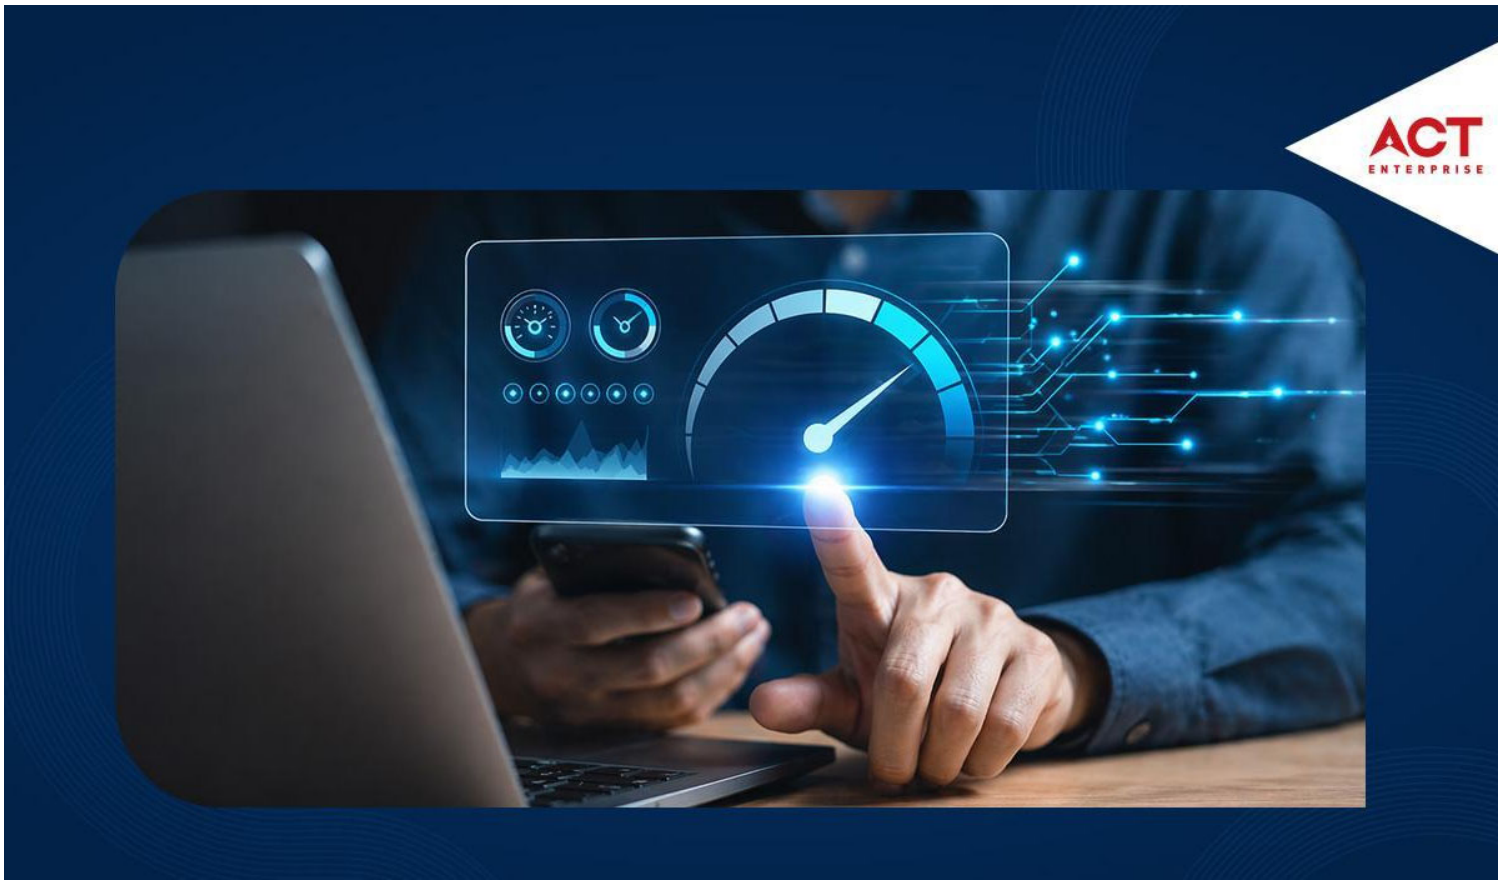

Why bandwidth alone can’t power agentic AI

The Connectivity Gap: Why Bandwidth Alone Can’t Power Agentic AI

Across the Indian enterprise landscape, a tech-led revolution is taking place. Executive conversations have moved on to the latest wave of AI: agentic systems. From healthcare providers to logistics operators, Fintech start-ups and retail majors, a variety of businesses are moving beyond simple chat-bots and assistants to AI agents.

AI agents refer to autonomous systems that don’t just respond to queries but actually reason through tasks, handle multi-step workflows, and make context-based decisions. They are expected to do all this with little or no human overseer, across hybrid operating environments including cloud.

Market analysts forecast that by 2035, agentic AI is likely to generate about 30% of revenue from enterprise software. Encouraged by this potential, Indian enterprises are leading in agentic AI deployment. They’ve made considerable investments based on the tech’s promise to reimagine how they work, are organised and innovate.

Today, India’s businesses stand at crucial juncture in this revolution where its long-term success will depend on multiple factors. One of these is their network infrastructure. Because, in the rush to deploy world-class SaaS and the most advanced AI models, many are hitting a physical wall.

Many Indian enterprises are forced to operate on a paradox: **cutting-edge intelligence systems running on fragile network infrastructure.**

This introduces an important challenge: If your network isn't aware of the application it's carrying, your AI is only as smart as your slowest packet.

Why an Enterprise-Grade Backbone is necessary

Many businesses assume that gaining more Mbps in internet speed is the solution to every digital lag. But when implementing AI agents, raw bandwidth is only a cosmetic improvement. What these agents actually need is predictability**.** This comes from guaranteed throughput.

To understand why, let's explore how the internet is set up and how data moves. This understanding is crucial as we move from casual browsing to running high-stakes applications powered by AI.

Many enterprises still rely on standard connectivity, which operates on a philosophy known as best-effort delivery**.** But often this actually becomes a bottleneck. Enterprises integrating agentic AI in their workflows must evaluate the limitations of this approach and why a better alternative – in the form of Guaranteed Throughput with Enterprise-Grade Connectivity – is needed. We suggest you make your assessments based on these questions:

a) **Consistency: Multi-Tenant Contention or Dedicated Capacity?**

Best-effort delivery relies on a shared infrastructure where performance is dictated by regional traffic patterns. For an enterprise, this introduces contention ratios—which means your bandwidth is not exclusively yours even during your peak business hours. Let's say, nearby businesses increase their load at the same time. This causes your throughput to drop, leading to unpredictable application performance.

In contrast, Guaranteed Throughput provides a 1:1 contention ratio, meaning your bandwidth is physically or logically set aside exclusively for your organisational needs. As a result, your data-intensive AI operations maintain a steady performance regardless of external network demand.

b) **Why Stabilise the AI Reasoning Loop?**

For enterprises deploying Agentic AI and Large Language Models, latency jitter is a complete no-no as it can cause significant technical risk.

But best-effort, basic connections often suffer from variable packet delivery times. The resulting jitter can disrupt AI reasoning loops, causing delays in API responses and bring down the user experience for customers or employees.

When an AI agent is in the middle of its reasoning stage, a millisecond of jitter isn't just a glitch—it becomes a broken thought process. Imagine a customer support scenario in which this minor hiccup could result in a frustrating interaction for the customer or even a failed transaction. The outcome can be quite adverse including potential loss of business opportunity.

In contrast, enterprise-grade connectivity focuses on direct, optimised routes to ensure it is consistently fast and lag-free, which is the basis of ultra-low, predictable latency. This stability is exactly what "thinking" models need as they demand high-frequency, two-way data exchange to stay coherent and responsive in their interactions.

c) Why Prioritise Traffic Optimisation?

If you're on a standard broadband link, your mission-critical AI API call is treated with the same urgency as a background email synchronisation. There is no measure used to protect or prioritise high-value packets. During periods of congestion, your most important data is just as likely to be dropped or delayed as a non-work-related, low-value data from ad-heavy web traffic.

In the Guaranteed Throughput approach, your IT team has the opportunity to define policies that prioritise critical AI applications, making sure these move to the front of the queue. This mechanism ensures that even if the network is under heavy load, your revenue-generating AI services remain unhampered by secondary traffic.

d) Business Continuity or Best Effort?

The best effort approach doesn't factor in accountability or reliability, which creates an important risk. The provider is not financially or legally bound to maintain specific performance levels. If the connection falters, your AI-driven operations simply stop, and the recovery timeline is at the provider's discretion.

This is why enterprise-grade services are a superior alternative as they are governed by stringent SLAs. This includes a contractual assurance of uptime (e.g., 99.99%) and specific performance metrics (like latency, packet loss, and jitter).

For an enterprise, these changes elevate the role of connectivity. From an IT overhead to be avoided or minimised, it can become a reliable asset wherein the provider is expected to meet stringent standards of performance and positively impact AI effectiveness.

Why Connectivity Needs Intelligence

In the past, the network was just a pipe without any abilities to think. It didn't matter if the pipe was carrying an email or a heart rate monitor's data. Today, the scenario has drastically altered as we have entered the era of application-aware networking**.**

In this networking approach, the main idea is to ensure critical applications receive essential resources before any other app. The network sorts, prioritises application traffic and determines the specific needs of resources like bandwidth, latency etc. Through regular monitoring of application needs, its settings can be adjusted accordingly. For instance, it would prioritise bandwidth for a client video interaction over casual browsing or data transfer.

This capability can become the backbone of agentic AI's effectiveness. Unlike chat-bots, AI agents thrive on real-time context of data rather than raw content. Therefore, for Agentic AI to function and handle core workflows, the network must be intelligent enough to recognise the intent of the data quickly.

This is where the difference between "being fast" and "being smart" becomes critical.

Modern connectivity solutions, like ACT's **Corporate Broadband with SmartWiFi**, help make this paradigm shift from reactive to proactive. Instead of waiting for a collision to happen, network solutions integrated with AI capabilities can help:

- **Identify Congestion:** Detect when a heavy video call or a massive AI data sync is about to start.
- **Enable Proactive Clearing:** intelligent access points designed to automatically shift frequencies or prioritise traffic before the lag occurs.
- **Provide Adaptable Signal Settings:** Constantly optimise the signal to ensure the AI agent has a clear, high-speed lane leading to the cloud.

The Network Infrastructure Gap: A Wake-Up Call for Indian Enterprises

India is home to some of the most digitally advanced businesses in the world, yet many are building their AI houses on brittle foundations. What is the point of investing millions in globally acclaimed SaaS apps if you end up running them on consumer-grade backbones?

This mismatch creates a fragile enterprise. You may have the most sophisticated AI agent in the world, but if its connectivity is inconsistent, the agent's performance will be mediocre.

On the other hand, if you want your AI agent to operate as desired in real time, you need an enterprise-grade network ecosystem that adapts in real time, too. In the age of autonomous agents, network intelligence is the real asset that will help you differentiate yourself. And your competitive advantage can start right at the access point.

If you are currently deploying AI agents, it is time to audit your connectivity and understand if it's ready for the agentic era. It is time to ask: Are we settling for best-effort, or are we enabling our AI with the guaranteed throughput it needs to actually perform?