



Holiday
Offer

Don't miss out! :-)

[AI](#) [APPS](#) [BIG TECH](#) [CLOUD](#) [CYBERSECURITY](#) [GOVERNANCE](#)

[HARDWARE](#) [SOFTWARE](#) [STARTUPS & SMES](#) [TELECOM](#)

[Home](#) [Expert Opinion](#) [India's Next Enterprise AI Challenge Will Be Operational, Not Technical](#)

Sep. 25, 2026 at 5:38 pm

EXPERT OPINION

India's Next Enterprise AI Challenge Will Be Operational, Not Technical

 **CXOtoday News Desk** 4 days ago



By Shomiron Das Gupta

The dominant assumption inside boardrooms is that AI adoption is still a technology race pick the right model, secure the right infrastructure, win. That assumption is now outdated.

Most large enterprises have already crossed the technology threshold. They have access to capable models, elastic cloud infrastructure, and a deep bench of AI-literate talent. What they lack is the ability to run AI consistently across functions, geographies, and teams.

Recent industry data shows that nearly 25% of technology services companies have moved AI experiments into production. That number is telling for what it implies about the other 75%. The constraint is not capability. It is execution.

The Technology Problem Is Largely Solved

Five years ago, access to advanced AI was a genuine differentiator. Building or licensing a capable model required capital, talent, and infrastructure that few organisations could assemble quickly.

That barrier has collapsed.

Cloud providers now offer AI-ready infrastructure as a default, not a premium tier. Foundation models are available on demand, priced by consumption rather than ownership. Integration tooling has matured to the point where connecting AI into existing enterprise software is a configuration exercise, not a multi-year engineering programme.

The result is a level playing field. Any competitor, in any sector, can now access broadly similar AI within months, sometimes weeks.

That changes where advantage actually sits.

It no longer comes from possessing AI. It comes from embedding it into daily work reliably, at scale, without breaking the processes that already function. A model that performs brilliantly in a proof of concept and then stalls at the handoff to a business unit delivers nothing.

That lesson is nowhere more visible than inside the security operations centre.

The SOC Is Where the Bottleneck Shows Up First

Most large enterprises have already layered AI-driven detection into their security stack. Anomaly detection, log correlation, and automated triage are no longer novel. Yet SOC teams report little relief from alert volume. The reason has less to do with model accuracy than with what happens after a system flags something.

A tool that surfaces a thousand anomalies a day is only useful if the workflow around it can absorb that output. Escalation paths were built for manual triage and have rarely been redesigned around machine-generated signal. Analysts still spend most of their time deciding what deserves attention rather than acting on confirmed threats, because the link between detection and response was never rebuilt.

Ownership gaps make this worse. Few organisations have settled who is accountable when a system recommends isolating a machine, blocking an account, or cutting off access — the analyst, the CISO's office, or the business unit disrupted by the action. That ambiguity becomes a liability the first time an automated response causes damage of its own, rather than preventing it.

Collaboration between analysts and AI remains especially undefined. Teams are often expected to trust prioritised alerts without clear rules for when to override, when to escalate further up, and when a false positive should feed back into retraining rather than simply being dismissed. Without that loop, the same noise resurfaces indefinitely.

A 2026 survey of technology executives found that only 17% of organisations have actually deployed AI agents to date, even though more than 60% expect to within two years — the largest gap between intent and deployment among all emerging technologies measured. Security operations sit squarely inside that gap: the appetite for automated detection and response is high, but the scaffolding to run it safely, with clear ownership and a verifiable audit trail, is still being built.

What works in one part of the SOC rarely transfers by simple replication either. A triage model tuned for cloud infrastructure alerts does not carry over cleanly to identity threats or insider risk, because the underlying data quality, escalation authority, and tooling differ across each domain.

What Comes Next

The enterprises that pull ahead over the next three years will not be the ones running the most sophisticated detection models. They will be the ones with the steadiest operations.

That means defining accountability for automated security decisions before an incident forces the question, not after. It means treating AI-driven SOC workflows with the same rigour as any other major process change. And it means building the capacity to keep functioning when a system misfires or is deliberately targeted.

Boards and regulators have already changed their line of questioning. They are no longer asking whether security teams use AI. They are asking how it operates, who answers for its decisions, and what happens when it fails and that scrutiny will only sharpen.

Clear sightlines into how AI performs across the enterprise, not just within one team, will soon be an expectation set from the top, not a technical dashboard buried in an ops review.

The unglamorous work of aligning functions, incentives, and ownership will decide who actually captures the value AI makes possible.

The next phase of enterprise AI will not be won by whoever adopts the sharpest technology.

It will be won by whoever runs it best.



(The author is Shomiron Das Gupta, Founder & CEO, Bloo.io, and the views expressed in this article are his own)

The CXOtoday App Is Here.

News, briefings and analysis for enterprise leaders
— now in your pocket.



Download Now on
Google Play

Coming Soon to
App Store

CXO
today