Machines Don't Watch Netflix

A Delphi Study-based Scenario Analysis of the Network for AI Ecosystem in 2035

Alex Chamoun and Robert Westring

LUNDS UNIVERSITET

LTH
LUNDS TEKNISKA HÖGSKOLA

Division of Production Management
Faculty of Engineering, LTH

Machines Don't Watch Netflix

A Delphi Study-based Scenario Analysis of the Network for AI Ecosystem in 2035

Published by
Division of Production Management
Faculty of Engineering LTH, Lund University
Box 118
SE-221 00 LUND
Sweden

Subject: Scenario Building and Strategy
Division: Production Management
Supervisor: Ola Alexanderson (LTH), ola.alexanderson@iml.lth.se and
Alex Wang (Ericsson), alex.c.wang@ericsson.com
Examiner: Izabelle Bäckström, izabelle.backstrom@iml.lth.se

Abstract

Mobile connectivity is set to change in the coming years due to AI generated data from agents and robotics. How this will develop remains uncertain since this type of data generation is not only a question of connectivity but also compute location and use cases. This creates a need for forward looking scenario analysis for ecosystem players so they can position themselves to win in a new paradigm.

The purpose of this research was to analyze how the introduction of artificial intelligence traffic would impact the information and communication technologies industry by 2035. It aimed to provide a prognosis of network traffic, scenarios that might play out and how the ecosystem might change in each scenario. The goal was not to predict a single correct future but to map the space of plausible ones and identify what conditions led to each scenario.

The research employed a two-round Delphi study as its primary method, gathering structured responses from 23 domain experts across academia, technology vendors, regulators, and industry associations. Industry reports were used as secondary data to contextualize and validate the expert findings.

Four scenarios looking toward 2035 were highlighted as possible futures, the Delphi panel was prompted on each scenario's plausibility. This resulted in one scenario with a strong positive consensus of 93 percent, two with 70 percent and one with no consensus splitting the panel at about 50 percent. Additionally, for each scenario: winners, losers, consequences and things that need to happen for the scenario to materialize were identified.

Three bifurcation points were found that have major implications for communication service providers (CSPs) and telecom equipment vendors. First, regulation is critical for CSPs. Market forces alone are not considered enough for them to gain traction. Second, the decoupling of infrastructure from economic value can be detrimental. If computing shifts to the edge but the economic value is still captured by the orchestrators, CSPs end up losing out, even though they are the ones supplying the compute. Third, geopolitical fragmentation makes data sovereignty a priority. This positions CSPs to displace global hyperscalers as primary compute providers.

Keywords: Delphi study, Scenario, Ecosystem, Telecom, Artificial intelligence, Information and communication technology

Preface

This master's thesis was written during the spring of 2026 marking the end of our master in electrical engineering. The thesis was written at the Division of Production Management, Faculty of Engineering at Lund University (LTH). It was done in collaboration with the Strategy & Portfolio at Ericsson's headquarters in Kista, Stockholm.

We want to direct our warm appreciation to our LTH supervisor Dr. Ola Alexanderson for his wise guidance and support. Your deep interest and knowledge in strategy and the wider technology sphere paired with your humor and wit have made the process fun and rewarding.

We also want to express our gratitude toward Alex Wang, our supervisor at Ericsson, for his encouragement, trust and the opportunity to carry out the work. His expertise and experience kept us on course to deliver a thesis with high industry relevance. Ericsson's Strategy & Portfolio team also deserves a big thanks for answering all our questions.

Our gratitude also goes out to everyone participating in the Delphi study. Your insights formed the foundation of our thesis. We deeply appreciate your contributions, we hope that you find the results interesting and that they might help you in your work.

Finally, we are very thankful for the support from our friends and families.

Thank you!

Alex Chamoun	Robert Westring
Lund, May 2026	*Lund, May 2026*

List of Abbreviations

AI	artificial intelligence
API	application programmable interface
AR	augmented reality
AV	autonomous vehicle
CAGR	compound annual growth rate
CAPEX	capital expenditures
CSP	communication service provider
EB	exabyte
GPU	graphics processing unit
ICT	information and communication technology
IoT	internet of things
KPI	key performance indicator
LLM	large language model
N4AI	Networks for AI
UL/DL	uplink-to-downlink
VR	virtual reality
XR	extended reality
TCO	total cost of ownership

List of Figures

List of Tables

Table of contents

1 Introduction

This chapter introduces the master's thesis by giving a background, formulating the problem and purpose. This forms the base of the research questions. The chapter ends with delimitations.

1.1 Background

1.1.1 Information and communication technology

Today's information and communication technology (ICT) has its roots in over 150 years of technological innovation and with that numerous shifts in needs and capabilities. First with manual switching, then mechanical, then digital and with the introduction of the internet, packet switching. Ericsson led the launch of the first 4G network in 2009 setting the stage for broad adoption of high speed mobile internet. The latest shift was introduced with 5G providing services for internet of things (IoT) (Ericsson, n.d.) Today the total number of mobile subscriptions outpace the world population by ten percent and over 70 percent of the world's population use the internet (Ritchie et al., 2023).

1.1.2 Artificial intelligence

Artificial intelligence (AI) is increasingly becoming a part of everyday life. Weekly active users of large language models (LLMs) have surpassed one billion worldwide (Mishra, 2025; Bellan, 2025). Organizations reporting AI usage rose from 55 percent in 2023 to 78 percent in 2024. The American Federal Drug Administration approvals of AI-enabled medical devices rose from six in 2015 to 223 in 2023 and Waymo's autonomous taxis provide over 150 000 rides per week. Sentiment from the general public is rising as optimism around AI is on the rise in previously skeptical countries. Agentic work is also progressing and scoring significantly higher than human experts when time is limited. That is, when given around two hours, agents outperform humans on certain tasks. However humans still outperform when spending more than eight hours on the problem (Yolanda Gil, Raymond Perrault, 2025).

1.1.3 Cloud and edge computing

Cloud computing provides on-demand access to shared computing resources over the internet, enabling organizations to scale training and inference workloads without owning hardware (Google Cloud, n.d.(b)). Edge computing refers to deploying compute, storage, and networking resources closer to end users or data sources rather than relying solely on a centralized infrastructure. Placing workloads at the edge directly addresses latency and bandwidth constraints (IBM Corporation, 2026). Processing data locally reduces the time required for round-trip communication and limits the volume of data that must traverse the network. The distinction is architectural. Cloud computing centralizes resources in large data centers, while edge computing distributes them to increase traffic efficiency and decrease latency.

1.1.4 Training and running AI models

Two important concepts when talking about AI models is the difference between training and inferencing. Training is in simple terms the equivalent to building the model, giving an input and desired output and thereafter adjusting the models weights to minimize the error. This part is computationally and energy intensive. The second part is inferencing where the model is making predictions based on unseen inputs. That is, it gives predicted answers inferred from the training data, ergo inferencing (Cloudflare, n.d.)

Foundation models are large, general purpose AI models trained on vast and diverse datasets and designed to perform a wide range of tasks such as language processing, vision, and reasoning. Their scale and broad training make them capable across use cases but also computationally intensive to train and run. These are the types of models commonly found in AI assistants like ChatGPT, Gemini, and Claude (Google Cloud, n.d.(a)). On the other side there are specialized models which are built and optimized for narrow tasks. They use targeted datasets to maximize accuracy in specific applications such as medical imaging or fraud detection. This focus typically reduces compute requirements and improves inference speed relative to generalized foundation models (NVIDIA, n.d.) and deployment often differs accordingly. Specialized models are generally more feasible to run locally or on the edge because they are smaller and less resource intensive. This supports distributed deployments where processing occurs closer to the data source which reduces latency and bandwidth usage while increasing privacy. Foundation models often rely on centralized cloud infrastructure due to their high compute and data intensity. However Yolanda Gil, Raymond Perrault (2025) states that the inference cost for a system performing at the level of GPT-3.5 dropped over 280-fold between November 2022 and October 2024. This greatly increases the possibility for advanced models to run locally or at the network edge.

1.1.5 Implications for communication networks

As the possibility for AI models to run centralized, on the edge and locally the network will have to adapt to these new needs. How this will develop remains uncertain since this type of distributed computing is not only a question of connectivity but also compute location and use cases. This also requires carrying out scenario analyses to anticipate how the ecosystem could develop, allowing the players to prepare and position themselves for various outcomes.

1.1.6 AI for networks and networks for AI

There are two categories of discussion when it comes to AI-impact on networks:

- AI for networks refers to the use of AI to automate and manage network infrastructure. For example using models to predict traffic loads or detect faults.
- Networks for AI (N4AI) refers to the adaptation of network infrastructure to meet the demands of AI workloads, such as providing low latency, high bandwidth or local computing.

1.2 Problem description

Mobile network infrastructure is entering a period of uncertainty. The early, but rapid, introduction of AI agents, autonomous systems and physical AI creates a different class of traffic. New traffic is in part machine generated and increasingly decoupled from human interaction, making it hard to predict. CSPs and telecom equipment vendors have historically derived their position from building and owning the network while monetizing it through subscriptions. This model is not necessarily something that will work for machine-to-machine traffic.

Finding new monetization paths requires first understanding how the ecosystem might develop. Which players will gain influence and under what conditions depends heavily on how the ecosystem evolves over the coming decade. This opens up the need for scenario building to map how the N4AI ecosystem may develop through 2035.

1.3 Purpose

The purpose of this master's thesis is to analyze how the introduction of AI traffic will impact the ICT industry by 2035. It aims to provide a prognosis of network traffic, scenarios that may play out and how the ecosystem may change in each scenario. The aim is not to forecast one definitive future, but to outline a range of credible possibilities and determine which conditions give rise to each of them.

1.4 Research questions

To address the problem, two research questions are examined. The first research question serves as the foundation for evaluating the second research question. They are as follows:

- **RQ1:** How will AI usage impact mobile network traffic until 2035?
- **RQ2:** What will the Network for AI ecosystem look like in 2035?

1.5 Delimitations

This report takes a forward looking perspective to 2035, analyzing mobile network traffic and the ICT and AI ecosystem. This is done through a global lens where focus is directed toward trends and roles in the ecosystem, not specific companies or events. A large emphasis is placed on insights from a Delphi panel, which serves as the foundation of the study. Additional sources are used where relevant to complement and validate the panel input. This report focuses solely on Networks for AI. While adjacent AI application areas are relevant, they fall outside the definition of N4AI and are therefore not part of the thesis.

2 Methodology

This chapter presents the methodological basis of the research project. It describes the research strategy, the research design and the data collection method. The chapter also illustrates the process of data analysis and discusses the academic quality of the paper.

2.1 Approaching the problem

To ensure a systematic and rigorous investigation of AI traffic and its impact on the N4AI ecosystem, this study follows the practices of the methodological principles for small-scale social research as defined by Denscombe (2017). In line with this framework, Denscombe (2017) considers it appropriate to use a mainly qualitative approach, aided by quantitative synthesis, to address the research questions. Given the shifting landscape of the ICT and AI sectors, the project as a whole is of exploratory nature.

The reasoning follows an abductive logic, which according to Hoek et al. (2005) is characterized by a continuous movement between the theoretical world and empirical world. Unlike a purely deductive approach that tests existing hypotheses, or an inductive approach that ignores prior theory, abduction allows this study to utilize existing frameworks, such as N4AI ecosystems, as a lens to interpret the expert insights gathered through the Delphi rounds (Hoek et al., 2005).

2.2 Research design

According to Höst et al. (2006), the most occurring designs of research methods are experiments, case studies, surveys, action research or mixed methods. Denscombe (2017) states that the use of mixed methods is a way to broaden the perspective on a research topic, as opposed to using a single method. A mixed method approach becomes useful since it increases the credibility of the results since the method is in line with the triangulation framework, where various methods of research are combined to complement each other (Jick, 1979). Furthermore, the analysis of the findings can get developed sequentially, as one method provides background information for the other. By using "combination research", the analysis becomes more complex, that is, becomes more time consuming. Moreover, as noted by Denscombe (2017), the scope of the research is adjusted so that the thesis can be completed within the available time frame.

The thesis employed a mixed-method research design with two components: industry reports and a survey, which were used in combination. This design was chosen to address the inherent complexity of the 10-year strategic horizon. Although the industry reports provide the necessary depth to understand the current organizational and technical baseline, they alone would not suffice to capture the broad, ecosystem-wide shifts required for an exploratory scenario analysis of how the future N4AI might look.

To gain a comprehensive understanding through the literature study, the PESTEL framework was used as it gained insights to Ericsson's current system and strategic posture. In this way, the literature study allowed for the construction of a robust Delphi survey which was essential for assessing the high levels of uncertainty in the future N4AI. This sequential approach was further strengthened by access to primary data from internal technology experts at Ericsson and external industry stakeholders, ensuring the naturally occurring context was fully captured (Denscombe, 2017). Moreover, the goal is to map the range of futures that could emerge given different combinations of driving forces.

2.2.1 The Delphi method

A Delphi study is an approach attempting to gather insights from experts in different, yet relevant areas on the same topic. It is an interactive approach involving two or more rounds where the knowledge and opinions of the experts independently are compiled to reach a consensus (Denscombe, 2017). The idea of the method is that participants answer to a round of surveys and subsequently, they receive the results of all the experts' responses in an iterative process. The aim is to converge towards a consensus among the experts. According to (Diamond et al., 2014), a consensus in a Delphi study is described in terms of a percentage agreement, where $\geq 75\%$ is valued as a strong agreement point on the issue in question. A vital part of the Delphi study is the encouragement of anonymity. In this way, aspects of group interaction such as dominance and pressure are drastically reduced (Skulmoski et al., 2007). Overall, the Delphi method is a suitable approach when exploring uncertain research areas, forecasting the future or gathering consensus about a matter that is widely spread (Okoli and Pawlowski, 2004).

2.2.1.1 Process

According to Denscombe (2017), the Delphi method is divided into seven steps, each followed accordingly in this research project:

1. Defining the problem
2. Tailoring the research design to the resources available
3. Establishing contacts with suitable experts
4. First round of questions
5. Initial analysis and feedback
6. Subsequent rounds (in this case one more round) of questions and feedback
7. Final report accessible to all experts

According to Effah et al. (2016), Delphi studies typically range from three to five rounds, but two rounds are generally considered appropriate for graduate research given practical limitations (Skulmoski et al., 2007). Two rounds can be sufficient when a solid knowledge base already exists in the research area in question (Iqbal and Pipon-Young, 2009). This project therefore conducted a two-round approach where the first round was grounded in literature and the second being synthesized group results. The participants were thereafter asked whether they agreed on the plausible scenarios built from round one - inviting the experts to explain their reasoning (Hsu and Sandford, 2007).

2.2.2 PESTEL Analysis

The PESTEL framework provides a structured way to categorize the external factors that shape an industry or ecosystem (Johnson et al., 2017). The acronym covers six dimensions: Political, Economic, Social, Technological, Environmental and Legal. In this study, PESTEL is not used as a standalone analytical tool but as an organizing lens for structuring the environmental analysis of factors influencing the N4AI ecosystem through 2035.

Specifically, the framework informed the design of the Delphi survey by ensuring that the question bank covered a broad range of external forces rather than narrowing too early to purely technical or commercial factors. It also provides a structure for presenting findings on which external conditions most significantly shape the plausible futures of the ecosystem (Yuksel, 2012).

PESTEL is a descriptive rather than predictive tool. It does not rank factors by importance or specify how they interact, which is why it is paired here with scenario building methodology to move from environmental mapping to structured future projections (Johnson et al., 2017).

2.3 Data collection

Data collection was conducted through two sequential online survey rounds, following the Delphi structure outlined above. The first survey questionnaire was designed using insights from academic literature and industry reports.

2.3.1 Delphi survey rounds

Expert participants were recruited through purposive sampling, which according to (Patton, 2002) aims to target individuals with domain-relevant expertise, in this case across academia, CSPs, technology vendors, hyperscalers and regulatory or public sector bodies. This was deliberately done to span stakeholder groups in order to capture a broad ecosystem perspective (Skulmoski et al., 2007). Both surveys were structured around four thematic areas: AI-driven traffic and emerging technologies, the N4AI ecosystem and factors influencing adoption and development such as compute architecture preferences. The second round was built using the answers from the first one and aimed to build plausible future scenarios for the N4AI ecosystem in a

2035 horizon. A full description of the scenarios and the construction methodology is provided in section 5.2.

Participants were solely contacted via email and informed that participation would take no more than 15-20 minutes each and that the answers would be anonymized and finally, that synthesized results would be shared upon completion (Hsu and Sandford, 2007). Round one gathered structured responses (and optionally open comments on each question) to gain a deeper understanding of the landscape of the expert opinions.

2.3.1.1 Scoring the answers in the Delphi Survey

Answers from the Delphi survey were coded numerically to enable comparison across questions. All three coding schemes follow the same weighted-sum formula:

$$S = \sum_i w_i \cdot n_i$$

where w_i is the weight assigned to response category i and n_i is the number of panelists who selected that category.

Ordinal scale (low to high): for questions with an ordered scale from low to high (e.g. *none, small, medium, high, massive*) or slow to fast (e.g. *very slow, slow, moderate, fast, very fast*), the lowest category receives weight $w = 0$ and each step upward adds one, reaching $w = 4$ at the highest category. For example, if two panelists answered *slow* ($w = 1$) and two answered *moderate* ($w = 2$), the score is:

$$S = (1 \times 2) + (2 \times 2) = 6$$

Bipolar scale (directional): for questions structured around two opposing poles, for example *centralized, slightly centralized, no impact, slightly distributed, distributed*, weights run from -2 at one extreme to $+2$ at the other, with zero as the neutral midpoint. A positive score indicates net movement toward one pole and vice versa. For example, if two panelists answered *centralized* ($w = -2$) and one answered *no impact* ($w = 0$) the score is:

$$S = ((-2) \times 2) + (0 \times 1) = -4$$

Winner-loser coding: for questions asking panelists to classify ecosystem actors as winners or losers, winners were coded $w = +1$ and losers $w = -1$. A positive sum indicates a relative winner and vice versa. For example, if three panelists voted for an actor a winner and two voted it a loser:

$$S = (1 \times 3) + ((-1) \times 2) = 1$$

Since $S > 0$, the actor is classified as a relative winner.

2.3.1.2 Thematic analysis

The open ended questions in round two were analyzed using thematic analysis. Responses were read and coded into recurring themes, no pre-defined category list was imposed. For each question, representative quotes were selected to anchor the themes in the panel's own language. A mention count was assigned to indicate how many respondents raised each theme, each respondent could mention multiple themes. These counts are a measure of importance within the panel, reflecting the relative weight of each theme for the specific scenario. Mentions can not be generalizable between scenarios since the number of respondents per question differ as indicated in each table description.

2.3.2 Secondary data

Secondary data primarily served two purposes: first, it supported the design of the survey, and second, it helped to contextualize the findings during the analysis (Saunders et al., 2019). Industry reports helped shape the understanding of AI networking trends, mobile traffic segments and the overall AI and telecom landscape. The aid from secondary sources provided a richer interpretation of expert responses as well as ensuring that the survey reflected both established theory and current industry discourse.

2.4 Work process

Throughout this process, the authors worked in close collaboration at every stage, with no major separation between their respective contributions. The construction of the first round of the survey was developed iteratively over several weeks. The starting point was the literature review, which was supplemented by discussions with both the thesis supervisor and internal Ericsson contacts. With this help, it was ensured that the questions reflected both academic grounding and industry relevance. From there, the authors went back and forth on question wording, structure and coverage before arriving at a first complete draft.

Before the draft was shared with the LTH supervisor, it was pilot tested with a small group of people outside the expert panel. The reasoning behind this intermediate step was to identify any issues and flag anything that felt unclear, ambiguous or poorly worded. Completion time was also recorded to get a rough estimate of respondent burden, with the consideration that experts engaging deeply with the subject matter may take longer than the pilot group.

Based on pilot feedback, a number of questions were revised for clarity before the survey was sent to the supervisor for a final review. The resulting questionnaire was then distributed to the expert panel for round one. Round two was designed following the synthesis of round one responses, with questions about scenarios built using the responses from round one.

Throughout the process, AI tools were used as a practical aid at various stages. This included refining and polishing written text for clarity, as well as automating aspects of the workflow such as formatting. The Google suite of apps was used for data collection and management.

2.5 Trustworthiness

Trustworthiness was assessed using the four criteria from Lincoln and Guba (1985), which was operationalized by Krefting (1991): credibility, transferability, dependability and confirmability (Shenton, 2004).

Out of the four criteria, credibility was the most actively managed. The composition of the panel played a vital part. That is, academics, CSPs, vendors, hyperscalers and regulators should be equally weighted in the expert panel. Furthermore, the thesis used triangulation in which the survey was cross-referenced against industry reports (Denscombe, 2017).

The transferability is moderate. The two research questions addressed in the thesis concern the broader telecom ecosystem, that is, they are not tied to a single company. Moreover, the findings are relevant to telecom vendors and CSPs beyond Ericsson. The collaboration with Ericsson gave context, but does not constrain the findings in such a way that limits applicability to the wider industry.

3 Theory

This chapter establishes the theoretical foundation of the study. It introduces two interconnected bodies of theory: ecosystem and network theory, which provides a lens for mapping actors and their interdependencies; and foresight and scenario building, which structures how plausible futures are projected. The chapter closes by showing how the two theoretical blocks form an analytical sequence that mirrors the logic of the study itself.

3.1 Ecosystem & network theory

In the N4AI ecosystem, there is a wide range of actors whose contributions are interdependent, that is, no single company delivers the full stack alone. In order to get an understanding of how all actors relate to each other, where power sits and what holds the ecosystem together, a theoretical lens is required that goes beyond simple supplier-customer relationships. This section introduces the frameworks used to map and analyze that structure.

3.1.1 Ecosystem mapping: the ecosystem pie model

The Ecosystem Pie Model, developed by Talmar et al. (2020), provides a structured way to visualize which actors are involved in delivering a value proposition, what each contributes and how value is captured. The model arranges actors as slices around a central value proposition, with each slice describing an actor's resources, activities, value addition and revenue logic.

For this study, the model is used as an analytical lens to map the N4AI ecosystem, in which it identifies the relevant actors, their roles and where the boundaries between them sit. It also helps surface where interdependencies are tight and where there is room for different actors to capture value.

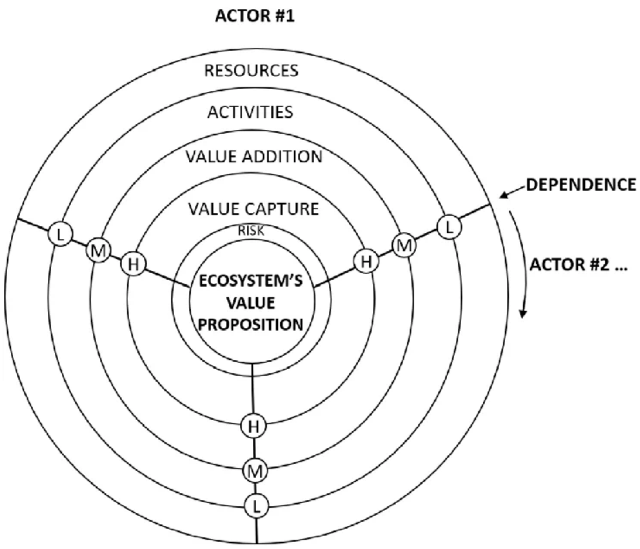

Figure 3.1 The Ecosystem Pie Model, adapted from Talmar et al. (2020)

3.1.2 Network ties and actor interdependencies

Beyond mapping which actor is in the ecosystem, it is equally vital to study how the actors are connected. Ahuja et al. (2012) describe organizational network ties as four categories: hierarchical ties, which reflect authority relationships; referential ties, which represent certification or endorsement; market ties, which are either competitive or transactional; and affective ties, which reflect partnerships and mutual commitment.

In the N4AI ecosystem, these distinctions are useful. A few examples follow below: a CSP buying network equipment from Ericsson is a transactional market tie. A joint development agreement between a hyperscaler and a telecom equipment vendor is an affective tie. A regulatory body sits closer to a referential relationship, it does not own anyone but shapes what everyone builds to. Recognizing these differences matters when analyzing who can capture value and under what conditions.

Network ties are also dynamic. Chen et al. (2022) identify three main drivers of network change: contextual factors such as regulatory shifts or technology shocks, actor-level decisions such as choosing to enter or exit a partnership and relational factors where existing ties shape which new ones form. All three are relevant in this case as the emergence of AI, as a dominant workload, is reshaping who partners with whom across the telecom and cloud industries.

3.2 Foresight & scenario building

To answer RQ2, a snapshot of today's ecosystem is required. This section covers the theoretical foundations for making structured projections about the future, including how expert knowledge is aggregated, how diffusion patterns shape technology adoption and how scenarios are constructed to explore possible futures.

3.2.1 Foresight studies

In contrast to forecasting, in which extrapolation from existing trends assume the future, foresight studies explores uncertainties, driving forces and alternative outcomes in contexts where the future is open (Bishop et al., 2007). When studying ecosystems, specifically telecom ecosystems, they do not follow from a simple extrapolation of today's infrastructure trends. The structure depends on choices, investments, partnerships and competitive dynamics that have not yet played out.

Foresight studies typically combine qualitative and quantitative methods, including expert panels, workshops, Delphi studies and scenario building (Popper, 2008). They tend to have time horizons of ten to twenty years, which aligns well with the 2035 horizon of this study. The goal is not to predict a single correct future but to map the space of plausible ones and identify what conditions lead to each.

3.2.2 The Delphi method as a foresight tool

One of the most common approaches in foresight research is in fact the Delphi method, which was originally developed as a forecasting tool (Rowe and Wright, 1999). It is a valuable tool in a foresight context since it aggregates the judgment of experts in a structured way. This is because the method reduces distortions from group dynamics, as mentioned before, as well as allowing for iterative refinement of views across rounds (Hsu and Sandford, 2007).

In foresights applications, the Delphi method is particularly useful in this case, as the data is limited and that perspectives need to be captured across different parts of a system (Okoli and Pawlowski, 2004). Furthermore, this method does not eliminate uncertainty, but it produces a strcutured picture of where experts agree, where they disagree and why. Thus, the method itself is informative for scenario construction.

3.2.3 Diffusion of innovations

Understanding how technologies spread through markets is relevant to any foresight study involving emerging technology. Rogers (2003) introduced the technology adoption lifecycle, describing how innovations diffuse through five successive adopter categories: innovators, early adopters, early majority, late majority and laggards. The rate of diffusion depends on five characteristics of the innovation itself: relative advantage over existing alternatives, compatibility with existing systems and values, complexity of use, ease of experimentation and observability of results.

In the perspective of this study, Roger's framework helps think through which parts of the ecosystem are likely to see rapid adoption and which face structural barriers. Edge inference infrastructure, for example, scores well on relative advantage but poorly on compatibility with existing telecom network architectures, which suggests a slower diffusion path than raw technical performance would imply.

Moore (2014) builds on Roger's model by identifying a chasm between early adopters and the early majority that many technologies fail to cross. Early adopters tolerate incomplete solutions and are motivated by the novelty of technology. On the other hand, the early majority wants proof of work, they want reliable work that fits into existing workflows. Crossing that gap requires a different value proposition than the one that attracted early adopters. For technologies in the ecosystem, this distinction between early adopter CSPs willing to experiment with AI networking and the broader market expecting proven, integrated solutions is directly relevant.

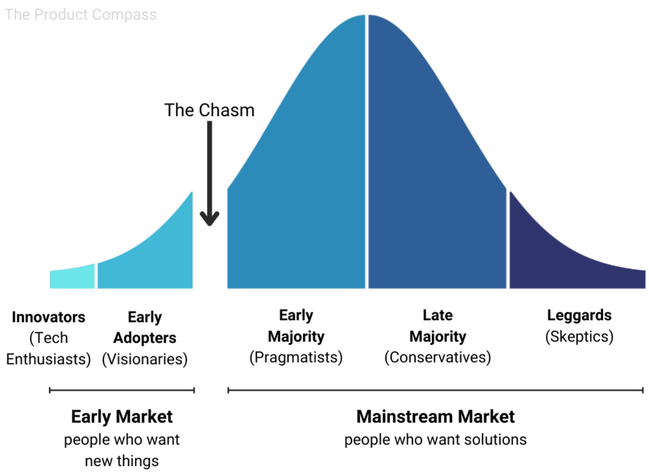

Figure 3.2 The technology adoption lifecycle and the chasm, adapted from Moore (2014)

3.2.4 The theory of scenario building

Scenario building is a structured method for constructing a small set of plausible future states that differ in meaningful ways (Schwartz, 1996). The goal is not to identify the most likely future but to map the range of futures that could plausibly emerge given different combinations of driving forces. Scenarios give decision-makers a way to stress-test strategies against uncertainty rather than optimizing for a single expected outcome.

In practice, the driving forces used to construct scenarios are identified through a qualitative reading of collected expert input and secondary data. Responses are grouped thematically, with recurring topics and tensions surfacing as candidate dimensions of uncertainty. Driving forces are then filtered based on two criteria: the degree of uncertainty surrounding their future development, and their independence from one another. Forces that are both highly uncertain and mutually orthogonal are selected as scenario axes.

The scenario cross is one of the most common approaches to scenario construction (Lindgren and Bandhold, 2009). The workflow is divided into two parts. First, identification of two driving forces that are highly uncertain and largely independent of each other. Then, using them as axes to generate four distinct scenarios completes the approach. The approach forces internal consistency within each scenario and produces outcomes that are genuinely different from each other rather than just variations on the single theme.

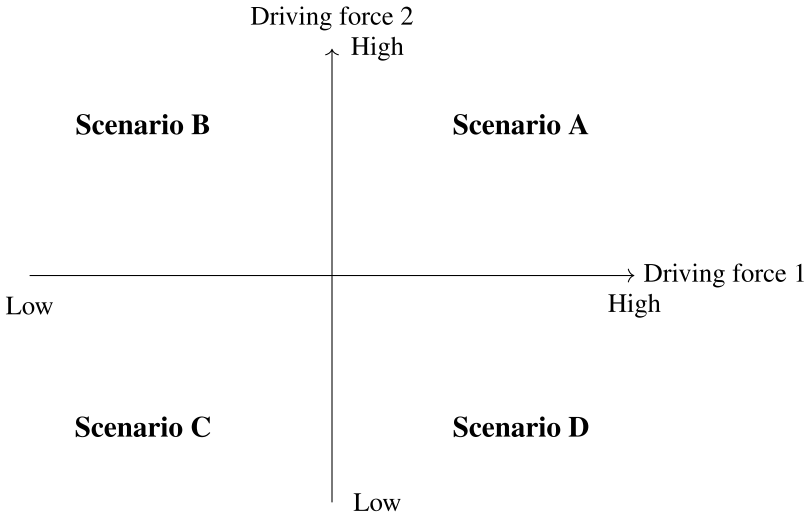

Figure 3.3 The scenario cross, adapted from Lindgren and Bandhold (2009)

3.3 Synthesized theoretical framework

The two theoretical blocks presented in this chapter have not been selected independently. They form a sequence that mirrors the analytical logic of the study itself. Understanding the N4AI ecosystem 2035 requires first knowing who the actors are and their dependencies, and then projecting how the actor landscape might evolve. Ecosystem and network theory, and foresight and scenario building each handle one of these questions. Figure 3.4 shows how they connect.

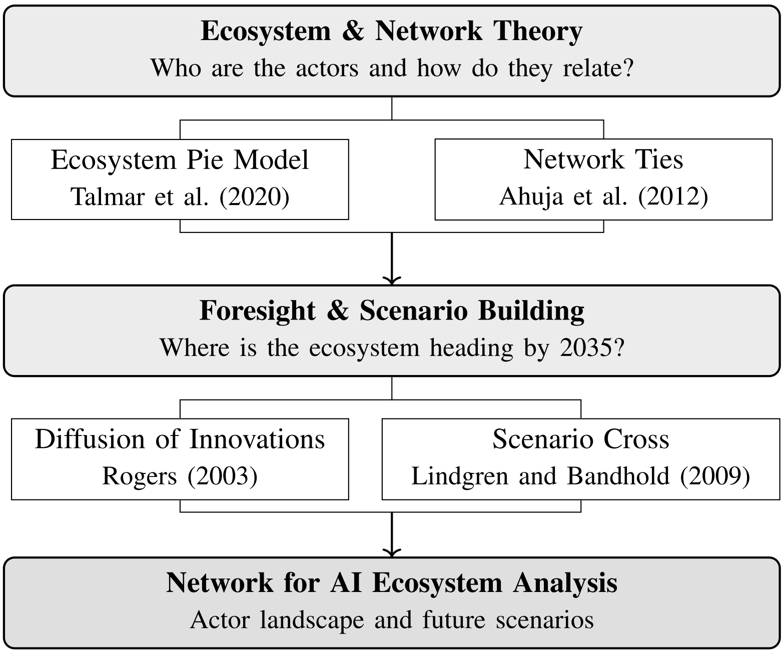

Figure 3.4 Synthesized theoretical framework

In section 3.1, the structural foundation is established. That section provides the tool to map who participates in the N4AI ecosystem, what each actor contributes and how they depend on each other. The Ecosystem Pie Model gives a visual and analytical structure for this mapping, while network tie theory adds nuance by distinguishing between different kinds of relationships: transactional, hierarchical, competitive and partnership-based.

The second layer, section 3.2, introduces time. Foresight and scenario building acknowledges that the actor landscape is not static. Diffusion theory helps explain which technologies and business models are likely to spread quickly and which face structural adoption barriers. The scenario cross then translates uncertainty about key driving forces into a structured set of plausible futures, giving the analysis a 2035 horizon rather than just a snapshot of today.

Together, the two layers form an analytical sequence that moves from structure to trajectory. Each layer depends on the one before it as it is not possible to build credible scenarios without first knowing who the actors are and how they relate. This sequence is not arbitrary, it reflects the order in which the research questions themselves are addressed.

4 Setting the stage

This chapter discuss the current and future telecom and AI industry, the change in traffic, technologies that will drive the change, the ecosystem, and future scenarios. This lays the groundwork for the Delphi survey questions and future scenarios.

4.1 The state of mobile connectivity

4.1.1 Coverage

For the past 10 years, 95 percent of the world population have had coverage through some 3GPP cellular technology. In 2016, 55 percent had LTE coverage, a figure that has now increased to 90 percent. 5G was introduced at scale in 2020 and today coverage sits at 60 percent (Ericsson, 2025h). Figure 4.1 shows this progression.

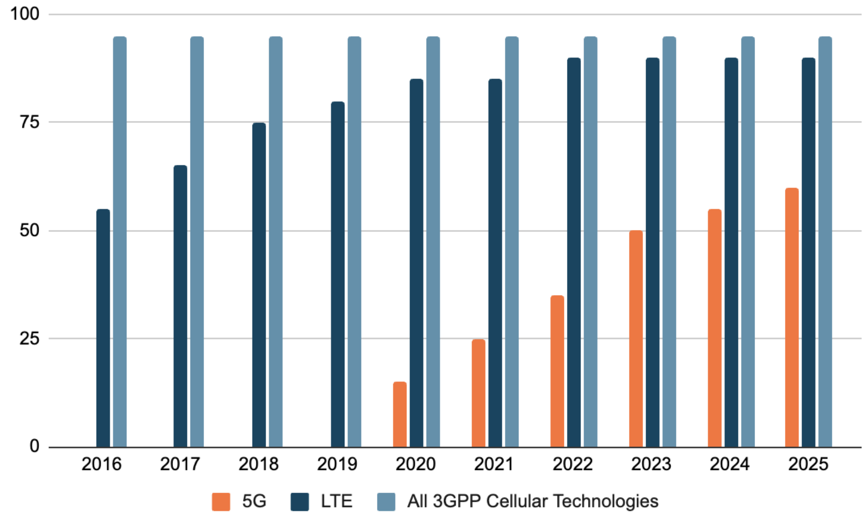

Figure 4.1 Population coverage for mobile connectivity over time in percent (Ericsson, 2025h).

4.1.2 Mobile network traffic

In 2025 the global mobile network traffic reached 197 exabyte (EB)/month split into three main categories; 2G/3G/4G, 5G, and fixed wireless access. 2G/3G/4G stands for the largest share at 81 EB/month, then 5G at 62 EB/month, at a close third fixed wireless access at 54 EB/month (Ericsson, 2025g). Total traffic has grown significantly and in just the past 10 years mobile network traffic has increased over 35 fold. From 5.7 EB/month in Q1 2016 to just above 200 EB/month in Q4 2025 (Ericsson, 2025a). 90 percent of this traffic is downlink, which is when a mobile device downloads data, and 10 percent is uplink where the device uploads data (Ericsson, 2025f).

4.1.3 Connected devices

Today, approximately 33.4 billion devices are connected globally, spanning fixed and mobile phones, laptops and IoT devices, see Figure 4.2. The breakdown is as follows: 0.8 billion fixed phones, 8.6 billion mobile phones, 1.7 billion laptops, PCs, and tablets, 17.5 billion short-range IoT devices and 4.9 billion Wide-Area IoT devices (Ericsson, 2025b). Of the 22.4 billion IoT devices, 4.5 billion use cellular connections (Ericsson, 2025c).

Devices connect either via fixed or mobile networks, this report focuses on the latter. Mobile subscriptions currently total 8.8 billion, with a regional breakdown provided in Figure 4.3.

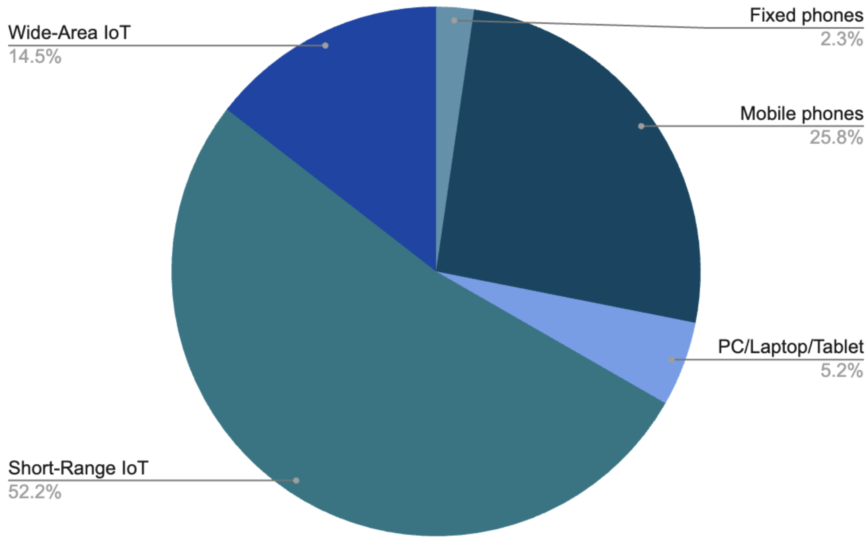

Figure 4.2 Percentage share of connected devices (Ericsson, 2025b).

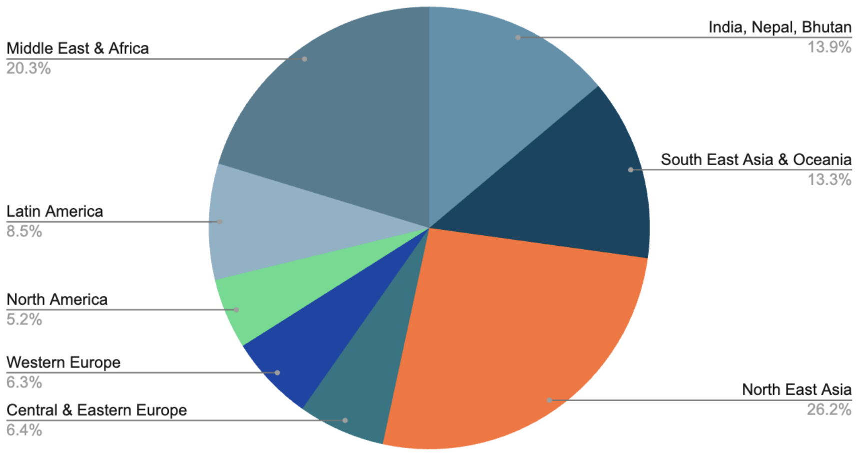

Figure 4.3 Mobile subscriptions by region (Ericsson, 2025d).

4.1.4 Mobile data traffic by application category

The type of data flowing through devices today is quite homogeneous with video making up 76 percent of total mobile data traffic, see figure 4.4. Video is followed by social networking at 8 percent, software updates at about 3 percent, and audio and web browsing both at about 1 percent (Ericsson, 2025e).

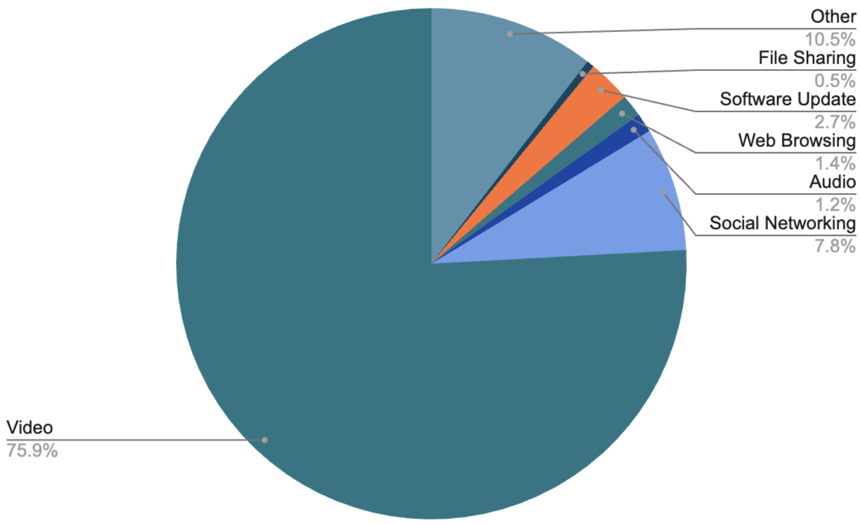

Figure 4.4 Mobile data traffic by application category (Ericsson, 2025e).

4.2 Traffic prognosis toward 2035

Ericsson (2026) states that:

> *"The consolidation of existing and emerging applications is predicted to drive traffic growth: downlink at 15 percent compound annual growth rate (CAGR) and UL* [uplink] *at a significant 30 percent CAGR."*

They also state that the uplink share in June 2025 was about 10 percent (Ericsson, 2025f). This enables extrapolating these numbers to 2035 yielding an estimated uplink of 38 percent. Ericsson projects global mobile network data traffic to grow from 196 EB/month in 2025 to 483 EB/month by 2031 (Ericsson, 2025g), implying an overall compound annual growth rate (CAGR) of 16.6 percent. Extrapolating to 2035 using this rate gives $196 \cdot 1.162^{10} \approx 882$ EB/month. This gives a very rough estimate of global mobile data traffic data of 882 EB/month by 2035, with uplink accounting for 38 percent.

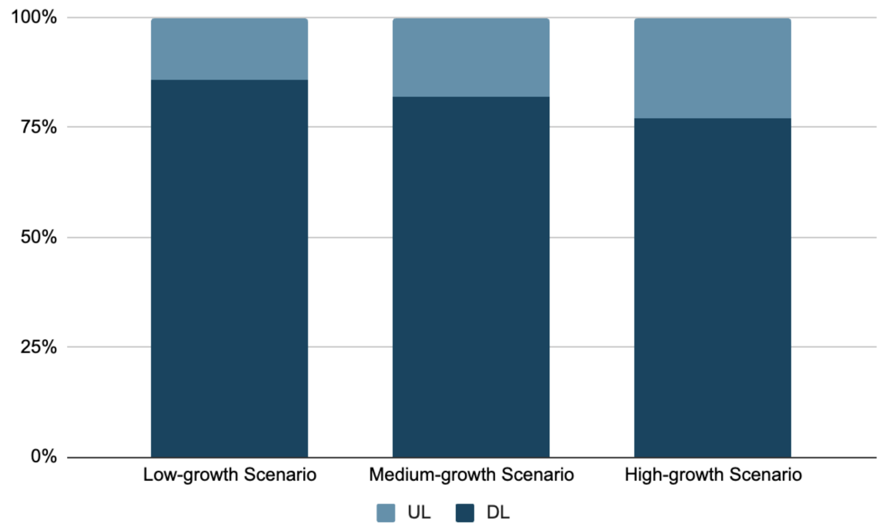

Figure 4.5 Traffic distribution by downlink and uplink for three traffic growth scenarios (Kalvin Bahia, Facundo Rattel, 2025).

Kalvin Bahia, Facundo Rattel (2025) expect the traffic to be 1000, 1250, and 1500 EB/month by 2035 in their low, medium, and high growth scenarios specified in figure 4.5 and 4.6. Much higher than the extrapolated data from Ericsson. They predict that the primary traffic drivers are 6G enhanced use cases like extended reality, autonomous vehicles, and holographic communication, see figure 4.6. 6G use cases are expected to be introduced around 2030.

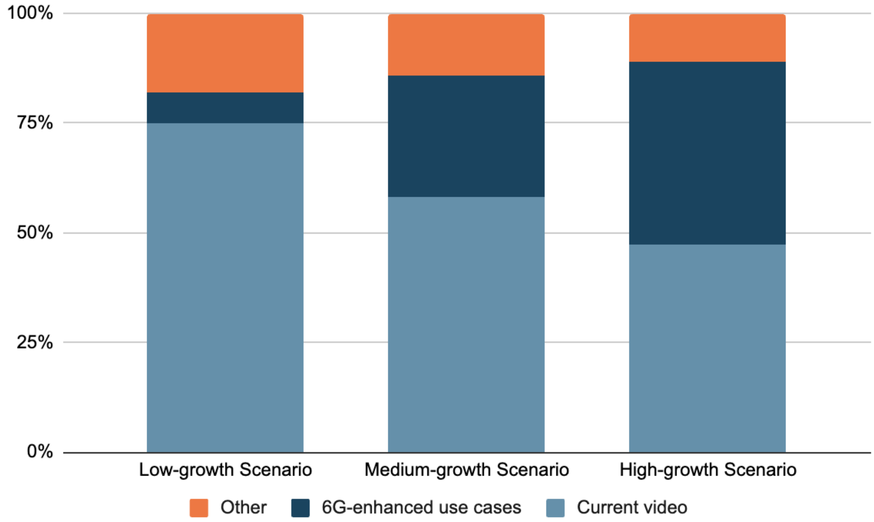

Figure 4.6 Traffic distribution by use case for three traffic growth scenarios (Kalvin Bahia, Facundo Rattel, 2025).

During Mobile World Congress 2026 the Nokia CEO, Justin Hotard described how the traffic will change from a linear easily predictable traffic pattern to a highly variable one. The section following his quote below is about trying to identify and understand the devices Hotard is referring to.

> *"It's about understanding the devices that are coming and delivering deterministic connectivity because what will matter to a drone, to a robot, to an autonomous vehicle is that it transmits the information it needs and it gets the intelligence back to take the action it needs at that particular moment of time. And that means a massive shift. It's a shift from an SLA or policy driven architecture to one of deterministic connectivity."* (Nokia, 2026)

4.3 AI & connectivity enabled technologies driving the change in network traffic

To analyze how the telecom industry will be impacted by AI it is required to understand which technologies are enabled by AI and how they will drive new traffic. Global X, Ericsson and McKinsey lists many technologies that are on the rise, see table 4.1 (Global X Management Company LLC, 2023; Global X Management Company LLC, 2024; Global X Management Company LLC, 2025; Global X Management Company LLC, 2026; Michael Chui, Mark Collins, Mark Patel, 2021; Ericsson, 2026). They categorized the technologies according to Rogers (2003) technology adoption lifecycle explained in section 3.2.3. To find those relevant to this study all technologies mentioned were filtered based on whether they are enabled by AI and connectivity. With this filter the technologies in table 4.1 were identified. Ericsson also discuss AI Agents but no adoption phase is mentioned (Ericsson, 2026). As seen in the table there is overlap between both where they are in their development and between technologies. As seen AI and healthcare are in multiple stages. Some were therefore grouped and some broken up as visualized in table 4.2, this was done by the authors in collaboration with experts at Ericsson.

Table 4.1 AI and connectivity enabled technologies.

Innovators	Early Adopters	Early Majority
Humanoids	Wearable medical devices	AI
Robotaxis	AI	Cloud gaming
Autonomous vehicles	IoT	IoT
GenAI	VR	
AR/VR	Robotics	
Healthcare wearable tech		
Neurology IoT Devices		
Cardiology IoT Devices		

Table 4.2 Technology Categories defined by the authors and experts at Ericsson.

Technology	Referenced Categories
AI Assistants	GenAI, AI, AI Agents
Software Based Agentic AI	GenAI, AI, AI Agents
Internet of Things	IoT
Smart Cities	IoT, City IoT
Healthcare AI	Healthcare wearable tech, Neurology/Cardiology IoT Devices
Autonomous Mobility	Robotaxis, Autonomous vehicles
Autonomous Robots	Humanoids, Robotics, Droids
Extended Reality	VR, AR, XR
Cloud Gaming	Cloud gaming

4.3.1 AI assistants

Charlotte Hu (2026) defines an AI assistant as follows:

> *"An AI assistant is an intelligent application that understands natural language commands and uses a conversational AI interface to complete tasks for a user. Many modern virtual assistants, such as Amazon's Alexa and Apple's Siri, rely on these capabilities to enhance user interactions."*

Today, Generative AI has 26 percent uplink traffic on average but it varies from app to app. ChatGPT, which accounts for 60 percent of AI traffic as of June 2025, had 29 percent uplink while DeepSeek and Microsoft Copilot had roughly 50 percent (Ericsson, 2025f). Costs of both training and running AI-models are drastically decreasing. The cost of training Chat GPT-4 has dropped from 45 million USD to only 2 million USD between 2023 and 2026, by 2030 it is expected to be below 0.1 million USD, see figure 4.7. Similarly, inferencing costs dropped by 90 percent in just 18 months from March 2023 to August 2024. At the same time the AI-assistant space is experiencing explosive growth, ChatGPT weekly active users grew from one to 800 million between November 2022 and September 2025, see figure 4.8 (Global X Management Company LLC, 2026).

Example: A user asks a smartphone assistant to summarize an email chain. The system parses the text and provides a summary.

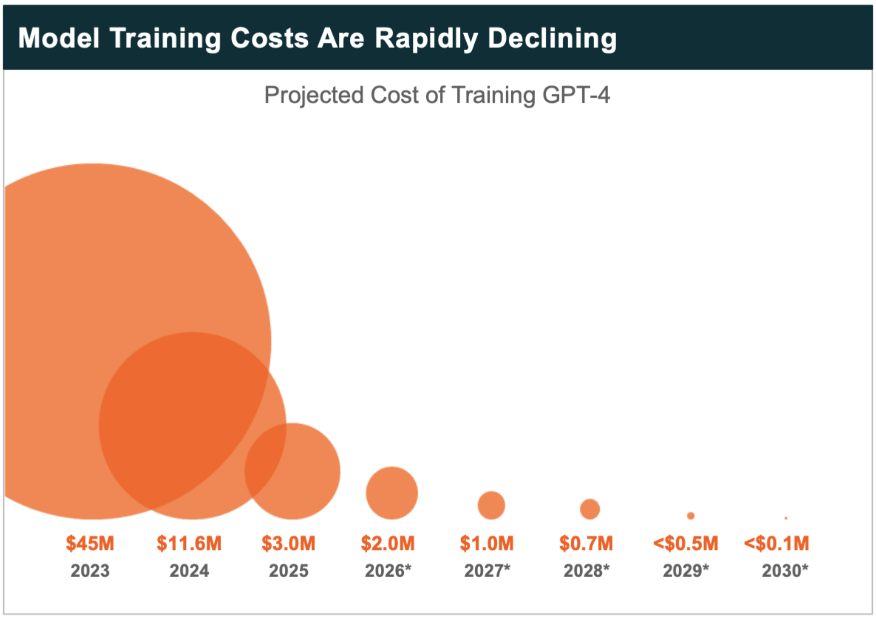

Figure 4.7 Cost of training and running AI-models over time (Global X Management Company LLC, 2026).

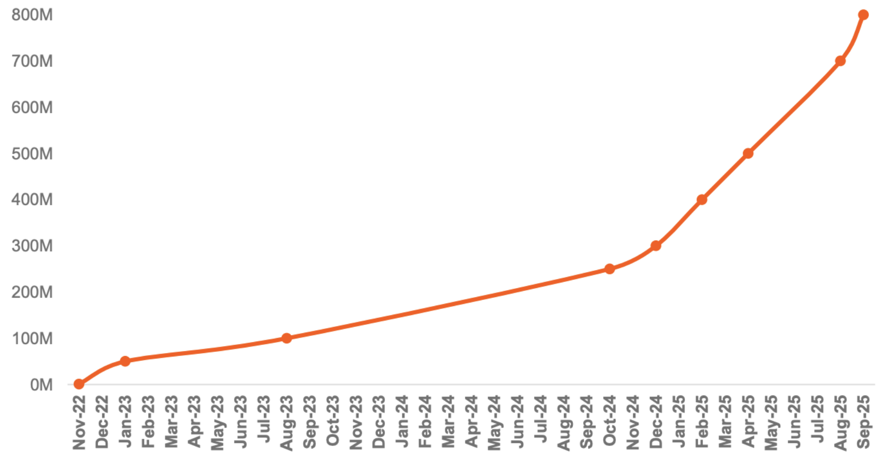

Figure 4.8 ChatGPT weekly active users (Global X Management Company LLC, 2026).

4.3.2 Software based Agentic AI

Stryker (2026a) defines agentic AI as follows:

> *“Agentic AI is an artificial intelligence system that can accomplish a specific goal with limited supervision. It consists of AI agents—machine learning models that mimic human decision-making to solve problems in real time. In a multiagent system, each agent performs a specific sub-task required to reach the goal and their efforts are coordinated through AI orchestration.”*

Global X Management Company LLC (2026) states that AI agents is set to reshape the 750 billion USD Cloud Computing industry due to an explosion in global data generation and app creation. Agentic AI systems are more prone to create sustained uplink pressure resulting in a more symmetric uplink/downlink ratio (Stephens, 2026).

Example: An autonomous coding agent identifies a bug. It writes a fix, runs tests, and submits a pull request without intervention.

4.3.3 Internet of Things

Michael Chui, Mark Collins, Mark Patel (2021) defines IoT as follows:

> *"The Internet of Things (IoT) refers to physical objects embedded with sensors that communicate with computers. The IoT enables the physical world to be digitally monitored or controlled."*

Ericsson (2026) anticipates that on-device AI will drive new applications resulting in a resurgence of IoT. Typically a single IoT device has a low data rate, uplink centric traffic making each one have a small impact on total traffic but with the possibility of having in the order of millions per square kilometer they can have a large total contribution (Gieske, 2024). Figure 4.9 shows how machine and IoT data is emerging as a large data category (Global X Management Company LLC, 2024). Michael Chui, Mark Collins, Mark Patel (2021) list the largest enablers in order from largest to smallest as follows: Factories, human health, worksites, city, retail environments, outside, home, vehicles, and office. The authors exclude human health, city, and vehicles to include them in other categories described below. Given the exclusion they estimate that IoT could enable 3.5 to 8.5 trillion dollars in value globally by 2030.

Example: A smart thermostat monitors home temperature. It adjusts heating based on occupancy patterns and external weather data to save energy.

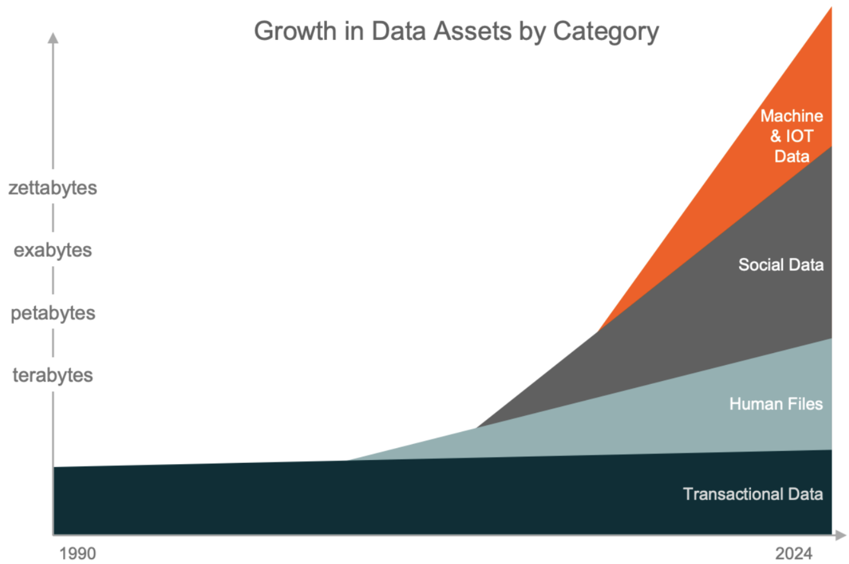

Figure 4.9 Data growth by asset category (Global X Management Company LLC, 2024).

4.3.4 Smart cities

Alice Gomstyn (2026) defines smart cities as follows:

> *"A smart city is an urban area where technology and data collection help improve quality of life as well as the sustainability and efficiency of city operations. Smart city technologies used by local governments include information and communication technologies (ICT) and the Internet of Things (IoT)."*

Michael Chui, Mark Collins, Mark Patel (2021) notes that by 2030 IoT in cities could reach 1-1.7 trillion USD. Since it is a subcategory of IoT, connectivity demands are similar.

Example: City sensors monitor traffic density. The system adjusts traffic light timing in real time to reduce congestion and vehicle emissions. At the same time, thousands of additional sensors transmit small data packets every hour to report electricity consumption for a metropolitan utility provider.

4.3.5 Healthcare AI

There are many use cases of AI in healthcare but the primary one discussed in this report is connected healthcare services or devices where analytics is done using AI. In the three latest reports from Global X healthcare AI is discussed (Global X Management Company LLC, 2024; Global X Management Company LLC, 2025; Global X Management Company LLC, 2026). In their 2026 edition the following two definitions are given:

> *"Smart Medical Devices: AI-powered medical hardware including wearables and surgical robotics that enhance patient outcomes."*
>
> *"Tech-Enabled Consumer Care: Solutions that improve access to care through virtual services, remote monitoring, and online pharmacies."*

They list smart medical devices at a 12 percent compound annual growth rate (CAGR) and tech enabled consumer care at 18 percent CAGR. This category is in large part driven by an aging world population where the 65+ age group is growing at three times the rate of younger cohorts. Total healthcare spending increase by three times between ages 45-64 to 85+, see figure 4.10. AI enabled device approvals in the United States have grown from 150 in 2022 to 350 in 2025 and wearable sensor revenue is set to double in the 2023 to 2032 time frame reaching almost 50 billion USD (Global X Management Company LLC, 2026). Furthermore Michael Chui, Mark Collins, Mark Patel (2021) prognosticates that the healthcare IoT category could unlock 0.5 to 1.8 trillion USD in economic value by 2030.

Example: A wearable patch tracks a patient's heart rhythm. It alerts a cardiologist if it detects signs of an irregular heartbeat.

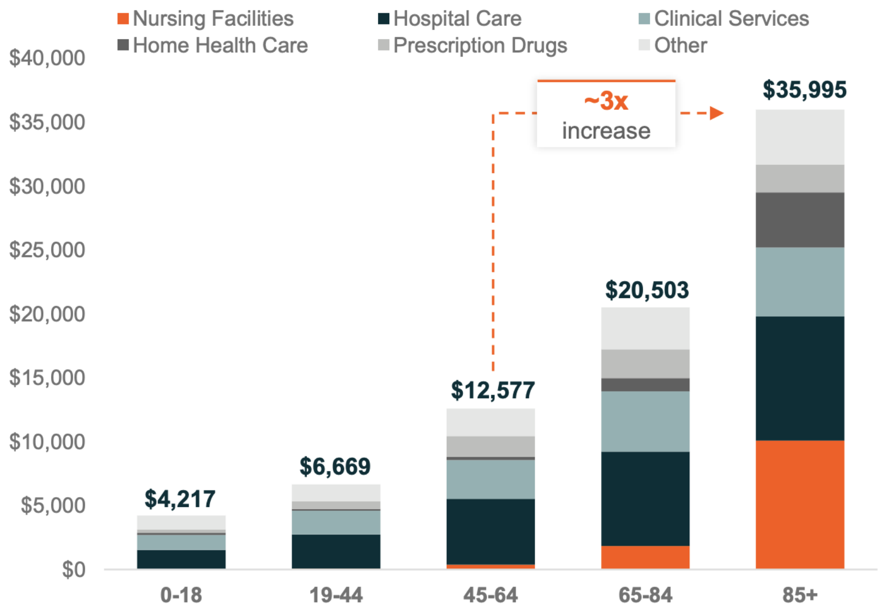

Figure 4.10 Average annual healthcare spend by age group (Global X Management Company LLC, 2026).

4.3.6 Autonomous mobility

McKinsey & Company (2025) defines autonomous mobility as follows:

> *"At its most basic level, an autonomous vehicle is one that uses advanced software, hardware, and services to operate by itself, with minimal human intervention. Self-driving vehicles stand to provide more affordable and convenient transportation options for passengers, freight, and urban mobility. But challenges remain, including those related to safety, economic viability, and regulation."*

There are six levels of assisted driving where level zero is the first, meaning that the vehicle has no systems assisting the driver what so ever. The highest, level 5, is defined as fully self driving. The autonomous vehicle (AV) can be sent out in any conditions without a driver. Some companies, like Waymo, already offer driverless taxi rides in select areas (McKinsey & Company, 2025). Ericsson states AVs as the ultimate uplink use case. Projected to grow at 19.9 percent CAGR growing from 68 billion USD in 2024 to 214 billion USD in 2030. Functions that are safety critical would be processed locally but telemetry, occurring almost at all times while the AV is operational needs about 1-10 Mbps and is not latency critical. Another is remote assistance where a human can operate the AV remotely needing several Mbps uplink and below 100 ms latency. Another future case would be using the AV as a sensor sending information about the surroundings, like weather and objects on the road, further increasing uplink (Riback, 2026). By spring 2025 Waymo reached one million autonomous rides per month, aiming for one million per week in 2026 (Waymo, 2025). Giattino et al. (2026) reports that the total number of autonomous kilometers driven in California is closing in on ten million as of December 2025, see figure 4.11.

Example: An autonomous vehicle drives itself to charge at night, picks up inbound packages in the early morning and delivers them during the day. By dinnertime it starts picking up food orders, and in the evening it picks up outbound packages.

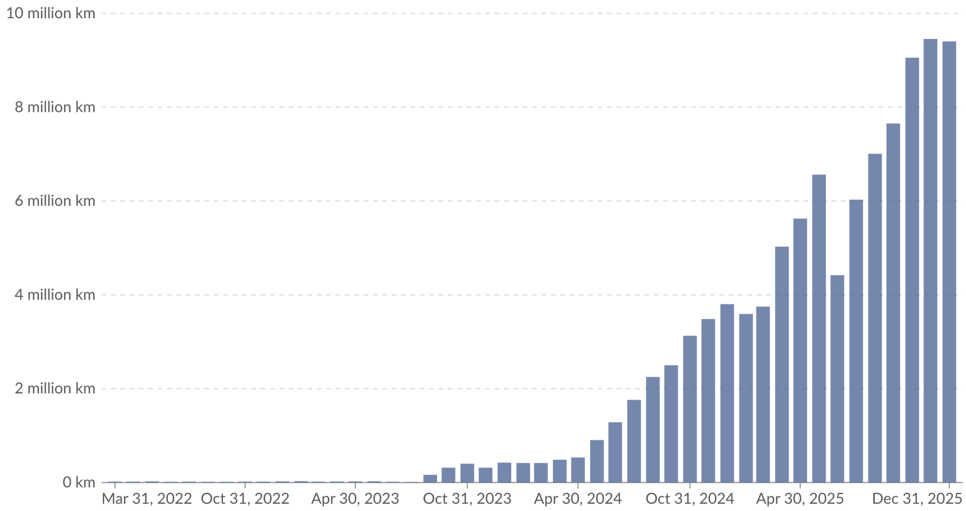

Figure 4.11 Total monthly distance traveled by passengers in California's driverless taxis (Giattino et al., 2026).

4.3.7 Autonomous robots

Stryker (2026b) defines physical AI as follows:

> *"Physical AI refers to artificial intelligence (AI) systems that operate in and interact with the physical world, rather than existing only in software or digital environments."*

Therefore, for the purposes of this study, the definition has been tightened. Surgical robots were moved to the health AI category, autonomous vehicles to the autonomous mobility category. Therefore the Autonomous robots category is defined as: Mobile physical robots, whether humanoid, quadrupedal, or otherwise, that use AI to perceive and interact with their environment.

Global X Management Company LLC (2026) lists physical AI as the next leap after Agentic AI. With a total estimated cost of 50 thousand USD per robot they are set to transform both industry and home applications. The humanoid market is massive, having a total addressable market of 4.85 trillion USD by 2035. Deployment is still early and is set to accelerate considerably after 2030, see figure 4.12. Uplink traffic for this application is heavy since a lot of data is sent for processing off device (Morris, 2026).

Example: A mobile robot navigates a warehouse floor. It identifies specific pallets, picks them up, and moves them to shipping docks. Another patrols the premises of an industrial site doing safety checks.

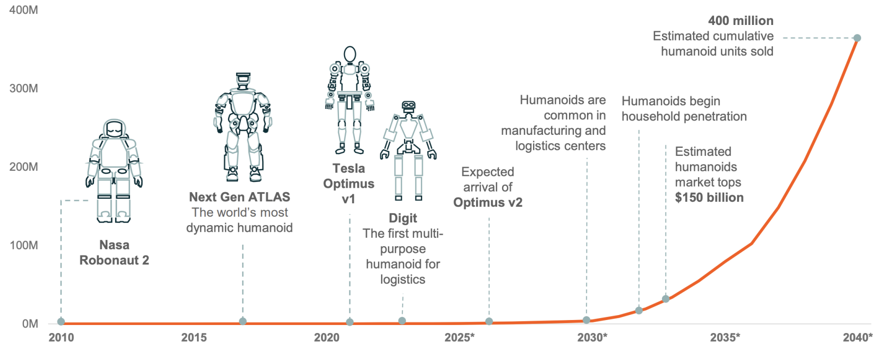

Figure 4.12 Cumulative Humanoid Unit Sales (Global X Management Company LLC, 2026).

4.3.8 Extended reality

Bernardo (2024) defines XR, VR, AR, and MR as follows:

> *"Extended reality (XR) is an emerging umbrella term for all immersive technologies, including virtual reality, augmented reality, and mixed reality.*
>
> *Virtual reality (VR) is a technology that creates a digitally simulated immersive environment or experience for users. It typically involves the use of specialised hardware and software to create a computer-generated 3D spatial, and possibly multi-sensorial, environment that can be interacted with in a seemingly real or physical way. VR allows users to experience and interact with a digital environment as though they were real.*
>
> *Augmented reality (AR) on the other hand is a technology that overlays digital information, such as text; images; sound; videos; or 3D models, onto the real-world environment. AR enhances the real world by adding computer-generated elements to it.*
>
> *Mixed reality (MR) systems are immersive technologies that bring physical objects into digital environments or digital objects into physical reality. One type of MR is Cinematic Reality, offering immersive 360 degrees viewing with live camera footage."*

Future advanced VR applications will need at least 1 Gbps downlink and 10 ms latency due to processing being done off-device (Mangiante et al., 2017). The story is however flipped with AR having real time understanding, if the AR glasses have an AI that is interpreting the world, high resolution video needs to be sent off device for processing. In that case about 1-5 Mbps uplink and 0.5 Mbps downlink is needed resulting in 88 percent uplink traffic requiring a 200 ms or lower response time (Riback, 2026). MarketsandMarkets (2026) estimates the market to grow from 41 billion in 2025 to 140 billion by 2032.

Example: A technician wears AR glasses that project digital repair instructions directly onto a physical engine during a maintenance task.

4.3.9 Cloud gaming

NVIDIA (2021) defines cloud gaming as follows:

> *"Video games are interactive, obviously. So, cloud gaming servers need to process information and render frames in real time. Unlike movies or TV shows that can provide a buffer — a few extra seconds of information that gets sent to your device before it's time to be displayed — games are dependent on the user's next keystroke or button press."*

Mobile and VR gaming, both requiring off-device computing if the user is playing graphic intensive games, is expected to grow to 340 billion USD by 2030 (Global X Management Company LLC, 2024).

Example: A person plays a highly graphic intensive game on a low-spec tablet since all compute is done in the cloud and the tablet only takes input and displays the gameplay.

4.3.10 Summary of technologies

The technologies set to drive AI enabled mobile network traffic discussed above are summarized and defined below:

- **AI Assistants:** AI systems that respond to user queries and generate outputs in real time.
- **IoT:** Connected physical devices that collect and exchange data, processed by AI.
- **Software-based Agentic AI:** Autonomous AI agents that plan and execute multi-step tasks with minimal human input.
- **Autonomous Mobility:** Vehicles or drones that navigate and make decisions without human control.
- **Smart Cities:** AI-driven coordination of urban infrastructure to improve efficiency and services.

- **XR (VR/AR):** Immersive technologies that blend or replace physical reality with digital environments.
- **Cloud Gaming:** Games rendered on remote servers and streamed to end-user devices in real time.
- **Healthcare AI:** Remote patient monitoring via wearables that continuously track vitals and flag anomalies to clinicians and remote operation of medical robots such as surgical robots.
- **Autonomous Robots:** Physical robots, whether humanoid, quadrupedal, or otherwise, that use AI to perceive and interact with their environment.

4.4 The network for AI ecosystem

4.4.1 The roles of today

The ecosystem was structured around roles rather than individual companies. As shown in table 4.3, twelve distinct roles are identified. Focusing on roles have been chosen due to three reasons. First, several companies have multiple roles at the same time. Apple is a device manufacturer, chipset manufacturer and app developer. Second, the number of companies across the global N4AI ecosystem runs into the hundreds. This makes company level analysis impractical for scenario construction. Third, roles are more stable: individual companies enter, exit and merge, but the functions persist. These roles are necessary as part of the construction of an ecosystem pie model described in section 3.1.1.

4.4.2 Revenue breakdown of today's roles

To get a grasp of the size of each role in the N4AI ecosystem the authors identified the top players in each role and added their revenue together, see figure 4.13. A total of 91 companies were found which is a fraction of the total but the authors estimates that this covers a considerable part of the total N4AI ecosystem revenue. In figure 4.14 there is a link to an interactive pie chart where all the identified companies and their revenue is visualized. This is an adjusted version of Jacobides et al. (2018) ecosystem pie model explained in section 3.1.1. Structured with companies in the outer ring, roles in the middle ring and at the center of the ecosystem pie, the N4AI value proposition. The size of each role being its total revenue. Revenue numbers are the latest published, mostly from 2025, but some are estimates. A full breakdown and sources can be found in A.4.

4.4.3 Network ties

To identify the interdependencies between roles in the ecosystem the authors estimated major ties between roles to visualize their position in the value chain resulting in figure 4.15. It illustrates the flow of money from consumers through the N4AI ecosystem. Different types of network ties are explained in section 3.1.2. The ties in figure 4.15 are transactional market ties where arrows represents an actor buying a product or service from another. They are layered in figure 4.15 to represent flows from top to bottom, with the top being consumers and the bottom being further and further down the value chain.

Table 4.3 Ecosystem roles in the AI and telecom ecosystem

Ecosystem role	Description	Activities
Hyperscalers	Large cloud providers that operate massive-scale data center infrastructure and offer compute, storage, and AI services globally (AWS, Azure, Google).	Build and operate data centers; provide cloud and AI services, develop proprietary infrastructure, enable enterprise platforms.
Telecom equipment vendors	Companies that design and supply the hardware and software that builds and runs communication networks (Ericsson, Nokia, Huawei).	Develop network hardware/software, deploy infrastructure, support network upgrades (5G/6G)
CSPs / Operators	Network operators that provide connectivity services to consumers and enterprises (Vodafone, Deutsche Telekom, etc.).	Operate networks, sell connectivity, manage spectrum, ensure coverage and performance.
AI model providers	Companies that develop and deploy foundational AI models via APIs or integration (OpenAI, Anthropic, etc.).	Train models, provide APIs, integrate AI
Device manufacturers	Companies that produce end-user hardware (Apple, Samsung, etc.).	Design and manufacture devices, manage supply chains, launch products
Chipset manufacturers	Companies that design processors and chips (Qualcomm, Apple, NVIDIA).	Design semiconductors, optimize performance, license IP
Chip foundries	Facilities that fabricate chips at scale (TSMC, Intel, etc.).	Manufacture chips, operate chip factories
App developers	Entities building software applications (Google apps, Zoom, Uber, etc.).	Develop and maintain apps, deploy on platforms
Content platforms	Services distributing digital content (Netflix, Meta, TikTok, etc.).	Create/distribute content, manage engagement
Infrastructure providers	Companies owning shared assets like towers and data centers (Equinix, American Tower, Cellnex).	Build and lease infrastructure
Regulators & Governments	Public bodies setting legal and policy frameworks.	Define regulation, allocate spectrum, enforce compliance
Consumers	End users of services and devices.	Consume services, generate data, drive demand

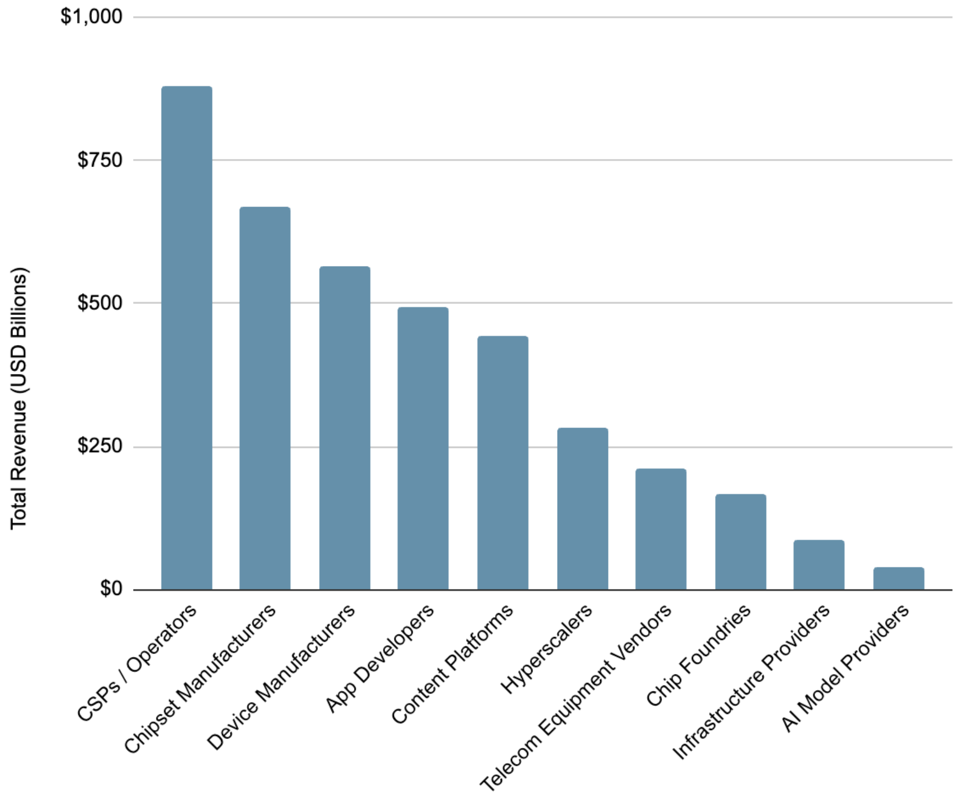

Figure 4.13 Revenue comparison of the N4AI roles in late 2025.

Figure 4.14 **Revenue breakdown of the N4AI roles and companies in late 2025. (Link to interactive chart)**

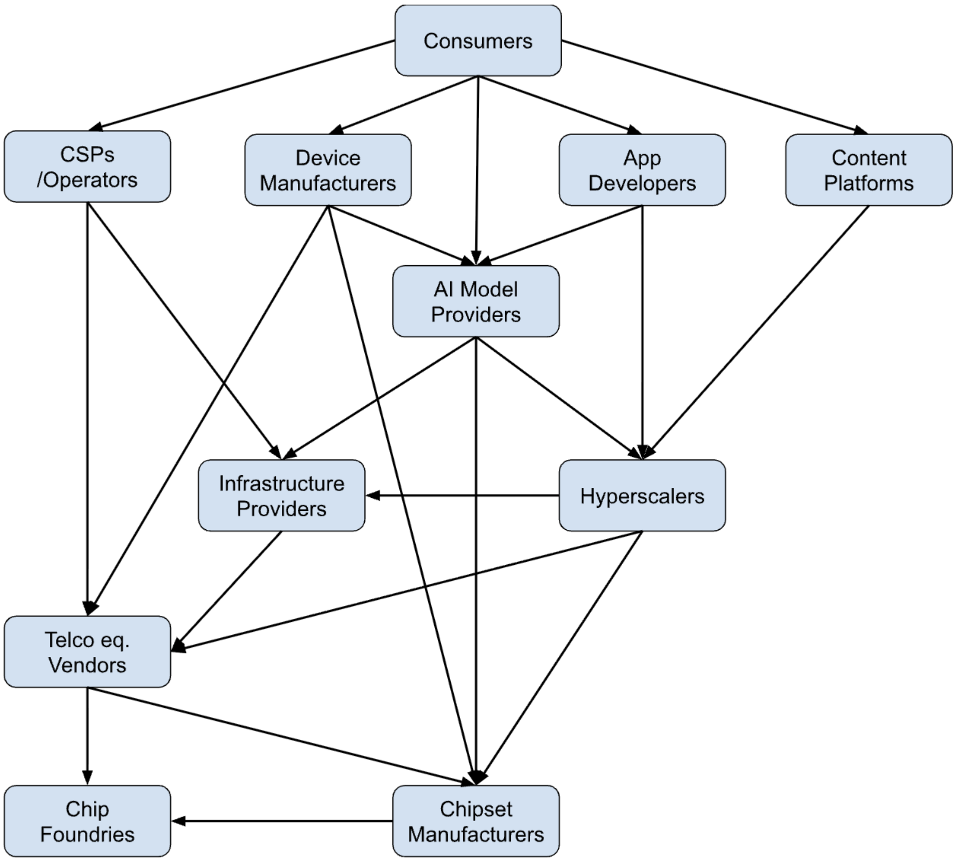

Figure 4.15 Money flow through the N4AI ecosystem, defined by the authors.

4.5 Macro factors

World Economic Forum (2026a) found that in the two year time frame primary risks that the world face are: geopolitical, technological, societal, and environmental. When looking at a 10 year time frame these risks shift to being primarily environmental, technological, and societal. Environmental making up four of the six top risks. The specified risks are listed in table 4.4. Economic risks does not find their way onto the top ten in either 2 or 10 year time frames, they have however climbed the most since the 2025 report. The risk of an economic downturn having gone from 19th to 11th place.

Table 4.4 Global risks ranked by severity in the short and long term. (World Economic Forum, 2026a).

Short term (2 years)	**Long term (10 years)**
1. Geoeconomic confrontation	1. Extreme weather events
2. Misinformation and disinformation	2. Biodiversity loss and ecosystem collapse
3. Societal polarization	3. Critical change to earth systems
4. Extreme weather events	4. Misinformation and disinformation
5. State-based armed conflict	5. Adverse outcome of AI technologies
6. Cyber insecurity	6. Natural resource shortages
7. Inequality	7. Inequality
8. Erosion of human rights and/or of civic freedoms	8. Cyber insecurity
9. Pollution	9. Societal polarization
10. Involuntary migration or displacement	10. Pollution

Geopolitical fragmentation is highly anticipated in the coming 10 years, see figure 4.16 (World Economic Forum, 2026a). This would reasonably result in increased data privacy concerns at a personal and national level, which in turn would lead to AI regulation. World Economic Forum (2026a) also emphasize this point saying that organizations face a rising pressure for both data flows and AI model behavior to be secure. Policy frameworks like EU's Gaia-X, the US Clarifying Lawful Overseas Use of Data (CLOUD) Act and privacy regulations such as the General Data Protection Regulation (GDPR) embeds trust, compliance and sovereignty into policy frameworks. These frameworks are reactive since AI and autonomous systems advance faster than existing regulations.

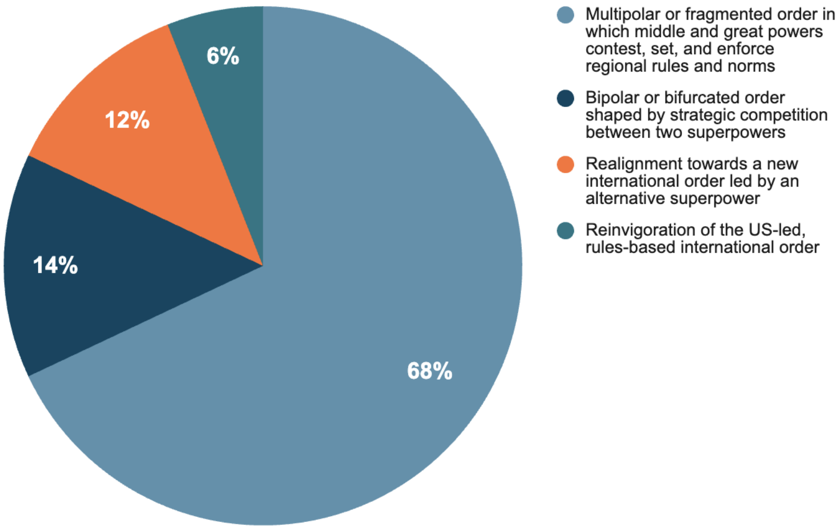

Figure 4.16 "Which of the following best characterizes the global political environment for cooperation on global risks in 10 years?" Question from World Economic Forum risks perception survey. (World Economic Forum, 2026a)

World Economic Forum (2026b) states the drivers of a sovereign AI infrastructure as distinct. They are: economic self-interest, data protection and geopolitical resilience. CSPs are positioned to control this sovereign infrastructure since they are seen as trusted, secure and local stewards. Already 81 percent of enterprises seek telcos as partners for sovereign AI adoption (World Economic Forum, 2026b).

Heavy infrastructure investment is being driven by AI. Noffsinger (2025) projects that with continued momentum a total of 5.2 trillion USD in AI-related data center capacity will be spent in the five years between 2025 and 2030. Omdia (2025) forecasts telecom capital expenditures (CAPEX) to grow at 3.6 spercent CAGR and reach almost 400 billion USD by 2030. 2028 and onward the 6G buildout will start and ramp up. Telecom operators are, as mentioned before, expected to expand into data centers and invest in specialized hardware. Driven by AI infrastructure, cloud services, and digital sovereignty policies.

On a parallel track the cost of training and inferencing is on a steep decline, figure 4.7, driven in part by the cost of compute. A fixed amount of computing power is getting 30 percent cheaper each year (Rahman, 2024).

With this backdrop, the authors identified six macro factors, one for each category in the PESTEL-model that is posed to have a large impact on the global telecom and AI sector:

- **Political - Geopolitical fragmentation:** The division of the global AI and ICT landscape into competing regional blocks driven by political tensions.

- **Economic - Infrastructure investment flows:** The direction and scale of capital allocated toward building network, compute, and AI infrastructure.
- **Social - Data privacy:** The degree to which consumers and regulators accept or resist how personal and sensitive data is collected, stored, and used.
- **Technological - Compute efficiency:** The ability to deliver more AI and network processing power per unit of energy and cost.
- **Environmental - Climate change:** The capacity of ICT and AI infrastructure to operate sustainably and withstand climate related disruptions.
- **Legal - AI regulation:** Policies that govern the development, deployment, and use of AI systems.

4.6 Scenario building

To build scenarios three orthogonal axes of change with mutually exclusive endpoints were identified based on the macro factors. The macro factors push the future back and forth along the three axis. For example, AI regulation could push the ecosystem toward resilience through mandates or investments may favor centralized compute. This was done to enable a set of outcomes in accordance with Schwartz (1996) method of scenario building and Lindgren and Bandhold (2009) scenario cross in section 3.2.4.

4.6.1 The three axes

4.6.1.1 Axis 1: Centralized vs. Distributed compute topology

Author definition:

- **Centralized compute topology:** AI training and inference is primarily performed in hyperscale cloud or large centralized data centers.
- **Distributed compute topology:** A significant share of compute is performed at the network edge or directly on devices (e.g., edge nodes, base stations, endpoints).

4.6.1.2 Axis 2: Concentrated vs. Distributed value capture

Author definition:

- **Concentrated value capture:** A small number of actors capture the majority of economic value in the ecosystem.
- **Distributed value capture:** Economic value is spread across many actors in the ecosystem (e.g., operators, infrastructure providers, device manufacturers vendors).

4.6.1.3 Axis 3: Cost efficiency vs. Resilience focus

Author definition:

- **Cost efficiency focus:** System design prioritizes cost efficiency and resource utilization.
- **Resilience focus:** System design prioritizes redundancy and operational continuity.

4.6.2 Possible Scenarios

Given the three axes there are eight distinct future scenarios that can be constructed. These are listed in table 4.5. By asking questions about the axes in round one of the Delphi study it is possible to identify the scenarios considered most likely and thereafter ask more detailed questions in round two.

Table 4.5 Scenario matrix of all possible scenarios given the three axes.

Compute Topology	Value Capture	System Design Focus
Centralized	Concentrated	Cost efficiency
Centralized	Concentrated	Resilience
Centralized	Distributed	Cost efficiency
Centralized	Distributed	Resilience
Distributed	Concentrated	Cost efficiency
Distributed	Concentrated	Resilience
Distributed	Distributed	Cost efficiency
Distributed	Distributed	Resilience

5 Empirics

This chapter presents the empirical findings of the two-round Delphi study. The first section covers the round one results, including panel composition, expert views on AI-driven traffic, ecosystem structure, and data sovereignty as a structural signal. These findings inform the scenario descriptions that follow. The chapter then presents the round two results, in which the expert panel validated and elaborated each scenario, and concludes by synthesizing what the combined rounds resolved and what remains open.

5.1 Delphi study: round one

To systematically track the outreach process, all potential participants were logged and assigned to one of six stakeholder categories reflecting their institutional affiliation and relationship to the research topic. Each contact was then assigned a stage label as the process progressed, allowing the pipeline to be monitored at every step, from initial outreach to confirmed participation or rejection.

Table 5.1 Outreach pipeline summary by stakeholder category. Total reached out to are those contacted (includes declined) plus those where the email was rejected and where it did no longer exist.

Stage / Bucket →	Academia	Regulator / institution	Business external	Technology external	Business internal	Technology internal	Row Total
Contacted	61	54	36	34	27	2	**214**
Accepted	15	7	2	2	3	1	**30**
Acceptance rate:	25%	13%	6%	6%	11%	50%	14%
Answered Round 1	10	7	1	2	2	1	**23**
Answer rate Round 1:	67%	100%	50%	100%	67%	100%	77%
Answered Round 2	6	4	0	2	1	1	**14**
Answer rate Round 2:	60%	57%	0%	100%	50%	100%	61%
Declined	14	11	3	3	0	0	**31**
Mail rejected	0	3	11	8	0	0	**22**
Does not exist	0	0	10	17	0	0	**27**
Total reached out to	**61**	**57**	**57**	**59**	**27**	**2**	**263**

5.1.1 Panel composition and interpretive scope

Initially, experts were recruited across six categories: academia, technology-external, business-external, regulators and institutions and technology- and business-internal Ericsson staff. The design followed two axes (internal versus external to the vendor and technology versus business) to keep commercial, technical and research perspectives from collapsing into a single institutional view. Regulators and industry associations were treated as a separate category because they set deployment conditions that no market actor controls.

Due to the low response rate from both internal and external technology/business stakeholders, the categories were redefined as follows: the original external category was consolidated into a single group, and the internal category into another. Furthermore, the former regulators and institutions category was divided into two separate groups: regulators/institutions and associations & partnerships. In conclusion, the initial categorization proved inadequate, so it was revised, see figure 5.1.

Round one gathered responses from 23 experts assembled through purposive sampling across academia, technology vendors, regulators and industry associations. Ten respondents were affiliated with academic institutions, three worked at Ericsson, three at other technology companies, three held regulatory or public-sector roles and the remainder represented foundations, industry associations, and public-private partnerships.

Across the full panel, 18 respondents identified their primary focus as technology and five as business or strategy. The panel was also highly experienced: 14 of the 23 respondents reported more than 20 years of professional experience in telecom or AI.

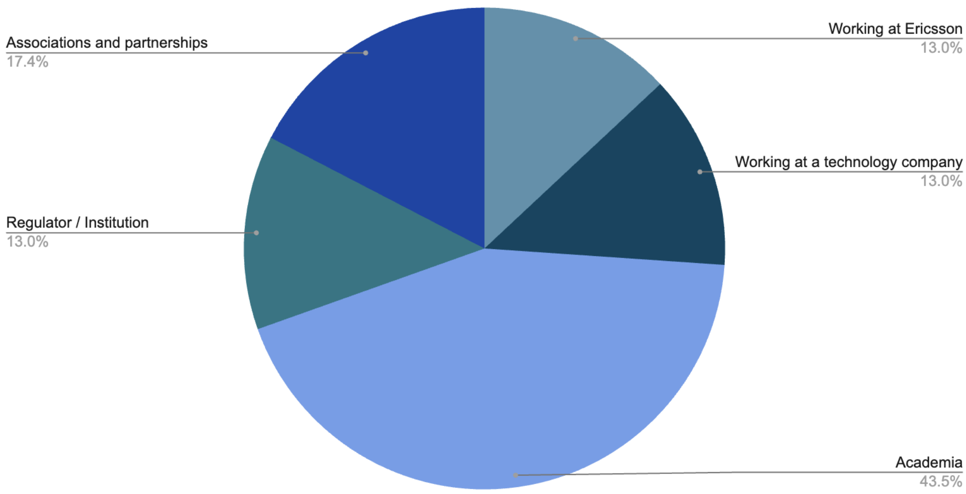

Figure 5.1 The distribution of experience among the experts. n = 23

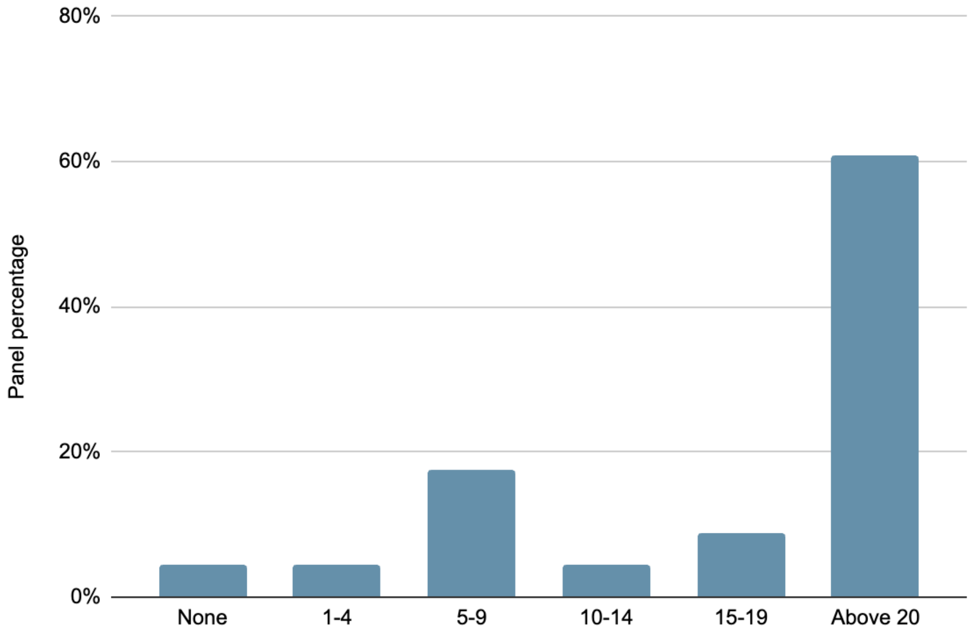

Figure 5.2 Years of professional experience among the experts.

The composition of the panel has interpretive consequences worth addressing before analyzing the findings. The panel distribution, as shown in figure 5.1, weighted toward technology expertise within the academia group, which was 44 percent of the panel group. Moreover, the academic framing is well positioned to assess infrastructure constraints, architecture trajectories and adoption timelines. However, it is less reliable on questions of operator business model viability, where experts with direct commercial experience tend to hold more grounded views. The near-absence of CSPs as a primary respondent segment means that questions about monetization and value capture were answered from the outside looking in. This limitation was explicitly noted in the round one qualitative analysis and was addressed in round two by centering those sessions more directly on actor-specific scenarios. The experts were also notably self-assessed as more knowing in telecom than in AI. Self-reported telecom knowledge scores ranged from three to 10, the majority clustered between eight and 10. AI knowledge scores were more spread, with a noticeable share of respondents rating themselves between five and seven. Following is a number of scoring tables with respect to AI-driven traffic, adoption rate, revenue potential, ecosystem impact and macrofactors of influence. See section 2.3.1.1 for the coding of the scores.

5.1.2 AI-driven traffic and network architecture

5.1.2.1 The uplink growth consensus

Regarding the question on how the uplink-to-downlink (UL/DL) ratio will shift by 2035, the panel showed moderate consensus around a meaningful but not dramatic change. Nine experts (39 percent) predicted a 30/70 split, six (26 percent) predicted 20/80 and three (13 percent) predicted 40/60. That is, combined, 78 percent of the panel expected uplink traffic to reach at least 20 percent of the total traffic, which is an increase from 10 percent today. The remaining respondents were more aggressive: three predicted 70/30 and one 60/40. Only one respondent declined to answer. The modal answer of 30/70 is consistent with the traffic forecasting material in section 4.2, where uplink growth is projected as the primary directional change associated with AI workload traffic. The outlier responses at 60/40 and 70/30 were accompanied in the qualitative comments by the argument that autonomous vehicle telemetry and physical AI applications would drive a more dramatic shift. These represent a minority but technically linked view. In fact, one of the respondent who answered 70/30 most probably meant to answer 30/70 as the motivation of the answer matched the consensus. The qualitative layer beneath the consensus in question reveals a structural disagreement about what will actually drive the numbers: where AI inference runs. If models process primarily in cloud data centers, mobile endpoints continuously transmit raw or compressed data upstream. If inference migrates to edge nodes or runs on-device, the data volume leaving the endpoint could remain relatively modest even as AI usage scales. One respondent articulated the conservative position clearly:

> "*I think that by 2035 we will have AI solutions that also operate on edge and do not require a lot of data being sent up.*"

The opposing view was equally represented:

> "*AI is likely to at least triple uplink traffic, as device-generated data will increasingly be sent to remote inference and training platforms whenever local AI capabilities are insufficient.*"

5.1.2.2 Technology traffic generation and adoption rates

The panel's assessment of individual technology categories reveal a clear pattern when synthesized. XR (VR / AR) was rated highest for expected mobile traffic generation, with 14 of 23 experts (61 percent) placing it at "high" or "massive". Cloud gaming and autonomous robots followed, with 48 percent and 39 percent respectively rating them "high" or "massive". Agentic AI and autonomous mobility sat in the middle range. Healthcare AI was rated lowest for mobile traffic impact, with 13 of 23 experts (57 percent) placing it at "small". Furthermore, according to the panel, healthcare AI will process largely locally or within controlled hospital networks rather than being processed over general mobile infrastructure.

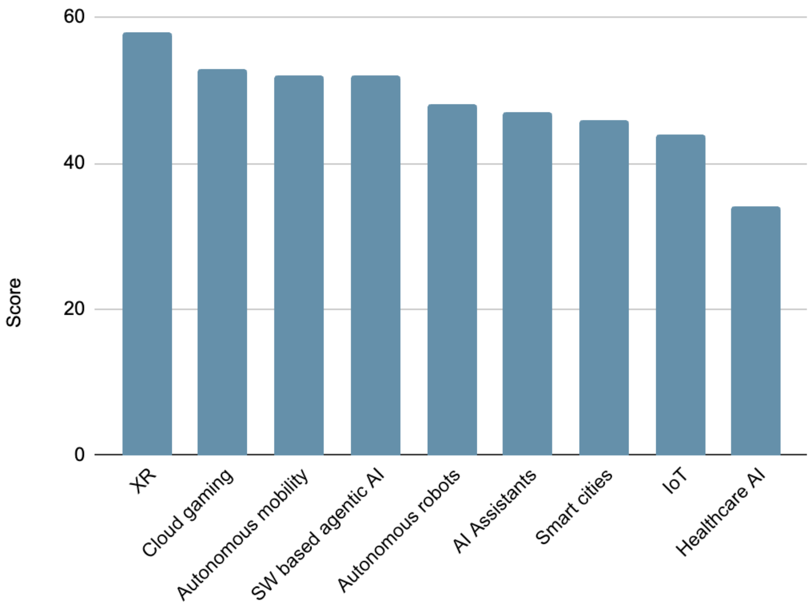

Figure 5.3 AI and connectivity enabled technologies traffic score.

The data gathered from the adoption rate viewed another perspective. According to the expert panel, AI assistants were the clear dominator with regards to the adoption rate, with 17 out of 23 (74 percent) votes on "fast" or "very fast". Software-based agentic AI followed closely, with 16 of 23 (73 percent) voting "fast" or "very fast" as well. At the other end, smart cities and IoT drew moderate or slow ratings from the majority of the panel, see figure 5.4.

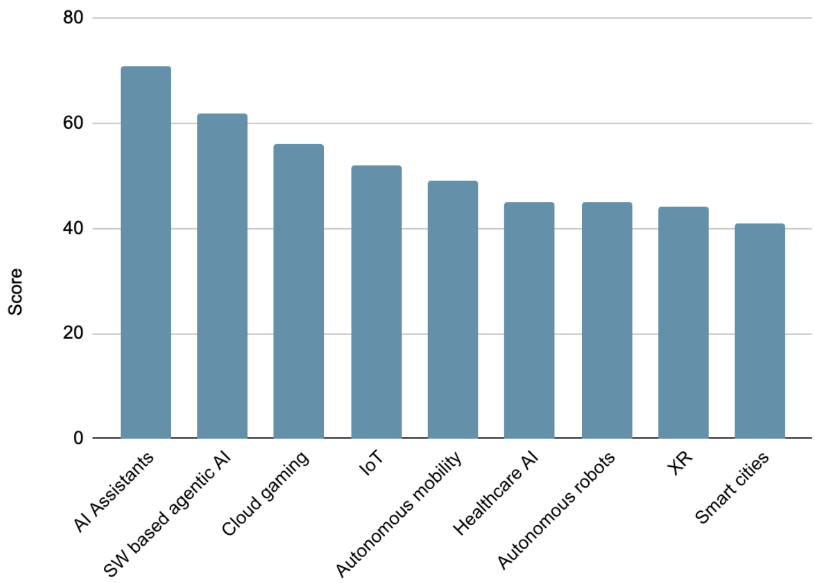

Figure 5.4 AI and connectivity enabled technologies adoption rate scores.

The pattern across both dimensions, traffic generation and adoption rate, corresponds well onto the software-first versus physically-constrained axis noted in the qualitative comments. One respondent provided the clearest articulation of this axis:

> "*Adoption rates will vary significantly depending on whether technologies are primarily software-driven or constrained by physical infrastructure, regulation and integration complexity. AI assistants and agentic AI will progress rapidly as software-first innovations. In contrast, autonomous mobility, smart cities, and healthcare AI face longer deployment cycles.*"

As addressed in section 3.2.3, innovations have different diffusion capabilities. Moreover, the technologies with the lowest infrastructure and regulatory barriers are tracking the steepest adoption curves, whereas those requiring physical rollout at scale will likely solely contribute meaningfully to network traffic in the 2030 and 2035 window, when the new ecosystem paradigm shifts have been more established. The revenue potential ratings broadly follow adoption confidence, with Agentic AI collecting 14 out of 23 respondents in the "high" or "massive" category (67 percent). Healthcare AI got 30 percent high or massive ratings, despite its traffic contribution is expected to be low. Furthermore, assuming the forward trajectory: when AI penetrates healthcare, the value creation will be large even if the network traffic is modest. Only nine percent voted "high" or "massive" on the revenue potential for Smart cities, which is consistent with the IoT cautionary signal in the qualitative comments.

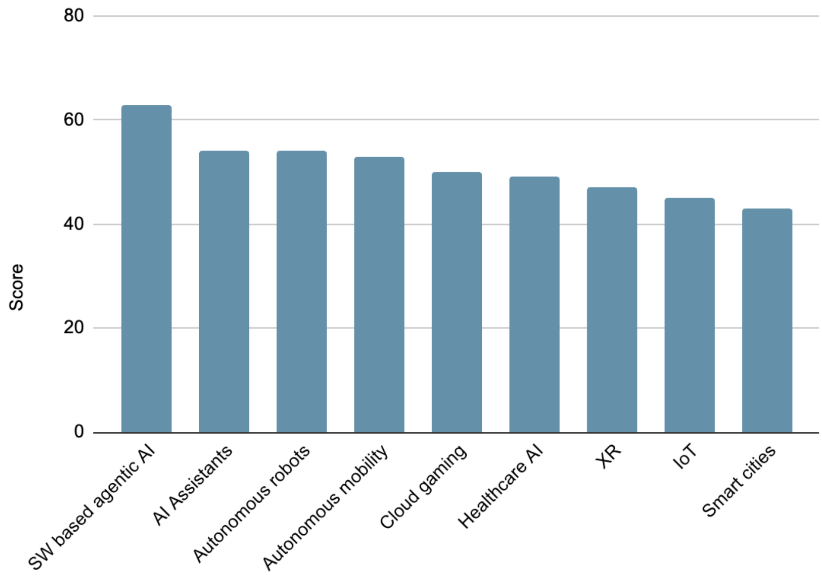

Figure 5.5 AI and connectivity enabled technologies revenue scores

5.1.2.3 The shift from human-centric to machine-centric traffic

The open-ended question - "*What single structural change will most reshape mobile network traffic by 2035?*" - created the most analytically developed material in the entire set of data. Experts clustered strongly around one theme: the transition from humans as the primary consumers of network services to AI agents and machines as a structurally distinct traffic class. One of the experts formulated it in the following way:

> "*The most significant shift will be the transition from humans as the primary consumers of network services to AI agents generating a substantial share of traffic. These agents will autonomously browse, query, stream, and interact with application programmable interfaces (APIs), other agents, and real-time data sources, fundamentally transforming network traffic characteristics, latency requirements, and service design.*"

This particular finding about the shift in traffic generation is in well accordance with the N4AI framing established in Chapter 1. The experts independently converged on this point as the structural change, which in turn acted as the foundation of empirical support for the framing choice and relevance of the scenarios built in section 5.2.

5.1.3 Ecosystem structure and value distribution

5.1.3.1 The current baseline: concentrated value at the top of the stack

The panel was almost unanimous on the current state, thus acting as the baseline scenario of today's description of the ecosystem. Today's compute topology was rated centralized by 22 of 23 respondents, 96 percent. Economic value capture was rated was concentrated by 21 of 23, 91 percent. Cost efficiency was rated as the dominant system logic by 17 of 23 experts, 74 percent. The three readings collectively describe the Efficiency Pipeline scenario from section 5.2, and could be framed as follows: AI training and inference run in hyperscaler cloud data centers. Value is captured by those who control the platform and compute layers. Finally, the organizing logic of the entire ecosystem is the cost optimization. The ratings regarding the N4AI ecosystem reinforced the current baseline, where the value is located at the top of the ecosystem stack. Hyperscalers were the most highly rated actor today, 86 percent rated "high" or "very high". By contrast, CSPs and operators were rated far lower with only 27 percent of high votes.

5.1.3.2 The 2035 trajectory

The panel's view, when asked about 2035, diverged. Regarding the compute topology, there was dominant consensus, 91 percent predicted the topology would become more distributed. Regarding the system focus, 61 percent predicted a shift toward greater resilience orientation. Furthermore, these two axes showcased clear directional consensus. However, there was an exception, only 48 percent predicted that the distribution of value would increase, 30 percent expected no change and merely 13 percent expected a more concentrated value capture. Looking ahead, The N4AI ecosystem could become more physically distributed, with compute spreading to edge nodes, while economic value capture could remain concentrated among those who control the software and platform layers. Looking at the qualitative data adds insights, mentioning that those controlling compute, platforms and user access will maintain or even extend their position:

> "*By 2035, influence in the ecosystem will increasingly be shaped by control over intelligence, compute, and data. AI model providers and hyperscalers will continue to hold significant power.*"

Furthermore, a minority of experts held the following position, that is, operators have a specific time window:

> "*Telecom operators cannot afford to sit idly like they did in 4G and 5G or they will once again be left behind.*"

The presented positions do not average out into a consensus view, which is further studied in the second round.

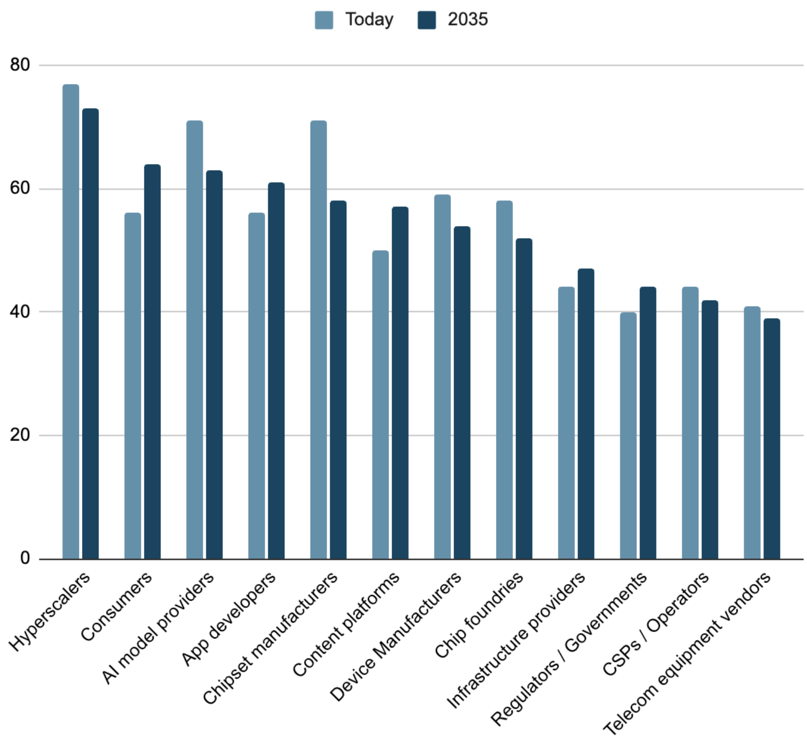

Figure 5.6 Ecosystem scores today and 2035.

5.1.4 Data sovereignty as a structural signal

An interesting finding from round one was that data sovereignty was a structural driver of CSP relevance. It was not prompted by the survey, but experts raised it because the theme seemed relevant to their responses. Among the macro factors, infrastructure investment flows were rated highest as the most influential factor, alongside geopolitical fragmentation.

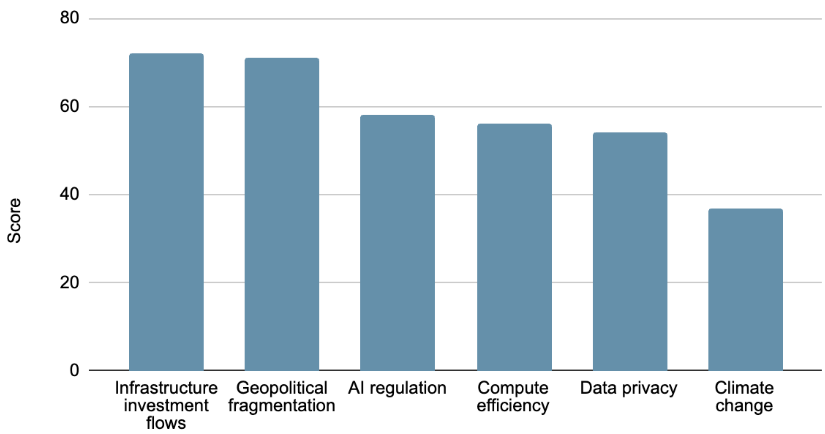

Figure 5.7 Factors of influence.

Following is three graphs displaying the panel's view on how each macro factor drives compute topology, value capture and system focus, respectively toward one or the other side of the axes.

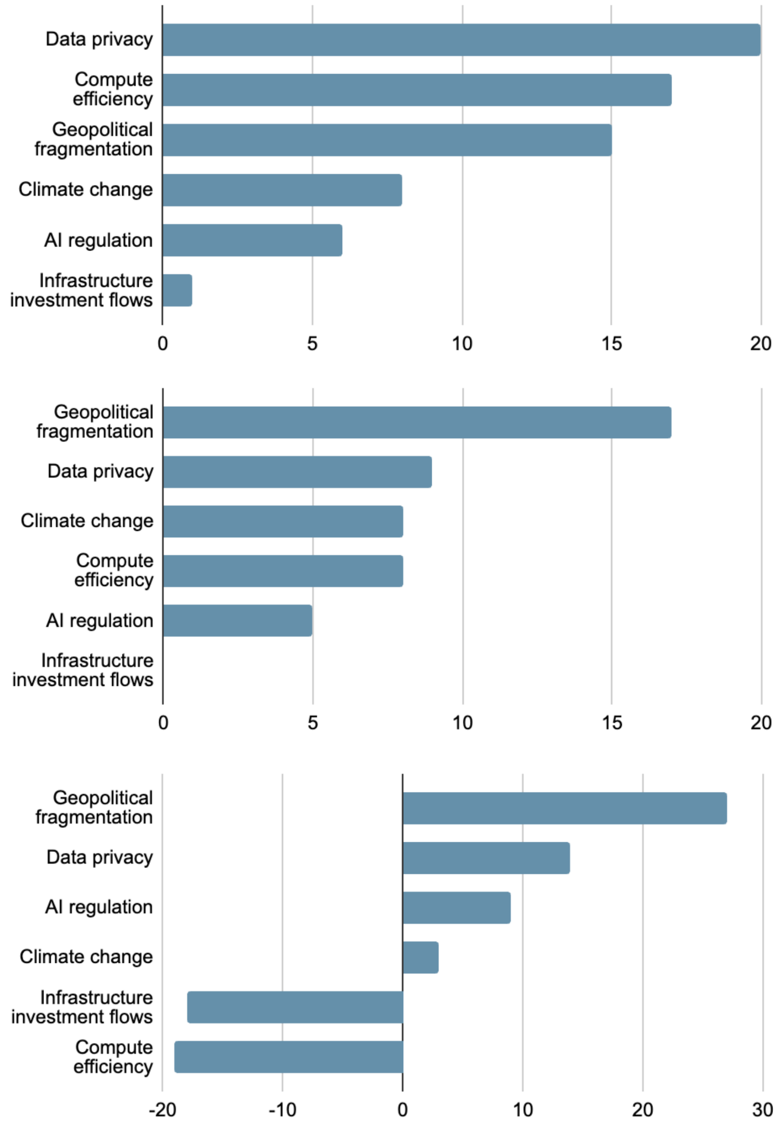

Figure 5.8 Expert panel assessments of how each macro factor influences compute topology (top: centralized–distributed), value capture (middle: concentrated–distributed), and system focus (bottom: cost efficiency-resilience). Each graph reflects aggregated Delphi panel responses across the identified macro drivers shaping AI infrastructure toward 2035.

5.1.5 Mapping panel views onto the scenario axes

The three axes developed in section 4.6.1 can be used to investigate where expert opinions clustered while also addressing what remains uncertain.

Table 5.2 Panel views mapped onto the three scenario axes.

Axis	Direction	Key uncertainty	Primary driver
1. Compute topology (Centralized → Distributed)	91 percent expect more distributed compute by 2035	Pace and degree of edge/device inference migration	Inference location
2. Value capture (Concentrated → Distributed)	48 percent expect more distributed value capture	Whether data sovereignty mechanisms activate	Regulatory and geopolitical developments
3. System focus (Cost efficiency → Resilience)	61 percent predict a shift toward more resilience	Net outcome between the competing pressures	Geopolitical fragmentation vs. infrastructure economics

5.1.6 What round one resolved and what it left open

Round one produced four clear findings. First, uplink traffic will grow meaningfully, with the 30/70 split as the modal expert prediction and the machine-centric traffic shift as the agreed structural explanation. Second, software-first AI applications (assistants, agentic AI) will reach commercial scale considerably earlier than physically-constrained categories (smart cities, autonomous mobility, healthcare AI). Third, today's ecosystem is effectively described by the Efficiency Pipeline scenario, with near-unanimous agreement on centralized compute, concentrated value and cost efficiency as organizing logic. Fourth, data sovereignty is the strongest identified mechanism by which this baseline could shift toward distributed value capture, with geopolitical fragmentation as the primary driver. The inference location variable, where AI inference will actually run by 2035, determines traffic and architecture outcomes, however this was not directly asked. The conditions under which CSPs could capture more value were identified qualitatively but not tested against structured scenarios.

5.2 Scenario descriptions

Which characteristics today's N4AI ecosystem has is clear with 96 percent answering centralized compute, 91 percent answering concentrated value capture and 74 percent answering a focus on cost efficiency. The panel was, from round one, overwhelmingly in agreement (91 percent) that the ecosystem in 2035 will have moved toward a more distributed compute topology. Therefore all future scenarios has this as a set parameter. The other two axes were not as clear. When asked about economic value capture, 48 percent of experts said the ecosystem is posed to move toward a more distributed value capture while 43 percent expected no change or more concentrated. The remaining nine percent did not answer. System focus also split the panel but not by as much. 61 percent expects a move toward more resilience and 26 percent expects it to be the same or more focus on cost efficiency. 13 percent did not answer when asked about system focus. Four distinct scenarios can be built using the Lindgren and Bandhold (2009) scenario cross described in section 3.2.4 yielding the scenario matrix in figure 5.9.

When describing each scenario the authors aimed to have the same structure for ease of comparison. Each having a catchy title and slogan as well as beginning with a deliberately vague description in past tense to get the panels imagination going without giving them answers to causes for why the scenario materialized or which consequences it has. Lastly ending with four bullet points listing differences in each scenario.

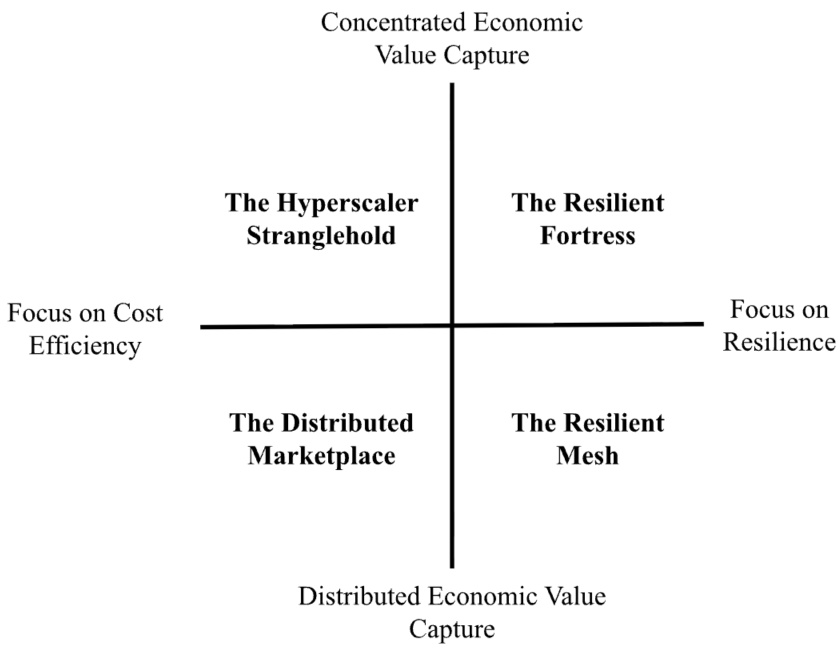

Figure 5.9 Scenario matrix with the four scenarios mapped onto each quadrant.

5.2.1 Today's baseline: efficiency pipeline

Centralized compute | Concentrated value capture | Cost efficiency focus

By 2026, the AI and telecom ecosystem had consolidated around a small number of dominant players. Hyperscalers and AI model providers owned much of the compute infrastructure, the developer ecosystems and the model deployment layer. As a result, capturing the majority of economic value generated by the AI boom. CSPs had been reduced to connectivity provision. Telecom equipment vendors aligned their roadmaps around solutions with low total cost of ownership to remain relevant in an environment driven by a reduction in operational expenditures. The ecosystem was highly efficient and value skewed toward those who controlled the value chain.

- **Market structure:** A small number of hyperscalers and AI model providers dominate the ecosystem, with CSPs and telecom equipment vendors competing on cost at the lower layers of the value chain.

- **Strategic logic:** Cost efficiency dominates across the ecosystem, with operators automating aggressively to reduce operational expenditures and procurement driven by lowest total cost of ownership.

- **Infrastructure characteristics:** AI training and inference are performed primarily in hyperscaler cloud data centers. CSPs provide connectivity but hold no meaningful stake in the compute layer.

- **Tensions:** Energy costs and latency demands are straining centralized architectures, geopolitical fragmentation is creating pressure for sovereign compute and regulatory scrutiny of value concentration is intensifying.

5.2.2 Hyperscaler stranglehold

More distributed compute | Concentrated value capture | Cost efficiency focus

By 2035, compute had distributed across edge nodes and regional infrastructure, driven by latency demands and energy constraints. Yet the economic structure of the ecosystem had shifted only modestly from the 2026 baseline. Hyperscalers and AI model providers had extended their influence, maintaining their grip on value capture despite no longer owning all the physical infrastructure. Efficiency remained the dominant strategic logic, with margin pressure continuing to push operators and telecom equipment vendors toward automation and standardization. The topology had opened up but the value distribution had not.

- **Market structure:** Value remains concentrated among hyperscalers and AI model providers. They had extended their reach across a more distributed infrastructure without ceding economic control.

- **Strategic logic:** Cost efficiency remained dominant, with commoditization and automation compressing margins further across infrastructure and connectivity.

- **Infrastructure characteristics:** Compute was deployed across edge and regional nodes, but control remained with a narrow set of dominant actors.

- **Tensions:** The growing gap between physical infrastructure ownership and economic value capture has created unresolved friction that neither market forces nor regulation have fully addressed.

5.2.3 Resilient fortress

More distributed compute | Concentrated value capture | More resilience focus

By 2035, the ecosystem's strategic logic had shifted from cost efficiency toward resilience. Compute had distributed, yet value had remained concentrated. Hyperscalers and AI model providers had positioned themselves as the trusted backbone of critical digital infrastructure. Resilience had become a premium product. Those who already controlled the value chain were best placed to deliver it. The ecosystem looked structurally different from 2026, but economic power had not meaningfully redistributed.

- **Market structure:** Value concentration has persisted and modestly intensified. Hyperscalers and AI model providers now compete on reliability and security credentials in addition to scale.
- **Strategic logic:** Resilience displaced cost efficiency as the primary investment driver. Procurement decisions across the ecosystem was increasingly shaped by risk management.
- **Infrastructure characteristics:** Compute became more distributed across edge and regional nodes with built in redundancy. However, operational control remained with a small number of dominant actors.
- **Tensions:** Ecosystem concentration created a systemic dependency, the ecosystem is more robust but remains exposed to a few dominant players.

5.2.4 Distributed marketplace

More distributed compute | More distributed value capture | Cost efficiency focus

By 2035, both compute and economic value had distributed across a broader set of actors than at any point in the prior decade. Edge infrastructure had become common. Competition for workloads previously dominated by hyperscalers had increased. Cost efficiency remained the governing logic. Competition was intense and differentiation was difficult. The ecosystem had opened up and cost competition was high.

- **Market structure:** The ecosystem had fragmented, with a wider range of actors competing across infrastructure and service layers.
- **Strategic logic:** Cost efficiency remained dominant, with intense competition across the fragmented landscape compressing margins and rewarding scale and automation.
- **Infrastructure characteristics:** Compute got broadly distributed across edge nodes and regional infrastructure. Ownership was spread across many actors rather than concentrated to hyperscalers.
- **Tensions:** Low margins reduce investment capacity, raising questions about who funds large scale infrastructure upgrades.

5.2.5 Resilient mesh

More distributed compute | More distributed value capture | More resilience focus

By 2035, compute had distributed toward the edge, value had spread across a broader set of players and resilience had become a priority. No ecosystem player controlled the entire value chain. Redundancy was treated as something of value rather than an overhead cost. The ecosystem had been redesigned to absorb shocks having continuity prioritized over cost reduction at every turn.

- **Market structure:** The ecosystem got broadly fragmented, each player held distinct but interdependent roles across the ecosystem.
- **Strategic logic:** Resilience was something worth paying extra for. Investments flowing toward continuity rather than only cost reduction.
- **Infrastructure characteristics:** Compute became more distributed with redundancy built in and ownership was spread across a broad set of actors.
- **Tensions:** A distributed and redundant system is expensive to operate which raise questions of economic sustainability.

5.3 Delphi study: round two

For round two the authors chose to focus on the scenarios described in the prior section. Round two was split into five parts: The first asking whether the Efficiency Pipeline is a good characterization of today's ecosystem and the other four sections, one for each future scenario, asking whether the scenario is plausible, what things need to happen for the scenario to materialize, what consequences follow given the scenarios and finally which roles described in section 4.4.2 will be winners and losers. Round two can be found in appendix A.3. There were a total of 14 respondents, resulting in an answer rate of 61 percent from the first round having a similar split between the subgroups. Each respondent can mention multiple themes and consequences seen in the tables of this section.

5.3.1 Today's baseline: efficiency pipeline

The panel was largely in agreement that the efficiency pipeline is a good characterization of today's N4AI ecosystem, see figure 5.10. Where the panel was not in agreement was that it was a more relevant description of the AI-part of the ecosystem but not the telecom ecosystem. Here one expert emphasized that energy cost or operating expenses is not a constraint as long as the business is sustainable:

> "*It does not matter if energy or other opex parameters are increasing as long as the business is sound (in general it will always increase).*"

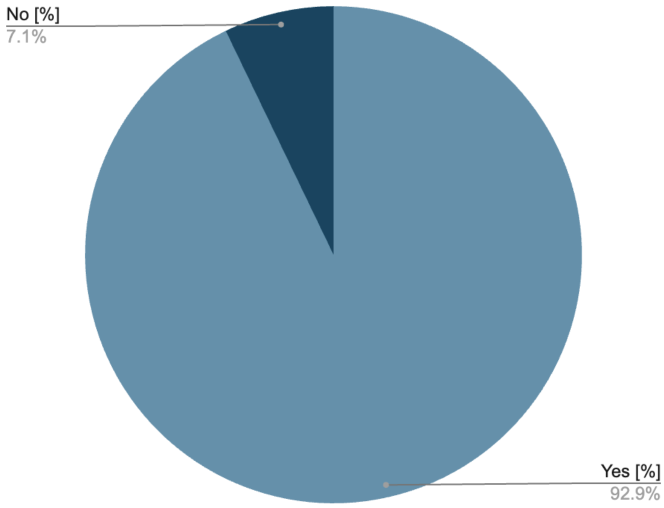

Figure 5.10 "Is The Efficiency Pipeline a fair characterization of the current AI and telecom ecosystem?" n = 14

5.3.2 Hyperscaler stranglehold

When asked if the Hyperscaler stranglehold is a viable scenario only one expert answered no, see figure 5.11. This is the scenario with the strongest consensus for being plausible at almost 93 percent.

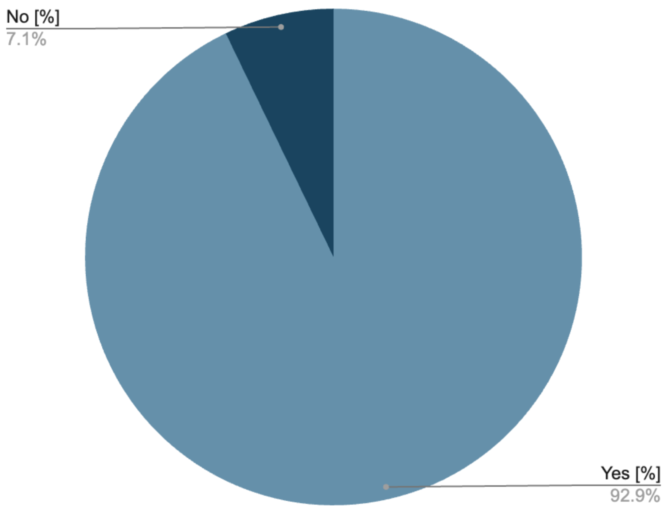

Figure 5.11 "Do you think The Hyperscaler Stranglehold is a plausible scenario?" n = 14

What needs to happen for the scenario to materialize?

This was the open question with the most responses, nine out of 14, producing the broadest spread of themes as seen in table 5.3. The most cited view was that this outcome follows if nothing changes: regulators need to act, or some external shock needs to shift power away from hyperscalers. One such thing could be regulations against hyperscalers. A second condition is the hyperscalers' ability to extend their control beyond their own data centers. If developers stay locked into hyperscaler ecosystems as compute moves to the edge, platform control follows the developers, not the hardware. That said, CSPs and telecom equipment vendors are not passive in this dynamic. Whether they can claim a meaningful position in the new ecosystem is a separate question, and respondents were skeptical, as one put it:

> *"CSPs will need to do what they do best, i.e. missing opportunities."*

Table 5.3 Themes obtained when expert panel was asked what needs to happen for the Hyperscaler Stranglehold to materialize. n = 11

What needs to happen	Quotes from panel	Mentions
Nothing changes, this is the current trajectory	“Nothing, this is the most likely scenario”; “... the default outcome if regulators don’t act”; “... not so different from the situation today... "the usual" needs to happen”; “This is the default option unless other major changes occur.”	4
An absence of regulation against hyperscalers	“No strong regulations against the hyperscalers”; “Regulation pushes sovereignty/diversification without breaking platform control...”; “... default outcome if government and regulators does not to force an alternative outcome.”	3
An extension of the hyperscalers control to the edge	“Hyperscalers successfully extend their control planes to the edge (same APIs/tooling/ops model at telco/metro edge so developers stay "inside" AWS/Azure/Google ecosystems)”; “Hyperscalers extending their edge closer to customers.”; “Hyperscalers/model providers successfully extend orchestration outward...”	3
A failure for CSPs to capture value	“CSPs will need to do what they do best, i.e. missing opportunities.”; “Telco value-capture attempts (e.g., network APIs) don’t mature fast enough...”	2
AI/data centralization reinforcing hyperscaler control	“AI tokenization could be the new monetization, this will inevitably need to stay in the data centers to maintain control.”; “Market concentration persists in cloud + AI platforms (dominant players keep scale advantages, reinforcing their ability to "own" orchestration even when infrastructure spreads out).”	2

Which are the major consequences given this scenario?

The hyperscaler stranglehold would according to the panel be detrimental to CSPs and telecom equipment vendors, figure 5.12 shows how the panel voted on winners and losers in the scenario, placing CSPs and telecom equipment vendors at the very bottom. Their greatest risk would be losing their strategic role. Geographic concentration of control is a second consequence. Some experts noted that economic value would flow out of Europe to the US. It would thereby pose a sovereignty and concentration risk. A slower market development due to CSPs not having reasons, or the money, to upgrade their networks since the value capture is in the hands of hyperscalers.

Table 5.4 Themes obtained when expert panel was asked about the major consequences of the Hyperscaler Stranglehold scenario. n = 10

Consequence	Quotes from panel	Mentions
CSPs margins compress and they lose their strategic role	"CSPs locked into a model of high cost and very low margins..."; "CSPs providing their real-estate but remain focused on connectivity."; "CSPs/operators and vendors face continued margin compression"; "Further margin compression in connectivity"	4
Europe's sovereignty decreases	"A high concentration and dependency on a few dominant non European player would for European countries be a problematic outcome both from a security perspective and from an economic perspective while most of the value would end up in the US."; "Sovereignty/regulatory friction persists rather than resolves..."; "Excessive power concentration on hyperscalers. Significant risk in failure situations."	3
Slow market development	"Slow market uptake"; "CSPs will not have business reasons to make major investments in their networks, so we will see a stagnation when it comes to new functionalities and capacity."	2
Value capture remains at orchestration/control-plane layer	"Value capture stays concentrated at the orchestration/control-plane layer, even with distributed edge hardware."; "A high concentration and dependency on a few dominant non European player would for European countries be a problematic outcome both from a security perspective and from an economic perspective while most of the value would end up in the US."	2

Who will be winners and losers?

No surprise that hyperscalers are winners in the hyperscaler stranglehold but it is interesting to see how the panel have ranked the other roles in the N4AI ecosystem.

There is an implicit assumption that AI model providers will be winners if hyperscalers stay in control and relative losers: infrastructure providers, CSPs/Operators and telecom equipment vendors, all operate in large part at the hardware level. It is also interesting to note that consumers and regulators/governments are voted as losers even though they have the power to change their fate, this can thereby be interpreted that staying idle makes both consumers and governments losers.

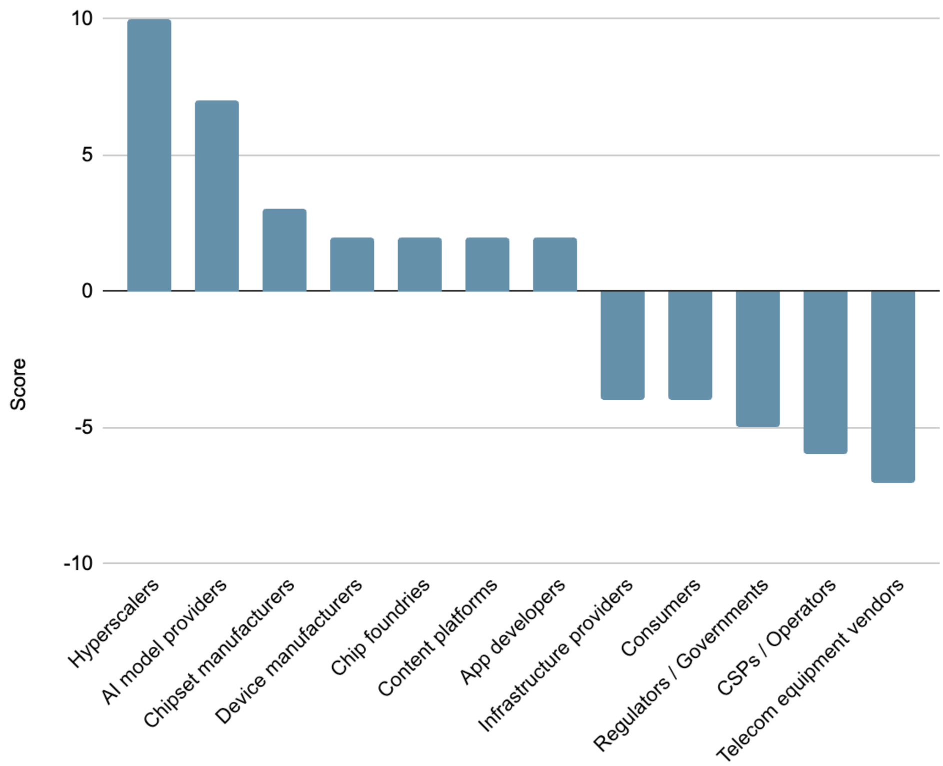

Figure 5.12 Winners and losers in the Hyperscaler Stranglehold scenario.

5.3.3 Resilient fortress

The resilient fortress differs sharply in plausibility from the hyperscaler stranglehold even though the only change is in system focus. Almost 70 percent state it as a plausible scenario and is thereby considered just below strong consensus, see figure 5.13.

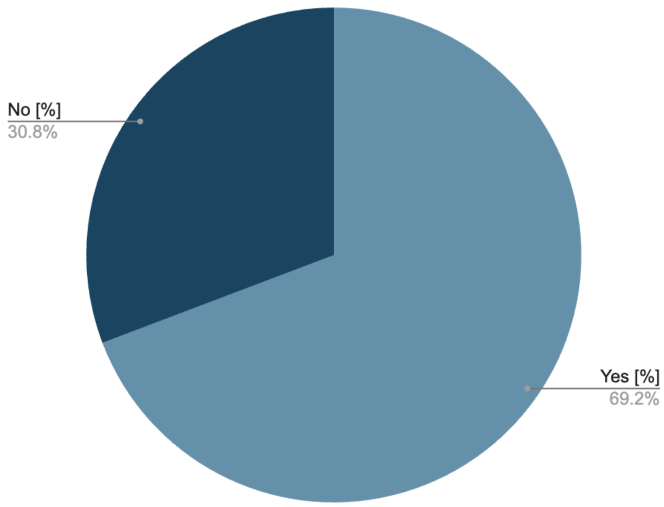

Figure 5.13 "Do you think The Resilient Fortress is a plausible scenario?" n = 13

What needs to happen for the scenario to materialize?

This question had a lower answer rate at only five respondents, there were however still some trends. The strongest being regulation that mandates resilience, implying this is nothing that hyperscalers will do unless forced to. It was also raised by one of the respondents that regulators would need to trust hyperscalers and AI model providers for the scenario to materialize. Another trend was resilience becoming a product for customers that need distributed and resilient compute.

Table 5.5 Themes obtained when expert panel was asked what needs to happen for the Resilient Fortress scenario to materialize. n = 5

What needs to happen	Quotes from panel	Mentions
Regulation that mandates resilience	“Regulators hard-code "continuity" into compliance.”; “Standardized resilience/security attestations”; “Governments demand data sovereignty...”	3
Resilience as a product	“Resilience becomes a differentiated, auditable product.”; “Successful B2C and B2B and critical applications that need distributed infrastructure with resiliency.”	2
Regulators place trust in hyperscalers and AI model providers	“Only if governments and regulators trust the hyperscalers and AI model vendors...”	1

Which are the major consequences given this scenario?

Given the high influence of few ecosystem players in this scenario it was mentioned by three of the respondents that concentration is an increased risk. Even though there is more focus on resilience the fact that a few large companies control that resilience makes it less so. It follows that competition and innovation stagnates in an environment with few large players is controlling the ecosystem. Lastly as a result of increased regulation mandating resilience, government control increases which in turn increases process overhead.

Table 5.6 Themes obtained when expert panel was asked about the major consequences of the Resilient Fortress scenario. n = 8

Consequence	Quotes from panel	Mentions
Concentration is an increased risk	"Reliance on major technology providers."; "Compute distributes, control doesn't. Systemic dependency risk rises."; "a small set of providers can afford the compliance and resilience engineering burden, reinforcing their "trusted backbone" status."	3
Competition and innovation will decrease	"Costly and lower productivity - other systems will overtake leadership"; "no innovations"; "High concentration will not create the needed incentives for the players to invest in resilience and on distributing the compute capabilities."	3
Government control increases	"More governance + process overhead across the stack."; "Government control"	2

Who will be winners and losers?

The same winners as in the Hyperscaler Stranglehold dominate the Resilient Fortress: hyperscalers and AI model providers, see figure 5.14. CSPs are the biggest losers, while telecom equipment vendors rank four positions higher than in the Hyperscaler Stranglehold scenario. Regulators and governments gain influence through their expanded role in mandating resilience, making them relative winners, though that influence does not translate into better outcomes for consumers. Software players rank worse than expected, with content platforms and app developers ranked third and fourth from the bottom.

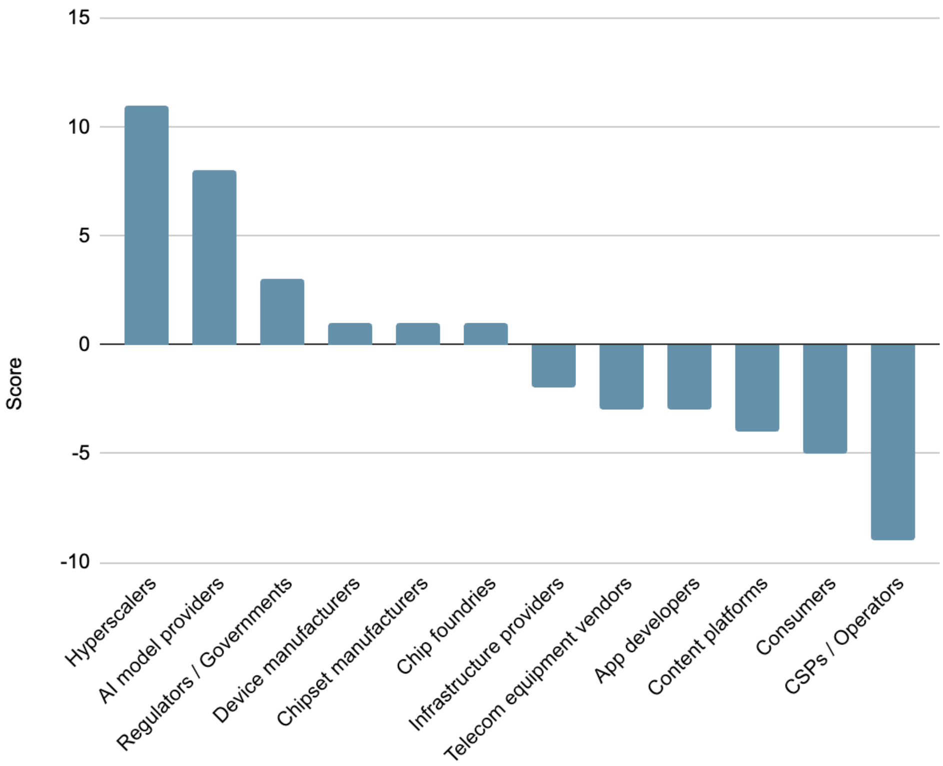

Figure 5.14 Winners and losers in the Resilient Fortress scenario.

5.3.4 Distributed marketplace

The distributed marketplace is the scenario with the lowest consensus having only 54 percent of the panel voting it as a plausible scenario as seen in figure 5.15.

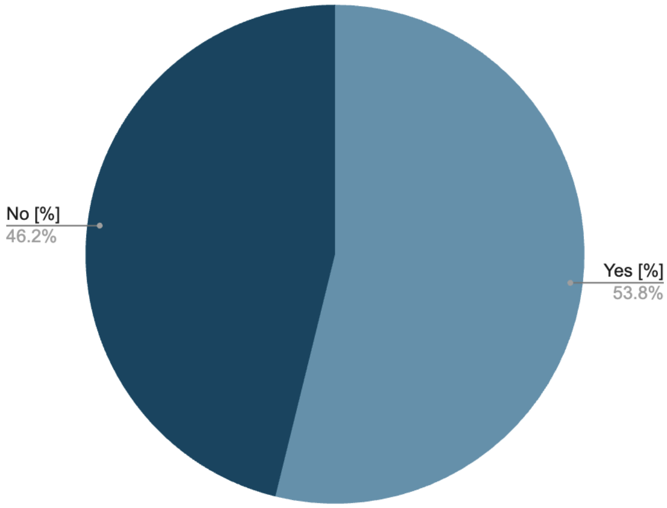

Figure 5.15 "Do you think The Distributed Marketplace is a plausible scenario?" n = 13

What needs to happen for the scenario to materialize?

Three trends each received two mentions, see table 5.7. The first is a multi-vendor environment, where edge infrastructure would need to be distributed across many actors. Working against this is the efficiency advantage larger players hold in integration and operations, a structural barrier that may partly explain why the scenario was rated as less plausible. Lowering costs/total cost of ownership (TCO) and increasing interoperability were identified as the conditions needed to make a multi-vendor environment viable.

Table 5.7 Themes obtained when expert panel was asked what needs to happen for the Distributed Marketplace scenario to materialize. n = 5

What needs to happen	Quotes from panel	Mentions
Multi vendor environment that prevents lock-in	"Edge build-out spreads to many owners (CSPs + regional providers + vendors) using COTS/cloud-native stacks"; "Multi-vendor environments expand despite an ongoing integration/operations "tax", keeping competition intense while preventing any single actor from rebuilding outsized margins through lock-in."	2
Cost/TCO needs to dominate	"Procurement remains dominated by cost/TCO logic..."; "Cost efficiency being dominant and coming before value in a distributed scenario."	2
Interoperability and standardization	"Interoperability/standardization improves enough to lower switching costs... making the infrastructure layer more substitutable and therefore more price-competitive"; "Workloads become more portable and contestable"	2

Which are the major consequences given this scenario?

Consequences from the efficient multi vendor environment pressing costs would be the risk of investments being reduced leading to under funding which lowers the innovation rate. Another that ecosystem complexity would increase worsening overhead. One benefit from this environment is listed by two respondents as being good for consumers, getting better services through increased competition, see table 5.8.

Table 5.8 Themes obtained when expert panel was asked about the major consequences of the Distributed Marketplace scenario. n = 7

Consequence	Quotes from panel	Mentions
Underinvestment risk	“thin margins reduce the ability to fund major upgrades”; “if financial sustainability is not in place, it will simply not happen”; “price competition intensifies –> ‘good enough’ offerings win”	3
Operational complexity increases across the ecosystem	“Operations and integration become a bigger tax (multi-vendor = more overhead)”; “increased complexity”	2
Competition and innovation will decrease	“Stagnation of innovation”; “This seems to be a lose-lose case for the "big guys". Both hyperscalers and CSPs potentially lose by being under extreme pressure. ”	2
Consumers benefit	“May be good for consumers”; “Consumers will benefit from more competition and hopefully better services”	2

Who will be winners and losers?

The only player clearly voted as a relative loser in this environment is the hyperscalers. The biggest winners are app developers and consumers, followed by regulators and governments. Telecom equipment vendors and CSPs are neither considered big winners or losers in this scenario.

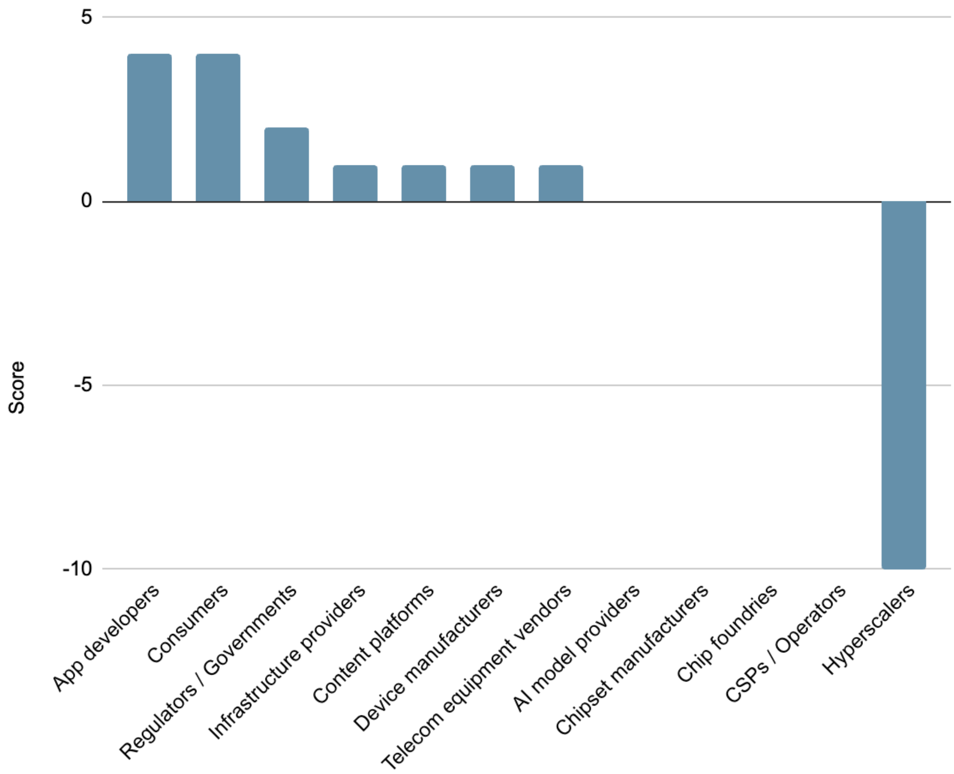

Figure 5.16 Winners and losers in the Distributed Marketplace scenario.

5.3.5 Resilient mesh

The resilient mesh was considered the second most plausible scenario having approximately 71 percent voting yes, see figure 5.17. Interestingly this was not the scnario voted most likely since from the first round 61 percent of the panel answered that they think the ecosystem would move toward more resilience and 91 percent answered that it would move toward a more distributed compute. Here is an interesting insight, the panel thinks that a larger emphasis will be put on resilience but not so that the ecosystem shifts entirely toward it. More resilient here does not imply that the primary focus will be resilience.

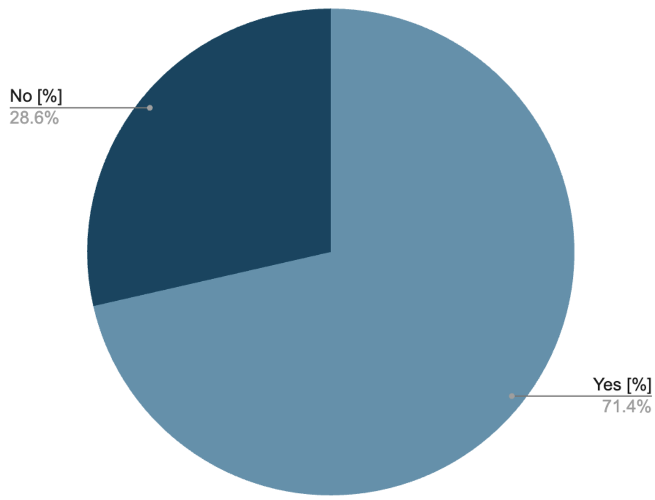

Figure 5.17 "Do you think The Resilient Mesh is a plausible scenario?" n = 14

What needs to happen for the scenario to materialize?

The Resilient Mesh, together with the Hyperscaler Stranglehold, was those who had the most answers for the qualitative questions, increasing the number of trends as seen in table 5.9. In the Resilient Mesh the most important trends is coordination and interoperability between ecosystem players since the ecosystem *"...collapses under coordination costs..."* if there is not an open, cross-domain collaboration. Another prediction is that for the Resilient Mesh to materialize open source models will have to win. Policy needs to reward resilience being pressured geopolitically, sovereignty becomes more important. There is also a trend mentioned by two respondents considering economic incentives to fund redundancy such as successful use cases that require resilience and customers willing to pay extra for it. One theme only mentioned by one respondent but which is highly relevant for CSPs and telecom equipment vendors is that in this environment CSPs would need to change how they operate to succeed. They are however, according to the panel, resistant to doing that.

Table 5.9 Themes obtained when expert panel was asked what needs to happen for the Resilient Mesh scenario to materialize. n = 7

What needs to happen	Quotes from panel	Mentions
Coordination and interoperability	"a fragmented ecosystem collapses under coordination costs unless interoperability becomes real."; "Federation standards for cross-domain operation..."; "open source models win"; "Standardization or defacto open collaboration across value chain."	4
Regulation rewarding resilience	"Policy + procurement that explicitly rewards diversification/resilience"; "... additional measures such as regulatory limitations in combination with requirements / limitations in public procurements will probably be needed..."; "fragmentation won't happen at scale purely from technical preference; it typically needs policy, procurement, and compliance that reward multi-provider and in-country capability."	3
Geopolitical pressure and sovereignty demand	"The driver for this scenario would be continued worsening geopolitical tension and especially for Europe a will to achieve increased sovereignty."; "For redundancy to be treated as value, buyers must consistently experience (or credibly fear) outages/shocks where the cost of interruption dominates marginal efficiency."	2
Economic incentives to fund redundancy	"Clear economic incentives to pay for redundancy (e.g., resilience SLAs, risk-based financing)"; "Highly successful use cases requiring distributed infrastructure"	2
A change in how CSPs operate	"CSPs will not succeed in this scenario if they don't change the way they operate, and we have seen once and again that they tend to be resistant to change."	1

Which are the major consequences given this scenario?

Consequences in the resilient mesh is a higher investment burden, that hyperscalers will be unwilling to give away control and that governance and coordination overhead is set to increase, all mentioned twice, see table 5.10. Another consequence, only mentioned by one respondent but relevant to consider as a regulator is that regions that force this scenario is set to have a competitive disadvantage.

Table 5.10 Themes obtained when expert panel was asked about the major consequences of the Resilient Mesh scenario. n = 8

Consequence	Quotes from panel	Mentions
There will be higher costs and investment burden	"Higher cost base (CapEx/OpEx) because redundancy and multi-location deployment become default rather than optional."; "There is an obvious risk that this scenario would slow down the implementation of AI in society a large and lead to higher cost of using AI in geographical areas that forces this solution."	2
Hyperscalers will resist giving away control	"Hyperscalers and Data Centers will want to hold onto the value - they will become the new centralized positions, connectivity will improve to reduce latency. But no way the hyperscalers will want to distribute control..."; "no way the hyperscalers will want to distribute control"	2
Governance and coordination overhead increases	"More coordination + governance overhead since cross-domain federation, access control, and orchestration across many actors become mandatory."; "increasing vendor-management complexity"	2
Regions forcing this model has a competitive disadvantage	"This could lead to lower economic growth and a competitive disadvantage in relations to other regions where a more cost effective implementation is allowed"	1

Who will be winners and losers?

This is the scenario where telecom equipment vendors and CSPs will have the largest relative gain. They are however not the biggest winners, those posed to win are AI model providers and infrastructure providers. This scenario has almost the opposite relative winners and losers compared to the Hyperscaler Stranglehold. In this scenario, mainly hardware comes out as winners and software as losers. Regulators and consumers are among the losers, see figure 5.18.

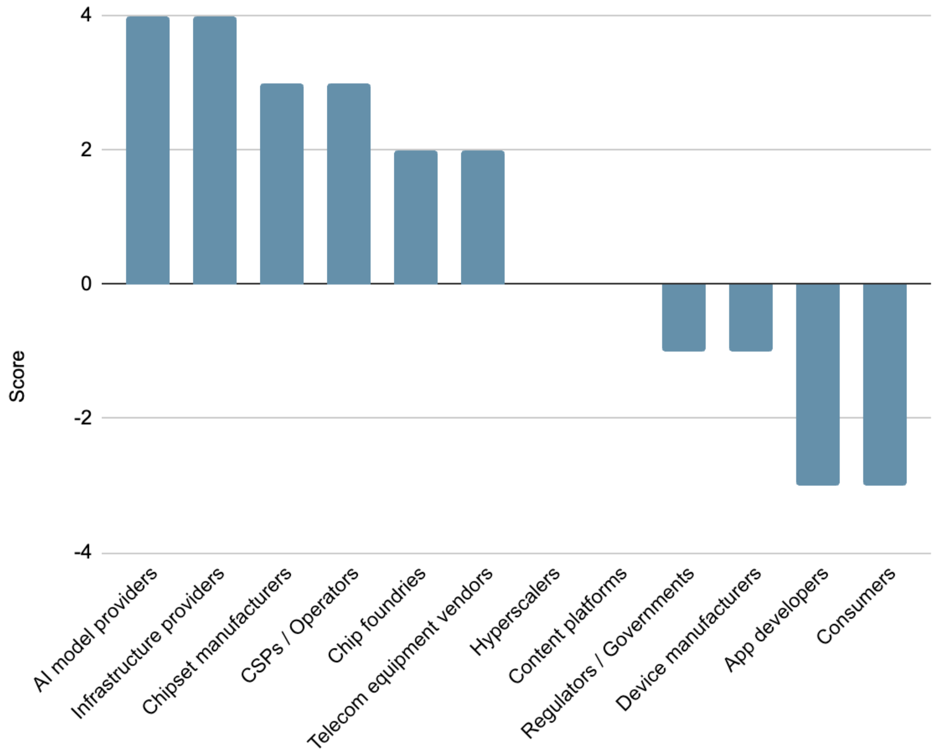

Figure 5.18 Winners and losers in the Resilient Mesh scenario.

5.3.6 Summary of round two analysis

In figure 5.19 a summary of the plausibility of the scenarios, according to the Delphi panel, are visualized. Table 5.11 lists the top three conditions that have to happen for the scenario to materialize as well as the top three consequences if it does. Table 5.12 lists the top three winners and losers in each scenario. For the Distributed Marketplace hyperscalers are the only ones listed as losers since it was the only role voted as a relative loser. Worth noting is that given our current trajectory, toward the hyperscaler stranglehold scenario, telecom equipment vendors and CSPs are considered to be the biggest losers. The only scenario where they were both voted as relative winners is in the resilient mesh, which is the scenario voted second most plausible (with a close third, see figure 5.19).

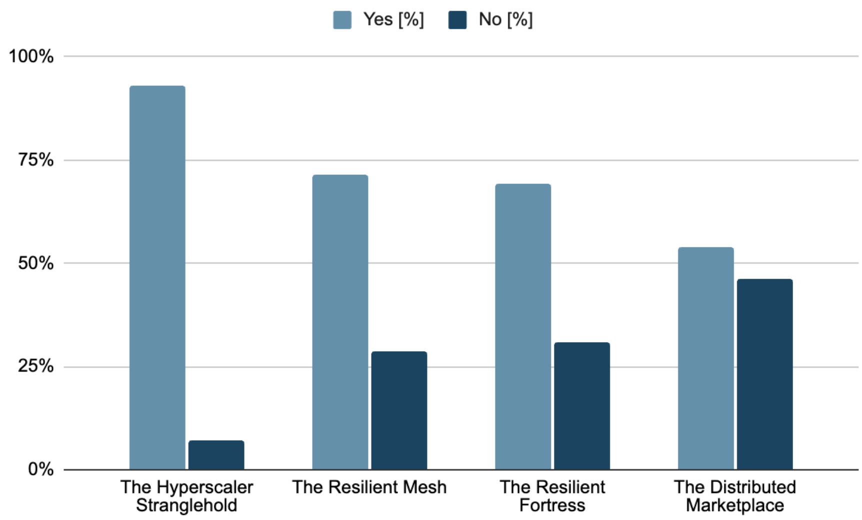

Figure 5.19 "Do you think this is a plausible scenario?"

Table 5.11 Summary of themes in the expert panel answers. Top three conditions for materialization, and consequences.

Scenario	What needs to happen	Consequences
Hyperscaler Stranglehold	• Nothing, this is the current trajectory • An absence of regulation against hyperscalers • An extension of the hyperscalers control to the edge	• CSPs margins compress and they lose their strategic role • Europe’s sovereignty decreases • Slow market development
Resilient Fortress	• Regulation that mandates resilience • Resilience as a product • Regulators place trust in hyperscalers and AI model providers	• Concentration is an increased risk • Competition and innovation will decrease • Government control increases
Distributed Marketplace	• Multi vendor environment that prevents lock-in • Cost/TCO needs to dominate • Interoperability and standardization	• Underinvestment risk • Operational complexity increases across the ecosystem • Competition and innovation will decrease
Resilient Mesh	• Coordination and interoperability • Regulation rewarding resilience • Geopolitical pressure and sovereignty demand	• There will be higher costs and investment burden • Hyperscalers will resist giving away control • Regions forcing this model has a competitive disadvantage

Table 5.12 Summary of the top three winners and losers in each scenario according to expert panel. In the distributed marketplace, hyperscalers were the only role voted as a net loser.

Scenario	Winners	Losers
Hyperscaler Stranglehold	• Hyperscalers • AI model providers • Chipset manufacturers	• Telecom equipment vendors • CSPs/Operators • Regulators & governments
Resilient Fortress	• Hyperscalers • AI model providers • Regulators & governments	• CSPs/Operators • Consumers • Content platforms
Distributed Marketplace	• App developers • Consumers • Regulators & governments	• Hyperscalers
Resilient Mesh	• AI model providers • Infrastructure providers • Chipset manufacturers	• Consumers • App developers • Device manufacturers

5.4 Findings from the two rounds

The individual round summaries in sections 5.1.6 and 5.3.6 report what each round found. This section aims to capture what the two rounds reveal when read together, specifically three cross-round patterns that become visible once synthesized. The first is regulation as a hidden precondition. Across all three alternative scenarios, some form of regulatory action is required for the scenario to materialize. The form differs in each scenario as *Resilient Fortress* requires mandated resilience requirements, *Distributed Marketplace* needs open standards, *Resilient Mesh* needs procurement policy that rewards multi-provider architecture. However, the pattern is the same, there is thus no scenario where CSPs or regional actors gain meaningful ground through market forces alone.

The second pattern is the decoupling of physical infrastructure from economic value. Round one flagged that hyperscalers were expected to hold their position even as compute distributed. Round two confirmed this concretely: in *Hyperscaler Strangle-*

hold, orchestration and platform control remain concentrated even as edge hardware expands; in *Resilient Fortress*, hyperscalers gain trusted backbone status precisely because of compliance demands. Who owns the compute matters less than who controls the orchestration layer. For CSPs, this is the structural problem.

The third pattern is data sovereignty. It was not prompted in round one but appeared organically in the qualitative data, alongside geopolitical fragmentation as a top macro driver. In round two it appeared as a condition for *Resilient Mesh* and a casualty of *Hyperscaler Stranglehold*.

6 Discussion

This chapter translates the study's findings into possible implications for the actors shaping the N4AI ecosystem. The discussion aims to ask what those findings mean for the strategic decisions facing CSPs and telecom equipment vendors. The chapter opens with the traffic shift and what it demands of network infrastructure, before turning to the strategic bind facing CSPs. It closes with an assessment of the study's trustworthiness.

6.1 The strategic bind for CSPs

6.1.1 The changing nature of traffic

The clearest point of convergence across the panel was the traffic shift: by 2035, network consumption will be driven more by machines than by humans. Today, mobile data is dominated by video streaming, producing an asymmetric 10/90 UL/DL ratio, as shown in section 5.1.2.1. AI agents browsing, querying, and interacting with real-time data sources will change that balance considerably.

The panel broadly expected uplink to reach at least 20 to 30 percent of total traffic by 2035. The scale of that growth depends largely on where inference happens. If foundation models stay centralized in hyperscaler clouds, robots and extended reality hardware will need to stream raw data continuously upstream. If inference moves to the network edge, local processing reduces that load substantially. The inference location question was one of the sharpest unresolved variables going into round two and remains open, since it is impossible to say with certainty whether inferencing will be efficient enough to run on device or at a base station within ten years.

Uplink growing from 10 percent to an estimated 30 percent breaks the engineering assumptions embedded in current network design. Antenna configurations, spectrum planning and network architecture will need to adapt to monetize the new traffic profile, increasing engineering complexity in the process. Successfully adapting to these changes is what will allow equipment vendors and CSPs to maintain their existing relevance.

The shift also introduces machine-generated traffic patterns, AI agents and physical AI, that differ structurally from human-generated traffic. Bursty, latency-sensitive signals will make flat-pricing and static capacity planning obsolete, pushing CSPs to rethink how they dimension and price their networks. For telecom vendors, this means designing infrastructure for a traffic profile that does not yet dominate but, as the findings suggest, probably will. The product roadmap question is whether to optimize for today's downlink-heavy traffic or tomorrow's more symmetric load, and when to make that shift.

6.2 Implications for telecom equipment vendors

According to the analysis done, telecom equipment vendors are identified as losers in the most probable scenario. Vendors defensible position lies in the interoperability and standards layer, specifically 6G architecture and edge orchestration. However, the chances are in the vendors' favor as long as they act before hyperscalers extend their control planes to the edge. As of today, the edge is open territory with no established long-time standard, which is why this moment is the most critical for equipment vendors like Ericsson. The actors who show up now in the standards bodies are the ones who will write the rules everyone else has to follow. The window for shaping the future standards is time-limited, early engagement in standards bodies (3GPP, ETSI, O-RAN Alliance) is a precondition for relevance in 2035.

6.3 Data sovereignty as a business opportunity

Data sovereignty was not explicitly prompted in the initial survey. It emerged on its own as a structural driver, which makes it one of the more credible findings. Experts across the panel flagged that geopolitical fragmentation and data privacy requirements will push back against the cost-efficiency logic that currently favors hyperscalers.

For CSPs, being locally regulated transitions from a compliance burden to a competitive differentiator in a world of geopolitical fragmentation. What needs to be true for CSPs to capture this opportunity is active engagement with public procurement requirements, investment in sovereign AI infrastructure and credible data governance frameworks instead of passive presence. In this case, CSPs could capture value if regulations require that data processing be carried out within the country where the data originates.

7 Conclusions

This chapter answers the research questions, reflects on the fulfillment of purpose and suggests future research building on the findings.

7.1 Answering the research questions

7.1.1 Answering RQ1

How will AI usage impact mobile network traffic until 2035?

The mobile network traffic in 2035 is set to become highly variable and nonlinear making it difficult to predict. Data sent over mobile networks is estimated to grow about five to seven times from the levels of today, from 200 EB/month to 1000-1500 EB/month. Uplink is still set to be a minority but will grow from 10 percent to around 30 percent of total traffic. The number of 4G and 6G connections is set to be about the same at two billion each and 5G connections at about five billion. The technologies driving the growth in traffic will be a mix of AI enabled technologies, contributing mostly to uplink, as well as continued growth in video streaming contributing to downlink.

These AI enabled technologies have been identified as: AI Assistants, Software Based Agentic AI, Internet of Things, Smart Cities, Healthcare AI, Autonomous Mobility, Autonomous Robots, Extended Reality, and Cloud Gaming. All having their own unique connectivity needs. According to the expert panel, extended reality, cloud gaming and autonomous mobility are the three technologies that will contribute the most to added traffic in 2035 and the technologies having the highest adoption will be AI assistants, software based agentic AI and cloud gaming. Technologies with the highest revenue potential is software based agentic AI, AI assistants and autonomous mobility.

Another insight from the panel is that AI generated traffic should be split into two categories: software driven AI traffic and physical AI traffic. Where software driven AI is set to generate the most new traffic until 2030 and physical AI driving the most new traffic in the 2030-2035 time span. Figure 7.1 gives a summary of the answer of RQ1.

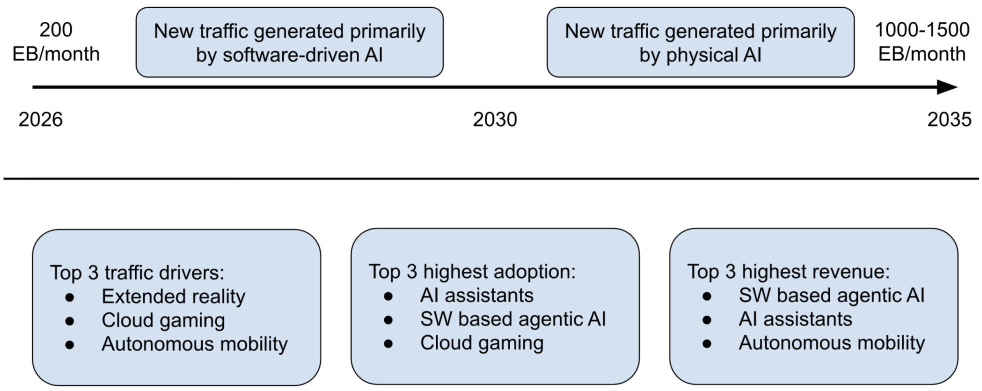

Figure 7.1 Summary of the answer to RQ1.

7.1.2 Answering RQ2

What will the Network for AI ecosystem look like in 2035?

7.1.2.1 Ecosystem

The ecosystem is expected to consist of the roles listed below with hyperscalers, consumers and AI model providers as those with the most influence. Telecom equipment vendors, CSPs and regulators/governments are those with the least influence. For the full ranking view figure 7.2. The factors influencing the ecosystem the most is expected to be infrastructure investment flows and geopolitical fragmentation, see figure 7.3. The three axes described in section 4.6.1 are influenced differently by each of the factors. Data privacy, compute efficiency and geopolitical fragmentation drives compute topology the most toward being more decentralized while none of the identified factors drive compute topology toward being more centralized. Geopolitical fragmentation strongly drives the economic value capture toward being more decentralized, none of the factors listed push toward a more concentrated value capture. Regarding the system focus, geopolitical fragmentation and data privacy are the two who push the most toward a more resiliency oriented focus while infrastructure investment flows and compute efficiency push toward more cost efficiency focus.

- **Hyperscalers:** Large cloud providers that operate massive-scale data center infrastructure and offer compute, storage, and AI services globally. (Amazon web services, Azure, Google)
- **Telecom equipment vendors:** Companies that design and supply the hardware and software that builds and runs telecom networks. (Ericsson, Nokia, Huawei)
- **CSPs / Operators:** Network operators that provide connectivity services to consumers and enterprises. (Vodafone, Deutsche Telekom, etc.)

- **AI model providers:** Companies that develop and deploy foundational AI models, made accessible via APIs or direct integration. (OpenAI, Anthropic, Open Weight, etc.)
- **Device manufacturers:** Companies that produce end-user hardware such as smartphones, laptops, and wearables. (Apple, Samsung, etc.)
- **Chipset manufacturers:** Companies that design the processors and chips powering devices, networks, and AI workloads. (Qualcomm, Apple, NVIDIA)
- **Chip foundries:** Facilities that fabricate chips at scale based on designs from chipset manufacturers. (TSMC, Intel, etc.)
- **App developers:** Companies or individuals that build software applications running on top of existing platforms and infrastructure. (Google (apps), Zoom, Uber, etc.)
- **Content platforms:** Services that produce or distribute digital content such as video, music, or games to end users. (Netflix, Meta, TikTok, etc.)
- **Infrastructure providers:** Companies that own and lease shared physical assets such as cell towers, data centers, and fiber networks to multiple tenants like telecom operators or hyperscalers. (Equinix, American Tower, Cellnex, etc.)
- **Regulators / Governments:** Public bodies that set the legal, policy, and spectrum frameworks within which the ecosystem operates.
- **Consumers:** End users who consume AI-powered services, connectivity, and devices.

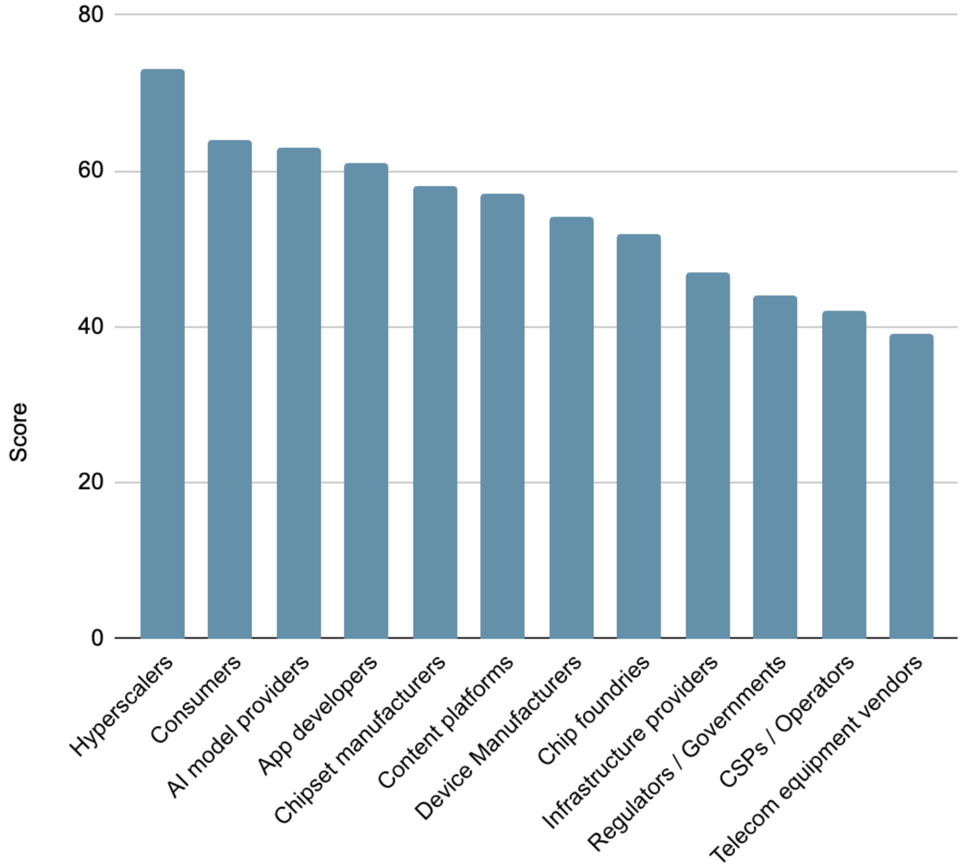

Figure 7.2 Expert panel ranking of roles with the most to least influence in the 2035 N4AI ecosystem. For an explanation of score see section 2.3.1.1

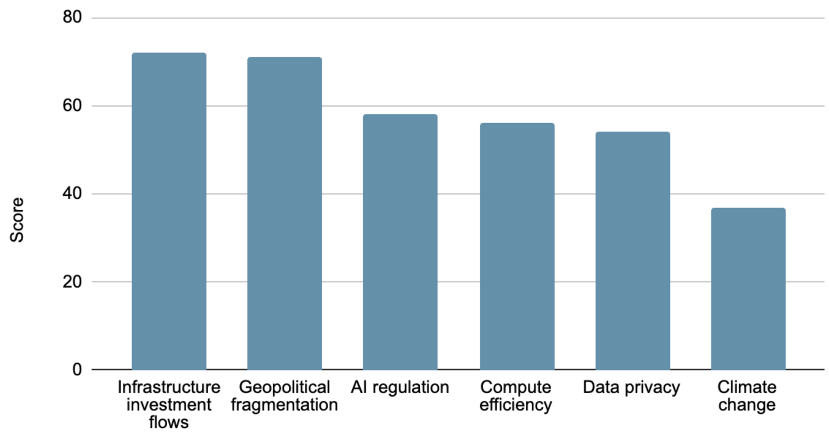

Figure 7.3 Expert panel ranking of factors of influence in the 2035 N4AI ecosystem. For an explanation of score see section 2.3.1.1

7.1.2.2 Scenarios

According to the expert panel the most plausible scenario is the hyperscaler stranglehold at 93 percent, holding a very strong consensus validating the scenario. The resilient mesh and the resilient fortress is 71 and 69 percent respectively, both having a moderate consensus. The Distributed Marketplace was the least convincing scenario, with only 54 percent of the panel rating it as plausible. This makes the consensus weak both on whether it is plausible or not plausible, see figure 7.4. Below is a description of what signifies the scenarios considered plausible, what needs to happen for the scenarios to materialize, what consequences follow from the scenario materializing and finally which ecosystem roles are set to be winners and losers.

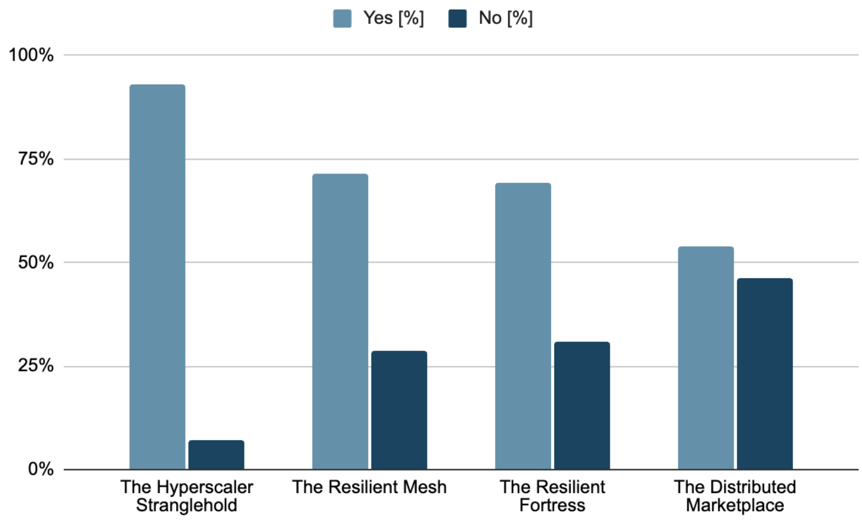

Figure 7.4 "Do you think this is a plausible scenario?"

Hyperscaler Stranglehold

More distributed compute | *Concentrated value capture* | *Cost efficiency focus*

The hyperscaler stranglehold is signified by further concentration of influence in the N4AI ecosystem having, unsurprisingly, hyperscalers as the biggest winners, followed by AI model providers and chipset manufacturers. Many of the experts express that this is the trajectory that the N4AI ecosystem is on and will materialize if nothing changes. Hyperscalers will also have to take control of edge computing and avoid regulation that might stop them gets put in place. Consequences for CSPs are dire since they risk to lose their strategic role in the ecosystem and when they lose, telecom equipment vendors also lose. Implications for the hyperscaler stranglehold are summarized in table 7.1.

Table 7.1 Summary of implications of the hyperscaler stranglehold.

What needs to happen	• Nothing, this is the current trajectory • An absence of regulation against hyperscalers • An extension of the hyperscalers control to the edge
Consequences	• CSPs margins compress and they lose their strategic role • Europe's sovereignty decreases • Slow market development
Winners	• Hyperscalers • AI model providers • Chipset manufacturers
Losers	• Telecom equipment vendors • CSPs/Operators • Regulators & governments

Resilient Mesh

More distributed compute | More distributed value capture | More resilience focus

The resilient mesh diverges the most from today moving away from the baseline in all three axes. Changes in the N4AI ecosystem is considerable and the experts predicts that much more influence will move to hardware companies like infrastructure providers, telecom equipment vendors and CSPs. What will need to happen for this to materialize is that a sovereignty demand, forced by governments, pull power away from global players, like hyperscalers, into the hands of more local ones. This forces coordination and interoperability, leading to higher costs and investment burdens. However, regions forcing this will have a competitive disadvantage. Implications for the resilient mesh are summarized in table 7.2.

Table 7.2 Summary of implications of the resilient mesh.

What needs to happen	• Coordination and interoperability • Regulation rewarding resilience • Geopolitical pressure and sovereignty demand
Consequences	• There will be higher costs and investment burden • Hyperscalers will resist giving away control • Regions forcing this model has a competitive disadvantage
Winners	• AI model providers • Infrastructure providers • Chipset manufacturers
Losers	• Consumers • App developers • Device manufacturers

Resilient Fortress

More distributed compute | Concentrated value capture | More resilience focus

The resilient fortress is similar to the hyperscaler stranglehold, the difference is that the system focus has shifted toward resilience rather than cost efficiency. Winners are hyperscalers, AI model providers and governments. Governments and regulators win because they force control over the ecosystem. Content platforms, consumers and CSPs are predicted to lose the most in this scenario. For this to happen, regulators will need to mandate resilience at the same time as they place trust in hyperscalers. This leads to an increased concentration risk and the rate of innovation is predicted to decrease. Implications for the resilient fortress are summarized in table 7.3.

Table 7.3 Summary of implications of the resilient fortress.

What needs to happen	• Regulation that mandates resilience • Resilience as a product • Regulators place trust in hyperscalers and AI model providers
Consequences	• Concentration is an increased risk • Competition and innovation will decrease • Government control increases
Winners	• Hyperscalers • AI model providers • Regulators & governments
Losers	• CSPs/Operators • Consumers • Content platforms

7.2 Fulfillment of purpose

The purpose of this master's thesis was to analyze how the introduction of AI traffic would impact the ICT industry by 2035. It aimed to provide a prognosis of network traffic, scenarios that could play out, and how the ecosystem might change in each scenario. The goal was not to predict a single correct future but to map the space of plausible ones and identify what conditions would lead to each scenario.

Multiple AI and connectivity enabled technologies were identified that would drive AI traffic and a projection of traffic was done using external sources. An expert panel was used to validate the projection of what mobile traffic might look like in 2035.

Ecosystem roles were identified and mapped using market ties and an ecosystem pie model. Factors of influence that is set to have a large impact on the N4AI ecosystem was identified. These factors were thereafter validated and ranked by an expert panel. The authors constructed eight distinct scenarios using three independent axes that the N4AI ecosystem could move back and forth along. Thereafter an expert panel determined which movements along the axes were the most likely. This resulted in four 2035 scenarios being highlighted as possible futures, the authors then prompted the expert panel once more. This resulted in one scenario with a consensus of 93 percent, two with a consensus of 70 percent and one with no consensus splitting the panel at about 50 percent. Additionally, for each scenario: winners, losers, consequences and things that need to happen for the scenario to materialize were identified.

7.3 Further reflections

7.3.1 Critique of the Delphi method

The method and research design were selected to match the ten-year horizon, where quantitative data on an emerging phenomenon like N4AI is inherently scarce. As discussed in the theory section, the goal was not to identify a single N4AI trajectory but to generate a set of plausible futures. Because the Delphi method asks experts to respond independently before seeing others' views, convergence and divergence became visible without anchoring effects distorting the early rounds. Presenting the scenarios built by the panel from round one, in round two, confirmed where consensus was strong, such as the shift toward a more distributed compute and where genuine uncertainty remained, particularly around who captures the economic value.

7.3.2 Panel composition and bias

The composition of the expert panel shapes the scenarios and this demands honest scrutiny. With 44 percent of respondents from academia, the panel was skewed toward technical expertise. This gave the study a solid grounding in infrastructure constraints and architecture trajectories-particularly the shift toward uplink-heavy traffic but it left a gap. CSPs were largely absent as primary respondents, so questions about operator business models, monetization, and value capture were not captured. The commercial realities of running an AI-driven network may be underrepresented compared to what practitioners inside those organizations would say.

7.3.3 Reflection on the scenario cross

Lindgren and Bandhold's scenario cross worked well for this kind of high-uncertainty problem, by orthogonalizing the axes: Compute topology, value capture and system focus, the method forces distinct, mutually exclusive futures rather than variations on a single trajectory. The obvious tradeoff is that a multi-actor ecosystem gets flattened onto two-dimensional planes. The transition from cost efficiency to resilience, for instance, is rarely binary in practice. But the forced choice is part of the point, it surfaces tensions, like those between geopolitical fragmentation and infrastructure economics, that a more gradual framing would obscure.

7.3.4 Reflection on ecosystem

When mapping the power dynamics of the current N4AI ecosystem, this study relied heavily on current revenue, as discussed in section 4.4.2. Revenue is a deliberately conservative lens: it fails to capture future expectations or forward-looking market valuations, but it shows where real money flows today, cutting through the AI hype that inflates many analyses of this space.

However, relying on current revenue alone may obscure structural shifts already in motion. A more complete picture would also track CAPEX and investment flows. Hyperscalers and cloud providers are already committing massive sums-projected at $5.2 trillion in AI-related data center capacity by 2030, as discussed in section 4.5. CAPEX trends might offer a leading indicator of where ecosystem dominance is heading that current revenue figures cannot capture.

7.3.5 Reflection on scenario building

7.3.5.1 The tension between cost efficiency and resilience

The most persistent tension in the scenario building was between cost efficiency and resilience. Two goals that pull against each other structurally. Today's ecosystem runs on an "Efficiency Pipeline" logic, focused on scale and cutting operational costs. The panel shifted this: 61 percent of experts predicted the system focus will move toward more resilience by 2035. And yet the *Hyperscaler Stranglehold* still drew the strongest consensus. The tension lies precisely there, most experts expect resilience to matter more, but also expect the architecture to remain locked in by concentrated economic power.

7.3.5.2 Paths to scenario materialization

The empirical findings point to regulatory intervention and open standardization as the main triggers for any scenario other than the *Hyperscaler Stranglehold*. That tells you something: market forces on their own are not enough to shift the power dynamics.

The inadequacy of organic market forces and technological lock-in
The strong consensus around the *Hyperscaler Stranglehold* as the default trajectory reflects how entrenched the hyperscalers' architectural position is. They control both the software ecosystems and the developer interfaces, which means that even as physical compute nodes push to the edge, customers and developers remain tied to proprietary orchestration platforms. The barrier to entry for alternative infrastructure providers is not merely high, it is structurally reproduced with each new deployment. This makes the panel's position, that regulatory mandates like data sovereignty requirements are a prerequisite for the *Resilient Fortress* or *Resilient Mesh* to emerge. Consistent with what the lock-in literature would predict. Dislodging a strong technological lock-in typically takes an external shock, whether geopolitical fragmentation or compliance legislation, rather than incremental competitive pressure.

Reshaping network ties for interoperability
A second path toward the more distributed scenarios, like the *Distributed Marketplace*, runs through open-source models and genuine interoperability. Applying Ahuja et al. (2012)'s network tie framework, this requires more than goodwill: overcoming the "integration/operations tax" that the panel identified cannot be done through competitive market dynamics alone. It requires CSPs and telecom vendors to form strategic partnerships and referential ties through standards development organizations. Without those collaborative structures enforcing cross-domain federation

and open standards, there is no structural reason to expect the ecosystem to land anywhere other than where the hyperscalers want it.

Timing and sequencing
The third dimension is timing, as Adner (2017) describes. Deploying edge nodes is not sufficient for CSPs to capture value if the architectural decisions have already been locked in by the time those nodes go live. The panel flagged this as a real risk: a passive CSP enters the distributed paradigm after the rules have been written. To claim a position in that future, traditional telecom actors need to act on their specific complementarities early (Jacobides et al., 2018), in particular, their role as trusted, locally regulated stewards of sovereign AI infrastructure and engage in standards development before hyperscalers extend their control planes to the edge.

7.4 Validity and reliability

This study addressed trustworthiness through four criteria: credibility, transferability, dependability and confirmability. Triangulating between the expert Delphi panel and secondary industry reports - analyst forecasts, operator disclosures, and standards documentation - strengthened credibility by grounding scenario elements in observable market evidence rather than expert opinion alone.

Dependability is harder to claim honestly here. The N4AI ecosystem is moving fast; the six months between this study's inception and its completion were not uneventful. AI capabilities, infrastructure investment patterns, and regulatory frameworks will probably shift and this thesis captures a particular moment rather than a stable state. Leading AI companies grew their revenues by orders of magnitude during the study period alone, a pace with no precedent in enterprise software history. The foundational tensions - between hyperscaler concentration and distributed alternatives, between cost efficiency and resilience - remain structurally sound.

7.5 Suggestions for future research

Throughout the work focus has been to define the most likely outcomes by starting with all possible outcomes and gradually removing those who are less likely. This approach has been a good way to answer the research questions but there are no concrete recommendations given to players within the ecosystem. There is great value in knowing what the future might look like and what to look out for that might push it toward one scenario or another. It is even more valuable to also know how to win in a given environment.

High value future research could therefore be to look at business models and value capture strategies in each of the defined scenarios. Every player in the today's ecosystem seeks to build products that are competitive but without knowing where the ecosystem is heading it is difficult to make informed decisions. This research lays the foundation of such product development or portfolio analysis. Analyzing value creation and value capture in ecosystems can be of interest. Another angle could be building business and revenue models specific to each scenario.

References

Adner, R. (2017). “Ecosystem as structure: an actionable construct for strategy”. *Journal of Management* **43**:1, pp. 39–58.

Ahuja, G., G. Soda, and A. Zaheer (2012). “The genesis and dynamics of organizational networks”. *Organization Science* **23**:2, pp. 434–448.

Alice Gomstyn, A. J. (2026). *What is a Smart City?* Accessed: 2026-03-22. URL: https://www.ibm.com/think/topics/smart-city.

Bellan, R. (Oct. 2025). *Sam Altman says ChatGPT has hit 800M weekly active users*. Accessed: 2026-02-13. URL: https://techcrunch.com/2025/10/06/sam-altman-says-chatgpt-has-hit-800m-weekly-active-users/.

Bernardo, V. (2024). *Extended Reality (XR)*. Accessed: 2026-03-22. URL: https://www.edps.europa.eu/data-protection/technology-monitoring/techsonar/extended-reality_en.

Bishop, P., A. Hines, and T. Collins (2007). “The current state of scenario development: an overview of techniques”. *Foresight* **9**:1, pp. 5–25.

Charlotte Hu Amanda Downie, M. F. (2026). *AI Agents vs. AI Assistants*. Accessed: 2026-03-22. URL: https://www.ibm.com/think/topics/ai-agents-vs-ai-assistants.

Chen, H., A. Mehra, S. Tasselli, and S. P. Borgatti (2022). “Network dynamics and organizations: a review and research agenda”. *Journal of Management* **48**:6, pp. 1602–1660.

Cloudflare (n.d.). *AI inference vs. training: What is AI inference?* Accessed: 2026-02-12. URL: https://www.cloudflare.com/learning/ai/inference-vs-training/.

Denscombe, M. (2017). *The Good Research Guide: For Small-scale Social Research Projects*. 4:1. Open University Press, Maidenhead. ISBN: 9780335241385.

Diamond, I., R. Grant, B. Feldman, P. Pencharz, S. Ling, A. Moore, and P. Wales (Apr. 2014). “Defining consensus: a systematic review recommends methodologic criteria for reporting of delphi studies”. *Journal of Clinical Epidemiology* **67**, pp. 401–409. DOI: 10.1016/j.jclinepi.2013.12.002.

Effah, E., Y. Hu, M. Shan, A. Chan, and Y. Le (Nov. 2016). “Application of delphi method in construction engineering and management research: a quantitative perspective”. *JOURNAL OF CIVIL ENGINEERING AND MANAGEMENT* **22**, pp. 991–1000. DOI: 10.3846/13923730.2014.945953.

Ericsson (2025a). *Ericsson Mobility Visualizer*. Accessed: 2026-03-22. URL: https://www.ericsson.com/en/reports-and-papers/mobility-report/mobility-visualizer?f=11&ft=1&r=1&t=8&s=4&u=3&y=2016,2025&c=5.

Ericsson (2025b). *Ericsson Mobility Visualizer*. Accessed: 2026-03-22. URL: https://www.ericsson.com/en/reports-and-papers/mobility-report/mobility-visualizer?f=15&ft=3&r=1&t=18&s=9,10,11,12,13&u=1&y=2025&c=4.

Ericsson (2025c). *Ericsson Mobility Visualizer*. Accessed: 2026-03-22. URL: https://www.ericsson.com/en/reports-and-papers/mobility-report/mobility-visualizer?f=10&ft=2&r=1&t=11,12,13,14,15,16,17&s=4&u=3&y=2025&c=4.

Ericsson (2025d). *Ericsson Mobility Visualizer*. Accessed: 2026-04-27. URL: https://www.ericsson.com/en/reports-and-papers/mobility-report/mobility-visualizer?f=1&ft=1&r=4,3,5,6,2,7,8,9&t=8&s=4&u=1&y=2025&c=4.

Ericsson (2025e). *Ericsson Mobility Visualizer*. Accessed: 2026-03-22. URL: https://www.ericsson.com/en/reports-and-papers/mobility-report/mobility-visualizer?f=10&ft=2&r=1&t=11,12,13,14,15,16,17&s=4&u=3&y=2025&c=4.

Ericsson (2025f). *GenAI Data Traffic Today*. Accessed: 2026-03-22. URL: https://www.ericsson.com/en/reports-and-papers/mobility-report/articles/genai-data-traffic-today-june-2025.

Ericsson (2025g). *Mobile Data Traffic Forecast*. Accessed: 2026-03-22. URL: https://www.ericsson.com/en/reports-and-papers/mobility-report/dataforecasts/mobile-traffic-forecast.

Ericsson (2025h). *Mobility Visualizer Data*. Accessed: 2026-03-22. URL: https://www.ericsson.com/en/reports-and-papers/mobility-report/mobility-visualizer?f=18&ft=2&r=1&t=1,2,19&s=14&u=4&y=2016,2025&c=6.

Ericsson (2026). *The Network for AI Experiences*. Accessed: 2026-03-22. URL: https://www.ericsson.com/en/reports-and-papers/white-papers/the-network-for-ai-experiences.

Ericsson (n.d.). *Ericsson history*. Accessed: 2026-02-12. URL: https://www.ericsson.com/en/about-us/history.

Giattino, C., E. Mathieu, V. Samborska, and M. Roser (2026). *Total Monthly Distance Traveled by Passengers in California's Driverless Taxis*. Accessed: 2026-03-22. URL: https://ourworldindata.org/grapher/passenger-miles-traveled-self-driving-taxis.

Gieske, J. (2024). *What is mMTC in 5G, and How Does It Benefit a Business?* Accessed: 2026-03-22. URL: https://nybsys.com/what-is-mmtc-in-5g/.

Global X Management Company LLC (2023). *Charting Disruption: Outlook for 2023 and Beyond*. Accessed: 2026-03-22. URL: https://assets.globalxetfs.com/files/ChartingDisruption2023.pdf.

Global X Management Company LLC (2024). *Charting Disruption: Outlook for 2024 and Beyond*. Accessed: 2026-03-22. URL: https://assets-cms.globalxetfs.com/Charting-Disruption-2024-Full-Report.pdf.

Global X Management Company LLC (2025). *Charting Disruption: Outlook for 2025 and Beyond*. Accessed: 2026-03-22. URL: https://go.globalxetfs.com/l/750543/2024-12-09/c943c5.

Global X Management Company LLC (2026). *Charting Disruption: Outlook for 2026 and Beyond*. Accessed: 2026-03-22. URL: https://sponsored.bloomberg.com/immersive/globalx/charting-disruption?utm_source=GlobalXETFs.com&utm_medium=Website&utm_content=ChartingDisruption.com&utm_campaign=GXChartingDisruption26.

Google Cloud (n.d.[a]). *What are foundation models?* Accessed: 2026-02-12. URL: https://cloud.google.com/discover/what-are-foundation-models.

Google Cloud (n.d.[b]). *What is Cloud Computing?* Accessed: 2026-02-12. URL: https://cloud.google.com/learn/what-is-cloud-computing.

Hoek, R. van, H. Aronsson, G. Kovács, and K. M. Spens (Feb. 2005). "Abductive reasoning in logistics research". *International Journal of Physical Distribution & Logistics Management* **35**:2, pp. 132–144. ISSN: 0960-0035. DOI: 10.1108/09600030510590318. eprint: https://www.emerald.com/ijpdlm/article-pdf/35/2/132/1102190/09600030510590318.pdf. URL: https://doi.org/10.1108/09600030510590318.

Höst, M., B. Regnell, and P. Runeson (2006). *Att genomföra examensarbete*. Studentlitteratur, Lund.

Hsu, C.-C. and B. A. Sandford (2007). "The delphi technique: making sense of consensus". *Practical Assessment, Research & Evaluation* **12**:10, pp. 1–8.

IBM Corporation (2026). *What is edge computing?* Accessed: 2026-02-12. URL: https://www.ibm.com/think/topics/edge-computing.

Iqbal, S. and L. Pipon-Young (2009). “The delphi method”. *The Psychologist* **22**:7, pp. 598–601.

Jacobides, M. G., C. Cennamo, and A. Gawer (2018). “Towards a theory of ecosystems”. *Strategic Management Journal* **39**:8, pp. 2255–2276.

Jick, T. D. (1979). “Mixing qualitative and quantitative methods: triangulation in action”. *Administrative Science Quarterly* **24**:4, pp. 602–611. DOI: 10.2307/2392366.

Johnson, G., R. Whittington, K. Scholes, D. Angwin, and P. Regnér (2017). *Exploring Strategy: Text and Cases*. 11th ed. Pearson, Harlow.

Kalvin Bahia, Facundo Rattel (Nov. 2025). *Vision 2040: Future Spectrum Needs*. Tech. rep. Accessed: 2026-03-22. GSMA. URL: https://www.gsma.com/connectivity-for-good/spectrum/gsma_resources/vision-2040-future-spectrum-needs/.

Krefting, L. (1991). “Rigor in qualitative research: the assessment of trustworthiness”. *American Journal of Occupational Therapy* **45**:3, pp. 214–222.

Lincoln, Y. S. and E. G. Guba (1985). *Naturalistic Inquiry*. Sage Publications, Newbury Park, CA.

Lindgren, M. and H. Bandhold (2009). *Scenario Planning: The Link Between Future and Strategy*. Palgrave Macmillan, New York.

Mangiante, S., G. Klas, A. Navon, G. Zhuang, J. Ran, and M. Dias Silva (2017). *VR is on the Edge: How to Deliver 360° Videos in Mobile Networks*. DOI: 10.1145/3097895.3097901. URL: https://dl.acm.org/doi/epdf/10.1145/3097895.3097901.

MarketsandMarkets (2026). *Augmented Reality and Virtual Reality Market Size, Share & Growth*. Accessed: 2026-03-22. URL: https://www.marketsandmarkets.com/Market-Reports/augmented-reality-virtual-reality-market-1185.html.

McKinsey & Company (2025). *What Is a Self-Driving Car?* Accessed: 2026-03-22. URL: https://www.mckinsey.com/featured-insights/mckinsey-explainers/what-is-a-self-driving-car.

Michael Chui, Mark Collins, Mark Patel (Nov. 2021). *The Internet of Things: Catching up to an accelerating opportunity*. Accessed: 2026-03-22. McKinsey & Company. URL: https://www.mckinsey.com/~/media/mckinsey/business%20functions/mckinsey%20digital/our%20insights/iot%20value%20set%20to%20accelerate%20through%202030%20where%20and%20how%20to%20capture%20it/the-internet-of-things-catching-up-to-an-accelerating-opportunity-final.pdf.

Mishra, R. (Dec. 2025). *Google’s Gemini Eating ChatGPT’s Lunch: Market Share Gain From 5% To 18% Is ’Clearest Signal’ That Alphabet Is Winning AI War*. Accessed: 2026-02-13. URL: https://finance.yahoo.com/news/googles-gemini-eating-chatgpts-lunch-163103026.html.

Moore, G. A. (2014). *Crossing the Chasm*. 3rd ed. Harper Business, New York.

Morris, I. (2026). *Ericsson has pinned its 6G hopes on AI’s rising uplink demands*. Accessed: 2026-03-22. URL: https://www.lightreading.com/6g/ericsson-has-pinned-its-6g-hopes-on-ais-rising-uplink-demands.

Noffsinger, J. (2025). *The Cost of Compute: A $7 Trillion Race to Scale Data Centers*. Accessed: 2026-03-22. URL: https://www.mckinsey.com/industries/technology-media-and-telecommunications/our-insights/the-cost-of-compute-a-7-trillion-dollar-race-to-scale-data-centers.

Nokia (2026). *Nokia MWC 2026 Press and Analyst Event*. Accessed: 2026-03-22. URL: https://www.youtube.com/watch?v=6z29KgZw-Ic.

NVIDIA (2021). *What Is Cloud Gaming?* Accessed: 2026-03-22. URL: https://blogs.nvidia.com/blog/what-is-cloud-gaming/.

NVIDIA (n.d.). *What Is Specialized AI?* Accessed: 2026-02-12. URL: https://www.nvidia.com/en-us/glossary/specialized-ai/.

Okoli, C. and S. D. Pawlowski (2004). "The delphi method as a research tool: an example, design considerations and applications". *Information & Management* **42**:1, pp. 15–29. ISSN: 0378-7206. DOI: https://doi.org/10.1016/j.im.2003.11.002. URL: https://www.sciencedirect.com/science/article/pii/S0378720603001794.

Omdia (2025). *6G and AI Investment to Drive Global Communications Industry Growth: Omdia Forecasts*. Accessed: 2026-03-22. URL: https://omdia.tech.informa.com/pr/2025/oct/6g-and-ai-investment-to-drive-global-communications-industry-growth-omdia-forecasts.

Patton, M. Q. (2002). *Qualitative Research and Evaluation Methods*. 3rd ed. Sage Publications, Thousand Oaks, CA.

Popper, R. (2008). "Foresight methodology". *The Handbook of Technology Foresight*, pp. 44–88.

Rahman, R. (2024). *Performance per dollar improves around 30% each year*. Accessed: 2026-03-22. URL: https://epoch.ai/data-insights/price-performance-hardware.

Riback, M. (2026). *How AI-powered devices will drive the shift to uplink-heavy networks*. Accessed: 2026-03-22. URL: https://www.ericsson.com/en/blog/2026/3/ai-powered-devices-drive-the-shift-to-uplink-heavy-traffic.

Ritchie, H., E. Mathieu, M. Roser, and E. Ortiz-Ospina (2023). "Internet". *Our World in Data*. https://ourworldindata.org/internet.

Rogers, E. M. (2003). *Diffusion of Innovations*. 5th ed. Free Press, New York.

Rowe, G. and G. Wright (1999). "The delphi technique as a forecasting tool: issues and analysis". *International Journal of Forecasting* **15**:4, pp. 353–375.

Saunders, M., P. Lewis, and A. Thornhill (2019). *Research Methods for Business Students*. 8th ed. Pearson, Harlow.

Schwartz, P. (1996). *The Art of the Long View: Planning for the Future in an Uncertain World*. Currency Doubleday, New York.

Shenton, A. K. (2004). "Strategies for ensuring trustworthiness in qualitative research projects". *Education for Information* **22**:2, pp. 63–75.

Skulmoski, G., F. Hartman, and J. Krahn (Jan. 2007). "The delphi method for graduate research". *JITE* **6**, pp. 1–21. DOI: 10.28945/199.

Stephens, R. (2026). *Agentic AI Will Make Uplink the Next Mobile Bottleneck*. Accessed: 2026-03-22. URL: https://www.globenewswire.com/news-release/2026/03/18/3257921/0/en/Agentic-AI-Will-Make-Uplink-the-Next-Mobile-Bottleneck.html.

Stryker, C. (2026a). *What is Agentic AI?* Accessed: 2026-03-22. URL: https://www.ibm.com/think/topics/agentic-ai#2095054954.

Stryker, C. (2026b). *What Is Physical AI?* Accessed: 2026-03-22. URL: https://www.ibm.com/think/topics/physical-ai.

Talmar, M., B. Walrave, K. S. Podoynitsyna, J. Holmström, and A. G. L. Romme (2020). "Mapping, analyzing and designing innovation ecosystems: the ecosystem pie model". *Long Range Planning* **53**:4, p. 101850.

Waymo (2025). *Delivering More for Our Riders in a Year of Incredible Growth*. Accessed: 2026-03-22. URL: https://waymo.com/blog/2025/12/2025-year-in-review/.

World Economic Forum (Jan. 2026a). *The Global Risks Report 2026*. Accessed: 2026-03-22. World Economic Forum. URL: https://www.weforum.org/publications/global-risks-report-2026/.

World Economic Forum (2026b). *The Strategic Role of Telecom Providers Across the AI Value Chain*. Accessed: 2026-03-22. World Economic Forum. URL: https://www.weforum.org/publications/the-strategic-role-of-telecom-providers-across-the-ai-value-chain/.

Yolanda Gil, Raymond Perrault (2025). *Artificial Intelligence Index Report 2025*. Tech. rep. Accessed: 2026-02-12. Stanford University. URL: https://hai.stanford.edu/ai-index/2025-ai-index-report.

Yuksel, I. (2012). "Developing a multi-criteria decision making model for pestel analysis". *International Journal of Business and Management* **7**:24, pp. 52–66.

A Appendix

A.1 Construction of Delphi survey

This appendix section reproduces the full questionnaire administered to the expert panel in Round 1 of the Delphi study. The survey was distributed via Google Forms. Questions are presented in the order respondents encountered them. Scale anchors, response options, and instructional text are reproduced verbatim where possible to allow replication or adaptation by future researchers.

A.2 Expert panel: Network for AI round one

This survey is part of a masters thesis written at Lund University (LTH) in collaboration with Ericsson.

Today, AI traffic represents only a small fraction of total network data. For example, generative AI (GenAI) alone accounts for just 0.06%. Yet as AI adoption accelerates broadly, that share is set to grow significantly. How this growth will evolve in the 10-year timeframe is however not clear. The purpose of our thesis is therefore to analyze how this growth will impact the telecom industry. Our research aims to answer how AI will impact global mobile network traffic, what the ecosystem may look like, and which new business cases that will emerge from the ecosystem in 2035.

During our research we have identified technologies that are likely to drive traffic, ecosystem players, and factors of influence. These will be elaborated on in the beginning of each section. With the help of a group of experts, of which you are one, we aim to better understand how the future telecom sector will evolve.

As a participant you will answer two surveys where the second one is based on the answers of the first. Each survey is expected to take approximately 20 minutes. Under each question you can elaborate on your answer if you'd like. This study is conducted in accordance with GDPR. Data will be stored securely, used solely for this research, and anonymized in all reporting.

Please respond before March 31st. Don't hesitate to reach out if you have any questions.

Thank you for your participation!

Robert Westring and Alex Chamoun Lunds Tekniska Högskola

ro4300we-s@student.lu.se | al4585ch-s@student.lu.se

A.2.1 Personal information

This information is only used for categorization and followup to round 2.

Q. Name & surname (will not be disclosed in reporting) *[Short answer]*

Q. Email address *[Short answer]*

Q. Which of the following best fits you? *[Working at Ericsson, Working at a technology company, Academia, Regulator / Institution, Other]*

Q. What is your primary focus area? *[Single choice – Technology, Legal, Business / Strategy, Other]*

Q. Years of experience in telecom or AI *[Single choice – 1-4, 5-9, 10-14, 15-19, Above 20]*

Q How knowledgeable would you say that you are in telecom? *Scale: 1 (Novice) – 10 (Expert)*

Q. How knowledgeable would you say that you are in AI *Scale: 1 (Novice) – 10 (Expert)*

A.2.2 Part 1: Traffic & Technologies

- **AI Assistants**: AI systems that respond to user queries and generate outputs in real time.
- **IoT**: Connected physical devices that collect and exchange data, often processed by AI.
- **SW-based Agentic AI**: Autonomous AI agents that plan and execute multi-step tasks with minimal human input.
- **Autonomous Mobility**: Vehicles or drones that navigate and make decisions without human control.
- **Smart Cities**: AI-driven coordination of urban infrastructure to improve efficiency and services.
- **XR (VR/AR)**: Immersive technologies that blend or replace physical reality with digital environments.
- **Cloud gaming**: Games rendered on remote servers and streamed to end-user devices in real time.
- **Healthcare AI**: Remote patient monitoring via wearables that continuously track vitals and flag anomalies to clinicians.
- **Autonomous Robots**: Physical robots, whether humanoid, quadrupedal, or otherwise - that use AI to perceive and interact with their environment.

A.2.2.1 Uplink/Downlink Ratio

Q. Today's mobile network traffic is approximately 10% uplink and 90% downlink. What do you think that ratio will be in 2035? *[Single choice]*

- 10/90 (no change)
- 20/80
- 30/70
- 40/60
- 50/50
- 60/40
- 70/30
- Prefer not to answer

Q. Comment *(open-ended)* *[Long answer]*

A.2.2.2 Mobile Traffic Generation

Q. How much mobile network traffic do you think these technologies will generate by 2035? *Scale: None – Small – Moderate – High – Massive – No answer*
None = negligible traffic contribution **Massive** = one of the dominant sources of mobile network traffic

	None	Small	Moderate	High	Massive
AI Assistants	○	○	○	○	○
IoT	○	○	○	○	○
SW-based Agentic AI	○	○	○	○	○
Autonomous Mobility	○	○	○	○	○
Smart Cities	○	○	○	○	○
XR (VR/AR)	○	○	○	○	○
Cloud Gaming	○	○	○	○	○
Healthcare AI	○	○	○	○	○
Autonomous Robots	○	○	○	○	○

Q. Comment on network traffic generation*(open-ended)* *[Long answer]*

A.2.2.3 Adoption Rate

Q. What do you think will be the adoption rate for the following technologies?
Scale: Very slow – Slow – Moderate – Fast – Very fast – No answer
Very slow = still niche or experimental by 2035 **Very fast** = already mainstream or widely deployed by 2035

	Very slow	Slow	Moderate	Fast	Very fast	No answer
AI Assistants	○	○	○	○	○	○
IoT	○	○	○	○	○	○
SW-based Agentic AI	○	○	○	○	○	○
Autonomous Mobility	○	○	○	○	○	○
Smart Cities	○	○	○	○	○	○
XR (VR/AR)	○	○	○	○	○	○
Cloud gaming	○	○	○	○	○	○
Healthcare AI	○	○	○	○	○	○
Autonomous Robots	○	○	○	○	○	○

Q. Comment on adoption rates *[Long answer]*

A.2.2.4 Revenue Potential

Q. What do you think the revenue potential is for the following technologies?
Scale: None – Small – Moderate – High – Massive – No answer
None = no meaningful commercial opportunity **Massive** = transformative, market-defining opportunity

	None	Small	Moderate	High	Massive	No answer
AI Assistants	◯	◯	◯	◯	◯	◯
IoT	◯	◯	◯	◯	◯	◯
SW-based Agentic AI	◯	◯	◯	◯	◯	◯
Autonomous Mobility	◯	◯	◯	◯	◯	◯
Smart Cities	◯	◯	◯	◯	◯	◯
XR (VR/AR)	◯	◯	◯	◯	◯	◯
Cloud gaming	◯	◯	◯	◯	◯	◯
Healthcare AI	◯	◯	◯	◯	◯	◯
Autonomous Robots	◯	◯	◯	◯	◯	◯

Q. Comment on revenue potential *[Long answer]*

A.2.2.5 Limiting Factors

Q. What do you think are the primary factors LIMITING the adoption of the following technologies? (Multiple choice possible)

	Regulation	Tech maturity	Infrastructure	Standardization	Consumer Demand	Data Privacy	Energy	Risk & Safety
AI Assistants	☐	☐	☐	☐	☐	☐	☐	☐
IoT	☐	☐	☐	☐	☐	☐	☐	☐
SW-based Agentic AI	☐	☐	☐	☐	☐	☐	☐	☐
Autonomous Mobility	☐	☐	☐	☐	☐	☐	☐	☐
Smart Cities	☐	☐	☐	☐	☐	☐	☐	☐
XR (VR/AR)	☐	☐	☐	☐	☐	☐	☐	☐
Cloud gaming	☐	☐	☐	☐	☐	☐	☐	☐
Healthcare AI	☐	☐	☐	☐	☐	☐	☐	☐
Autonomous Robots	☐	☐	☐	☐	☐	☐	☐	☐

Q. Comment on limiting factors *[Long answer]*

A.2.2.6 Driving Factors

Q. What do you think are the primary factors DRIVING the adoption of the following technologies? (Multiple choice possible)

	Regulation	Tech maturity	Infrastructure	Standardization	Consumer Demand	Data Privacy	Energy	Risk & Safety
AI assistants	☐	☐	☐	☐	☐	☐	☐	☐
IoT	☐	☐	☐	☐	☐	☐	☐	☐
SW-based Agentic AI	☐	☐	☐	☐	☐	☐	☐	☐
Autonomous Mobility	☐	☐	☐	☐	☐	☐	☐	☐
Smart Cities	☐	☐	☐	☐	☐	☐	☐	☐
XR (VR/AR)	☐	☐	☐	☐	☐	☐	☐	☐
Cloud gaming	☐	☐	☐	☐	☐	☐	☐	☐
Healthcare AI	☐	☐	☐	☐	☐	☐	☐	☐
Autonomous Robots	☐	☐	☐	☐	☐	☐	☐	☐

Q. Comment on driving factors *[Long answer]*

A.2.2.7 Missing Items

Q. Are there any factors you think is missing from the driving/limiting tables? *[Long answer]*

Q. Is there any AI-technology you think is missing? *[Long answer]*

A.2.3 Part 2: Ecosystem

"The ecosystem" refers to the global telecom and AI ecosystem.

Below are the ecosystem roles we have identified, their explanations and examples:

- **Hyperscalers:** Large cloud providers that operate massive-scale data center infrastructure and offer compute, storage, and AI services globally. (AWS, Azure, Google)
- **Telecom equipment vendors:** Companies that design and supply the hardware and software that builds and runs telecom networks. (Ericsson, Nokia, Huawei)
- **CSPs / Operators:** Network operators that provide connectivity services to consumers and enterprises. (Vodafone, Deutsche Telekom, etc.)

- **AI model providers:** Companies that develop and deploy foundational AI models, made accessible via APIs or direct integration. (OpenAI, Anthropic, Open Weight, etc.)
- **Device manufacturers:** Companies that produce end-user hardware such as smartphones, laptops, and wearables. (Apple, Samsung, etc.)
- **Chipset manufacturers:** Companies that design the processors and chips powering devices, networks, and AI workloads. (Qualcomm, Apple, NVIDIA)
- **Chip foundries:** Facilities that fabricate chips at scale based on designs from chipset manufacturers. (TSMC, Intel, etc.)
- **App developers:** Companies or individuals that build software applications running on top of existing platforms and infrastructure. (Google (apps), Zoom, Uber, etc.)
- **Content platforms:** Services that produce or distribute digital content such as video, music, or games to end users. (Netflix, Meta, TikTok, etc.)
- **Infrastructure providers:** Companies that own and lease shared physical assets such as cell towers, data centers, and fiber networks to multiple tenants like telecom operators or hyperscalers. (Equinix, American Tower, Cellnex, etc.)
- **Regulators / Governments:** Public bodies that set the legal, policy, and spectrum frameworks within which the ecosystem operates.
- **Consumers:** End users who consume AI-powered services, connectivity, and devices.

A.2.3.1 Ecosystem Influence Today

Q. Assess how much influence you think each role has in the ecosystem TODAY.
Scale: None – Low – Moderate – High – Very high – No answer

	None	Low	Moderate	High	Very high	No answer
Hyperscalers	◯	◯	◯	◯	◯	◯
Telecom equipment vendors	◯	◯	◯	◯	◯	◯
CSPs / Operators	◯	◯	◯	◯	◯	◯
AI model providers	◯	◯	◯	◯	◯	◯
Device manufacturers	◯	◯	◯	◯	◯	◯
Chipset manufacturers	◯	◯	◯	◯	◯	◯
Chip foundries	◯	◯	◯	◯	◯	◯
App developers	◯	◯	◯	◯	◯	◯
Content platforms	◯	◯	◯	◯	◯	◯
Infrastructure providers	◯	◯	◯	◯	◯	◯
Regulators / Governments	◯	◯	◯	◯	◯	◯
Consumers	◯	◯	◯	◯	◯	◯

Q. Comment on ecosystem influence today *[Long answer]*

A.2.3.2 Ecosystem Influence in 2035

Q. Assess how much influence you think each role will have in the ecosystem in 2035. *Scale: None – Low – Moderate – High – Very high – No answer*

	None	Low	Moderate	High	Very high	No answer
Hyperscalers	○	○	○	○	○	○
Telecom equipment vendors	○	○	○	○	○	○
CSPs / Operators	○	○	○	○	○	○
AI model providers	○	○	○	○	○	○
Device manufacturers	○	○	○	○	○	○
Chipset manufacturers	○	○	○	○	○	○
Chip foundries	○	○	○	○	○	○
App developers	○	○	○	○	○	○
Content platforms	○	○	○	○	○	○
Infrastructure providers	○	○	○	○	○	○
Regulators / Governments	○	○	○	○	○	○
Consumers	○	○	○	○	○	○

Q. Comment on ecosystem influence in 2035 *[Long answer]*

Q. Is there any role in the ecosystem you think is missing? *[Long answer]*

A.2.4 Part 3: Factors of Influence

Below are the macro factors we have identified and their explanations:

- **Geopolitical fragmentation:** The division of the global technology and telecom landscape into competing regional blocks driven by political and economic tensions.
- **Infrastructure investment flows:** The direction and scale of capital allocated toward building network, compute, and AI infrastructure.
- **Data privacy:** The degree to which consumers and regulators accept or resist how personal and sensitive data is collected, stored, and used.
- **Compute efficiency:** The ability to deliver more AI and network processing power per unit of energy and cost.
- **Climate change:** The capacity of network and AI infrastructure to operate sustainably and withstand climate-related disruptions.
- **AI regulation:** Government and supranational policies that govern the development, deployment, and use of AI systems.

A.2.4.1 Macro Factor Influence on the Ecosystem

Q. Rate how much influence you think these macro factors will have on the ecosystem in 2035. *Scale: None – Low – Moderate – High – Very high – No answer*

	None	Low	Moderate	High	Very high	No answer
Geopolitical fragmentation	○	○	○	○	○	○
Infrastructure investment flows	○	○	○	○	○	○
Data privacy	○	○	○	○	○	○
Compute efficiency	○	○	○	○	○	○
Climate change	○	○	○	○	○	○
AI regulation	○	○	○	○	○	○

Q. Is there any macro factor you think is missing? *[Long answer]*

Q. Comment on macro factors *[Long answer]*

A.2.5 Part 4: Compute Topology

A.2.5.1 Compute Topology Characterisation

Q. Choose what you think best characterizes the compute topology today.
Centralized compute topology: *AI training and inference are primarily performed in hyperscale cloud or large centralized data centers.*
Distributed compute topology: *A significant share of compute is performed at the network edge or directly on devices (e.g., edge nodes, base stations, endpoints).*

	Centralized	Distributed	No answer
Today	○	○	○

Q. Choose what you think best characterizes the change in compute topology until 2035.

	More centralized	No change	More distributed	No answer
2035	○	○	○	○

Q. Comment on compute topology *[Long answer]*

A.2.5.2 Macro Factor Drivers of Compute Topology

Q. To what extent do the macro factors drive compute toward a centralized or distributed topology?

	Centralized	Slightly centralized	No impact	Slightly distributed	Distributed	No answer
Geopolitical fragmentation	○	○	○	○	○	○
Infrastructure investment flows	○	○	○	○	○	○
Data privacy	○	○	○	○	○	○
Compute efficiency	○	○	○	○	○	○
Climate change	○	○	○	○	○	○
AI regulation	○	○	○	○	○	○

Q. Comment on compute topology drivers *[Long answer]*

A.2.6 Part 5: Economic Value Capture

A.2.6.1 Value Capture Characterisation

Q. Choose what you think best characterizes the economic value capture today.
Concentrated value capture: *A small number of actors capture the majority of economic value in the ecosystem.*
Distributed value capture: *Economic value is spread across many actors in the ecosystem (e.g., operators, cloud providers, hardware vendors, application providers).*

	Concentrated	Distributed	No answer
Today	○	○	○

Q. Choose what you think best characterizes the change in economic value capture until 2035.

	More concentrated	No change	More distributed	No answer
2035	◯	◯	◯	◯

Q. Comment on economic value capture *[Long answer]*

A.2.6.2 Macro Factor Drivers of Value Capture

Q. To what extent do the following factors drive economic value capture toward being concentrated or distributed?

	Concentrated	Slightly concentrated	No impact	Slightly distributed	Distributed	No answer
Geopolitical fragmentation	◯	◯	◯	◯	◯	◯
Infrastructure investment flows	◯	◯	◯	◯	◯	◯
Data privacy	◯	◯	◯	◯	◯	◯
Compute efficiency	◯	◯	◯	◯	◯	◯
Climate change	◯	◯	◯	◯	◯	◯
AI regulation	◯	◯	◯	◯	◯	◯

Q. Comment on economic value captures drivers *[Long answer]*

A.2.7 Part 6: System Focus

A.2.7.1 System Focus Characterisation

Q. Choose what you think best characterizes the system focus today.
Cost efficiency focus: System design prioritizes cost efficiency and resource utilization.
Resilience focus: System design prioritizes redundancy and operational continuity.

	Cost efficiency focus	Resilience focus	No answer
Today	◯	◯	◯

Q. Choose what you think best characterizes the change in system focus until 2035.

	More focus on cost efficiency	No change	More focus on resilience	No answer
2035	◯	◯	◯	◯

Q. Comment on system focus *[Long answer]*

A.2.7.2 Macro Factor Drivers of System Focus

Q. To what extent do the following factors drive system design toward a cost efficiency-oriented or resilience-oriented approach?

	Efficiency-oriented	Slightly efficiency-oriented	No impact	Slightly resilience-oriented	Resilience-oriented	No answer
Geopolitical fragmentation	◯	◯	◯	◯	◯	◯
Infrastructure investment flows	◯	◯	◯	◯	◯	◯
Data privacy	◯	◯	◯	◯	◯	◯
Compute efficiency	◯	◯	◯	◯	◯	◯
Climate change	◯	◯	◯	◯	◯	◯
AI regulation	◯	◯	◯	◯	◯	◯

Q. Comment on compute architecture drivers *[Long answer]*

A.2.8 Part 7: Open-Ended Closing Question

Q. What technological or structural change do you think will most reshape mobile network traffic by 2035? *[Long answer]*

A.3 Expert panel: Network for AI round two

A.3.1 Section 1: Intro

This is the second round of a Delphi study done as part of a masters thesis at Lund University at the Faculty of Engineering. The thesis is in collaboration with Ericsson and is conducted during the spring of 2026. Part two contains questions about four future scenarios built based on the answers from round one.

Detailed answers and analysis will be shared with all respondents when the masters thesis is completed in early June.

Please respond before April 24th. Do not hesitate to reach out should you have any questions.

Thank you for your participation!

Robert Westring & Alex Chamoun Lunds Tekniska Högskola

ro4300we-s@student.lu.se | al4585ch-s@student.lu.se

Q. Please enter your email address:

A.3.2 Section 2: Panel information

Below are distributions of expert segments and the panels years of experience from round 1.

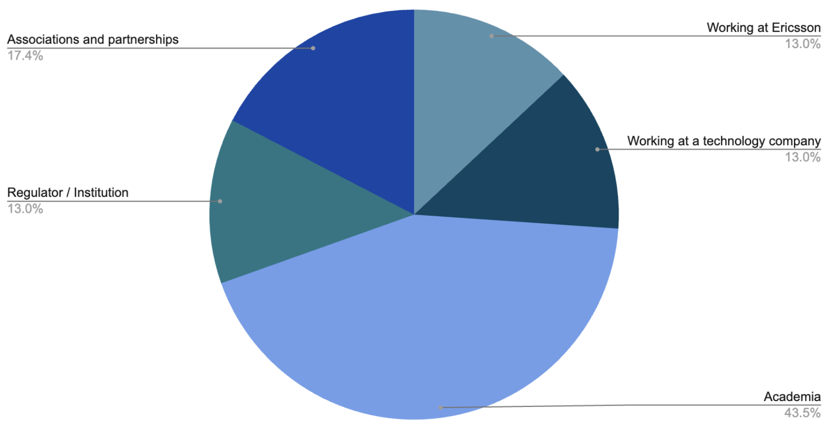

Figure A.1 Expert segments

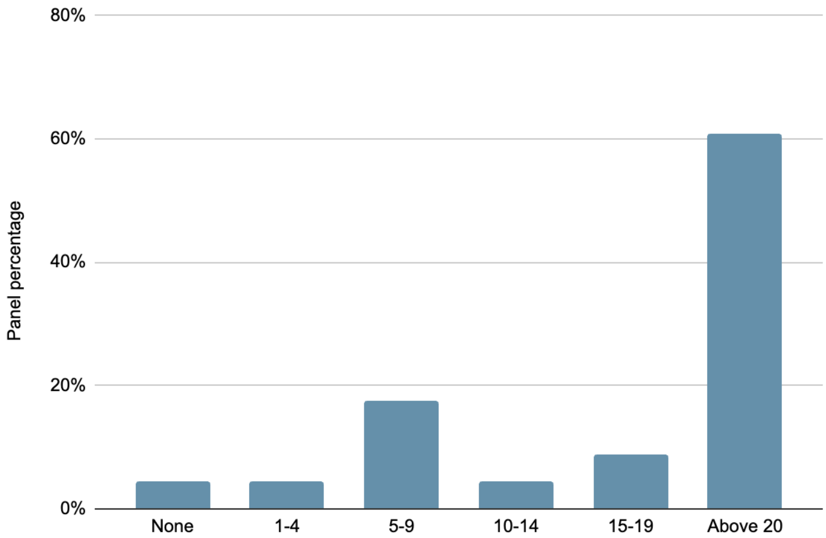

Figure A.2 The panel's years of experience

A.3.3 Section 3: Scenarios

The aggregated results from round 1 suggest that the following four scenarios are the most plausible, alongside a baseline characterization of todays Network for AI ecosystem.

The panel was overwhelmingly in agreement that the ecosystem in 2035 have moved toward a distributed compute topology. Therefore all future scenarios has that as a set parameter.

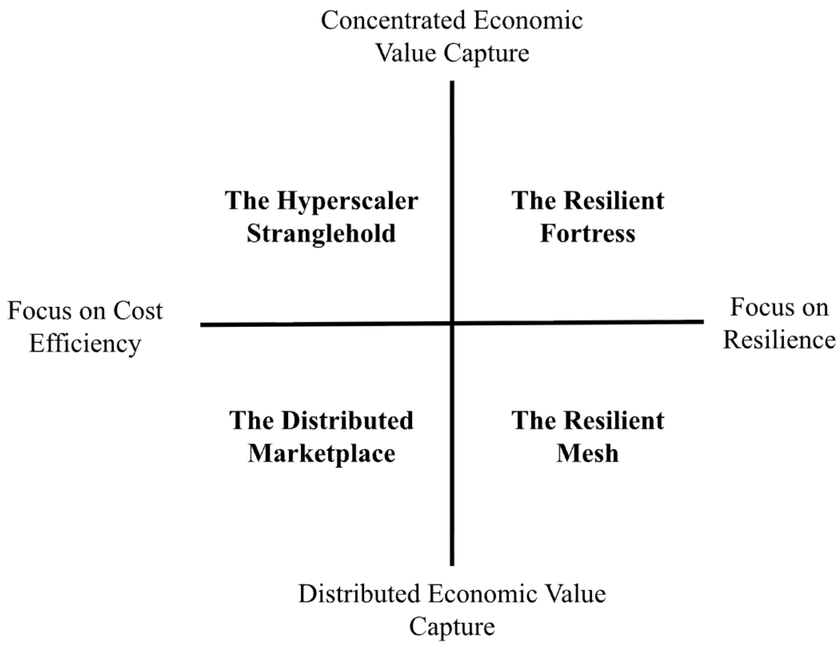

Figure A.3 Scenario matrix

Today's Baseline: The Efficiency Pipeline

Centralized Compute | Concentrated Value Capture | Cost Efficiency Focus

"Scale wins, pipes lose."

By 2026, the AI and telecom ecosystem had consolidated around a small number of dominant players. Hyperscalers and AI model providers owned much of the compute infrastructure, the developer ecosystems, and the model deployment layer. As a result capturing the majority of economic value generated by the AI boom. CSPs had been reduced to connectivity provision. Telecom equipment vendors aligned their roadmaps around solutions with low total cost of ownership to remain relevant in an environment driven by a reduction in operational expenditures. The ecosystem was highly efficient and value skewed toward those who controlled the value chain.

- **Market structure:** A small number of hyperscalers and AI model providers dominate the ecosystem, with CSPs and telecom equipment vendors competing on cost at the lower layers of the value chain.
- **Strategic logic:** Cost efficiency is the dominant imperative across the ecosystem, with operators automating aggressively to reduce operational expenditures and procurement driven by lowest total cost of ownership.
- **Infrastructure characteristics:** AI training and inference are performed primarily in hyperscale cloud data centers. CSPs provide connectivity but hold no meaningful stake in the compute layer.

- **Tensions:** Energy costs and latency demands are straining centralized architectures, geopolitical fragmentation is creating pressure for sovereign compute, and regulatory scrutiny of value concentration is intensifying.

Q. Is "The Efficiency Pipeline" a fair characterization of the current AI and telecom ecosystem? [Yes/No]

Q. If no, please elaborate: [short answer text]

A.3.4 Section 4: Scenario 1 The Resilient Mesh

More Distributed Compute | More Distributed Value Capture | More Resilience focus

"No single point of failure, neither technical or economic."

By 2035, compute had distributed toward the edge, value had spread across a broader set of players, and resilience had become a priority. No ecosystem player controlled the entire value chain. Redundancy was treated as something of value rather than an overhead cost. The ecosystem had been redesigned to absorb shocks having continuity prioritized over cost reduction at every turn.

- **Market structure:** The ecosystem is broadly fragmented, each player holding distinct but interdependent roles across the ecosystem.
- **Strategic logic:** Resilience is something worth paying extra for. Investments flowing toward continuity rather than only cost reduction.
- **Infrastructure characteristics:** Compute is distributed with redundancy built in and ownership is spread across a broad set of actors.
- **Tensions:** A distributed and redundant system is expensive to operate which raise questions of economic sustainability.

Q. Do you think this is a plausible scenario? [Yes/No]

Q. If yes, what do you think are the primary things that need to happen for this scenario to materialize? [Long text answer]

Q. What do you think are the major consequences given this scenario? [Long text answer]

Q. Which ecosystem roles will be winners given this scenario? (Pick three)

- ☐ Hyperscalers
- ☐ Telecom equipment vendors
- ☐ CSPs / Operators
- ☐ AI model providers
- ☐ Device manufacturers
- ☐ Chipset manufacturers
- ☐ Chip foundries
- ☐ App developers
- ☐ Content platforms
- ☐ Infrastructure providers
- ☐ Regulators / Governments
- ☐ Consumers

Q. Which ecosystem roles will be losers given this scenario? (Pick three)

- ☐ Hyperscalers
- ☐ Telecom equipment vendors
- ☐ CSPs / Operators
- ☐ AI model providers
- ☐ Device manufacturers
- ☐ Chipset manufacturers
- ☐ Chip foundries
- ☐ App developers
- ☐ Content platforms
- ☐ Infrastructure providers
- ☐ Regulators / Governments
- ☐ Consumers

A.3.5 Section 5: Scenario 2 The Hyperscaler Stranglehold

More Distributed Compute | Concentrated Value Capture | Cost Efficiency focus

"The pipes got smarter. The landlords got bigger."

By 2035, compute had distributed across edge nodes and regional infrastructure, driven by latency demands and energy constraints. Yet the economic structure of the ecosystem had shifted only modestly from the 2026 baseline. Hyperscalers and AI model providers had extended their influence, maintaining their grip on value capture despite no longer owning all the physical infrastructure. Efficiency remained the dominant strategic logic, with margin pressure continuing to push operators and telecom equipment vendors toward automation and standardization. The topology had opened up but the value distribution had not.

- **Market structure:** Value remains concentrated among hyperscalers and AI model providers. They have extended their reach across a more distributed infrastructure without ceding economic control.
- **Strategic logic:** Cost efficiency has remained as the dominant imperative, with commoditization and automation compressing margins further across infrastructure and connectivity.
- **Infrastructure characteristics:** Compute is deployed across edge and regional nodes, but control remain with a narrow set of dominant actors.
- **Tensions:** The growing gap between physical infrastructure ownership and economic value capture has created unresolved friction that neither market forces nor regulation have fully addressed.

Q. Do you think this is a plausible scenario? [Yes/No]

Q. If yes, what do you think are the primary things that need to happen for this scenario to materialize? [Long text answer]

Q. What do you think are the major consequences given this scenario? [Long text answer]

Q. Which ecosystem roles will be winners given this scenario? (Pick three)

- ☐ Hyperscalers
- ☐ Telecom equipment vendors
- ☐ CSPs / Operators
- ☐ AI model providers
- ☐ Device manufacturers
- ☐ Chipset manufacturers
- ☐ Chip foundries
- ☐ App developers
- ☐ Content platforms
- ☐ Infrastructure providers
- ☐ Regulators / Governments
- ☐ Consumers

Q. Which ecosystem roles will be losers given this scenario? (Pick three)

- ☐ Hyperscalers
- ☐ Telecom equipment vendors
- ☐ CSPs / Operators
- ☐ AI model providers
- ☐ Device manufacturers
- ☐ Chipset manufacturers
- ☐ Chip foundries
- ☐ App developers
- ☐ Content platforms
- ☐ Infrastructure providers
- ☐ Regulators / Governments
- ☐ Consumers

A.3.6 Section 6: Scenario 3 The Distributed Marketplace

More Distributed Compute | More Distributed Value Capture | Cost Efficiency focus

"More players, thinner margins, no clear winner."

By 2035, both compute and economic value had distributed across a broader set of actors than at any point in the prior decade. Edge infrastructure had become common. Competition for workloads previously dominated by hyperscalers had increased. Cost efficiency remained the governing logic. Competition was intense and differentiation was difficult. The ecosystem had opened up and but the prize was a race to lower costs.

- **Market structure:** The ecosystem has fragmented, with a wider range of actors competing across infrastructure and service layers.
- **Strategic logic:** Cost efficiency remains dominant, with intense competition across the fragmented landscape compressing margins and rewarding scale and automation.
- **Infrastructure characteristics:** Compute is broadly distributed across edge nodes and regional infrastructure. Ownership is spread across many actors rather than concentrated to hyperscalers.
- **Tensions:** Low margins reduce investment capacity, raising questions about who funds large scale infrastructure upgrades.

Q. Do you think this is a plausible scenario? [Yes/No]

Q. If yes, what do you think are the primary things that need to happen for this scenario to materialize? [Long text answer]

Q. What do you think are the major consequences given this scenario? [Long text answer]

Q. Which ecosystem roles will be winners given this scenario? (Pick three)

- ☐ Hyperscalers
- ☐ Telecom equipment vendors
- ☐ CSPs / Operators
- ☐ AI model providers
- ☐ Device manufacturers
- ☐ Chipset manufacturers
- ☐ Chip foundries
- ☐ App developers
- ☐ Content platforms
- ☐ Infrastructure providers
- ☐ Regulators / Governments
- ☐ Consumers

Q. Which ecosystem roles will be losers given this scenario? (Pick three)

- ☐ Hyperscalers
- ☐ Telecom equipment vendors
- ☐ CSPs / Operators
- ☐ AI model providers
- ☐ Device manufacturers
- ☐ Chipset manufacturers
- ☐ Chip foundries
- ☐ App developers
- ☐ Content platforms
- ☐ Infrastructure providers
- ☐ Regulators / Governments
- ☐ Consumers

A.3.7 Section 7: Scenario 4 The Resilient Fortress

More Distributed Compute | Concentrated Value Capture | More Resilience focus

"Power stayed concentrated. The stakes just got higher."

By 2035, the ecosystem's strategic logic had shifted from cost efficiency toward resilience. Compute had distributed, yet value had remained concentrated. Hyperscalers and AI model providers had positioned themselves as the trusted backbone of critical digital infrastructure. Resilience had become a premium product. Those who already controlled the value chain were best placed to deliver it. The ecosystem looked structurally different from 2026, but economic power had not meaningfully redistributed.

- **Market structure:** Value concentration has persisted and modestly intensified. Hyperscalers and AI model providers now compete on reliability and security credentials in addition to scale.
- **Strategic logic:** Resilience displaced cost efficiency as the primary investment driver. Procurement decisions across the ecosystem is increasingly shaped by risk management.
- **Infrastructure characteristics:** Compute is distributed across edge and regional nodes with built in redundancy. However, operational control remain with a small number of dominant actors.
- **Tensions:** Ecosystem concentration creates a systemic dependency, the ecosystem is more robust but remains exposed to a few dominant players.

Q. Do you think this is a plausible scenario? [Yes/No]

Q. If yes, what do you think are the primary things that need to happen for this scenario to materialize? [Long text answer]

Q. What do you think are the major consequences given this scenario? [Long text answer]

Q. Which ecosystem roles will be winners given this scenario? (Pick three)

- ☐ Hyperscalers
- ☐ Telecom equipment vendors
- ☐ CSPs / Operators
- ☐ AI model providers
- ☐ Device manufacturers
- ☐ Chipset manufacturers
- ☐ Chip foundries
- ☐ App developers
- ☐ Content platforms
- ☐ Infrastructure providers
- ☐ Regulators / Governments
- ☐ Consumers

Q. Which ecosystem roles will be losers given this scenario? (Pick three)

- ☐ Hyperscalers
- ☐ Telecom equipment vendors
- ☐ CSPs / Operators
- ☐ AI model providers
- ☐ Device manufacturers
- ☐ Chipset manufacturers
- ☐ Chip foundries
- ☐ App developers
- ☐ Content platforms
- ☐ Infrastructure providers
- ☐ Regulators / Governments
- ☐ Consumers

A.4 Revenue figures used for ecosystem analysis

Revenue numbers in table A.1 were collected using annual reports and financial data websites for public companies and public estimates for private companies.

Table A.1 Global Telecom Ecosystem: Companies and Revenues (FY2024/2025, USD Billions)

Category	Company	Revenue
AI Model Providers	Anthropic	14.0
AI Model Providers	Google DeepMind (Alphabet)	1.2
AI Model Providers	OpenAI	25.0
App Developers	Airbnb	12.2
App Developers	Apple Services	109.0
App Developers	Google (apps)	303.0
App Developers	Shopify	11.6
App Developers	Uber	52.0
App Developers	Zoom	4.7
CSPs / Operators	3	9.4
CSPs / Operators	AT&T	125.6
CSPs / Operators	Bharti Airtel	21.0
CSPs / Operators	Deutsche Telekom	139.0
CSPs / Operators	Etisalat (e&)	19.3
CSPs / Operators	MTN Group	13.6
CSPs / Operators	NTT (NTT Docomo)	43.0
CSPs / Operators	Orange	41.0
CSPs / Operators	Reliance Jio	125.0
CSPs / Operators	SKTelecom	11.6
CSPs / Operators	Singtel	13.3
CSPs / Operators	SoftBank	12.6
CSPs / Operators	Stc (Saudi Telecom)	20.8
CSPs / Operators	TMobile	66.2
CSPs / Operators	Telefonica	35.1
CSPs / Operators	Telenor group	7.6
CSPs / Operators	Verizon	138.2
CSPs / Operators	Vodafone	37.4
Chip Foundries	GlobalFoundries	6.8
Chip Foundries	Intel Foundry	17.8
Chip Foundries	Samsung Foundry	12.6
Chip Foundries	TSMC	122
Chip Foundries	UMC	7.4
Chipset Manufacturers	AMD	34.6
Chipset Manufacturers	Apple Silicon	25.0

Continued on next page

Table A.1 – continued from previous page

Category	**Company**	**Revenue (USD B)**
Chipset Manufacturers	Arm Holdings	4.7
Chipset Manufacturers	Broadcom	68.3
Chipset Manufacturers	Intel	49.1
Chipset Manufacturers	Marvell Technology	8.2
Chipset Manufacturers	MediaTek	18.9
Chipset Manufacturers	Micron Technology	58.0
Chipset Manufacturers	NVIDIA	215.9
Chipset Manufacturers	Qualcomm	44.3
Chipset Manufacturers	SK Hynix	68.7
Chipset Manufacturers	Samsung Semiconductor	72.0
Content Platforms	Alphabet / YouTube	40.0
Content Platforms	Amazon Prime Video	14.0
Content Platforms	Disney+ & Hulu	5.4
Content Platforms	LinkedIn (Microsoft)	17.1
Content Platforms	Meta Platforms	199.0
Content Platforms	Naver (Webtoon, etc.)	8.2
Content Platforms	Netflix	45.2
Content Platforms	Paramount Plus	30.0
Content Platforms	Peacock	1.6
Content Platforms	Pinterest	4.2
Content Platforms	Spotify	17.2
Content Platforms	TikTok (ByteDance)	23.0
Content Platforms	Twitch (Amazon)	1.8
Content Platforms	Warner Bros. Discovery (Max)	37.3
Device Manufacturers	Apple	282.0
Device Manufacturers	HP Inc.	55.3
Device Manufacturers	Motorola (Lenovo)	69.1
Device Manufacturers	Samsung Electronics	158.0
Hyperscalers	Amazon Web Services (AWS)	128.7
Hyperscalers	Google Cloud (Alphabet)	59.0
Hyperscalers	Microsoft Azure	80.0
Hyperscalers	Oracle Cloud	15.9
Infrastructure Providers	American Tower	10.6
Infrastructure Providers	Cellnex	4.9
Infrastructure Providers	Cogent Communications	1.0

Continued on next page

Table A.1 – continued from previous page

Category	Company	Revenue (USD B)
Infrastructure Providers	Crown Castle	4.3
Infrastructure Providers	Cyrus One	1.2
Infrastructure Providers	Digital Realty	6.3
Infrastructure Providers	Equinix	9.2
Infrastructure Providers	IHS Towers	1.9
Infrastructure Providers	Iron Mountain	6.9
Infrastructure Providers	Lumen Technologies	10.8
Infrastructure Providers	NTT Data Centers	25.1
Infrastructure Providers	SBA Communications	2.8
Infrastructure Providers	Zayo Group	1.1
Telecom Eq. Vendors	Ciena	5.1
Telecom Eq. Vendors	Cisco	56.7
Telecom Eq. Vendors	CommScope	5.3
Telecom Eq. Vendors	Ericsson	25.7
Telecom Eq. Vendors	Fujitsu (networking equipment)	1.0
Telecom Eq. Vendors	Huawei	51.0
Telecom Eq. Vendors	Juniper	5.1
Telecom Eq. Vendors	Keysight	5.8
Telecom Eq. Vendors	NEC	24.3
Telecom Eq. Vendors	Nokia	21.6
Telecom Eq. Vendors	Red Hat OpenShift	8.0
Telecom Eq. Vendors	Ribbon Communications	1.0

A.5 Anonymized list of participants

Below is the anonymized list of participants. Those marked in bold answered both round one and two, those who are not bold only answered round one.

Academia

- **Professor, Electronics engineering**
- **Professor, Information science and engineering**
- **Professor, Data science**
- **Professor, Communication systems**

- **Postdoc, Information science and engineering**
- **Professor, Communication technologies**
- Professor, Computer science
- Professor, Artificial intelligence robotics and cyber security
- Associate Professor, Communication and electrical engineering
- Associate Professor, Computing science and algorithms

Associations & partnerships

- **Executive, Telecom Infra Project**
- **Head of Networks, Industry association**
- Executive, AI Commission
- Executive, AI Sweden

Regulator & Institution

- **Expert, Swedish Post and Telecom Authority**
- **Manager, Swedish Post and Telecom Authority**
- Innovation Executive, Standards body

Working at Ericsson

- **Strategy development manager**
- **Senior developer**
- Key account CTO

Working at a technology company

- **Senior manager, Dell Technologies**
- **Senior Leader, Samsung**
- Executive, Equinix
