

# What Happens When CXOs Speak Honestly About AI Adoption

Pune has seen every major technology wave up close.

From large ERP transformations to cloud migrations, from automation to analytics, the ecosystem here has built, delivered, and scaled systems for some of the world's most demanding enterprises. This is a city shaped less by hype and more by execution.

So when a room full of CXOs gathered for the SEAP roundtable titled **“AI Is Scaling Fast, Outcomes Are Not,”** there was very little appetite for buzzwords.

The opening sentiment was blunt:

**“We’re doing more AI than ever before and still struggling to see outcomes that actually stick.”**

That statement didn't come from skeptics.

It came from leaders already investing in AI across product companies, services firms, regulated industries, and enterprise operations.

What followed was not a discussion about models or tools.

It was a candid comparison of lived experience: where AI stalled, where it quietly worked, and why the gap between excitement and execution keeps widening.

## When Technology Moves Faster Than Readiness

One leader shared a story that resonated immediately with the room.

Years ago, their organization undertook a large-scale **ERP transformation**. The promise was sweeping change: efficiency, visibility, and control. What landed at the ground level was far less dramatic: *new screens, same old work*.

The technology wasn't the problem.

The strategy never made it into the workflow.

As the discussion unfolded, it became clear why this story mattered. Many leaders have seen the same pattern repeat with AI today. Ambition is set at

the top. Pilots are built in the middle. But at the edges—where people actually work—behavior doesn't change.

The group agreed that AI adoption is not stalling because the technology is weak.

It's stalling because organizations are not ready in the ways that matter.

*"We didn't fail because the model was wrong. We failed because five teams had five different definitions of success."*

— Product leader, enterprise software company

Readiness, as discussed in the room, wasn't about tooling. It showed up as alignment across **people, process, and data**, and most failures traced back to cracks between those layers.

### **When Measurement Stops Reflecting Reality**

As the conversation shifted to ROI, a quieter tension surfaced.

Everyone agreed that AI must be measured. But many admitted that measurement itself had become part of the problem. Several leaders described a familiar anti-pattern: new KPIs created to justify AI investments, rather than to understand real impact.

Productivity numbers looked impressive on slides, but didn't translate into meaningful operational change. Teams were told to be "more productive," without clarity on what would change—or how success would be recognized.

*"We created new KPIs to justify AI. Six months later, we couldn't explain whether anything had actually improved."*

— Engineering leader, services firm

The group converged on a sharper principle: **if AI outcomes can't be measured using metrics the business already trusts, the ROI is questionable by definition.** Real progress came from grounding AI impact in existing productivity, quality, and operational benchmarks and being honest when those metrics didn't move.

### **Why Calling Everything "AI" Erodes Trust**

One of the most animated discussions in the room centered on a deceptively simple question: *what do we actually mean by AI?*

Some participants used AI as a catch-all for automation. Others drew a clear distinction between deterministic systems, where outputs are predictable, and probabilistic systems, where outcomes must be interpreted.

The group didn't chase a perfect definition. But they did agree on something practical:

### **Definitions act like contracts.**

They set expectations for certainty, risk, and accountability.

When everything is labeled AI, certainty is oversold. And when certainty breaks, trust drops sharply.

*"Calling everything AI made people expect certainty. When outcomes varied, trust collapsed instantly."*

— Analytics leader, regulated industry

This is where the "AI gold rush" shows its cost—not in failed pilots, but in lost confidence.

### **The Moment That Changed the Room**

The most powerful moment of the session didn't come from a framework or a statistic.

It came from a story.

A participant described a real-world scenario in which an AI system and a human expert reached different conclusions in a high-stakes decision. The AI output prompted a second review. The outcome changed.

No one framed it as "AI versus human."

Instead, the room leaned into a deeper question:

### **How do we design systems where trust is earned, not assumed?**

The answers were practical:

- data quality as a non-negotiable foundation
- explainability that helps users understand *why*, not just *what*
- clear guardrails and escalation paths
- human-in-the-loop workflows as a deliberate adoption strategy

*“Trust wasn’t about accuracy. It was about knowing when the system would stop and ask for help.”*

— Operations leader, high-stakes domain

Trust, the group realized, isn’t built in policy documents. It’s built into everyday workflows.

### **AI Doesn’t Deploy. It Diffuses.**

As the discussion moved toward change management, parallels with past technology waves became impossible to ignore.

AI adoption behaves less like a rollout and more like diffusion. It spreads unevenly. It accelerates where workflows adapt. It stalls where behavior doesn’t change.

Leaders emphasized the importance of:

- giving ground-level users time to adapt
- Investing in training and feedback loops
- aligning terminology across business and technical teams
- designing user experience as carefully as the model itself

One sentiment captured the shift well:

**“AI becomes real only when it changes how people work, not when it’s installed.”**

### **Ending With the Hard Question**

The session didn’t close with predictions about the future of AI. It ended with a question about discipline:

**If you had to cut 50% of your AI initiatives tomorrow, what would you keep and why?**

The room went quiet. And that silence said more than any roadmap could.

*“The hardest decision isn’t funding AI. It’s deciding which AI projects to shut down.”*

— CXO, mid-to-large enterprise

## From AI Excitement to AI Results

The SEAP CXO Roundtable didn't produce a checklist or a playbook. What it produced was shared clarity.

AI outcomes don't scale with enthusiasm.

They scale with **readiness, trust, measurement, and change**.

AI is scaling fast.

Now organizations must learn how to scale confidence with it.

SEAP will continue this conversation through follow-up sessions focused on practical action, experience sharing, and collective learning across the ecosystem.

## About the Moderator

**Rohit Virmani** is a senior product leader in enterprise SaaS, working at the intersection of AI, product strategy, and large-scale adoption. He focuses on translating AI capability into measurable business outcomes, particularly in complex, regulated, and multi-stakeholder environments, where execution, trust, and change management matter more than hype.