NOKIA

Global network traffic report

Understanding the growing impact of advancing technologies on future networks

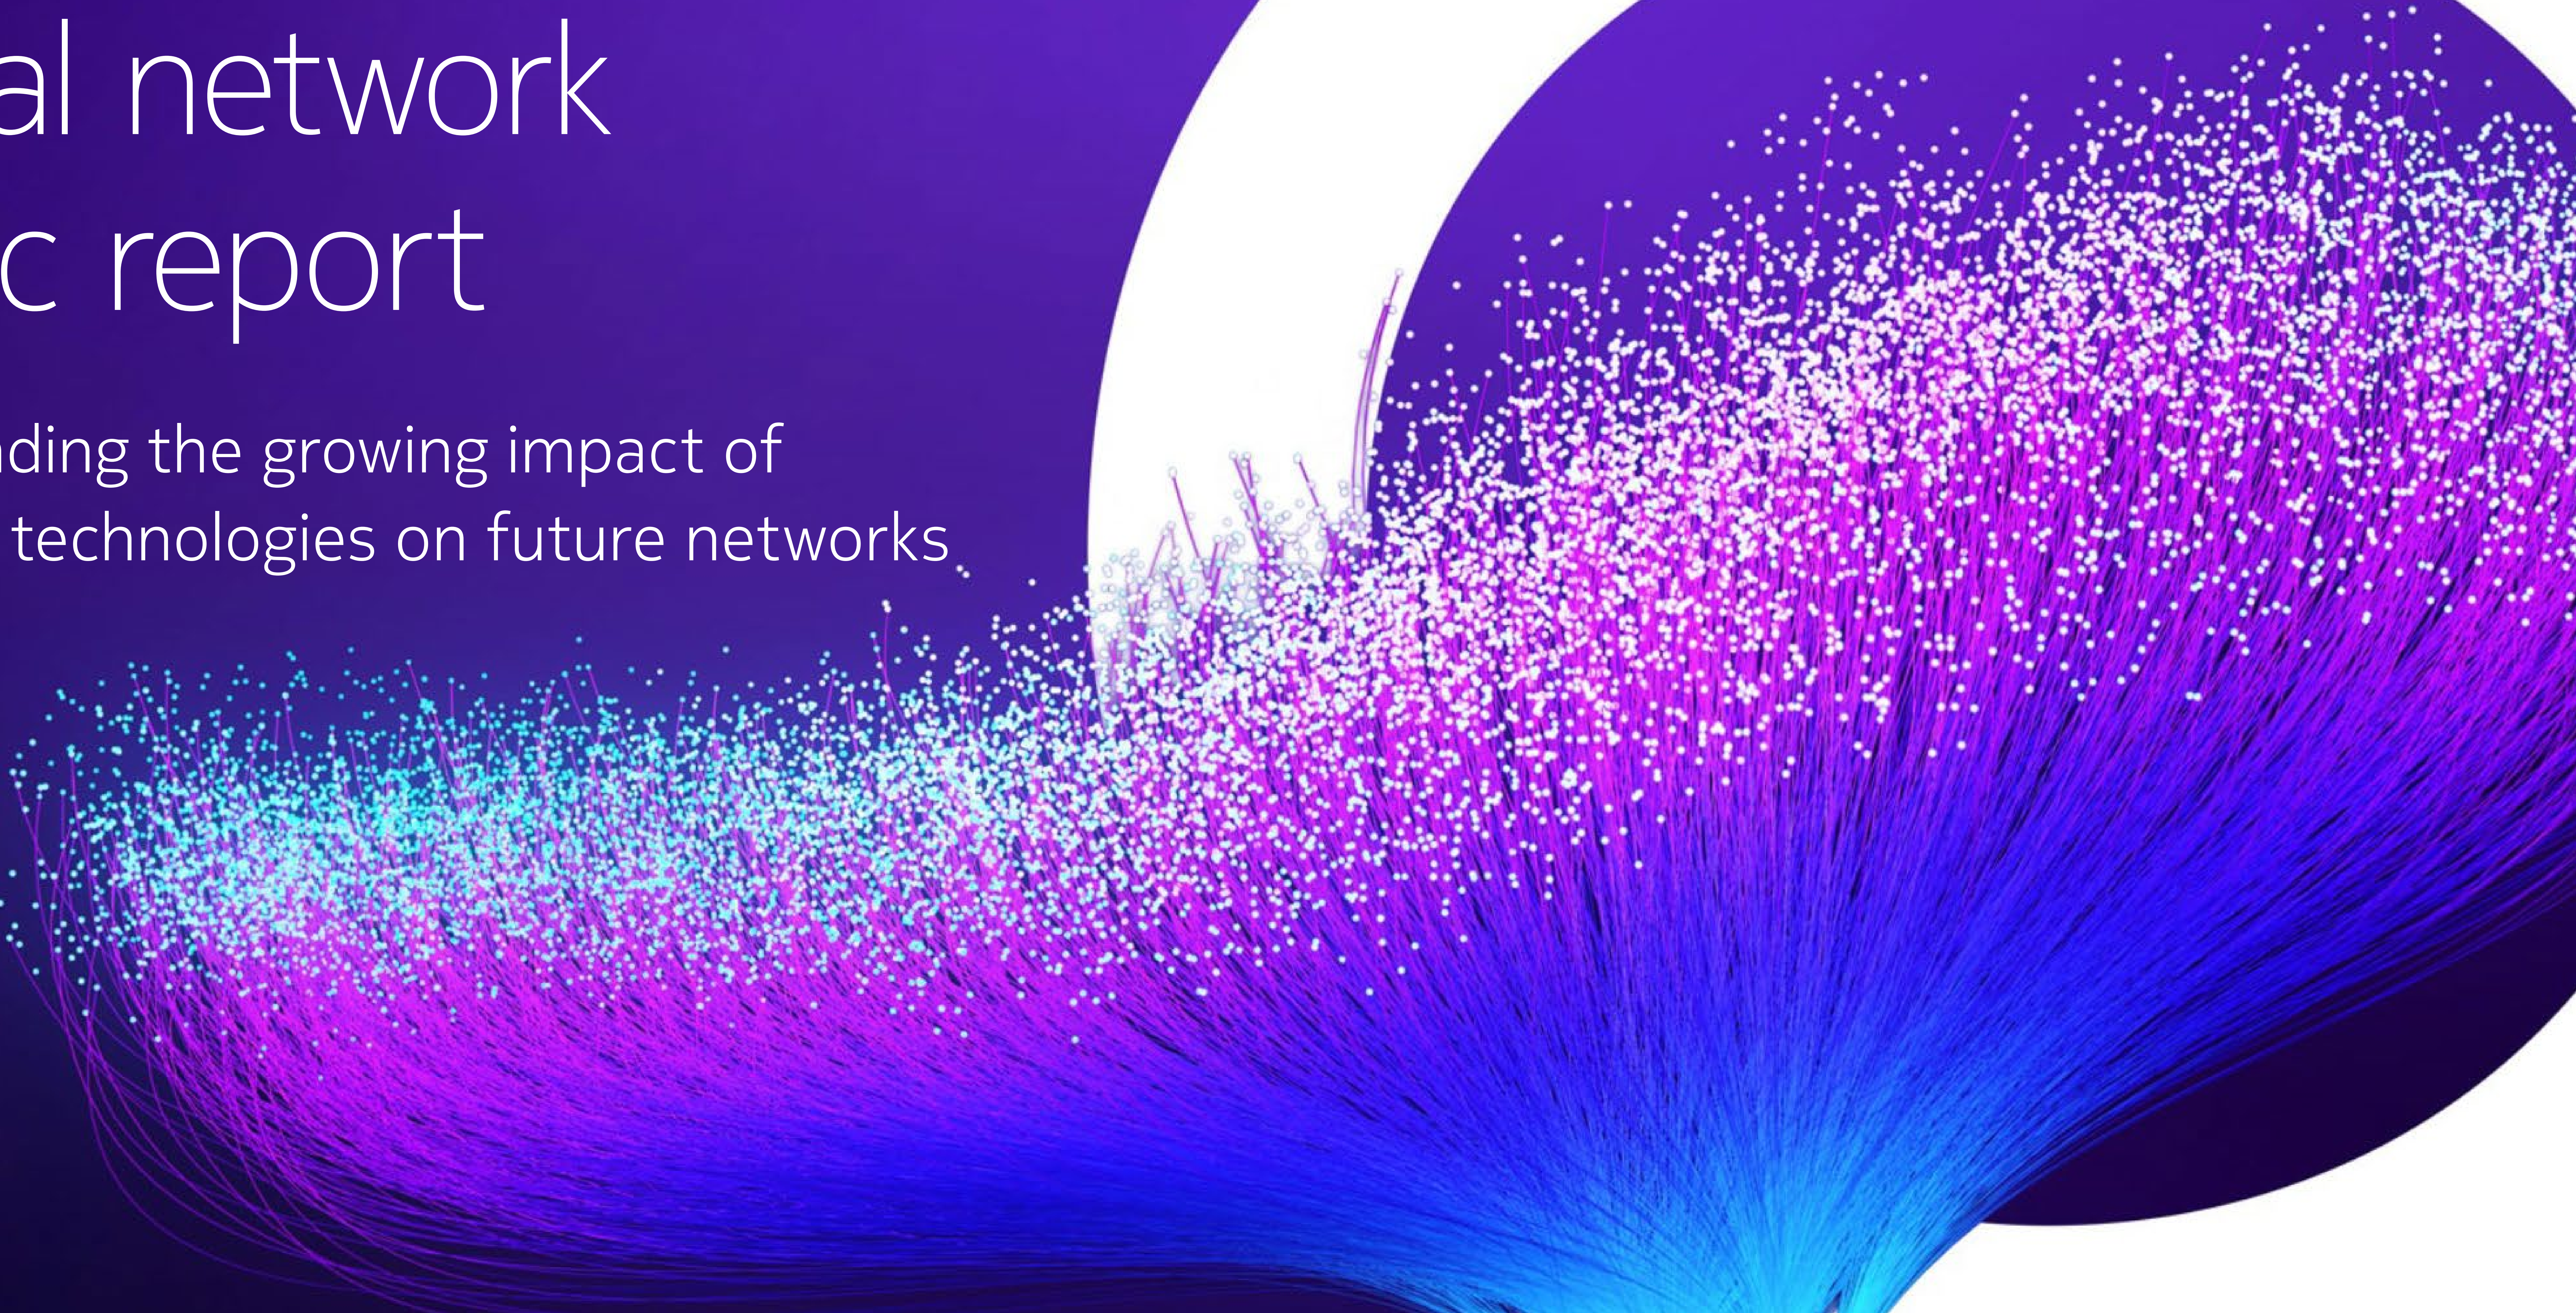

Introduction

Global network traffic is changing in character, not just in volume. It is no longer shaped only by people streaming video over home broadband. Today, it is being driven by three converging forces: consumers are shifting toward more immersive, interactive and increasingly uplink-heavy uses; enterprises and industrial sites are extending operations over wide-area connectivity rather than keeping activity on-site; and AI systems are generating traffic autonomously, with machine-to-machine exchanges flowing across metro and core networks and between large compute locations.

At the heart of this transformation is the rise of connected intelligence—a distributed fabric of humans, devices and AI systems continuously sensing, learning, and acting across networks. This connected intelligence amplifies traffic diversity, making it less asymmetric, more latency-sensitive and increasingly dependant on AI-driven automation.

This report explores how these forces are redefining where traffic flows, how it behaves, and where networks will feel the greatest pressure first. Its goal is not only to describe growth but also to explain why the nature of that growth is changing—and why that difference matters for operators, cloud providers, and large enterprises as they plan capacity and architecture for the decade ahead.

Explore the findings

Our methodology

Nokia, through its Bell Labs Consulting arm, forecasts global network traffic on an annual basis to assist network providers in their decision-making and planning. Our forecast covers from 2024 through 2034. It focuses exclusively on traffic that traverses service provider and enterprise wide-area network (WAN) infrastructure, including public internet, managed and private WAN services, and movement between data centers. It does not attempt to forecast traffic that stays inside a private, single campus or local network, because the scaling, routing, resiliency and interconnection requirements are different.

We model traffic in three domains. The first is consumer, which includes both mobile access and fixed access, including fixed wireless access. The second is enterprise and industrial, which covers wide-area connectivity that supports knowledge work, automation, machine vision, robotics coordination, field support and industrial IoT.

The third is AI, which includes not only AI that people directly invoke. For example, assistants, copilots or media generation in which AI systems trigger other AI systems and move data autonomously across networks.

The forecast is presented as a set of demand scenarios rather than a single prediction, reflecting variations in key drivers: the pace of AI assistance adoption, the transition of industrial processes to wide-area networks, and the evolution of consumer applications toward continuous, interactive engagement.

This year's report defines three scenarios - conservative, moderate, and aggressive. Our goal is to present scenarios that fall within a realistic range of possible outcomes, encouraging stakeholders to plan across the full spectrum of high-impact demand possibilities.

Total global traffic projected to grow 3-7x through 2034

Global WAN traffic, EB/month

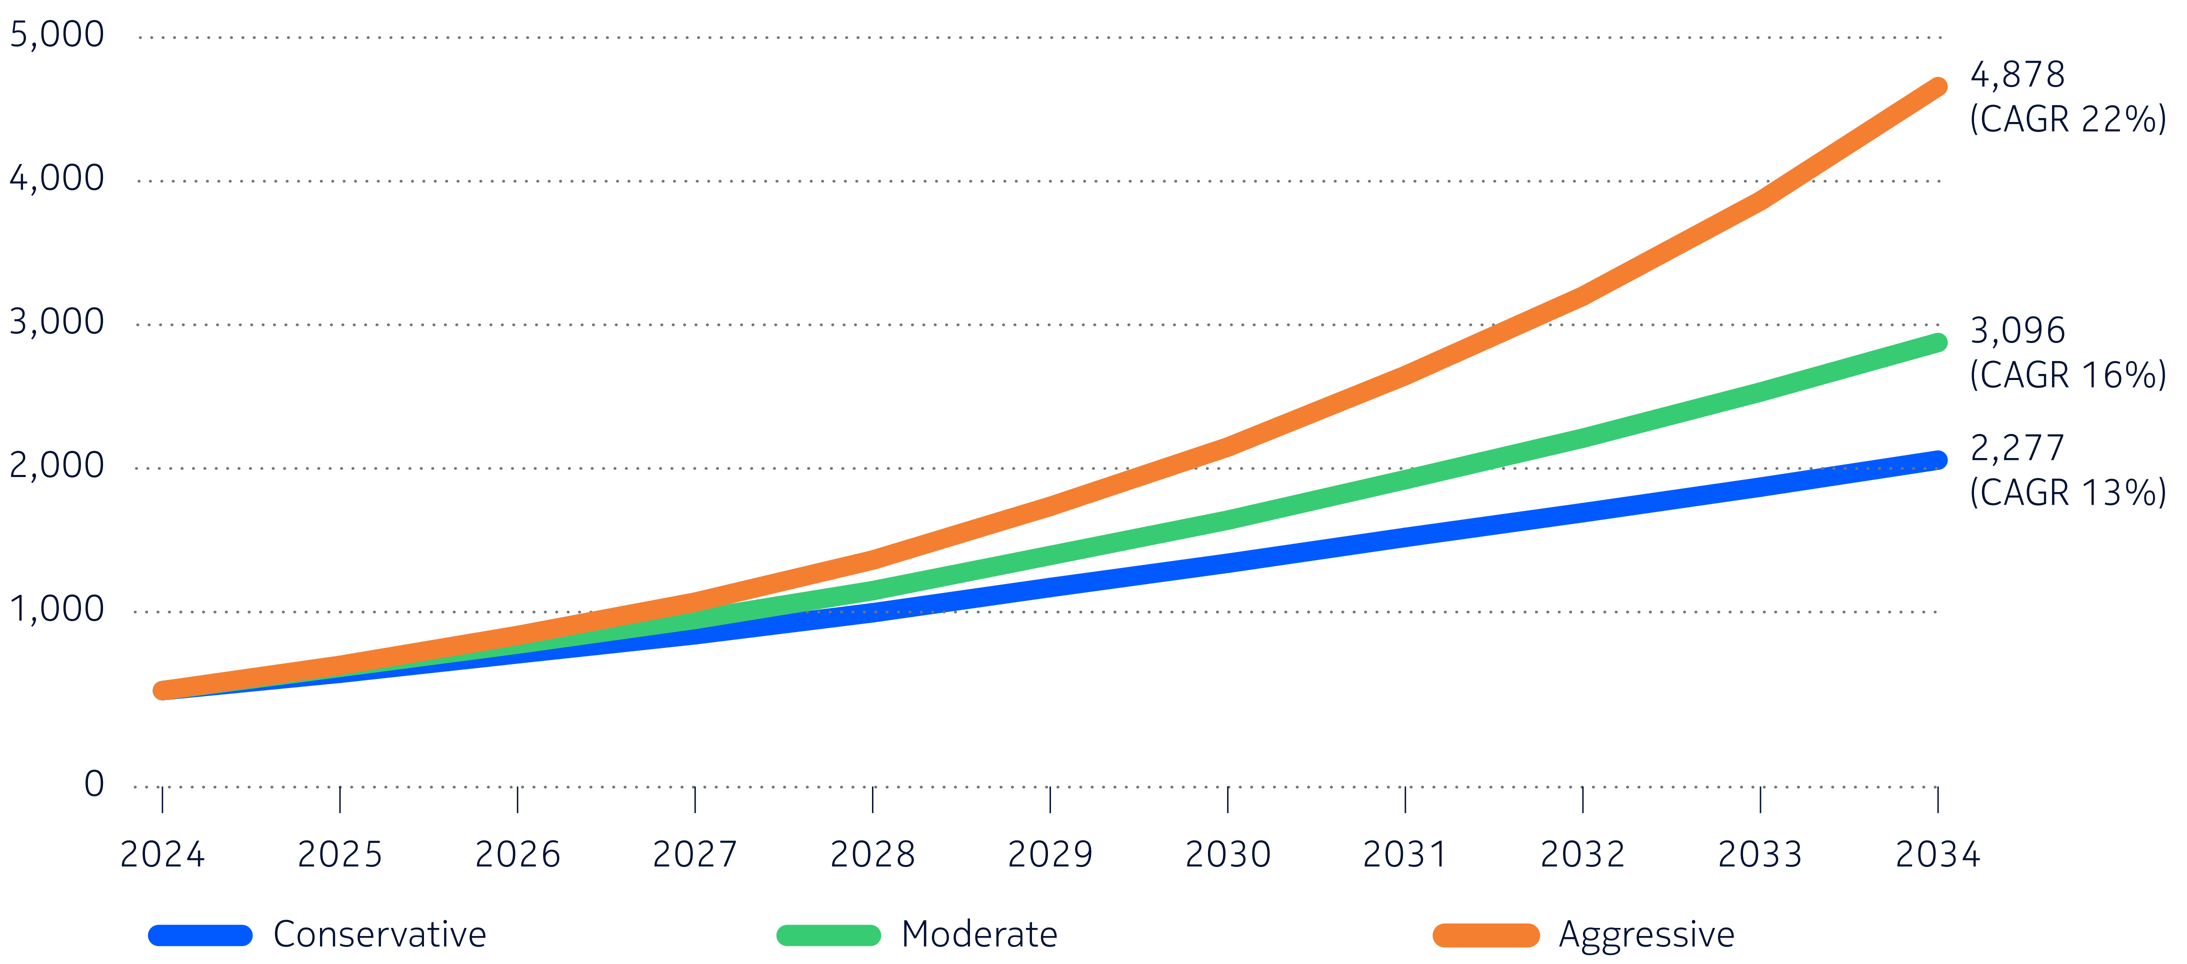

Services WAN traffic share

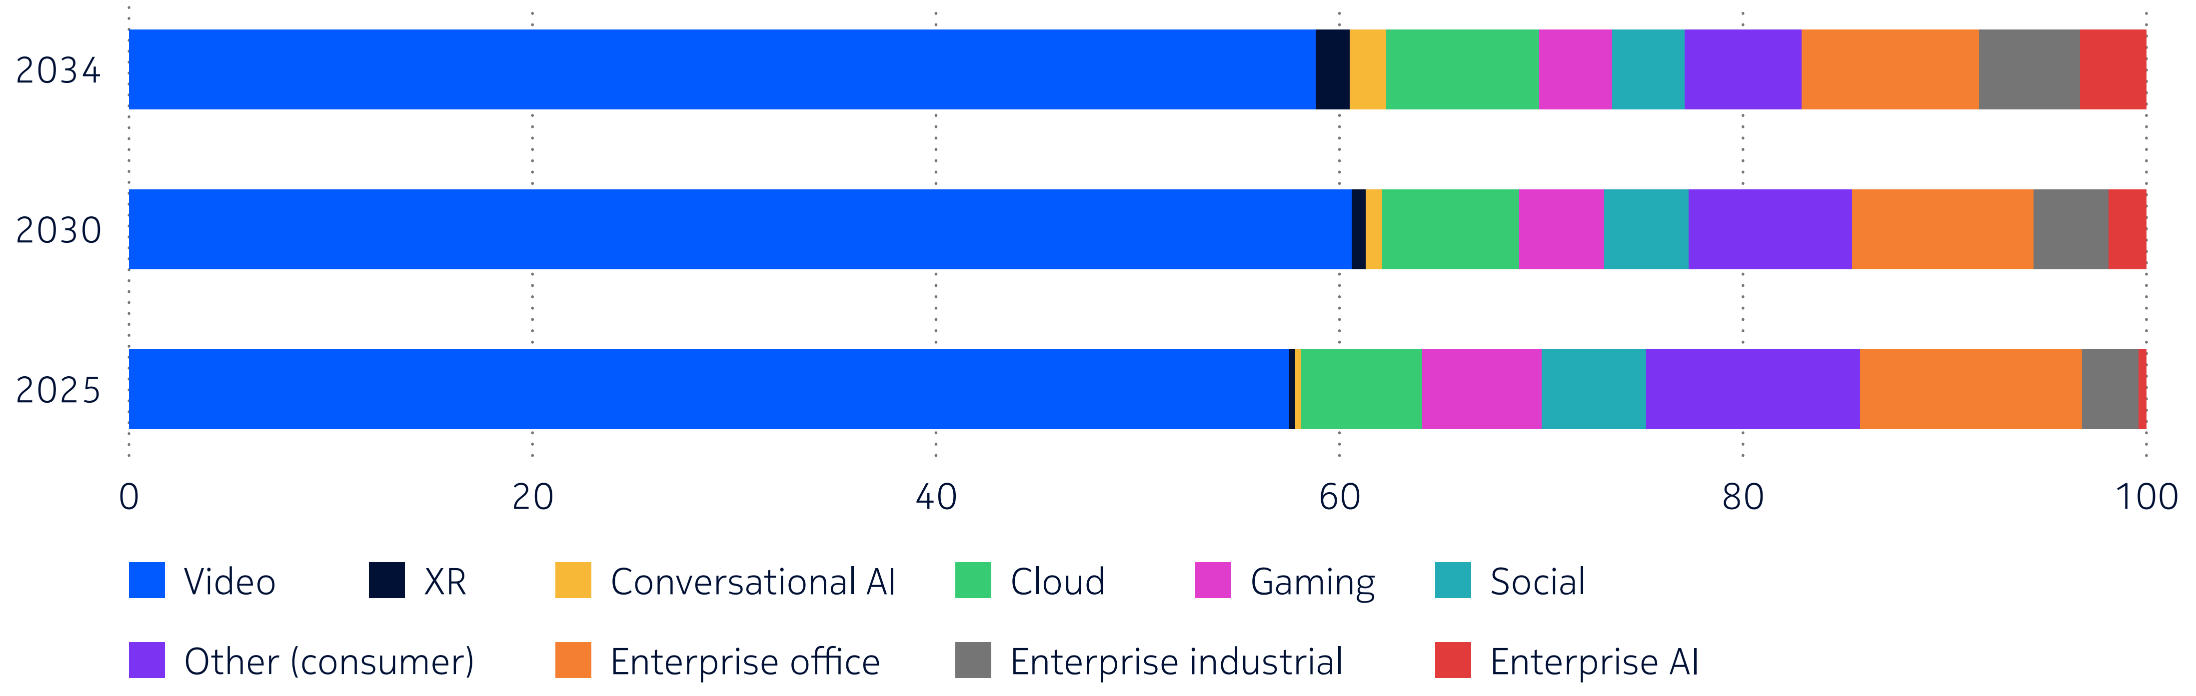

Global WAN traffic share by domain (moderate scenario)

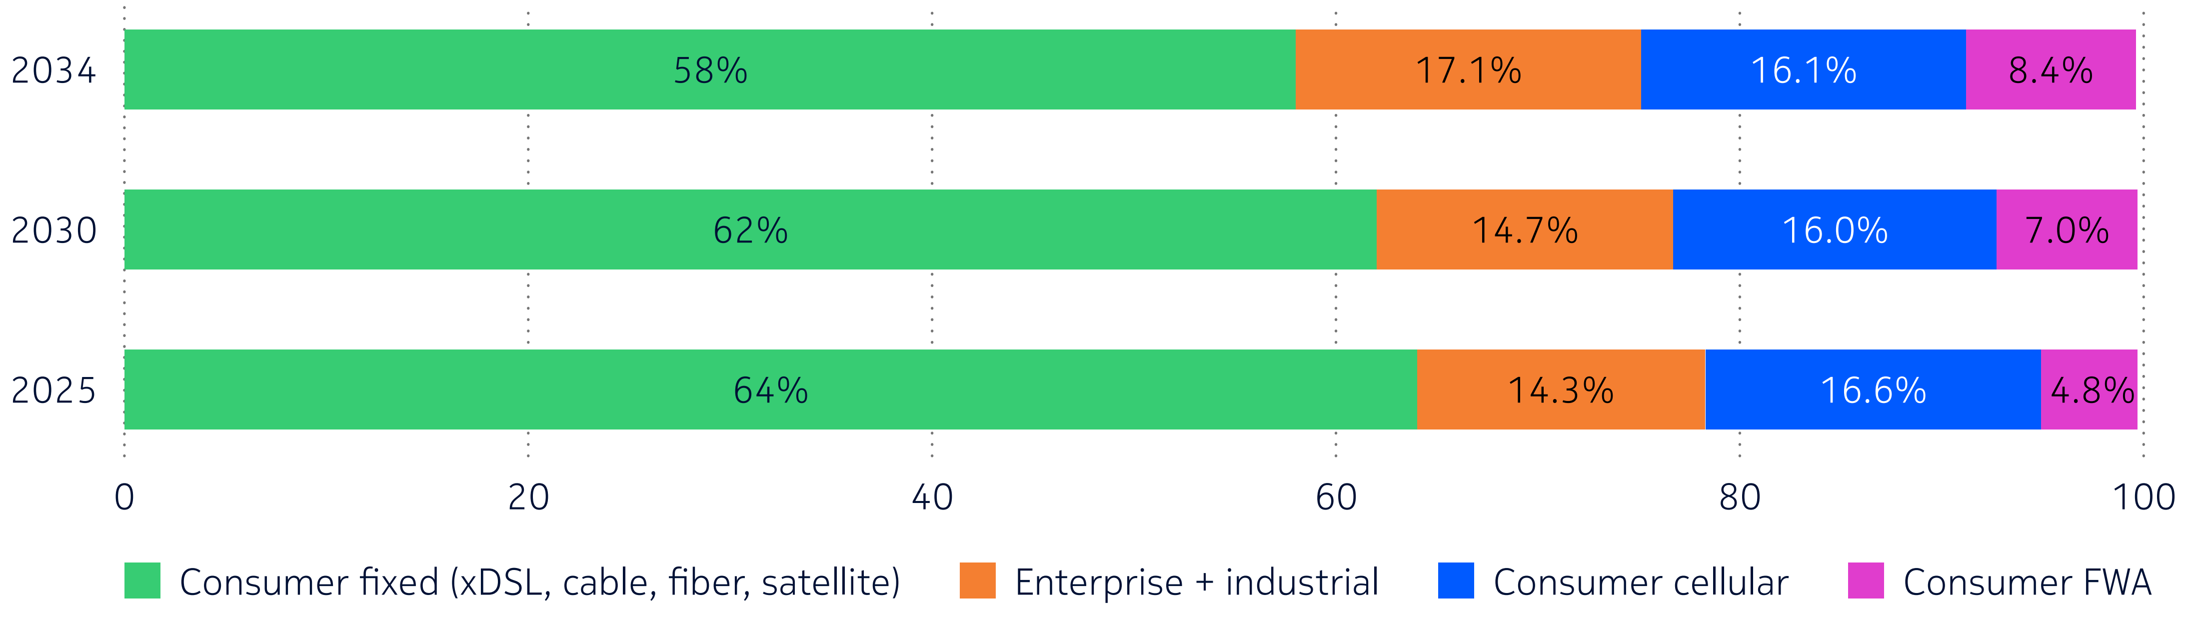

Projections for

2034

Global WAN traffic
2,277-4,878 EB/month
CAGR of 13%-22%

Video remains the anchor of consumer demand, while AI is the largest incremental driver—shifting behavior toward uplink creation and increasing movement across inter-DC corridors.

The new AI traffic paradigm

What was once dominated by activities like content streaming and data synchronization has evolved into a new era of AI-driven traffic. Today, interactions range from individuals engaging with AI assistants and copilots to autonomous AI systems communicating with each other to collect data, coordinate tasks and execute actions. This shift spans both direct and indirect domains of impact—creating entirely new traffic categories while simultaneously amplifying existing ones across consumer and enterprise networks.

Direct AI traffic

This is traffic generated during interaction of user or system with AI applications and services.

Consumer direct AI traffic includes traffic from AI chatbots, AI-assisted cloud applications like fitness or shopping, AI-powered gaming, and AI powered XR applications.

Enterprise direct AI traffic is generated by use cases that enhance workplace productivity, such as support, information retrieval, and content creation powered by generative AI. It also stems from use cases that boost operational efficiency, including predictive maintenance, autonomous operations, video and image analytics, and immersive media applications.

Indirect AI traffic

This AI-driven traffic is indirect, arising from AI algorithms that enhance user engagement and, in turn, increase traffic from non-AI applications.

Consumer indirect AI traffic includes net difference between the non-AI applications traffic with and without using AI recommending models. The examples of such applications are video streaming, social media, audio streaming, and online marketplaces.

Enterprise AI traffic has no indirect component, as it is largely determined by business needs and the efficiency of application or system design, not by recommendation-based activity.

Impact of on-device AI on WAN traffic

On-device AI has revolutionized how users interact with their devices, delivering smarter, faster and more personalized experiences. From voice assistants and AI-enhanced photo editing to real-time translation and productivity tools like email assistance, these features boost convenience and productivity while reducing reliance on external networks.

As hardware and AI models advance, more tasks powered by generative AI, such as small-scale image and video generation, voice synthesis, and offline content summarization, are expected to shift on-device. This transition will offload some content creation from the cloud, reducing cloud data traffic while improving efficiency. However, due to limits in power, battery life and storage, on-device AI will remain a complementary tool, with complex computations and heavy loads still handled in the cloud. Even so, increased on-device activity will indirectly drive higher regular traffic, particularly uplink traffic from content creation, requiring telecommunications providers to monitor trends and scale networks accordingly.

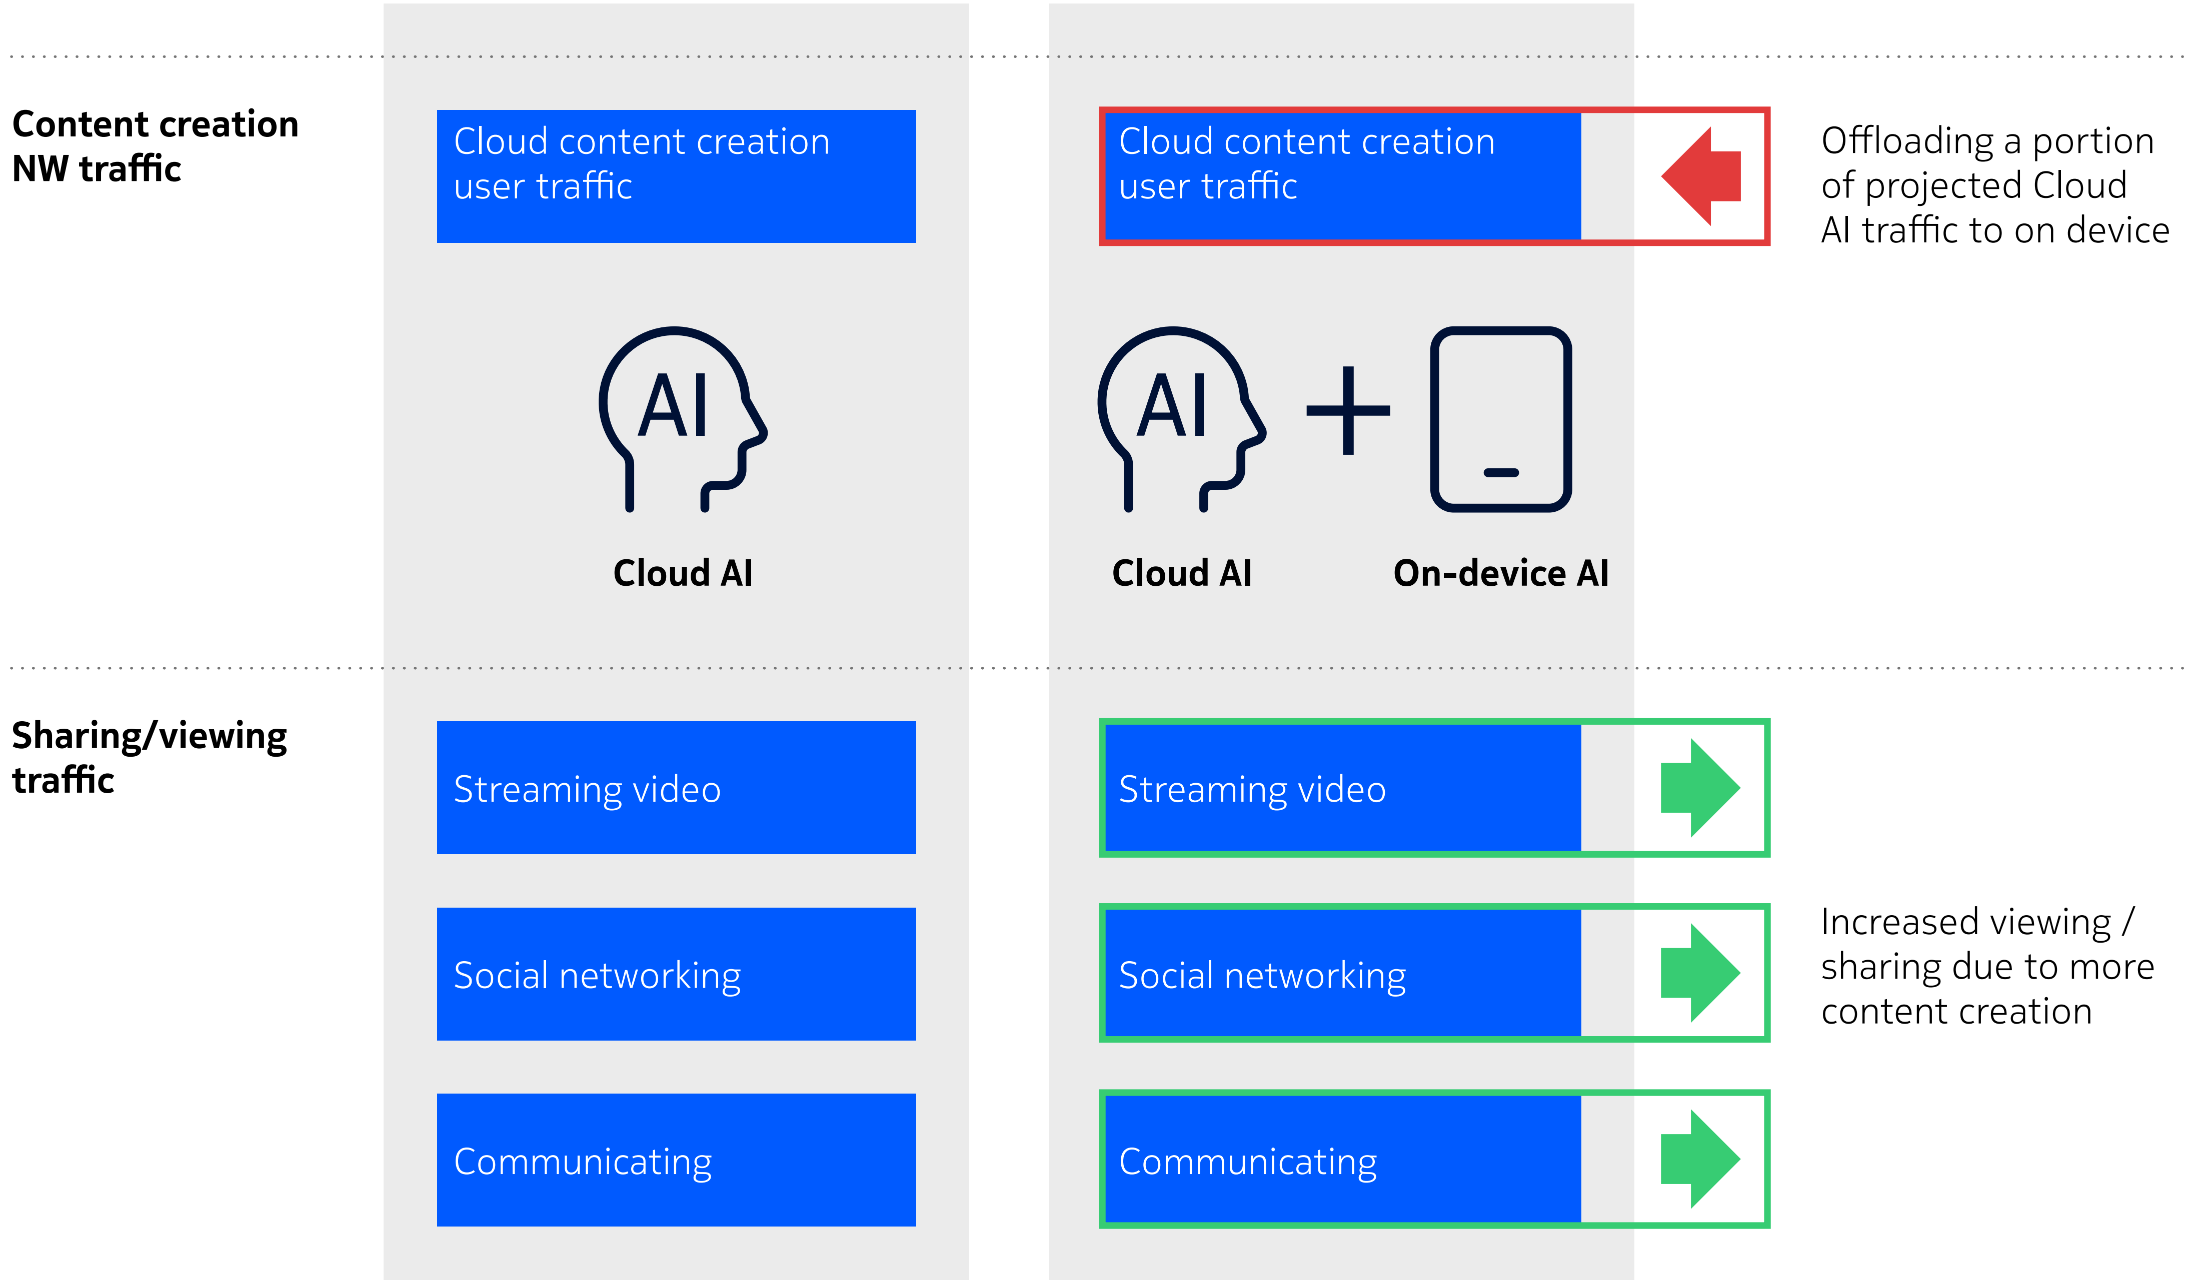

Global WAN AI traffic projections

AI traffic is set to become a major share of global WAN traffic, not just an incremental overlay. It is projected to reach 921 EB/month by 2034, growing at a CAGR of 23% and representing approximately 30% of total global WAN traffic. Figure 1 shows that traditional consumer and enterprise WAN traffic keeps rising steadily, but AI-related traffic—both consumer AI (direct and indirect) and enterprise and industrial AI— accelerates much faster and becomes the primary growth engine over the forecast period.

Consumer direct AI adds new demand on top of the non-AI applications and creates more upstream and interactive traffic.

Enterprise and industrial AI traffic is projected to grow at a CAGR of 48%, reaching 101 EB/month by 2034. Growth is driven by AR applications adoption, agentic-AI-powered industrial automation, agentic AI assistants, and increasing use of sensor data in AI applications. While traditionally uplink-heavy, enterprise AI traffic is expected to shift toward downlink dominance in the next decade, led by immersive AI applications.

By the end of the period, AI traffic is no longer a niche category sitting at the edges of the network. It is a core driver of WAN scaling, capacity planning, and revenue opportunity for telecommunications providers.

Figure 1: WAN AI traffic, global, moderate view, EB/month

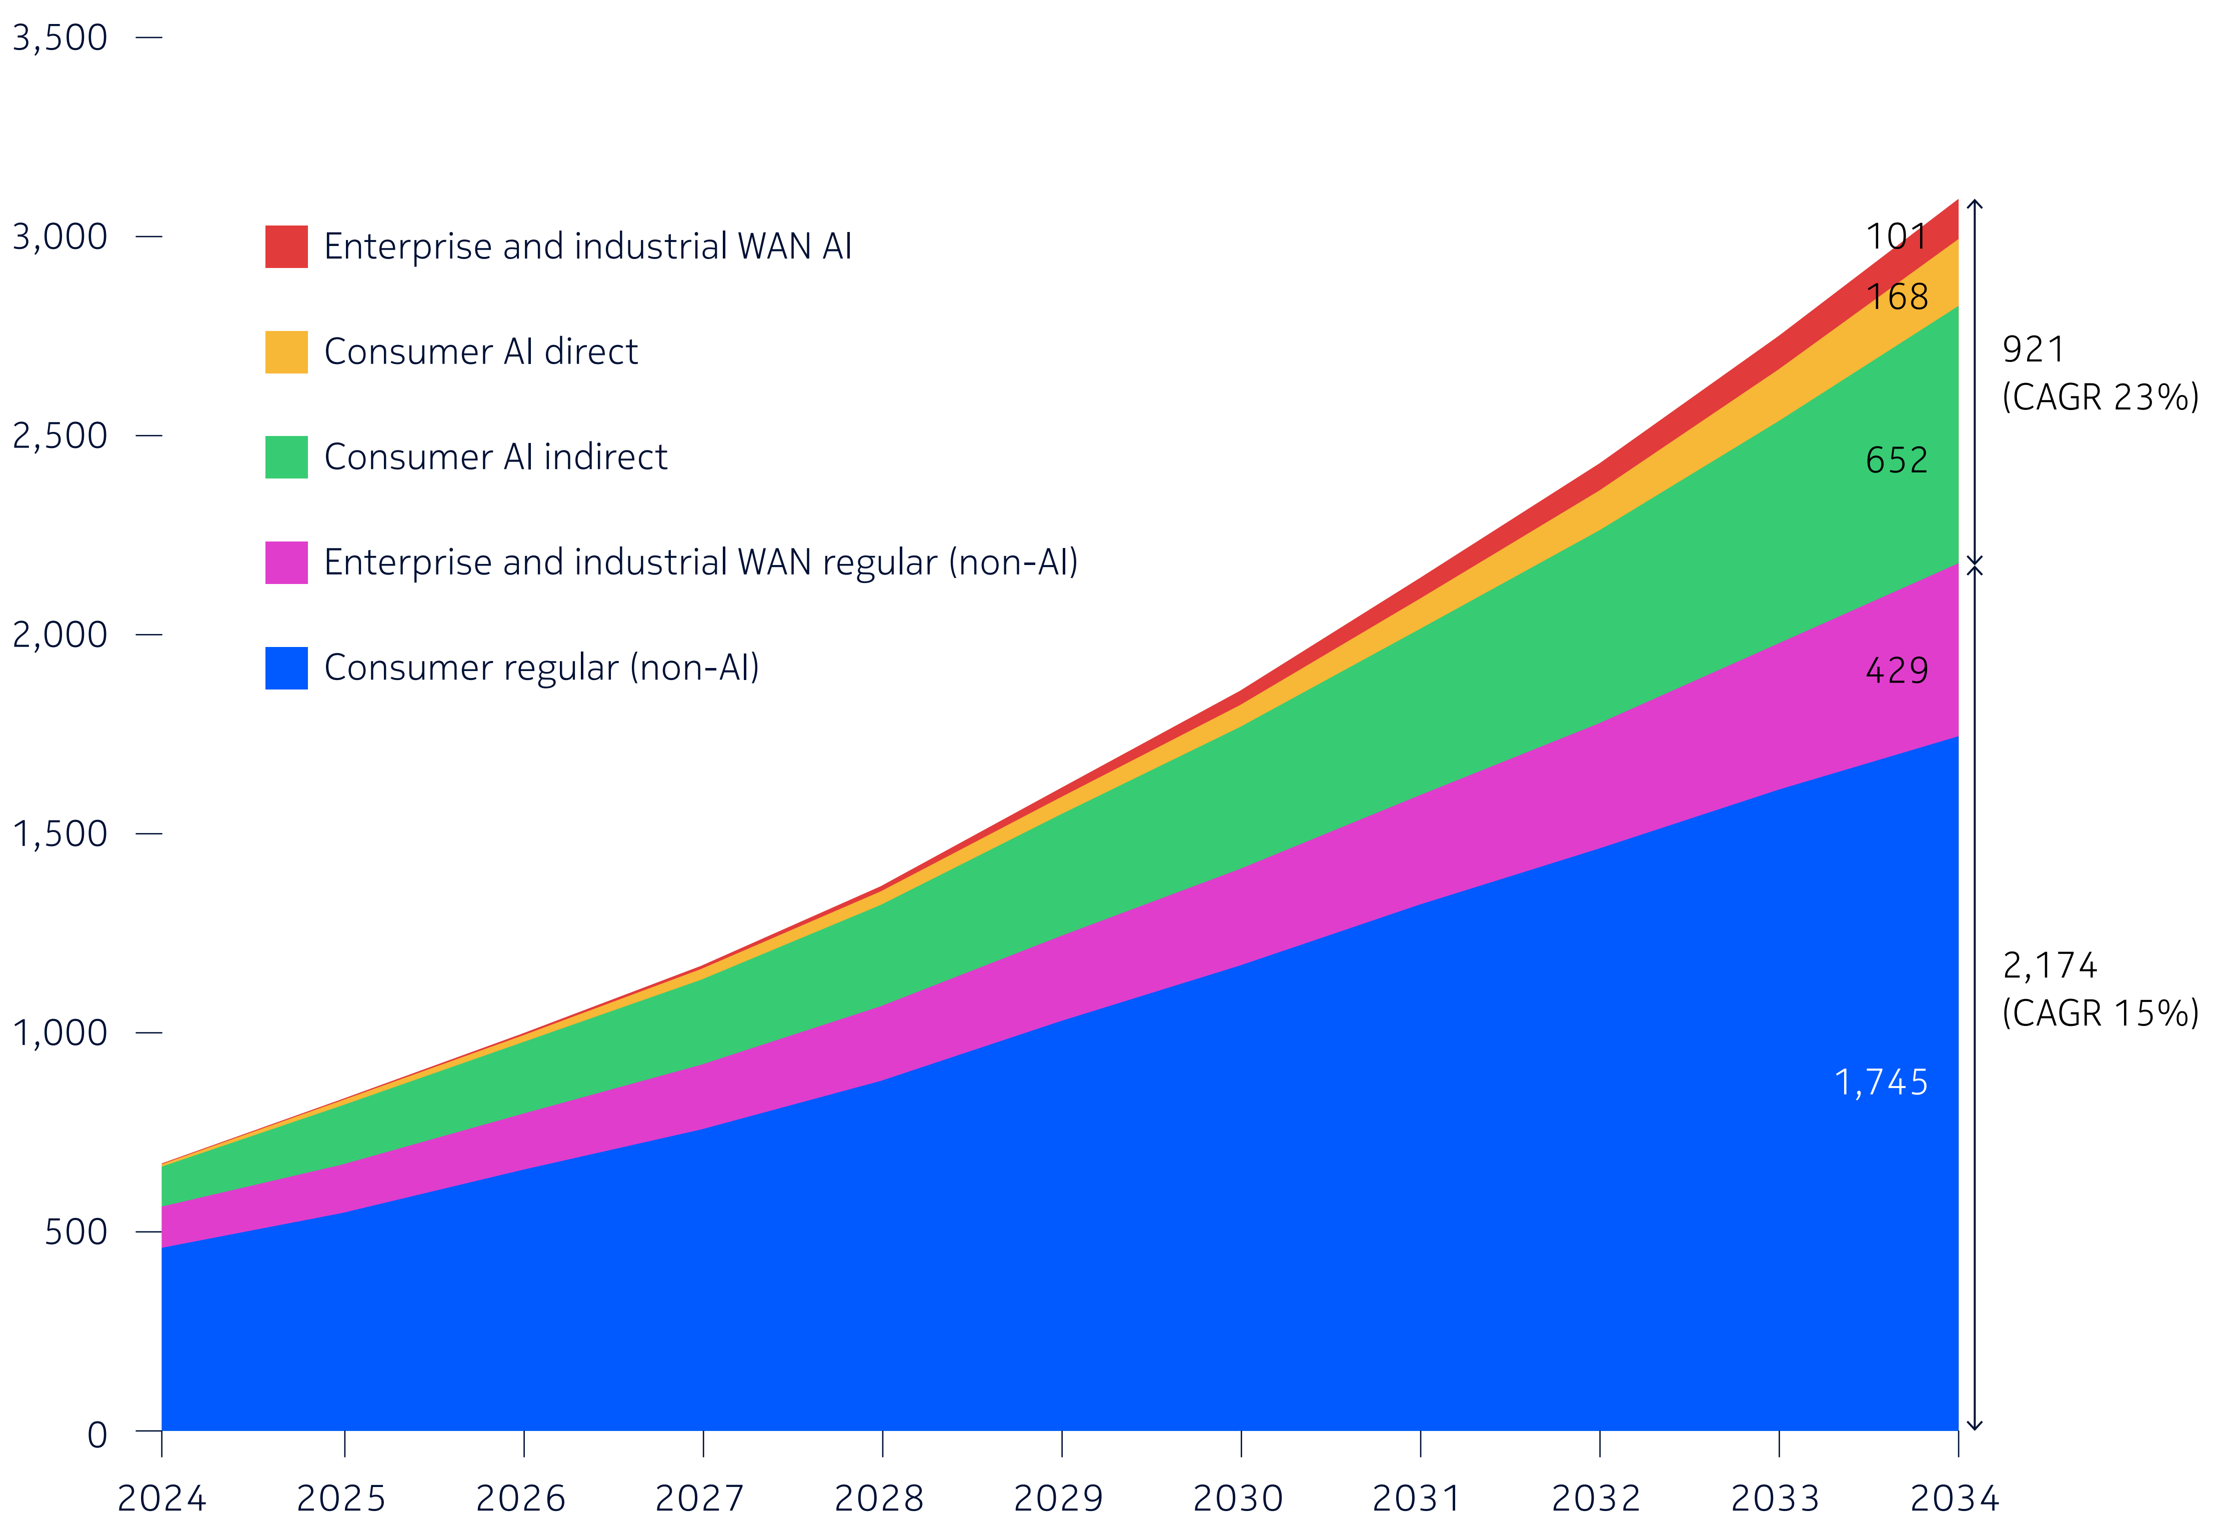

Consumer

Consumer demand remains the largest share of global WAN traffic through the forecast period, but what people are doing on the network is changing. AI-assisted creation and agentic assistants are shifting behavior from passive viewing to more active, uplink-heavy usage like sharing, collaborative editing and live interaction. Video remains the biggest contributor to overall volume, and ongoing codec improvements help keep individual streams manageable even as devices continue to push toward higher resolutions and frame rates. At the same time, multi-modal experiences—blending video, voice, gaming and real-time interaction—are driving tighter expectations for latency and jitter. Networks are being pushed not just to deliver more data, but to deliver it predictably. This continues the pattern we described last year: traditional drivers like video and gaming are now tightly coupled with emerging AI-driven use.

Consumer mobile traffic growth

Total consumer mobile traffic is projected to reach 364–633 EB/month by 2034, growing at a 12–18% CAGR, driven largely by video and social media usage. Video traffic will remain the core growth engine, contributing nearly three-fourths of total mobile data, supported by AI-powered recommendations, GenAI-led content creation, and adoption of high-resolution, large-screen and foldable devices. However, advancements in video codecs will help moderate traffic growth by improving compression efficiency even as video quality and viewing time rise. Social networking will maintain a steady share of around 9% through the period.

In the later years, AI-driven content creation, cloud-based personal assistants, and AR/VR will emerge as new growth contributors, collectively adding around 8% of total traffic by 2034.

Meanwhile, 5G Advanced and early 6G deployments will ease bandwidth constraints and enable richer, high-quality streaming and immersive user experiences.

Figure 2: Consumer mobile traffic, EB/month

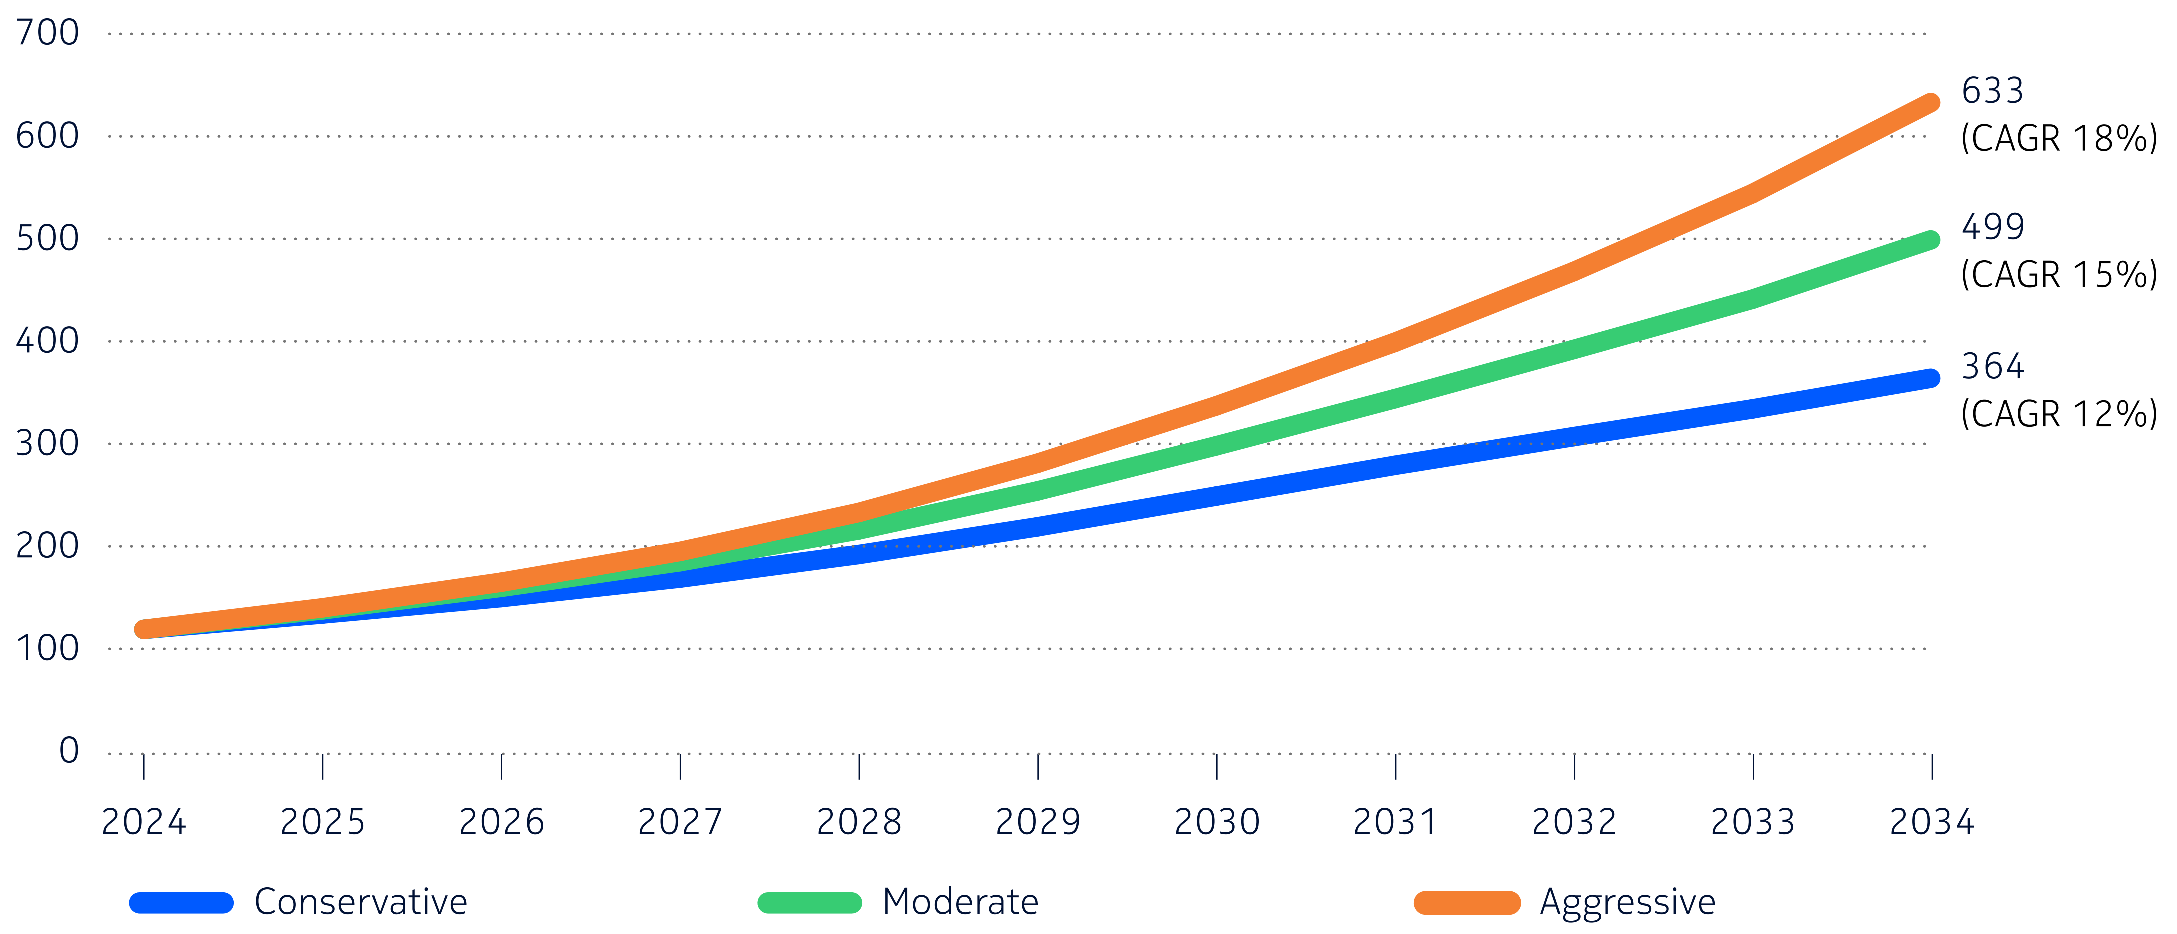

Figure 3: Application share of consumer mobile traffic, moderate

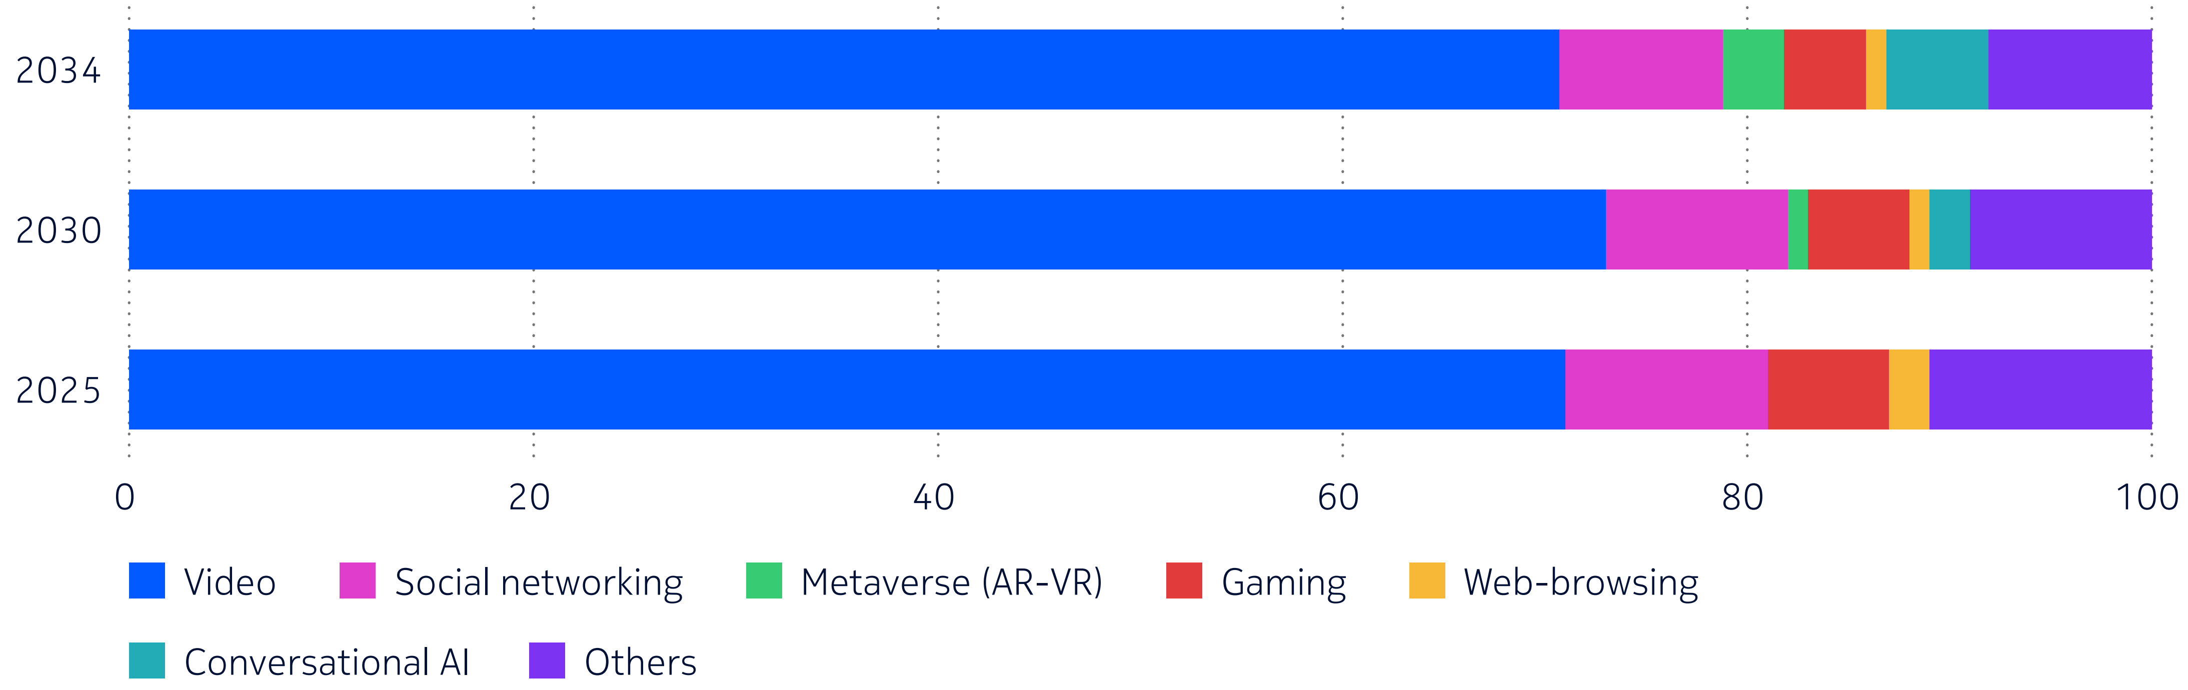

Consumer mobile network evolution: 4G to 5G to 6G

The shape of mobile traffic over the next decade is determined less by how many people are connected and more by what networks they are connected to. 5G becomes the dominant access layer before the end of this decade, and 6G begins to appear in the early 2030s. As that transition plays out, more of the world's mobile traffic rides on higher-throughput, lower-latency, more deterministic radio and transport infrastructure.

That shift matters because the jump from 4G to 5G, and then from 5G toward 6G, does two things at once. First, it unlocks uses that were previously impractical on mobile, such as high-resolution live video, cloud gaming with real responsiveness, spatial or AR overlays in real time, and AI assistance that feels immediate rather than buffered. Second, it normalizes that behavior. What begins as a premium experience for a subset of users becomes table stakes for everyone on upgraded radio layers. In other words, traffic density follows the network upgrade. Subscriber growth alone does not explain the curve; capability does.

Figure 4: Mobile traffic, regional, moderate scenario, EB/month

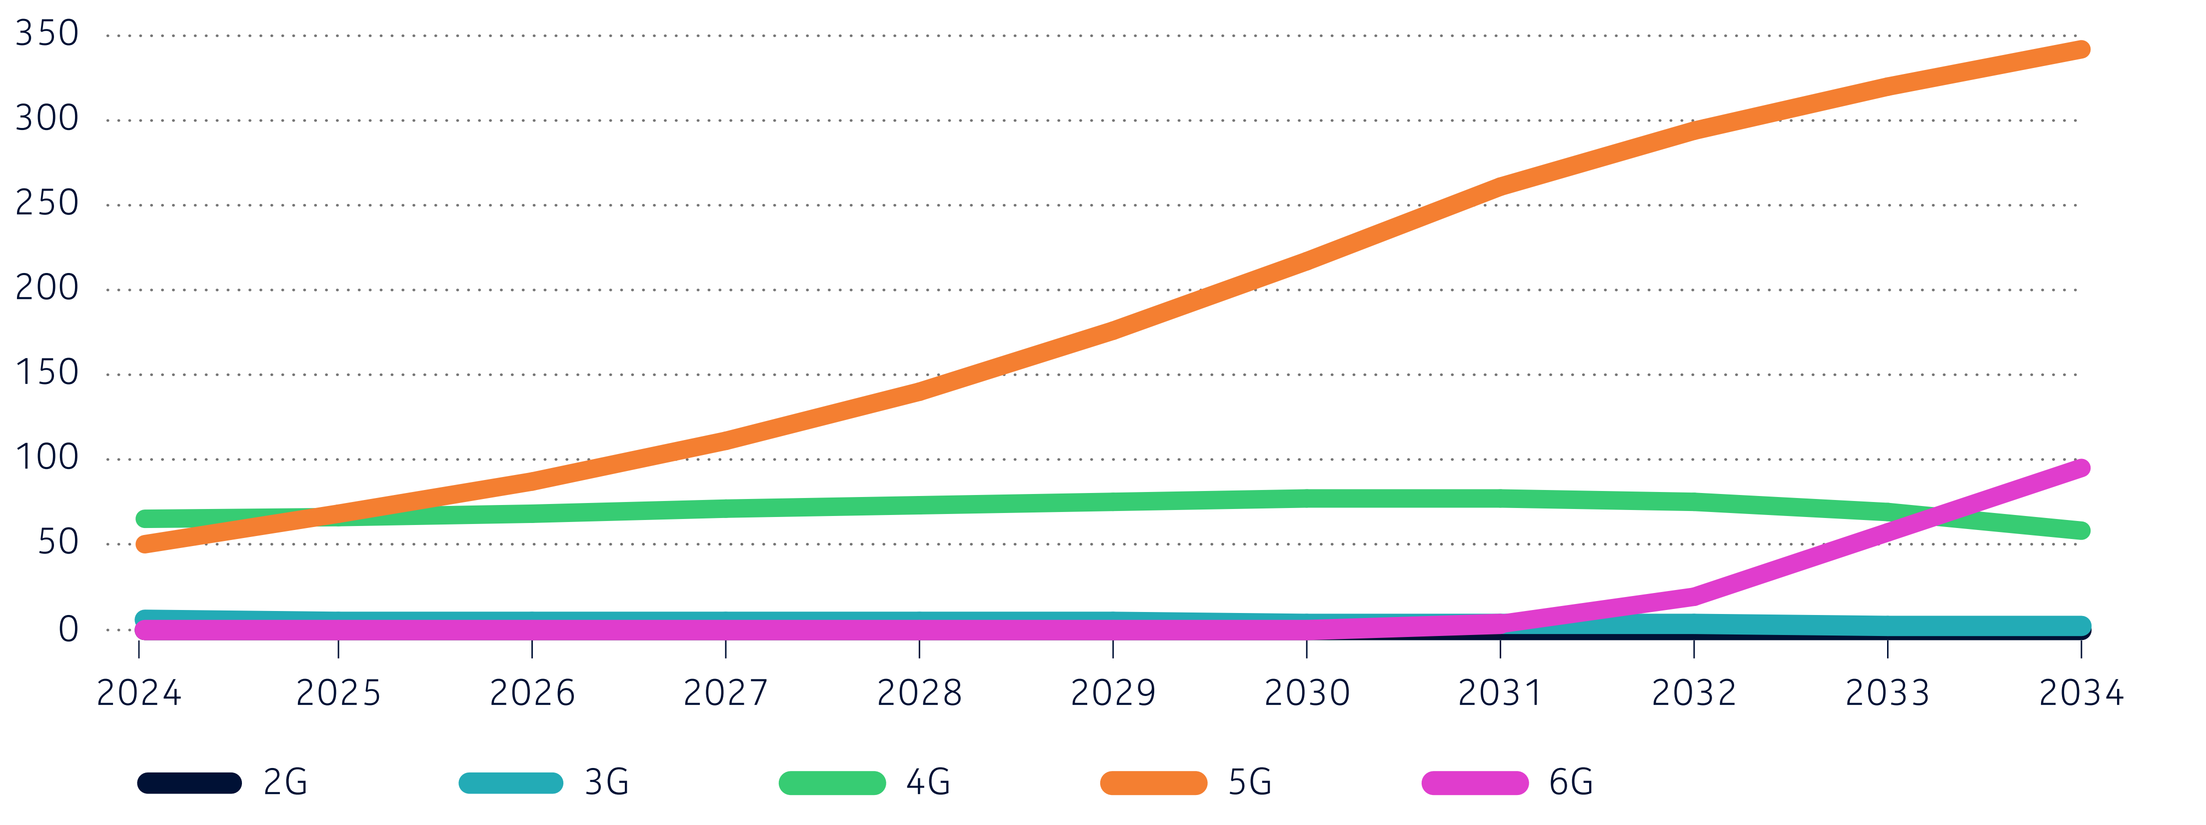

Figure 5: Normalized 5G+6G subscription increase (2024–2034)

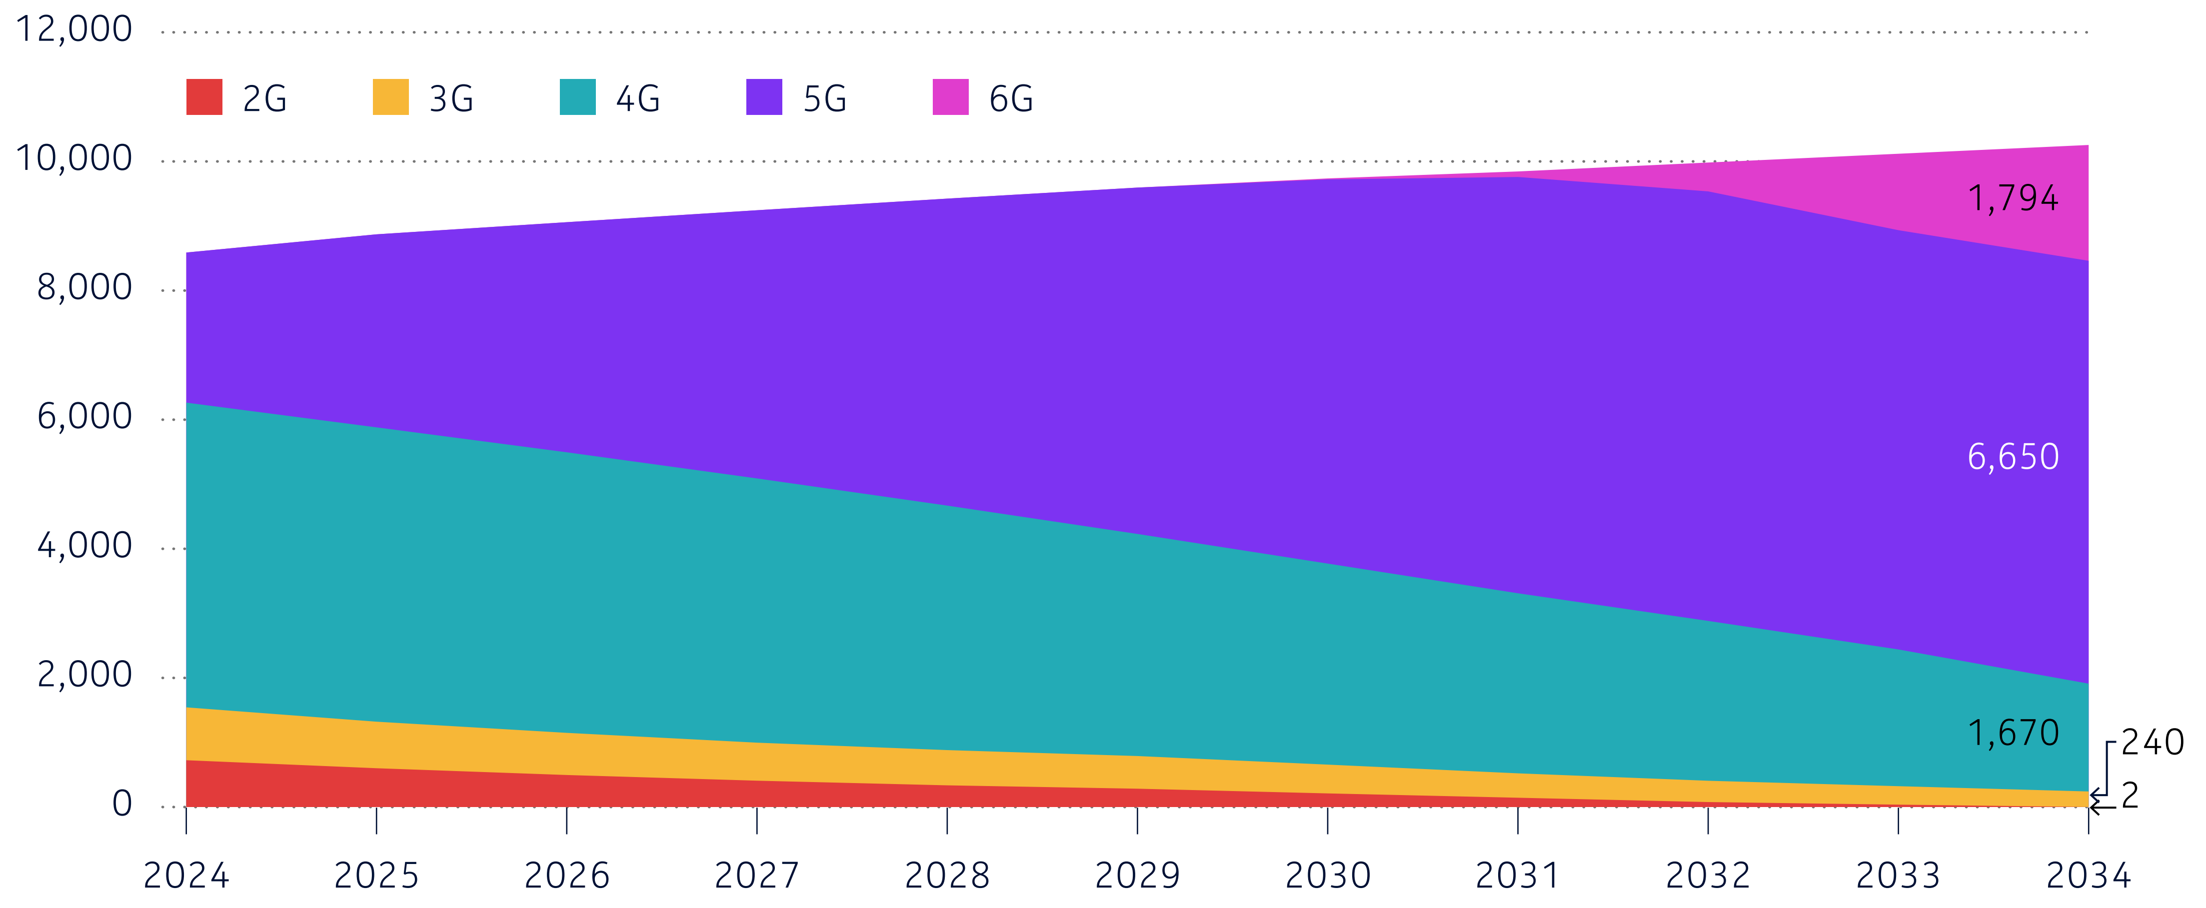

Consumer mobile traffic by region

In the near term, mature markets (NA, China, WE) are projected to experience slower traffic growth due to the saturation of mainstream applications and early-stage adoption of generative AI and extended reality (XR). However, in the mid-to-long term, these regions are expected to see sustained growth driven by the mass adoption of generative AI and accelerated XR adoption, alongside a shift to higher video resolutions.

Conversely, emerging markets will demonstrate steady growth. In the near term, this growth is fueled by digital inclusion, 5G rollout, and increased service adoption, with a gradual shift from SD to HD video for premium users. The mid-to-long term will see continued expansion in traditional services, generative AI gaining mass traction, and HD becoming the norm, although XR traffic is expected to have limited impact.

Through 2034, APAC, India, and MEA are expected to record strong compound annual growth rates (CAGRs) and drive global traffic growth. China will continue to contribute significantly in volume, despite a lower CAGR resulting from market maturity and high existing penetration.

Figure 6: Mobile traffic, regional, moderate, EB/month

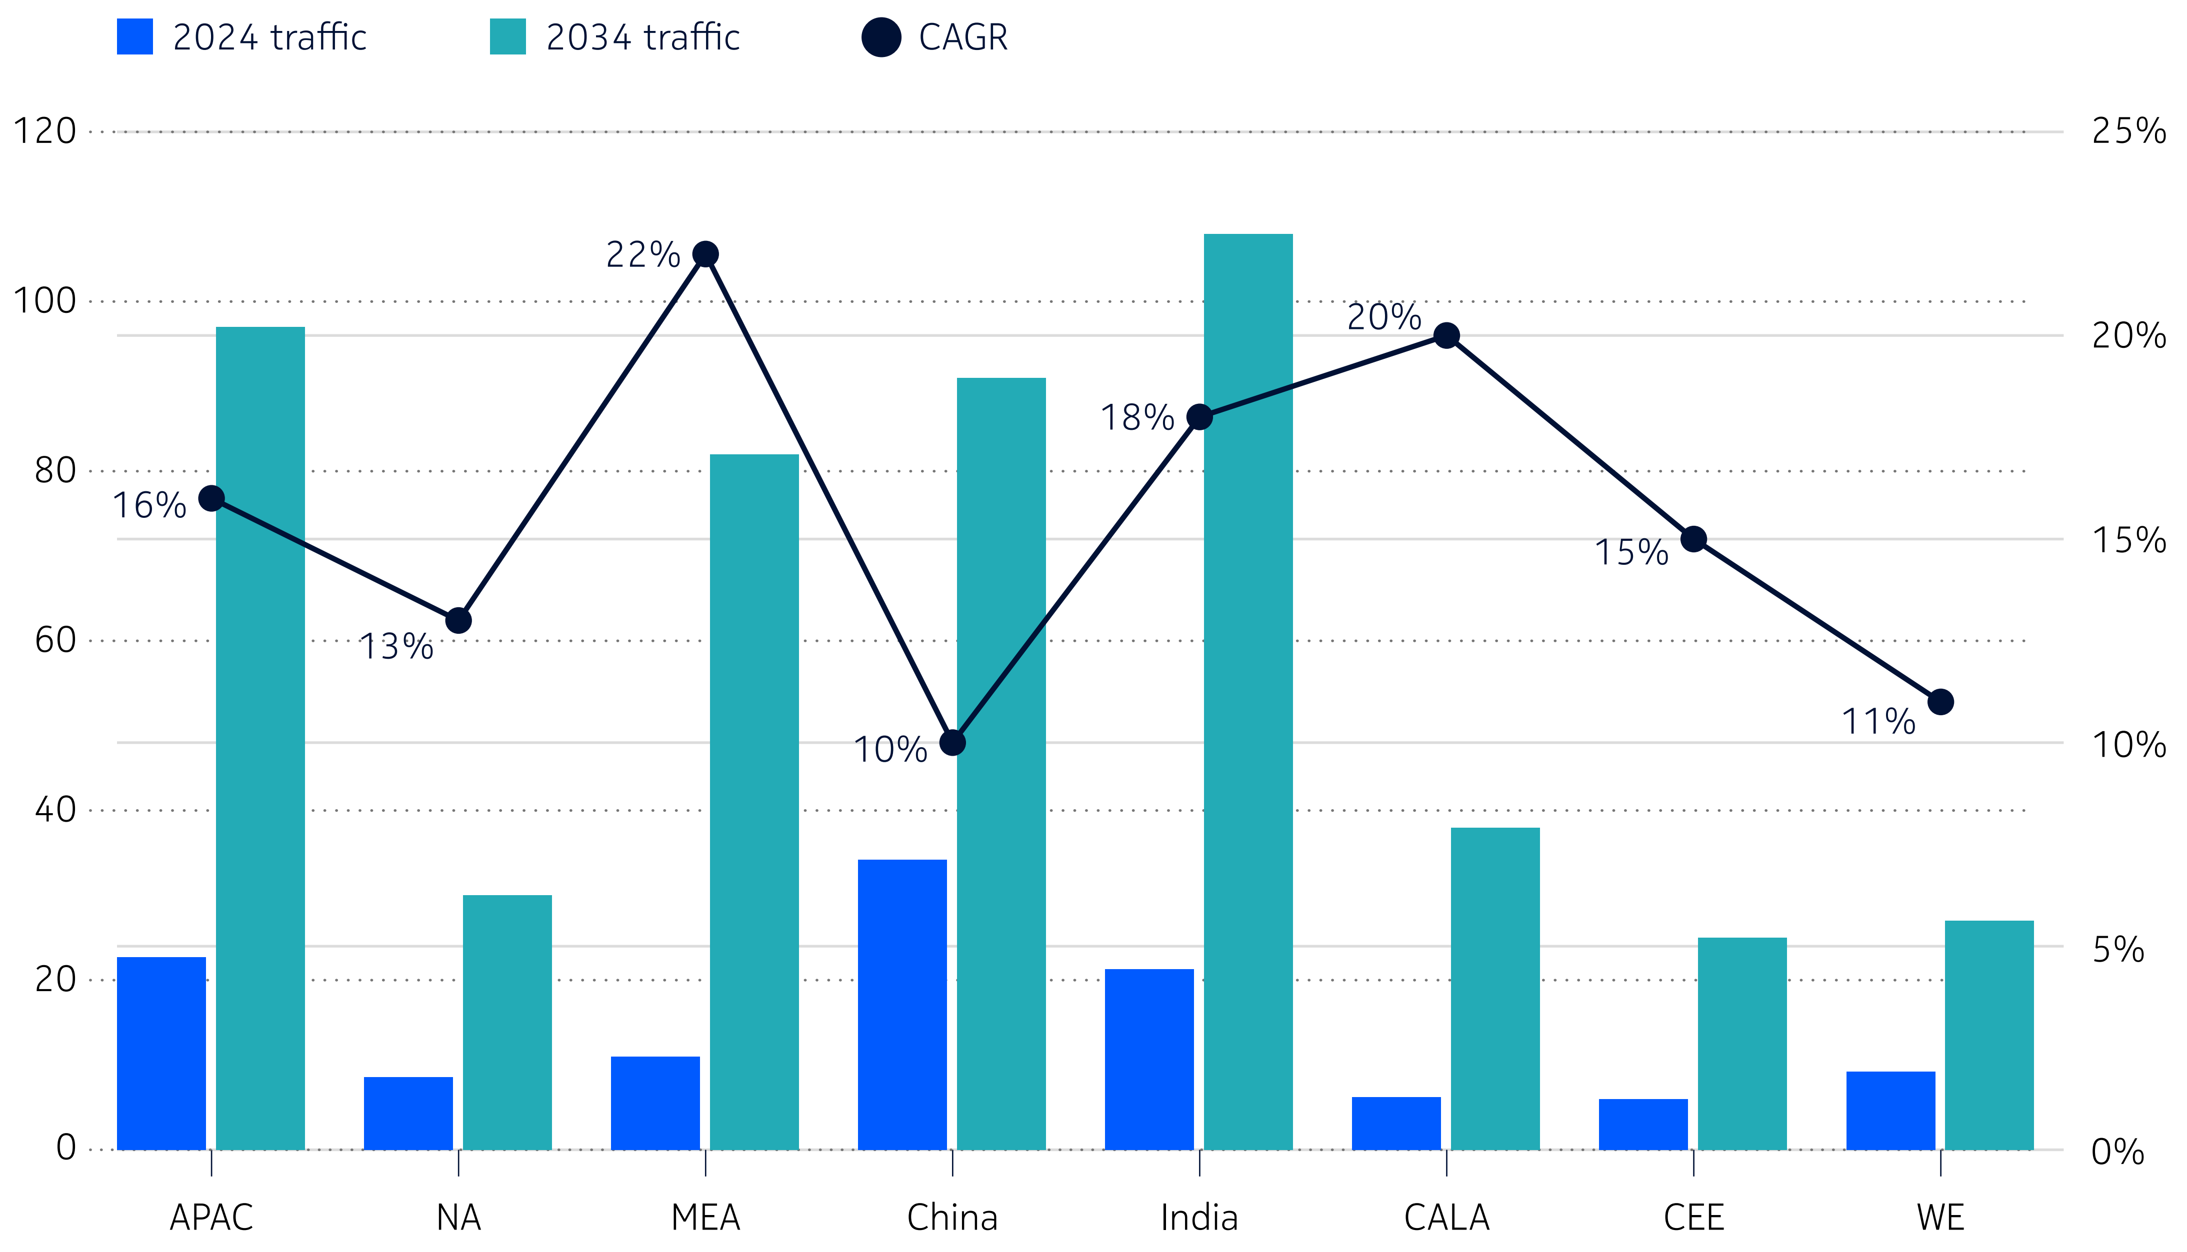

Consumer cellular including FWA

Consumer cellular traffic, encompassing both mobile and fixed wireless access (FWA), is forecast to grow at a CAGR of 14% to 23% between 2024 and 2034, depending on the scenario (conservative, moderate or aggressive).

Under a moderate growth scenario, total consumer cellular traffic is expected to reach 759 EB/month by 2034, representing an 18% CAGR. FWA is a key driver of this expansion. By 2034, FWA is projected to contribute 261 EB/month to the total cellular traffic under the moderate scenario, demonstrating its rapid growth and expanding share relative to traditional mobile traffic (499 EB/month). The steady expansion of FWA is expected to play a major role in driving the overall rise in consumer cellular data consumption.

Figure 7: Consumer cellular traffic, including FWA, EB/month

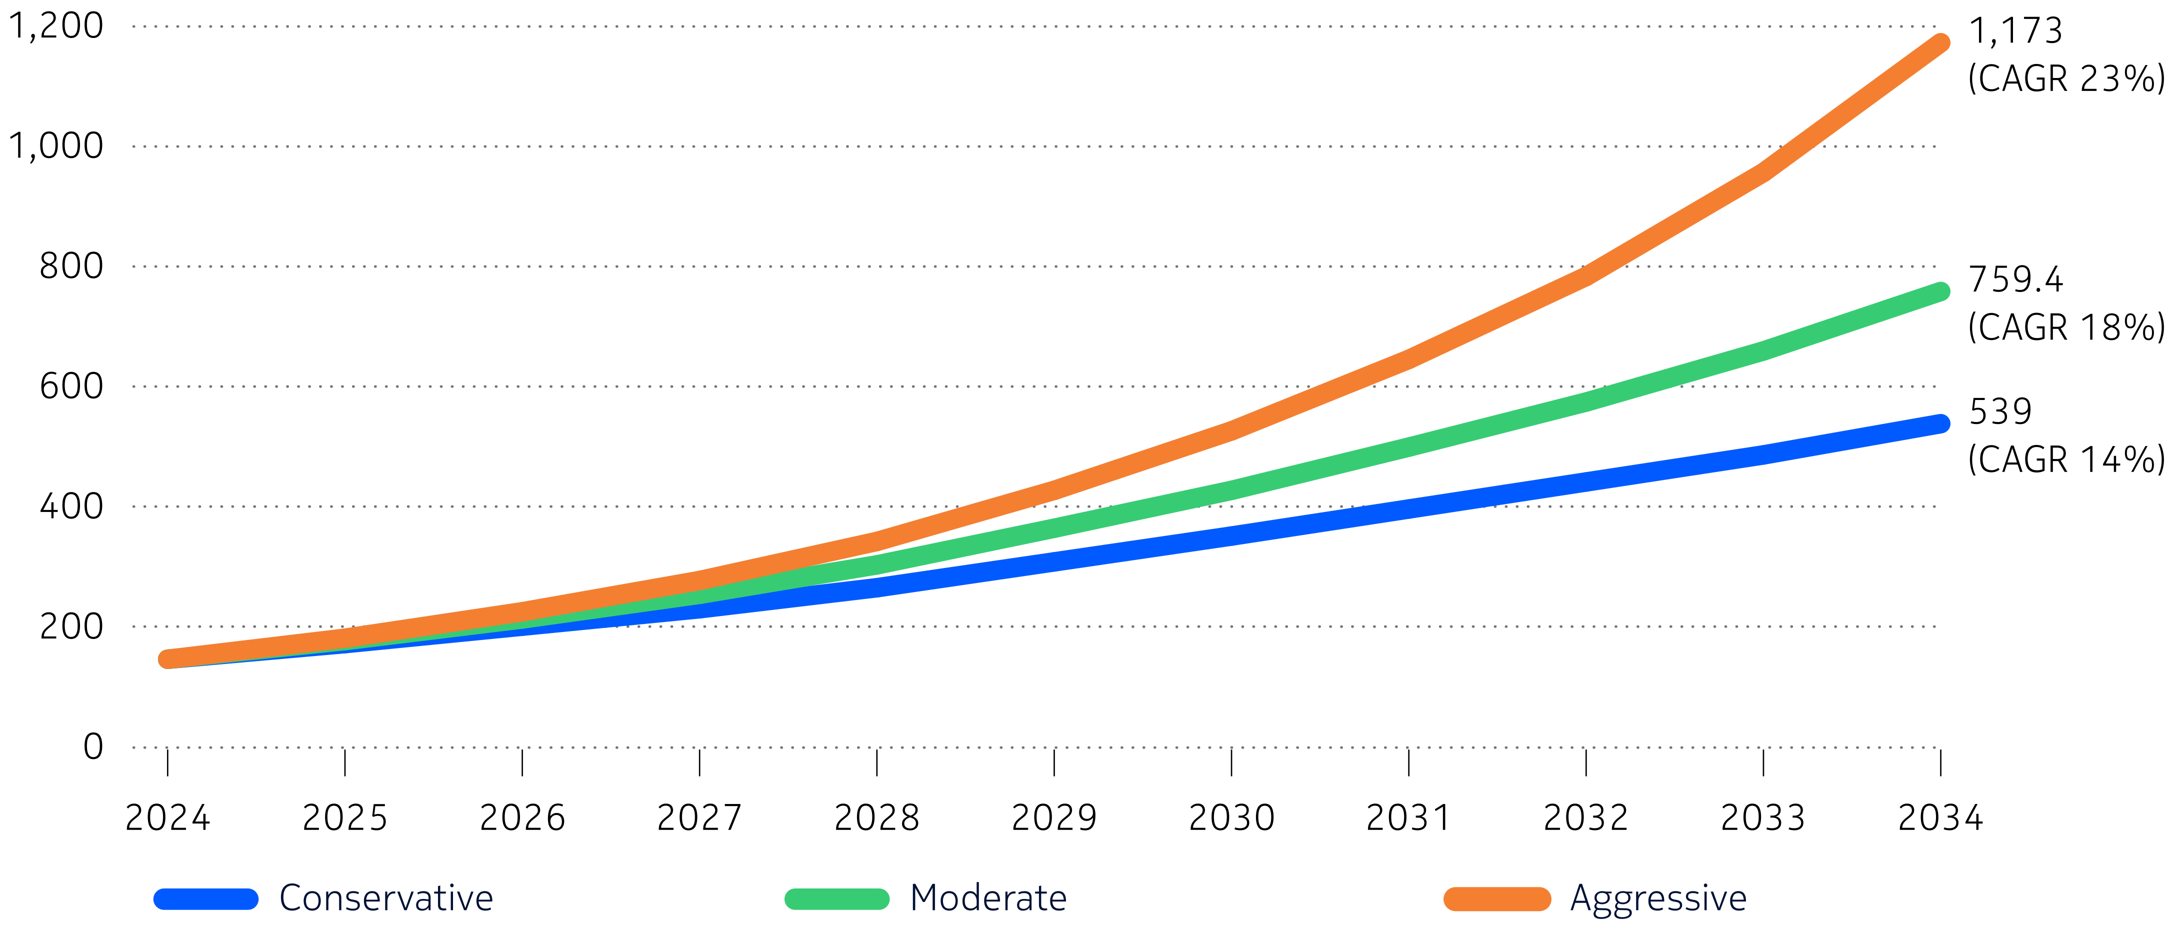

Figure 8: Consumer cellular traffic, moderate, including FWA, EB/month

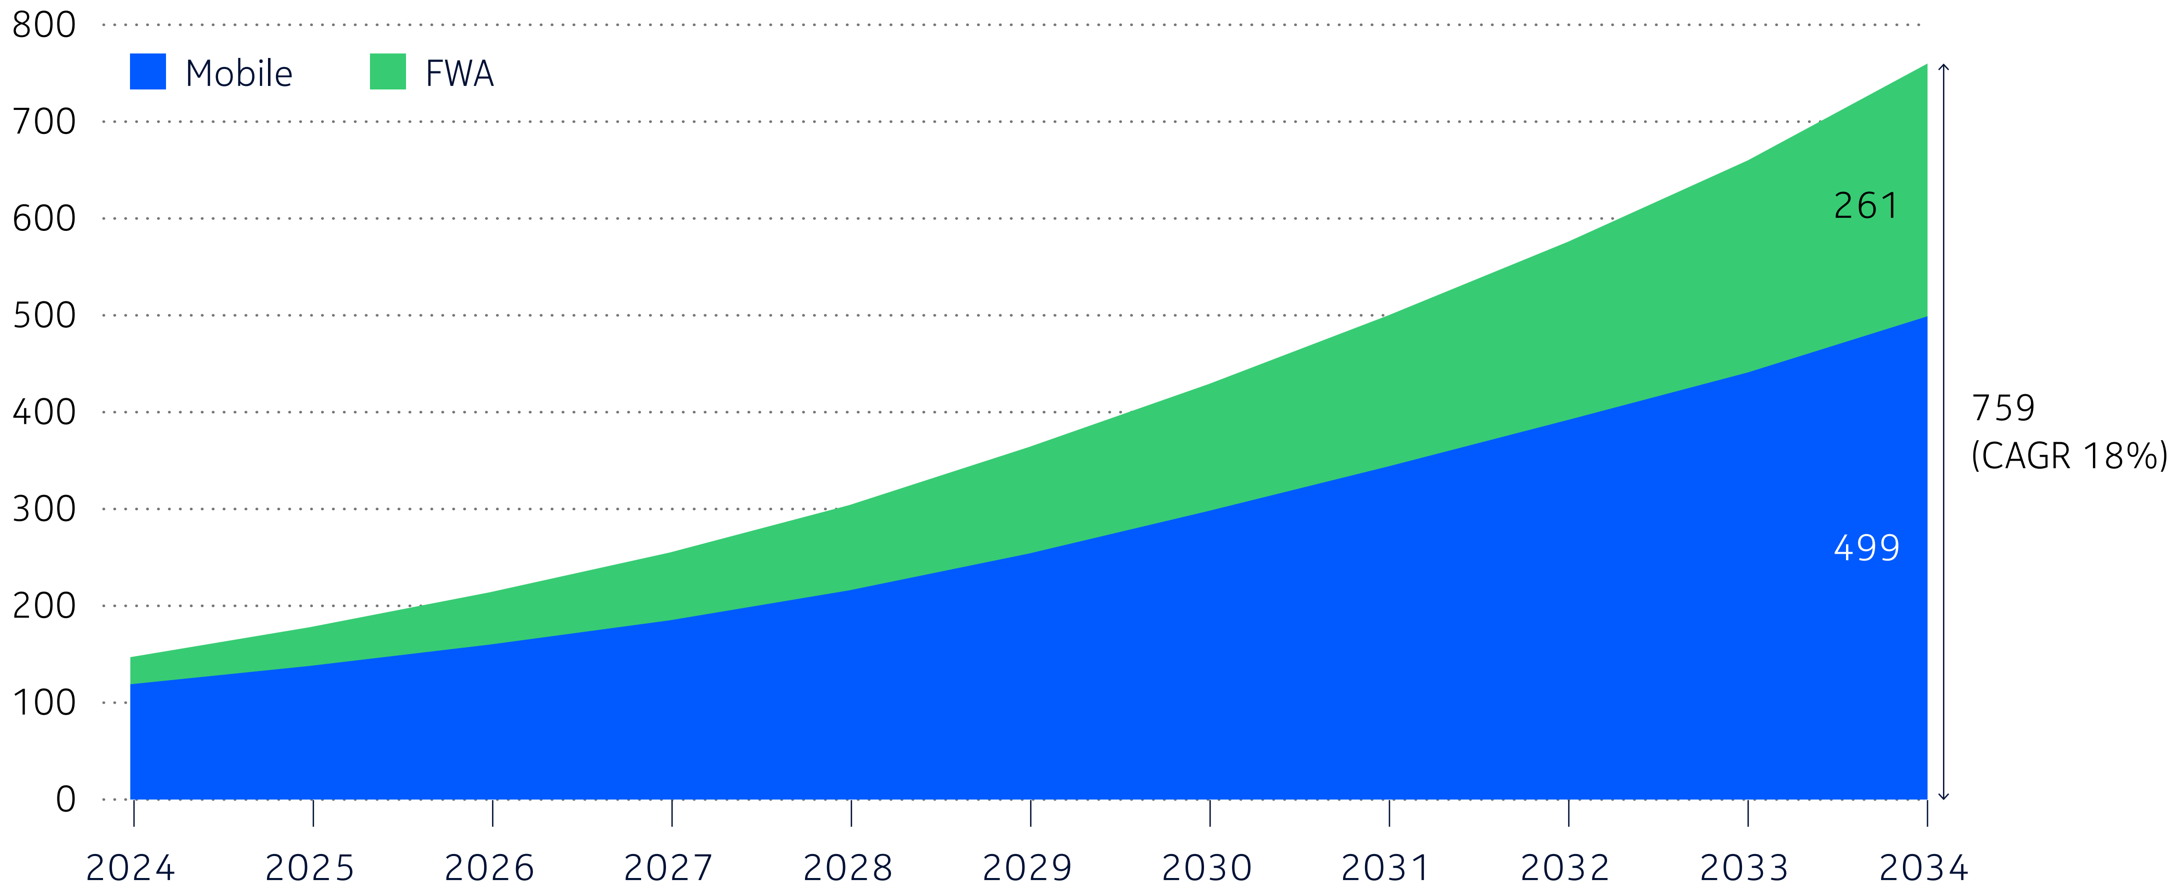

Consumer fixed traffic

Global consumer fixed traffic (figure 9), excluding fixed wireless access, keeps rising through 2034 to roughly 1,405 to 2,791 EB/month, which corresponds to a 13–22% CAGR. However, the increase in traffic is gradual rather than explosive, as advancements in video compression technologies continue to optimize per-stream bitrates, while the demand for higher-resolution formats like 8K remains relatively constrained.

While traditional video streaming will continue to account for a substantial portion of traffic, emerging applications such as cloud services, VR gaming, and AI-enhanced experiences are poised to claim a growing share. These applications differ from traditional watch-only content by introducing interactive elements that demand more uplink capacity, shorter round-trip times to the cloud, and continuous data exchange between homes and cloud infrastructure. This shift will gradually make household traffic more symmetric, as uplink usage becomes increasingly important.

Absolute volume of VR traffic stays low due to high costs and limited comfort of VR headsets as well as a focus on AI over full-immersion VR. Meanwhile, AI and AR applications are expected to grow across areas like work, fitness and home management. Despite higher uplink traffic share (up to 40%), their overall network impact will stay modest due to low per-user bitrates.

Figure 9: Consumer fixed traffic, global, excluding FWA, EB/month

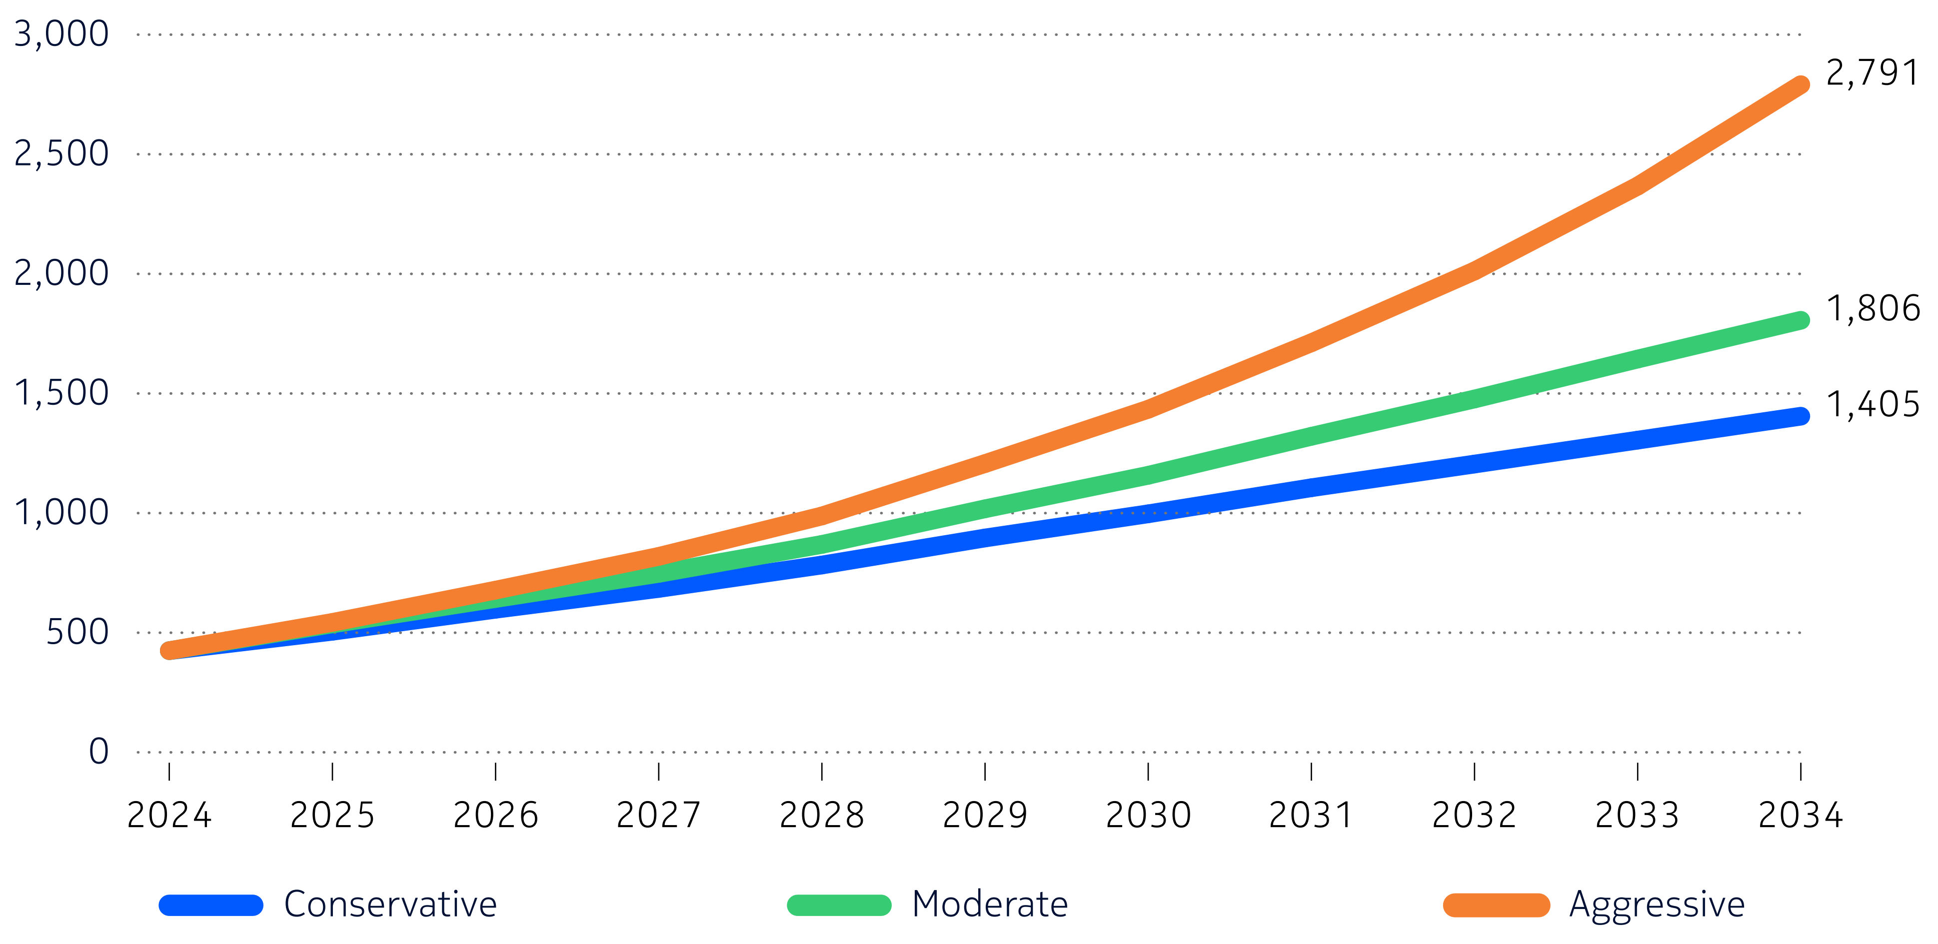

Figure 10: Share of services, moderate

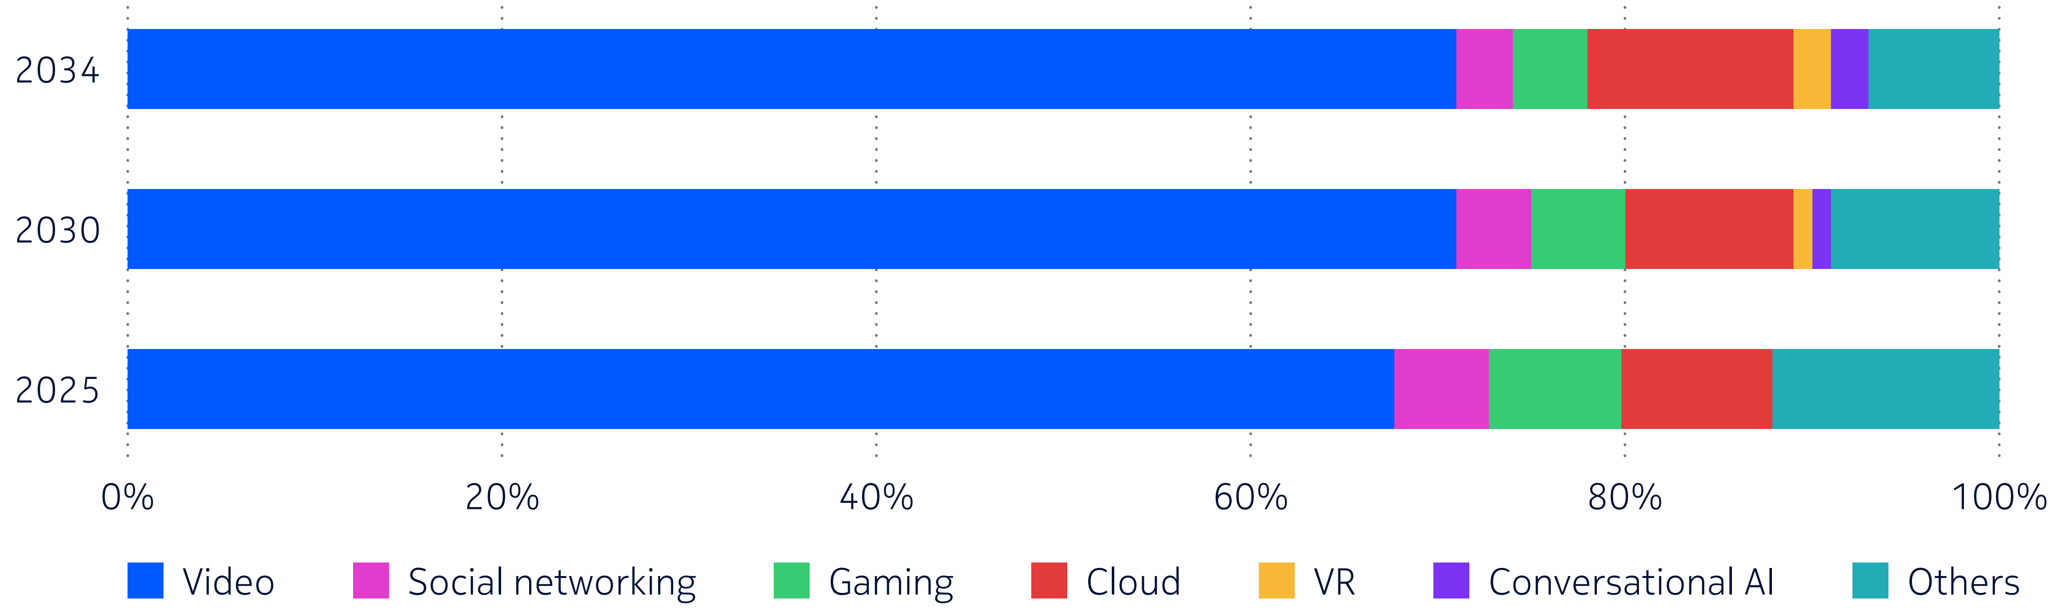

Consumer fixed traffic by region

Regional consumer fixed traffic is broken out in Figure 11, which shows how volume is distributed by 2034 in the moderate scenario. The leading region is China, which is projected to generate on the order of 500 EB/month of fixed consumer traffic by 2034. North America follows at roughly 430 EB/month, and the broader Asia-Pacific region outside China (including high-growth markets such as India and Southeast Asia) contributes on the order of 370+ EB/month.

This regional view tells two important stories. First, total demand is not concentrated in a single geography; high-capacity home access is now driving extremely large volumes in multiple regions at once. Second, the mix is shaped by both maturity and expansion.

In regions where high-speed internet is not yet widespread, consumer application developers are incentivized to prioritize bandwidth efficiency to ensure smooth user experiences, so per-user traffic consumption stays low even in high-speed hotspots. Conversely, in areas with widespread reliable, fast internet and higher manpower costs, like North America, developers may prioritize the speed of delivering software features over bandwidth optimization. Thus, the per-user traffic will be high even in areas with poor internet quality. In fast-scaling markets across Asia, traffic ramps as the critical mass of households moves onto high-capacity last-mile access, including fiber and densified wireless alternatives.

Figure 11: Consumer fixed traffic, moderate, excluding FWA, EB/month

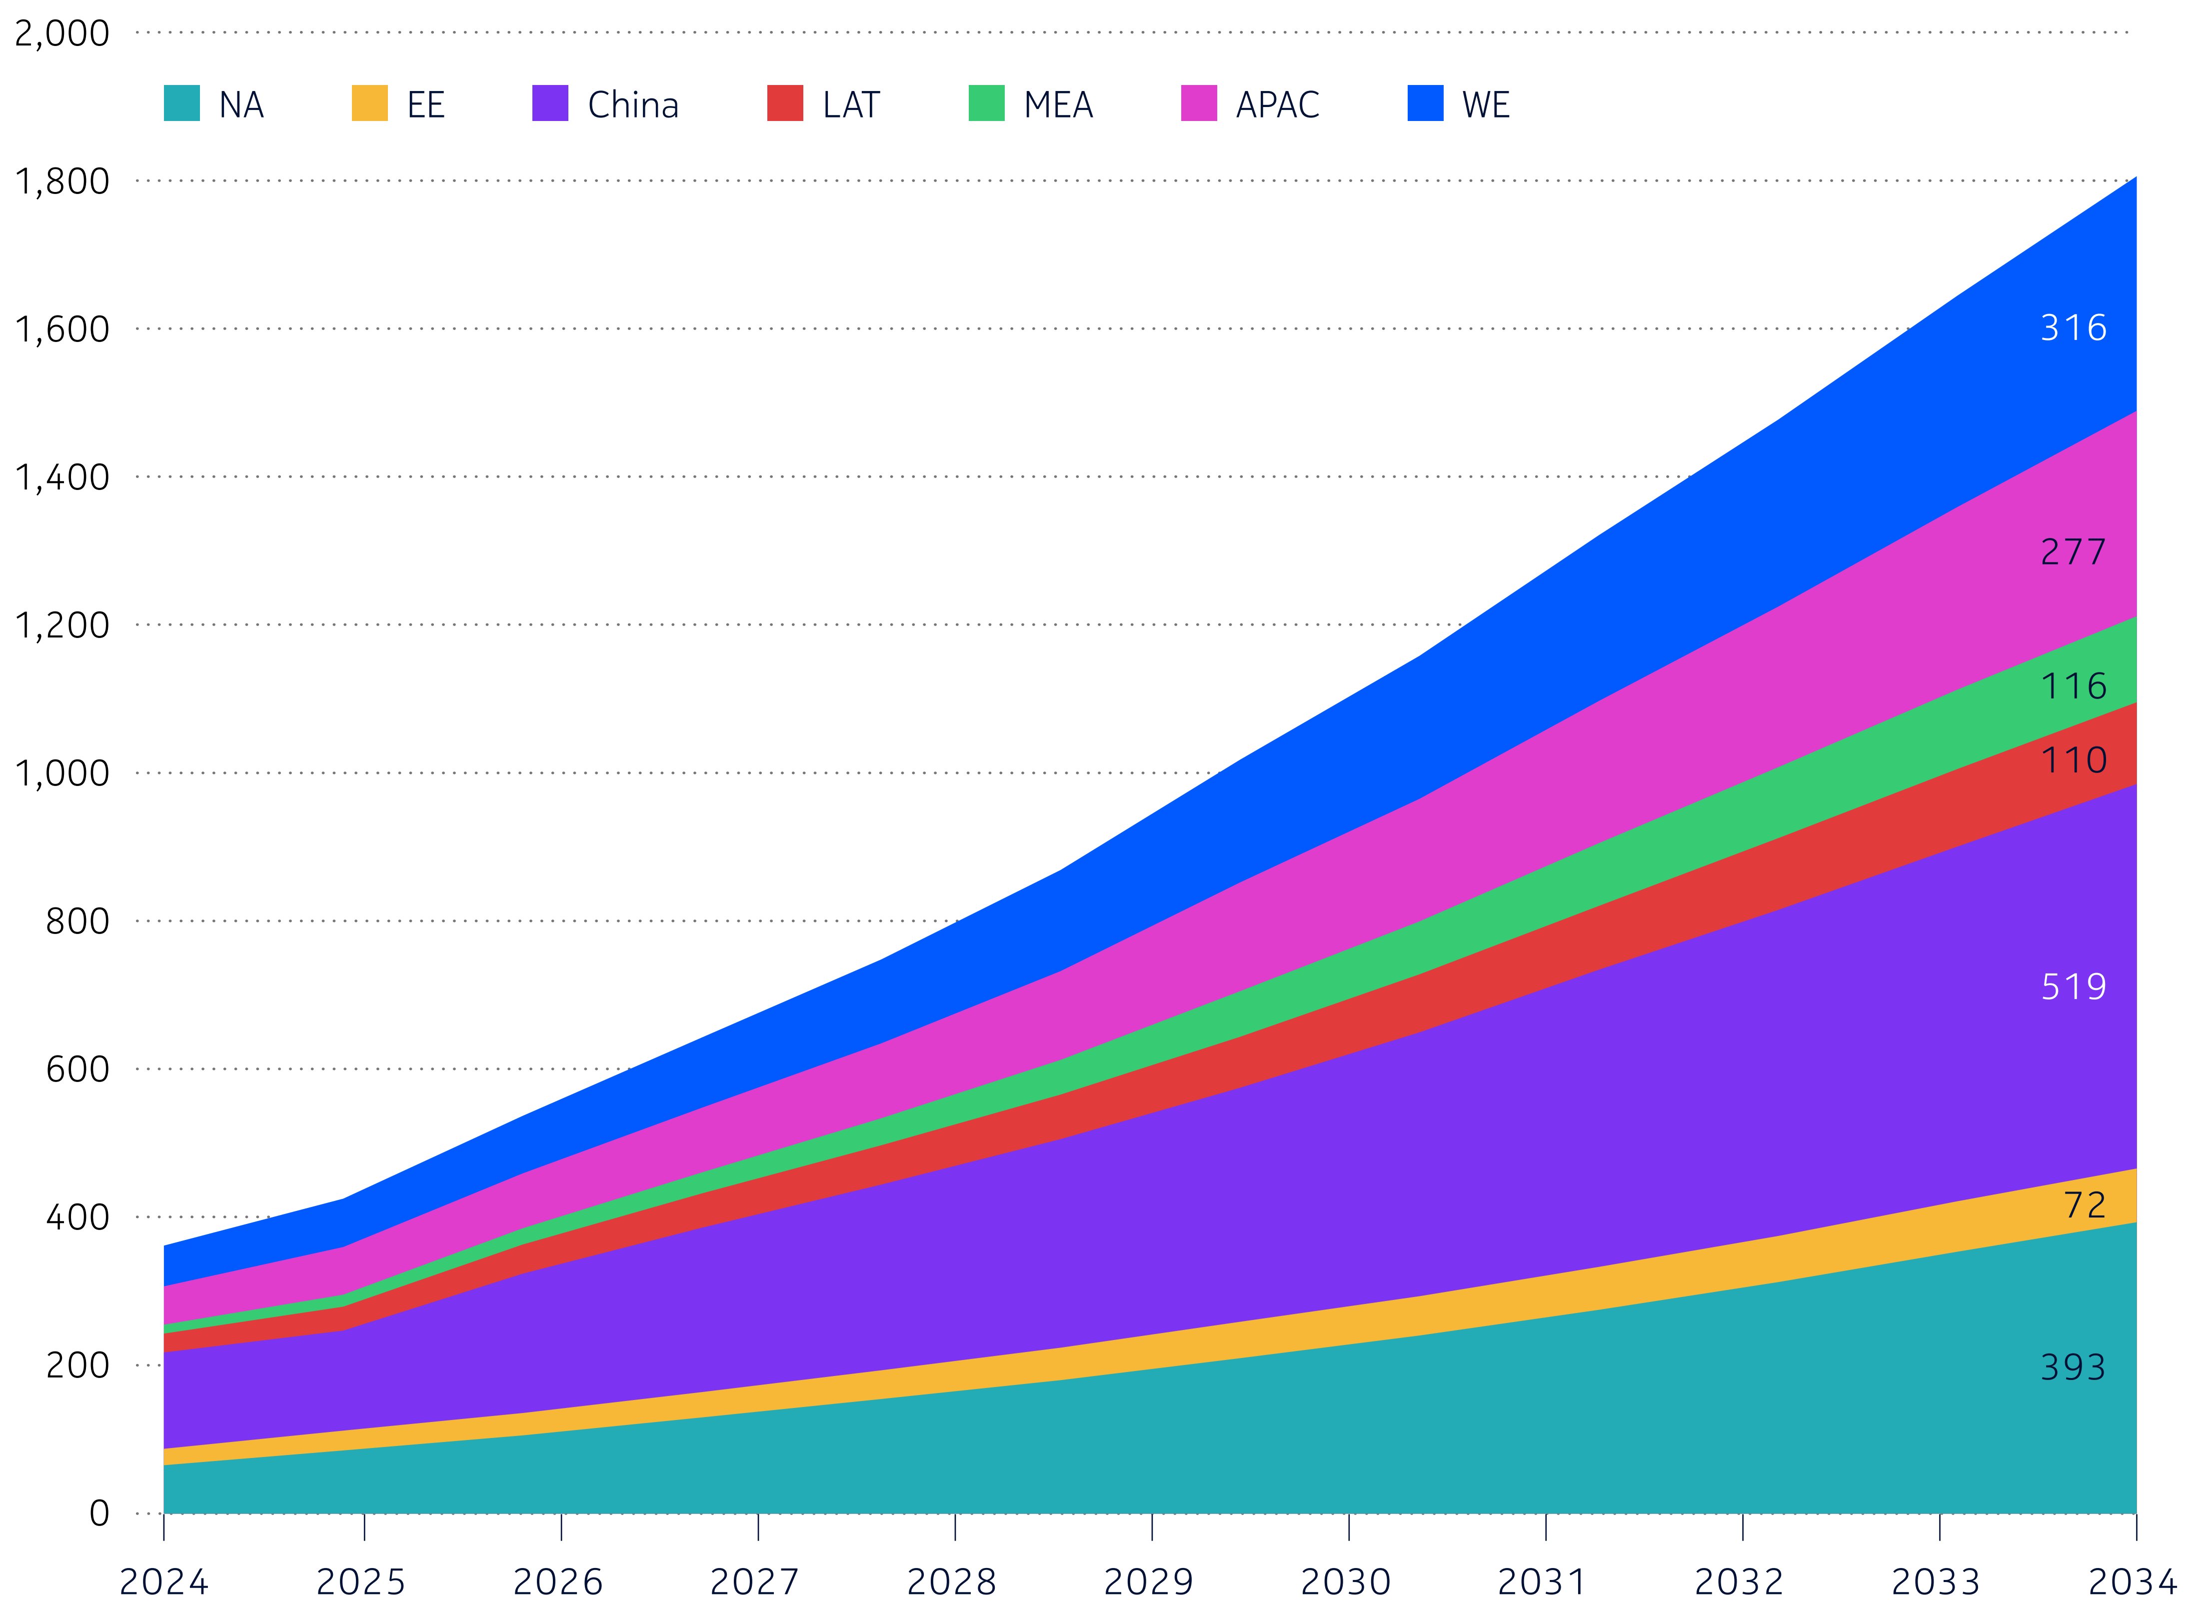

Enterprise and industry

Enterprise and industrial traffic is now growing as fast as—and in many cases faster than—consumer traffic. Just as consumers drove the first big wave of internet scale, production environments are now driving the next one. This isn't about basic office collaboration anymore. It's about operations. Factory lines are streaming machine vision data to the cloud. AI copilots are assisting personnel in real time. Field teams are using AR instead of manuals. Robots are coordinating across sites. Industrial systems are continuously sending telemetry over the WAN instead of keeping it on-site. This shift makes wide-area connectivity part of the core production workflow. It's no longer an optional layer or an IT add-on. The next two pages break down where this traffic is coming from and how quickly it's accelerating.

Enterprise and industrial traffic

Enterprise and industrial traffic is set to scale rapidly over the next decade as more operations, machines and workers become digitally connected. Figure 12 shows our forecast for global enterprise and industrial traffic, including fixed wireless access. Across all scenarios—conservative, moderate and aggressive—traffic rises steadily, year over year through 2034. The moderate case more than doubles over the period, and the aggressive case pushes to 913 EB/month by 2034, driven by pervasive automation, high-resolution video, AI-driven analytics, and remote access to industrial systems.

A notable trend is the shift in the composition of enterprise traffic. Traditional enterprise traffic, such as cloud tools and data sharing, will continue to dominate, accounting for around 57% of traffic in the moderate scenario by 2034. However, the share of industrial traffic will grow significantly—with a CAGR of 5% for IoT traffic and 50% for XR traffic between 2024 and 2034.

The total traffic volume in 2034 will largely depend on the types of use cases adopted by enterprises, which in turn will be influenced by the quality of internet infrastructure available in different regions. For instance, some enterprises may rely on low-speed AR devices or AR applications running on consumer-grade devices like tablets, smartphones or glasses. On the other hand, industries with access to high-quality internet may adopt high-definition and immersive XR technologies for more demanding tasks, such as construction design, product design, medical tasks and employee training.

Figure 12: Enterprise and industrial traffic, global, EB/month

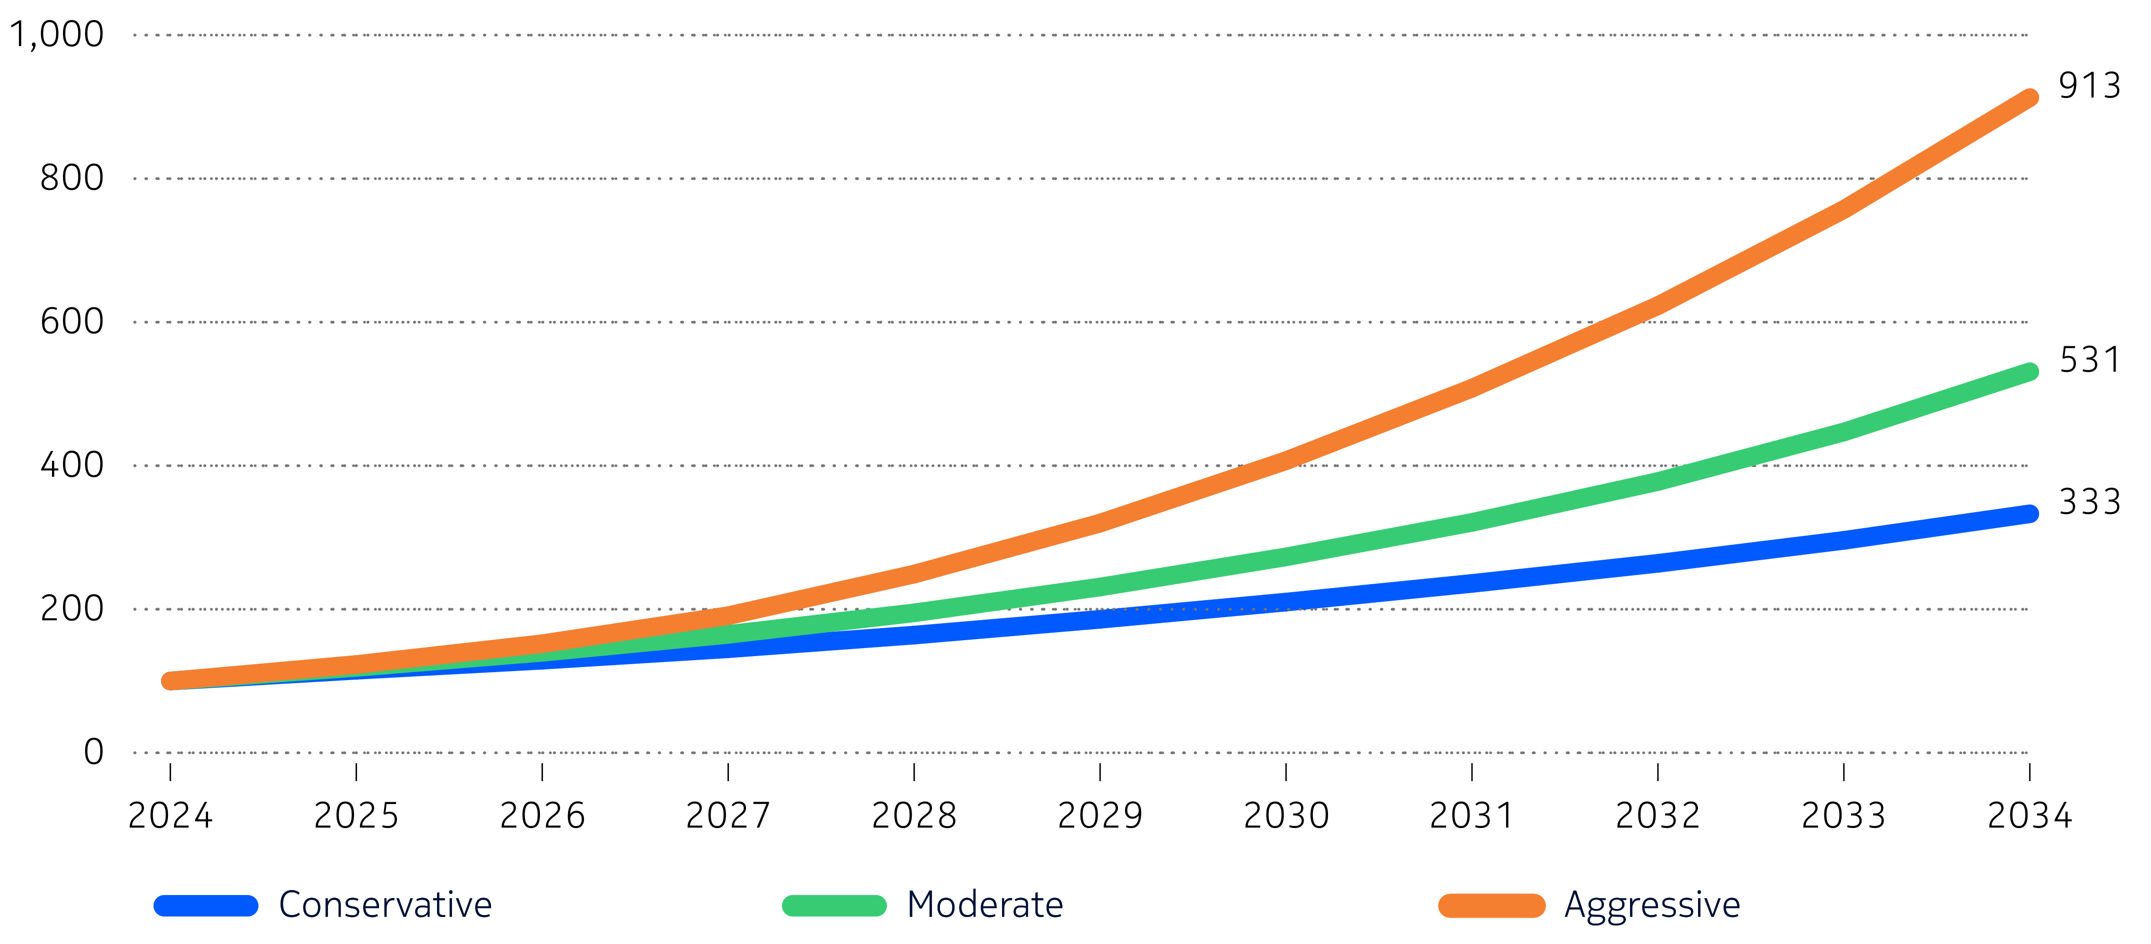

Figure 13: Enterprise WAN services traffic, share, 2034

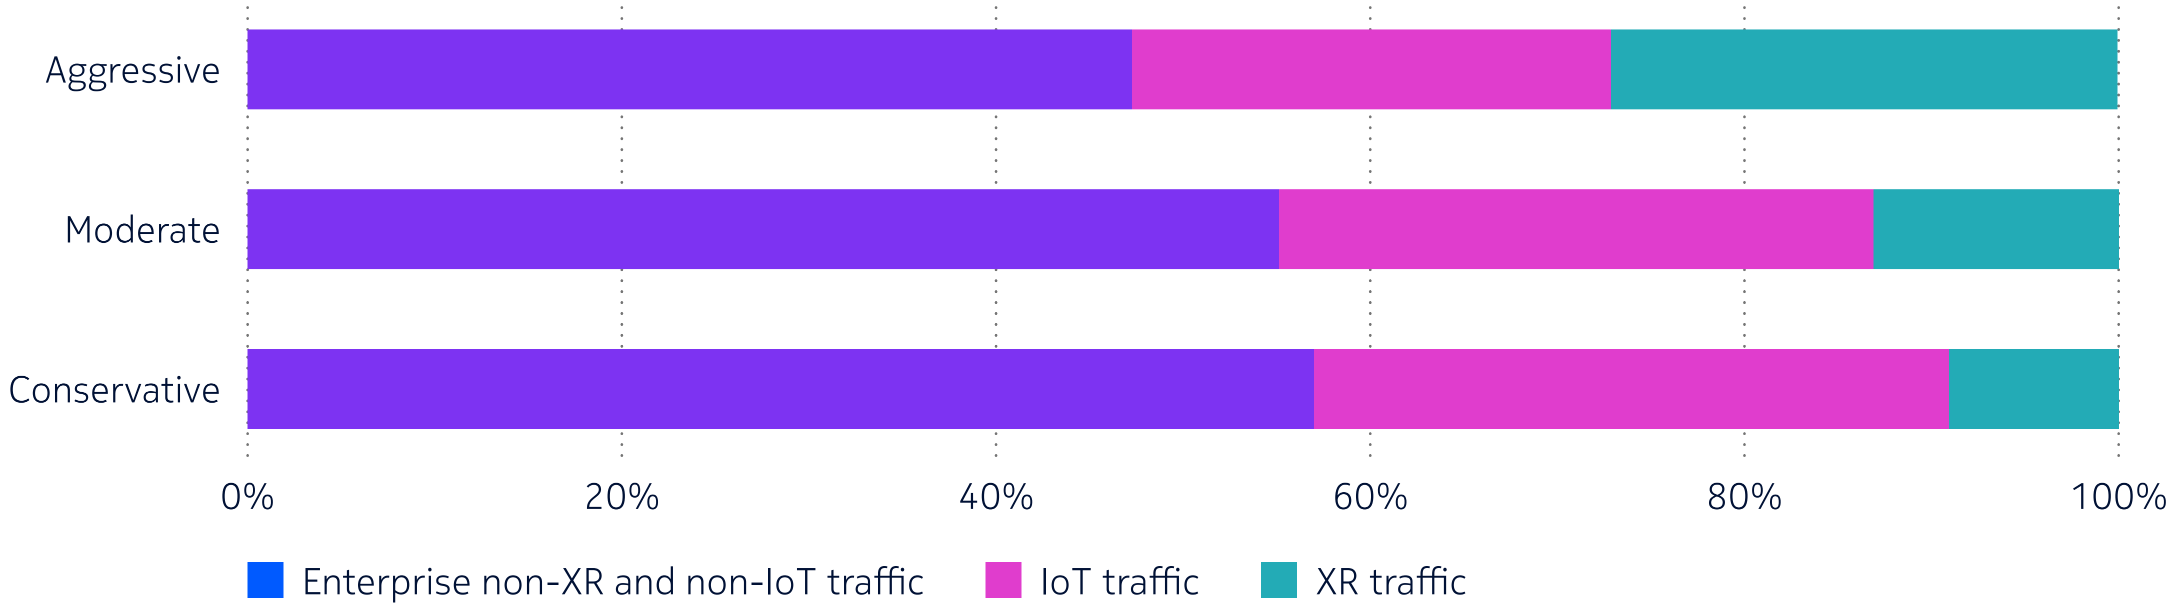

Enterprise and industrial traffic by region

Enterprise and industrial traffic rises in every region through 2034, and the curve in Figure 14 bends upward more sharply after 2030. What's driving that acceleration is the availability of high-speed internet, scaling of AI workloads, video inspection, sensor data, and automated control systems inside enterprises.

APAC is the largest contributor throughout the forecast period and remains the main engine of global growth. North America and Western Europe also grow quickly and represent a major share of total traffic by the end of the period as cloud, AI and private wireless use expand. Latin America, the Middle East and Africa, and Eastern Europe start from a smaller base but increase steadily rather than staying flat.

The result is that enterprise and industrial traffic becomes a global story. APAC leads, but every region is contributing meaningful growth and adding pressure to backbones and interconnect capacity.

Figure 14: Enterprise traffic, global, EB/month

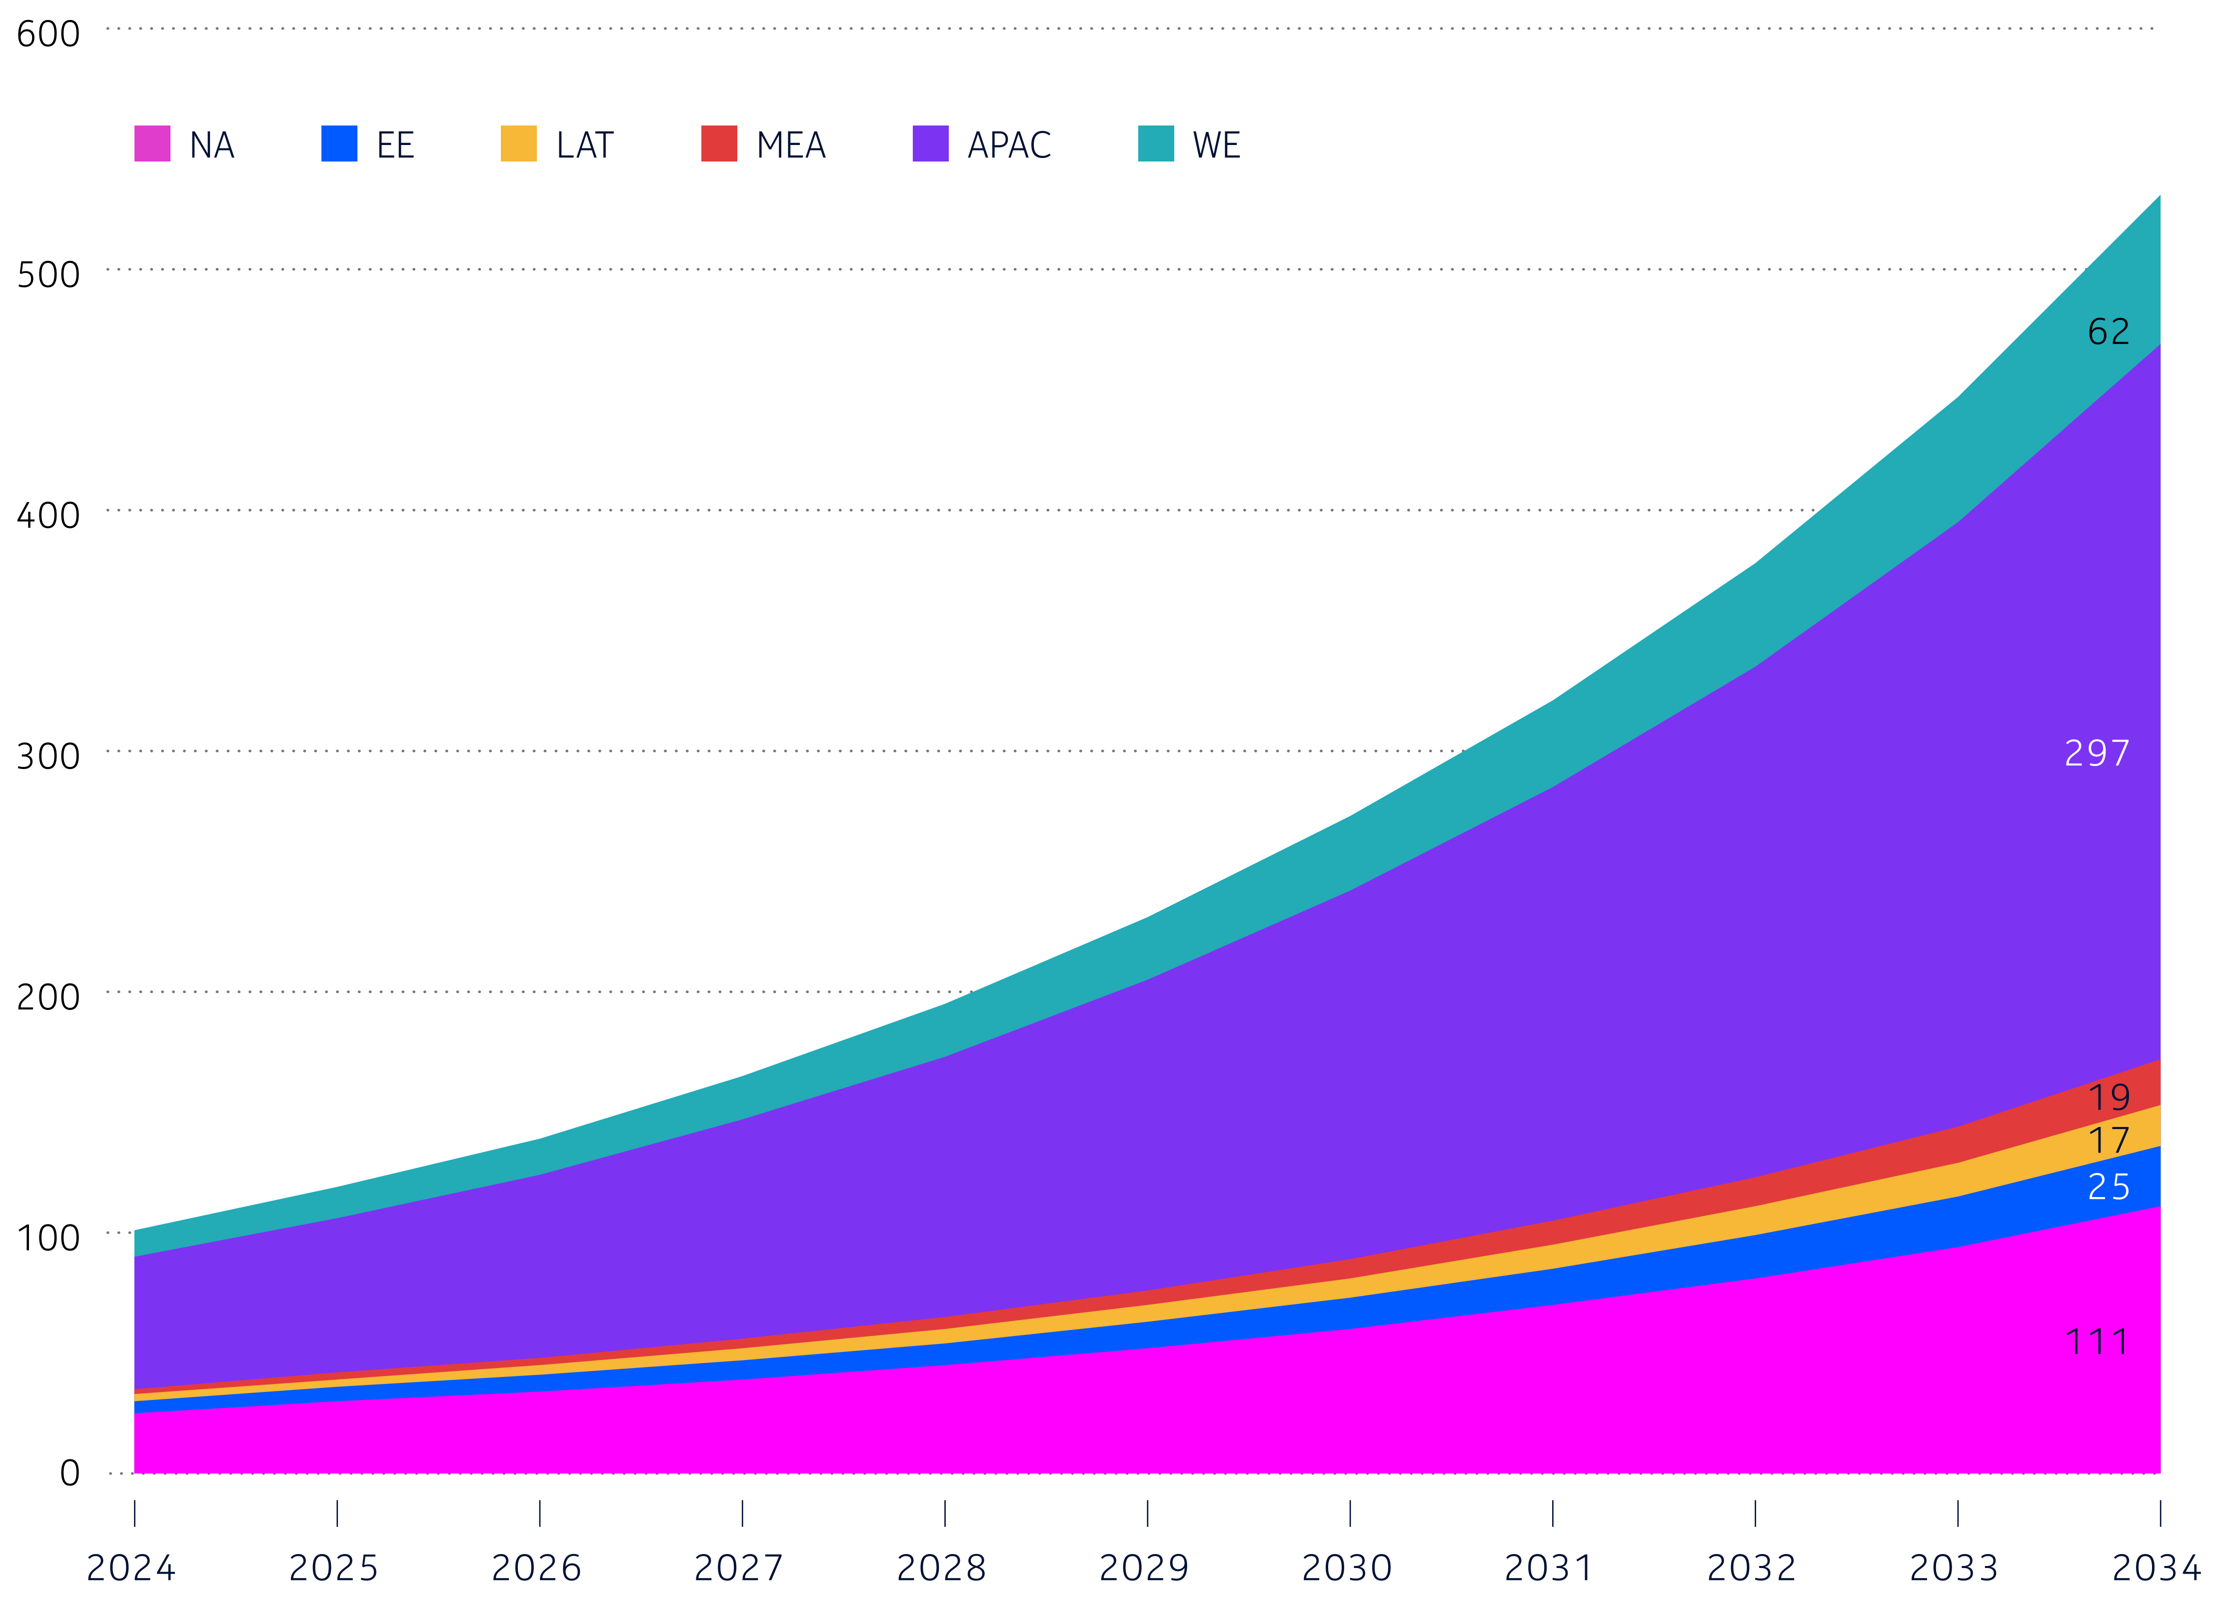

Impact of AI on networks

Consumer and enterprise user-generated AI traffic imposes a substantial impact on the wide area network (WAN) by adding AI workloads to be processed by data centers across the WAN. AI traffic does not stay inside one data center. It moves across edge, metro, core and cloud infrastructure, driving dense lateral flows and new capacity demands. The explosion of agentic-AI applications further fuels it by inducing extra machine-to-machine (M2M) traffic in the background. The following section shows our forecast of agentic-AI traffic, how AI traffic moves through the network, and explains how these traffic types are reshaping network design.

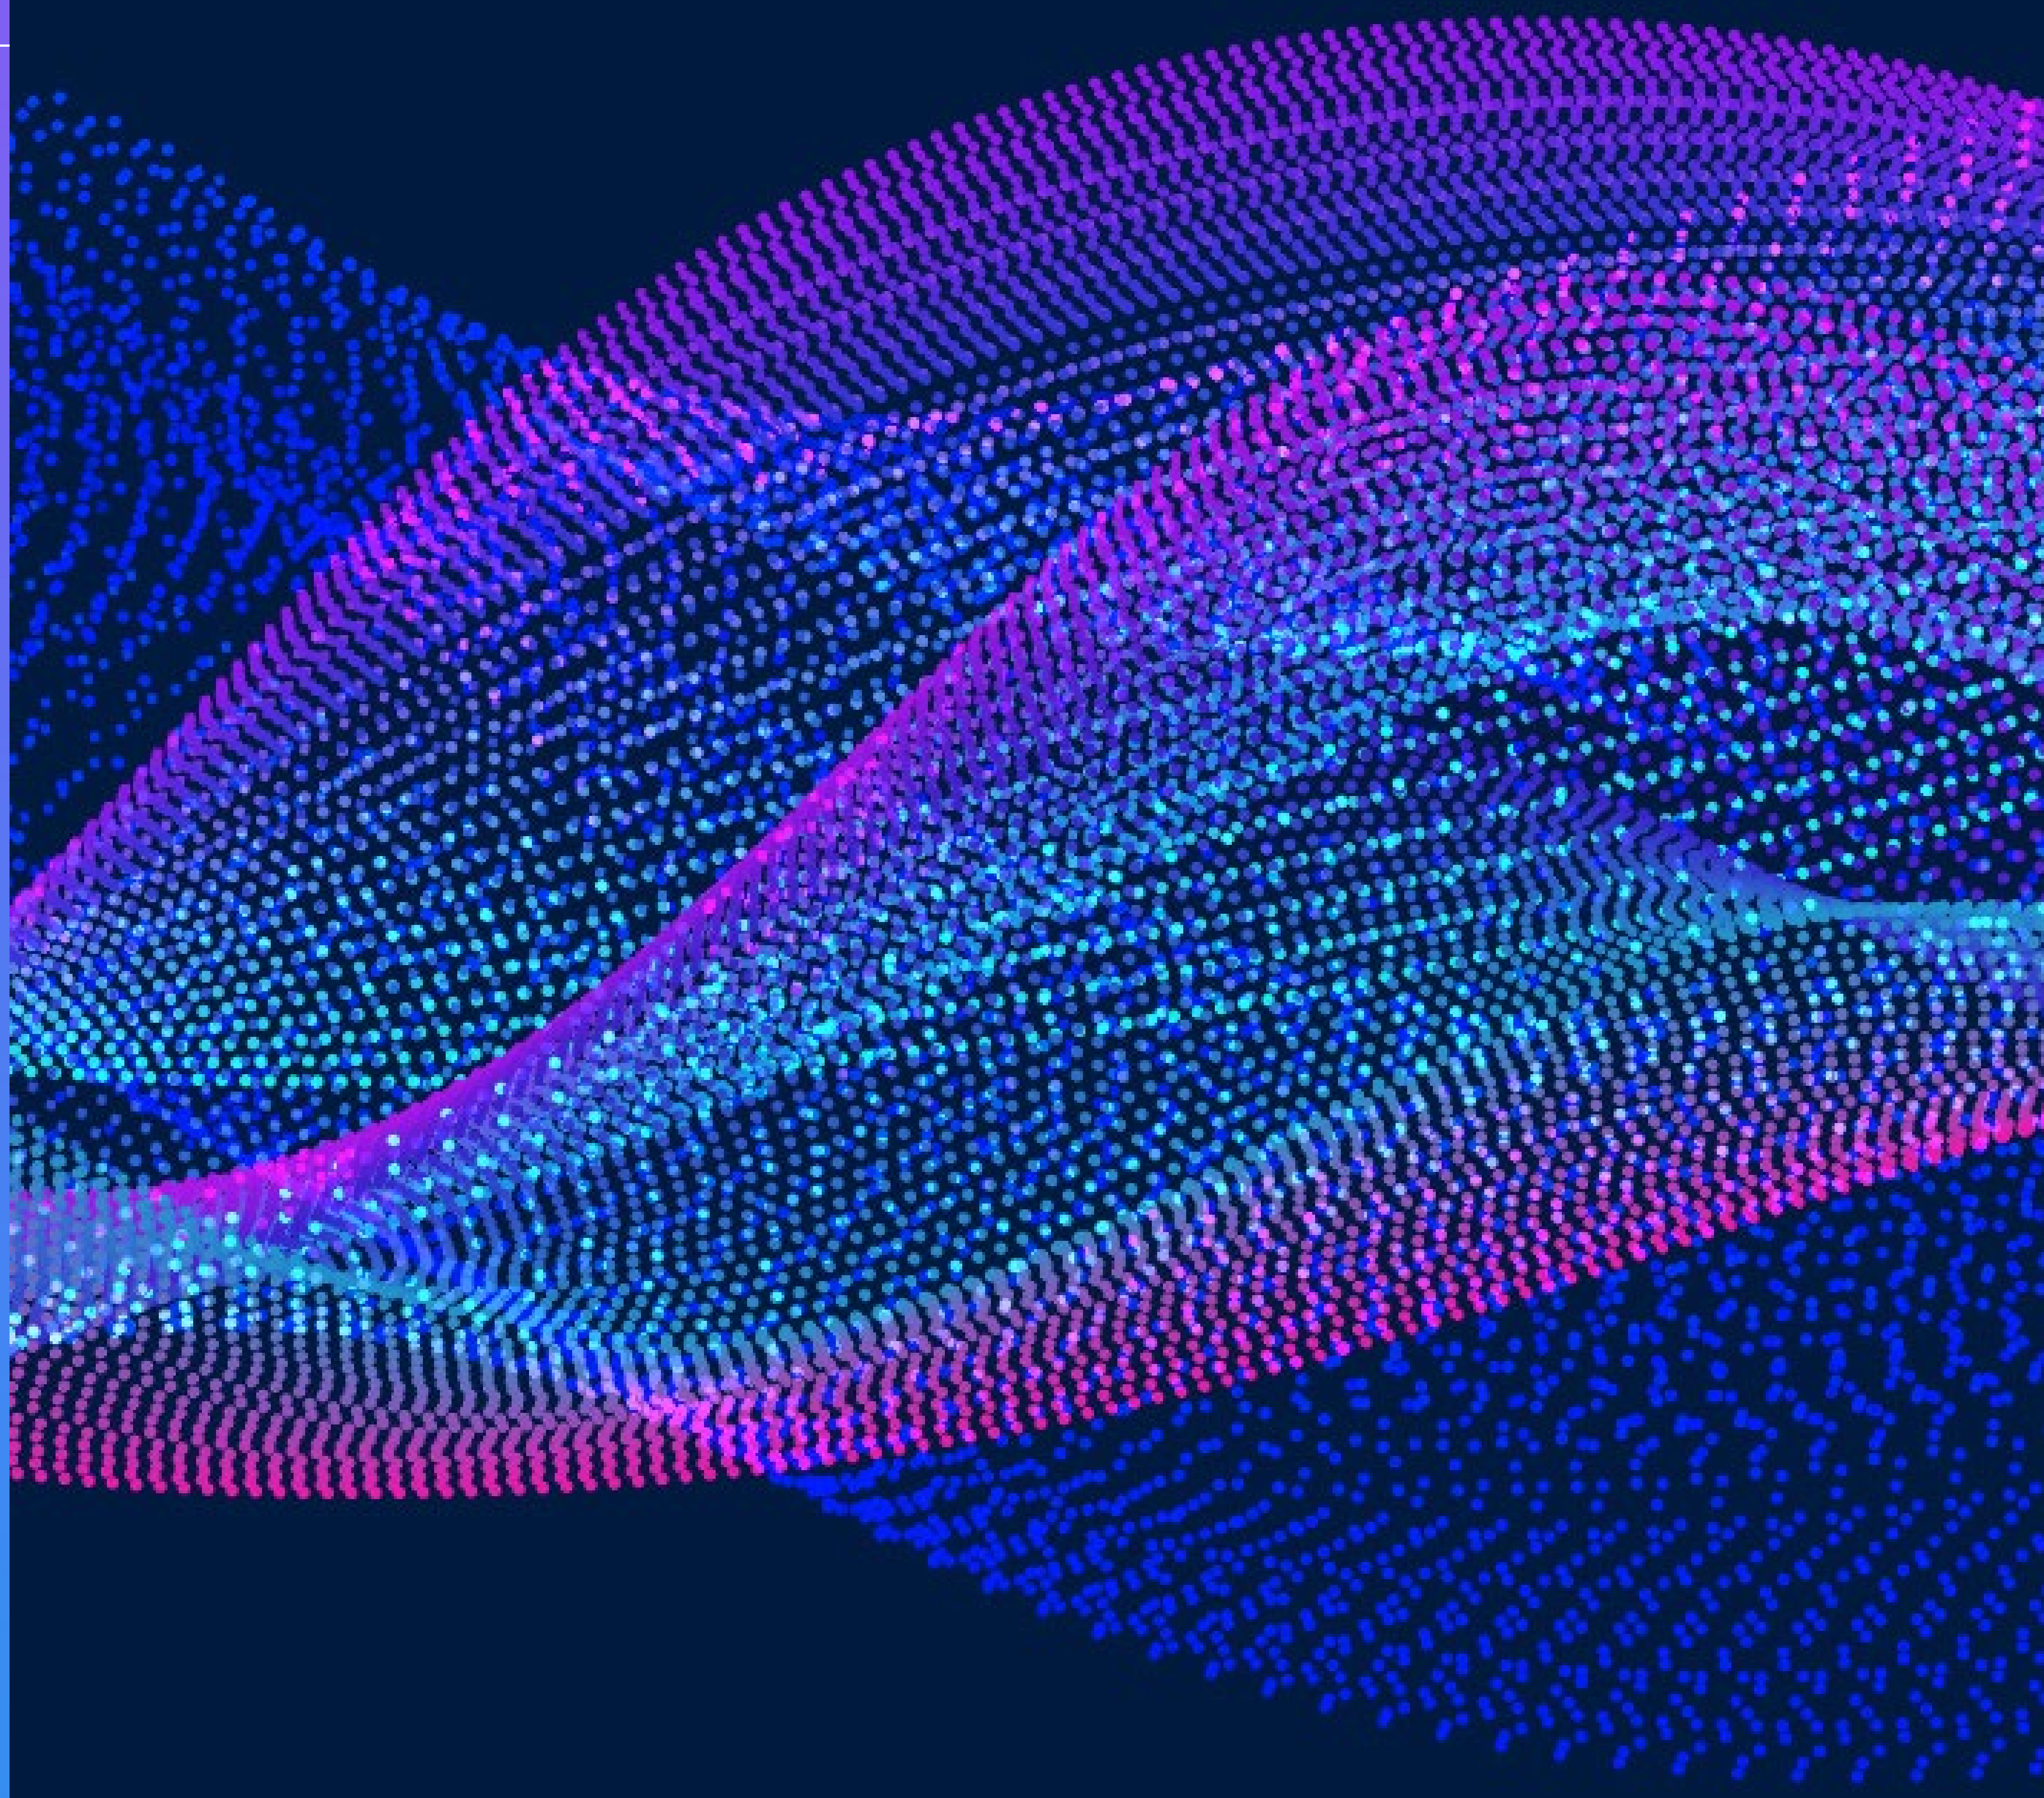

Machine-to-machine traffic induced by agentic AI

Agentic AI adds a new layer of interaction—introducing autonomy, planning, and retries—and spans LLM, RAG, and generative flows such as content creation and personal assistance. It doesn't just increase intelligence; it turns task execution into an autonomous loop, dramatically increasing machine-to-machine interactions and background traffic. These interactions include LLM prompts, web searches, database queries, and the safety checks agentic-AI applications run to produce responses for users.

The global agentic-AI M2M traffic is expected to increase from 66 EB/month in 2025 to 537 EB/month in 2034 (CAGR of 26%), which will result in a significant increase in the amount of AI inference traffic that is transmitted across the wide area network (WAN).

Figure 15: WAN agentic-AI traffic, global, moderate, EB/month

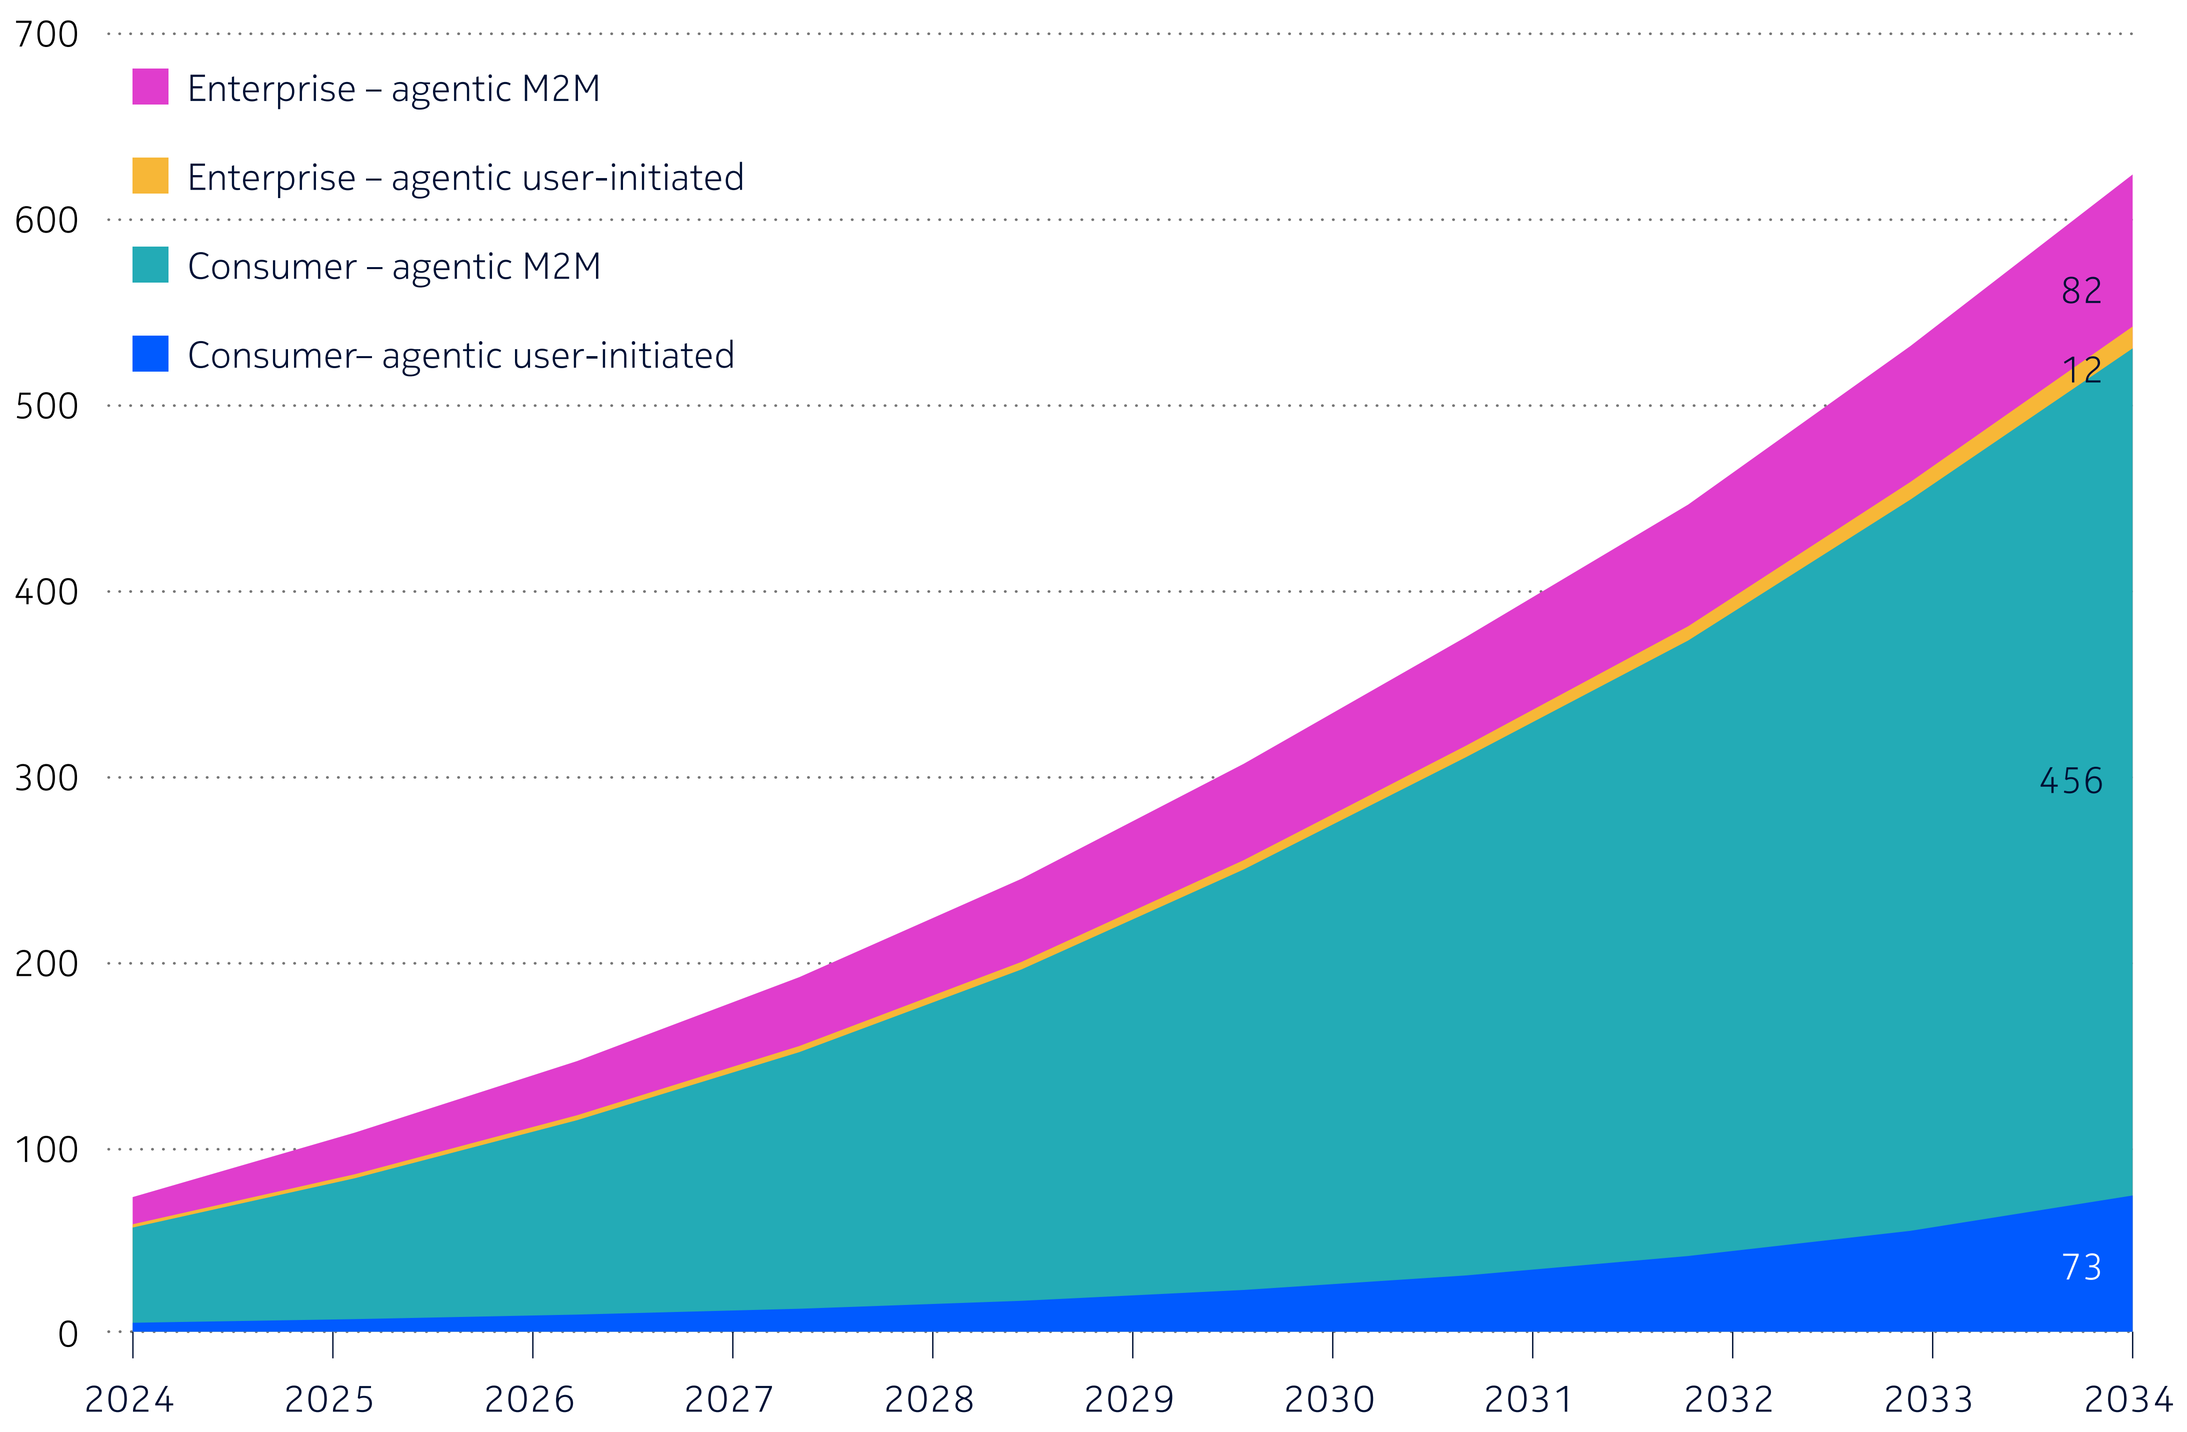

AI traffic over inter-DC links

AI traffic isn't just creating more demand inside data centers; it's pushing a sustained surge of traffic between them. AI inferencing traffic—both user-initiated and agentic-AI-induced machine-to-machine—moving over inter-data-center links grows at a 20.3% CAGR through 2034. Operators, cloud providers, and large enterprises should plan for persistent scale-out of interconnect capacity, not occasional bursts.

Here's what's happening in Figure 16: AI inference and training generate bursts at the metro edge, then data is rapidly replicated and moved across metro cores and regional sites for scaling, storage, policy enforcement, and resilience. Workloads hand off between hyperscalers, second-tier clouds, telecommunications provider infrastructure, and enterprise data centers. Along the path, agents and RAG systems continually call models, move context, sync state, and pull data from other zones.

The result is that a single AI session can traverse several inter-DC links in sequence. Each hop multiplies total carried traffic—especially inside and between metros and regions. By 2034, this becomes a dominant driver of WAN engineering: networks must deliver high-capacity, low-latency, policy-compliant connectivity across many different data-center owners, not just bigger pipes to a single cloud.

Figure 16: AI traffic impact on the WAN in 2034 (over inter-DC link 1-12)

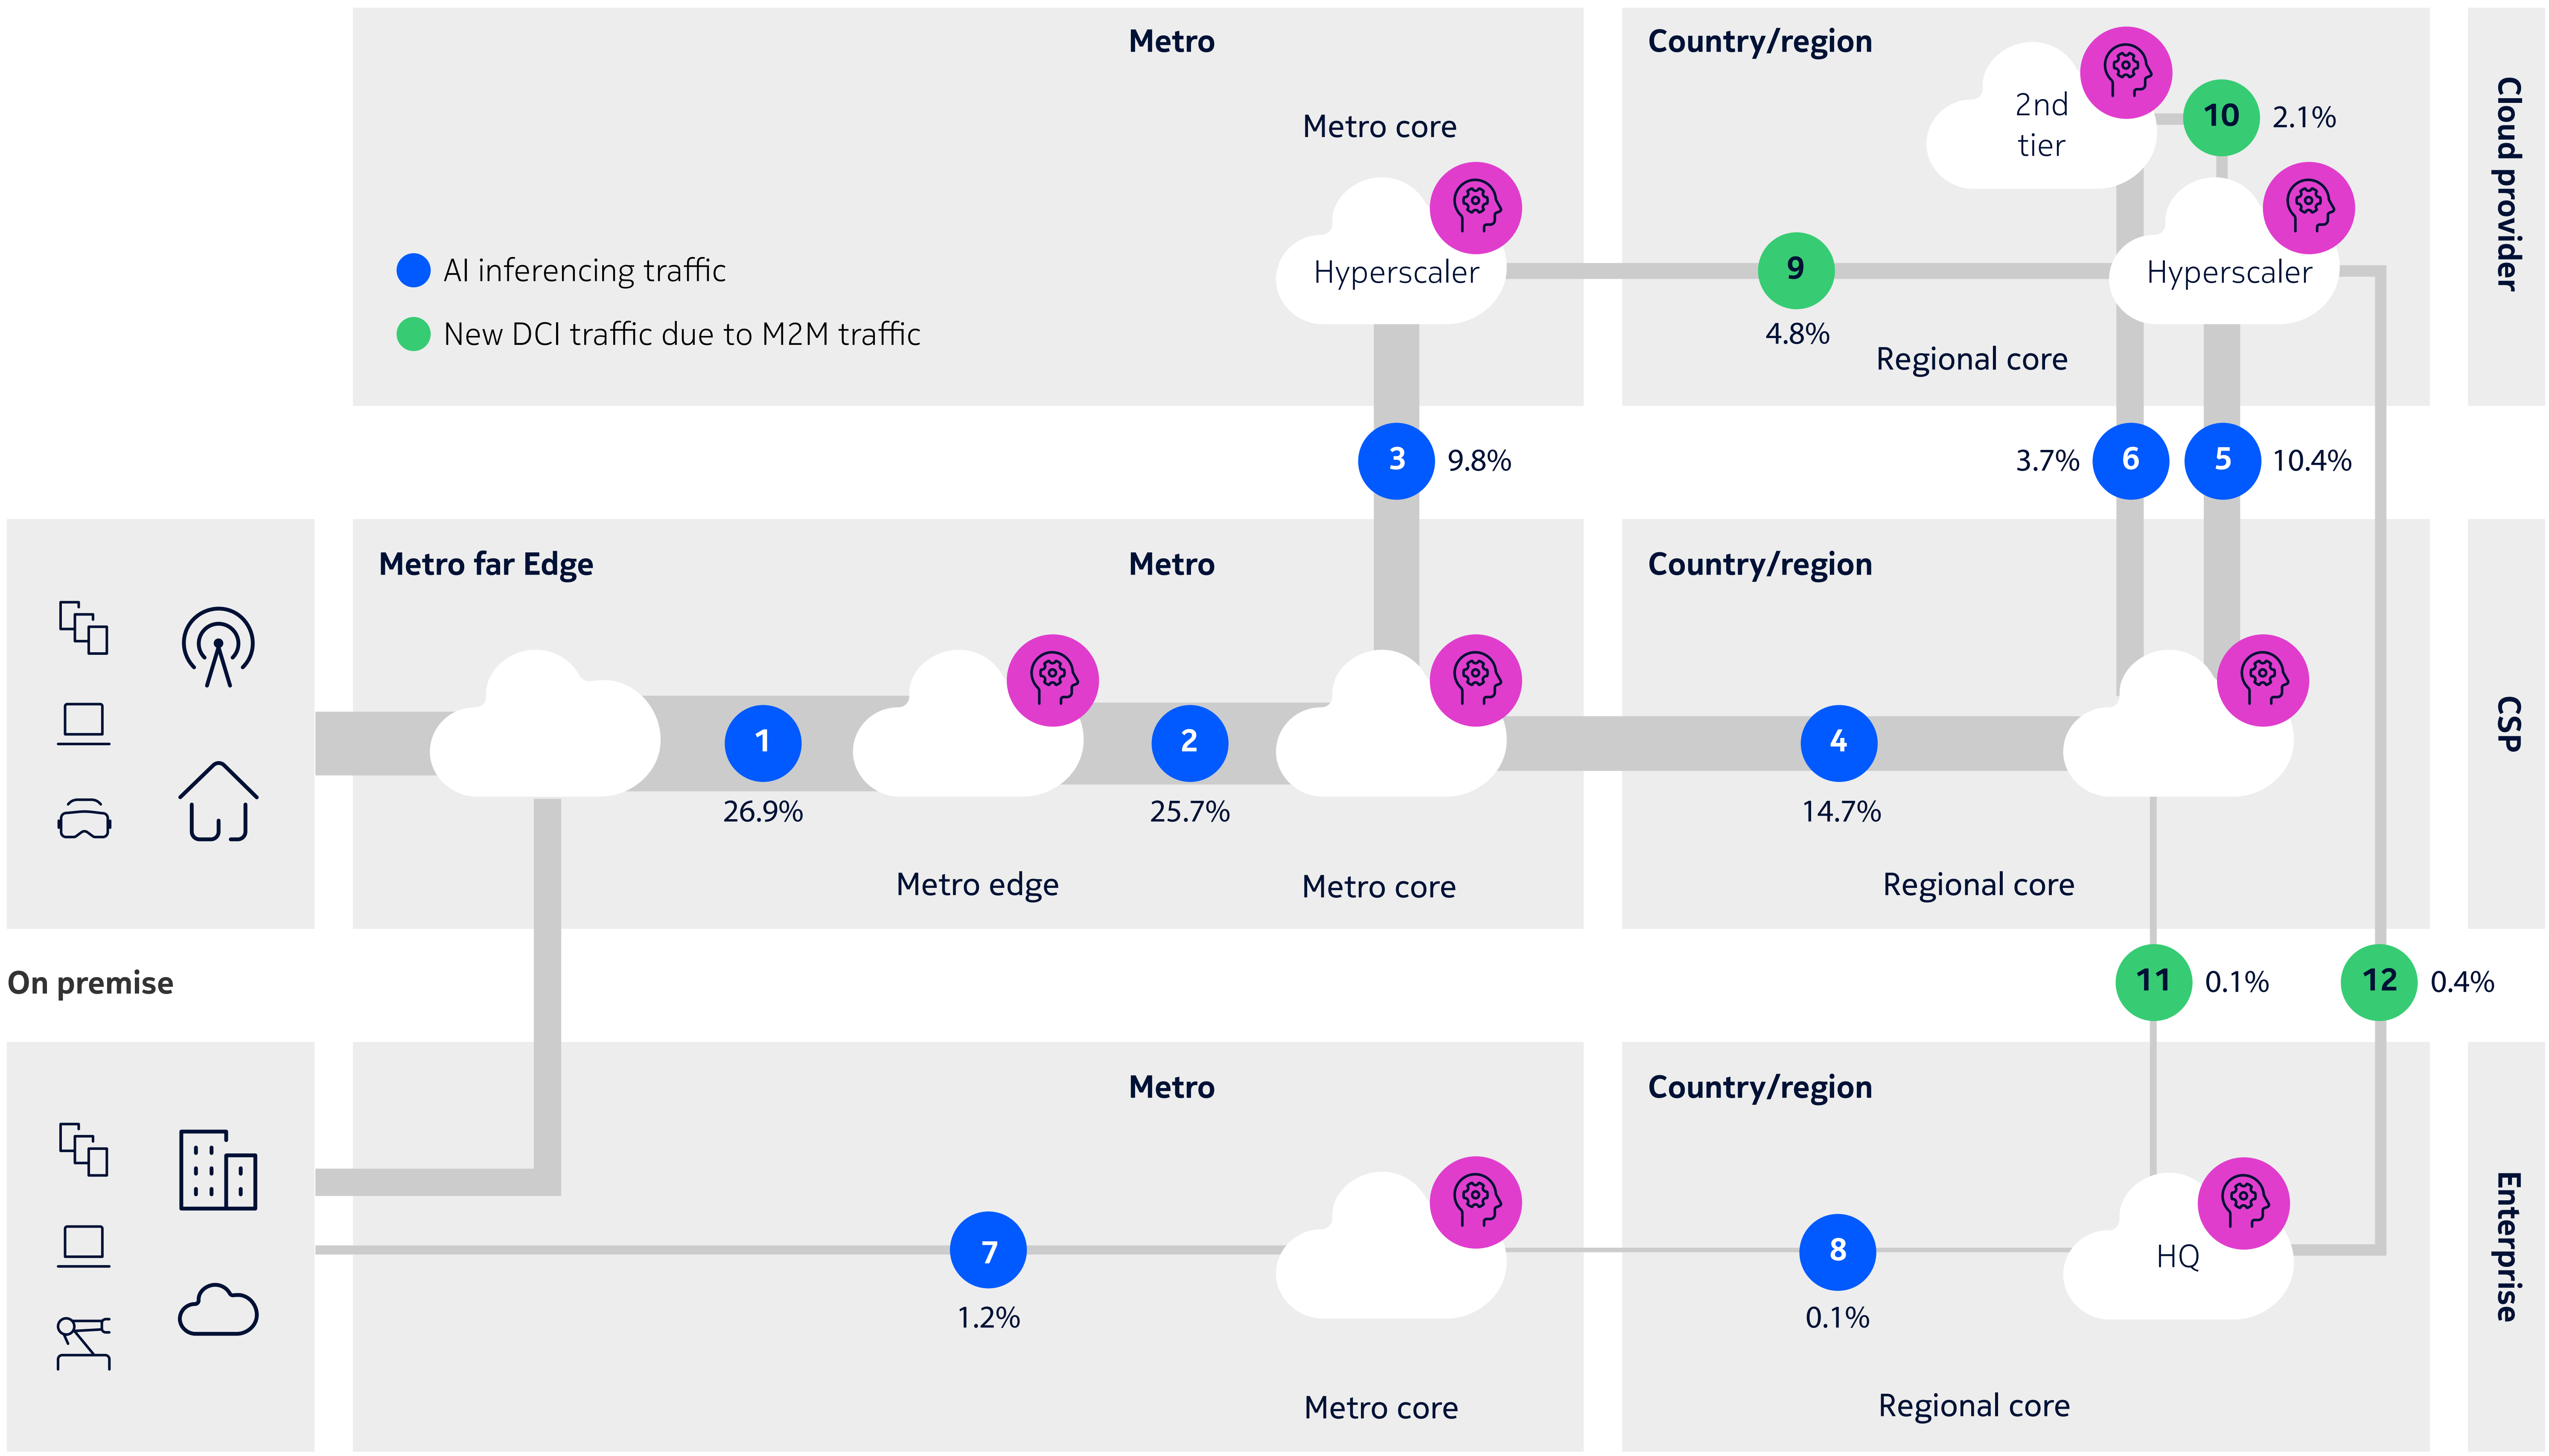

User source generated AI inference traffic reaching 921 EB/month in 2034, creating 3.5x traffic of 3,260 EB/month over the WAN inter DC links, 67.4% of which is over telecommunications provider transport (1, 2, 4), 23.8% over telecommunications provider-cloud IDC links (3, 5, 6) and 7.0% over Cloud-cloud IDC links (9, 10).

Our conclusions

Global network traffic is now scaling in three reinforcing ways: consumer, enterprise and AI. Total global WAN traffic is projected to grow roughly three to seven times by 2034, reaching about 2,277 to 4,878 EB/month, at a compound annual growth rate between 13% and 22%. Video remains the single largest source of volume, but AI is now the biggest new driver of growth.

AI changes how traffic is generated, where it flows and when it peaks. It increases uplink use in the home, it injects automation and machine vision into industrial sites, and it multiplies east–west movement between data centers.

Enterprise and industrial traffic is no longer a secondary story. It is growing as fast as or faster than consumer traffic, led by AI copilots, robotics, agentic industrial automation, AR experience, and always-connected operations.

By 2034, APAC carries the largest share of this load, reflecting large-scale industrial digitalization. At the same time, AI traffic has become its own category: both human-initiated prompts and fully autonomous agent-to-agent activity are already visible in the WAN and are projected to reach hundreds of exabytes per month each.

This AI workload does not stay inside one data center. It drives persistent demand for high-capacity, low-latency inter-data center connectivity across metro edge, metro and regional core.

Networks now must scale for this pattern: more symmetric traffic, tighter latency and massive interconnect. That is the design point for the next decade.

Global network traffic is projected to grow roughly 3x to 7x through 2034.

Figure 18: Global WAN traffic, EB/month

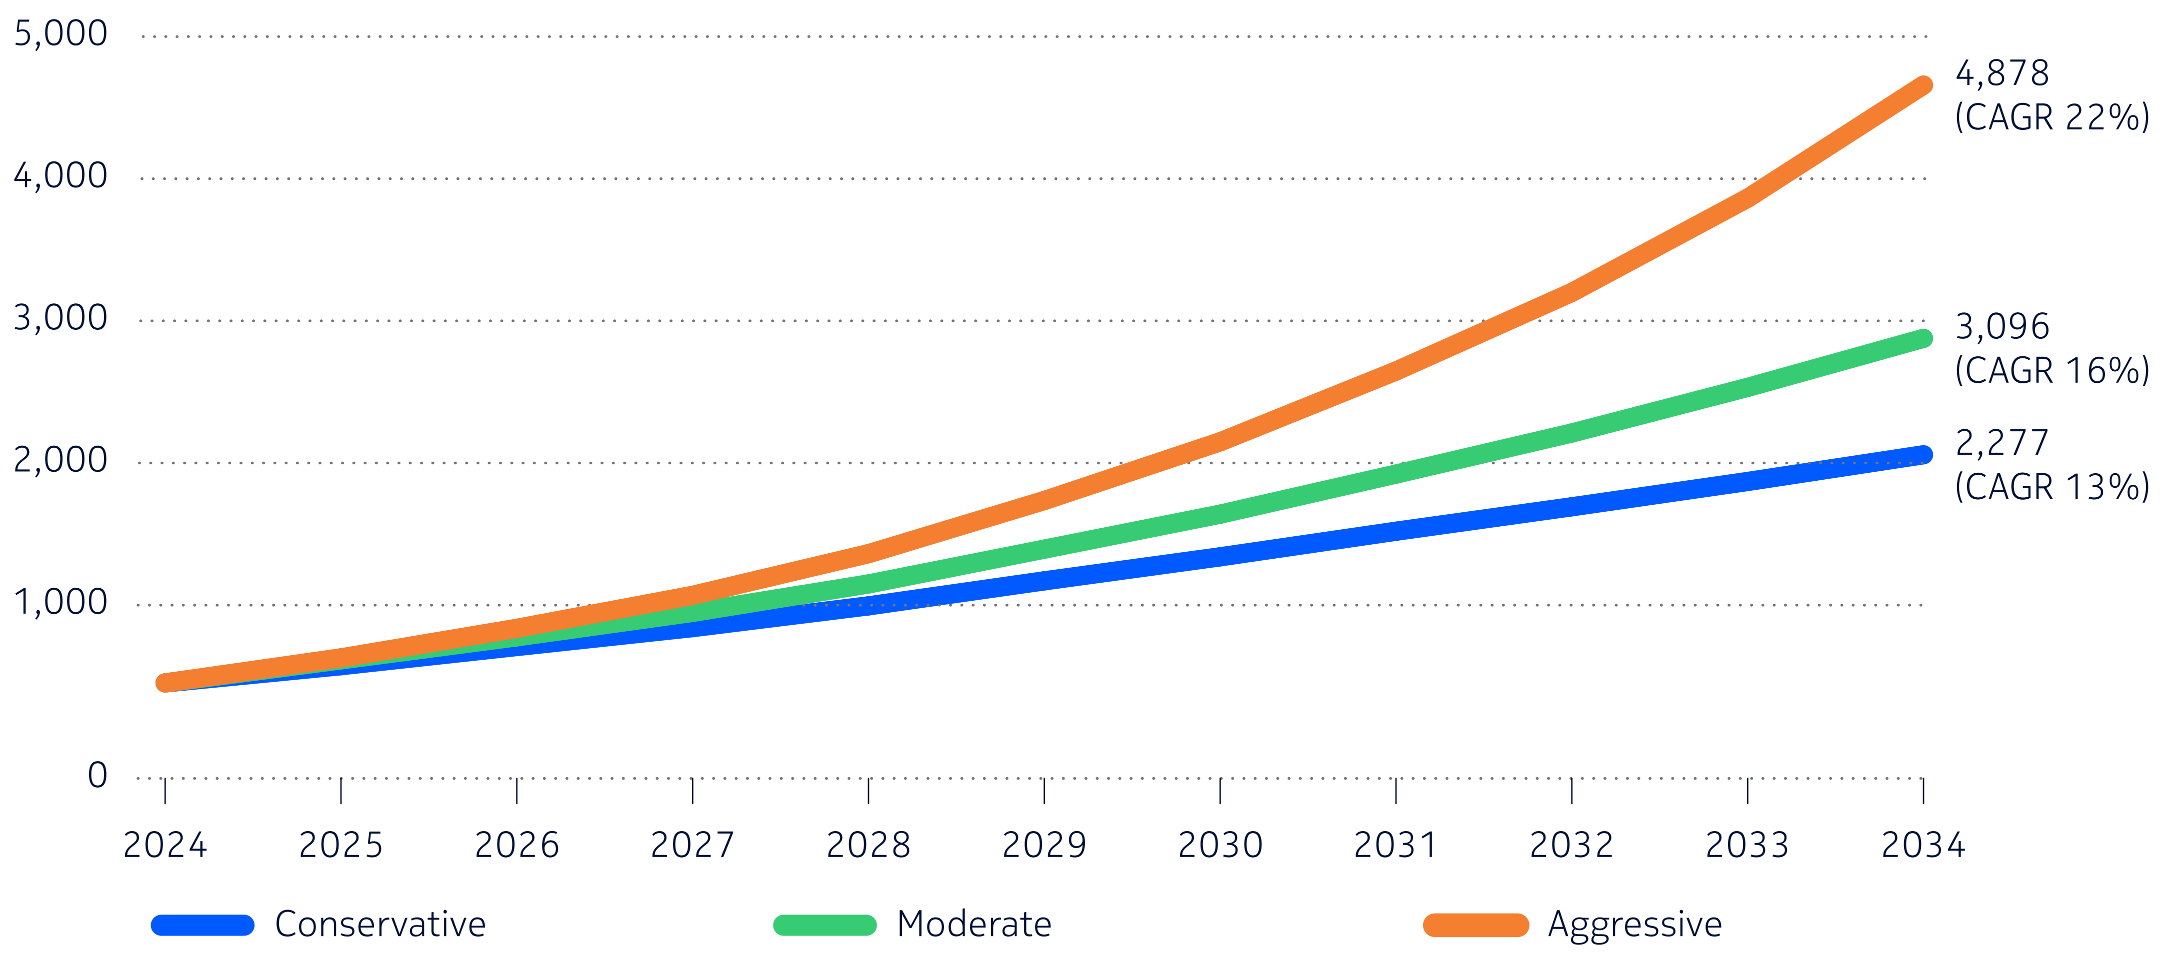

Global telecom bandwidth demand is projected to increase at a CAGR of 13%-22% from 2024 to 2034 to reach 2277 to 4878 exabytes per month.

References

Our global network traffic forecast combines operator-reported traffic data, public datasets, national statistics, market analysts reports on device shipments and proprietary network measurements. We normalize these inputs into a single global WAN model that covers Internet, video and managed/enterprise WAN usage from 2024 through 2034, excluding campus-only LAN traffic.

We incorporate demographic, infrastructure and digitalization data from national and regional statistical agencies (for example, the U.S. Bureau of Labor Statistics, U.S. Census, China Statistical Yearbook and agencies in South Africa, Brazil and Poland), along with telecom market baselines from GSMA Intelligence and GlobalData.

Our modeling of AI's impact on network demand integrates external research on AI workload behavior, data center interconnect patterns, and inference scaling. This foundation is supplemented by projections for user-driven AI traffic and agentic machine-to-machine traffic. We validate these projections against real-world customer network observations and employ scenario forecasting (conservative, moderate, aggressive) to account for uncertainties in the adoption of AI, private WAN, industrial automation, and Fixed Wireless Access (FWA).

Nokia OYJ
Karakaari 7
02610 Espoo
Finland

Tel. +358 (0) 10 44 88 000

CID:214296

nokia.com

At Nokia, we create technology that helps the world act together.

As a B2B technology innovation leader, we are pioneering the future where networks meet cloud to realize the full potential of digital in every industry.

Through networks that sense, think and act, we work with our customers and partners to create the digital services and applications of the future.