dynatrace | CIO&LEADER | aws

FUTURESCAPE 2025

The Intelligent Enterprise Playbook

Driving Business Outcomes through AI, Observability, and Automation

dynatrace | CIO&LEADER | aws

FUTURESCAPE 2025

The Intelligent Enterprise Playbook

Driving Business Outcomes through AI, Observability, and Automation

FUTURESCAPE 2025

PUBLISHER & EDITORIAL DIRECTOR
Vikas Gupta

COO & ASSOCIATE PUBLISHER
Sachin Nandkishor Mhashilkar

EDITORS
R. Giridhar, Jatinder Singh

RESEARCH & WRITING
Deepak Kumar, Balaka Baruah Agarwal, Musharrat Shahin, Jagrati Rakheja

PROJECT
Vandana Chauhan, Rajiv Pathak, Jagdish Bhainsora, Dipanjan Mitra, Vaishali Banerjee, Reetu Pande, Snehal Thosar, Shabana Shariff

PARTNERSHIPS
Hafeez Shaikh, Sourabh Dixit

OPERATIONS
Neelam Adhangale, MP Singh, Amit Singh, Satish Chaudhari, Sampath Kumar, Himanshu Kumar, Dipti Gamre

VIDEO
Sunil Kumar

DESIGN
Shokeen Saifi , Manish Kumar

FUTURESCAPE 2025
The Intelligent Enterprise Playbook

PUBLISHED AND PRINTED BY
9.9 Group Pvt Ltd
121, Patparganj, Mayur Vihar
Phase 1, New Delhi-110 091

ISBN
9789361944307

CONTENT

A Collective Vision for Intelligent Transformation

FUTURESCAPE 2025 is more than a book—it is a strategic blueprint for technology leaders navigating the next wave of intelligent digital transformation. Conceived for CIOs and senior IT leaders in India's large and mid-sized enterprises, it distills the hard-earned lessons and forward-looking strategies of those who have been shaping the future of enterprise technology.

This initiative, led by CIO&Leader—the flagship B2B technology platform of the 9.9 Group—in collaboration with Dynatrace, a global leader in AI-powered observability and automation, brings together the perspectives of 35 senior technology leaders across industries. The conversations that inform this work were not about abstract trends. They were about execution: how to align observability, automation, DevOps maturity, and AI-driven operations directly with measurable business outcomes.

How FUTURESCAPE 2025 Was Built

Our process was designed to mirror the rigor of strategic planning. We began by identifying the core domains where enterprises are investing—or will invest significantly—in the near term:

- Observability and AI-powered insights in live production environments

- Scaling DevOps maturity to balance speed with stability
- Platformization and automation to deliver IT that is directly aligned with business priorities
- Linking digital experience to quantifiable business outcomes

For each area, select CIO advisory groups engaged in structured, multi-stage discussions. These were not passive panels—they were working sessions that:

- **Explored the landscape:** Clarifying drivers, emerging use cases, and practical considerations.
- **Addressed challenges:** Sharing peer-to-peer solutions and implementation frameworks.
- **Prioritized insights:** Distilling the most relevant, high-impact takeaways for Indian enterprises.

Every participant contributed something essential—whether a complete transformation journey, a hard-won lesson, or an innovative approach to a stubborn challenge. Dynatrace leaders complemented these discussions with global and India-specific insights drawn from real-world deployments.

How to Read This Book

FUTURESCAPE 2025 is designed to be both a strategic compass and an operational playbook.

- If you are exploring a domain for the first time, our lead articles offer a business-relevant grounding in the concepts and their enterprise implications.
- If you are a seasoned practitioner, you will find benchmarks, frameworks, and peer experiences that challenge your assumptions and expand your toolkit.

Each section follows a consistent structure:

- **Lead Article –** distilled from collective discussions, capturing the shared intelligence of the group.
- **CIO Contributions –** direct accounts from practitioners, covering challenges, successes, and lessons learned.
- **Dynatrace Perspective –** best practices and insights from global enterprise implementations.

Acknowledgments

Initiatives of this scale require both rigor and collaboration. We are grateful to the CIOs who contributed their time and expertise; to Dynatrace's leadership for bringing global perspective with local relevance; and to the 9.9 Group editorial team for facilitating, synthesizing, and shaping these discussions into the work you now hold.

As Indian enterprises accelerate toward AI-enabled, adaptive, and business-aligned digital operations, FUTURESCAPE 2025 offers a roadmap for leaders determined not just to respond to the future, but to shape it. We hope you find it as actionable, relevant, and inspiring as it was for us to create.

"FUTURESCAPE 2025 is more than a book—it is a strategic blueprint for technology leaders navigating the next wave of intelligent digital transformation."

R. Giridhar
Group Editor
9.9 Group

Jatinder Singh
Executive Editor
CIO&Leader

The Intelligent Enterprise Playbook

As organizations embrace cloud-native architectures, data overload threatens progress. To thrive, they must become intelligent enterprises—transforming data into real-time, intelligent action.

The digital landscape across the Asia Pacific and Japan (APJ), especially India, is undergoing rapid transformation. India is leading the charge in digital adoption with innovations like Unified Payments Interface (UPI), which recorded 12 billion monthly transactions by June 2025, and widespread cloud adoption. With over 400 million active Fintech users and a nationwide 5G rollout expected by year-end, India is becoming a global testbed for mobile commerce, streaming, and IoT.

The startup ecosystem is booming too—with over 115 unicorns—driving artificial intelligence (AI) and software-as-a-service (SaaS) innovation across sectors. But as organizations shift to cloud-native architectures, they encounter growing complexity and data overload. Without effective tools, this data becomes a burden, slowing progress.

To thrive, businesses must become intelligent enterprises, turning data into intelligent, real-time actions. Achieving this shift depends on the combined power of AI, observability, and autonomous intelligence. Observability is no longer just an IT concern—it's a business-critical function. It connects system behavior to outcomes like customer experience, compliance, and revenue.

The third-generation Dynatrace platform is built to leverage the significant value of observability data in context, feeding AI models that generate timely, actionable insights. Dynatrace Davis® AI, the purpose-built foundation for agentic AI, delivers deterministic root-cause analysis, auto-remediation, and natural language insights. The Dynatrace Grail data lakehouse unifies data at a petabyte scale, making faster, more accurate decisions possible. Dynatrace is the leading AI-powered observability platform, enabling meaningful outcomes for enterprises across APJ and India.

Dynatrace collaborates with Amazon Web Services to help customers innovate confidently with AI and Agentic solutions, reduce operational costs, improve user experience, and cut mean time to resolution (MTTR). In India, Dynatrace is making strategic investments to deepen its presence and support enterprises with integrated observability, security, and business analytics, empowering them to navigate complexity and scale AI with confidence.

"The Intelligent Enterprise Playbook" offers a strategic roadmap for tech leaders embracing this transformation. It is a call to action to embed AI, automation, and observability into your operations, building a resilient, autonomous enterprise where intelligence is built in. By doing so, you'll unlock new levels of agility, innovation, and sustainable growth. The future of the intelligent enterprise is already here — and Dynatrace is here to help you seize it.

"A key fallout of cloud-native architectures is growing complexity and data overload. To thrive, businesses must turn data into intelligent, real-time actions—by combining the power of AI, observability, and autonomous intelligence."

Arun Balasubramanian
Managing Director, India & SAARC,
Dynatrace

Rafi Katanasho
APJ Chief Technical Officer and
VP - Solution Engineering, Dynatrace

AI Applications: From Pilots to Enterprise-wide Transformation

How enterprises are scaling AI from isolated initiatives to trusted, business-aligned systems embedded across workflows and decision-making.

Executive Summary

As enterprises across industries embrace digital transformation, AI is moving from experimentation to execution. What was once confined to innovation labs and pilot projects is now showing up in production workflows, customer journeys, risk models, compliance frameworks, and employee productivity platforms. Yet while the ambition is high, the path to scale remains complex—and uneven.

This chapter explores how leading CIOs and digital leaders are translating AI promise into enterprise value. It reflects a shift in posture: from opportunistic deployments to purposeful architecture. From isolated use cases to platform strategies. From proof-of-concept wins to cross-functional alignment.

The conversations across CIO panels revealed a consistent insight: AI is no longer an edge initiative—it is becoming a core capability. But for that transition to succeed, CIOs must lead with clarity, governance, and a commitment to business alignment. Technical talent and tooling are necessary but not sufficient. The organizations getting AI right are those treating it as a systems-level transformation, not just a model or algorithm.

One of the biggest unlocks is the maturity of foundational technologies. Cloud-native infrastructure, scalable data platforms, and enterprise-ready ML ops toolchains have made it possible to train, test, and deploy AI models faster—and more securely—than ever before. At the same time, new paradigms such as generative AI and large language models (LLMs) are creating new demand vectors across customer service, knowledge management, and decision support.

However, challenges remain. Many enterprises still struggle with fragmented data landscapes, legacy application bottlenecks, and under-prepared governance frameworks. Business teams often lack the interpretability and confidence to act on AI recommendations. And while automation potential is high, the change management required to realize that potential is often underestimated.

CIOs are responding by focusing on value pathways rather than technical proofs. This means designing AI initiatives that are measurable, explainable, and aligned to real business outcomes. It also means engaging stakeholders early, educating leadership on model limitations, and embedding AI into the workflows and platforms that people already use.

The role of AI is also shifting—from being a tool for prediction or classification to becoming a co-pilot in everyday decision-making. Whether in finance, HR, supply chain, or customer service, enterprises are beginning to reimagine how work is performed—with AI not replacing humans, but augmenting their capabilities. The rise of agentic AI—systems that can take action, not just provide answers—is accelerating this trend.

Security and compliance are taking center stage, especially in regulated sectors. Leaders are embedding explainability, fairness checks, and model observability into their MLOps pipelines to ensure that AI is not just fast, but trustworthy. This is especially critical as AI moves closer to customer-facing or compliance-sensitive domains.

Finally, the discussion is moving from tools to operating models. Enterprises are building AI Centers of Excellence (CoEs), forming cross-functional squads, and investing in data literacy and AI fluency across departments. CIOs are rethinking talent models—not just hiring data scientists, but reskilling analysts, product owners, and frontline managers to work in AI-infused environments.

This chapter captures the evolving mindset of enterprise AI: from a collection of use cases to an operating layer. From narrow efficiency gains to systemic transformation. And from experimentation to trust-based scale.

"AI is no longer an edge initiative—it is becoming a core capability. The organizations getting AI right are those treating it as a systems-level transformation, not just a model or algorithm."

The organizations leading the charge aren't necessarily the ones with the most models deployed—they're the ones that understand what's worth modeling, how to operationalize it, and how to do so in a way that creates sustainable business value.

Key Drivers: What's Fueling Strategic AI Adoption

Enterprise adoption of AI is no longer driven by novelty—it's driven by necessity. As digital ecosystems expand and business environments grow more complex, AI offers the analytical, predictive, and decision-making horsepower that traditional systems cannot match. The following drivers are propelling AI into the center of enterprise strategy:

Business Model Disruption Demands Intelligence at Speed
Industries are being reshaped by digital-first competitors, regulatory shifts, and rising customer expectations. In this context, AI enables faster response cycles—whether through automated decisions, predictive insights, or intelligent customer engagement. Enterprises are turning to AI to spot trends early, respond in real time, and personalize at scale.

Data Maturity Has Reached a Tipping Point
Many organizations have spent the last decade building data lakes, integrating systems, and modernizing infrastructure. That groundwork is now paying off. With cleaner, more accessible, and better-governed data, AI can move from theory to production, accelerating adoption across functions like finance, risk, supply chain, and customer service.

Cloud-Native and SaaS Architectures Are Lowering the Barrier to Entry
AI no longer requires specialized infrastructure. With cloud-native ML platforms, pre-trained models, and API-first AI services, teams can experiment and deploy quickly. This accessibility is fueling a wave of domain-specific applications—from fraud detection to sentiment analysis—without deep in-house AI engineering.

Generative AI Is Opening New Frontiers
The advent of large language models and generative AI has shifted the conversation. These technologies are unlocking use cases in content creation, summarization, conversational interfaces, and knowledge retrieval, enabling automation of tasks previously considered too complex for traditional systems.

Workforce Augmentation Is Becoming a Strategic Priority
Enterprises are looking to AI to amplify human capability, not just replace tasks. From decision support in underwriting to smart assistants for sales and HR, AI is being integrated into workflows to free up human attention for higher-value work. The focus is shifting from automation to collaboration between humans and machines.

Rising Pressure on Productivity and Efficiency
In an environment of tight budgets and rising expectations, AI offers a path to optimize resources, reduce manual effort, and scale operations without linear increases in headcount. Use cases in document processing, reconciliation, claims handling, and ticket triage are delivering clear ROI.

Increased Demand for Real-Time, Predictive, and Adaptive Systems
Static reporting and batch analytics no longer suffice. Businesses want real-time insights and forward-looking intelligence. AI enables dynamic pricing, proactive maintenance, fraud prediction, and inventory optimization—capabilities that create agility in volatile markets.

Regulatory and Risk Functions Need Continuous Intelligence
As compliance becomes more data-driven and scrutiny intensifies, AI is being deployed in risk scoring, anomaly detection, and audit analytics. It helps ensure not just speed, but control—surfacing insights that humans might miss across thousands of transactions or logs.

AI Is Becoming Embedded in Core Platforms
Modern enterprise applications—from CRM to ERP to ITSM—are increasingly offering built-in AI features, making adoption easier and more pervasive. Enterprises are leveraging these native capabilities to jumpstart value realization.

Executive Sponsorship Is Stronger Than Ever
The conversation around AI has moved from the IT lab to the boardroom. With CXOs now championing AI initiatives, there's greater alignment between business priorities and AI strategy, unlocking cross-functional investment and faster decision-making.

Implementation Challenges: What's Slowing AI Value Realization

While AI adoption is accelerating, enterprise-scale value realization remains elusive for many organizations. Technical capability is just one part of the equation; realizing sustained impact from AI also requires governance, change management, and alignment across business and IT. Below are the key challenges that surfaced through the panel discussions:

Fragmented and Poor-Quality Data
AI performance is only as good as the data feeding it. Many enterprises still struggle with data silos, inconsistent formats, and limited lineage, making it hard to train, test, and scale models effectively. Even with modern data platforms in place, ensuring data readiness and contextual richness remains a bottleneck.

> "Technical talent and tooling are necessary but not sufficient. CIOs must lead with clarity, governance, and a commitment to business alignment."

Lack of Business Ownership and Alignment
AI projects often start in tech teams but fail to gain traction with business users. Without clear ownership, sponsorship, and embedded metrics tied to business outcomes, AI pilots stagnate or fail to scale. True adoption requires co-ownership between business and IT, not just buy-in.

Trust and Explainability Gaps
Stakeholders are reluctant to act on AI recommendations when models behave like black boxes. Explainability, traceability, and transparency are essential—especially in regulated sectors. Without them, user confidence falters and adoption stalls, even if model accuracy is high.

Skill Shortage Challenges Across the Value Chain
While many organizations have data scientists, fewer have the cross-functional talent needed to operationalize AI. MLOps engineers, data stewards, model validators, and domain interpreters are in short supply. This creates gaps between development and deployment, delaying time to value.

Technical Debt and Legacy Constraints Issues
AI models require access to real-time data and modern interfaces, which are often hard to extract from legacy core systems. Integrating AI with monolithic applications, manual workflows, or fragmented APIs slows adoption and creates friction for automation.

Inconsistent MLOps Maturity
Many enterprises lack mature model lifecycle management practices. Without standardized pipelines for versioning, retraining, monitoring, and governance, AI models become unreliable or outdated quickly. This undermines both performance and credibility over time.

Security, Privacy, and Ethical Risks
As AI systems touch sensitive data and influence decisions, privacy controls, data minimization, and ethical guardrails become critical. Yet many organizations still lack formal frameworks for AI governance, leaving them exposed to risk—even when intentions are sound.

Over-Reliance on Vendors and Black-Box Tools
Some teams depend too heavily on off-the-shelf AI tools or external vendors, resulting in limited customization, poor interpretability, and lock-in risks. This can hinder agility, reduce control, and make it difficult to adapt AI to evolving business needs.

Cost of Scaling
Even when initial AI use cases succeed, scaling them across geographies, business units, or customer segments brings infrastructure and cost challenges. Without clear ROI metrics and prioritization frame-

works, enterprises struggle to make a business case for broad rollout.

Misaligned Expectations and Hype Cycles
AI is often overhyped in executive conversations, leading to misplaced bets or premature scaling. The absence of clear value frameworks and realistic timelines results in disillusionment—and in some cases, backlash from internal stakeholders.

These challenges don't signal failure—they highlight the growing pains of a transformative capability. The organizations that succeed with AI aren't those with perfect data or mature tooling, but those that build cross-functional muscle, define clear value pathways, and treat AI as a long-term operating capability.

The Way Forward: From Tactical Wins to Enterprise-Scale AI

To move beyond isolated use cases and realize sustainable value from AI, enterprises must think in terms of platforms, people, and purpose. Scaling AI is not about multiplying models—it's about embedding intelligence into the workflows, tools, and decision-making structures that run the business. The organizations leading the way are taking the following steps:

Define a Clear AI Strategy Anchored to Business Outcomes
Success starts with alignment. Leading CIOs are partnering with functional leaders to co-create AI roadmaps tied to revenue growth, cost reduction, or customer experience KPIs. Rather than chasing trends, they're building focused portfolios of use cases that matter most to the business—and defining success in measurable terms.

Invest in MLOps for Scalability and Governance
AI cannot scale without engineering discipline. Organizations are investing in robust MLOps pipelines—including model versioning, testing, explainability, monitoring, and rollback mechanisms. This ensures that models remain performant, secure, and compliant over time.

Build Cross-Functional AI Squads and CoEs
The best AI teams aren't just data scientists—they're cross-functional squads that bring together business analysts, product owners, data engineers, risk managers, and ops teams. Many enterprises are formalizing this through AI Centers of Excellence (CoEs) that provide governance, shared tooling, and reusable accelerators across business units.

Prioritize Data Readiness and Contextualization
Clean, accessible, and well-labeled data is the bedrock of AI. But it's not just about volume—it's about context. Enterprises are

> "Enterprises are beginning to reimagine how work is performed—with AI not replacing humans, but augmenting their capabilities."

investing in metadata management, master data programs, and semantic layers that make data understandable and usable for AI across domains.

Scale AI Through Platform Thinking
Instead of building one model at a time, forward-looking organizations are investing in AI platforms that offer reusable components, modular architectures, and enterprise-grade APIs. This enables faster development, consistent standards, and reduced duplication across teams.

Lay Focus on Explainability and User Trust
No AI system can thrive without human trust. Leaders are integrating explainability features into models and dashboards, allowing business users to understand how predictions were made—and to challenge or override them when necessary.

Incorporate AI into Everyday Tools and Processes
Rather than asking users to adopt new tools, successful deployments embed AI into existing systems—CRM, ERP, HRMS, service portals—via co-pilots, bots, and nudges. This lowers the barrier to adoption and ensures AI becomes part of the user's flow of work.

Embed Risk and Ethics from the Start
Trustworthy AI requires proactive risk management. Enterprises are adopting AI governance frameworks that cover bias detection, fairness audits, consent management, and regulatory compliance—often in collaboration with legal, audit, and security teams.

Build AI Literacy Throughout the Enterprise
AI is no longer the domain of specialists alone. Organizations are rolling out training programs, simulations, and sandbox environments to help business teams understand AI basics, interpret outputs, and work confidently with machine-generated insights.

Start Small—but Design for Scale
The journey to enterprise AI often begins with focused use cases—but those use cases must be designed with scale in mind. Reusable datasets, modular pipelines, and flexible governance are key to ensuring that early wins become enterprise capabilities.

The most successful AI adopters treat AI not as a one-off project but as a strategic discipline that evolves with the business. They focus not just on what AI can do, but on how the organization must change to unlock its full potential.

From Insight to Impact: AI as a Core Business Capability

AI is no longer confined to innovation labs or niche pilots. It is becoming a foundational capability that shapes how modern enterprises sense, decide, and act—in real time, at scale, and with confidence. But this shift doesn't happen by accident. It requires intentional design, cultural alignment, and long-term commitment.

Across industries, leading CIOs are repositioning AI as a strategic operating layer—not just a tool for analytics, but a driver of continuous learning, customer-centricity, and business resilience. They are embedding AI into core platforms, workflows, and customer touchpoints—not to replace humans, but to amplify their ability to make faster, smarter, and more consistent decisions.

This new model of enterprise AI is defined not by the number of models in production, but by how well AI is woven into the day-to-day rhythm of the business. In risk, AI enables proactive detection and early mitigation. In customer service, it augments agents with insights, summaries, and recommendations. In finance and HR, it accelerates reconciliation, forecasting, and policy compliance.

Crucially, AI maturity is not measured solely by technical sophistication—it's measured by trust, adoption, and impact. That means building systems that are transparent, fair, and easy to engage with. It means giving users—not just data scientists—the confidence to rely on AI for meaningful decisions. And it means making AI responsive to change—whether that's a shift in regulation, customer behavior, or business strategy.

Enterprises that succeed with AI do so because they treat it not as a magic bullet, but as a strategic capability that evolves. They combine engineering excellence with strong governance. They balance speed with safeguards. And they focus relentlessly on outcomes—not just predictions.

As the role of AI continues to expand—from recommendation engines to generative copilots, from ML-powered operations to autonomous agents—the real differentiator will be how well enterprises connect AI to value.

That's the path forward: not AI for AI's sake, but AI as an intelligent, ethical, and enterprise-wide enabler of business transformation.

AUTHORS

Enterprise GenAI: Strategy, Oversight, and Guardrails Over Hype and Haste

Anil Kuril
CISO & Head - Data Protection Office, Union Bank of India

Why AI Demands Data Discipline and Iterative Precision

Dr Pankaj Dikshit
Chief Technology Officer & EVP, Government e-Marketplace

Embedding AI with Purpose: A Playbook for Scalable, Responsible Impact

Gaurav Kataria
Vice President - Digital (Manufacturing) & CDIO, PSPD, ITC

AI Belongs in the Fabric of Business, Not Just the Stack

Metessh D Bhati
EVP & Chief Digital and AI Officer, Protean eGov Technologies

Observability and Strategy Transform AI from Opaque to Trusted System

Prasad Rao
Senior Director - IT, Vee Healthtek

Scaling AI with Guardrails, Not Guesswork

Praveen Shrikhande
Chief Digital & Information Officer, Aditya Birla Fashion & Retail

AI Leadership Demands Action, Not Hesitation

Sandeep Soman
Senior Director, IT, Global Pharmaceutical Major

Smarter AI Starts with Trust, Alignment, and Business Pull

Sankaranarayanan Raghavan
Chief Technology and Data Officer, IndiaFirst Life

To Scale AI, Build Platforms—Not Just Models!

Shankar G. Rao
Vice President, Chief Information Officer & Chief Digital Officer, Bosch

AI That Lasts Starts with Value, Not Hype

Sunil Mishra
Chief Information Officer, Kotak Securities

Why AI Fails as an IT Project and Succeeds as Strategy

Vamsi Krishna Ithamraju
Chief Technology Officer, Axis AMC

AI Is a Business Transformation Journey, Not Just a Tech Deployment

Venkat Krishnan V
Chief Information Officer, Karnataka Bank

AI Applications: Orchestrating Intelligent Automation Across the Digital Enterprise

Nalin Agrawal
Director - Solutions Engineering, Dynatrace

Vilas Landge
Business Head - Strategic Accounts, Dynatrace

Enterprise GenAI: Strategy, Oversight, and Guardrails Over Hype and Haste

Enterprise GenAI demands governance, strategy, and continuous human-in-the-loop validation.

Our journey with AI and GenAI has taught us that expectation management is as critical as the technology itself. In boardroom conversations, there is often a perception that AI is a plug-and-play tool—especially because public GenAI tools like ChatGPT make it appear so effortless. But in a highly regulated environment like ours, the reality is far more complex. As a CIO or CISO, I often find myself explaining why enterprise AI adoption is not as straightforward as it seems.

One of the first things we had to clarify—both internally and with stakeholders—was that GenAI is not the same as traditional AI. Classical AI has long existed in banking, powering rule-based systems and analytics. GenAI, however, is non-deterministic, context-driven, and inherently less predictable. It requires a completely different strategic lens.

Before implementation, we established an Innovation Hub to evaluate emerging technologies. An early GenAI experiment revealed unreliable and inconsistent outputs—even in controlled settings—highlighting the need for a more strategic, enterprise-grade approach. We responded by building a structured framework across three parallel tracks:

- A dedicated Center of Excellence (CoE)
- A use case–driven model
- A platform-based approach

Given our regulatory obligations, we intentionally limited GenAI use to internal-facing applications.

A key principle of our GenAI adoption is the human-in-the-loop approach, embedded across all workflows to ensure accuracy and accountability. Our field staff use GenAI to address queries related to products, processes, HR, and in some cases, customer interactions. However, every GenAI response is reviewed by a human before action is taken.

We've also embedded an explainability framework, which has helped identify gaps and refine response quality. In the early stages—even with Retrieval-Augmented Generation (RAG)—we saw only around 50% accuracy. But through iterative feedback and improvement, we've increased this to approximately 70–75%.

One of the biggest challenges we face is data governance and model auditability. Most organizations still lack robust frameworks to identify and mitigate vulnerabilities in GenAI systems. While the promise is immense, enterprise GenAI is not the same as personal use. It requires quality data, explainability, human oversight, and a governance structure that aligns not only with business needs, but also with ethical and regulatory standards.

"We quickly realized GenAI isn't plug-and-play—it demands strategy, guardrails, and human oversight."

Anil Kuril
Chief Information Security Officer & Head - Data Protection Office,
Union Bank of India

Why AI Demands Data Discipline and Iterative Precision

Noisy data breaks AI. Discipline and iteration build trust.

"AI success begins with trustworthy data, purposeful use cases, and continuous iteration."

Dr Pankaj Dixit
Chief Technology Officer & EVP,
Government e-Marketplace

Over the past 15 months, we have been on a journey to bring meaningful AI applications onto our platform. The first step was to build a customer-facing chatbot. It quickly became clear that general-purpose LLMs were not enough and we needed a retrieval-augmented generation (RAG) model built on our platform-specific knowledge. But building that knowledge base was no small feat as the data had to be cleaned, sanitized and validated.

The problem is that AI doesn't correct data—it amplifies it. So, if your input is flawed, your output will be flawed too. After much iteration and user feedback, we reached a 90–95% accuracy level and felt confident moving it from beta to production.

From there, we expanded into voice-enabled ticketing. This use-case turned out to be far easier, thanks to existing APIs for voice translation and simpler backend integration. Users could now raise support tickets and receive updates across multiple languages through conversational AI.

The next big frontier was to move beyond pre-login information to authenticated, user-specific queries—orders, tenders, bids. We are piloting with security and privacy top of mind. In parallel, we are exploring conversational BI, which remains our most complex and challenging use case. The idea of letting users query large databases in natural language and auto-generate dashboards sounds great, but making it work with thousands of columns, live data and multiple layers of context is incredibly hard.

We are not yet ready for beta—but we are getting closer.

We have also been experimenting with code generation and automated documentation, but those models need strict guardrails and validation before deployment. Through all this, one thing has become clear: the success of AI depends entirely on the quality and readiness of data.

Another key takeaway is governance. From hallucinations to data privacy, control is non-negotiable. Explainability is another focus—we want new team members to understand how models work, even after people move on. Building AI is not about plugging in a model—it's about understanding your data, aligning with business needs, and iterating relentlessly.

Embedding AI with Purpose: A Playbook for Scalable, Responsible Impact

People, process, data, and governance shape meaningful enterprise AI.

Our AI journey has progressed well beyond isolated experimentation and is now on the path to enterprise-wide adoption. Real AI adoption depends on getting three things right—people, process, and technology. Our strategic focus is on the value we extract from digital investments—whether to improve efficiency and profits, grow revenue, or drive societal impact such as sustainability, worker safety, or better outcomes for farmers.

In a manufacturing and supply chain–heavy setup like ours, everything begins with data. On the technology front, investing in foundational data infrastructure is a prerequisite. Without clean, accessible, and well-governed data, no AI layer can deliver value. We have traditionally stored data on-premises. Only the datasets needed for MLOps or training are moved to the cloud. This hybrid model gives us the best of both worlds—security and cost efficiency—while preserving OT-IT segregation, governance, and layered security.

Process alignment is equally critical. Every AI project is jointly owned by a business and a tech leader to ensure accountability and problem-solution alignment. Often, AI insights require process changes—and that only works if business stakeholders are aligned from day one.

Then comes the people aspect. We've made training and change management a priority—from the boardroom to the shop floor—to enable innovation and build trust in the system.

One key lesson is never to underestimate the cost of scale. We've seen great pilots fail because scalability wasn't factored in early enough. Now, every proof of concept is evaluated for both business impact and scalability from the beginning.

Operational challenges such as model drift, hallucinations, and black-box logic are real. So, every project begins with a clear objective, owner, and success metric.

Societal impact is deeply woven into our AI strategy. Whether it's improving worker safety, reducing water consumption in manufacturing, or optimizing energy in our mills, we are using AI as a lever for sustainability and shared value.

Our journey is about more than technology. It's about embedding AI into how we work—responsibly, at scale, and with lasting impact.

"Think impact, start small, scale responsibly—that's the AI playbook we follow."

Gaurav Kataria
Vice President - Digital (Manufacturing) & CDIO, PSPD, ITC

AI Belongs in the Fabric of Business, Not Just the Stack

Intelligent enterprises embed AI into core systems, values, and behaviors.

"AI needs to be embedded in the architecture, platforms, and principles—and absorbed by the people."

Metesh D Bhati
EVP & Chief Digital and AI Officer,
Protean eGov Technologies

AI is no longer just a technology to be integrated—it must become a foundational capability within the enterprise. True value is unlocked when AI is embedded deeply into the core of systems, workflows, and decision-making processes. The shift from integration to immersion must not just be AI-enabled but AI-native—systems that learn continuously and adapt intelligently.

Many organizations still treat AI like any other software product—something to install, configure, and plug into existing systems. But AI demands a different approach—one that starts with reimagining platforms and infrastructure to enable intelligent workflows and real-time decisions. It requires moving from digital-first to AI-first thinking.

A major opportunity lies in the data enterprises already possess. Instead of choosing off-the-shelf large language models (LLMs), organizations gain more by building small, domain-specific models trained on proprietary data. However, they must overcome challenges such as the lack of standardized ontologies, missing semantic layers, and inconsistent data governance.

Governance must also be rethought. Compliance—especially with evolving regulations like the DPDP Act—should not be treated as a checklist item. Data governance must be compliant by design, with ethics, explainability, and accountability built into the architecture from the ground up.

Equally critical is absorption. It is not enough to adopt AI—the enterprise must absorb it. That means aligning AI with people, ethics, and evolving business models. Intelligence must not only live in systems but also shape how decisions are made and trust is built across stakeholders.

At an operational level, deploying AI presents complex, practical challenges. Latency in data ingestion, drift in model performance, and weak alignment between ML metrics and business KPIs are common friction points. Many monitoring systems fail to reflect the domain-specific context in which AI operates. To overcome this, mature teams are building observability layers to track drift, ensure lineage, and generate insights across the full AI lifecycle.

Ultimately, AI success hinges on mindset, architecture, and accountability—and how deeply it is absorbed into the fabric of the enterprise.

Observability and Strategy Transform AI from Opaque to Trusted System

AI earns trust through transparency, observability, compliance, and cultural alignment.

As a healthcare revenue cycle management automation company, we operate in a heavily regulated domain. My foremost concern is safeguarding data privacy and ensuring compliance while embracing AI to drive innovation.

A core challenge we face is managing apprehension within the organization with employees constantly asking if they can use tools like ChatGPT, Gemini or Copilot. While these tools hold promise, there is fear about violating compliance norms or risking data privacy. That fear is valid—especially when clients and government authorities scrutinize every step we take with sensitive data.

Our operations rely heavily on both internally developed legacy systems and new applications tailored to client requirements. These systems, though valuable, often operate in silos, making it hard to draw unified insights or modernize effectively. Bridging that divide is essential to optimize performance and unlock real-time business insights.

The first step for us has been data classification—clearly separating compliance-bound data from less sensitive information. By doing this, we can explore automation, analytics, and AI use cases in a secure, structured way. We have realized that even non-critical data, when harnessed properly with the right tools, can yield significant value in simplifying manual tasks and improving efficiency.

That said, AI implementation is not just about adding intelligent layers. It is about confronting deep-rooted issues with AI's 'black box' behavior. Outputs from large language models (LLMs) can be unpredictable—bias, hallucinations, or performance drift are real concerns. Moreover, as these models evolve, the complexity of data and interactions make real-time tracking and accountability difficult.

To tackle this, we are actively evaluating enterprise-grade observability platforms. Our aim is to gain full-stack visibility—from infrastructure health to token usage, latency, and model drift—across our AI and cloud-native environments. The goal is clear: tie AI performance metrics back to user experience and business outcomes.

Our focus is to build a culture of responsible AI adoption wherein teams trust the systems they use, data is protected, and every AI insight translates into tangible business value. It is not just about embracing AI—it's about doing it securely, strategically and transparently.

"AI trust begins with transparency—and that starts with data."

Prasad Rao
Senior Director - IT,
Vee Healthtek

Scaling AI with Guardrails, Not Guesswork

You can't scale AI through isolated brilliance—it needs structure, standards, and shared accountability.

"AI at scale needs governance, observability, and business-aligned orchestration—not silos and guesswork."

Praveen Shrikhande
Chief Digital & Information Officer,
Aditya Birla Fashion & Retail

At Aditya Birla Fashion and Retail, we believe the enterprise AI journey is no longer just about experimentation—it's about orchestration. While the company has embraced AI across functions like merchandising, marketing, and supply chain, we are clear that scaling such efforts requires more than just scattered pilots or domain-specific enthusiasm. It calls for a deliberate framework that balances innovation with oversight.

Teams would build AI models in silos— on notebooks, in sandbox environments—but without lifecycle visibility, MLOps discipline, or data lineage. The outcome was often interesting in isolation, but fragile or inconsistent when it needed to be scaled for production.

To address this, our team has adopted a federated AI model—allowing business units the freedom to innovate while providing a shared backbone for governance, data access, security, and model operations. It's about setting guardrails, not roadblocks. The company developed centralized data platforms and standardized MLOps pipelines to help teams deploy, monitor, and version AI models in a consistent, auditable way.

A critical component of this journey is AI observability, which must extend beyond infrastructure health or uptime—it needs to monitor how models perform in the real world and track for model drift, latency issues, hallucination risks, and responsiveness to prompts. And we have to look at how those affect business KPIs, not just system metrics.

Infrastructure decisions also play a critical role. While we generally supports a cloud-first approach for flexibility and rapid scaling, sensitive workloads involving personally identifiable information (PII) or high-security classifications would need to be treated differently.

Another recurring challenge is cross-functional alignment. AI can't succeed if it's treated as a data science-only problem. You need product managers, compliance officers, business leaders, and IT architects working together. This requires establishing governance councils that bring together stakeholders from across the business to evaluate AI use cases, assess risks, and ensure ethical boundaries are observed.

Crucially, the real opportunity lies in institutionalizing AI as a capability, not j ust a technology. It's not about replacing people—it's about augmenting decisions. But that augmentation has to be explainable, measurable, and relevant to business context.

For CIOs embarking on similar journeys, it is important to move fast, but not loose. Don't wait to perfect the stack before starting, but don't launch AI programs without basic accountability structures either. Build shared platforms. Create sandboxes—but connect them to real-world feedback loops. And above all, make AI a business conversation, not just a technical one.

AI Leadership Demands Action, Not Hesitation

AI rewards experimentation, mindset shifts, and human-centric, cloud-first adoption.

Our AI journey has made one thing clear: the real challenge is not technology—it's cultural. As models evolve and innovation outpaces implementation, we are constantly balancing between investing in current solutions and evaluating what's next. From managing hallucinations and retraining cycles to ensuring data readiness and model auditability, AI adoption remains a dynamic endeavor that requires clarity, confidence, and a mindset shift.

Our focus has been on user empowerment, enriching user experience and productivity gains for the workforce. It is not towards making people redundant but instead to augment them in their jobs to make them 'AI Powered'. That distinction is critical, especially in an industry like pharma, where any mention of AI often triggers fear around job displacement. We're using AI to enhance employee experience—building solutions that augment human expertise, not automation for its own sake.

We embed AI into existing enterprise systems where data already lives. This minimizes governance friction and increases data reliability. Where we need to pull from multiple sources, we use a centralized data lake with real-time integrations. We have also made sure business users are part of the loop and understand the positive business outcomes that will come with AI adoptions. Once they understand that data quality drives AI accuracy, they become champions—proactively maintaining the information their tools rely on.

We are heavily invested in Cloud with all modern global application on public cloud via Saas/Paas/Iaas Models. On-prem does not offer the flexibility or speed we need to stay current. Our systems are already designed with security, privacy, and data isolation in mind—fundamentals that give us the confidence to scale cloud-first. In addition, we leverage the best cloud features for cybersecurity and data protection.

We face pushback on ROI as AI does not offer linear returns like traditional IT investments, and business leaders often ask for hard numbers upfront. But this is a paradigm shift—comparable to certain extent with the internet or PC revolution but AI is much beyond. There is a constant struggle between born analog and born digital companies in New Technology Adoptions. To mitigate the same, we have created small, agile teams with the freedom to experiment, mindset to innovate and confidence to fail but learn.

I also urge our AI partners to build more guardrails, offer better grounding, deliver consistency and accuracy. Teams must upskill for AI as it will touch every role, and those who embrace it will lead the transformation.

"AI is moving fast, and we cannot afford to sit on the sidelines. We must experiment boldly, embrace uncertainty, and build the future before AI passes us by."

Sandeep Soman
Senior Director, IT,
Global Pharmaceutical Major

Smarter AI Starts with Trust, Alignment, and Business Pull

Business-led AI builds momentum when trust precedes technology and ROI.

"Start small, and let the business lead the AI shift."

Sankaranarayanan Raghavan
Chief Technology and Data Officer, IndiaFirst Life

AI at IndiaFirst Life has transitioned from a top-down enforcement approach to a function-led, demand-driven model. We have deployed two sets of Large Language Models (LLMs): one built using Vertex AI, which sits on our lakehouse with native monitoring, and another custom setup layered on ChatGPT. As more use cases emerge, I see a growing need for unified AI observability, and we're actively working toward that.

In the early days, even predictive models faced resistance. Adoption picked up only when leadership made it clear that AI was not about replacing people, but about enabling them to work smarter. Once that trust was built, teams began to see the ROI and the broader value.

What excites me now is the shift in who's driving the conversation. Just a few months ago, it was management asking questions about missing out on GenAI. Today, business functions are approaching us with ideas—asking how they can use AI to improve compliance, automate processes, or personalize engagement. That transition from tech-push to business-pull is powerful.

Functions are now thinking more deeply about how AI can reduce bias, improve transparency, and drive operational excellence. In sales, for example, high attrition means we often onboard new talent. AI helps ramp them up quickly by guiding conversations with the right pitch and product knowledge. That's vital in a long-term business like life insurance, where trust drives revenue.

ROI is no longer the starting point, as efforts have shifted toward learning and experimentation. That doesn't mean we ignore the fundamentals—you still need clean data, a secure pipeline, and strong DevOps practices.

As a regulated organization, we remain focused on responsible AI use. Teams are committed to being 100% compliant—meeting tax obligations, staying within regulatory boundaries, and aligning performance metrics appropriately.

Begin with a horizontal use case—training, internal search, or customer service. These help test the environment for data governance, change management, and security. Sometimes, being a late adopter is an advantage—you benefit from the latest tools and avoid early missteps. AI is not just about technology; it's about creating a future where insurance feels intuitive, personalized, and human-first.

To Scale AI, Build Platforms—Not Just Models!

Platform thinking turns scattered AI efforts into scalable enterprise strategy.

Our AI journey began with scattered innovation— pockets of experimentation led by enthusiastic developers across manufacturing plants and functions. Many of these teams were building promising models independently. While there was no shortage of talent or intent, these efforts were happening in silos and the value they created was limited to local use cases. It lacked the scale or governance to deliver enterprise-level impact.

Recognizing this, we knew it was time to evolve from isolated innovation to institutional transformation. Our current approach is built around a hybrid AI operating model. This means that while innovation continues to originate within the domains—close to the business in plants and functional teams—we have added centralized support from the CIO's office. We provide foundational elements such as shared infrastructure, common data architecture, guardrails for responsible AI, and a central platform that enables model reuse. It is a federated delivery model to balance agility and alignment.

A key pillar is our Office of Responsible AI to ensure every model is tested not just for performance, but for explainability, bias, transparency and compliance. AI governance must be embedded from the start.

We have also introduced a FinOps layer to manage AI-related costs. AI infrastructure is expensive and without centralized oversight, there will be duplication in tools, cloud environments or models. We monitor usage centrally and maintain a shared repository to enable reuse. It is not just about saving costs; it is about scaling smartly.

Data security and privacy guide our architectural decisions. Sensitive data remains on-premises while other data—once tokenized or wrapped—can safely reside on the cloud.

We are adopting the LANG framework for prompt engineering to bring structure to how prompts are designed, factoring in logic, audience, output type, goals and constraints. Strategically, our AI programs align with three pillars: Grow, Operate, and Experience.

The path to meaningful AI adoption is about combining domain-led innovation with centralized guardrails. It's not just about deploying models—it's about building a scalable, secure, and responsible AI ecosystem that delivers value across the enterprise.

"AI at scale needs structure, guardrails, and a platform mindset."

Shankar G Rao
Vice President,
Chief Information Officer &
Chief Digital Officer, Bosch

AI That Lasts Starts with Value, Not Hype

Real-world AI impact begins with clear value, not hype cycles.

"Start small, think smart: scaling AI begins with business ownership—and data readiness."

Sunil Mishra
Chief Information Officer,
Kotak Securities

Our AI journey has been as much about understanding real-world constraints as it has been about chasing innovation. Scaling AI in a heavily regulated, customer-centric business like ours is about getting the basics right—from data quality to business alignment.

When we started exploring AI use cases we had practical considerations, such as, how much data and compute will be needed and whether it will work across all use cases. User experience is another challenge. The quality of responses, maintaining the context of previous questions and capturing user feedback meaningfully impacts adoption. But the biggest hurdle is to link AI to business value— unless we define the business problem clearly from the start, translating a model's performance into actual ROI is tough.

That is why I have learned to treat every AI project like a business initiative, not just a tech one. If there isn't a strong use case backed by a business owner—whether in operations, compliance, or customer service—it's unlikely to succeed.

When it comes to data, we have a hybrid approach. For public or marketing use cases we use cloud-based LLMs. But anything involving customer transactions, that is a hard no. We are exploring our own versions of AI agents and training smaller, more focused models internally.

Cost is another key factor. Setting up your own GPU infrastructure is expensive unless you have a clear roadmap of use cases. And scalability must be baked in from the start. We have seen promising PoCs fail to scale simply because the cost of deployment was too high.

We've also realized that adoption doesn't start with complex AI—it starts with small wins. Automating document summaries, generating research insights, or simplifying internal workflows might seem minor, but they build confidence and prove the value quickly. In regulated industries like finance, AI can help create outputs, but humans still need to review and approve to ensure compliance.

We're still learning, still experimenting—but with each iteration, we're getting closer to practical, responsible AI.

Why AI Fails as an IT Project and Succeeds as Strategy

AI thrives with strategic alignment, cloud agility, and governance-first approach.

Organizations adopting AI are realizing that cloud-first strategies offer unmatched agility, especially given how rapidly the underlying AI infrastructure landscape is evolving. However, regulatory constraints are a critical factor. For companies in highly regulated industries like BFSI, cloud adoption is contingent on data residency requirements. If a necessary cloud service is not available in the local region or availability zone (like AWS or Azure), teams often choose to wait rather than risk non-compliance.

Security and access control are equally important. We are using models like CASB (Cloud Access Security Broker) to manage permissions granularly ensuring that different users can access only the features relevant to them. This protects sensitive operations while enabling flexibility in user experiences.

On-premise AI infrastructure will only be considered when there is long-term clarity on use cases, talent availability and ROI. Until then, the cloud remains the platform of choice for its scalability, flexibility and cost-effectiveness—provided it aligns with strict security and compliance requirements.

One of the most critical learnings has been around organizational structure. We found that when AI initiatives were initially run purely as IT projects, it failed to deliver results. We experienced poor adoption. There was also limited business alignment and therefore, proofs of concept did not succeed. But after we repositioned it as a strategic business program, the shift began to happen. Our AI initiatives are now run by leadership teams—reporting directly to the CEO—and are prioritized based on function-specific needs and ROI. This top-down sponsorship gives accountability, drives adoption and helps to keep the initiatives grounded in business outcomes.

A major risk comes from the emergence of shadow IT, where non-IT departments set up their own infrastructure to run AI experiments. This results not only in duplication of cost but also in significant security and compliance risks. These teams lack the capability to manage patches, updates or audits. When something goes wrong, the burden falls back on IT, exposing the organization to governance issues and potential regulatory fallout.

"Treating AI as just another IT project is a recipe for failure. AI must be embedded within strategic business priorities, backed by governance and accountability."

Vamsi Krishna Ithamraju
Chief Technology Officer, Axis AMC

"AI is not just tech—it's about aligning data, processes, and goals, and astute change management."

Venkat Krishnan V
Chief Information Officer,
Karnataka Bank

AI Is a Business Transformation Journey, Not Just a Tech Deployment

AI works when business metrics, not buzzwords, guide execution.

In a highly regulated sector like banking, the primary concern around AI adoption isn't just implementation—it's about maintaining control, ensuring governance, and securing data. While incoming data can be managed, the real risk lies in what might inadvertently go out. To address this, the bank has built strong guardrails, including the conceptualization of an "AI Gateway," modeled after API gateways and firewalls.

The AI journey is not merely technical—it requires organizational alignment, astute change management, and a clearly articulated vision to ensure stakeholders understand what AI can do and how it can deliver measurable impact.

Strategically, the bank—now over a century old—is undergoing a large-scale reengineering effort, placing AI at the core of its architectural transformation. AI is no longer a nice-to-have but a foundational capability to modernize everything from fraud detection to core banking operations. In risk and fraud management, AI is used to reduce false positives from vast volumes of transaction data.

Decisions around cloud versus on-premise deployment depend heavily on data classification. For highly sensitive personal data, on-premise remains the default. For less sensitive workloads, the cloud offers agility. The bank is also investing in robust governance mechanisms to ensure full visibility and control over all AI-related data flows.

Another focus is decoupling monolithic systems by optimizing core banking workflows. For example, updating customer contact details—traditionally handled via CRM—can now be executed by intelligent, channel-agnostic AI agents.

Crucially, every AI initiative must demonstrate a return on business metrics—whether it's improving the cost-to-income ratio, increasing profitability per customer, or enhancing customer experience. These KPIs keep AI initiatives anchored in real business value.

An often overlooked challenge is the lack of clarity in understanding AI itself. Terms like AI, machine learning, and automation are often used interchangeably, leading to confusion and inflated expectations. To counter this, the bank is building internal frameworks to define the scope and application of each technology precisely, ensuring better understanding and smoother adoption across the organization.

AI Applications: Orchestrating Intelligent Automation Across the Digital Enterprise

Scalable, responsible AI demands more than innovation. It requires real-time observability, intelligent automation, and governance aligned to business outcomes across hybrid environments.

Enterprise AI adoption has entered a pivotal phase. No longer confined to labs or pilot projects, AI is being embedded into core operational and customer-facing systems. Across sectors—from banking and manufacturing to healthcare and retail—organizations are embedding predictive models, generative AI, and intelligent agents into business workflows to drive productivity, agility, and top-line growth.

However, this AI expansion is not without its challenges. Enterprises are increasingly focusing on responsible AI for transparency, fairness, explainability, and regulatory compliance. Concerns over data privacy, prompt manipulation, hallucinations, and model drift have made AI observability a top priority. AI observability monitors, analyzes, and explains the behavior, performance, and cost of AI applications, models, agents, and infrastructure.

Furthermore, the decision of whether to host models on-premises, in hybrid setups, or on public cloud hinges on factors like data sovereignty, latency, compute availability, and control. Additionally, organizations are realizing that AI applications cannot succeed in silos.

Fragmented adoption across departments, legacy infrastructure, and talent gaps often stall enterprise-scale deployments. The need of the hour is a unifying layer of intelligent observability that offers visibility not just into system health, but into AI model behavior, usage, and business impact in real time.

> "As AI moves into production, enterprises need observability not just for systems but for model behavior, performance drift, and decision transparency."

Nalin Agrawal
Director, Solution Engineering,
Dynatrace

Vilas Landge
Business Head, Strategic Accounts, Dynatrace

"AI success requires balance: innovation at the edge and governance at the center. Dynatrace enables this balance with AI-powered observability."

CIO Challenges: Fragmentation, Trust, and Governance Gaps

Despite AI's promise, CIOs and technology leaders face multiple friction points in translating experimentation into enterprise value. These include the following:

- **Fragmented AI Initiatives.** AI projects are often scattered across business units and functions, leading to duplication, inconsistency, and a lack of shared learning. Federated models need centralized oversight to align with enterprise goals.
- **Limited Trust in AI Decisions**. Black-box behavior in large language models (LLMs) and predictive models creates a trust deficit among business users, especially in regulated domains like finance, pharma, and insurance.
- **Operational Monitoring Gaps.** Traditional monitoring tools fall short in detecting AI-specific issues such as model drift, response latency, hallucinations, or tokenization failures. Without real-time diagnostics, issues escalate before mitigation, negatively impacting employee and customer experiences.
- **Cloud vs. On-prem Dilemmas.** Teams must weigh the agility and scalability offered by cloud against privacy, cost, and regulatory constraints that sometimes favor on-premises or hybrid deployments.
- **Ambiguity in Business ROI.** Linking AI outputs to business KPIs, such as revenue uplift, operational cost savings, or customer retention, remains a challenge, hindering the justification of further investments.
- **Security and Governance Gaps.** Prompt injection, unauthorized model access, or exposure of personally identifiable information (PII) in generative responses demand robust guardrails and observability at runtime.

Enterprises want AI that is not only performant and scalable but also explainable, governable, and cost-effective—and this is where traditional tooling falls short. These challenges highlight a critical gap: Traditional monitoring and governance tools aren't designed to handle AI's unique complexities. Dynatrace addresses this through a platform-first approach to unified observability.

Operationalizing AI at Scale

Dynatrace brings a unified platform approach to operationalizing AI applications, with full-stack, AI-powered observability extending from infrastructure and applications to model behavior, user interaction, and business outcomes. As AI applications evolve, real-time observability becomes non-negotiable. Dynatrace enables teams to monitor not just application telemetry, but also AI workloads—tracking inference latency, model availability, data pipeline health, and GPU utilization in real time. This provides the foundational reliability needed to scale AI into production.

Going further, the platform offers AI model observability capabilities that can detect performance drift, hallucination trends, and prompt anomalies. This helps

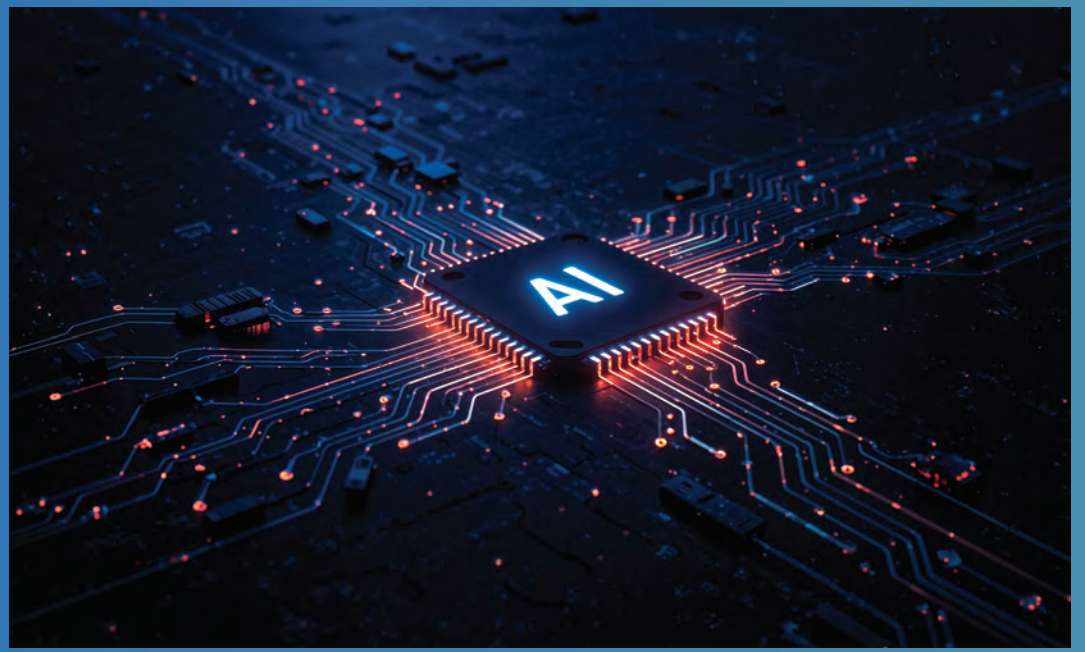

pre-empt incidents that could impact decision-making or customer trust. Business-facing dashboards link model outputs to KPIs like cost per decision or SLA compliance— bridging the gap between technical insights and strategic impact. It is relevant to underscore the importance of a single pane of glass for hybrid, multi-cloud AI environments. Many enterprises use a mix of cloud providers, open-source libraries, and proprietary LLMs.

Dynatrace brings unified visibility and control across these landscapes, while respecting data sovereignty and privacy mandates. Whether AI models are hosted in AWS, private cloud, edge devices, or diverse infrastructure environments, the platform provides consistent observability, governance, and policy enforcement. Another core strength of Dynatrace is automated remediation.

Using causal AI, the platform doesn't just alert teams to anomalies—it can trigger workflows, scale compute nodes, restart model APIs, or flag corrupt datasets autonomously. This turns AI observability into a closed-loop system where detection is followed by intelligent action. The platform's AI-powered root cause analysis provides a path to Agentic AI-powered automated remediation, preventative action, and continuous optimization.

The platform also aligns with emerging AI regulations by enabling explainability and auditability. Organizations can track which model versions were used in which context, who accessed what prompts, and what data was ingested, ready for AI audits, risk reviews, or forensic investigations.

Embedding AI Responsibly: Best Practices for Scale and Trust

Based on insights from global implementations and diverse enterprise environments, the following best practices can help organizations embed AI effectively, driving performance, trust, and business alignment at scale:

- **Integrate AI Observability from Day One.** Monitoring should not be an afterthought. Instrument models, APIs, and data pipelines from the prototype stage to detect performance or ethical risks early.
- **Link Model Output to Business KPIs.** Establish clear metrics for success—be it customer churn reduction, fraud detection accuracy, or agent productivity—and track them continuously.
- **Use Causal AI for Explainability.** Move beyond black-box inference. Leverage platforms that provide root-cause visibility into why a model behaved the way it did.
- **Adopt a Federated Yet Governed Architecture.** Encourage innovation at the edge, but standardize core services like observability, compliance checks, and deployment policies through a central Center of Excellence.
- **Secure the Prompt Layer.** Treat LLM prompts and completions as sensitive data. Monitor them for anomalies, injection attacks, and content governance.
- **Automate Model Remediation.** Use observability data to auto-trigger actions like model retraining, cache invalidation, or workload scaling, without manual intervention.
- **Ensure Hybrid Cloud Compatibility.** Design observability to work seamlessly across cloud and on-prem setups, respecting data residency and availability zone requirements.

By making AI performance observable, reliable, and explainable, Dynatrace enables enterprises to embed intelligence across the business—with confidence, compliance, and scale.

Improving DevOps and Application Engineering: From CI/CD Pipelines to Business-Aligned Platforms

DevOps is evolving into a business-aligned, intelligent, and resilient operating model.

Executive Summary

In today's hyper-connected digital economy, the role of DevOps has evolved from a technical function to a strategic enabler of innovation, resilience, and trust. What began as a movement to accelerate software delivery has now matured into a broader discipline—one that governs how enterprises build, operate, and continuously improve the digital systems at the heart of their business.

This chapter brings together insights from technology leaders across industries—financial services, insurance, automotive, and engineering services—who are reimagining DevOps as a system of accountability, not just automation. While continuous integration and delivery remain foundational, the new mandate is clear: DevOps must deliver visibility, security, and business relevance—at scale and in real time.

The complexity of modern enterprise environments is reshaping the DevOps agenda. Hybrid architectures are now the norm. Applications span cloud-native microservices, on-premise legacy platforms, third-party SaaS tools, and heavily regulated partner ecosystems. Code is deployed faster than ever—but so are misconfigurations, integration failures, and observability blind spots. In this reality, speed without context is risk, and automation without alignment can do more harm than good.

Panelists agreed that traditional DevOps tools and workflows—designed for static environments and siloed teams—are no longer enough. The next generation of DevOps must be intelligent, adaptable, and deeply integrated with the business. This means embedding observability earlier in the software lifecycle, tying alerts to user journeys, and prioritizing remediation based on business impact rather than just error rates or CPU metrics.

A recurring theme in the discussion was the need to shift-left—moving observability, compliance, and security into pre-production phases. Too often, organizations discover issues only after users are affected. By making telemetry, context, and RCA (Root Cause Analysis) available during build and test stages, enterprises can detect and prevent failures before they hit production. This reduces downtime, accelerates release cycles, and protects brand trust.

Equally important is the human dimension. DevOps maturity is not just a function of platform choices—it is driven by culture. Silos between developers, operations, and business teams remain a critical barrier. RCA becomes a blame game. Alerts flood inboxes but trigger no meaningful action. To overcome this, organizations are investing in shared dashboards, cross-functional CoEs, and joint incident reviews, so that all stakeholders see the same picture and act in concert.

Security and compliance are also front and center. With data privacy laws tightening and cloud footprints expanding, CIOs are embedding security telemetry, policy enforcement, and audit trails directly into CI/CD pipelines. This ensures that releases are not only fast, but also defensible—aligned with both internal policies and regulatory expectations.

Yet amid all this transformation, cost and scalability remain pressing concerns. Full-stack observability and intelligent automation require significant investment. Panelists emphasized the importance of targeted instrumentation—focusing efforts on high-risk, high-impact flows—rather than trying to monitor everything. This makes the economics of DevOps more sustainable, especially in industries with high transaction volumes or seasonal traffic spikes.

Ultimately, what's emerging is a new kind of DevOps—one that's not just about code, but about confidence. Confidence that systems will perform. That incidents will be caught early. That releases won't break compliance. And that when something does go wrong, teams will know why—and what to do next.

This chapter sets the stage for that transformation. It distills real-world lessons, strategic shifts, and implementation insights from leaders who are not only scaling DevOps, but reshaping it for the next decade. The message is clear: DevOps is no longer the backend of digital strategy. It is the execution engine of enterprise trust.

Key Drivers: What's Powering the Next Generation of DevOps

As enterprises confront the dual demands of speed and reliability, DevOps is undergoing a strategic evolution. No longer focused solely on CI/CD and release velocity, the new DevOps agenda is shaped by deeper imperatives: resilience, risk management, business alignment, and user trust. This section explores the key drivers shaping that transformation, as surfaced through the CIO panel's collective insights.

Shift-Left Observability: Catching Failures Before They Scale
Panelists agreed that the earlier observ-

> "Shift-left observability means catching failures before they scale into business disruptions."

ability is embedded in the development lifecycle, the greater its impact. Too many organizations only monitor production—by which point failures have already affected users, breached SLAs, or triggered compliance incidents. The next-gen DevOps mindset is about proactive detection during build, test, and pre-production phases. This shift reduces firefighting, shortens RCA cycles, and increases release confidence.

Full-Stack Telemetry and Context-Rich RCA

Modern environments are fragmented by design—spanning microservices, legacy applications, cloud-native workloads, and partner APIs. In this landscape, traditional alerting and log analysis fall short. CIOs are investing in unified, full-stack telemetry that captures not only logs and traces, but also user behavior, third-party latency, and journey-level signals. AI-powered root cause analysis is helping teams connect seemingly unrelated events—speeding up resolution and reducing false positives.

"From automation to orchestration, from tools to trust—that's the DevOps transformation."

DevOps with Business Context: Aligning Delivery with Impact

A recurring insight was the need to connect technical metrics to business outcomes. Panelists described how they are now mapping system events to customer journeys, transaction flows, and revenue drivers. When a service fails, teams don't just ask "What broke?"—they ask, "What was the user trying to do?" This business-aware DevOps helps prioritize fixes based on value, not just noise.

Embedded Security and Compliance in CI/CD Pipelines

As data regulations tighten and security threats grow, DevOps is becoming a frontline function in risk mitigation. Leaders are embedding security checks, policy enforcement, and audit logging directly into CI/CD pipelines. This "DevSecOps by default" model ensures every release is not just faster, but safer—meeting both internal controls and external compliance mandates.

Resilience as a Design Principle

Downtime is no longer measured in minutes—it's measured in lost trust. Enterprises are now treating resilience as a product feature, not a recovery tactic. That means designing for failure—through service isolation, automated fallback, dynamic scaling, and self-healing logic. Panelists described how intelligent auto-remediation, powered by rules or AI, is

reducing dependency on manual intervention during incidents.

Interoperability Across Legacy and Modern Systems
Legacy systems are still core to many enterprises, but integrating them into modern DevOps pipelines is a growing challenge. Panelists highlighted the need for adaptive instrumentation that can extract telemetry from systems not designed for observability—bridging the gap between COBOL and Kubernetes, mainframes and microservices.

Data Pipeline Scalability for Observability at Volume
High-frequency businesses—like capital markets or insurance—generate massive telemetry volumes. Without scalable architectures, observability can become a bottleneck. Leaders are investing in Kafka pipelines, non-SQL data lakes, and stream processors to handle real-time data ingestion, correlation, and long-term retention without sacrificing performance.

Intelligent Automation: From Alerting to Action
Automation is no longer a luxury—it's a necessity. But the bar is rising. Organizations are shifting from "auto-alerts" to contextual, auto-executing remediation routines. Whether it's scaling queues, clearing caches, or restarting services, automation now must include governance, risk thresholds, and decision logic.

DevOps Culture and Cross-Team Collaboration
Even the best tools fail in the absence of shared ownership. Panelists stressed the importance of cultural alignment—bringing developers, SREs, infrastructure, and business teams onto a common platform. CoEs, shared dashboards, and RCA rituals are helping break silos and foster a DevOps culture built on trust and transparency.

Observability ROI: Monitoring What Matters Most
Finally, cost discipline is a growing concern. Observability tools can be expensive—especially when priced by host, session, or data volume. CIOs are moving toward value-based instrumentation, focusing deep observability where business risk is high and using AI to model behavior where full coverage is impractical.

Together, these drivers reflect a profound shift: from DevOps as a toolchain to DevOps as a trust chain—spanning code, context, compliance, and customer experience.

Implementation Challenges: What's Slowing Down DevOps Maturity

While the promise of DevOps is well understood—faster releases, greater stability, and better collaboration—many enterprises continue to struggle with operationalizing that vision at scale. From legacy integration to organizational silos, the barriers to DevOps maturity are as much cultural and architectural as they are technical.

Based on the panel discussions, here are the key obstacles holding back the evolution of DevOps from pipeline efficiency to strategic enablement:

Fragmented Tooling and Data Silos
A common challenge is tool sprawl. Different teams use different platforms for monitoring, deployment, logging, and alerting—resulting in fragmented visibility and inconsistent RCA workflows. Without an integrated toolchain, teams operate with partial context and lose time reconciling data instead of resolving issues.

Lack of Business-Aligned RCA
While technical alerts may indicate infrastructure or application failures, they often lack business context—making it difficult to prioritize incidents or quantify impact. Teams can trace a CPU spike or an API timeout, but not whether it disrupted a high-value user journey or compliance-critical transaction. This disconnect erodes trust and delays resolution.

Observability Gaps in Legacy Systems and SaaS Platforms
Legacy systems remain opaque, and many third-party SaaS providers limit telemetry access. These blind spots make it difficult to build an end-to-end observability model. In industries like manufacturing or capital markets, legacy-heavy environments create observability debt that drags down incident response and DevOps performance.

Manual RCA and Reactive Incident Management
Despite investments in observability, many organizations still rely on manual log scrapes, siloed tickets, and war-room escalations. RCA becomes a time-consuming negotiation rather than an intelligent process. Alert storms often lead to confusion, not clarity—wasting cycles and delaying recovery.

Poor Integration of Security and Compliance into DevOps
In many enterprises, security controls are bolted on at the end of the pipeline rather than integrated into the CI/CD flow. This leads to delays, rework, or release rollbacks when vulnerabilities are discovered late. Regulatory controls like data retention or audit trails are inconsistently applied, creating risk exposure.

Alert Fatigue and Noisy Dashboards
A high volume of alerts—many of them redundant or low-priority—overwhelms teams. Without proper alert tuning or correlation logic, meaningful signals are buried under noise. This leads to desensitization, delayed action, and missed critical incidents.

Resource Overheads of Observability
As observability scales, so does its infrastructure footprint. High-volume telemetry collection imposes processing, storage, and cost overheads. Without efficient data pipelines or retention policies, observability can become a bottleneck—affecting performance, budget, and usability.

Organizational Silos and Misaligned Metrics
Dev, Ops, Infra, Security, and Business teams often operate with different goals, tools, and vocabularies. This leads to handoff issues, unclear ownership, and misaligned incident priorities. Without shared dashboards or incident playbooks, coordination suffers.

Skills Gaps and Change Resistance
Adopting DevOps practices requires cross-functional skills—in automation, observability, cloud-native architectures, and agile governance. Many teams lack the expertise to fully leverage new platforms. Moreover, legacy mindsets often resist the transparency and ownership that modern DevOps demands.

Significant Governance and Cost Management Gaps
Finally, without strong governance, DevOps investments can spiral—tools proliferate, automation gets duplicated, and telemetry becomes unfocused. CIOs emphasized the need for value-based governance, where instrumentation, automation, and observability are aligned with risk and return, not just ambition.

These challenges don't diminish the value of DevOps—they underscore the complexity of making it real. To move forward, enterprises must treat DevOps not just as a delivery pipeline, but as a disciplined system of collaboration, insight, and trust.

Observability Gaps in Legacy Systems and SaaS Platforms
Legacy systems remain opaque, especially those built on older programming languages like C++ or in environments where the source code is no longer available. Instrumenting these systems for observability is often impractical or impossible. Similarly, many SaaS platforms offer limited or no internal telemetry, creating blind spots in end-to-end monitoring. This lack of visibility severely hampers RCA and proactive remediation efforts.

Cost of Observability in Parallel Environments
While observability is critical, scaling it across environments—development, staging, UAT, and production—can be prohibitively expensive. Tool licensing models based on host count, data volume, or session metrics make full coverage financially unsustainable. As a result, organizations often under-instrument in pre-production environments, where early detection would have the most impact, leading to higher costs downstream in the form of production incidents.

> "Speed without context is a risk. Automation without alignment creates chaos."

From Automation to Intelligent, Trusted Delivery

DevOps is entering a new phase—one defined not just by how quickly software is shipped, but by how reliably it performs, how transparently it operates, and how intelligently it adapts. The goal is no longer speed for its own sake. It is trusted delivery: systems that are observable, secure, business-aware, and capable of responding in real time to user and market needs.

Based on the panel discussions, here's how forward-looking enterprises are evolving their DevOps practices from tactical automation to intelligent, outcome-oriented systems:

Build for Observability from the Ground Up
Rather than retrofitting monitoring tools post-deployment, mature organizations are now embedding observability into the development lifecycle. This means instrumenting code as it's written, mapping services to business transactions, and using shared metrics across development, operations, and business teams.

Some are adopting experience-level objectives (XLOs)—metrics that link system health directly to user or customer impact. This reframes observability from a technical overhead to a business-critical function.

Unify Toolchains and Context Across the Stack
Enterprises are working to eliminate tool sprawl by consolidating their DevOps environments around integrated platforms—capable of end-to-end telemetry, deployment, automation, and governance. More importantly, they are investing in cross-functional dashboards that present unified views of infrastructure, application behavior, user journeys, and security posture.

The goal is to break down silos—not just between tools, but between teams and their understanding of what's happening.

Automate RCA and Enable Closed-Loop Remediation
The shift from reactive to intelligent DevOps hinges on automated root cause analy-

sis and remediation. Leading organizations are deploying AI/ML models to detect anomalies, correlate metrics, and propose or even trigger fixes—whether it's restarting a container, scaling a service, or routing traffic dynamically.

This closed-loop automation reduces mean time to resolution (MTTR), minimizes escalation churn, and frees up engineering time for innovation.

Prioritize What to Observe—and What to Ignore
As telemetry volumes grow, smart organizations are resisting the urge to "monitor everything." Instead, they're practicing selective instrumentation—focusing visibility on critical user journeys, compliance-sensitive flows, and high-risk APIs.

This not only optimizes observability costs but sharpens operational focus. When teams know where the business risk lies, they can tune alerts, SLOs, and response protocols more effectively.

Make DevOps Business-Aware and User-Centric
DevOps teams are increasingly being asked to explain system behavior in terms of business outcomes—missed SLAs, revenue loss, or compliance exposure. To meet this need, they're incorporating business telemetry into their pipelines: tagging logs by customer, mapping endpoints to journey stages, and using real-user monitoring (RUM) to trace outcomes back to technical events.

The result is more actionable insights—and a faster bridge between tech and business priorities.

Shift Compliance Left with Secure-by-Design Pipelines
Enterprises in regulated industries are embedding security and policy enforcement directly into their DevOps workflows. This includes automated scans, policy-as-code, SBOM validation, and audit trail generation—all as part of CI/CD.

By treating compliance as a shared responsibility, not a separate function, organizations can reduce risk without slowing delivery.

Invest in Skills and Shared Accountability
The most impactful DevOps transformations are as much cultural as they are technical. Leading organizations are creating Centers of Excellence, running cross-team workshops, and building shared incident review processes. They're upskilling teams not only in observability tools, but also in journey mapping, SLA design, and risk modeling.

In this model, DevOps becomes a shared discipline—bridging the gaps be-

"DevOps maturity isn't about doing more—it's about doing what matters."

tween development, operations, security, and business.

Engineer for Resilience, Not Just Recovery

Recovery is important—but resilience is better. Enterprises are now designing systems that anticipate and absorb failure, using patterns like circuit breakers, graceful degradation, and canary releases. This allows services to stay up—or fail safely—even under load or during outages.

By integrating these patterns into the DevOps lifecycle, organizations shift from reaction to readiness.

The shift to intelligent, trusted DevOps doesn't happen overnight. It requires architectural shifts, cultural change, and strategic alignment. But as the panelists showed, the organizations making these moves are not only delivering faster—they're delivering smarter, safer, and more sustainably.

From Pipelines to Platforms: DevOps as a Strategic Operating Model

For years, DevOps was viewed through the lens of tools, scripts, and release cycles—a tactical practice designed to bridge development and operations. But as digital systems become central to how enterprises deliver value, DevOps is no longer a backend function. It is fast becoming a strategic operating model—governing how organizations design, deliver, and sustain their digital products in an always-on world.

What emerged from the panel conversations was clear: the future of DevOps lies not in faster pipelines alone, but in the platformization of trust, observability, and intelligence. In this new paradigm, DevOps is not merely a set of workflows—it is the architecture of adaptability.

Enterprises are increasingly treating their DevOps stack as a platform for continuous decision-making, where data, context, and automation converge to drive business outcomes. This platform spans development, testing, deployment, observability, security, and governance. It integrates seamlessly with business telemetry, regulatory controls, and customer experience metrics. It supports experimentation—but with guardrails. Innovation—but with accountability.

This shift is not just technical—it is philosophical.

Where older models prized velocity, the new model values clarity and coherence. Teams no longer ask, "How quickly can we deploy?" They ask, "How confidently can we improve the system?" Success is measured not just in MTTR or uptime, but in reduced churn, better NPS, faster onboarding, and fewer regulatory escalations.

The panelists underscored a key idea: DevOps maturity is not about doing more. It's about doing what matters—consistently, observably, and securely. That means investing in telemetry not to flood dashboards, but to illuminate blind spots. It means automating RCA not to reduce headcount, but to scale decision-making. It means surfacing journey-level insights—not just system health—so teams can optimize for outcomes, not just availability.

This model also creates leverage. With the right DevOps foundation, enterprises can:

- **Launch features faster without increasing risk**
- **Catch regressions before they go live**
- **Improve audit readiness with embedded controls**
- **Empower developers with feedback loops that teach, not just alert**
- **Align engineering and business priorities through shared metrics**

Importantly, this model is extensible. As GenAI, adaptive security, and autonomous operations enter the mainstream, DevOps will serve as the integration layer for intelligence. It will monitor AI behavior, manage API dependencies, and govern how new code blends into regulated, customer-facing environments.

But this evolution requires commitment. Organizations must invest not just in platforms, but in people. They must build cross-functional fluency—where devs understand infra constraints, ops understand journey flows, and business leaders understand the impact of each build. They must design governance that guides without stifling, and choose metrics that reflect business intent—not just system load.

As the chapter concludes, one truth stands out: DevOps is no longer the engine room—it is the control room. It determines how reliably value is delivered, how quickly risk is mitigated, and how well digital systems adapt to changing market demands.

For enterprises seeking not just agility but resilience, not just releases but relevance, DevOps is the platform on which digital trust is built.

AUTHORS

Root cause analysis needs business mapping to deliver value

Ananth Subramanian
Sr. Executive Vice President & Head IT, Kotak AMC

Engineering Observability Needs Engineering Intelligence

Aravindan Raghavan
Global Head - Business Excellence, CISO & DPO, Quest Global

Shift Observability Left— Where Resilience Begins

Dr Pawan Kumar Sharma
Chief Information Security Officer, Tata Motors

When Systems Talk, Let Humans Listen Last

Gaurav Duggal
Senior Vice President - AI & IT, Jio Platforms

If Business Can't See It, It Doesn't Exist

Goutam Datta
Chief Information & Digital Officer, Bajaj Allianz Life Insurance

Observability isn't just a Tool—It's a way of life

Rohit Kilam
Chief Technology Officer, HDFC Life Insurance

Improving DevOps and Technology Engineering

Sampath Manickam
CTO - Technology Infrastructure, National Stock Exchange of India

Redefining DevOps with Intelligent Observability

Abhishek Sharma
Business Head, Dynatrace

Mithun Gangadhariah
Lead Solutions Engineer, Dynatrace

Root cause analysis needs business mapping to deliver value

Insights matter only when they reflect real-world business consequences.

"You can't fix what you don't understand—and what you haven't mapped. Observability must span tech and business, or it's just noise."

Ananth Subramanian
Sr. Executive Vice President & Head IT, Kotak AMC

At Kotak AMC, we operate in a world where milliseconds matter—and where one failed transaction could mean lost revenue, a missed NAV, or a customer lost to a competitor. In this environment, observability isn't just a DevOps concern—it's a business continuity imperative.

One of the biggest issues I've seen across organizations—and we've experienced it ourselves—is that technical observability without business context is incomplete. Yes, you can monitor infrastructure, trace APIs, and detect anomalies. But when something breaks, what does it mean for the business? Is it impacting a critical investor transaction? Is a distributor unable to complete a purchase? That's the layer we need to surface.

We've moved toward a model where user journey observability is stitched end-to-end—from the front-end interaction, through the middleware and API layers, all the way to the RTA and back-office systems. This is especially critical for us because many of our touchpoints span third-party services, SaaS applications, and partner systems that we don't fully control. If a single link in that chain is delayed or fails, the customer experience takes a hit—and our teams need to know exactly where and why.

We've also faced the challenge of RCA complexity. Often, issues manifest downstream, but the root cause lies upstream—in data mismatches, version misalignments, or middleware slowness. This is where AI-powered correlation across logs, traces, and events has helped, especially in pinning down problems that would otherwise require hours of manual effort and inter-team coordination.

And then there's the human layer. I believe strongly that technical support teams must understand the business flows they're safeguarding. A DevOps engineer might resolve an alert quickly—but without understanding how that alert maps to a failed fund switch or an unprocessed SIP, we miss the point. We've started sensitizing our support teams to the business processes, so they don't just triage issues—they interpret them.

On the governance side, cost is a very real concern. With observability platforms typically priced per host, session, or agent, coverage can't be blanket—it must be intentional. We're increasingly focusing observability where the risk is highest: transaction-heavy flows, partner API chains, and investor-facing channels. At the same time, we keep iterating on dashboards that serve both tech and business audiences—so that we're all looking at the same reality.

Ultimately, observability for us is about speed, precision, and alignment. Speed to detect and fix issues. Precision in knowing where and why. And alignment between IT and business is important so that RCA doesn't just explain the past—it prevents future disruptions too.

Engineering Observability Needs Engineering Intelligence

AI-driven observability must adapt to compliance, context, and complexity.

At Quest Global, we operate in a uniquely complex environment—serving engineering clients across aerospace, rail, automotive, and industrial sectors, each with its own toolchains, compliance needs, and architectural constraints. In such a world, observability isn't just a platform—it's a puzzle.

One of the toughest challenges we face is managing observability in siloed, heterogeneous environments, where each client may use a different set of platforms, security postures, and access restrictions. This makes traditional, one-stack-fits-all monitoring irrelevant. AI-driven RCA must be context-aware, capable of navigating the dimensional complexity of varied configurations and fragmented ownership.

Another dimension is regulatory compliance observability. In aerospace, for instance, design-time decisions must comply with industry-specific safety standards. If a compliance violation is caught late—say, at the prototype or pre-production stage—it can derail timelines and budgets. That's why we're looking at ways to embed compliance observability into DevOps pipelines, especially in domains where rework is expensive and delays are unacceptable.

We've also recognized the need to mine past RCAs for patterns. Often, recurring failures trace back to similar root causes—but the tribal knowledge isn't always codified. By layering AI over our observability data, we're building a system that proactively surfaces RCA patterns, enabling faster resolution and reducing cognitive load on engineering teams.

What's different in engineering services is that DevOps here isn't only about code-to-deploy—it's about code-to-compliance-to-client-handoff. Every touchpoint is potentially a handover to another organization or system outside our direct control. That makes dependency observability—across tools, APIs, data layers, and validation workflows—an absolute must.

We're also exploring automated anomaly detection that understands real-world signals—not just telemetry but also external triggers like customer escalations, audit outcomes, or regulatory updates. For us, observability must go beyond logs and metrics—it must ingest external context to remain relevant.

Ultimately, the goal isn't just uptime. It's engineering assurance—knowing that what we build, monitor, and deliver stands up to the quality and compliance expectations of the world's most demanding industries. And for that, observability has to be intelligent, contextual, and engineered to adapt.

"In a multi-domain, multi-client world, no RCA is one-size-fits-all. Observability must adapt to the complexity of engineering services."

Aravindan Raghavan
Global Head - Business Excellence, CISO & DPO, Quest Global

Shift Observability Left—Where Resilience Begins

Observability must start pre-production to prevent failure and ensure trust.

"You can't afford to wait for production to find problems. In regulated, engineering-driven businesses, pre-production is where observability must lead."

Dr Pawan Kumar Sharma
Chief Information Security Officer,
Tata Motors

In the automotive and manufacturing world, system complexity is not just a function of software—it's tightly intertwined with hardware, compliance mandates, and decades-old legacy systems. I've seen how critical it is to embed observability as early as possible in the software lifecycle.

Too often, organizations treat observability as a production concern. The reality is, by the time you're putting out fires in production, the damage—operational, reputational, even regulatory—is already done. That's why my mantra is simple: shift observability left.

This means designing an observability strategy that starts in pre-production. In a manufacturing environment, where software increasingly drives core operations, connected vehicle platforms, and IoT-linked diagnostics, the margin for error is razor-thin. Whether it's a performance issue, an integration failure, or a downstream latency choke, catching it early is non-negotiable.

One area that remains a challenge is working with fragmented tooling—tools that monitor specific stacks or layers without the context to tie events back to business impact. Root cause analysis suffers when signals are scattered across multiple systems — only to find the issue stemmed from a legacy hardcoded delay, like a 30-second thread sleep which may take weeks to diagnose. That kind of inefficiency is what modern observability can eliminate, if applied correctly.

Another real-world pressure is justifying ROI for observability investments. Whether you're on-prem, hybrid, or SaaS-first, costs add up fast when licensing, infrastructure, and skilled resource needs are factored in. The key is understanding how observability reduces unplanned downtime, accelerates root cause resolution, and supports compliance and risk mitigation.

Let's also not forget cybersecurity and compliance. As observability platforms evolve, they must interlink with security telemetry—especially in a post-DPDP, zero-trust world.

One of my biggest lessons? Observability isn't about monitoring more things—it's about guessing less. The less time people spend chasing symptoms, the more they focus on outcomes. And when you apply that mindset early in the lifecycle, you don't just prevent incidents—you build confidence.

In high-performance, compliance-bound environments like ours, experience begins not at the front end, but at the far left of your pipeline. And that's exactly where your observability must begin too.

When Systems Talk, Let Humans Listen Last

AI resolves routine DevOps issues; humans handle meaningful escalations only.

At Reliance Jio Platforms, scale is not just a characteristic—it's a defining constraint. With a digital ecosystem that spans hundreds of millions of users, thousands of services, and billions of daily telemetry signals, DevOps cannot function as a manual, reactive model. For us, the only sustainable path forward is one where observability is not just intelligent, but automated—where systems handle the routine, and people focus on the exceptional.

Jio Platforms has built a highly evolved approach to incident management, driven by AI and issue categorization that prioritizes autonomy over intervention. Incidents are dynamically classified based on pattern recognition, persistence, and impact. The vast majority of what would traditionally be considered "alerts"—momentary latency spikes, packet drops, retry attempts—are absorbed and resolved by automation before they ever surface on a dashboard.

Where patterns persist, or when impact expands in scope or severity, the system escalates intelligently—first to enriched diagnostics and correlation layers, and only then to human teams. In this model, an engineer rarely touches an incident unless AI has already filtered, suppressed, enriched, and attempted remediation. This automation strategy is powered in part by ATK—a platform we use to encode and orchestrate its first-response reflexes. Through ATK, predefined workflows are triggered based on observed conditions, e.g., rolling back faulty releases, adjusting runtime configurations, isolating noisy services, or rebalancing infrastructure loads. These actions are not only automated but traceable, auditable, and dynamically updated based on operational learnings.

We leverage ATK as more than a remediation engine—it's an evolving DevOps brain. With each incident, the system learns, adapts, and expands its scope of autonomy. Over time, engineers shift from firefighting to curating the conditions under which the system decides to act. It's not just that we fix faster—we fix smarter, and we know exactly when to escalate and when not to.

Another breakthrough in Jio's model is its emphasis on client-side observability. Many system degradations don't originate in the backend. They start at the edge—within the browser, inside a mobile cache, or during a device-network handoff. These failures rarely show up in logs or server traces, but they distort the user experience. To close this gap, Jio instruments the client environment, analyzes behavioral anomalies, and builds baselines of "normal" user flows. When those flows deviate. the system responds as if it were a backend incident. We don't wait for logs to tell us there's a problem—We let user behavior tell us.

By combining automated RCA, self-healing actions, client-side intelligence, and platformized decision logic, we have turned DevOps into an architecture of trust. Engineers' time is spent improving system design, refining suppression logic, and working on issues that truly require human reasoning.

The future of DevOps is not about adding more dashboards or writing better scripts. It's about building systems that can sense, act, and learn with minimal handholding.

"AI handles our first line of defense. If an alert still needs a person, it's already escalated beyond the trivial."

Gaurav Duggal
Senior Vice President - AI & IT,
Jio Platforms

If Business Can't See It, It Doesn't Exist

Democratized observability drives trust, speed, and outcome-focused collaboration.

"DevOps is about faster releases—shared accountability. When business teams see the same metrics as tech, silos collapse and trust scales."

Goutam Datta
Chief Information & Digital Officer,
Bajaj Allianz Life Insurance

At Bajaj Allianz Life, we've long moved past the idea that digital experience is about UI or page speed. For us, experience is measured by what the customer doesn't notice—no delays, no disruptions, no confusion. That invisible smoothness is what observability must enable. But here's the reality: for a long time, observability was an IT-only concern. And that created more problems than it solved.

We had a classic disconnect: issues were being fixed by tech, but business teams never saw what was going wrong—or being made right. So even when we solved problems fast, business was left out of the loop. That creates friction, mistrust, and confusion during every incident review or RCA.

We changed that by making core process observability visible to business stakeholders. We started surfacing experience dashboards for them—showing real-time signals on transaction drops, conversion lag, queue buildups, and system responsiveness. Suddenly, business leaders didn't have to ask us what went wrong. They could see it themselves.

This created a new kind of partnership. Not only were we resolving issues faster, we were now collaborating on which metrics mattered most. And more importantly, we were learning that the most valuable observability doesn't start with "What went wrong?" It starts with "What outcome are we trying to protect?"

Of course, all this is easier said than done. We work in a complex hybrid ecosystem—legacy systems, new-age APIs, SaaS platforms, partner integrations, and highly variable customer journeys. And observability must stitch across it all. Without unified context, signals become noise, and alerts become distractions.

Another challenge we face is diagnosing issues at the edge—at the partner layer, in B2B2C journeys, or in hybrid mobile-desktop flows. When a distributor or agent can't complete a transaction, we need to know why—immediately, and without relying on logs we don't control.

We're also investing in dashboards built for journey health, not just system uptime. And we're moving toward self-healing patterns wherever we can—auto-scaling queues, clearing stale sessions, restarting hung services—before the user even notices a glitch.

But perhaps the biggest learning has been cultural: observability must be democratized. If the only people who understand the system are in the NOC, you haven't built resilience—you've created bottlenecks. The more we bridge the view between tech and business, the more proactive, accountable, and aligned we become.

When everyone can see what's happening, you don't need to explain your SLA—you just deliver it.

Observability isn't just a Tool—It's a way of life

True observability unites systems, silos, and teams—before failure hits.

Most IT leaders are betting big on multi-cloud , SaaS services, and microservice architectures to achieve scale.

The rapid adoption of these technologies poses a unique set of challenges.

- Today's systems are modular and scattered across cloud(s) and SaaS services. The SaaS services used in CRM, marketing, and claim stacks are totally blackboxes that offer very little to no telemetry data for observability and management. One of the challenges is the visibility gap created by SaaS applications.
- All systems work flawlessly in isolation with their own infra monitoring, logging, and application performance dashboards, but downhill descent starts when you try to integrate them. The integration points between applications create huge blind spots, and when something breaks, since you can't fix what you can't see, SLA breaches happen while teams play the guessing game, as no one knows the whole story.
- The lost time due to the delay in remediation results in frustration and finger-pointing among IT and business teams.
- The bug that creates havoc in production originate in lower environments and could have been caught during pre-production.

As the CTO of HDFC Life, it is my responsibility to earn and maintain the trust of millions of our customers.

What we're building in HDFC Life is a "single pane of glass" observability platform.

It is our one-stop solution to all our observability needs. It not only stitches the end to observability across all logs, traces, and metrics across our entire hybrid landscape but also effectively overlays it with business context.

When a transaction fails anywhere across the landscape, It alarms the L1 team of the occurrence of the incident.

The L1 then has a look at a single dashboard and check what has failed and where (whether it was a frontend latency issue, a payment gateway timeout, or a database pool exhaustion triggered by a faulty code commit).

As a wise man once said, the more you sweat during peace, the less you bleed in war.

We are building a culture where hard work and trust among peers are rewarded.

Since observability is extended to non production environments, there are no defensive RCAs , no drama, or pointing fingers

Any misbehaviour of service is caught and then collaboration and ownership kicks in to fix it before it is live in production.

"Your stack may be modern, but if your teams are stuck in silos and RCA takes days, you're not doing DevOps—you're doing patchwork."

Rohit Kilam
Chief Technology Officer,
HDFC Life Insurance

Improving DevOps and Technology Engineering

DevOps at scale demands predictive observability, precision, and resilience.

"When the scale is billions of transactions a day, DevOps must evolve from agile delivery to real-time, predictive orchestration and fail-safe precision."

Sampath Manickam
Chief Technology Officer - Technology Infrastructure, National Stock Exchange of India

In the world of live streaming and ultra-low latency platforms, speed isn't a luxury—it's the baseline. Every transaction, every log entry, every deployment can ripple through millions of user experiences even at micro-seconds. At this level of operation, DevOps isn't just about agility or automation, It becomes the cornerstone of precision, predictability and resilience powering platforms built for scale, trust and real-time impact.

In such environments, observability becomes the nervous system of technology platforms. It must be embedded into the architecture from design to deployment. The standard dashboards and rule-based alerts don't suffice when workloads are bursty, throughput is extreme, and latency budgets are measured in microseconds. These systems are engineered to detect anomalies, initiate autonomous correction, and feed real-time intelligence back into operational pipelines—long before the end user is ever aware.

Telemetry is foundational, but high-scale environments generate overwhelming volumes of it. Logs, metrics, and traces can become clutter unless ingestion and storage are engineered with surgical precision. Drawing from my experience in serving multiple leadership roles across mission-critical organizations like telecom, OTT and BFSI industries, I've led teams to architect real-time pipelines, layered data streams, and telemetry strategies that eliminate noise. Because at this scale, observability isn't just about seeing more—it's about elevating clarity at scale

Self-healing Is the new standard as downtime isn't tolerated; remediation must happen in less than milliseconds. With rule-based AI-augmented frameworks, edge systems now reallocate memory, tune buffers, and reroute network loads on their own—no human intervention is needed. Stability is no longer reactive—it's built into the fabric of our platforms. Resilience isn't an outcome. It's a prerequisite.

Modern technology ecosystems span monoliths, microservices, legacy systems, and third-party APIs—making full-stack observability a critical design principle. Standard monitoring approaches often fall short in such diverse environments. Leading platforms now adopt custom instrumentation, deep log correlation, and open standards like OpenTelemetry to ensure every integration is transparent, every signal actionable, and every byte traceable. This approach transforms observability into a strategic layer of resilience, trust, and operational clarity.

DevOps is Governance Architecture at planetary scale. DevOps must be governed like core infrastructure. Every change, every deployment, every rollback must be traceable, reversible, and explainable; and configuration drift is immediately detected. In a highly regulated environments, operational events are treated not merely as logs but as auditable artifacts. This governance-first approach ensures that trust and traceability are engineered into the platform, standing alongside throughput as non-negotiable pillars of reliability.

Redefining DevOps with Intelligent Observability

DevOps teams need AI-driven observability for enabling faster releases, unified telemetry, proactive remediation, and business-aligned engineering outcomes.

The software-driven enterprise is under relentless pressure to innovate faster, release more frequently, and deliver uncompromised quality and security across every digital touchpoint. In this environment, DevOps has become the default operating model for agile delivery, but the complexity of cloud-native architectures, containerized applications, hybrid environments, and third-party integrations is testing its limits.

Modern DevOps now requires more than just automation pipelines and version control. It demands real-time feedback loops, intelligent observability, and secure-by-design practices across the software delivery lifecycle.

The shift-left movement—bringing testing, security, and performance considerations into development earlier—is gaining traction, but organizations struggle to implement it holistically. Meanwhile, AI and machine learning are being looked at as force multipliers —capable of accelerating root-cause analysis, improving release velocity, and enabling smart automation. However, AI maturity varies widely, and many enterprises still operate with fragmented visibility and reactive tooling.

Against this backdrop, observability is evolving from a backend function into a strategic DevOps enabler. It is not just about identifying what went wrong after a release, but about creating a live control system that guides design, development, deployment, and operations in near-real time.

> "We're helping teams shift left with observability built into the pipeline, so issues are caught before they ever reach production."

Mithun Gangadharaiah
Lead Solutions Engineer,
Dynatrace

Breaking Through Complexity: The Leadership View

CIOs and DevOps leaders face several converging challenges as they scale digital delivery:

- **Production Blind Spots.** Complex distributed systems often mask the

Abhishek Sharma
Business Head – Strategic,
Dynatrace

"A tool is only as powerful as its adoption is. We see mature large enterprises taking major steps here, such as setting up centers of excellence (CoEs) for observability and building a culture of democratizing the monitoring across business-tech functions. This is accelerating business outcomes by better leveraging a powerful monitoring strategy with a world-class observability platform."

root causes of failures, leading to long mean time to resolution (MTTR) and finger-pointing between infrastructure, development, and application teams.

- **Tool Sprawl and Data Fragmentation.** Organizations often use multiple disconnected monitoring tools across environments—on-prem, cloud, SaaS, LLMs, Agents, and APIs—making unified insights elusive and slowing incident response.

- **Legacy Constraints in a Modern Stack**. Engineering teams operate across systems built in different generations—legacy systems and microservices running alongside Kubernetes and microservices. Observability must bridge these layers seamlessly.

- **Delayed Detection and Costly Fixes.** Issues that could have been caught during pre-production testing often emerge in production, leading to customer impact, revenue loss, and firefighting.

- **Cultural Silos and Adoption Barriers.** While DevOps practices aim to break down silos, many organizations still have limited collaboration between developers, site reliability engineers (SREs), quality assurance (QA), and security teams. Tools are underutilized due to a lack of contextual relevance or skill gaps.

- **Security and Compliance Complexity.** As software supply chains expand and vulnerabilities evolve, integrating security into DevOps pipelines without slowing them down has become a significant operational and governance challenge.

- **Limited Time to Innovate.** Many teams spend disproportionate time triaging incidents and chasing manual fixes, leaving little bandwidth for innovation and long-term value creation.

Converging Insights, Accelerating Value with Dynatrace

Dynatrace provides a comprehensive, AI-powered observability platform that addresses these challenges head-on—enabling engineering teams to move from reactive troubleshooting to proactive, business-aligned DevOps. The importance of unified telemetry cannot be emphasized enough.

Dynatrace stitches together logs, metrics, traces, topology, and user experience into a single context-rich data model. This allows organizations to eliminate blind spots across production, pre-prod, and legacy environments—while keeping agent overhead minimal through intelligent placement and auto-discovery.

In DevOps pipelines, Dynatrace supports shift-left observability by embedding visibil-

ity directly into CI/CD stages. Faulty builds or degraded services are flagged in pre-production, enabling rollback or correction before customer-facing impact. This helps reduce release anxiety and accelerates deployment cycles with confidence.

Dynatrace's AI capabilities are a key differentiator. Unlike traditional tools that rely on static thresholds or brute-force alerting, Dynatrace's deterministic AI understands the underlying system topology and dependencies. This enables precise root-cause identification within seconds—even in highly dynamic environments.

The platform also plays a critical role in bridging business and engineering. Dynatrace links technical metrics with service-level objectives (SLOs), conversion rates, failure costs, and user impact—enabling teams to prioritize actions based on business value. An efficient and scalable DevOps approach—enabled by Dynatrace's automation and intelligence—frees up valuable time for innovation. Teams spend less effort performing manual tasks and more time delivering value, ultimately helping businesses exceed customer expectations and gain a competitive advantage.

By integrating observability, automation, and AI into a highly cohesive DevOps toolchain, Dynatrace helps organizations improve not just speed of delivery, but also code quality and engineering throughput. It aligns to innovate faster while maintaining high trust and low risk. Dynatrace also supports the evolution of a more productive and collaborative culture.

By eliminating silos and providing a single source of truth across development and operations, the platform enhances communication and shared accountability. This builds a high-trust environment that enables faster resolution, more agile decision-making, and continuous improvement.

Security is built into the platform through runtime vulnerability analysis, SBOM tracking, and log-based data sensitivity detection. This observability supports compliance with regulatory frameworks like GDPR, DPDP, and ISO, without adding friction to development workflows.

Finally, Dynatrace supports both SaaS and on-prem deployment models, providing flexibility for enterprises with strict data residency or control requirements. Whether an organization is a cloud-native startup or a heavily regulated BFSI major, Dynatrace adapts to the operating model with consistency and scale.

Foundations of DevOps Success

From Dynatrace's field experience across global and Indian enterprises, the following best practices consistently emerge as drivers of DevOps excellence:

- **Instrument Pre-production Environments.** Avoid production surprises by monitoring lower environments, such as user acceptance testing (UAT) and staging, with the same rigor as live systems. Many critical issues surface earlier, but only if appropriately observed.
- **Unify Telemetry into a Single Pane.** Replace fragmented dashboards with a centralized observability layer. Integrate across CI/CD, infrastructure, application, and business metrics. Automate SLO Validation and Governance. Use Dynatrace to enforce SLOs at each pipeline stage—Automate approvals, rollbacks, and alerts based on predefined business and risk thresholds.
- **Drive Cross-team Adoption.** Build a center of excellence (CoE) for observability. Include developers, ops, QA, and security stakeholders. Invest in enablement and role-specific dashboards.
- **Prioritize Business Context in Root-Cause Analysis.** Ensure root-cause analysis highlights impact on users, transactions, or SLAs — not just technical faults. This helps align teams and decision-making.
- **Embed Security into Observability.** Use observability data to monitor real-time vulnerability exposure, misconfigurations, and sensitive data usage, and automate reporting for compliance.
- **Plan for Scale and Retention.** Choose a data architecture that supports high ingestion volume, long-term retention, and flexible analytics without compromising cost or performance.

With Dynatrace, observability becomes a strategic foundation for modern DevOps, enabling faster releases, smarter operations, and a direct line of sight between engineering outcomes and business performance.

Business Observability for the Intelligent Enterprise: From Monitoring Signals to Strategic Insights

To support tomorrow's digital-first, AI-led enterprises, observability must be embedded with a unified, strategic vision.

Executive Summary

As businesses accelerate their adoption of cloud-native architectures, AI workloads, and real-time digital services, observability has emerged as a strategic enabler of trust, performance, and resilience. Static monitoring no longer suffices in today's dynamic, hybrid environments. Enterprises must move beyond fragmented alerts and siloed dashboards to adopt a more unified, intelligent, and outcome-driven observability model.

India's digital landscape presents unique complexity: legacy systems still power critical operations while modern digital applications demand agility, interoperability, and continuous availability. In this hybrid ecosystem, identifying root causes and preventing disruption requires a real-time view that spans infrastructure, application code, user behavior, and business impact.

Observability, when implemented effectively, enables this view. It collects rich telemetry (logs, metrics, traces, and beyond), stitches them into meaningful signals, and supports early detection, faster resolution, and continuous optimization. More importantly, it allows IT and business teams to ask not just what went wrong, but what it affected—and what to do next.

Today, observability is no longer just a toolset. It's becoming a board-level priority, directly linked to customer experience, revenue protection, compliance, and operational efficiency. CIOs now view observability as the nervous system of the digital enterprise: embedded, intelligent, and proactive.

This chapter distills perspectives from three expert panels of CIOs and digital leaders across India, synthesizing their views into a playbook for building business-centric observability in the AI-powered enterprise.

Key Drivers: What's Powering the Rise of Observability

Based on the panel discussions, the following are the primary drivers driving observability adoption in modern enterprises:

Digital-first business environment is raising the bar on experience.
The shift toward digital-first operations has fundamentally transformed how organizations conduct business. In highly competitive and saturated markets, delivering exceptional digital experiences is no longer optional, it is a key differentiator. This shift has created a pressing need for end-to-end performance visibility across all customer-facing touchpoints.

Customer trust in digital platforms has become a measurable business asset. To protect that trust, organizations must adopt proactive approaches to issue detection and resolution. This demands a unified observability strategy that seamlessly covers web, mobile, APIs, and backend systems—ensuring that customer experiences remain seamless, consistent, and high-performing across every digital channel.

Technological complexity is compounding observability challenges.
The exponential growth in system complexity has introduced observability challenges that traditional monitoring approaches are ill-equipped to handle. For instance, the shift from monolithic applications to distributed microservices has dramatically increased the number of components that must be monitored and correlated. Additionally, organizations operating across multiple cloud providers and on-premises environments require unified visibility into these disparate systems.

As AI becomes increasingly embedded into core business processes, organizations also need new monitoring capabilities to evaluate AI model performance, data integrity, and business impact. The coexistence of legacy applications with modern cloud-native systems further complicates observability, demanding monitoring strategies that can span multiple technology generations.

"In today's hybrid ecosystem, identifying root causes and preventing disruption requires a real-time view that spans infrastructure, application code, user behavior, and business impact."

Moreover, the proliferation of APIs as critical business interfaces necessitates comprehensive monitoring of third-party dependencies and integration points to ensure reliability, performance, and security across the digital ecosystem.

Operational efficiency goals are driving investment.
Organizations are adopting observability to achieve operational excellence and reduce IT operational costs. Business demands for faster issue resolution are creating pressure to invest in tools that can quickly identify root causes across complex systems. Mean Time to Resolution (MTTR) has become a key performance indicator, and observability platforms can have a direct impact.

The shift from reactive to predictive maintenance requires analytics that can identify potential failures before they affect customers. This turns IT into a strategic enabler, capable of preventing disruption. Observability also supports cloud cost optimization and infrastructure efficiency. This is a major concern as cloud spending grows and CFOs push for return on tech investments.

Reducing manual intervention through intelligent alerting and automated remediation frees skilled staff to focus on innovation. Given the shortage of experienced talent, observability helps teams do more with less.

Regulatory and compliance requirements are expanding.
The complexity of digital operations is mirrored by a growing wave of regulatory requirements and industry standards, making observability a compliance necessity.

From enforcing SLAs and performance metrics to fulfilling audit trail requirements and supporting risk management frameworks, observability is now a tool for governance. It also helps enterprises adhere to data protection regulations like GDPR, CCPA, and India's DPDP Act by enabling real-time user activity tracking, anomaly detection, and proactive threat mitigation.

In this context, observability is not just about performance—it is fundamental to enterprise compliance, security, and risk posture.

Business leaders want better tech-business visibility.
C-suite executives increasingly expect IT teams to correlate system health with business outcomes. It's no longer sufficient to report uptime or technical metrics—executives want to understand how technology failures impact revenue, customer satisfaction, or operations.

This has led to a growing demand for business impact correlation, where observability tools can translate system performance into business-relevant insights. Predictive insights are also gaining traction, with leaders seeking early warning systems to prevent disruptions.

Boards and risk committees now require regular reporting on IT resilience, system performance, and business continuity—further cementing observability as an enterprise-level concern.

Vendor ecosystem maturity is accelerating adoption.
The observability vendor ecosystem has matured significantly. Unified platforms are replacing fragmented tools, reducing complexity and improving ROI.

Cloud-native and SaaS-based solutions have lowered the barriers to implementation by eliminating infrastructure requirements and offering scalability, faster deployment, and continuous updates. Built-in AI and ML capabilities are enabling teams to detect patterns and generate insights with less manual effort.

Open-source solutions have also gained momentum, offering cost-effective, customizable tools for organizations with budget constraints or niche requirements. Vendor competition has improved both innovation and pricing.

Observability powers AIOps and DevSecOps evolution.
As companies transition to cloud-native systems, AI is being used to automate DevSecOps processes. Observability feeds these systems with clean, contextual data, enabling automation of monitoring, testing, SLA tracking, security, and incident response.

It plays a pivotal role in building intelligent IT operations, supporting smarter decision-making and reducing manual oversight across the stack.

Agile and DevOps transformation requires integrated observability.
The rise of DevOps and agile methodologies has redefined observability needs. With frequent releases and CI/CD pipelines, teams need visibility during development, testing, and deployment, and not just in production.

> "AI and cloud platforms generate massive telemetry volumes—from serverless functions and Kubernetes pods to model inference and drift data. Without proper observability, telemetry costs spiral out of control."

"Shift-left" observability is gaining momentum, embedding monitoring early in the development cycle. Site Reliability Engineering (SRE) practices require real-time alerting tied to service-level objectives (SLOs), while agile teams rely on continuous feedback loops to refine features and performance.

Cross-functional collaboration depends on shared visibility, making observability a unifying platform for development, operations, and business teams.

Implementation Challenges: What's Slowing Down

Observability Maturity

Panelists highlighted several persistent barriers that prevent organizations from maximizing the value of observability.

Misalignment with business context limits strategic value
As cloud and AI adoption scales, many infrastructure and application teams continue to track low-level technical KPIs without aligning them to business outcomes. This limits strategic value. In AI-driven environments, where model performance, latency, or drift can impact revenue or compliance, observability must connect technical signals with business metrics to remain relevant.

The ground-level data disconnect
Cloud architectures often abstract infrastructure layers, while AI models rely heavily on high-quality, real-time data. Without the ability to go deep that is ingesting and integrating data from the ground up, enterprises risk corrupting downstream analytics and predictions. Even the most advanced observability tools fail if data ingestion points lack integrity and visibility. Closing the gap between high-level oversight and ground-level data fidelity is crucial.

Tool sprawl and overlapping dashboards create noise
Cloud-native teams often deploy multiple observability tools across workloads, each with its own dashboards and alert mechanisms. Add AI pipelines to this mix, and the noise multiplies. Redundant data streams and uncoordinated alerts create confusion, reduce responsiveness, and obscure true root causes. A unified observability fabric is essential to declutter and centralize insights.

False positives erode trust in alerting systems
Dynamic cloud environments and self-learning AI models constantly shift baselines. Without adaptive alerting and intelligent thresholds, teams face a flood of false positives. This alert fatigue reduces trust and delays critical responses. Observability must evolve with context, integrating AI-powered anomaly detection to remain effective.

Manual setup burden slows down innovation
Instrumenting new cloud services or AI pipelines often involves manual setup—configuring agents, telemetry, and monitoring logic. This slows down releases and burdens DevOps teams. As organizations embrace platform engineering and MLOps, observability must become plug-and-play, with automation and auto-instrumentation built in.

Pre-production blind spots increase deployment risks
Most observability tools begin monitoring after deployment. In AI/ML and cloud-na-

tive DevOps pipelines, this creates blind spots around model behavior, security, and performance pre-release. Bugs, bias, or instability may only be caught post-production. Embedding observability earlier in the lifecycle is critical for safe and resilient deployments.

Fragmented data silos delay root cause analysis
Cloud-based observability often separates logs, metrics, traces, and model telemetry across services and platforms. This fragmentation makes it difficult to correlate issues across layers, especially when troubleshooting AI performance degradation tied to infrastructure or data pipeline problems. Unified, cross-domain observability enables faster, more accurate root cause analysis.

High cost of telemetry without governance
AI and cloud platforms generate massive telemetry volumes—from serverless functions and Kubernetes pods to model inference and drift data. Without proper governance, telemetry costs spiral out of control. Organizations must define data collection boundaries, retention policies, and sampling strategies to ensure observability remains cost-effective and scalable.

Incomplete coverage results in critical blind spots
In the hybrid, multi-cloud AI ecosystem, any monitoring gaps—whether at the edge, in data pipelines, or within black-box ML models—can cause business-impacting disruptions. Observability must be comprehensive, covering everything from API endpoints and model behavior to the underlying cloud infrastructure, to deliver the reliability today's enterprises require.

The Way Forward: Observability into the Enterprise DNA

As modern enterprises evolve toward highly distributed, cloud-native, and AI-augmented environments, CIOs are redefining what they expect from observability platforms. It's no longer about basic visibility—it's about actionable intelligence, operational foresight, and alignment with strategic outcomes. Here's how CIO expectations are shifting:

Move from Reactive to autonomous operations
Today's CIOs want observability platforms to do more than raise red flags. The ambition is to create systems that can learn from past incidents, identify patterns in real time, and autonomously initiate remediation without human intervention.

Proactive and Preventive Interventions
Rather than waiting for incidents to occur, CIOs expect observability to act as an early warning system—flagging anomalies

> "By embedding observability into the fabric of enterprise systems and workflows, CIOs are laying the groundwork for resilient, autonomous, and experience-centric digital operations"

before they impact end users or business processes. The ability to anticipate known risks and take preventive action is a key.

Business-Centric Visibility and Impact
Observability must connect the dots between system behavior and business value. CIOs are demanding dashboards and insights that highlight how infrastructure health impacts transaction rates, revenue, customer satisfaction, and compliance.

Embedded Security and Continuous Risk Awareness
With attack surfaces expanding, CIOs expect observability platforms to provide real-time, contextual visibility into vulnerabilities and exposures—detecting threats like zero-day exploits the moment they emerge and pinpointing impacted services.

Developer-First Enablement
As software velocity increases, developers need deep, real-time visibility into live systems to debug, optimize, and secure code safely—without compromising uptime or user experience. Observability must integrate into CI/CD pipelines and workflows.

Monitoring the AI Layer Itself
With AI agents now embedded in customer-facing and operational systems, CIOs are asking: Are these agents behaving reliably, ethically, and cost-effectively? Observability must extend to AI operations, including detecting anomalies in bot behavior or drift in ML models.

From Insight to Impact: Observability as a Strategic Operating Layer

To meet these expectations, enterprises are re-architecting their observability strategies—integrating automation, AI, and cross-platform integration. While the journey is in the process, several tactical and strategic interventions are already underway:

Embedding Self-Healing Capabilities
Enterprises are automating responses to known, repeatable issues like temporary file cleanup, memory optimization, and dynamic resource scaling, laying the groundwork for truly autonomous operations.

Configurable, Context-Aware Workflows
Modern observability platforms are being equipped with decision-tree logic and conditional workflows that can respond to incidents based on context, triggering diagnostics, controlled restarts, or routing escalations intelligently.

Internal Platforms for Unified Observability
Many organizations are building internal, tool-agnostic platforms that unify telemetry (logs, metrics, traces) across hybrid environments. These platforms deliver persona-based insights to developers, SREs, and business stakeholders alike.

Signal-to-noise optimization
To reduce alert fatigue and improve focus, teams are refining telemetry pipelines by setting thresholds, eliminating redundancy, and ensuring only meaningful, actionable alerts reach operations teams.

Secure software development integration
Security observability is being integrated earlier into the development cycle through automated scans, memory diagnostics, and vulnerability detection—enabling faster, safer code releases.

AI-Augmented insights, not yet fully autonomous
While fully autonomous, agentic AI is still a future goal, many organizations are already using AI and ML to detect anomalies, predict failures, and enhance root cause analysis, making observability smarter and more predictive.

Focus on People and Process Maturity
Observability success doesn't depend on tools alone. Organizations are investing in people, upskilling teams in SRE, DevSecOps, and data literacy—while fostering a culture of shared accountability across roles.

Conclusion

Observability is becoming a strategic operating layer that connects technology to business outcomes in real time. For enterprises navigating the complexity of AI, cloud, and digital acceleration, it provides the intelligence to respond faster, act smarter, and operate with confidence.

By embedding observability into the fabric of enterprise systems and workflows, CIOs are laying the groundwork for resilient, autonomous, and experience-centric digital operations.

The journey is incremental, but the direction is clear: observability will define how the intelligent enterprise sees, understands, and evolves.

AUTHORS

Observability That Misses the Process Misses the Point

Aashish Kshetry
Vice President - Systems, Asian Paints

Observability Is a First Principle of Digital Trust

Abhijit Chakravarty
Executive VP - Networks & Cyber Security, Kotak Mahindra Bank

Observability Is the Next Frontier of Public Sector Efficiency

Ashwini Kumar Pandey
Chief Information Security Officer, Punjab National Bank

Observability Must Move from Reactive Logs to Real-Time Intelligence

Munish Blaggan
Head Technology - Customer Engagement Platform, ICICI Bank

Observability Is the Bridge Between IT, Business, and the Unknowns.

Naresh Choudhary
Senior Vice President, Infosys

Observability Is the Pulse of a Multi-Business Digital Core

Rejin Surendran
Global Chief Information Officer, Wipro Enterprises

Observability isn't a tool, it's a mindset

Sajith Chakkingal
Head of Global Technology Services, TMF Group

Observability Must Make Sense to the Business—Not Just to the Dashboard

Sudesh Puthran
Chief Technology Officer, CreditAccess Grameen

Enabling Business Observability at Enterprise Scale

Shakti Agarwal
Lead Solution Consultant, Dynatrace

Praveen Mahajan
Regional Director- India, Dynatrace

Observability That Misses the Process Misses the Point

True observability tracks business outcomes, not just system uptime metrics.

At Asian Paints, where B2B and B2C supply chains intertwine across a vast network of dealers, warehouses, and customer channels, observability must do more than track systems—it must understand outcomes. The real challenge isn't just technical—it's operational. It's about how well IT visibility maps to business performance.

In theory, we talk about full-stack observability. However, most tools still operate from a technology lens. From a business standpoint, observability must track the entire process—how long an invoice takes to process, when an order gets fulfilled, or where a transaction stalls across people and systems.

We need to emphasize the growing gap between technical and business observability. While APM platforms and infrastructure monitors provide valuable system-level insights, they often leave out key parts of the value chain—especially human workflows or legacy dependencies embedded deep in the supply chain. As a result, an application might be healthy, but the customer may still be stuck.

It is also important to flag the challenge of fast-changing business needs, especially in a consumer-centric company. The observability requirements of the business can evolve every week. And if the tools or telemetry pipelines can't keep up, people start creating shadow processes or informal diagnostics to fill the gaps.

That's where business observability needs to differentiate itself—not just by capturing more data, but by staying agile to business context. For example, Asian Paints may need one type of visibility during seasonal sales spikes, and another during logistics rollouts or ERP transitions. Static dashboards don't work in these scenarios. What's needed is adaptive, user-aligned observability that changes with the business.

There is promise in solutions that go beyond traditional IT metrics and start incorporating process telemetry—tracking not just system errors but completion times, handoffs, user actions, and exception patterns. These are the signals that indicate whether a process is working, not just whether a server is up.

At the same time, we need to be pragmatic about limitations. Legacy systems, varied tech stacks, and multiple stakeholder touchpoints make perfect visibility difficult to achieve. But the goal is not perfection—it's progressive alignment.

Ultimately, observability must answer questions business leaders actually care about: Are orders being fulfilled on time? Are partners experiencing transaction friction? Is customer satisfaction trending with service performance? If your observability doesn't help answer those questions, it's not business observability—it's just monitoring.

"Business observability isn't about uptime dashboards—it's about whether your order shipped, your invoice cleared, and your customer stayed."

Aashish Kshetry
Vice President - Systems,
Asian Paints

Observability Is a First Principle of Digital Trust

Observability drives digital trust by uniting resilience, security, and insight.

"Observability is not just a visibility function—it's a control plane for how we deliver, defend, and de-risk the digital enterprise."

Abhijit Chakravarty
Executive Vice President - Networks & Cyber Security, Kotak Mahindra Bank

At Kotak Mahindra Bank, we operate in a high-stakes, high-frequency environment where resilience is non-negotiable and latency is unforgiving. In this context, observability is not a post-incident analysis tool—it's a first principle of trust. It governs how we build, secure, and operate every layer of the digital stack—from network infrastructure and application workloads to transaction flows and cyber defense systems.

One of the biggest challenges I see is that observability is still too often treated as a back-office function—owned by a few engineers, understood only at the technical layer, and surfaced when something breaks. That's no longer enough. For us, observability has become a strategic discipline—as integral to reliability as it is to security.

What we're doing is building observability into the system at design time. It's not just about monitoring endpoints or application latency. It's about stitching together real-time, contextual telemetry across infrastructure, APIs, application logic, and user behavior—so we don't just detect anomalies, we understand them in the context of business workflows.

This matters especially in hybrid environments like ours. We've got systems on-prem, in private cloud, across SaaS, and exposed to customers through mobile and web. If a payment doesn't go through, I want to know: was it a network flap, an API timeout, a rate-limit event, or something malicious? And I want that answer in real time—with enough context to act, not guess.

Security is a key part of this. Observability feeds into our cybersecurity fabric—not just as log data for post-facto forensics, but as live signals for threat correlation, anomaly detection, and policy enforcement. For example, if a high-risk user suddenly initiates an unusual transaction pattern, we want that flagged not just in a SOC alert, but in our operational view of service health. Because sometimes, what looks like a system degradation is actually a security event in disguise.

We're also pushing for more integration between IT operations and business observability. It's not enough to know that CPU usage spiked. We want to know which customer journey that spike disrupted, what revenue it affected, and whether it created regulatory exposure. That level of correlation requires more than tooling—it requires a mindset shift across teams.

Of course, challenges remain. Observability at scale generates enormous data volumes. We're constantly tuning what we collect, how we store it, and who gets access. We want to avoid alert fatigue, but we also don't want to miss silent failures. And we need systems that are smart enough to suppress noise, but sensitive enough to flag risk.

Ultimately, the goal is to make observability not just the eyes and ears of the enterprise—but also its central nervous system. Because when your systems are too complex to watch manually and too critical to fail quietly, you need observability that's intelligent, secure, and aligned with how your business thinks.

Observability Is the Next Frontier of Public Sector Efficiency

Observability enables transparency, accountability, and efficiency in public sector IT.

At Punjab National Bank, we run one of the largest, most diverse IT landscapes in the country—spanning branches, core banking, middleware, digital channels, and regulatory platforms. The complexity is enormous, and so is the expectation: performance, uptime, and auditability at every step. In this context, observability is not a luxury—it's a strategic enabler of service delivery and compliance.

In public sector banking, we often deal with systems that have evolved over decades. We can't afford to discard them overnight, yet we must integrate them with modern digital interfaces and deliver a seamless experience to customers and regulators alike. This means our observability goals are multi-layered—we need visibility into not just application health, but also data flows, backend jobs, batch dependencies, and regulatory handshakes.

One of our biggest challenges is the asynchronous nature of many of our systems. For example, a service request initiated at the branch or digital interface may go through multiple back-office systems before it's fulfilled. If a delay happens, we need to pinpoint where exactly it occurred—was it a job scheduler? A network node? A data mismatch? That kind of diagnosis is difficult without well-integrated observability pipelines.

I've come to believe that the real power of observability lies in process correlation. It's not just about metrics or uptime. It's about seeing the transaction the way a user sees it—across channels, across systems, and across delays. If a pension payment gets delayed, it's not enough to know that the server was up. We need to know whether the file was generated, the signature verified, the payment dispatched, and the acknowledgment received.

The second big aspect is governance. Observability, for us, also has a strong compliance angle. We need to demonstrate that our systems not only work—but work securely, consistently, and within regulated SLAs. That's why we are looking at observability as a foundation for audit readiness—capturing event trails, access patterns, and exception handling automatically, without burdening our teams with manual logs or checklists.

In public sector setups, cost and capacity are always critical considerations. We can't afford to flood our infrastructure with telemetry or over-provision our monitoring stack. So we're taking a more targeted approach to observability—identifying critical journeys, high-volume transactions, and sensitive data paths, and focusing our efforts there.

We're also pushing for more awareness and skill-building across our IT and operations teams. Observability is not just for NOC engineers—it's a shared responsibility. Everyone—from infra admins to application developers to business users—needs to understand how their piece of the puzzle contributes to overall service health.

"In complex, federated environments like ours, observability must bring transparency not just to systems—but to decisions, delays, and deviations."

Ashwini Kumar Pandey
Chief Information Security Officer,
Punjab National Bank

Observability Must Move from Reactive Logs to Real-Time Intelligence

Observability must evolve into AI-driven, real-time, business-aligned intelligence.

"The future of observability is not just full-stack—it's full-scope, with AI stitching together patterns that humans would otherwise miss."

Munish Blaggan
Head Technology - Customer Engagement Platform, ICICI Bank

At ICICI Bank, we manage technology infrastructure that touches millions of customer interactions every day. It's not just about uptime anymore—it's about detecting, responding to, and even predicting disruptions before they affect user experience or regulatory commitments. That's where observability is moving—from static dashboards to real-time, AI-enriched decision support.

Traditional monitoring gives us health metrics, but it doesn't always tell us the full story. A CPU spike might get flagged, but what's the impact? Is it slowing down a high-value journey like loan disbursement or UPI payments? Or is it isolated to a low-volume, non-critical task? We need that context immediately, and at scale. That's why we've begun leaning into AI-based correlation—not just for speed, but for precision.

One of the critical gaps we're trying to close is what I call the "visibility-to-action" delay. Too often, even when systems raise alerts, it takes time to sift through logs, check dependencies, and coordinate teams. By then, the customer may already be impacted. Our goal is to shrink that window to seconds. We want observability platforms that not only surface issues but also suggest remediation, backed by learning from past patterns.

AI plays an important role here. For instance, if a particular batch job fails intermittently during month-end load spikes, the system should recognize that pattern and pre-emptively flag similar risk conditions. In many cases, AI now helps us identify latent issues—signals that don't trip traditional alerts but combine to form meaningful insights when stitched together. That's something humans simply can't do at enterprise scale, especially in hybrid cloud environments.

Another area where we're evolving is mapping observability to service-level commitments. Uptime alone doesn't cut it. We're tracking journey-level SLAs—how long a particular transaction takes, whether a service chain performs within tolerance levels, and how downstream APIs behave under variable loads. This moves observability closer to business reality.

But technology alone isn't the answer. Observability must be embedded into the way teams work—not just something a central NOC owns. Developers, infra leads, and business units all need a common understanding of what system health means in their context. We're working to build cross-functional dashboards, where alerts are not just technical anomalies but mapped to business KPIs like turnaround time, NPS, or revenue leakage.

As we grow, the volume of telemetry will only increase. We're already taking a tiered approach—deep instrumentation where the business impact is highest, lighter signals elsewhere. AI helps us prioritize where to look, when to act, and how to escalate.

Observability Is the Bridge Between IT, Business, and the Unknowns.

AI-powered observability bridges systems, business impact, and operational foresight.

In today's hyper-distributed enterprise ecosystems, observability is evolving from a technical necessity to a strategic enabler. We have learnt that no two incidents behave the same, and no single tool or static rulebook can anticipate the ripple effects of failure across systems. That is why our approach to observability must be rooted in making the unknown visible—and actionable, proactively.

By integrating AI-powered business analytics, customer journey mapping, and event correlation, we must enable leaders to make real-time decisions that influence business outcomes—not just uptime.

In modern enterprise ecosystems, failures rarely occur in isolation. A minor latency in a single microservice can have a magnified impact across APIs, message queues, and third-party integrations—ultimately disrupting customer experiences. That is why observability strategies need to go beyond surface-level symptoms. These must focus on tracing the full consequence chain, connecting events to system-level signals and further to business processes. By designing platforms that start resolution from business impact rather than technical analysis, teams must be empowered to respond proactively with precision, speed, and relevance.

Tool based analytics and isolated RCA are too slow for today's dynamic environments. Embedding predictive AIOps into the observability fabric is crucial. AI needs to correlate data across layers—network, infrastructure, application, and user behavior—learning new and emerging patterns. The next time a similar event occurs; the system does not have to start from zero—it must start with foresight.

This shift transforms operations teams from firefighters into strategists, enabling automated remediation and causal AI-driven alerting. It also puts teams on a forward-looking automation maturity journey towards AI-led autonomous operations across the landscape.

Observability without a business context is just noise. A failed background job at 2 a.m. might appear trivial—until you realize it feeds the morning MIS or supports a campaign launch. Platforms should be designed to map technical events to business events / processes such as billing cycles, product launches, compliance windows—so resolution begins with understanding the impact.

In noisy IT environments, context-aware observability is the key to clarity. As enterprises move towards a unified observability fabric that standardizes telemetry across disparate systems they must deliver persona-based dashboards for developers, SREs, and business leaders alike.

This democratization ensures every stakeholder sees what matters—in their own language, not just in logs or metrics.

"The complexity is no longer in the code—it is in the interconnections. That is where observability and AI must converge."

Naresh Choudhary
Senior Vice President,
Infosys

Observability Is the Pulse of a Multi-Business Digital Core

Observability unifies diverse systems with AI for business-aligned clarity.

"You can't run a digital enterprise without knowing how the blood flows. Observability gives us that pulse—and AI keeps it steady."

Rejin Surendran
Global Chief Information Officer,
Wipro Enterprises

At Wipro Enterprises, we oversee a diverse mix of businesses—, consumer care & lighting and industrial manufacturing—each with its own customer needs, operational models, and digital maturity. This diversity means that there's no single view of performance, no universal baseline. Yet as a global CIO, I need confidence that our systems are not only running but performing with purpose.

That's where observability comes in—not just as a monitoring function, but as a strategic mechanism to unify visibility across silos. In today's multi-business digital cores, you can't afford to rely on gut feel or anecdotal feedback. You need telemetry—real-time, contextual, and connected.

What we've learned is that observability isn't about instrumenting everything. It's about instrumenting what matters.

But scale brings complexity. A single process, like order-to-cash or invoice reconciliation, may cut across ERP, workflow, API gateways, mobile apps, and third-party services. In the absence of comprehensive observability across all layers, responses are limited to addressing symptoms rather than identifying and resolving underlying causes. That's why we've begun leveraging AI-enabled correlation engines to identify patterns across logs, metrics, and user events. These tools don't just tell us what failed—they help us connect technical signals to business disruptions.

AI also helps us prioritize incidents based on real-world impact, not just on alert thresholds. An error in a high-volume B2B billing module deserves more attention than a latency spike in a sandbox test app. Instead of drowning in logs, we want the system to surface the five things we must care about today—and why.

Another dimension we're focused on is bringing business teams into the observability loop. Our goal is to democratize insights, so functional leaders can see the health of their own processes—without needing a technical background. If there's a delay in a plant dispatch, or a spike in returns processing time, business users can see it live and work with IT to address the root cause. This collaborative visibility builds trust, reduces friction, and speeds up recovery.

Of course, observability must also adapt to change. As we onboard new platforms, migrate workloads, or acquire new businesses, our observability model needs to keep up. That's why we treat observability not as a static dashboard, but as a living capability, continuously evolving with the enterprise.

In a distributed, multi-business environment like ours, AI and observability go hand in hand. One tells us what's happening, the other tells us what to do next. And when both are aligned with business context, we move from being reactive operators to proactive enablers of transformation.

That, to me, is the real value of observability—not just seeing more, but knowing better.

Observability isn't a tool, it's a mindset

Observability with AI ensures global service trust, clarity, and resilience.

At TMF Group, we deliver business-critical services including HR, payroll, compliance, legal and accounting for clients across more than 87 jurisdictions. That means our digital systems carry the weight of regulatory deadlines, payroll cycles and real-time commitments in dozens of different contexts. For us, observability is not just a platform, it's a risk-control layer, a trust mechanism and a key part of how we serve clients without disruption.

One of the biggest changes we've made is shifting from traditional infrastructure monitoring to end-to-end journey observability. Our services often involve workflows that touch multiple applications, some internal, some third-party, some running in cloud-native containers and others in legacy environments. If something breaks, it's not enough to know which server or container triggered an alert. We need to know which client, country and business function is impacted, and what the potential fallout is.

That's where AI is becoming indispensable. Our environments generate vast amounts of logs and telemetry across applications, cloud platforms and integration layers. Using AI to correlate patterns across these layers helps us avoid alert fatigue and focus on what truly matters. We've seen incidents where early symptoms would have gone unnoticed in isolation, but when viewed together, pointed to a brewing failure. That kind of foresight is only possible when AI augments our visibility with intelligent context.

We're also looking at how observability can improve SLA adherence. If payroll processing in a particular country takes 20 minutes longer than expected, or if a statutory report isn't delivered on time, we want to know before it becomes a client issue. Observability, in that sense, helps us stay on the front foot not just monitoring performance, but predicting risk.

Another critical shift has been cultural. In large global environments, teams often operate in regional bubbles. One of our goals has been to bring observability into the shared vocabulary across Dev, Ops, Security and Business. It's not enough for engineers to see metrics; business stakeholders need to see how system health connects to service reliability. That's why we've started building dashboards not just for IT, but for country managers and client delivery heads, giving them real-time assurance.

But observability at a global scale isn't free. We're mindful of cost, complexity and context. We don't want to instrument everything blindly. We want to prioritise the systems and journeys that are mission-critical. AI helps here too, by identifying which workflows show early signs of fragility and which endpoints are bottlenecks under load.

Observability for us is not about chasing incidents—it's about building a system that self-monitors, self-learns, and knows when to ask for help. We're moving in the direction where automation handles the routine, and human teams focus on what truly demands expertise.

"When you operate globally, you can't chase every incident manually. Observability and AI must work together to ensure trust at scale."

Sajith Chakkingal
Head of Global Technology Services,
TMF Group

Observability Must Make Sense to the Business—Not Just to the Dashboard

Observability must translate system signals into real-world, business-aligned outcomes.

"In a field-intensive, high-touch business like ours, observability must connect digital signals to real-world outcomes—and AI helps bridge that gap."

Sudesh Puthran
Chief Technology Officer,
CreditAccess Grameen

At CreditAccess Grameen, our technology landscape supports one of the most complex and human-centric business models in financial services. We serve millions of rural and semi-urban women across India, and much of our operations happen far from the reach of core banking infrastructure—often at the edge, in field locations, and under challenging network conditions. In such a model, observability isn't just about system uptime—it's about service continuity, performance, scale, field agent productivity, and customer trust. (We process around 1 lakh++ loan applications a day !!! and hence performance and scale is very imp.

Our digital stack is evolving rapidly, with mobile-first apps, cloud-native services, and integration with mutiple fintech API`s. But even with the most advanced systems, problems can arise where visibility is weakest—at the last mile. If a mobile app fails to sync in a low-connectivity area, or if a digital payment stalls at an intermediary node, the customer doesn't see a backend glitch—they just see a broken promise. That's why we've started rethinking observability—not just as a backend function, but as a field-aware, outcome-driven capability.

We're working to stitch together telemetry from application, infrastructure, and user behavior—but the real challenge is interpretation. That's where AI becomes essential. We're using AI models to look for early indicators of friction—slower app response in a geography, increasing sync retries, or a spike in manual overrides by agents. These may not always raise traditional alerts, but they often point to a service degradation waiting to happen.

If a region is seeing slower transaction throughput, or if a cloud integration layer is under stress, we want the system to tell us—not just what's wrong, but what might go wrong, and what to do about it.

Another area we're focused on is business alignment. As a financial inclusion institution, our business users care about metrics like loan disbursement timelines, repayment syncs, and agent productivity. If observability doesn't help answer those questions, it's irrelevant. That's why we're building business-facing dashboards, where system health is presented not in CPU or API metrics, but in terms of field-level success rates and process compliance.

This also improves accountability. When observability is democratized—shared with product teams, field operations, and business heads—it becomes a collaborative tool rather than a diagnostic afterthought. People don't just react to issues—they anticipate them, flag them, and learn from them.

Of course, cost is a factor too. We can't afford to over-instrument in every environment, especially when dealing with remote regions. We're using a tiered approach, where critical workflows receive deep observability, and non-critical ones are monitored using AI-based anomaly detection to keep telemetry lean.

Enabling Business Observability at Enterprise Scale

Today, organizations need to transform observability into a strategic advantage by aligning full-stack telemetry with business outcomes, AI-driven insights, and automated action.

"Modern enterprises need real-time observability to understand incidents, their business impact, and actionable responses. Deep, contextual, and automated insights drive faster decisions and resilience, keeping business outcomes at the core."

Shakti Agarwal
Lead Solution Consultant,
Dynatrace India

The modern digital enterprise is undergoing a fundamental shift. As cloud-native architectures, hybrid infrastructures, microservices, and now LLMs and agents proliferate, IT environments are becoming more dynamic, distributed, and interdependent than ever before.

The pace of change—driven by digital transformation, customer expectations, and regulatory shifts—has elevated the need for real-time visibility into every layer of the digital stack. However, observability is no longer just about infrastructure uptime or application health. The focus has shifted toward tying technical performance directly to business outcomes—conversion rates, revenue generated, customer churn, business failure codes, and more.

There is also a clear transition underway in how observability platforms are consumed. Organizations are increasingly moving toward SaaS-based delivery models that offer massive scalability, AI-powered analytics, usage-based pricing, and an all-in-one platform. At the same time, expectations from observability platforms are expanding to include root-cause detection, prediction, and auto-remediation—preferably in real time. CIOs and IT leaders now view observability not as a reactive tool, but as a strategic enabler of business agility, resilience, and growth.

Key Challenges in Achieving Effective Observability

Despite the criticality of observability, many enterprises continue to face significant challenges in implementation and value realization.

A persistent issue is fragmentation across tools, data sources, and teams. IT departments often juggle multiple monitoring solutions that operate in silos, creating data dissonance and making it difficult to stitch

Praveen Mahajan
Regional Director – India,
Dynatrace

> "The difference between successful and failed observability implementations often comes down to adoption. It's not just about deploying tools; it's about how well teams embrace and use them across the business."

together a complete picture of business service health. This fragmentation leads to data silos, alert fatigue, lack of collaboration, and longer mean time to identify (MTTI) and mean time to resolve (MTTR).

Another significant gap lies in the disconnect between technical telemetry and business impact. Metrics such as CPU utilization or error rates, while helpful to engineers, are not meaningful to business stakeholders unless they can be contextualized in terms of revenue loss, business failures, or customer churn. This misalignment prevents CIOs from influencing strategic decisions and justifying technology investments.

Scalability is a parallel concern. With increasing data volumes from distributed systems, legacy monitoring infrastructures struggle to ingest, process, and correlate observability signals in real time. Additionally, organizations operating in regulated sectors (e.g., BFSI, healthcare, and aerospace) must grapple with compliance, data privacy, and operational governance—while still maintaining agility and responsiveness.

Cultural and operational hurdles further compound these technical challenges. Observability initiatives often lack centralized ownership, and adoption across teams can be uneven. The absence of intuitive dashboards and actionable insights undermines trust in the platform, limiting its role to a back-office utility rather than an enterprise-wide capability.

A Unified and Intelligent Approach to Modern Observability

Dynatrace approaches observability as a unified, intelligent, and automation-ready platform built for the complexity of modern enterprise environments. At the core of its solution is full-stack observability that extends from infrastructure and applications to real user monitoring, synthetic testing, business metrics, and security telemetry—all visualized in a single, contextual interface.

Unlike many AI tools that struggle with inaccuracies and hallucinations, Dynatrace's Davis AI engine empowers teams to not only detect what went wrong during an incident, but also understand why it happened—enabling faster resolution before it impacts customers or revenue.

Dynatrace combines the powers of predictive, causal, and generative AI. Its causal AI capability analyzes observability and security data in the context of topology information. Generative AI creates queries, notebooks, and dashboards to simplify analytics and provides workflow and automation recommendations. With predictive AI capability, you can now forecast any data or events being ingested into Dynatrace, allowing organizations to take proactive measures. Capacity planning is another use case where predictive AI functionality can be extremely useful.

Generative AI further helps by bringing in natural language insights, significantly improving operational efficiency, providing smart auto-remediation suggestions, and more.

To close the business-IT alignment gap, Dynatrace enables the creation of role-

based dashboards tailored for diverse personas—from developers and SREs to business leaders. Stakeholders at every level can see the metrics that matter most to them—be it SLA compliance, service uptime, or conversion performance.

On the automation front, Dynatrace supports closed-loop remediation via workflow orchestration. This includes triggering actions like restarting services, scaling infrastructure, or throttling traffic based on predefined policies. These capabilities extend across hybrid and multi-cloud environments, with integrations into major CI/CD pipelines to support shift-left observability and DevSecOps practices. Its AI capability goes a step further, enabling AI systems to make decisions and take actions on behalf of a user or system.

The platform also embeds runtime security observability, thereby enabling application teams to detect vulnerabilities in real time, block real-time attacks, understand their security posture concerning cloud and Kubernetes setups, perform threat hunting, and more. With this security intelligence, the platform serves as a key enabler of digital trust and risk management.

To support rapid deployment and scale, Dynatrace offers a SaaS delivery model with in-region data centers, granular data access controls, and regulatory alignment. This addresses the dual need for innovation velocity and compliance assurance—particularly relevant for sectors with stringent regulatory requirements like BFSI, pharmaceuticals, and automotive.

Operationalizing Observability for Enterprise Impact

From Dynatrace's experience working with leading enterprises, several best practices have emerged for implementing observability as a business enabler:

- **Focus on Business-Critical Services.** Rather than attempting blanket coverage, organizations should begin by instrumenting services that directly affect revenue, customer experience, churn, or operational stability.
- **Establish a Center of Excellence (CoE).** Observability should be led by a cross-functional team that includes application owners, infrastructure engineers, business analysts, and SREs. A centralized CoE helps drive standardization, governance, and adoption.
- **Map Technical KPIs to Business Outcomes.** Verify that business metrics and failure codes are being captured. Use dashboards that surface meaningful, decision-grade insights.
- **Automate Wherever Possible.** Use workflows to act—automatically resolve known issues, create auto-tickets in ITSM solutions, collect forensic data, and apply gating, setting the stage for agentic AI in the future.
- **Invest in Continuous Enablement.** Adoption is not a one-time event. Regular training, enablement programs, and feedback loops are critical to embedding observability into the organization's DNA.
- **Align Observability to Compliance and Security Goals.** Incorporate data sensitivity monitoring, anomaly detection, and policy enforcement into observability workflows to meet internal and regulatory requirements.

Through its platform and field expertise, Dynatrace enables organizations to move beyond operational monitoring and toward business observability—connecting the dots between system behavior and enterprise value. In a digital-first economy, this is not a luxury—it is a strategic and business necessity.

As enterprises evolve in complexity and scale, the role of observability will transcend traditional boundaries. Business observability enabled by platforms like Dynatrace empowers CIOs to move from reactive firefighting to proactive, outcome-driven leadership. By aligning technical performance with strategic business metrics, organizations can unlock faster innovation, reduce risk, and deliver exceptional customer experiences.

In this new era, observability is not just a tool — a catalyst for digital transformation, competitive advantage, and sustained growth.

Elevating Digital Experience: From Reactive Monitoring to Real-Time Business Value

In tomorrow's digital-first, AI-led enterprises, observability must be embedded with a unified, strategic vision.

Executive Summary

In an economy increasingly defined by milliseconds and mobile screens, digital experience has moved to the frontlines of business strategy. No longer a downstream concern of UI or performance monitoring, it now functions as a brand promise: one that must be fulfilled consistently, intuitively, and securely—whether for customers navigating a payment flow or employees resolving a critical task in a complex global enterprise.

Today, experience defines enterprise value. Whether engaging customers through a mobile app or supporting employees in distributed environments, organizations are expected to deliver intuitive, responsive, and seamless digital journeys. These experiences now play a central role in how users perceive trust, convenience, and brand relevance.

This chapter, grounded in insights from three in-depth panel discussions with technology leaders across industries, explores how Indian enterprises are reimagining digital experience as a driver of growth, trust, and resilience. The transformation is underway—from reactive incident response to proactive, AI-powered orchestration across fragmented technology landscapes.

What emerges clearly is a new mandate for CIOs: to deliver experience at scale while aligning deeply with business outcomes. Experience today spans mobile-first consumer journeys, real-time risk detection, frictionless employee support, and hyper-personalized services—all underpinned by infrastructure, data, and applications that are themselves evolving rapidly. Meeting this mandate demands not only technical modernization, but a philosophical shift: from system-centricity to user-centricity.

Yet the challenges remain formidable. Fragmented tools and siloed systems continue to impede visibility. Business teams struggle to connect performance insights to revenue or customer retention. AI and GenAI investments are rising, but concerns around ROI, security, and regulatory complexity persist. And with a growing dependence on external ecosystems—SaaS platforms, fintech APIs, digital partners—organizations must now manage trust and uptime far beyond their own walls.

Several cross-cutting needs stand out. Enterprises require unified observability that connects infrastructure telemetry with business events and user behavior. They seek real-time dashboards that serve both CIOs and CMOs—tracking not just latency and crashes, but journey drop-offs, consent failures, and missed conversion triggers. They need systems that move from insight to action—automating remediation, scaling intelligently, and adapting in-flight experiences. Most importantly, they need governance models that balance personalization with privacy, agility with accountability.

Panelists have shared real-world accounts of making this shift—from implementing predictive observability platforms and event-driven architectures, to embedding AI into fraud prevention, compliance, and experience tuning. There's growing maturity in the ecosystem too, with partners contributing more than tools—bringing implementation expertise, industry benchmarks, and patterns that accelerate adoption.

The strategic direction is now clear: Digital experience is no longer an output of transformation. It is the operating system of the intelligent enterprise. Whether serving a retail banking customer, an energy distribution partner, or an internal developer, the quality of the digital interaction is the enterprise's value, delivered in real time.

As one speaker aptly summarized, "Digital experience is the new balance sheet." It captures not just how well systems run, but how deeply an organization understands and supports its users at every touchpoint. The organizations that can sense, respond, and optimize at this level—consistently and securely—are not just building better interfaces. They're building lasting business advantage.

This chapter sets the foundation for that journey, drawing from both the pain points and breakthroughs of a diverse panel of CIOs. It frames the challenges, surfaces the imperatives, and charts a path from reactive monitoring to experience-driven value creation—anchored in observability, automation, and aligned intent across business and technology teams.

Key Needs and Drivers: What's Fueling the New Digital Experience Imperative

As organizations accelerate their digital transformation journeys, delivering a high-quality digital experience has moved from a functional goal to a foundational business requirement. Users—whether customers or employees—expect seamless, fast, personalized, and secure interactions at every digital touchpoint. This expectation is reshaping enterprise priorities

"The strategic direction is clear: Digital experience is no longer an output of transformation. It is the operating system of the intelligent enterprise."

and elevating experience to the top of the CIO agenda.

The discussion among industry leaders highlights several converging drivers behind this shift.

From Transactional Interfaces to Seamless Journeys
Users no longer engage through a single channel or touchpoint. They move fluidly across mobile apps, websites, kiosks, chatbots, and physical spaces. This omnichannel reality demands seamless orchestration of journeys—not just individual features.

Experience is no longer about building an app. It's about engineering connected moments that span systems, screens, and even partners. Every interaction must feel contextual, consistent, and complete.

Unified View of the User: The 360-Degree Mandate
Many organizations still operate with fragmented customer data, stored across CRMs, ERPs, legacy systems, and bespoke platforms. This makes personalization difficult and limits real-time responsiveness.

The ability to assemble a dynamic, unified view of the user—including preferences, transactions, behaviors, and intent—is now central to delivering high-quality digital experiences. Customer Data Platforms (CDPs) and event-driven architectures are key enablers here.

Mobile-First and AI-Ready Architectures
India's digital adoption is largely mobile-led. In BFSI, retail, and utilities, mobile apps are now the primary service channel. Expectations around speed, simplicity, and reliability are higher than ever—and often unforgiving.

To meet these expectations, enterprises are adopting cloud-native, API-first, composable architectures that support not just performance, but continuous improvement. These architectures also need to be AI-ready—capable of ingesting data, generating insights, and taking action in real time.

Experience as a Promise, Not Just a Feature
Digital experience is now deeply tied to trust, value, and brand identity. Users expect that a company will "know them," value their time, and act in their interest. This means personalization must be balanced with privacy, and convenience must never compromise control.

Especially in regulated industries, experience must reflect security by design, compliance awareness, and transparent data handling. Consent management and data minimization are becoming experience-critical features, not backend responsibilities.

Real-Time Observability Across Systems and Partners
Modern digital journeys traverse not just internal systems, but fintech APIs, cloud platforms, SaaS tools, and third-party gateways. The need for unified observability—across infrastructure, application performance, and user experience—is critical.

Leaders now demand real-time dashboards that provide business context: not just latency or CPU usage, but drop-offs in account opening flows, transaction failures, or delayed SMS messages that breach service expectations.

Preventive and Predictive Detection over Reactive
Proactive detection is no longer a luxury—it's expected. Whether it's a mobile app crash or a payment gateway slowdown, users expect zero friction. The shift from reactive incident response to preventive, self-healing systems is accelerating.

AI-powered root cause analysis, anomaly detection, and automated remediation are becoming standard components of the digital experience stack. The goal: detect before disruption, and fix before fallout.

Experience Metrics Must Map to Business Outcomes
Organizations are increasingly linking digital experience metrics with business KPIs—conversion rates, NPS, churn, and revenue.

"Experience is no longer about building an app. It's about engineering connected moments that span systems, screens, and even partners. Every interaction must feel contextual, consistent, and complete."

It's no longer sufficient to monitor uptime; CIOs want to know where and why users abandon carts, delay decisions, or disengage entirely.

This requires platforms that surface not just technical anomalies but business-impacting insights—bridging the gap between application logs and boardroom strategy.

Embedded Trust, Compliance, and Security
Data privacy regulations like India's Digital Personal Data Protection Act (DPDP) are reshaping the digital experience agenda. Enterprises must now embed consent management, auditability, and minimal data exposure directly into user flows.

Beyond compliance, trust itself is becoming a differentiator. Secure, predictable, and respectful digital interactions are now core to user satisfaction and brand value.

Internal Experience Matters Just as Much
Many panelists highlighted that digital experience isn't just customer-facing. For large IT services firms, banks, and manufacturing units, internal users—employees, agents, and partners—rely on streamlined digital tools.

Improving the internal experience unlocks productivity, reduces fatigue, and strengthens service delivery. Whether it's resolving IT tickets faster or onboarding a new employee, these journeys are just as important as customer ones.

Observability as a Native, Not Add-On, Capability
Perhaps most crucially, observability must now be embedded into the DNA of digital strategy—not layered on as an afterthought. It should be integrated into application lifecycles, tied to experience-level objectives (XLOs), and shared across business and IT.

As digital expectations continue to grow, observability is becoming the nervous system of the enterprise—connecting performance, experience, and outcome into a continuous feedback loop.

The Panel's Mandate
The message from the panel clear: Digital experience is not a destination—it is a system to be architected, measured, secured, and continuously evolved. The enterprises that can unify these needs—across people, platforms, and partnerships—will not just respond to market shifts. They will shape them.

Implementation Challenges: What's Hindering the Shift to Experience-centric Enterprise Models

While the ambition to deliver superior digital experiences is nearly universal, the execution is riddled with challenges. From fragmented technology ecosystems to cultural resistance and regulatory complexity, CIOs face a diverse set of hurdles that can stall—or even reverse—digital momentum.

The following themes reflect the core operational and strategic challenges that surfaced in the discussion, cutting across sectors and enterprise sizes.

Fragmented Tools and Data Silos
Most organizations today rely on a patchwork of tools for infrastructure monitoring, application performance, digital analytics, and business reporting. These tools often operate in isolation, leading to disjointed visibility and data overload.

This fragmentation hampers root cause analysis, slows decision-making, and creates "blind spots" across the user journey—especially when systems or APIs span multiple departments, partners, or cloud providers.

Poor Correlation Between Technical and Business Metrics
Many IT teams can track uptime, latency, or error rates—but struggle to connect those signals to business outcomes like conversion, churn, or revenue leakage.

Without the ability to tie technical performance to business KPIs, CIOs face challenges in securing investments, justifying transformation, or making experience metrics part of strategic dashboards. Experience remains isolated from value.

Monolithic Systems and Legacy Drag
Despite investments in modernization, most core applications in large enterprises remain legacy-bound. Modularization efforts often stall due to architectural complexity, vendor lock-in, or change management fatigue.

This results in "all-or-nothing" modernization cycles, where even simple experience changes—like a PIN reset or onboarding

tweak—require full-stack upgrades. Agility is lost, and time-to-value is extended.

Operational Complexity in Distributed Environments
Modern environments span on-premise systems, hybrid clouds, SaaS platforms, and third-party integrations. Observability and experience monitoring across this mix is challenging—especially when providers don't expose telemetry or support open standards.

The result: IT teams spend more time stitching together insights than acting on them. Alert fatigue, inconsistent SLAs, and lack of transparency become chronic pain points.

Reactive Incident Management and Manual Resolution
While observability platforms have matured, many organizations still rely on manual ticketing, siloed dashboards, and tribal knowledge to resolve issues.

This reactive posture delays response, erodes trust, and strains IT operations. Automated diagnostics, AI-driven anomaly detection, and closed-loop remediation remain underutilized—especially in high-stakes systems like banking, manufacturing, and critical public infrastructure.

Inadequate Coverage of Ecosystem Dependencies

Today's digital journeys often involve third-party systems—fintech gateways, government APIs, payment platforms, and partner portals. Most observability tools are limited to internal environments, offering little visibility into external performance degradations or integration failures.

Without a full picture of the ecosystem, incident response becomes guesswork. Root cause blurs across organizational boundaries.

Compliance and Security Complexities
With the rollout of India's Digital Personal Data Protection Act (DPDP) and increasing regulatory scrutiny, organizations must design experience flows with consent, auditability, and minimal data exposure in mind.

But many tools and architectures were not built with this by default. This leads to retrofits, manual reviews, and policy ambiguity—undermining trust and compliance agility.

Skills Gaps and Change Resistance
Even where tools exist, adoption lags. Teams often lack the skills to configure observability platforms fully, extract business insights, or integrate them across silos.

At the same time, cultural resistance persists—especially where observability is seen as "extra work" or as exposing performance gaps. Without upskilling, collaboration, and change enablement, tools remain underused and their impact muted.

Vendor Overlap and Tool Proliferation
With multiple teams buying tools independently—analytics from marketing, observability from IT, compliance from GRC—enterprises face rising costs and overlapping capabilities. Tool rationalization becomes difficult.

Instead of an integrated experience ecosystem, organizations find themselves managing a sprawl of dashboards, licences, and data models—none of which deliver end-to-end clarity.

"Reactive postures delay response, erode trust, and strain IT. Automated diagnostics, AI anomaly detection, and closed-loop remediation remain underused in high-stakes sectors."

Lack of Preventive Design and Automation
Finally, few organizations have moved beyond proactive monitoring to true prevention. Very few environments use automation to not just detect and diagnose issues, but to fix them before they impact users.

Preventive design—such as auto-scaling, intelligent throttling, and dynamic routing—is rare. The result is firefighting, not foresight.

The Panel's Mandate
The cost of inaction can be high. The consequences of these challenges are real: degraded trust, lost revenue, regulatory penalties, operational burnout, and strategic inertia. As one participant put it, "We know what needs fixing—but stitching it together is the hard part."

Experience-first enterprises must now focus not just on what they want to build—but on what's getting in the way. The next section outlines how some are already doing that—and what others can learn from their path forward.

The Way Forward: From Tactical Fixes to Experience-led Transformation

The path to elevating digital experience is no longer about cosmetic redesigns or dashboard upgrades—it is a strategic shift that requires rethinking architecture, operations, measurement, and culture. As panelists across industries agreed, the enterprises that will succeed are those that embed observability, automation, and user context into every layer of the digital stack.

The shift is already underway. Here's what leading organizations are doing—and what others can learn from their direction.

Unify Observability Across Layers and Journeys
The fragmentation of tools, metrics, and teams has long been a barrier to experience excellence. The next wave of transformation involves breaking down these silos through full-stack, real-time observability platforms.

This means stitching together data from infrastructure, applications, APIs, and end-user interactions—whether through Real User Monitoring (RUM), synthetic tests, or distributed tracing. More importantly, it means surfacing business-impacting insights from that data: where users drop off, which journeys are breaking, and what's causing revenue leakage.

Make Experience Metrics Boardroom-relevant
To move experience up the priority ladder, CIOs are shifting from technical KPIs to business-linked outcomes—conversion rates, SLA violations, onboarding completion, or NPS changes triggered by app friction.

Leading organizations are building shared dashboards that business and IT can both understand and act on. These "joint taskboards" bridge the long-standing divide between backend telemetry and front-end outcomes, helping teams speak the same language and prioritize what truly matters.

Shift from Reactive to Predictive—and Preventive
Reactive response is no longer sufficient in environments where downtime or drop-offs can destroy trust instantly. CIOs are increasingly investing in predictive capabilities—powered by AI and anomaly detection—that allow teams to act before users are affected.

The next frontier is preventive automation: systems that self-heal, auto-scale, or throttle intelligently based on real-time load or failure signals. Organizations are beginning to design for recovery before failure, not just response after impact.

Make Automation Intelligent and Actionable
It's not enough to observe issues—teams need tools that can also act. This means integrating observability platforms with orchestration tools, configuration managers, and service desks for closed-loop automation.

From real-time issue isolation to workload rebalancing or even rollback of a faulty release, intelligent automation reduces manual resolution time and lowers the risk of cascading failure. Done right, it frees up human capital for higher-value work.

Embed Trust and Compliance by Design Itself
Trust is now a cornerstone of experience. Enterprises are moving away from bolt-on compliance to native governance—embedding consent frameworks, audit trails, and DPDP-compliant data flows into every user journey.

This also includes deploying dynamic consent interfaces, visibility into how data is used, and the ability for users to control their own information. In highly regulated industries, this is not just compliance—it's brand protection.

Extend Experience Visibility to the Entire Ecosystem
As digital journeys increasingly rely on third parties—fintech APIs, identity services, content delivery networks, etc.—organizations must extend observability beyond their own walls.

Leading enterprises are using open standards, telemetry integration, and partner SLAs to gain insight into external choke points. They're asking not just "Is my system up?" but "Is the entire customer jour-

ney performing as intended—even the parts I don't control?"

Build Digital Experience into Strategy—Not Just Tooling
One recurring theme was clear: observability must be native to the enterprise's digital operating model, not an add-on. It needs to be embedded in development pipelines, release processes, compliance frameworks, and even go-to-market planning.

When experience is treated as a strategic input—not a post-launch report card—organizations can build digital journeys that are faster, safer, and smarter from day one.

Reskill for an Experience-Centric Culture
Achieving all of this requires more than technology—it requires new skills and mindsets. IT teams must learn to connect performance to business value. Business teams must understand system dependencies and operational realities.

Upskilling across observability, AI, automation, and CX engineering is critical. Just as importantly, organizations must foster a culture of shared ownership—where digital experience is everyone's job, not just the CIO's.

Integrate GenAI and AI Observability Thoughtfully
As enterprises experiment with GenAI for chat interfaces, personalisation, and content generation, new complexity emerges around explainability, model behavior, and trust. CIOs are now exploring how to observe AI systems just like any other production workload—tracking latency, drift, prompt injection, or unexpected outputs.

AI observability platforms are helping detect when models behave unexpectedly or when biased or irrelevant outcomes risk user trust. As AI becomes part of the experience layer, observing how it performs—and how it impacts the journey—is no longer optional.

> "The next frontier is preventive automation: systems that self-heal, auto-scale, or throttle intelligently on real-time load or failure signals. Organizations are shifting focus to recovery before failure, not just response after impact."

Justify Experience Investments with Stronger Business Cases
One of the recurring asks from panelists was for clearer frameworks to quantify ROI from observability and digital experience platforms. Business leaders want to understand how investments in tools, data, or automation translate into measurable gains—reduced MTTR, increased retention, or compliance readiness.

Organizations are now building internal playbooks that map specific interventions to outcomes, using journey-level metrics and before–after comparisons to tell the story. A well-articulated business case not only secures buy-in but accelerates adoption across teams.

The Panel's Mandate
The way forward is not about doing more of the same with better tools. It's about re-architecting experience as a shared priority across technology, business, compliance, and customer functions. That means eliminating silos, aligning metrics to value, and designing systems that anticipate and adapt in real time.

The organizations leading this charge are turning digital experience into a living, intelligent system—one that senses, responds, and evolves as fast as users expect it to. For them, experience isn't an output. It's the new operating principle.

From Insight to Impact: Digital Experience as the New Strategic Operating System

Digital experience is no longer a surface-layer concern—it is the connective tissue linking user trust, operational resilience, and business performance. What began as a conversation about performance metrics has evolved into a strategic dialogue about how enterprises sense, respond, and deliver value in real time. Elevating digital experience requires more than great design—it demands unified observability, preventive automation, intelligent orchestration, and a culture of shared accountability. The CIOs leading this shift aren't just responding to change—they're architecting a future where experience is measurable, adaptable, and inseparable from enterprise success.

The next wave of growth and loyalty will come from organizations that treat experience as a living system—constantly monitored, intelligently optimized, and deeply aligned with user needs and business outcomes.

Those who succeed will not be the ones with the most tools or the latest interfaces, but those who build adaptive, insight-driven systems that respond in real time to human needs. Elevating the digital experience is not a one-time initiative—it is the new operating system for enterprise success.

AUTHORS

The Key Is to Deliver Experience Without Overengineering

Amol Pai
Chief Technology Officer, Jio BlackRock AMC

Experience Is a Business Priority, Not a Tech Afterthought

Anjani Kumar
Chief Digital and Information Officer, Ather Energy

Making Experience Observable from the Inside Out

Ashish Desai
Joint President and CIO - Textiles Business, Aditya Birla Group

Elevating Experience Starts with Business Context

Harsh Jha
Group Head of Technology, Nuvama

Scaling Experience Begins with Control, Context, and Confidence

Kunal Dhingra
Chief Technology Officer, RBL Bank

From Data to Decisions—Engineering Digital Experience That Learns

Manish Malik
Executive Director (Information Systems) Indian Oil Corporation Limited

Digital Experience Design Is Grounded in Trust and Transparency

Om Prakash Seth
Chief Information Officer, IDBI Bank

It's about Connecting Internal Experience to External Excellence

Vikas Dureja
Vice President & Global IT Leader, HCL

From Performance to Business Impact: Rethinking the Digital Experience

Meghana Shroff
Senior Customer Success Engineer, Dynatrace

Faisal Shaikh
Regional Director, Dynatrace

The Key Is to Deliver Experience Without Overengineering

Impactful digital experience comes from focused delivery, not overengineering everything.

"One doesn't need to modernize the entire stack to make an impact. Focused, scalable solutions win over overengineered transformation."

Amol Pai
Chief Technology Officer,
Jio BlackRock AMC

In the banking industry, customer expectations around digital experience are not just high—they're evolving faster than most architectures can keep up with. Users want control, speed, and personalization at their fingertips, and often, the pressure is to match or beat the digital journeys offered by fintechs and consumer tech players.

But as I've experienced firsthand, meeting those expectations doesn't always mean tearing down and rebuilding your entire IT landscape. One of the recurring traps I've seen across organizations—including our own in earlier phases—is the tendency to over-engineer solutions in the name of digital experience.

Take a simple customer need—say, allowing someone to reset their card PIN through the app. In theory, it's a lightweight feature. But because of how traditional banking systems are structured, delivering that experience often requires tapping into multiple modules across the core platform, customer authentication stack, risk engines, and notification layers. That "small" feature can quickly turn into a large-scale project involving weeks—or months—of work. It's in these moments that the experience ambition collides with architectural reality.

This is where CIOs need to exercise judgment. The goal isn't just to modernize everything—it's to identify where value can be unlocked with minimal disruption. Sometimes, that means building experience layers on top of legacy cores. Other times, it means isolating a micro-journey that can be containerized, API-enabled, and delivered independently. We've learned to segment experience delivery from system modernization—and that's made a huge difference.

Another dimension to this is the complexity of the customer journey itself. Most journeys don't begin and end in one application. A single interaction might span multiple backend systems, third-party integrations, risk and compliance checks, and real-time messaging services. Orchestrating all of that to feel smooth to the end user is a challenge—but it's also where a lot of the magic happens.

We're also seeing our technology partners evolve in how they engage. Earlier, vendors would provide tools or platforms and leave it to us to figure out the implementation. Now, they come in with design patterns, industry benchmarks, and playbooks from similar rollouts.

One area I'm cautious about is peer benchmarking. Of course, we all want to keep up with the market. But jumping into transformation just because a competitor launched something flashy can backfire.

Experience Is a Business Priority, Not a Tech Afterthought

Digital experience succeeds when led by business intent, not tools.

Across the industries I've worked with, I've observed a recurring misconception that digital experience is something you can "deploy" through tools.

But experience isn't a product of technology alone. It's deeply embedded in how a business operates and how decisions are made. It reflects priorities, processes, and culture—not just code.

Too often, organizations invest in the latest buzzwords and platforms, hoping these will resolve longstanding pain points. But without clearly defined problems or purposeful outcomes, the effort remains superficial.

The Challenge of Fragmentation

A big hurdle in delivering meaningful digital experience is fragmentation—not just of systems, but of accountability. Very often, infrastructure belongs to one team, while applications belong to another. Security may be managed separately, and user metrics might sit with the marketing.

Each group sees only part of the picture. The result is that there are silos of monitoring, but no true visibility.

Start with the Journey

To cut through the noise, begin with a specific user journey: onboarding, billing, claims, etc. Then ask: Where are the delays? Where do users drop off? What's the cost of failure in this moment?

Only after you've answered those questions should you bring in tools—observability, automation, and performance tuning.

The goal isn't total visibility. It's intelligent observability—the ability to surface what matters, when it matters, without overwhelming the teams that rely on it.

Complexity Is the New Normal

In today's hybrid environments, users aren't always customers on a website. They could be field technicians working in harsh conditions, or engineers in a control room.

So, meeting the bar requires real-time insight—not just into your systems, but across partner ecosystems, edge devices, and operational constraints.

Improving digital experience is not a tech upgrade—it's a leadership decision. It requires cross-functional alignment; disciplined prioritization; and a cultural shift toward intentional design—at the problem-definition stage.

Co-Creation is the Future

When the business owns the experience and IT co-creates it, you unlock something powerful: not just speed, not just uptime, but a digital journey that feels seamless—not because it's fast, but because it's been thought through, end to end.

That's the kind of experience people remember, That's what makes digital transformation real.

"Observability tools alone doesn't make a business observant. True experience must be designed with business goals in mind."

Anjani Kumar
Chief Digital and Information Officer,
Ather Energy

Making Experience Observable from the Inside Out

True experience requires observability that maps system signals to journeys.

"Monitoring system health isn't enough—we must monitor business journeys. That's where the real experience lives."

Ashish Desai
Joint President and CIO - Textiles Business, Aditya Birla Group

As someone leading digital transformation in a large, complex organization, I've learned that elevating digital experience isn't about pushing out shiny new interfaces. It's about what happens behind the screen. It's about the reliability of APIs, the behavior of backend systems, the responsiveness of teams, and the alignment between technology and business processes.

At the heart of great experience is end-to-end observability—but most enterprises still don't have it. We might monitor individual systems or applications, but we lack a complete picture that connects the infrastructure layer to the actual user journey. If a customer drops off during onboarding, is it because the server was slow? Because the KYC gateway didn't respond? Because a rule engine took too long? Without integrated visibility, we're left guessing.

This fragmentation creates firefighting. IT gets called when something breaks, but without context. Business teams escalate based on outcomes, but they don't always know where the root issue lies. The result is delay, blame-shifting, and sometimes lost opportunities. We've learned that the only real way forward is to connect signals across the stack—app, infra, network, APIs—and make them business-aware.

What's equally important is how that insight is operationalized. It's not enough to have dashboards. We need to integrate observability into day-to-day workflows—issue detection, RCA, capacity planning, even change approvals. Our goal is to shift left—to surface potential problems earlier, reduce MTTR, and move from reactive resolution to proactive prevention.

There's also a cultural shift underway. In the past, experience was seen as a "front-end" responsibility—UI/UX, page load speeds, and so on. But we've realized that true experience is cross-functional. It's everyone's job. Infra teams must ensure uptime and scalability. Developers must build fault-tolerant, modular code. Business teams must define journeys that are efficient and intuitive. And IT operations must monitor it all in context.

One example we saw was in loan processing. A new digital initiative failed to meet expectations—not because the tech didn't work, but because one downstream system (hosted externally) kept timing out intermittently. Users experienced slowness and abandonment. Traditional monitoring didn't catch it, because each individual system showed "green." It was only after introducing journey-based observability that we saw the pattern. That's when it clicked: monitoring system health isn't enough—we must monitor business journeys.

Another lesson has been around tool sprawl. Many teams buy their own platforms—marketing has analytics, IT has infra monitors, DevOps has CI/CD tools, and so on. But the lack of a unified source of truth creates friction. We're now ensuring the tools we use speak to each other, and to us in a language we can act on.

Elevating Experience Starts with Business Context

Observability must translate tech metrics into meaningful business context and value.

In a customer-facing business like ours, where financial services and capital markets move in real time, the experience we deliver to users isn't just about convenience. It's about credibility. Every glitch, every delay, every broken journey affects not just the transaction, but trust. And once trust erodes, rebuilding it is far harder than preventing the issue in the first place.

I believe that we need a mindshift that observability isn't an IT initiative, it's a business enabler. We can't treat digital experience as something that gets measured after deployment or only when something goes wrong. It has to be part of the strategy, part of the design, and part of how we operate every day.

For that to happen, observability needs to speak the language of the business. It's not just about CPU utilization or error rates. It's about whether onboarding journeys are completing, whether customers are dropping off during KYC, or whether an advisory dashboard is rendering the right insights fast enough for our relationship managers. These are not "technical issues", they are business blockers, and they need to be treated that way.

What I often see is that while organizations have invested in various tools like APM, analytics, infrastructure monitoring, they still struggle to derive insights that are usable across functions. The signal-to-noise ratio is poor. Everyone has dashboards with tons of data, but few can tell a clear story about what's actually affecting users or revenue, which we call 'insight'. That's where the real opportunity lies: in intelligent observability that prioritizes what matters most.

Another challenge is the over-indexing on incident response, rather than preventive operations. Most environments still function with a "wait for alert, then fix" approach. But in a hyper-competitive, always-on world, that's too slow. We need to shift left, not just in DevOps pipelines, but in experience assurance overall. Can we identify the leading indicators of friction? Can we predict drop-offs or failures before they escalate? That's where AI and automation come in and not to replace human judgment, but to amplify it.

Security and compliance are also core to experience. A customer who feels their data isn't safe won't stick around, no matter how slick the interface is. We've built guardrails that align with regulatory expectations, but more importantly, we're trying to make those guardrails invisible to the user, so security doesn't create friction, but quietly reinforces trust.

For me, the ultimate goal is real-time, business-aware observability, where we can correlate what's happening in the system to what it means for the user and for the business. That's the gap we're working to close. And as we do, we're moving from firefighting to foresight and building experiences that not only function well, but feel effortless.

"Observability must speak the language of the business. Only then can it become a true enabler of digital experience."

Harsh Jha
Group Head of Technology,
Nuvama

"We're not just managing systems—we're curating journeys. Experience needs to scale, stay compliant, and still feel personal."

Kunal Dhingra
Chief Technology Officer,
RBL Bank

Scaling Experience Begins with Control, Context, and Confidence

Scalable experience demands observability that aligns control, context, and confidence.

In large-scale IT services and enterprise environments, digital experience has to deliver at three levels: performance, personalization, and predictability. But doing all three at scale—without compromising compliance or control—is where things get tricky.

From my perspective, personalization is no longer a luxury—it's expected. Whether the end user is an employee working remotely, a client engaging with our service platform, or a partner interacting through an API, people want contextual, responsive digital experiences that feel tailored to their needs. And the only way to deliver that reliably is through a solid platform strategy.

What that means practically is moving away from siloed solutions and toward platforms that are built to scale horizontally across services, geographies, and use cases. When you operate in a highly regulated sector like IT services or financial operations, you can't afford to have fragmented observability, inconsistent governance, or overlapping tooling. We've had to work toward consolidation and centralization, not just to cut costs, but to gain meaningful control.

At the same time, agility can't be sacrificed. Business units want speed. Clients want innovation. Developers want autonomy. And compliance teams want audit trails. Balancing all these demands means building guardrails, not gates—and embedding those into the platform itself.

One of the biggest challenges we've faced is around connecting infrastructure observability to actual user experience. Traditionally, we'd measure server health, application uptime, or network latency—but we couldn't always see how those translated into friction for the end user. Now, we're pushing for journey-based monitoring—where we look at experience end to end, from click to transaction to outcome.

This also means surfacing the right insights at the right level. A service desk operator might need alerts about queue times. A platform engineer might care about memory usage. But the business leader needs to know: how many users are abandoning this service? Why is productivity dropping in this geography? We've realized that observability has to be multi-layered—contextualized for each role, but stitched together from a common source of truth.

Security, too, is fundamental. But compliance also cannot be allowed to slow down innovation. So we're integrating real-time security observability—not just for attacks or intrusions, but for behaviors and patterns that hint at risk.

Looking ahead, I see the role of IT shifting from system owner to experience orchestrator. We're not just managing servers—we're curating journeys. And that means our platforms, our culture, and our metrics all need to evolve. If the digital experience isn't improving every month, something's wrong. And observability gives us the feedback loop to make sure we're always getting better.

From Data to Decisions—Engineering Digital Experience That Learns

Experience is engineered by observability that sees, learns, and adapts.

In an organization of our size and complexity, where services are delivered across a vast and distributed network, digital experience is not just a technical challenge—it's a systems thinking challenge. It involves aligning processes, data, people, and platforms in a way that allows the organization to not just serve users—but to learn from them in real time.

One of the biggest experience gaps we've faced historically isn't in the UI or feature set. It's in the disconnect between user behavior and system intelligence. We might know what services were accessed, but not why users dropped off. We could track uptime, but not intent. That's why we've had to move away from thinking about experience as a delivery problem, and start thinking about it as a data problem.

To do this, we've focused heavily on building event-driven architectures—where every action, transaction, or trigger leaves a trail that can be analyzed, aggregated, and acted upon. This gives us visibility not just into what's happening, but into what's likely to happen next. It also helps us identify which interactions matter most, so we can focus on optimizing journeys that move the needle for both users and the business.

Another area where we've invested is personalization. In our environment, users may not be digitally native, but they are digitally fluent—meaning they have expectations around speed, relevance, and support. Delivering personalized experiences in this context requires more than frontend logic. It means understanding the user's lifecycle, integrating data from multiple back-end systems, and responding to them with contextual nudges, pre-filled forms, or proactive guidance.

But none of that works without robust observability. You can't personalize what you can't see. That's why we've prioritized tools and platforms that offer end-to-end, real-time insight—not just into technical metrics, but into behavioral and process metrics. For instance, how long does a user spend stuck in a form? At what point do they abandon a service? What system or rule is creating hidden friction?

We've also had to shift the way we think about modernization. In large organizations, a rip-and-replace approach is rarely feasible. Instead, we've adopted a composable approach, where we modernize selectively, based on journey-critical systems. Some core elements may remain untouched, while others are exposed via APIs or re-platformed entirely. The key is to ensure that the user doesn't feel the seams, even if the backend is stitched together.

What gives me optimism is the growing alignment between business and IT on the importance of experience as a differentiator. We're no longer just talking about "projects"—we're talking about journeys, personas, and outcomes. That shift in language is powerful. It means we're not just building systems. We're designing experiences that learn, adapt, and continuously improve.

"You can't personalize what you can't see. Experience starts with architecture, but succeeds with observability."

Manish Malik
Executive Director (Information Systems), Indian Oil Corporation Limited

Digital Experience Design Is Grounded in Trust and Transparency

Trust is the experience—secured, localized, and built into systems.

"Experience is not a surface issue—it's a systems issue. And trust is the true measure of how well the system is working."

Om Prakash Seth
Chief Information Officer,
IDBI Bank

In today's environment, especially in sectors like public utilities and government-linked services, digital experience is about far more than interface design or technical uptime. It's about trust. And trust, once lostbreached, is difficult to rebuild.

For us, digital experience starts with a very simple but powerful question: Can the user customer rely on us? That reliability is measured in multiple ways—accuracy of information, consistency of service, data protection, and transparency of interaction. When someone uses our platform to apply for a subsidy or track their electricity usage, they aren't just performing a transaction. They're putting faith in a system to respond fairly, securely, and without unnecessary friction.

What this means in practice is that we cannot separate experience from governance. Consent, privacy, auditability—these are not legal checkboxes to be handled downstream. They are now central components of user trust. We are designing our systems and flows to incorporate privacy and security by design, so that data handling is visible, explainable, and minimal. In fact, our principle is: don't collect what you don't need, and don't store what you can't secure.

Another challenge we've addressed is scale. We serve millions of users, many of whom are in low-connectivity regions, using basic devices, or interacting in regional languages. This requires deep localization of experience—not just in language, but in design sensibility. For example, a three-step journey on a desktop might need to be a single-tap interaction on a mobile device with intermittent signal.

We've also invested in context-aware support. Instead of showing static FAQs or sending users into long support loops, we aim to anticipate what they might struggle with based on where they are in the journey. That means combining analytics with service logic, and using that insight to trigger helpful nudges—before the user gets frustrated or drops off.

Security is another layer of the experience that users don't always see—but definitely feel when it's absent. We've had to build strong, non-intrusive authentication and fraud detection mechanisms that work silently in the background. If the user is being protected, they shouldn't have to think about it. But they should feel it—in the form of seamless, confident interaction.

What's helped us in this journey is recognizing that digital experience is not a surface issue. It's a systems issue. It touches everything—from backend performance and inter-departmental workflows, to regulatory frameworks and citizen expectations. Our observability tools help us connect the dots: from system anomalies to user impact, from failed API calls to disrupted services. And they do so in real time, with actionable insights.

It's about Connecting Internal Experience to External Excellence

Empowered teams create better outcomes—internal experience drives external excellence.

In the realm of IT services, the emphasis often falls on client delivery, system uptime, and service level agreements (SLAs). However, what frequently remains underappreciated is the experience of the individuals responsible for these achievements: our internal users. For our organization, enhancing the digital experience is an endeavor that commences internally. When engineers, support teams, and operations staff can perform their duties efficiently, confidently, and without friction, it has a direct and positive impact on the client's experience.

One of the most significant challenges we face is the establishment of real-time visibility across our distributed teams and systems. With thousands of users, hundreds of applications, and continuous changes permeating the environment, minor glitches can have disproportionate effects. A delay in a deployment pipeline, an outage in a developer sandbox, or a permissions lag in onboarding—these issues may not make headlines, but they disrupt work, erode morale, and create downstream impacts for our customers.

Therefore, we have intentionally extended our observability efforts to encompass not only production environments but also engineering platforms and internal processes. We are meticulously tracking user interactions across crucial internal systems, including request portals, development environments, and incident dashboards, to identify and address any obstacles that impede productivity. This comprehensive approach provides valuable insights into both the health of our infrastructure and the overall employee experience and productivity.

We have also recognized that an increase in tools does not necessarily equate to greater control. At one juncture, we utilized multiple monitoring platforms, each providing different segments of data. However, without integration, the overall picture remained fragmented. Consequently, we have transitioned to a more integrated observability layer that unifies user behavior, system events, and business impact. Now, when an internal service fails, we can ascertain not only the occurrence of the failure but also its impact on users, their objectives at the time, and the implications for our operations.

Security is naturally a core concern as well. We're working in a sensitive environment with enterprise clients. So whatever we build—internally or externally—must meet compliance standards without slowing people down. That means building secure-by-design experiences that enable productivity while keeping risks in check.

It is evident that the distinction between internal and external experiences is becoming increasingly blurred. Inefficiencies in our internal processes inevitably impact client satisfaction, delivery schedules, and innovation cycles. Therefore, it is imperative for Chief Information Officers to scrutinize internal operations with the same rigor applied to external client engagements. The optimal client experience is rooted in a team that is empowered and adequately supported.

"The best client experience begins with an empowered, supported team. Internal journeys matter just as much as external ones."

Vikas Dureja
Vice President & Global IT Leader, HCL

From Performance to Business Impact: Rethinking the Digital Experience

Enterprises need to be empowered to deliver flawless digital experiences by unifying real-time observability with business intelligence, automation, and mobile-first performance.

"We're helping Indian enterprises move from fragmented monitoring to an all-in-one platform that connects front-end experience with backend performance, business metrics, and user behavior."

Meghana Shroff
Senior Customer Success Engineer, Dynatrace

In today's digital-first economy, a brand's performance is defined by the experience it delivers—across mobile apps, websites, portals, and embedded services. Consumers and employees alike expect seamless, responsive, and personalized interactions across devices and platforms. As digital experiences become the primary touchpoint between companies and customers, the margin for error has all but disappeared.

However, delivering exceptional digital experiences has grown more complex. The application landscape is fragmented across hybrid and multi-cloud environments, LLMs, Agents, microservices, APIs, and third-party integrations. The volume and velocity of data from user interactions, application telemetry, and infrastructure logs have exploded. At the same time, mobile-first user behavior—especially in regions like India—demands not only flawless functionality but also optimization across low-bandwidth and high-latency environments.

This convergence of user expectations and architectural complexity has given rise to the next evolution of monitoring: real-time, business-aware digital experience observability. The ability to understand and optimize every digital journey is becoming a strategic differentiator for enterprises across sectors—from banking and retail to telecom and insurance.

What's Holding Enterprises Back from Experience-Led Growth

Despite the growing urgency, CIOs and digital leaders continue to face several

Faisal Shaikh
Regional Director,
Dynatrace

"Digital experience is the new balance sheet. Every slow page, every broken journey is no longer just a tech issue—it's a business risk."

persistent challenges in elevating digital experiences:

- **Fragmented Visibility.** Traditional tools monitor individual components but fail to connect the dots across user journeys, APIs, application tiers, and infrastructure layers. This leads to blind spots and delayed diagnosis of issues that affect customers.
- **Lack of Business Context.** Metrics like page load time or server CPU usage are disconnected from business KPIs such as conversion rate, NPS, or drop-offs. This makes it hard for IT teams to prioritize and for business leaders to make confident decisions.
- **Reactive Operations.** Without predictive insights, teams are stuck in a firefighting mode. Problems are often discovered by users before internal teams notice, eroding trust and brand equity.
- **Mobile Complexity.** With India's mobile-dominant user base, performance challenges are amplified in real-world scenarios—high device fragmentation, varied network quality, and app compatibility issues.
- **Data Overload and Dashboard Fatigue.** While dashboards exist, they are often static, overloaded with irrelevant alerts, and fail to empower business and IT users with actionable, real-time intelligence.
- **Compliance and Trust.** With new data privacy and digital security regulations, digital experience efforts must also ensure secure data handling, traceability, and audit readiness.

Connecting Experience to Outcomes: A New Approach

The Dynatrace platform takes a fundamentally different approach by unifying all critical layers of digital experience—front-end interaction, application performance, back-end services, and business outcomes—into a single intelligent observability framework.

At the core is causal AI, which not only identifies issues across the full stack but also understands the "why" behind them. This provides precise root-cause analysis without the noise or false positives that plague legacy monitoring tools.

Real User Monitoring (RUM) captures every user interaction in real time—clicks, hovers, taps, and load times—providing deep visibility into individual journeys. Synthetic monitoring supplements this with continuous, automated test journeys across geographies and networks to ensure uptime, responsiveness, and SLA compliance. Application performance monitoring (APM), integrated with distributed tracing, gives granular insights into how code, services, and third-party components contribute to user experience. Further, Dynatrace enables teams to find and resolve issues in real time with logs, traces, and metrics automatically in context.

It is essential to emphasize the shift from technical dashboards to business-aligned observability. Dynatrace links performance metrics with outcomes such as NPS, transaction values, and customer sentiment, allowing IT and business teams to speak a common language. This shift is already helping leading enterprises in India reduce app crashes, improve conversion, and personalize digital touchpoints.

Digital experience is no longer a UX or IT metric—it is now the new balance sheet. Dynatrace enables CXOs to track perfor-

mance with the same rigor as financial KPIs, using real-time, AI-powered dashboards. The platform supports compliance and data privacy mandates by embedding security intelligence across monitoring layers, supporting trust and resilience.

Moreover, Dynatrace enables proactive action. Instead of reacting to failures, organizations can predict degradation based on user behavior patterns, release cycles, and real-time telemetry. This supports faster innovation cycles and prevents business impact.

Practical Strategies to Accelerate Digital Experience Maturity

From our deep engagement with Indian enterprises, we can highlight the following best practices for organizations seeking to elevate their digital experience maturity:

- **Anchor Digital Experience to Business Metrics.** Start by identifying which user journeys drive business outcomes—account opening, checkout, and fund transfer—and map observability to those journeys.
- **Leverage log data.** Teams can automatically extract business data from log files, in context.
- **Prioritize Mobile Experience Monitoring.** Especially in mobile-first markets, invest in crash analytics, performance baselines for low-bandwidth regions, and AI-driven prioritization of user-impacting bugs.
- **Break the Dashboard Paradigm.** Replace static dashboards with dynamic, role-based views that provide business-aligned insights and guide action, not just observation.
- **Invest in AI for Real-time Decisions.** Use AI-powered root-cause analysis and anomaly detection to reduce MTTR and elevate IT from reactive support to a predictive business enabler.
- **Create Shared Responsibility.** Digital experience is not owned by IT alone. Establish cross-functional squads, including product owners, marketers, SREs, and CX leaders, to monitor and improve journeys collectively.
- **Automate and Scale.** Automate experience validation, synthetic testing, and self-healing mechanisms to support growth without overwhelming teams.

By aligning observability with user experience, business performance, and operational agility, Dynatrace empowers enterprises to move from measuring digital experience to mastering it. In an environment where milliseconds can mean millions, this capability is no longer optional—it is transformative.

"The greatest danger in times of turbulence is not the turbulence; it is to act with yesterday's logic."

—Peter Drucker

Index

Presented in Collaboration with Dynatrace

In a world where software powers everything, Dynatrace transforms the complexity of modern digital ecosystems into powerful business assets. As the leading AI-powered observability platform, Dynatrace enables organizations to analyze, automate, and innovate—faster and more intelligently.

With the rise of generative and agentic AI, the need for precision, governance, and real-time insights has never been greater. Dynatrace is built for dynamic, cloud-native environments and seamlessly integrates with AWS to provide context-rich observability and intelligent automation across every layer of your stack. Whether you're deploying LLMs, managing autonomous agents, or scaling AI workloads, Dynatrace delivers the clarity and control needed to ensure performance, reliability, and trust.

From real-time discovery and continuous topology mapping to release validation and runtime threat detection, Dynatrace empowers teams to deliver flawless digital experiences while accelerating cloud operations.

As an AWS Advanced Technology Partner with seven AWS competencies—including

"Whether you're deploying LLMs, managing autonomous agents, or scaling AI workloads, Dynatrace delivers the clarity and control needed to ensure performance, reliability, and trust."

Cloud Operations—Dynatrace helps organizations align with the AWS Well-Architected Framework, delivering high performance, security, and cost optimization across AI-driven workloads.

For more information on Dynatrace, see AWS monitoring (dynatrace.com).

FUTURESCAPE 2025 is a definitive thought leadership initiative designed for IT decision-makers in large and mid-sized Indian enterprises as they chart their course toward intelligent digital transformation.

Produced by CIO&Leader, the flagship B2B enterprise technology publication of 9.9 Group, in association with Dynatrace—a global company that empowers organizations to analyze, automate, and innovate faster by leveraging AI-powered insights to drive their business forward—this book captures the collective expertise of senior enterprise IT leaders.

Its content is based on a series of deliberations involving these technology leaders and senior executives from Dynatrace, moderated by the editors of 9.9 Group.

As enterprises accelerate toward platformization, automation, and business-aligned IT operations, FUTURESCAPE 2025 serves as both a practical and strategic guide, helping leaders reimagine observability, scale DevOps maturity, align digital experience with business outcomes, and unlock AI's full potential in production environments.

9.9 Group is an innovative media organization that engages the B2B technology community through cutting-edge content and forward-looking engagement formats.

dynatrace | CIO&LEADER | aws

9 789361 944307

₹7,500