NOKIA

Global network traffic report

2026



Contents

01 Executive summary
The big numbers behind the 2035 forecast

02 Definitions and forecast lenses
Traffic scope, the AI and access views, and demand scenarios

AI traffic

03 AI traffic outlook
AI traffic across consumer and enterprise

04 Agentic AI
Agents that execute multi-step workflows

05 Physical AI
Machines that act in the physical world

06 The network impact of AI traffic
AI traffic across the network

Global access traffic

07 Global access traffic
Global access traffic projections across consumer and enterprise

08 Consumer access traffic
Mobile, FWA and fixed access traffic, globally and by region

09 Enterprise access traffic
Enterprise office and industrial traffic outlook

10 Methodology
Forecast approach, scope and assumptions

1

Executive summary



Executive summary

Artificial intelligence is rapidly becoming the defining force in global traffic. What began as a wave of consumer chatbots and generative AI tools is evolving into a far broader ecosystem of AI applications — from autonomous software agents that act on our behalf to robots, vehicles and machines that sense and interact with the physical world. Each of these applications generates and consumes data in ways that place new demands on the networks that connect them.

Our analysis indicates that by 2035, global traffic will grow nearly fourfold, and nearly half of that traffic will be driven by AI. Volume, however, is only part of the story. AI behaves differently from the applications that shaped today's networks. This behavior will only diverge further from the norm as physical and agentic AI fuel a surge of autonomous robots and agents with unique network requirements.

AI can generate intermittent bursts of traffic or require continuous connectivity, depending on the application. Many applications will be uplink intensive. The majority of AI traffic will be invisible to end users, remaining within the confines of the cloud. Simultaneously, a growing share of AI processing will migrate from the central cloud toward the edge, giving rise to a distributed AI infrastructure, otherwise known as the AI grid.

These shifts show just how closely the futures of AI and connectivity are intertwined. As AI evolves, so must networks. Advanced connectivity will provide the capacity, low latency, uplink performance and distributed intelligence AI needs to reach its full potential. This report examines in detail the forces driving AI traffic growth and what this means for tomorrow's networks.

Executive summary

By 2035, global traffic will quadruple to 3,789 EB/month - nearly half of which will be driven by AI

It is clear AI will be the key engine driving future traffic growth, spurred by large-scale consumer and enterprise adoption, media-rich AI sessions, and the rise of agentic and physical AI.

Agentic AI and physical AI will join GenAI as sources of major traffic growth

Agentic and physical AI's contribution to total traffic today is small, but they will become two of the fastest growing application categories over the next decade, respectively growing to 160x and 105x of the 2026 levels by 2035.

A distributed AI infrastructure will emerge as traffic begins migrating from the central cloud to the edge

More-efficient models and the need for lower latency and data sovereignty are driving AI inference and intelligence closer to the network edge, altering traffic flows in the process. By 2035, we project edge and metro core will service 36% of all AI-native traffic, giving rise to what the industry calls the AI grid.

AI behaves differently, changing the fundamental nature of traffic

AI follows much different traffic patterns than traditional applications, and individual AI services also vary significantly in their traffic characteristics. Depending on the application, AI can produce extremely bursty traffic, demand continuous connectivity or generate bulk uploads of data. Some services tolerate delay, while others demand instantaneous action. Future networks will need to adapt to all these traffic contingencies.

AI will demand much greater uplink performance from the network

AI is calling into question the asymmetric nature of traditional access networks. Both consumer and enterprise AI are becoming more uplink intensive as applications increasingly sense their surroundings and stream raw data into the cloud for analysis. For cells where AI usage is highest, uplink traffic will account for 24% of total access traffic. This trend is most acute with physical AI, where more than 90% of all traffic originates in the uplink.

The terms "cellular" and "mobile" will no longer be synonymous

Fixed wireless access (FWA) will account for more than one third of consumer cellular traffic in 2035, turning cellular operators into home broadband providers. Fueling this trend is more data consumption per FWA connection and the growing installed base of FWA globally.



2

Definitions and forecast lenses

Two vantage points from which to view traffic patterns

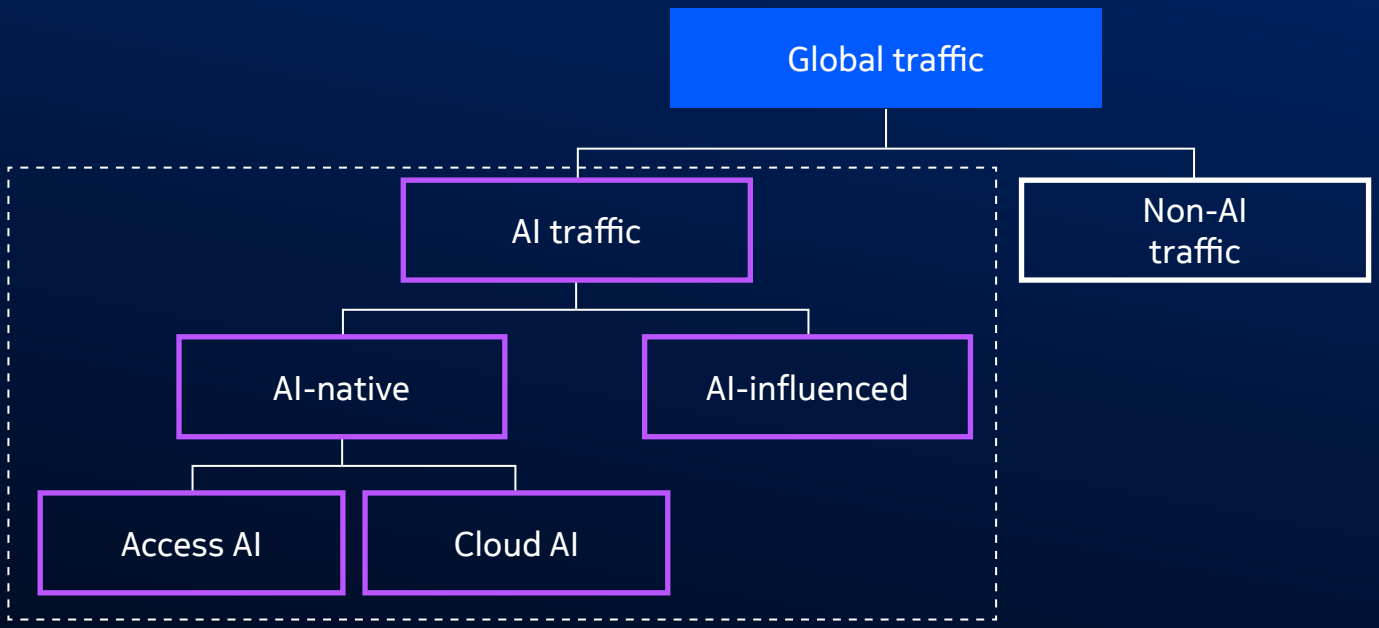
We have developed two approaches, or forecast lenses, for understanding the factors determining traffic characteristics and their impact on global networks.

The first lens categorizes global traffic into AI and non-AI traffic. This approach allows us to focus squarely on how AI drives traffic into the access network and within the cloud. AI traffic is then further split into AI-native and AI influenced traffic.

The second lens allows us to focus on global access traffic trends independently of AI. It looks at the same global traffic total but separates it into two different categories: 1) traffic generated within the cloud, and 2) global access traffic, which is traffic exchanged by consumers and enterprises with the access network.

By adopting both these lenses, we are able to see the full picture of the forces driving traffic trends. On the following pages, we provide more detailed definitions of each lens as well as an explanation on how we account for different demand scenarios in our forecasts.

Lens 1: AI view



For our AI lens, we categorize global traffic into AI traffic and non-AI traffic.

AI traffic

We split AI traffic into two categories: AI-native traffic and AI-influenced traffic.

AI-native traffic is traffic generated by, driven by or dynamically managed through AI. We divide that traffic into two areas:

- **Access AI traffic:** Traffic exchanged across the access network when consumers and enterprises use AI-enabled applications and services. Within access AI traffic we further delineate consumer and enterprise:
  - **Consumer** includes AI-driven consumer experiences, such as AI assistants, visual AI, content creation and editing, and AI-enhanced gaming and XR.
  - **Enterprise** spans AI-enabled business and industrial applications, such as productivity assistants, industrial IoT, computer vision, robotics, autonomous systems and digital twins.
- **Cloud AI traffic:** Traffic generated within the cloud when access AI traffic triggers AI workflows. It is the traffic exchanged between cloud-based AI models, agents, tools and data sources that work together to complete a request. Each exchange is counted once. It includes both agentic workflow traffic and conventional workflow traffic (fixed, predefined steps triggered by AI inference). Processing inside an AI platform, such as GPU-to-GPU communication, prompt preprocessing and safety filtering, is not included.

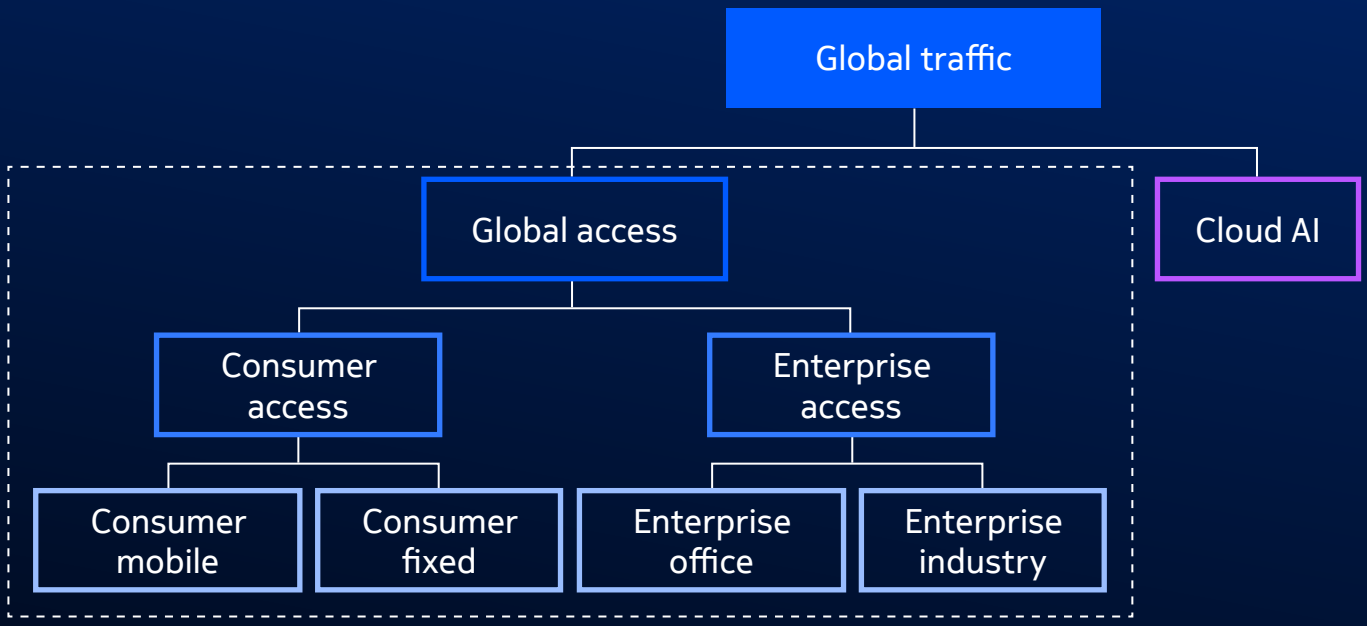
AI-influenced traffic is incremental consumer traffic on the access network generated when AI recommendations change behavior. For instance, consumers may watch more videos or buy more goods than they otherwise would, due to AI-based recommendations. Only the additional traffic generated from these recommendations is counted, not the application's total traffic.

Outside this forecast

Traffic outside of our forecast includes:

- AI traffic that remains on-device and on-premise
- Traffic that is not exchanged with the access network and is not a cloud AI workflow, such as:
  - Web crawling to collect training data
  - The transfer of training datasets between data centers
  - Token traffic on scale-up, scale-out and scale-across fabrics
  - Replication, synchronization, backup and cache fill between data centers of any type

Lens 2: Global access traffic view



For the second lens, we classify global traffic into global access traffic and cloud AI traffic. This total traffic remains the same as in the AI view.

Global access traffic

Traffic exchanged with the access network by users, connected devices and applications, wherever they are, including homes, offices and industrial sites. It is counted once at the point of exchange with the access network, regardless of the access technology used (mobile, fixed broadband, FWA, leased line or managed/private WAN). Traffic which remains on-device and on-premise is excluded. Access AI and AI-influenced traffic are included here, not added on top.

- **Consumer access traffic:** Consumer mobile traffic from smartphones, tablets and other devices over cellular networks, and consumer fixed traffic over fiber, cable, DSL and fixed wireless access (FWA).
- **Enterprise access traffic:** This includes **enterprise office traffic**, generated by office users accessing cloud applications, collaboration platforms, enterprise software and AI-enabled productivity services; and **enterprise industry traffic**, generated by industrial devices and systems such as connected machines, sensors, cameras, robots and autonomous systems, as well as immersive applications including AR/VR, remote assist and digital twins.

Demand scenarios

Recognizing the uncertainty associated with technology adoption and service evolution, our consumer and enterprise access forecasts are presented under three demand scenarios — conservative, moderate and aggressive. These scenarios reflect different assumptions regarding service adoption, application evolution, enterprise digital transformation and industrial automation, providing a realistic range of traffic demand.

3

AI traffic outlook

Overall global traffic will continue its steady pace of growth over the next decade, but AI traffic will expand much faster, claiming an ever-larger share of global totals. By 2035, AI-native and AI-influenced traffic combined will account for 44% of global traffic, according to our projections.

Consumer AI is shifting from lightweight text-based interactions toward media-rich experiences such as image and video generation and live visual AI from phones and wearables, which consume far more data per interaction. Enterprise AI, however, will outpace consumer AI, led by industrial IoT devices such as cameras, robots, autonomous vehicles and drones.

These trends won't just generate massive volumes of traffic; they will change the nature of traffic itself. AI applications fundamentally behave differently than traditional applications. One thing is for certain. In a decade's time, AI will produce large-scale shifts in global traffic patterns. Our networks will need to adapt.

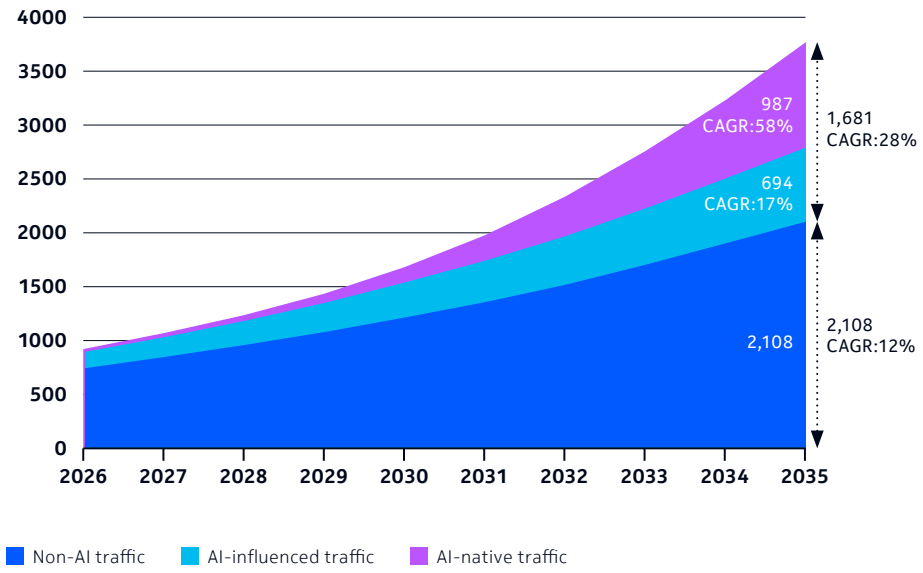
Global traffic will nearly quadruple by 2035, led by a massive surge in AI traffic

There is no question that AI will become the dominant driver of global traffic in the next decade. While overall global traffic will increase 4x to 3,789 EB/month in 2035, we anticipate 44% of this data movement will come from AI-native and AI-influenced traffic, which combined will grow at a rate of 28% annually.

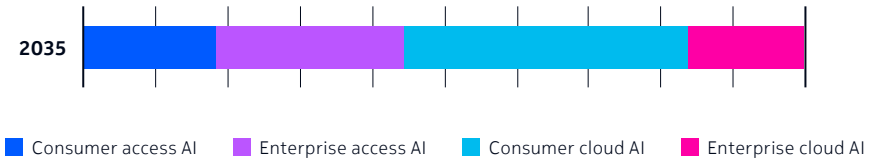
As AI services evolve into media-rich, real-time and autonomous experiences, we project AI-native traffic to grow to 60 times its 2026 levels at a 58% CAGR. Total AI-native traffic will reach 987 EB/month in 2035, accounting for 26% of global traffic. Consumer access and enterprise access AI traffic respectively will account for 181 EB/month and 258 EB/month of this traffic. But the biggest source of AI-native growth, 548 EB/month, will come from cloud AI traffic as agentic AI kicks off a wave of cloud-based interactions and exchanges. AI-influenced traffic will add another 694 EB/month to the traffic load in 2035, accounting for 18% of total global traffic.

We explore each of these trends individually in the subsequent sections of this report, showing not only how AI applications are changing the volumes of network traffic but also how this traffic behaves.

Global traffic (EB/month)



AI-native traffic share in 2035



AI-native traffic is the new growth engine

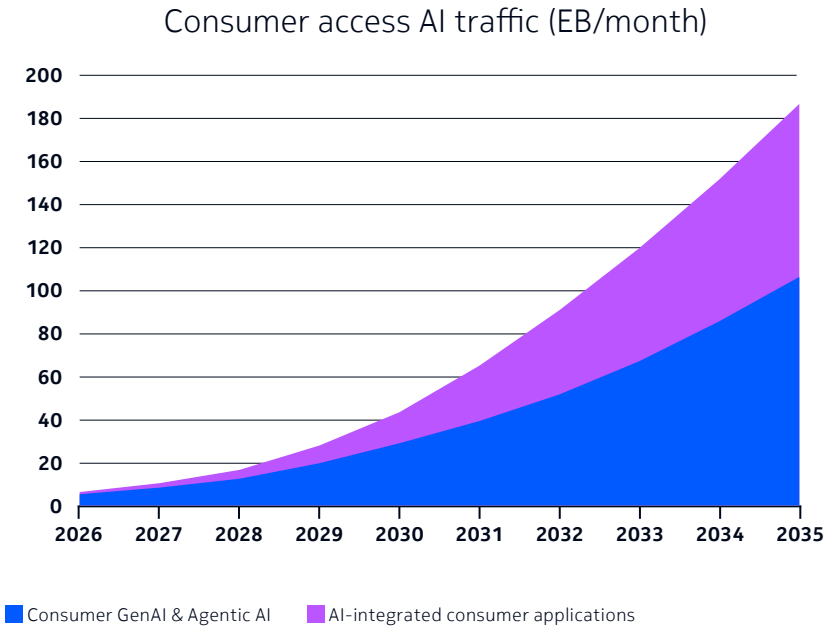
Media-rich, industrial, physical and agentic AI workloads will increase both traffic volume and network complexity.

Consumer access AI traffic will increase 42% annually, led by media-intensive AI

The evolution of consumer AI from predominantly text-based interactions to media-rich experiences, combined with widespread adoption, will be key factors driving consumer AI growth. Consumer access AI traffic will rise to 181 EB/month in 2035, increasing at a 42% CAGR.

Within consumer AI traffic, frequency of use will not directly translate into data volume. Text interactions may account for the largest share of sessions, but their lightweight payloads will generate limited traffic. Visual applications — including image and video generation, editing and enhancement — will consume substantially more data per interaction. Live visual AI will add further demand as smartphones and wearable devices increasingly stream camera feeds for real-time analysis. We project that 103 EB/month of consumer AI traffic in 2035 will come from GenAI and agentic AI use, fueled by increased adoption, enhanced application resolution and longer session durations.

A further 78 EB/month will come from AI-integrated consumer applications, where AI enhances the core experience without becoming the primary user interface. Examples include AI-enhanced XR, AI gaming and AI-assisted digital marketplaces. As AI becomes embedded across more consumer applications, these services will make an increasing contribution to consumer AI traffic.



Consumer AI traffic shifts from text interactions to media creation

Text-based AI will remain popular but generate light traffic. AI creation and editing of images, video and documents, and live visual AI from phones and wearables will drive traffic volume.

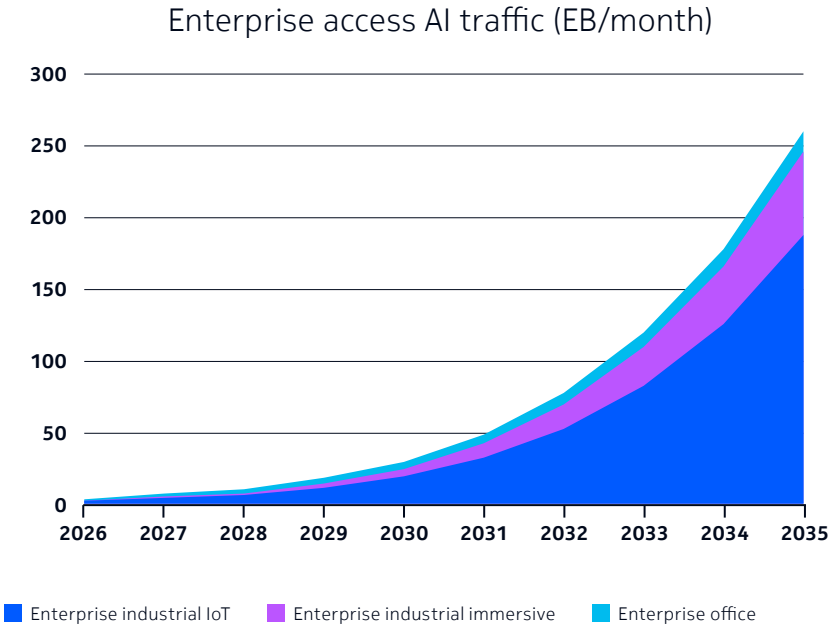
Enterprise access AI traffic will outpace consumer, driven by industrial IoT

While much attention is focused on consumer AI, enterprise AI is quickly becoming the bigger growth engine in access AI traffic. We project enterprise AI will overtake consumer AI in traffic load in the latter part of our forecast period, growing at a 59% CAGR to reach 258 EB/month by 2035. This growth is spread across three categories, each with distinct traffic characteristics.

Industrial IoT AI will be the largest contributor, accounting for 70% of enterprise AI traffic by 2035. It encompasses AI-enabled cameras, humanoid robots, autonomous vehicles, drones, sensors, industrial equipment and other connected devices and assets deployed across diverse industry verticals including manufacturing, energy and utilities, transportation, smart buildings, healthcare and retail (we cover many of these use case in the physical AI section). These machines transmit substantial volumes of video, telemetry, sensor and perception data over the uplink channel, while the downlink typically carries smaller volumes, consisting of alerts, control commands and instructions. This results in a strongly uplink-skewed traffic profile.

Industrial immersive AI is another major contributor to enterprise AI traffic. It encompasses applications such as AR/VR-based training, remote assistance, digital twins and collaborative design, where AI enhances immersive and interactive industrial workflows. Video-centric AR remote assistance can be uplink-heavy due to live video streaming, while collaborative design generally exhibits a more balanced uplink-downlink profile.

Enterprise office AI includes applications such as AI copilots, coding assistants, document analysis and AI-assisted content creation. Traffic characteristics will vary by application and interaction, with uplink driven by prompts, document uploads and code submissions, and downlink by AI-generated responses, code and content.



Industrial AI pushes the network in the uplink direction

Cameras, autonomous vehicles, drones and humanoid robots will drive traffic volumes and rely heavily on the uplink. Meanwhile many immersive apps will require varying uplink ratios.

AI traffic differs not only in volume, but in behavior

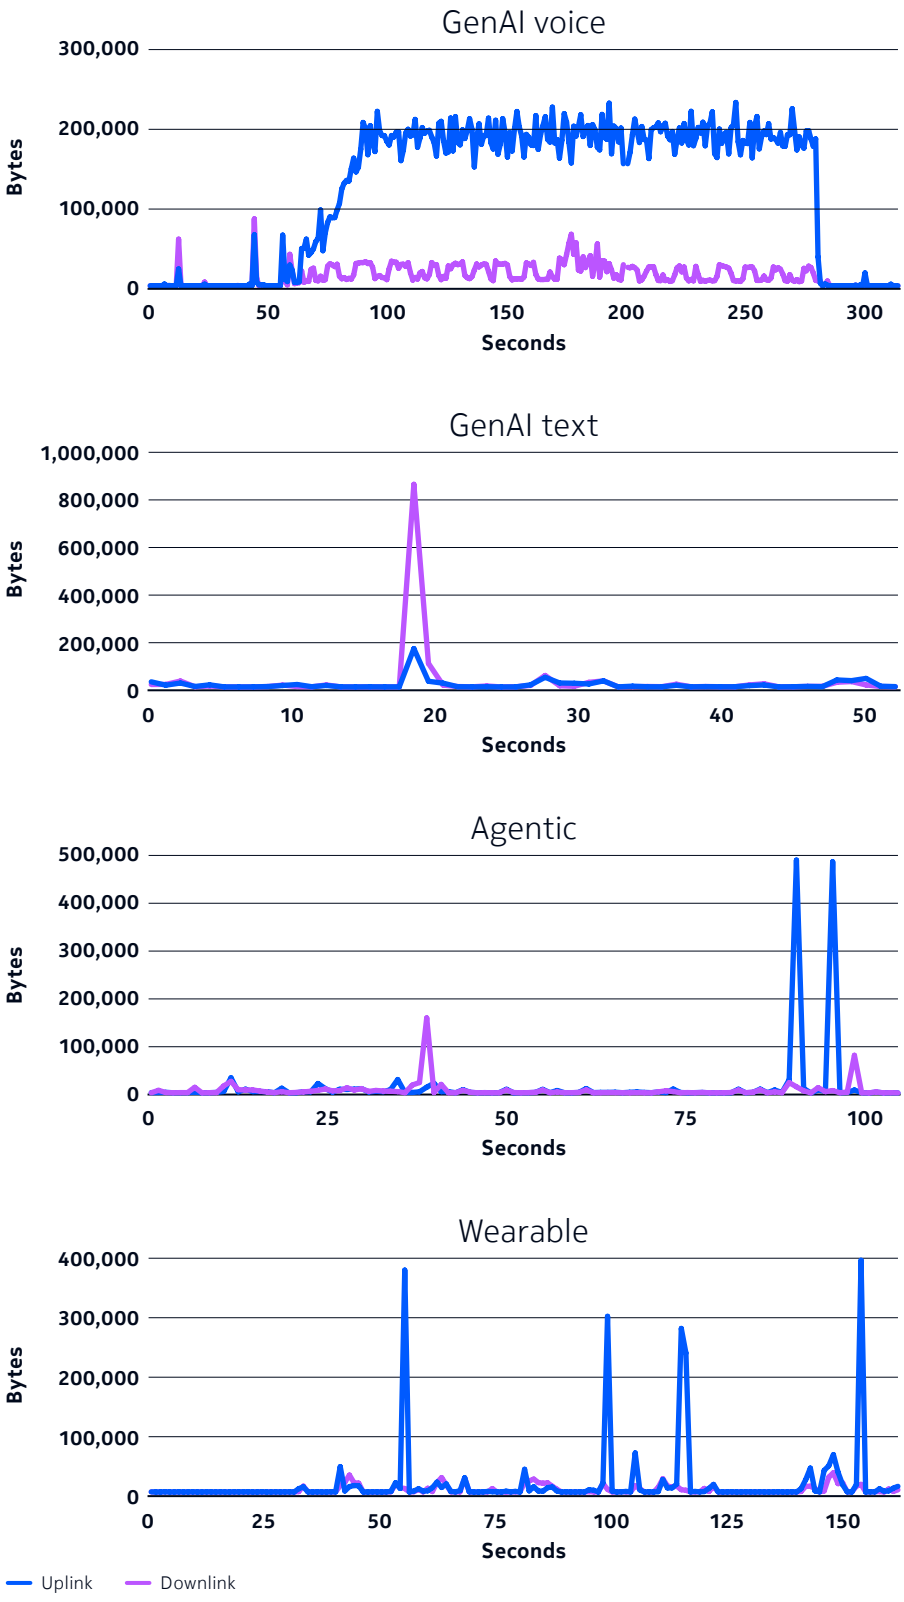
AI traffic exhibits diverse characteristics. Not all AI traffic follows the same patterns or has identical network requirements. Unlike general web traffic, which typically consists of short bursts and is strongly downlink-focused, AI traffic can vary significantly depending on the use case.

For instance, some AI applications can place heavy demands on the uplink when large video files or other content are uploaded for inference. Other applications may produce downlink bursts or mimic streaming patterns, depending on the use case. GenAI traffic for voice-based applications, such as conversational AI, typically involves continuous two-way media streams, characterized by long sustained sessions dominated by medium-sized packets.

In contrast, GenAI chat or text-based applications are mostly quiet. This lack of activity is punctuated at intervals by sharp bursts of traffic as responses are streamed back following a prompt. Meanwhile, agentic tasks exhibit longer sessions of traffic exchanges as AI agents plan, call tools and report progress. The link between the AI agent interface and the network rarely remains idle for extended periods.

Wearable devices, such as smart glasses, present a different pattern. These devices often experience long periods of silence interrupted by short, periodic spikes in traffic as they capture and process visual scenes. Data is typically transmitted in large, coalesced frames, indicating bulk transfers of image data from camera uploads.

These examples illustrate just a few use cases, each placing different demands on the network. As agentic workflows and physical AI machines scale, sustained streams, sharp bursts and heavier uplink will increasingly shape network planning.



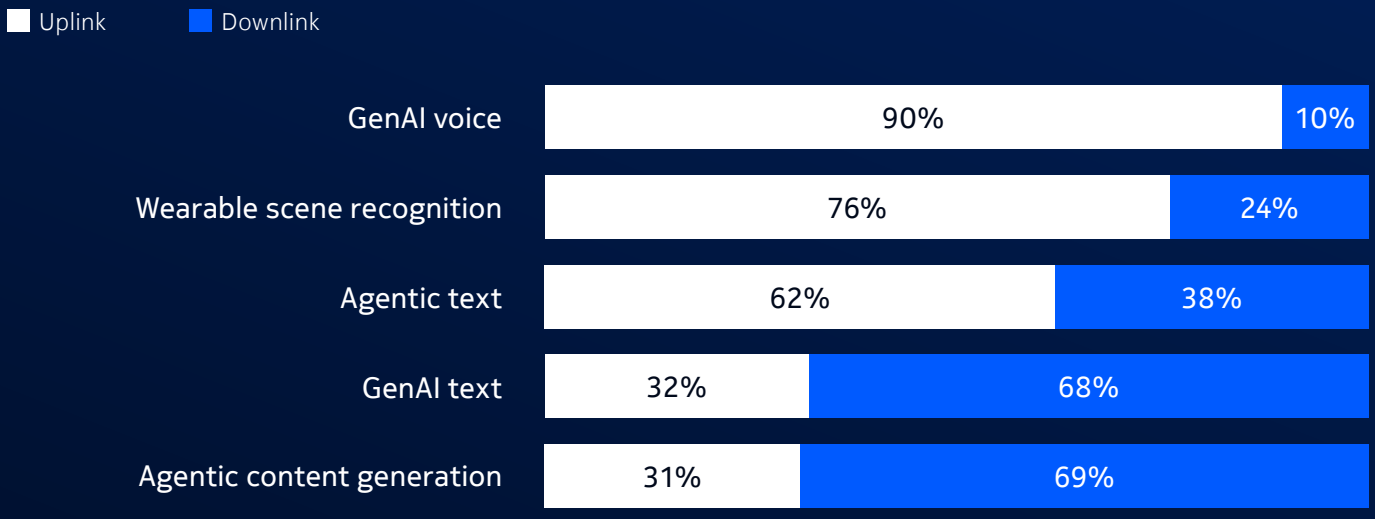
AI traffic will place increasing stress on the network uplink

AI traffic is predominantly uplink heavy. As the chart below shows, certain AI applications will send the majority of their traffic in the upstream direction. GenAI voice-conversation sessions will be 90% uplink, while smart glasses streaming video into the cloud will demand a 3-to-1 uplink-downlink ratio.

That said, specific AI use cases will also place significant demands on the downlink network. Rich-media GenAI apps will return large video, images and audio files after receiving simple text prompts from the user. Even chat applications will generate bursts of downlink traffic in response to prompts.

These trends have major implications for the network. Future network architectures must prioritize uplink capacity, while supporting short, high-intensity data transfers in both directions.

Share of session bytes, uplink vs downlink



4

Agentic AI

Agentic AI is beginning to change the nature of network traffic. Consumers and enterprises are increasingly turning to AI to assist them in their daily and professional lives — not only by automating tasks, but also through direct interaction with agents for research, content creation, planning and complex workflows. These agents are generating a new layer of traffic, much of it invisible to the end user.

What might appear to be a simple task to the user may be the result of dozens of interactions among orchestrators, other agents, data sources and external services. This is expected to drive a substantial increase in agentic workloads over the next decade, making agentic AI one of the principal contributors to future global traffic growth.

Agentic AI's impact on traffic

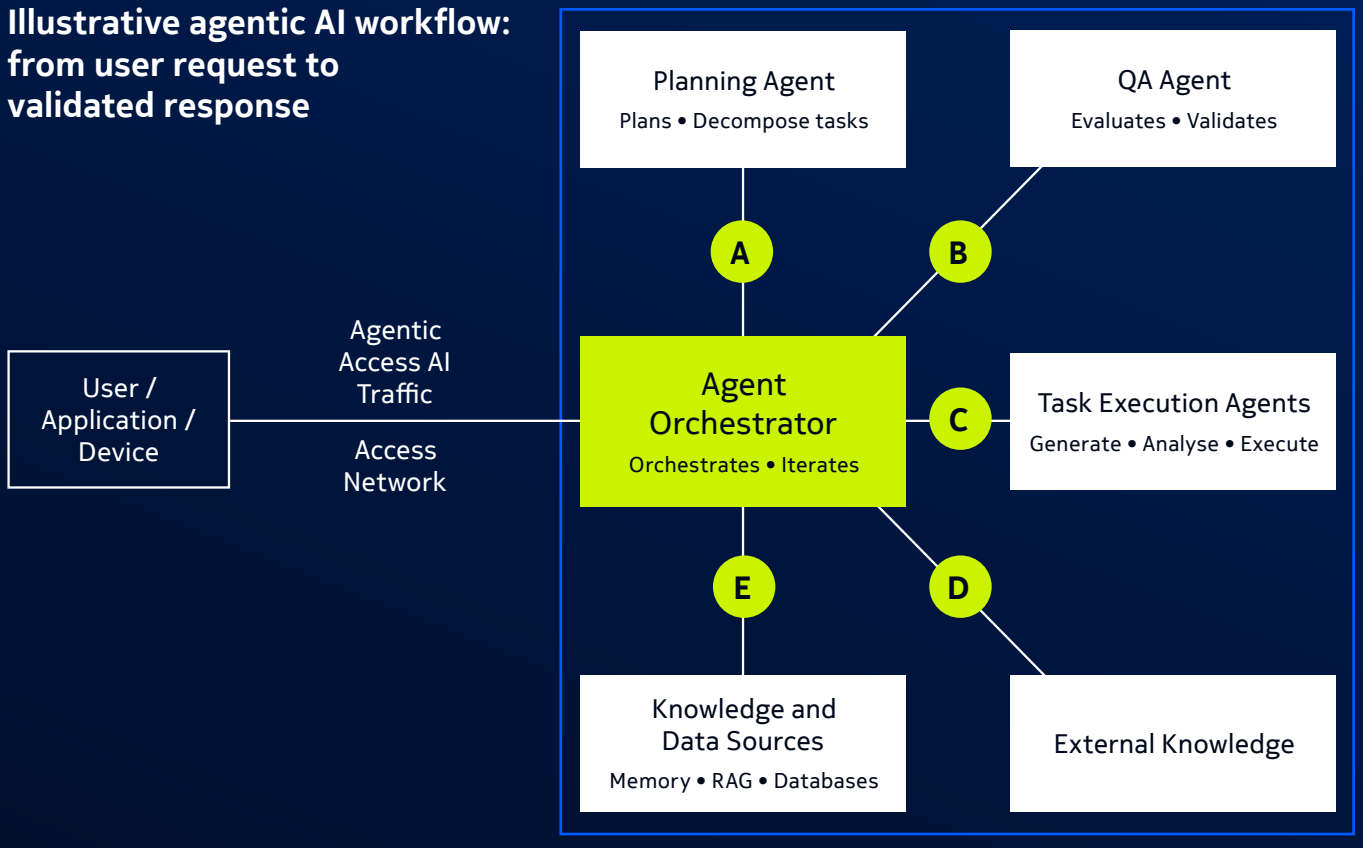
Before diving into our analysis, it's important we explain how agentic AI's orchestration process generates traffic across services. At its core, orchestration turns a single task into a series of interactions among agents, tools and data sources.

Agentic AI initiates a multi-step workflow in which an orchestrator plans tasks, invokes specialist agents and external tools, retrieves information from knowledge sources, and validates results. This process can involve repeated exchanges between the orchestrator, agents, data sources and external services, with the orchestrator only delivering a response to the user when the task is completed. These iterative workflows generate additional traffic based on the number of orchestration cycles and the type and size of data exchanged in each interaction.

Links A through E in the illustrative architecture below represent potential interactions within the workflow. An interaction is considered agentic when the orchestrator selects and triggers that interaction based on an earlier result. The orchestrator may repeat interactions over one or more links across multiple cycles, depending on task complexity and the information or validation required. Not every cycle uses every link, and some interactions may occur in parallel.

We define agentic AI traffic as a combination of agentic access AI traffic and agentic cloud AI traffic. Agentic access AI traffic includes traffic generated when applications or devices interact with cloud-based AI agents, as well as access traffic generated by agentic orchestration executed on-device or on-premises. Agentic cloud AI traffic includes traffic generated through cloud-based agentic orchestration across links A–E.

Illustrative agentic AI workflow: from user request to validated response



Agentic AI will grow 160x by 2035, becoming a key driver of network demand

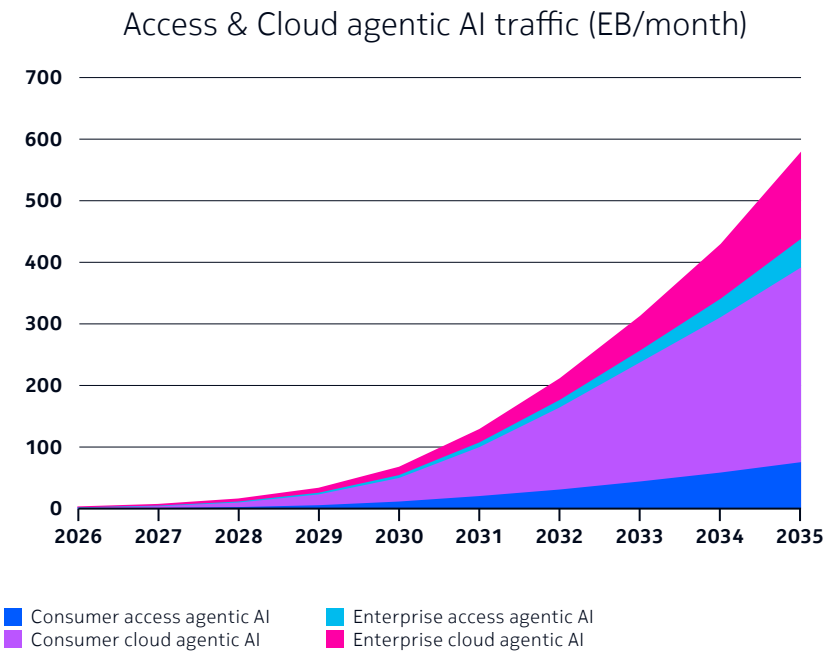
Today agentic AI traffic is small, but that is set to change in a dramatic way over the next decade. We project agentic AI traffic to reach 585 EB/month by 2035, growing at an astonishing 76% CAGR. Agentic AI traffic will not only become a major portion of all AI traffic, but also a significant percentage of total global traffic.

Much of this traffic boom, however, will be happening behind the scenes. Of the 585 EB/month, the lion's share will be agentic cloud AI traffic — 462 EB/month — which will originate inside the network, generated by cloud agentic orchestration. The remaining 123 EB/month will be agentic access AI traffic, which will be exchanged over the access network.

Some consumer and enterprise office applications will execute workflows automatically, but most agentic tasks will be user-initiated, especially in the consumer domain. Users will increasingly delegate complex media-rich tasks to AI agents that involve planning, information retrieval, tool invocation and iterative execution. This will significantly increase cloud agentic AI workflow traffic versus conventional AI inference.

Enterprise industrial AI follows a different pattern. As agentic capabilities become embedded in industrial systems, deployments span a spectrum of workflows — from event-driven or conditional execution (such as responding to equipment anomalies, quality deviations or maintenance needs) to continuous autonomous coordination and process optimization. Consequently, agentic AI traffic growth depends not only on industrial AI deployments, but also on the frequency and mode with which agentic workflows are invoked.

Depending on where agents, data sources and external services are hosted, agentic exchanges may traverse multiple network links. The resulting traffic, along with other AI traffic carried across these links, is captured in our AI traffic network impact analysis in the next section.



Agentic AI traffic projected to reach

585

EB/month by 2035

The vast majority of agentic traffic will originate inside the network

Consumer and office agentic traffic scales with adoption and usage. Industrial agentic traffic is operationally driven, scaling via triggered or continuous workflows. Once initiated, a single request fans out into multiple service-to-service exchanges.

5 Physical AI

Physical AI traffic refers to access AI traffic generated by intelligent physical systems that perceive their environment, make decisions, and trigger or execute actions in the physical world. This traffic is unique among the AI categories. While GenAI and agentic AI traffic are largely session-based, physical AI devices like autonomous vehicles, drones and robots will often remain continuously connected to the network, as will cloud cameras integrated into physical systems. And rather than just receive data, these machines will upload a disproportionate amount of data, transmitting feeds from cameras and a host of other sensors to the cloud.

This constant communication brings true contextual awareness, fusing computer vision, spatial sensing, and environmental telemetry to construct comprehensive models of surroundings while allowing multiple machines to act as synchronized swarms. As physical AI scales, so will its contribution to network workloads, becoming one of the fastest growing traffic sources over the next decade.

The expanding role of intelligent machines

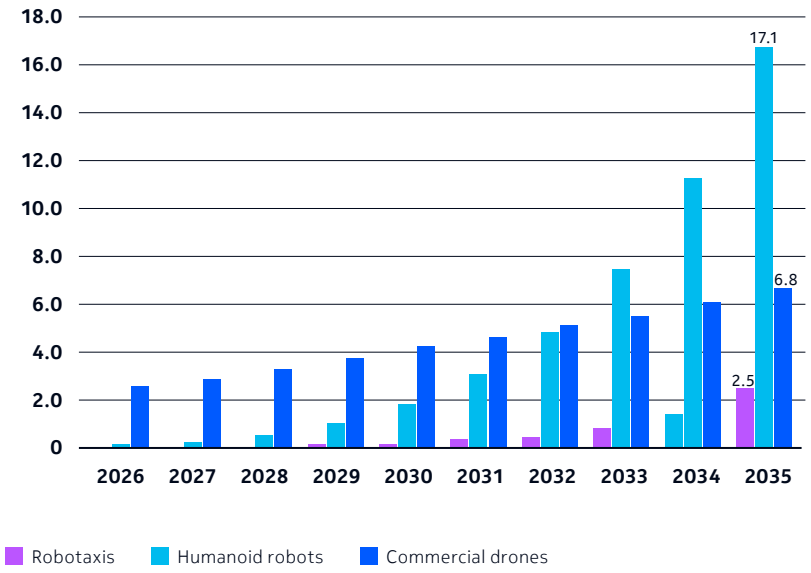
Before diving into our physical AI traffic analysis, it's important to understand the size and scale of the future intelligent-machine market. Today, physical AI deployments are limited, but we expect them to expand significantly by 2035.

The installed base of cameras alone — including surveillance, monitoring, automatic number plate recognition and body-worn cameras — will reach approximately 1.7 billion in a decade. Increasingly, these cameras will be integrated into cloud-analytics-driven physical systems, turning visual detections into actions such as triggering alarms, closing doors or gates, coordinating autonomous surveillance and enabling automated product sorting and handling.

Rapid growth will also come from machines that move through and interact with their environments, led by robots, drones and autonomous vehicles. This chart shows three such autonomous machine categories where we expect to see substantial growth in installed base through 2035: humanoid robots, commercial drones and robotaxis.

Beyond these use cases, billions of connected sensors will form the sensing and control layer of physical AI — capturing real-world conditions and generating contextual awareness. This awareness will form the basis for AI-driven decisions and actions.

Robotaxis, humanoid robots and drones: Installed base (million)



Humanoid robots reach

17.1m

by 2035

Rapid growth in three emerging physical AI devices.

Humanoid robots will reach 17.1 million units by 2035, while commercial drones and robotaxis will grow to 6.8 million and 2.5 million, respectively.

Cameras, robots, AVs and drones will trigger a wave of physical AI traffic

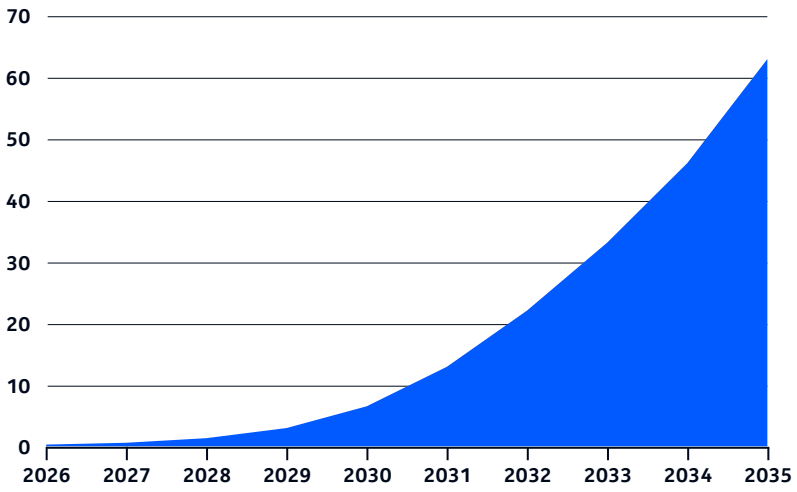
We project physical AI traffic to grow around 105x by 2035, reaching approximately 67 EB/month at a 68% CAGR. Growth of this magnitude is typically driven by a new class of connected endpoints, for instance, smartphones in the 2000s. In this case, this new class brings together autonomous machines — humanoid robots, autonomous vehicles and commercial drones moving from pilot deployments to large-scale commercial operations — with cloud cameras integrated into physical systems. Together these two categories of devices will carry a deluge of network traffic in their wake.

For the autonomous machines, each unit is a bandwidth-intensive, always-connected data source. While they run their core control loops onboard, network connectivity enables entire fleets to be supervised, coordinated and improved. Therefore, traffic grows with every new unit and every hour in service.

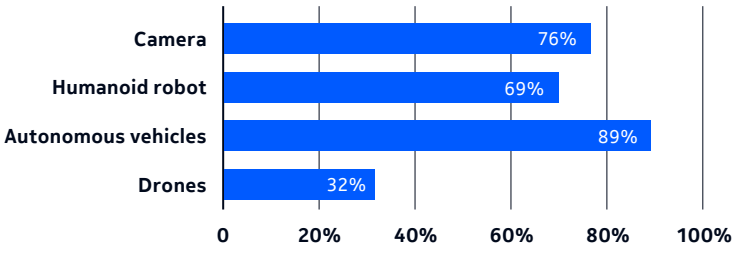
That hour in service is not the same for every autonomous machine. For instance, robotaxis, autonomous trucks and buses can operate for extended hours each day and transmit a significant amount of data to the cloud for fleet management, remote operations and model improvements. Meanwhile privately owned autonomous cars typically remain parked most of the time and may share only selected data with the cloud. Therefore, our projected autonomous vehicle CAGR of 89% is a weighted average across these vehicle classes, reflecting their rollouts and bandwidth demands.

For cameras, traffic grows as more of them are integrated into cloud-analytics-driven physical systems. Cameras will contribute the largest share of physical AI traffic because of the scale of their deployments and the substantial video traffic generated for cloud analytics.

Physical AI traffic (EB/month)



Physical AI traffic CAGR by headline use case, 2026–2035



Physical AI traffic scales in multiple dimensions

Every new machine added to a fleet and every hour it remains in service compounds traffic growth. Camera analytics will add to the load as they are integrated into physical AI systems.

Physical AI will place new demands on the uplink

Physical AI will not only boost the traffic load traversing networks, but it will also alter the character of that traffic. Physical AI workloads are overwhelmingly uplink intensive. In fact, we project that by 2035, more than 90% of physical AI traffic will be carried over the uplink channel.

Most classic IoT is light on data. Meters, trackers and sensors send small readings a few times a day. Physical AI brings a different kind of device into the IoT mix: data-intensive machines that see, move and act, while sending rich data upstream. The downstream traffic physical AI generates is relatively light, consisting of control commands, operational instructions, alerts, AI-generated insights and updates.

For instance, when cloud camera analytics spots an intruder, it directs a fleet of pan-tilt-zoom (PTZ) cameras to follow the intruder: multiple video streams go up, but only a short command comes back. Or when a docked drone launches, responding to an emergency call, it streams live video to the command center and uploads evidence footage after landing. But over the downlink channel, it receives little more than its flight plan.

As physical AI scales up, the current asymmetrical nature of access connections will need to be revisited. Cells built for downlink-led consumer demand will need to be optimized for the bursty uplink peaks that these autonomous machines require.

What is sent and what is received

Sent up: heavy and frequent

- **What it sees:** video, images and 3D scans for real-time world models
- **What it does:** status, position and results shared across the fleet
- **What it learns:** events and incidents for model retraining
- **When it needs help:** live views for remote operators

Received down: light and occasional

- **What to do:** commands, missions and routes
- **What to know:** maps, alerts and AI insights
- **How to think:** new AI models, software and firmware, often loaded at the dock or depot

Uplink defines the load

90%+

uplink in 2035 industrial physical AI across headline use cases

6 The network impact of AI traffic

AI is poised to reshape not just the services we use but the networks that carry them. AI won't only create traffic volume. It will demand new capabilities of the network from the edge to the regional core.

A single AI request rarely stays put. Once it enters the network, it can set off a cascade of exchanges among models, agents, tools and data sources, multiplying traffic as it travels across network tiers. As AI applications demand different levels of compute, sovereignty and performance, a distributed AI infrastructure — or AI grid — will emerge. AI's increasing contribution to uplink traffic will also call into question the heavily asymmetric nature of today's networks.

In this section, we examine how these shifts will impact the network in three parts. We start with how AI traffic will be amplified across the network. Next, we will explore the implications of a distributed AI infrastructure. We will conclude with an analysis of the growing stresses AI is placing on the cellular uplink.

Once AI traffic enters the network it will more than triple

By 2035, 1,134 EB/month of AI traffic will be exchanged with the access network, representing AI-influenced traffic and access AI traffic combined. This, however, is only part of the journey.

To illustrate this journey, we have created a reference architecture providing a framework for modeling AI traffic (a detailed description of this framework is in the methodology chapter). Once it enters the network, AI traffic will traverse multiple links between network tiers and service providers. This access traffic often triggers additional cloud AI traffic, which will travel over more network links. This will result in the network seeing 3.3 times the traffic exchanged with the access network.

As AI-influenced traffic is conventional traffic driven by AI, it will mostly be served from caches close to users. It will only cross a few network links, so its amplification across the network will stay modest.

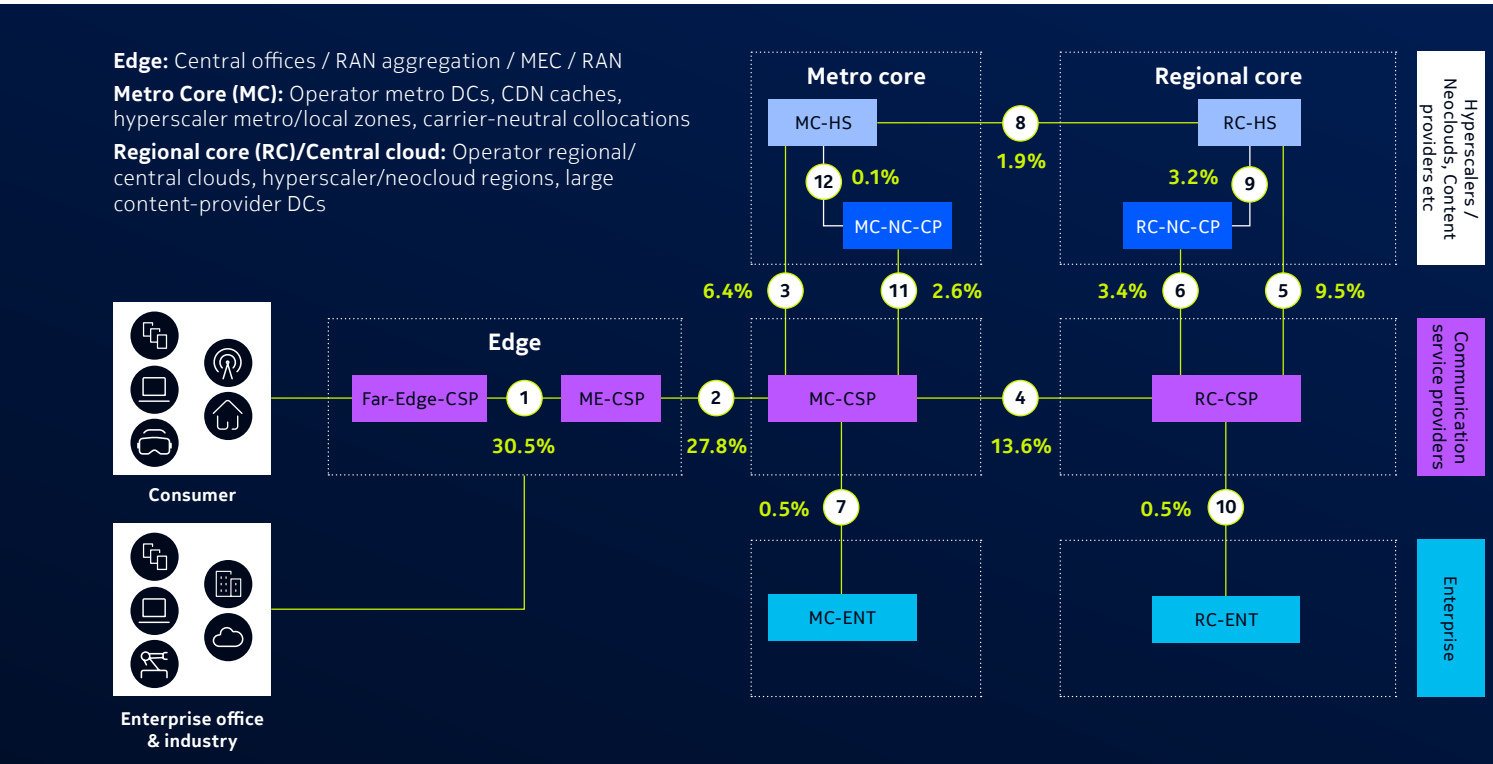
Access AI traffic, however, travels much further. In 2035, a major share of AI services will still be hosted in the regional core or the central cloud. This means a significant share of AI requests will cross the edge, the metro core and the regional core with the traffic load increasing on each link. Once at the processing location, an access AI request often triggers additional cloud AI traffic to complete requests. This cloud AI traffic will cross still more links. Each step in this chain adds to the traffic the network carries.

If we count all this traffic on every link it crosses — AI-influenced, access AI and the cloud AI it triggers — then we find that the 1,134 EB/month exchanged with the access network becomes 3,720 EB/month across network links. More than 70% of this traffic will run on the CSP transport links (links 1, 2 and 4 in the chart).

These projections provide a basis for translating AI demand into network capacity requirements and supporting scenario-based planning as AI deployment patterns evolve.

Advanced connectivity is central to AI delivery, not peripheral to it

We project more than 70% of AI's total network traffic will traverse CSP transport networks.



AI processing will increasingly move toward the edge, giving rise to a distributed AI infrastructure

We see plenty of evidence of the emergence of a distributed AI infrastructure over the next decade as AI processing begins migrating closer to the network edge. This distributed infrastructure, also known as the AI grid, will bring compute and network resources closer to where the data is created and used. By 2035, we project that the edge and metro core will service 36% of all AI-native traffic.

Which aspect of the distributed AI infrastructure services a particular workload is determined by latency and compute requirements, compute availability and data sovereignty. (Please see the methodology section for a detailed description of the three network tiers.)

The edge tier will service 9% of AI-native traffic, addressing the need for closed-loop, real-time interactions and latency-sensitive cloud rendering, particularly for immersive services. It will also handle compact consumer AI workloads, such as image enhancement and optical character recognition. In industrial IoT, the edge enables real-time camera analytics and alerts, as well as time-critical AI for autonomous systems, and robot and drone teleoperation and perception.

The metro core will handle 27% of AI-native traffic. It will serve distributed workloads that need lower latency than the regional core can provide or more compute power than available at the edge. The metro core supports hybrid XR, remote assistance, live video analytics for industry and security, and the coordination of drone, robotaxi and humanoid robot fleets. Distillation and quantization also bring lightweight conversational and generative AI closer to users and enterprises, meeting latency, compute, sovereignty, privacy and compliance needs.

The regional core will remain the primary inference tier, servicing 64% of AI-native traffic. It will host frontier conversational and multimodal assistants, enterprise copilots and compute-intensive GenAI services. These workloads depend primarily on compute capacity and access to large datasets, making centralized GPU infrastructure well suited to their processing requirements. For industrial IoT and autonomous systems, the regional core will support OEM and fleet-wide analytics, cross-site analysis, long-term storage and review of camera and drone footage, routine robot telemetry, and the processing of historical data for model evaluation and improvement.

AI-native traffic volume distribution by network tier, 2035



Edge — 5–15 ms

Closed-loop, interactive cloud-rendered immersive services; compact consumer AI; and real-time IoT inference for cameras, drones and autonomous systems.

Metro core — 15–50 ms

Distributed processing for hybrid XR, live multimodal AI, metro-wide IoT analytics and coordination, lightweight distilled and quantized generative AI models, and data-locality requirements.

Regional core — deep

Frontier assistants, enterprise copilots, compute-intensive generative AI and centralized analytics.

Note: Percentages are AI-native traffic volume share by the tier where AI processing runs, not compute, energy or token share. Token volumes describe inference intensity; traffic reflects the data exchanged and where it is served.

Increases in overall uplink demand will be modest, but they mask local AI-driven constraints

While overall mobile uplink traffic will certainly grow in the next decade, on average we don't project any big surges in the portion of traffic carried over the uplink. But this trend doesn't hold true when examining cells where AI use is most intense. Where GenAI, agentic AI and especially physical AI use is high, we see uplink traffic ballooning.

Cellular uplink share is increasingly shaped by three traffic categories: traditional consumer cellular applications, consumer GenAI and agentic AI access traffic, and cellular IoT, which is increasingly driven by industrial and physical AI.

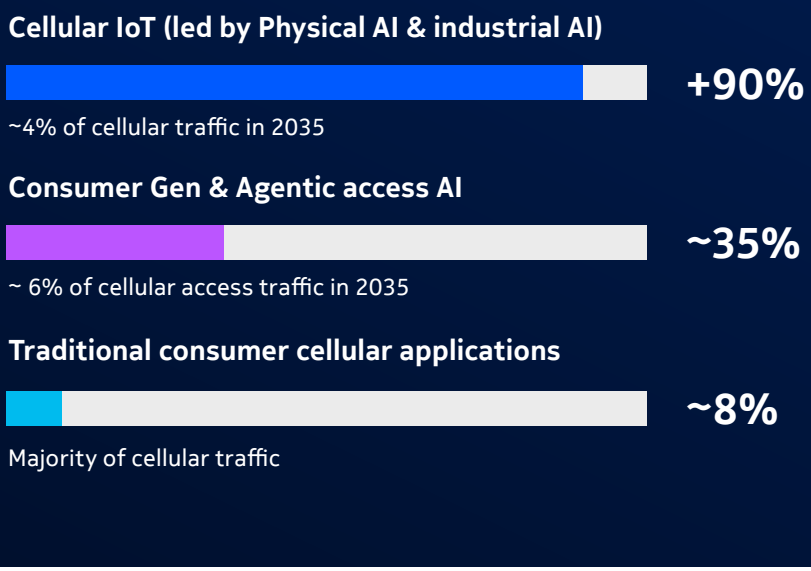
Traditional consumer cellular applications, which account for most cellular traffic, typically generate about 7%-8% uplink traffic, translating into an uplink-downlink ratio (UL: DL) of 1:12. In contrast, consumer GenAI and agentic access AI will have a more varied traffic mix, with uplink accounting for 25%–40% of total traffic. Cellular IoT remains the most uplink-intensive category, with more than 90% of its traffic traveling upstream. Cellular IoT is also shifting away from sensor readings toward data-rich AI workloads. By 2035, most cellular IoT traffic is expected to be AI-driven, led by physical AI systems such as autonomous vehicles, humanoid robots, drones and cloud camera analytics that trigger physical actions, alongside other industrial AI.

The average uplink traffic for a particular area, therefore, depends primarily on the relative contribution of these three categories. Markets with higher concentrations of physical AI and industrial AI, as well as consumer GenAI agentic AI applications, will exhibit comparatively higher uplink shares. These characteristics are most evident in mature markets, where we project aggregate cellular uplink to reach approximately 13% (UL: DL ratio 1:7) of total traffic by 2035. Regions with lower adoption of these uplink-intensive applications will correspondingly exhibit lower uplink shares.

However, relatively stable aggregate uplink proportions can mask increasingly heterogeneous demand at the cell level. There will be urban cells connecting high concentrations of cloud cameras, robotaxis and drones, or cells in pedestrian-heavy areas using AI-enabled visual inferences and other upload-intensive applications. These will become localized uplink hotspots, resulting in significantly higher uplink demand on individual cells versus what we see in regional averages.

13% **Network wide average UL of cellular traffic in mature markets, in 2035**

Uplink share by traffic segments (2035–mature markets)



Busy cells will face uplink capacity challenges

As uplink demand rises it will place increasing stress on the cellular network. To illustrate this trend, it's important to understand the impact of increasing uplink share on a cell's performance.

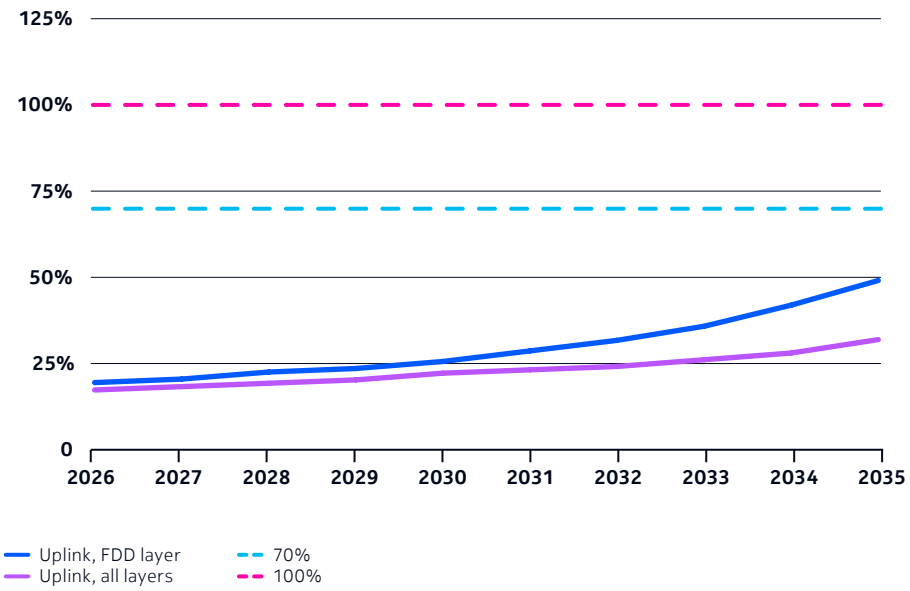
Uplink share describes the traffic mix on a cell, not the pressure placed on the network. What matters is how much of a cell's uplink capacity this traffic uses during busy hours. Above 70% utilization, we see latency and jitter start to rise, impacting many applications including AI.

Uplink capacity* grows more slowly than traffic. While frequency-division duplex (FDD) networks allocate equal spectrum to uplink, most newer networks use time-division duplex (TDD) schemes to allocate spectrum. This generally apportions about one fifth of available bandwidth to the uplink. Consequently, the uplink share of network capacity will continue to fall over the next decade, creating an uplink crunch.

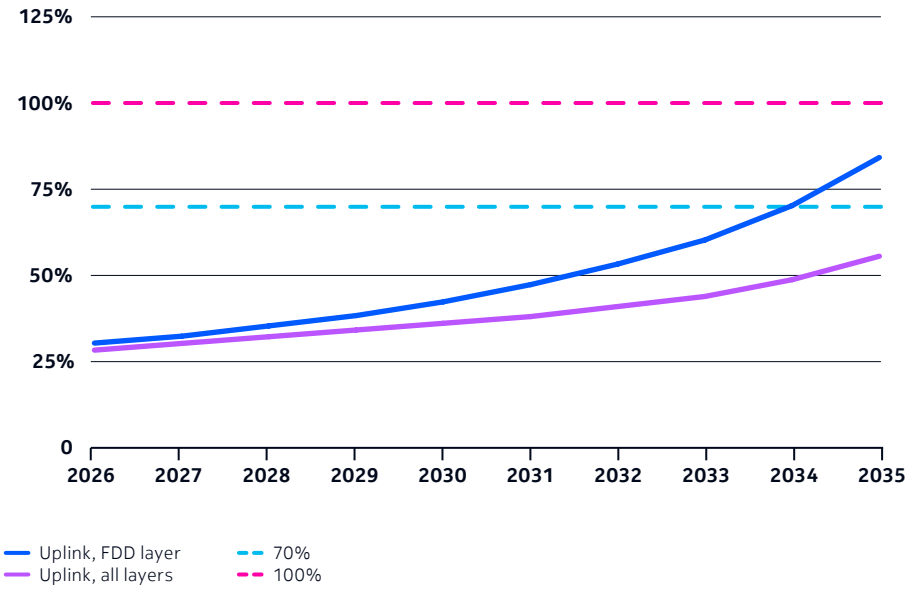
In the average sector, a 13% uplink share in 2035 is still manageable. This translates into an overall uplink load of about 33% capacity and closer to 50% on its FDD layer, which a large portion of devices rely on.

When we look at busy cell sectors, however, these increased uplink demands are far less manageable. A busy sector will reach about 56% overall capacity by 2035, but its FDD layer will breach 70% around 2034, putting the cell in the utilization zone impacting network performance.

Average sector: busy-hour load on capacity



Busy sector: busy-hour load on capacity



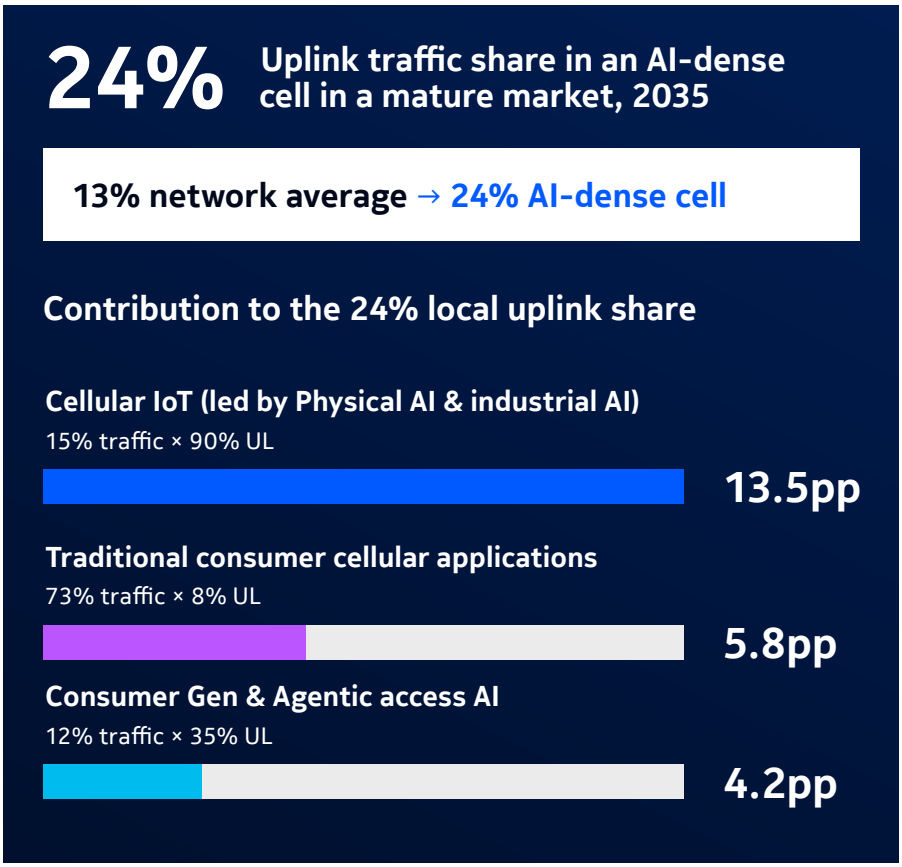
*The capacity model is based on current and anticipated future spectrum holdings for a median mature market operator and include enhancements due to 6G in a staggered manner, based on current visibility. Results may vary for different operators

AI-dense cells can nearly double uplink share to 24%

AI-dense cells are not average cells. Locations with concentrated physical AI and uplink-heavy GenAI services can develop a very different traffic mix. These workloads upload video, sensor data, multimodal prompts, telemetry and execution feedback.

To illustrate the impact of AI-dense cells we created a mature-market dense-cell scenario. In this scenario, cellular IoT — including emerging physical AI workloads — will reach 15% of traffic, while consumer GenAI and agentic AI will account for 12%. This is much higher than what we see at the average cell. With uplink shares of approximately 90% and 35% respectively, this mix raises the overall uplink share from around 13% to approximately 24% (UL: DL ratio 1:3) locally.

The implications of this are networks face a concentration challenge rather than a uniform shift of traffic. Moreover, simultaneous video uploads, sensor events, live visual-AI sessions and machine interactions can create short, synchronized traffic bursts, pushing instantaneous uplink demand beyond the level implied by the 24% average. Capacity planning must therefore account for both AI-intensive locations and burst-driven local peaks rather than relying on aggregate traffic averages.



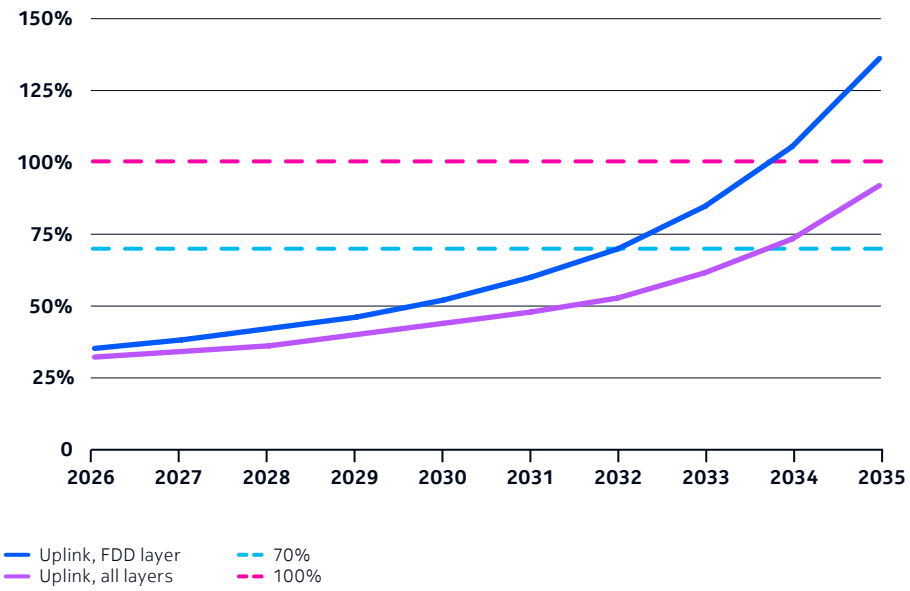
AI-dense sectors get congested faster

A 24% uplink share in a cell highlights a real capacity problem. In AI-dense sectors, average uplink load will reach about 92% of capacity by 2035. Meanwhile the FDD layer in AI-dense sectors will cross the 70% threshold around 2033 and become completely full by 2034.

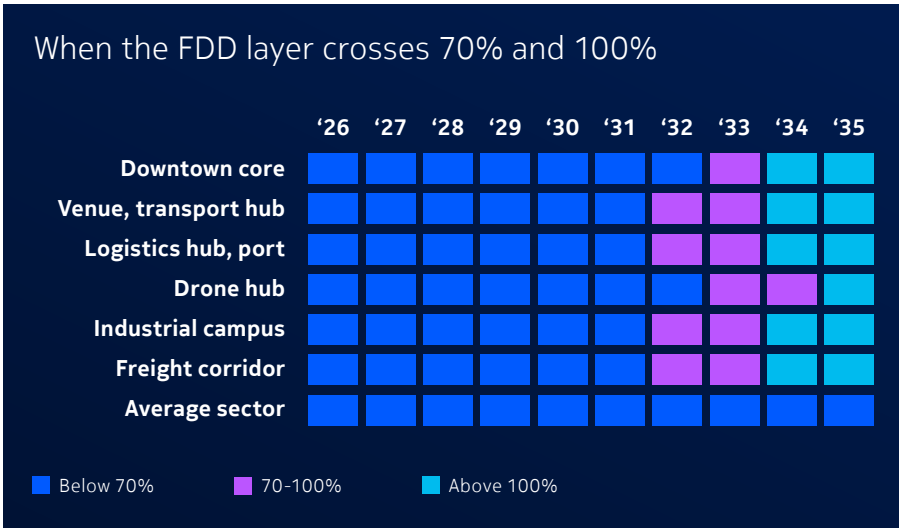
As the FDD layer will be the hardest hit, this could have significant consequences for physical AI and industrial IoT applications. Robotaxis, humanoid robots, drones and IoT modules depend largely on FDD. Without additional capacity-related measures these machines would be severely impacted. Venues, transport and logistics hubs and ports, drone hubs, industrial campuses and freight corridors could potentially breach their uplink capacities between 2032 and 2033.

Uplink capacity will need to be revisited in these locations. Key mitigation strategies include, but are not limited to, spectrum refarming, enhanced uplink optimization, new cell sites, cell splitting and AI-RAN.

AI-dense downtown cell (24% share: load on capacity)



When the FDD layer crosses 70% and 100%



7

Global access traffic

Global access traffic will grow 3x to 5x in the next decade. This growth will be dominated by familiar applications, such as video, social media and workplace productivity tools. Enterprise traffic will increase more quickly than consumer traffic, but AI is the fastest growing category.

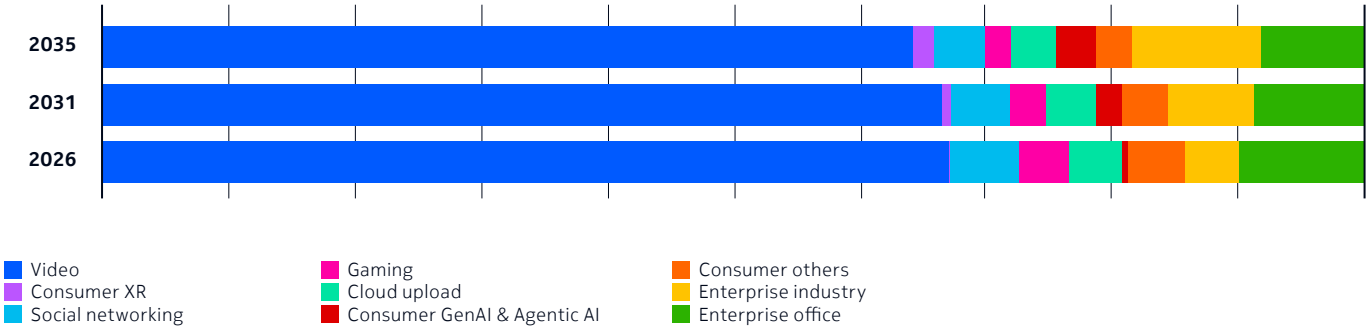
Global access traffic trends

The world's access networks are entering a decade of sustained, broad-based growth. We project global access traffic to increase three- to five-fold by 2035, reaching between 2,498 EB/month and 4,657 EB/month. In our moderate scenario, access traffic will climb to 3,242 EB/month, at a steady rate of 15% a year.

Every category of traffic will contribute to this growth. Consumer fixed will remain the largest share at 2,121 EB per month, while consumer mobile will reach 524 EB/month. Enterprise will grow faster than both, reaching 597 EB/month. Entertainment, led by video and social media, will continue to fill the network, while consumer GenAI and agentic AI will emerge as the fastest-rising traffic category.

The total numbers, though, are only part of the picture. Growth will not be distributed evenly across all networks. On the consumer side, mobile, fixed wireless and fiber traffic will increase at varying rates. Meanwhile, emerging and developed markets will grow at different paces. Looking beyond the consumer, the office will remain the foundation of enterprise traffic, but industry will steadily claim a larger share.

Global access traffic share by services



AI spans several service categories
Consumer GenAI and Agentic AI are shown separately but represent only part of access AI traffic. (See the AI sections for details.)

Global access traffic to grow between 3x and 5x through 2035

Total traffic in 2035 will reach

2,498–4,657

EB/month

Our moderate projections put traffic at

3,242 EB/month in 2035, representing a healthy CAGR of

15%

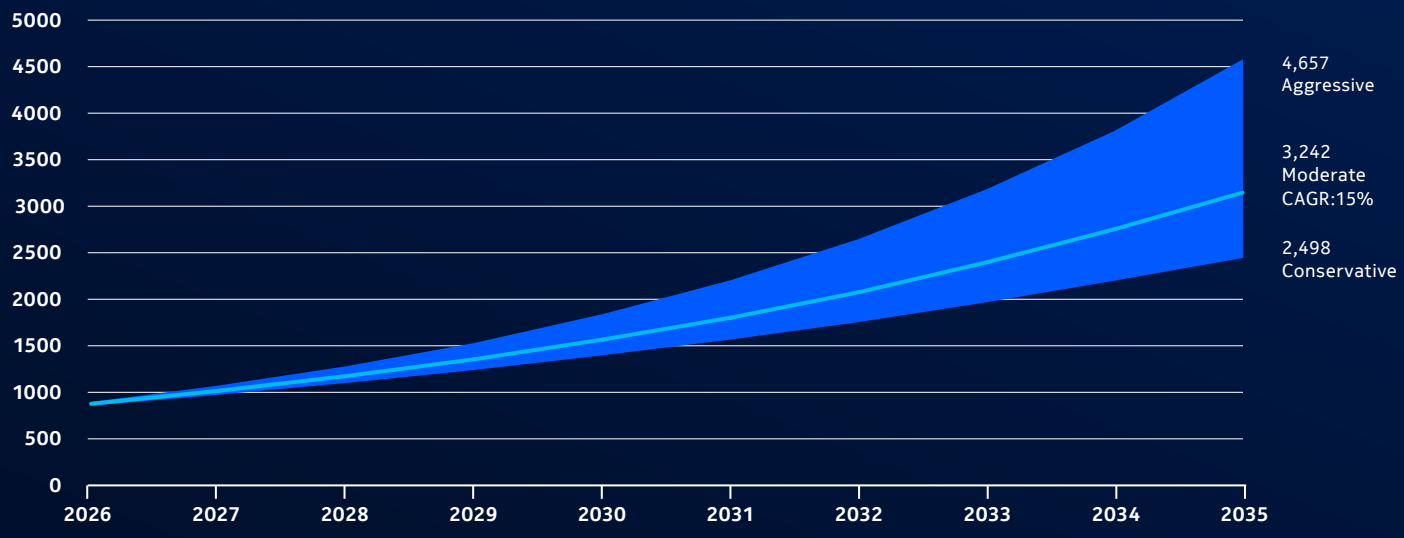
Projections for 2035

2,121
EB/month
Consumer fixed
15% CAGR, includes FWA

524
EB/month
Consumer mobile
14% CAGR

597
EB/month
Enterprise
18% CAGR

Global access traffic (EB/month)



8

Consumer access traffic

Consumer access traffic has always been the bedrock of the internet, driven by our increasing appetite for entertainment, information and community. Many of the trends driving this growth are familiar: the proliferation of video streaming, rising home-broadband adoption and increasing data consumption per user. However, we see that new sources of demand are emerging.

Nowhere will those shifts be more apparent than on the cellular network. By 2035, GenAI, agentic AI and XR will provide new avenues of mobile traffic growth, particularly in developed markets. At the same time, “cellular” will no longer be synonymous with “mobile” as fixed wireless access assumes a third of overall consumer cellular traffic.

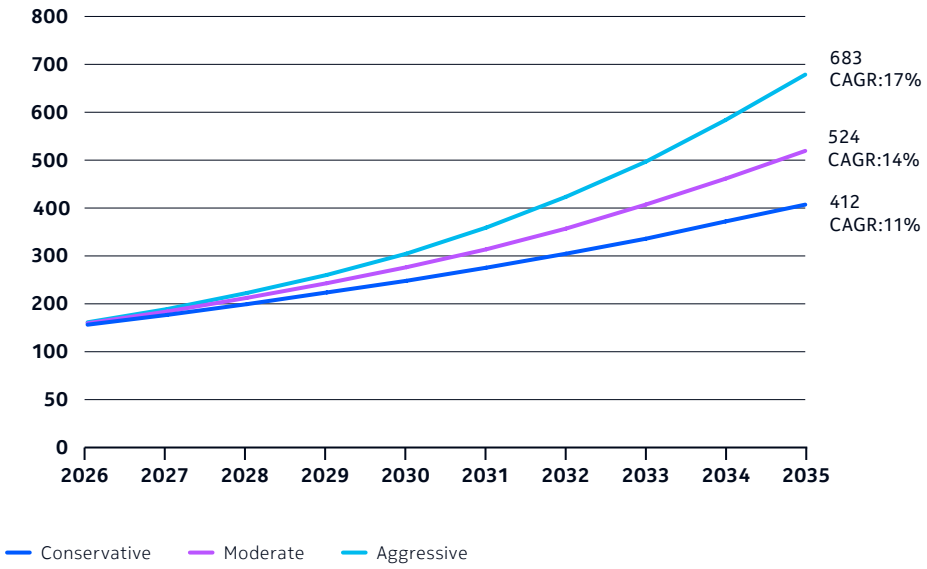
Video and social will continue to dominate consumer mobile traffic, but AI's contribution will grow

Under our moderate projections, consumer mobile traffic will reach 524 EB/month by 2035, growing at a CAGR of 14%.

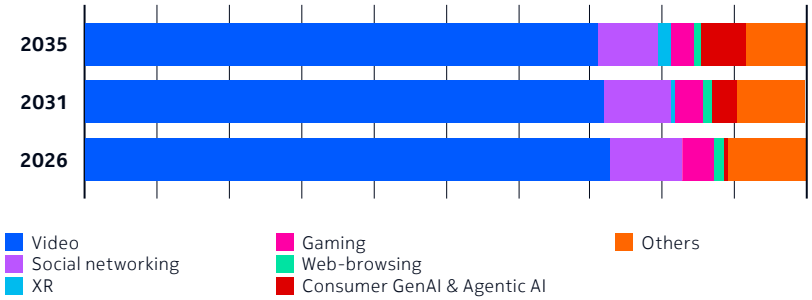
Video streaming and social networking will retain the largest share of consumer mobile traffic as consumption continues to rise. Video growth will be supported by expanding connectivity and service adoption in emerging markets, increasing consumption per user, and a shift toward higher-quality content viewed on larger, higher-resolution smartphone displays. GenAI will further fuel this trend, encouraging video content creation and consumption across categories like entertainment, education and gaming.

Consumer GenAI and agentic AI will be among the fastest-growing application categories, supported by new capabilities and deeper integration into smartphones and everyday applications. AI usage is projected to increase faster on smartphones than on PCs, while XR and cloud gaming will create further avenues for consumer mobile traffic growth.

Consumer mobile access traffic (EB/month)



Consumer mobile access traffic share by services



AI spans several service categories
Consumer GenAI and Agentic AI are shown separately but represent only part of access AI traffic. (See the AI sections for details.)

Traditional media still fills the pipe

Streaming and social networking will retain the largest share of consumer mobile traffic through 2035, while consumer GenAI and agentic AI are the fastest-growing categories (CAGR: 42%).

Sustained double-digit mobile traffic growth across all eight regions

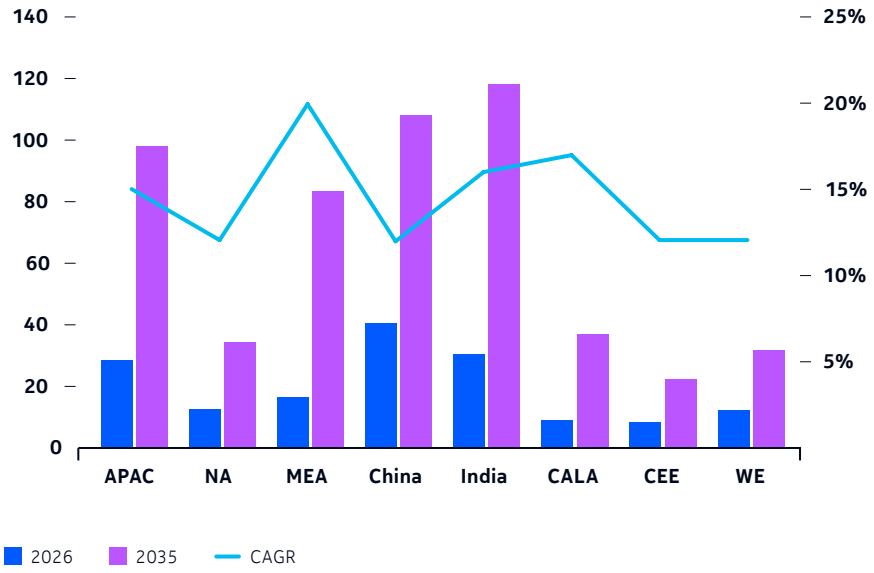
We see significant variation in consumer mobile access traffic growth across the eight regions we track, reflecting differences in market maturity, smartphone penetration and digital-service adoption.

Emerging regions are projected to record higher growth rates as both mobile internet penetration and data consumption per user continue to increase. This is further amplified by a higher reliance on mobile networks as the primary means of data consumption, given limited fixed broadband availability. These markets will rely primarily on traditional services like video streaming and social networking to drive consumer mobile access traffic over the next decade, as service adoption expands with increasing market penetration.

In contrast, developed markets like North America and Western Europe are expected to grow more moderately, toward higher usage per subscriber rather than new user additions. Here, higher adoption of GenAI and agentic AI, along with other next-generation services like XR and cloud gaming, will fuel the upward trend in mobile data consumption.

India, Asia Pacific (APAC) and the Middle East and Africa (MEA) will remain the largest contributors to global consumer mobile traffic growth, supported by large subscriber bases, expanding 5G coverage and increasing consumption of video and AI-enabled services. India is expected to surpass China early next decade in mobile data traffic because of a greater reliance on the mobile network for data consumption. China will continue to generate one of the highest traffic volumes globally, though the maturity of its mobile market will moderate its growth.

Consumer mobile access traffic and CAGR by region (EB/month, %)



12%–20% CAGR, depending on region

Higher growth in emerging regions like India, APAC, and MEA will be fueled by expanding 5G, rising data consumption and limited fixed broadband. Developed markets will see moderate growth from higher per-user usage.

FWA will account for more than one third of consumer cellular traffic in 2035

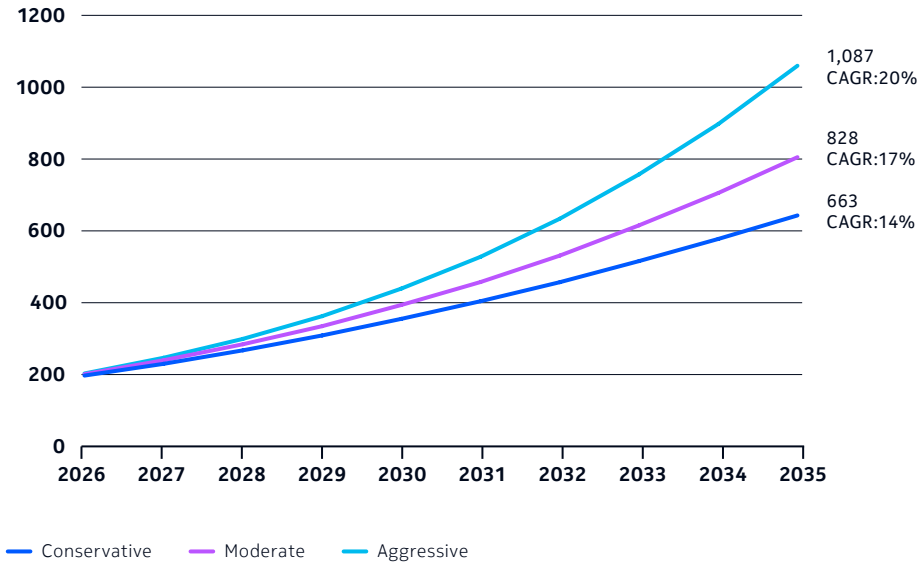
Cellular carries more than just mobile traffic. Fixed wireless access (FWA) is increasingly driving up cellular data consumption — a trend that will become more apparent over the next decade. Under our moderate-scenario projections, consumer FWA will carry 304 EB/month, boosting total consumer cellular traffic to 828 EB/month in 2035. This means FWA will account for more than a third of all consumer cellular data.

APAC is expected to generate over half of global FWA traffic, supported by large-scale deployments across the region, particularly in high-growth markets such as India. North America is projected to be the second-largest contributor, accounting for approximately 17% of global FWA traffic, supported by significantly higher traffic consumption per FWA connection.

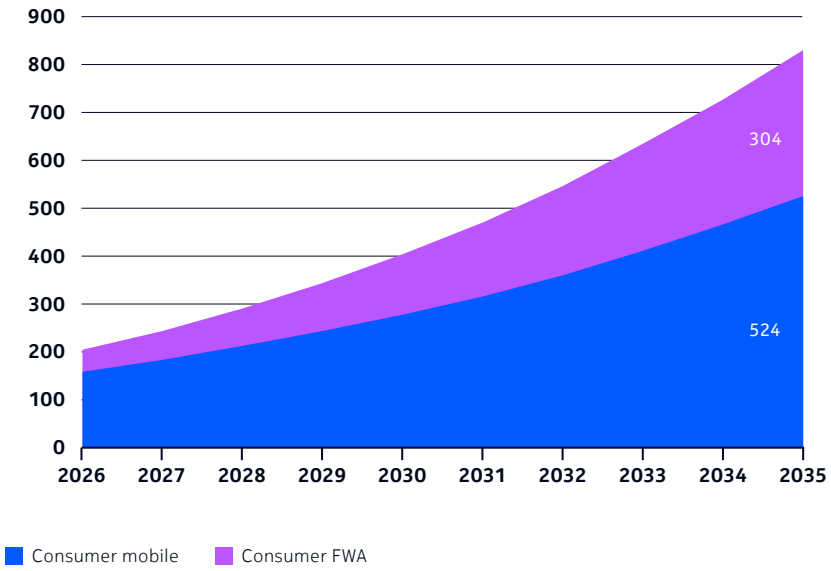
17%
Consumer cellular CAGR including FWA

>1/3
FWA share in consumer cellular traffic in 2035

Consumer cellular access traffic including FWA (EB/month)



Consumer cellular access traffic including FWA (EB/month)

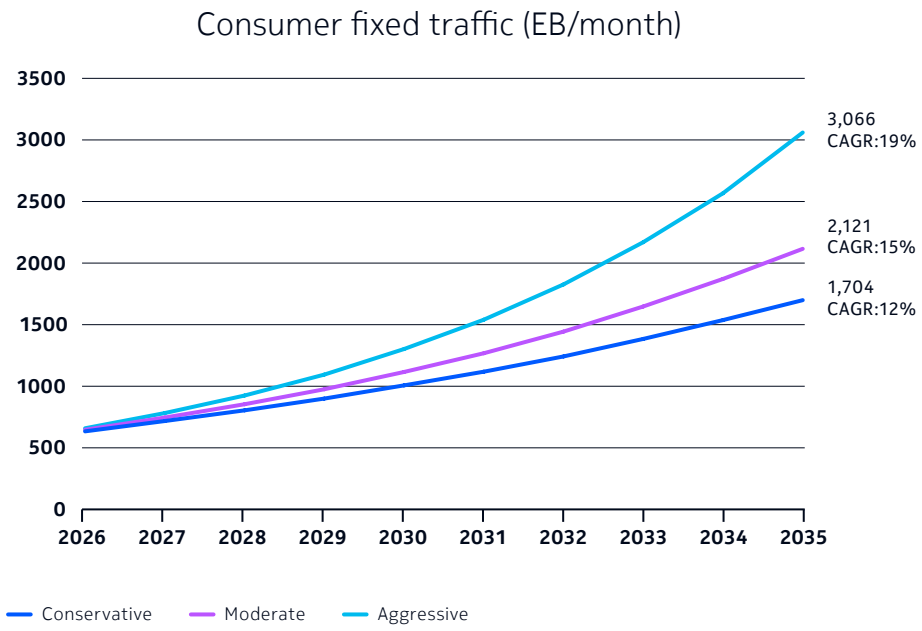


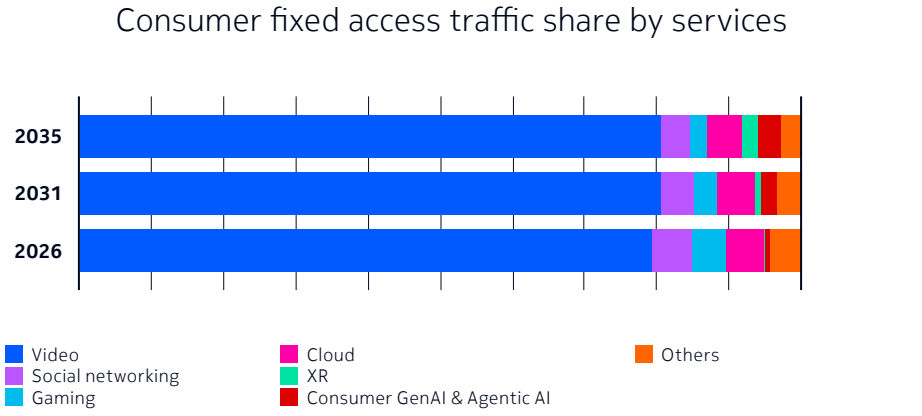
OTT and expanding broadband access will drive consumer fixed data traffic beyond 2,000 EB/month

According to our moderate-scenario projections, consumer fixed traffic, including FWA, will reach 2,121 EB/month by 2035, growing at a 15% CAGR. Growth will be driven by broadband expansion in emerging markets and rising data consumption per household globally. Fiber deployment, higher broadband speeds, improved home Wi-Fi, larger screens and higher video resolutions will continue to increase traffic demand.

The shift from broadcast television to OTT streaming is one of the key drivers of fixed broadband traffic. Unlike broadcast TV, OTT delivers an individual IP stream to each connected device. At the same time, viewing habits are shifting from shared family viewing to personalized, on-demand consumption, with multiple household members increasingly streaming different content simultaneously across smart TVs and Wi-Fi-connected personal devices. Combined with expanding content libraries and higher-quality content, these trends will continue to drive downstream traffic despite ongoing improvements in video compression technologies.

Video will remain the dominant contributor to fixed broadband traffic. However, the traffic mix will gradually diversify as AI applications, cloud storage and synchronization, gaming, home surveillance and other connected household services account for a growing share of total traffic.





AI spans several service categories
Consumer GenAI and Agentic AI are shown separately but represent only part of access AI traffic. (See the AI sections for details.)

Regional consumer fixed traffic: China maintains scale while APAC delivers large incremental traffic growth

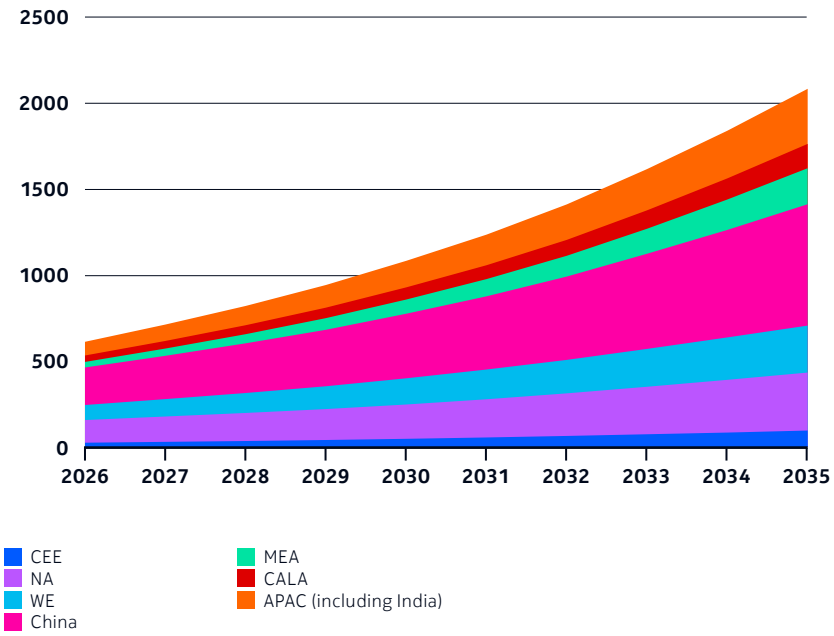
China is expected to remain the largest contributor to global consumer fixed access traffic throughout the forecast period, supported by near-universal FTTH deployment, high broadband penetration, and organic growth in household data consumption.

North America will retain the second-largest traffic volume, driven by the world's highest traffic consumption per broadband subscription. This reflects a combination of high OTT penetration, a large share of video viewing on internet-connected TVs, one of the highest concentrations of connected devices per household, widespread cloud sync service adoption, and a greater prevalence of individual streaming across multiple screens.

While higher growth rates are projected in emerging regions, particularly the Middle East and Africa (MEA), the Asia Pacific (APAC) region, including India, will account for a larger share of incremental global traffic growth owing to its substantially larger broadband subscriber base, increasing household data consumption and cheap data subscriptions.

The regional traffic mix will become progressively more balanced over the forecast period as emerging markets expand their contribution to global traffic. Fiber deployment and broadband adoption will continue to accelerate across emerging markets, narrowing the gap with mature regions.

Regional share: consumer fixed access traffic (EB/month)



9

Enterprise access traffic

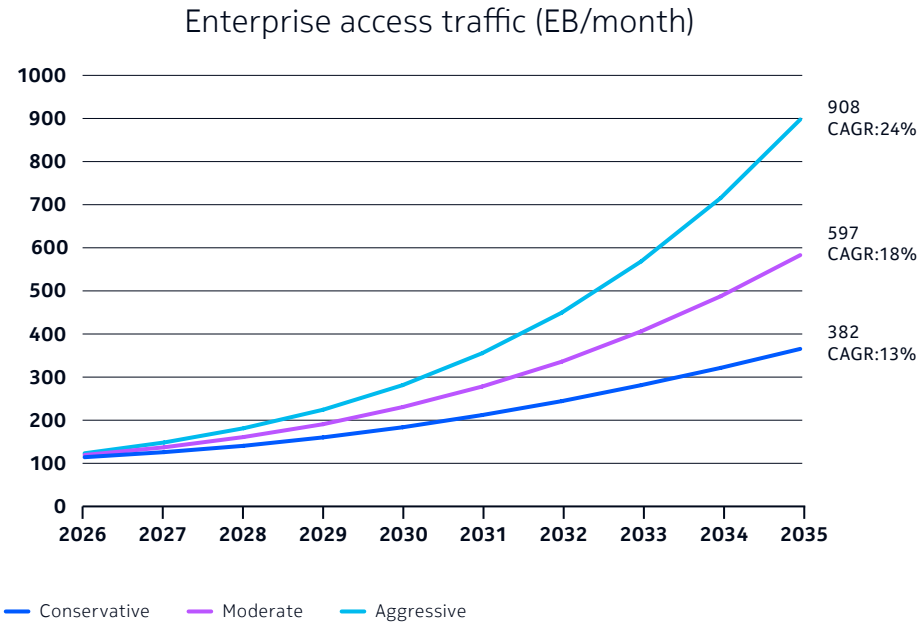
In the enterprise, we will see a significant shift in the balance of access traffic, as industrial traffic grows faster than office traffic. Driven by expanding IoT, more data-intensive connected systems and growing XR adoption, industrial traffic is set to overtake enterprise office traffic in the next decade. However, there will be plenty of growth in the enterprise office space as well. Increased AI usage, a larger digital workforce and richer digital workflows will nearly triple enterprise office traffic in the next 10 years.

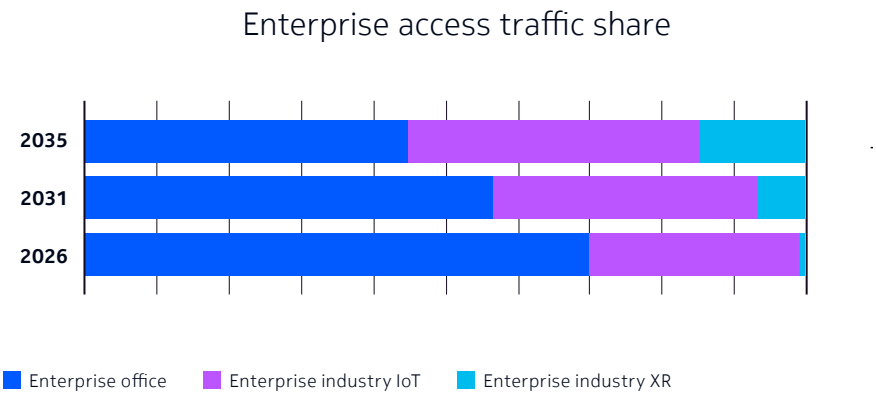
Industrial IoT and XR will account for half of enterprise traffic

Enterprise traffic is projected to grow steadily throughout the forecast period as organizations continue to digitize business processes, modernize enterprise IT and expand connectivity across operational environments. But the faster growth will come from industrial applications. Expanding IoT deployments, rising traffic per device and increased AR and VR adoption will give a sizable lift to industrial-application usage. Combined, IoT and XR will account for 55% of enterprise access traffic by 2035, exceeding enterprise office traffic.

Enterprise office traffic will by no means decline. Growth will remain robust, reflecting continued reliance on cloud-based business applications, collaboration services and digital workflows.

Owing to the pace of enterprise digital transformation, we project that global enterprise traffic will reach between 382 EB/month (based on our conservative model) and 908 EB/month (our aggressive model) by 2035, corresponding to a CAGR of 13–24% over the forecast period.





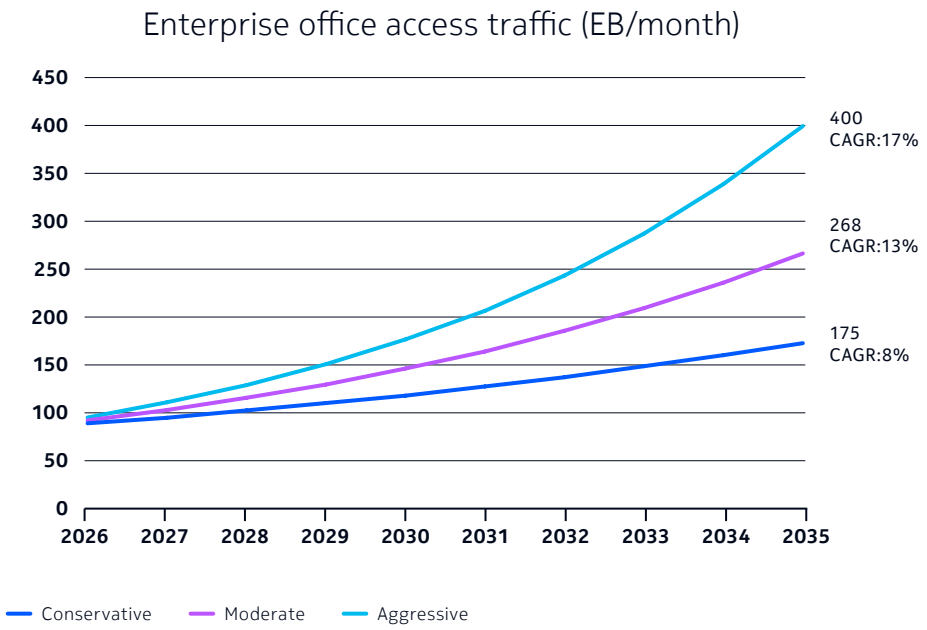
Advanced connectivity will power digital offices and connected industries

Enterprise traffic is projected to grow steadily (at a 13–24% CAGR), driven by digital transformation, cloud-based applications and rapid industrial connectivity expansion across diverse sectors.

Enterprise office traffic outlook 2026–2035

Over the next decade, enterprise office access traffic will change in both volume and character. Under our moderate-growth scenario, we project enterprise office traffic will grow to 268 EB/month in 2035, growing at a 13% CAGR.

These increases are driven by two factors. First, the office workforce will return to pre-pandemic levels. This, coupled with the trend of increased digitalization in the workplace, will result in increased traffic per employee. Second, AI will drive significant traffic growth. We expect enterprise AI traffic to balloon in the next decade due to wider adoption, more usage time and richer per-session payloads. Specifically, AI traffic will increase to 11 times its 2026 levels by 2035 with a CAGR of 31%. While non-AI office workloads will continue their upwards trajectory, office AI is growing 3x faster.



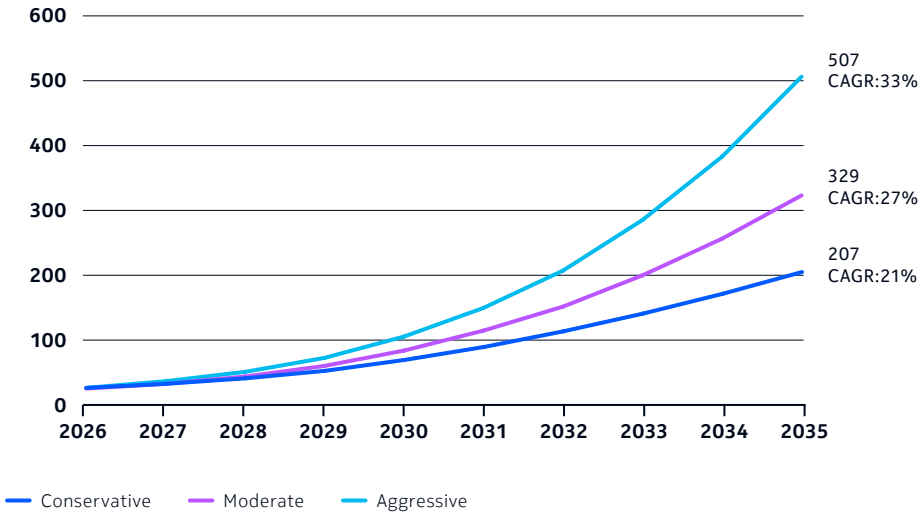
The proliferation of high-bandwidth devices will lead to big spike in industrial traffic

Our moderate-scenario projections show global industrial access traffic increasing at a CAGR of 27% reaching 329 EB/month by 2035, driven by growth in connected endpoints and increasing traffic intensity per device. This traffic spurt will be concentrated among higher-bandwidth endpoints — particularly security and monitoring cameras, outdoor and transport surveillance systems, body-worn and imaging devices, autonomous vehicles, drones, and humanoid robots — which generate a substantial share of overall volume. Sensors, meters and tags will continue to dominate the installed base but contribute less traffic because of their lower per-device data exchange.

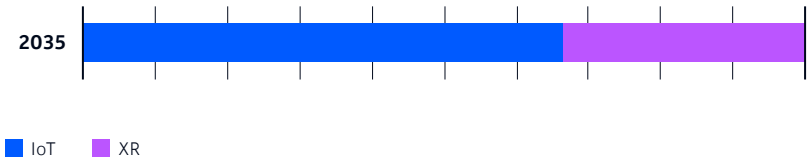
XR, including both AR and VR, will be the fastest growing segment of industrial traffic. Though AR and VR apps contribute minuscule traffic today, by 2035, we project nearly 25% of all industrial traffic will originate or be delivered to an XR device.

VR is the heavier traffic generator, with use cases such as immersive training, simulation and design review requiring high, sustained bit rates on higher-fidelity standalone and tethered headsets. AR is the broader, lighter layer, providing remote assistance, hands-free guidance, and digital-work instructions across field and factory operations. It's worth noting that while AR produces lighter per-stream traffic loads today, its formats are rapidly becoming richer and more immersive. This means per-device traffic on AR devices will only keep climbing.

Enterprise industry access traffic (EB/month)



Enterprise industry traffic: IoT and XR share



A handful of device types drive the bulk of industrial IoT volume

Cameras, vehicles, robots, drones — not the billions of sensors and meters — will generate the bulk of industrial IoT volume, with immersive XR forming a distinct bandwidth-intensive layer of its own.

10 Methodology

How we forecast is just as important as what we forecast. The following two pages detail our methodology. First, we describe how we model data and the sources from which we collect our data. Next, we provide a more granular look at the reference architecture we use as a framework for our analysis in chapter 6.

How we forecast

Nokia's Global Network Traffic Report projects network traffic from 2026 to 2035 to help communications service providers, hyperscalers and neoclouds plan their networks. It covers traffic exchanged with access networks by consumers and enterprises, and the cloud AI traffic that this access traffic triggers, with a dedicated view of how AI is reshaping demand.

The forecast is built bottom-up. Consumer and enterprise traffic is modeled by application across eight regions, based on adoption, usage and the volume of data exchanged per session. Industrial and physical AI traffic is modeled using the installed base of connected devices and machines and the traffic generated by each. Agentic AI traffic included in the forecast is estimated based on the frequency of agentic workflow execution and the data exchanged during each workflow. The forecast draws on third-party industry and analyst research and on public data from operators, vendors and regulators, and is reviewed by Nokia subject-matter experts. Nokia's controlled lab tests of AI applications provide insights into traffic characteristics, including uplink-downlink mix and burstiness. The forecast also assesses where AI traffic is served across the network and how increasing uplink demand could constrain capacity in the busiest locations.

The forecast is developed under three demand scenarios — conservative, moderate and aggressive — with moderate as the headline scenario. The full range is presented for global access traffic and selected results, illustrating how outcomes vary under different assumptions. The forecast provides a repeatable framework that is updated annually as adoption, technology and deployment patterns evolve.

Our reference architecture: How AI traffic distributes across the network

We model the distribution of AI traffic across telecom and cloud infrastructure domains by identifying where the traffic is served, which infrastructure domain hosts the serving location and the network paths traversed by this traffic.

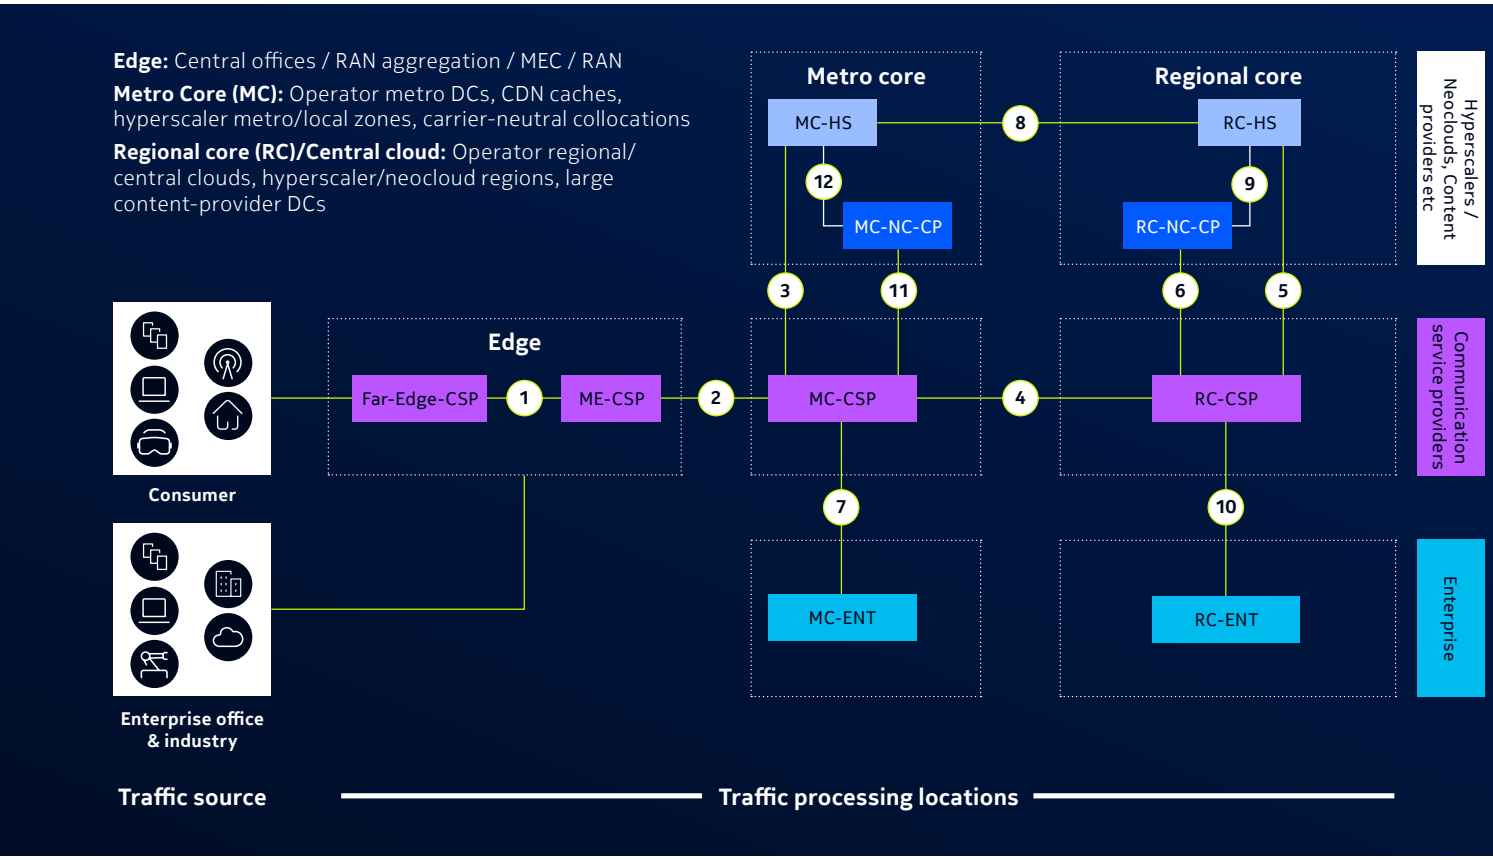
This chart shows a reference architecture that provides a framework for modeling AI traffic. It comprises three network tiers — edge, metro core and regional core/central cloud — along with four infrastructure provider categories: 1) hyperscalers (HS), 2) neoclouds, content providers and others (NC-CP), 3) communication service providers (CSPs), 4) and enterprises (ENT). This architecture allows us to characterize AI traffic along three dimensions: the network tier serving the traffic, the infrastructure provider hosting the AI service and the network links traversed.

AI service placement is determined through a byte-weighted assessment of AI deployment architectures, considering workload characteristics, latency and compute requirements, compute availability, data sovereignty and cacheability.

Each traffic flow is mapped to the network tier serving the traffic and the provider hosting the AI service. Based on this placement, the methodology determines the network links traversed (IDs 1–12) and the traffic volume carried on each link.

Aggregate network traffic is calculated across all modeled links. As a result, a single AI traffic flow may contribute to the traffic volume on multiple links along its network path.

A special case is cloud AI traffic, which represents traffic generated among cloud-based models, tools and data sources to support AI workflows, including both agentic and conventional, non-agentic processing. Cloud AI traffic —may traverse multiple network links and contribute to aggregate network traffic, depending on the location of the participating services across infrastructure providers and network tiers.



Nokia OYJ
Karakaari 7
02610 Espoo
Finland

Tel. +358 (0) 10 44 88 000

CID:215578

nokia.com



About Nokia

Nokia is a global leader in connectivity infrastructure, enabling the AI supercycle. We build networks that connect intelligence and deliver a comprehensive connectivity solution across fixed, mobile, IP, optical, core, and data center networks. With security by design, AI-driven innovation, and AI-optimized architecture, Nokia delivers trusted infrastructure that enables AI to scale and advances connectivity to secure a brighter world.

Nokia is a registered trademark of Nokia Corporation. Other product and company names mentioned herein may be trademarks or trade names of their respective owners.

© 2026 Nokia