transactions on Network and Service Management*, vol. 20, no. 1, pp. 502–520, March 2023, doi: 10.1109/TNSM.2022.3216588.
- [5] Report ITU-R M.2410-0: "Minimum requirements related to technical performance for IMT-2020 radio interface(s)".
- [6] M. A. Uusitalo et al., "Ultra-Reliable and Low-Latency 5G Systems for Port Automation," in *IEEE Communications Magazine*, vol. 59, no. 8, pp. 114–120, August 2021, doi: 10.1109/MCOM.011.2001060.
- [7] D. S. Fowler et al., "A 5G Automated-Guided Vehicle SME Testbed for Resilient Future Factories," in *IEEE Open Journal of the Industrial Electronics Society*, vol. 4, pp. 242–258, 2023, doi: 10.1109/OJIES.2023.3291234.
- [8] E. A. Oyekanlu et al., "A Review of Recent Advances in Automated Guided Vehicle Technologies: Integration Challenges and Research Areas for 5G-Based Smart

- Manufacturing Applications," in IEEE Access, vol. 8, pp. 202312-202353, 2020, doi: 10.1109/ACCESS.2020.3035729.
- [9] Guanjin Huang, Meng Ye, Huaxu Zhou, Jianhui Li, Gengmiao Wang, and Xuhui Zhang. 2021. 5G network architecture and networking analysis of smart power plant. In Proceedings of the 2021 3rd International Conference on Information Technology and Computer Communications (ITCC '21). Association for Computing Machinery, New York, NY, USA, 120–126. <https://doi.org/10.1145/3473465.3473486>.
- [10] N. Ioannou, D. Kokkinis, D. Katsianis, and D. Varoutas, "Comparative Techno-Economic Evaluation of 5G Infrastructure Sharing Business Models in European Rural Areas", ITS 32nd European Conference 2023, Madrid, Spain, 19th, 20st June, 2023. Available online: <http://hdl.handle.net/10419/277973>.
- [11] Walia, J. S., Hämmäinen, H., Kilkki, K., Flinck, H., Yrjölä, S., & Matinmikko-Blue, M. (2021). A Virtualization Infrastructure Cost Model for 5G Network Slice Provisioning in a Smart Factory. Journal of Sensor and Actuator Networks, 10(3), 51. <https://doi.org/10.3390/jsan10030051>
- [12] Chiha Ep Harbi, A., Van der Wee, M., Colle, D., & Verbrugge, S. (2020). Network slicing cost allocation model. Journal of Network and Systems Management, 28(3), 627–659. <https://doi.org/10.1007/s10922-020-09522-3>
- [13] European Commission. (2024, December 18). Finalisation of the mobile cost model for roaming and the delegated act on a single EU-wide mobile voice call termination: SMART 2017/0091 [Report]. Shaping Europe's digital future. Retrieved June 10, 2025, from <https://digital-strategy.ec.europa.eu/en/library/finalisation-mobile-cost-model-roaming-and-delegated-act-single-eu-wide-mobile-voice-call-0>
- [14] GSMA Intelligence. (2021, January). The economics of mmWave 5G [Report]. GSMA Intelligence. Retrieved June 18, 2025, from <https://www.gsmainelligence.com/research/research-file-download?id=59768858&file=210121-Economics-of-mmWave.pdf>
- [15] Eurostat. (2024, February 1). Electricity prices – bi-annual data (from 2007 onwards) (nrg_pc_204): Reference metadata in Single Integrated Metadata Structure (SIMS). Eurostat. Retrieved June 10, 2025, from https://ec.europa.eu/eurostat/cache/metadata/en/nrg_pc_204_sims.htm