AI Leaders

Expert Exchange

Executive summary
March 4, 2026

kyndryl

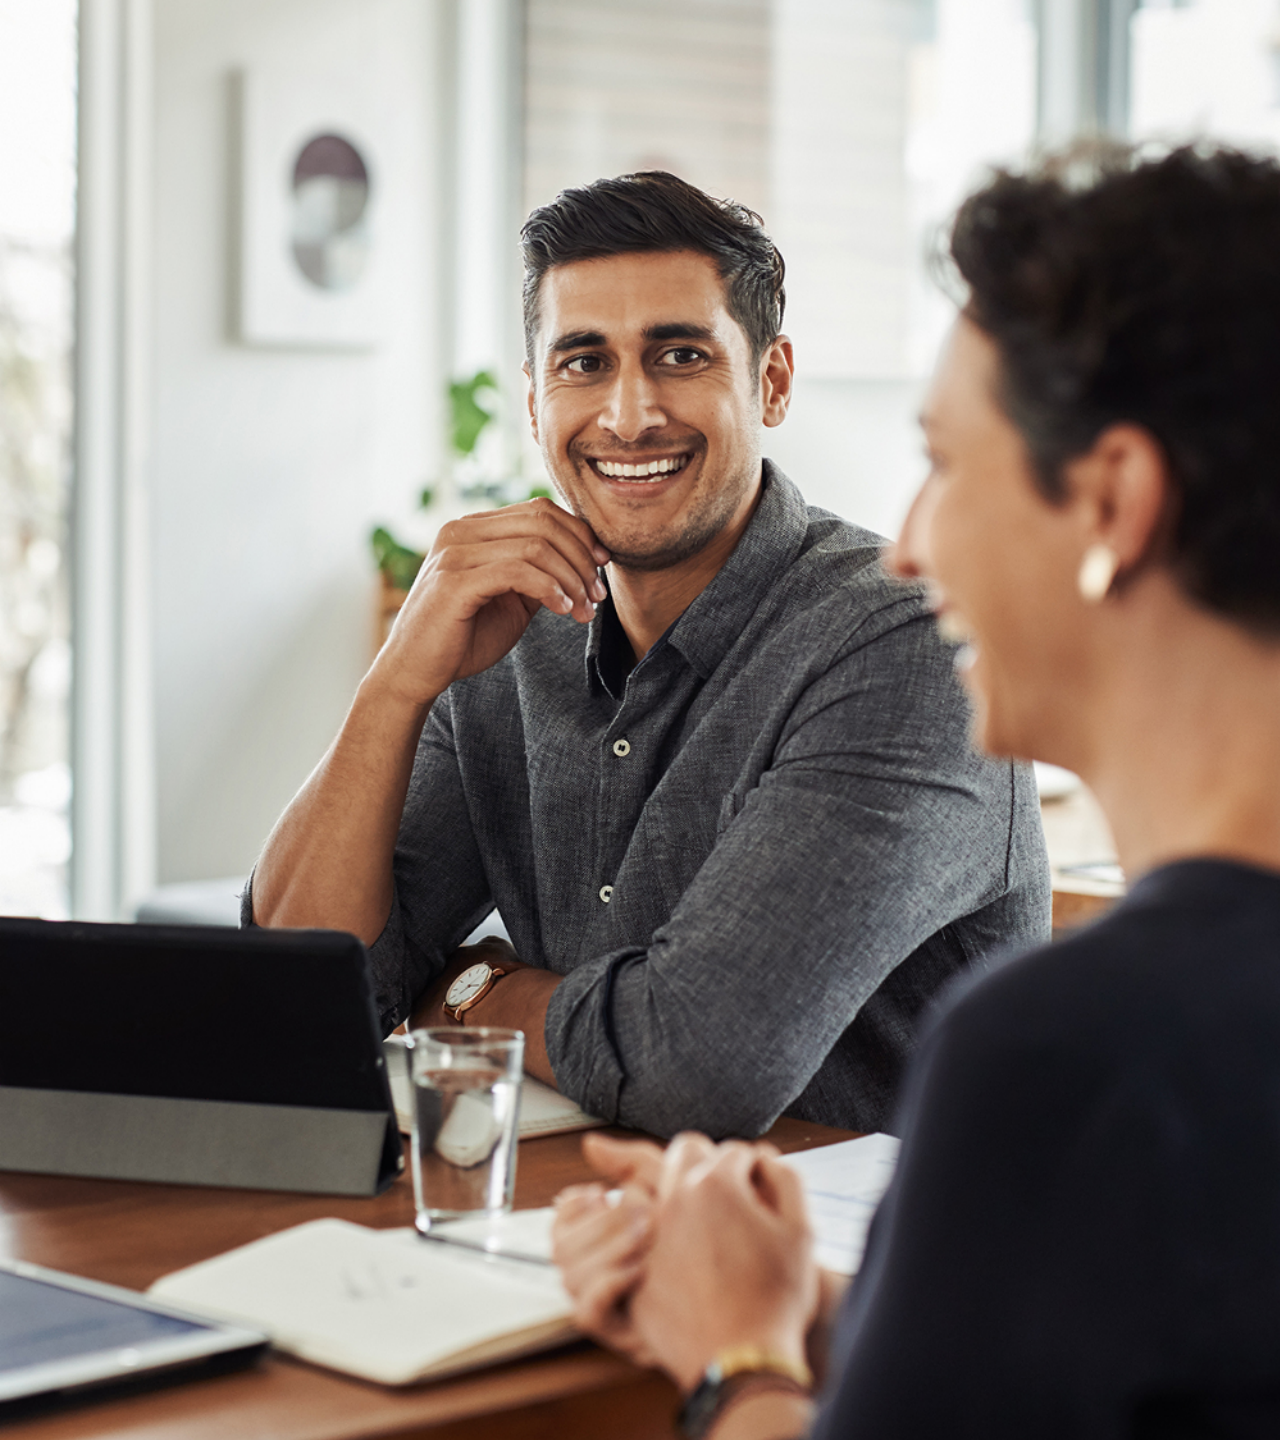

Overview

Senior technology leaders recently met at the Kyndryl AI Leaders Expert Exchange to discuss AI adoption, security, and practical deployment strategies. Key topics included managing security risks in rapid code generation, ensuring human oversight in automated workflows, and weighing “build versus buy” decisions when evaluating vendor solutions.

Host

Rob Reynolds
Vice President, Data and AI
US Consult

Key topics

Kyndryl AI Leaders Expert Exchange: Executive brief

Key takeaways

- Leaders aligned on rising security and operational risk from rapid AI-driven development
- Human oversight remains essential to maintain context, quality and regulatory control
- Vendor and architecture decisions must balance speed, compliance and long-term flexibility

Implications

- Security risk now scales with automation velocity, not headcount
- Operating model discipline is as critical as AI tooling
- Architecture choices determine future agility and control

Key topics

Security and risk

- AI accelerates code creation faster than traditional security models can absorb
- Leaders split on retrofitting controls versus rethinking the SDLC end-to-end
- Embedding security and architecture rules into AI inputs reduces downstream risk

Human oversight

- Automation without intervention quickly degrades business context
- Regulated industries require tighter human control thresholds
- Maturity path: full review → monitored confidence-based oversight

Vendor and architecture

- Demand for experimentation conflicts with enterprise compliance needs
- Debate: rigorous evaluation versus rapid adoption to avoid market lag
- Agnostic architectures protect against lock-in and enable platform switching

“The development time is shorter, but the risk is higher because humans cannot keep up with the rate at which automated tools turn out code.”

— AI Leaders Expert Exchange Member

Security and risk mitigation in automated code generation

- Participants overwhelmingly agreed that rapid code generation introduces unprecedented security challenges, prompting a debate between relying on existing testing procedures versus reimagining the software development process.

- Some expressed deep concern that traditional security teams cannot keep pace with the sheer volume of code generated by automated systems. They noted that independently testing identical code through multiple distinct security scanners yielded hundreds of newly introduced vulnerabilities each time, highlighting significant enterprise blind spots.

- Countering this alarm, others argued that these fundamental vulnerabilities already exist with human developers. They suggested organizations should rely on their established testing workflows to mitigate these issues rather than viewing automated generation as a threat requiring totally new security paradigms.

- One member offered a diverging view, arguing that simply retrofitting new tools into legacy testing processes is entirely insufficient. Instead, they proposed that organizations must completely reimagine their software development frameworks from initial requirements to final deployment to ensure automation is safely integrated.

- Addressing the broader impact, another executive questioned whether automated tools actually deliver meaningful speed improvements for complex enterprise applications. They noted that while simple tasks accelerate, traditional enterprise development remains highly complex, requiring continuous staff education to safely manage new security vulnerabilities applications. They noted that while simple tasks accelerate, traditional enterprise development remains highly complex, requiring continuous staff education to safely manage new security vulnerabilities.

> "The development time is shorter, but the risk is higher because humans cannot keep up with the rate at which automated tools turn out code."
>
> — AI Leaders Expert Exchange Member

Balancing automation with human oversight

- The group reached a strong consensus that human intervention remains critical to operational success, though participants differed on how to best implement this oversight across varying organizational maturity levels.

- Kyndryl Host Rob Reynolds emphasized that unchecked automated systems quickly lose business context and introduce unwanted errors if left to build independently. He advised organizations to mandate continuous collaboration between workers and systems to ensure final outputs remain highly accurate and contextually relevant to the business.

- Highlighting varying industry risk tolerances, a member noted that highly regulated organizations cannot irresponsibly deploy fully automated operations into live production environments. They suggested that oversight procedures must remain highly flexible, adapting specifically to the experience level of the staff and the strictness of the surrounding regulatory environment.

- One member shared a progressive maturity framework for gradually reducing manual oversight over time. They recommended starting with complete human review to build operational confidence scores, eventually transitioning to a monitored model where humans only review the final outputs once the system consistently proves its reliability.

- Others cautioned that relying too heavily on external vendors for automated application support creates massive organizational blind spots. They argued that organizations must hold their senior internal developers fully accountable for totally understanding the generated code before assuming any vendor modifications are safe to implement.

- Expanding on the need for continued human involvement, Kyndryl Host Rob Reynolds noted that automated testing demands high-quality initial inputs from human workers. He explained that true acceleration requires humans to write exceptionally clear initial requirements, which the automated system then uses to create rigorous test cases and self-healing frameworks.

> "We should use automated tools to amplify our people, not replace them, because the longer you let the system operate without human intervention, the further it loses context."
>
> — AI Leaders Expert Exchange Member

Vendor selection and strategic architecture

- The conversation revealed diverging strategies for acquiring external technology, contrasting the desire for exhaustive internal vendor testing against the immediate business need to deploy enterprise-grade solutions quickly.

- Members highlighted the severe operational bottleneck caused by employees requesting numerous unapproved experimental tools. They questioned how compliance-focused organizations could safely balance the overwhelming volume of employee requests to experiment while still enforcing rigid enterprise security standards across the business.

- To successfully address restrictive vendor lock-in, several executives advocated for building agnostic architectural frameworks that allow companies to seamlessly switch between different foundational technologies. A healthcare executive noted this flexible infrastructure is crucial for protecting sensitive medical records, as it strictly controls user access.

- One member shared a highly comprehensive evaluation strategy, recommending that organizations start broadly by testing multiple competing tools in secure, isolated environments. Following rigorous legal, internal licensing and strict security reviews, the organization eventually narrows the broad selection down to one or two scalable enterprise platforms.

- Conversely, one member forcefully argued against prolonged internal testing periods, asserting that resource-constrained companies absolutely cannot afford to waste months evaluating rapidly changing tools. They firmly advised selecting a compliant enterprise platform quickly and deploying it immediately to avoid significantly falling behind market advancements.

- In further discussion of corporate acquisition strategies, a senior executive emphasized that contemporary organizations should adopt a more nuanced approach rather than limiting themselves to the binary decision of building or buying. They suggested companies critically evaluate whether they explicitly need to buy, build, or extensively extend existing vendor platforms by integrating them deeply with proprietary corporate data.

> "Organizations should universally buy foundational tools for general enterprise use, and build customized solutions only where they possess a unique competitive advantage."
>
> —AI Leaders Expert Exchange Member

Agentic AI at scale: The CD&AI scorecard

Agentic AI delivers CD&AI success only when speed, trust and governance are engineered together.

Time-to-value and scale

(From pilots to enterprise workflows)

What CD&AI leaders care about:
Speed from concept to production; ability to scale beyond isolated use cases

What the research shows:

- AI-native enterprises move beyond disconnected pilots by embedding AI at the core of workflows, enabling faster sensing, decision-making and action
- Agentic AI enables continuous execution across systems — but only when designed for enterprise-wide orchestration, not one-off copilots

Implication: Speed comes from designing agentic AI as a repeatable operating model, not as experiments.

Reliability of outcomes

(Sustaining trust in AI decisions)

What CD&AI leaders care about:
Consistency of decisions; confidence in AI-driven outcomes over time

What the research shows:

- Agentic AI introduces new risk surfaces — including adversarial prompts, poisoned data and opaque decision paths that traditional application security does not address
- Static, perimeter-based controls fail when AI systems are adaptive, autonomous and continuously learning

Implication: Outcome reliability requires continuous controls embedded across the AI lifecycle, not post-hoc validation.

Governance and auditability by design

(Scaling AI without losing trust)

What CD&AI leaders care about:
Explainability, audit readiness, regulatory alignment

What the research shows:

- Policy as code translates business rules and regulatory requirements into machine-enforceable guardrails that govern how AI agents act
- This creates traceable, reviewable, and explainable agent decisions — critical for regulated and mission-critical environments

Implication: Governance must move from documentation to execution-level enforcement.

Sustainable operating model

(Autonomy with control, humans with agents)

What CD&AI leaders care about:
Adoption at scale; clear decision rights; manageable operating cost

What the research shows:

- AI-native organizations are built around human–agent collaboration, with clear boundaries on what agents can do independently and when humans intervene
- Treating security and governance as a final checkpoint increases friction and cost; agentic AI requires security and policy embedded from design through runtime

Implication: The winning model balances autonomy and oversight — enabling scale without chaos.

kyndryl

Additional Resources:

Security

- Redefining security for the agentic AI era

AI and policy as code

- Journey to AI Native
- Policy As Code - Article
- Policy As Code Homepage

The AI Expert Exchange is hosted by Kyndryl. Please contact Rob Reynolds (rob.reynolds@kyndryl.com) with any questions about Kyndryl or this Exchange.