

WORLD'S TOP AI MEDIA POWERHOUSE



MAY 2026 Edition



**India is Not
Ready for
Mythos**

**Amaravati
Makes a
Quantum Leap**


Amaravati's new open-access testbed aims to help India design, test, and scale quantum hardware. It serves as a prototype for a broader network of quantum testbeds across India.

**The Last Mile of AI Runs
Through Tredence**

The future will not be defined by who has access to AI. It will be defined by who can make AI work, at scale, and in context.



THE COMPANY TURNING AI CHAOS INTO ENTERPRISE OPPORTUNITY



From left to right: Shub Bhowmick (CEO), Sumit Mehra (CTO) and Shashank Dubey (CRO)

Every industry reaches a point at which the assumptions on which it was built begin to fail. For global IT services, that moment arrived in early 2026.

For decades, the industry thrived on scale, talent, and cost efficiency. Then came systems that could think, act, and execute independently. The shift had been building in the background. In early February, it became impossible to ignore.

The trigger came from Anthropic.

The company released a suite of autonomous AI agents capable of executing end-to-end workflows across legal review, financial analysis, software development, compliance, and data research—work that once required entire teams of specialised professionals.

The market's reaction was immediate. The Nifty IT Index fell nearly 6% in a single session—its steepest drop

since the COVID-19 pandemic, erasing billions in market value. Investors were not reacting to a product launch alone. They were reacting to what it implied for the economics of the entire outsourcing industry.

India's \$300 billion IT industry, built on the proposition of cost arbitrage through cheaper skilled talent, was suddenly facing a threat from computer programmes promising faster, cheaper, and more efficient execution.

"The core problem is structural. AI is bringing a seismic shift in tech services by reducing project timelines and automating tasks, putting direct pressure on the headcount-based outsourcing model. What a team of 500 engineers once delivered across 18 months, clients are now asking whether 50 engineers and a fleet of AI agents can deliver in six. The answer, increasingly, is yes," says Pavan Nanjundaiah, Head of Tredence's AI studio—the company's R&D division.

This disruption to the delivery model is also dismantling the commercial logic of IT services itself. For years, enterprise technology engagements were priced around inputs—headcount, hours, infrastructure. Now clients increasingly care about measurable outputs like customer satisfaction scores, inventory costs, and revenue impact. When execution becomes faster and cheaper, effort becomes a less significant differentiator. Impact becomes the real measure of value.

Yet, even as AI systems dramatically improve productivity across writing, coding, analytics, and customer service, embedding them into real enterprise environments remains uneven. Governance gaps, legacy systems, and organisational inertia continue to slow adoption.

Analysts believe this tension between AI capability and enterprise reality will determine the next phase of the AI



AI is bringing a seismic shift in tech services by reducing project timelines and automating tasks, putting direct pressure on the headcount-based outsourcing model. What a team of 500 engineers once delivered across 18 months, clients are now asking whether 50 engineers and a fleet of AI agents can deliver in six. The answer, increasingly, is yes.

Pavan Nanjundaiah,
Head of Tredence's AI Studio



economy. What is disappearing is not the need for tech services altogether, but a specific layer within it—the layer responsible for translating instructions into execution.

What becomes more valuable instead is everything around that layer: context, integration, governance, domain expertise, and the ability to make AI systems function reliably inside real businesses.

That is the space Tredence has spent the last 13 years preparing for. The \$300 million data and AI firm has quietly built its entire business model around turning insights, models, and copilots into production-grade decisions and measurable business outcomes.

For Tredence, it was never about the headcount pyramid or billable-hours model. Instead, it invested in domain depth, proprietary accelerators, technical alliances with hyperscalers and

AI labs, and long-term relationships with Fortune 500 clients who need someone to make their AI actually work.

The Racetrack Builders

By early March, Hyderabad has already stopped pretending it is spring. At 40 degrees Celsius, the city's famous summer has arrived ahead of schedule, shimmering off glass buildings and emptying footpaths by noon.

Inside Tredence's new office, Co-Founder and CRO, Shashank Dubey, is at an all-hands, taking questions. The heat outside is nothing compared to what employees keep circling back to: will Tredence feel the burn of this new wave of AI agents? The coding tools, the co-working platforms, the systems that seem to be eating through the IT industry from the bottom up?

Dubey does not flinch.

Large language models, he says, are like Formula One cars. Extraordinarily powerful. But a Formula One car cannot drive on a normal road. It needs purpose-built infrastructure, precise maintenance, the right circuit, and the right governance around it. A Ferrari costs between \$15 and \$20 million to build. The racetrack it runs on costs \$200 million.

"The more Ferraris get made, the more racetracks need to get built," Dubey argues.

What he is describing is the enterprise AI deployment gap—the yawning chasm between what a foundation model can do in a controlled demonstration and what a large retailer or healthcare provider can actually use on Monday morning. Between a model that works in a lab and AI that actually changes how a business operates lies the graveyard of most enterprise AI initiatives.

To be clear, Tredence is not your traditional tech services firm. While others build the AI stack from the bottom up, Tredence works top-down, capturing value where it matters most: the decision layer.

In NVIDIA CEO Jensen Huang's five-layer AI cake, the decision layer is an important sub-layer that sits between models and applications. It is where intelligence becomes context-aware action, where workflows are transformed, and where measurable outcomes are realised. Most firms stop at building the stack. Tredence monetises it.

Through its last-mile approach, Tredence delivers value at this layer and monetises the AI stack top-down.

A pricing algorithm, for example, is only as useful as the data it ingests, the constraints it understands, and the decisions it is allowed to influence. A supply chain model must account for local realities, logistics infrastructure, regulatory environments, and even weather patterns that generic systems often miss.

The Domain Moat

If the first wave of AI rewarded technical capability, the current one rewards domain knowledge. Large models are trained on vast amounts of generalised data. They are remarkably good at abstraction. What they lack is context.

"A supply chain optimisation model built for a US retailer does not work for a company in the APAC market with different cold-chain infrastructure and different supplier dynamics. Knowing that difference, and understanding the granular, operational detail of how a specific industry actually works, is the kind of domain knowledge that is Tredence's USP," says Shub Bhowmick, the Co-Founder and CEO.

Having worked with Fortune 500 companies across verticals, Tredence has built something that foundation models cannot replicate on their own: a library of industry-specific context, accumulated through engagement and iterations, and refined into frameworks.

The domain depth is further strengthened by practitioners with operational expertise across industries

“

Large language models are like Formula One cars. Extraordinarily powerful. But a Formula One car cannot drive on a normal road. It needs purpose-built infrastructure, precise maintenance, the right circuit, and the right governance around it. The more Ferraris get made, the more racetracks need to get built.

Shashank Dubey,
Co-Founder and CRO at
Tredence



such as retail, CPG, and healthcare. For Tredence, the expertise in the room is as important as the technology running on the servers.

Beyond vertical expertise, Tredence has spent years building what it calls an accelerator ecosystem—over 100 proprietary tools deeply integrated with Google Cloud, Databricks, and Snowflake.

Think of accelerators as pre-built, already proven, components—data connectors, industry-specific models, governance frameworks—that eliminate the need to rebuild common infrastructure from scratch on every engagement. Where a traditional implementation might spend the first six months simply laying foundational plumbing, Tredence's accelerators compress that work into weeks, letting teams get to the actual business problem faster.

But compressing deployment time only matters if the data foundation underneath it is solid. In most enterprises, that foundation remains a major issue.

"While data may mean different things to different people, and those ambiguities

corrupt everything downstream, Tredence starts by building what it calls an ontology: a precise, industry-specific business dictionary that creates a shared language before any AI is deployed. From there, solutions are co-developed in a live sandbox, with business teams in the room and real data in play," says Arvind Ramachandran, Director, Data Engineering at Tredence.

Ramachandran says that results are visible immediately; tweaks happen on the spot, and the distance between concept and production collapses from months to weeks.

Punching Above Its Weight

For a firm making bold claims about where the industry is headed, the financials are the most persuasive argument. Tredence's revenue has grown 8x in five years, at a clip exceeding 40% annually.

Although retail and CPG have been the biggest growth drivers so far, CEO Shub Bhowmick says he is seeing increasing momentum in healthcare and life

sciences, as well as travel and hospitality.

A case in point is a leading multinational pharma company that relied on Tredence's experience in breaking down data silos and used predictive maintenance to spot issues early, keeping equipment running smoothly. The result is 15% more time with machines running, 12% increase in output, and 7% lower maintenance costs, improving both efficiency and reliability.

This laser-sharp focus on outcomes and impact is doing something else too: opening doors that were previously closed to a firm of Tredence's size. Bhowmick calls it the "David and Goliath moment"—Tredence trumping some of the large, established IT consulting majors to win deals. Last year, the company won multiple deals in the \$10-15 million range, ticket sizes that have historically been the exclusive territory of Accenture, Deloitte, and global system integrators.

Human in The Loop

Large deals are one problem solved. The harder one is finding the people to deliver them.



In this uber-competitive space, we are not just competing for the right talent, we are also competing for potential. The dynamic nature of this market has made us reconsider what really makes a great AI professional. It's good to have people who can code or crunch numbers, but we are also attaching a lot of importance to people who can think, connect, and solve problems humans actually have.

Rekha Nair,

Chief HR Officer at Tredence



“

For their models to create enterprise value, they need partners who bring domain expertise. The relationship is symbiotic, not predatory. Anthropic and OpenAI need us as much to get deeper into businesses.

Sumit Mehra,
Co-Founder and CTO
of Tredence



Even as AI displaces knowledge workers at scale, over 1.6 million AI-specific positions sit unfilled globally against only 5,18,000 qualified candidates. The conventional response would be to poach. Tredence made a different bet.

In 2020, it launched TALL—the Tredence Academy of Leadership and Learning—built on the premise that AI-ready talent simply cannot be acquired from a shrinking external pool; it must be grown from within.

The programme runs foundational tracks for fresh graduates, lateral programmes in data science and engineering, and vertical-specific streams in retail, CPG, and healthcare—industries where understanding the business problem is as important as knowing the tooling.

“In this uber-competitive space, we are not just competing for the right talent, we are also competing for potential. The dynamic nature of this market has made us reconsider what really makes a great AI professional. It’s good to have people who can code or crunch numbers, but we are also attaching a lot of importance to people who can think, connect, and solve problems humans actually have,” says Rekha Nair, the Chief HR Officer at

Tredence.

After all, technical skills can be taught. What remains irreplaceable, even with AI, are the human qualities: curiosity, empathy, and the drive to solve real problems. Those are the things that can’t be learned from a certification. The agents may be writing the code. But someone still has to decide what the code should do.

Today, Tredence is 4,000+ people strong, spread across 19 locations—a truly global resource base that brings together industry expertise, engineering depth, and client proximity to across markets.

What Survives

The doomsday narrative for tech services is seductive, but probably overstated.

There is also the question of what the foundation model companies actually need. Anthropic and OpenAI can build increasingly capable models, but they cannot, and do not, have deep expertise in pharmaceutical manufacturing, logistics or fast-fashion retail replenishment cycles.

“For their models to create enterprise value, they need partners who bring

domain expertise. The relationship is symbiotic, not predatory. Anthropic and OpenAI need us as much to get deeper into businesses,” says Sumit Mehra, Co-Founder and CTO of Tredence.

The firms that will survive this transition are those that understand what AI genuinely cannot replace: domain fluency to know what questions to ask, credibility to sit in a boardroom and explain the answer, and engineering rigour to make the whole system run reliably at scale. That combination lives at the intersection of human expertise and machine capability, which is exactly where Tredence has spent the last 13 years positioning itself.

What seems increasingly clear is that the future will not be defined by who has access to AI. That access is becoming ubiquitous.

It will be defined by who can make AI work, at scale, and in context.

Disruptions like Anthropic’s will keep arriving, sharp and unexpected. But so will the optimism that builds companies in the first place. And in the space between AI’s promise and real-world outcomes, there is still plenty of road left to lay. 🚀