F5 regional CXO roundtable series

New York edition

Architecting the AI-enabled enterprise

Key takeaways | June 17, 2025

Lessons from the New York CXO roundtable

An actionable path for scaling AI for business outcomes

Executive summary

The F5 regional CXO roundtable in New York convened senior executives from leading global financial and technology institutions to explore strategies for scaling artificial intelligence (AI) with trust, agility, and measurable impact. Held in a market defined by significant operational and regulatory complexity, the discussion moved decisively beyond experimentation, focusing on the strategic, governance, and leadership imperatives required to integrate AI as a transformative force across enterprises.

The roundtable crystallized a shared understanding: scaling AI is a strategic and organizational priority, not solely a technological one. Successful adoption demands reimagined operating models, enterprise-wide talent development, and architectures designed for transparency, security, and resilience. These elements are critical to embedding AI as a core driver of business value and competitive differentiation.

Insights from live polling during the session revealed clear alignment among participants:

- **Decentralized innovation with robust governance:** A majority favored empowering business units to drive AI use case prioritization, supported by centralized governance to ensure strategic coherence and risk mitigation.
- **Real-time observability:** Most executives emphasized the need for continuous monitoring of AI systems' interactions with enterprise platforms, highlighting observability and control as essential for trust and performance in live environments.
- **Advancing AI maturity:** The majority of organizations represented are actively scaling AI adoption, signaling an urgent shift from pilots to enterprise-wide deployment to sustain market leadership.

To support this shift, the group outlined a structured approach around five focus areas:

1. **Strategic imperatives**: Align AI to enterprise goals and balance value creation across efficiency and growth. Discussions surfaced seven imperatives, including securing LLM access, deploying circuit breakers for agentic AI, embedding explainability, and redesigning talent models for an AI-augmented workforce.

2. **Critical challenges**: Address fragmentation in governance, data quality, vendor transparency, and talent readiness. Participants cited siloed ownership, uncurated data, opaque SaaS integrations, and cost unpredictability as major scale inhibitors.

3. **Implementation plan**: Design scalable AI systems built on secure, modular, and interoperable foundations. Emphasis was placed on centralized governance, federated execution, API-first integration, and telemetry-driven oversight.

4. **Success metrics:** Measure business value, operational impact, adoption, and regulatory auditability, not just model performance. Leaders advocated for KPIs tied to ROI, trust, and explainability across functions.

5. **Next steps:** Institutionalize hybrid governance, invest in workforce readiness, and embed explainability and safety into every deployment, from internal workflows to customer-facing use cases.

This session reaffirmed that the enterprises leading the AI curve are not those with the most models, but those with the clearest vision, strongest governance, and boldest commitment to responsible innovation.

1. Strategic imperatives for AI adoption

The roundtable surfaced seven forward- looking priorities to help enterprises scale AI responsibly, securely, and at speed.

1.1 Govern LLM access as a first-class security perimeter

Insight: Training and fine-tuning models is no longer the challenge—governing who can query them is. As LLMs ingest sensitive assets like system architecture, credentials, and customer data, access control becomes the most exposed risk vector. Participants highlighted multiple cases where internal teams fed critical enterprise data into AI systems without understanding the consequences. In heavily regulated environments, such actions have direct compliance implications.

Recommendation
Treat LLMs as privileged interfaces and apply security and governance controls at the access and prompt layer—not just the model layer.

Actions
- Implement role-based access and intent-aware prompt controls.
- Tokenize sensitive inputs and encrypt AI query interactions.
- Monitor usage in real time and align access with compliance standards (e.g., SR 11-7, GDPR equivalents).

1.2 Build circuit breakers to contain agentic AI risk

Insight: Agentic systems, where AI agents interact, act autonomously, or call downstream APIs, pose new types of risk. Participants compared this to flash crashes in algorithmic trading, where runaway logic caused systemic breakdowns. Without guardrails, agent-based AI could create "crashes in logic," leading to decisions that spiral outside expected behavior or business policy.

Recommendation

Deploy circuit breakers and safety protocols across agentic AI systems to detect, contain, and shut down anomalous interactions in real-time.

Actions

- Define logic thresholds and failure conditions for agents.
- Simulate multi-agent failures and agent-to-agent escalation scenarios.
- Require kill switches, shutdown triggers, and oversight rules in all internal and vendor-driven agentic systems.

1.3 Make auditability and accreditation core to AI lifecycle

Insight: As AI enters high-stakes use cases such as onboarding, compliance checks, and credit decisions, auditability and explainability are becoming essential. However, the increasing reliance on foundational models (often external and opaque) makes it hard to trace how an outcome was generated. Leaders voiced concern over being unable to reproduce outputs or satisfy regulators in the event of an audit.

Recommendation

Develop an enterprise-wide model accreditation framework with layered explainability, lineage tracking, and reproducibility safeguards.

Actions

- Track data provenance and maintain version-controlled model repositories.
- Embed explainability dashboards across all business-integrated models.
- Align internal audit trails with external regulatory requirements (e.g., SEC, SR 11-7).

1.4 Enable federated innovation on centralized foundations

Insight: Organizations are struggling with fragmented AI efforts—business units moving fast without architectural alignment, while central teams struggle to keep up with governance. Mature firms have addressed this through federated AI governance, where centralized teams provide common tooling, platforms, and policies, while business teams drive execution with agility.

Recommendation
Adopt a hybrid model with centralized CoE for governance and decentralized innovation led by business units.

Actions
- Establish a central AI CoE with platform ownership and governance authority.
- Let business units own problem statements, datasets, and KPIs.
- Promote reuse of approved APIs, models, and security patterns across units.

1.5 Redesign talent models for an AI-augmented enterprise

Insight: AI is eliminating entry-level tasks at scale, including analysts, legal researchers, junior ops—all historically essential for developing context and judgment. Without action, this shift risks long-term capability erosion. Leaders warned of a "death of apprenticeship," and acknowledged that new roles must emerge to supervise, guide, and partner with AI.

Recommendation
Reimagine career paths and redefine roles to integrate human expertise with AI supervision, not just automate tasks.

Actions
- Launch AI fluency and prompt engineering programs enterprise wide.
- Redesign roles around human-in-the-loop responsibilities and AI oversight.
- Embed AI usage and literacy into KPIs, especially for emerging and frontline talent.

1.6 Build consumer trust through transparency and ethics

Insight: As AI becomes embedded in customer journeys, from onboarding to personalization, trust is under pressure. Participants noted a growing consumer reluctance to engage with AI systems when data usage isn't clear. Opaque AI can undermine brand trust, invite regulatory scrutiny, and slow adoption, especially in finance, health, and insurance.

Recommendation
Treat transparency and ethical use as core design requirements, not optional add-ons, in all customer-facing AI deployments.

Actions
- Define ethical guardrails (e.g., opt-outs, human fallback, data limits).
- Apply privacy-preserving techniques like homomorphic encryption.
- Publish AI usage disclosures and allow customers to review or escalate decisions.

1.7 Fund AI with strategic intent and enterprise ownership

Insight: AI investments today are often siloed, shared between CIOs, CTOs, and BUs, resulting in overlapping projects and misaligned priorities. Some firms have addressed this by appointing a Chief Transformation Officer to align AI to business strategy. Others use a two-tier budget model: central funding for tooling and governance, decentralized funding for business unit use cases.

Recommendation
Anchor AI as a strategic business investment, with clear ownership and ROI frameworks at the enterprise level.

Actions

- Define budget accountability and executive sponsorship for AI.
- Allocate centralized funds for platform, security, and compliance.
- Let business units prioritize high-impact AI projects with outcome-specific metrics (e.g., CX, fraud, efficiency).

2. Critical challenges

While the ambition to scale AI was unanimous, leaders surfaced deep-seated organizational and operational barriers that continue to slow progress. These challenges must be addressed head-on to avoid fragmentation, risk, and stalled momentum.

2.1 Siloed accountability across business and control functions

Mitigation: Establish cross-functional AI governance councils that include legal, risk, security, HR, and data leaders to drive consensus.

Action: Form an enterprise AI decision forum with predefined voting rights and escalation paths; assign a cross-functional program leader to align priorities.

2.2 Fragmented and uncurated data foundations

Mitigation: Invest in building centralized, curated, and governed data layers accessible across business units.

Action: Standardize data pipelines and metadata schemas; appoint data product owners in each business unit with clear stewardship responsibilities.

2.3 Vendor-led AI with limited transparency or control

Mitigation: Mandate transparency, auditability, and opt-out provisions from SaaS vendors embedding LLMs.

Action: Renegotiate SaaS contracts to include model explainability rights, data usage policies, and BYOK (Bring Your Own Key) provisions.

2.4 Cost unpredictability driving AI infrastructure reassessment

Mitigation: Map workloads to cost sensitivity and explore hybrid or on-prem deploymentsfor high-volume inference.

Action: Conduct a cost-control audit of cloud AI usage; identify candidates for repatriation based on performance, sovereignty, and financial volatility.

2.5 Absence of standardized AI governance in absence of regulation

Mitigation: Proactively create internal accreditation, explainability, and risk scoring frameworks tailored to your business.

Action: Establish an internal AI governance framework aligned to SR 11-7 principles, covering model approval, monitoring, and audit readiness.

3. Implementation plan: Scalable AI architecture

Scaling AI, as participants emphasized, requires clear design principles, resilient architecture, and seamless integration into enterprise systems.

3.1 Design principles

Modularity: Architect AI systems using interchangeable models, agents, and components to enable agility, experimentation, and failover.

Interoperability: Design AI workflows to integrate seamlessly with existing business systems, SaaS platforms, and infrastructure using standardized APIs.

Security: Embed zero-trust principles, encryption, and prompt-level controls across all layers—from training data to LLM deployment.

Observability: Build real-time visibility into model behavior and decision logic, with full lineage and traceability to support audits and trust.

3.2 Architecture components

Data layer

- Unified and governed data pipelines shared across business units.
- Metadata tagging, real-time validation, and lineage tracking.
- Integration with enterprise data catalogs for discoverability and access control.

AI layer

- Modular LLMs and agents with embedded explainability and safety mechanisms.
- Agent supervision via escalation thresholds and circuit breakers.
- Model accreditation workflows with version control and usage logs.

Governance layer

- Centralized AI CoE for tooling, risk, and compliance oversight.
- Model approval policies aligned with regulatory expectations (e.g., SR 11-7).
- Continuous evaluation of ethical guardrails and audit readiness.

Integration layer

- API-first architecture for AI-to-system interoperability.
- Compatibility with legacy and cloud-native environments.
- Auditable input-output flows connecting prompts to outcomes across systems.

3.3 Implementation steps

- Stand up a centralized AI CoE to define guardrails, tooling standards, and compliance workflows.
- Build governed, real-time data pipelines with lineage and metadata tagging.
- Deploy modular, explainable LLMs and agents with safety and drift-monitoring built-in.
- Integrate AI via secure, observable APIs into existing business workflows and SaaS platforms.
- Embed audit trails and telemetry across models, prompts, and decisions to ensure traceability.

4. Success metrics

Participants emphasized that success must be defined beyond model accuracy. Metrics must capture business value, trust, auditability, and enterprise-wide adoption.

- **Business value creation**: Revenue uplift from AI-driven personalization, faster time-to-market, and process automation.
- **Operational efficiency**: Reduction in manual work, increased agent-led execution, and shorter resolution cycles.

- **AI adoption across roles**: AI fluency KPIs integrated into employee performance, including prompt engineering and agent oversight.
- **Regulatory audit readiness**: Reproducible model outputs, traceable decisions, and explainability dashboards aligned with SR 11-7.
- **Innovation scalability**: Number of production-grade AI use cases deployed; rate of agent reuse across business units.

5. Next steps

From pilot to production, these priorities define what leaders must act on to scale AI responsibly.

- **Elevate AI to a board-level strategic agenda**: Position AI as a business transformation priority, not a technical initiative. Ensure ownership rests with C-level leaders.
- **Institutionalize hybrid governance models**: Centralize platform governance, compliance, and risk via an AI CoE. Empower business units to innovate under shared architectural and ethical standards.
- **Establish an AI Architecture Review Board (ARB)**: Implement a formal ARB to review AI initiatives for risks such as algorithmic bias, model drift, data privacy, and operational dependencies, ensuring mitigation strategies are in place before deployment.
- **Operationalize explainability and auditability**: Deploy real-time dashboards that show model rationale, training lineage, and data inputs, enabling regulatory and internal transparency.
- **Embed safety mechanisms into all agentic deployments**: Design for fail-safes by default, circuit breakers, shutdown triggers, and escalation workflows must be standard across agents.
- **Redesign talent architecture around AI fluency**: Develop AI literacy programs, redefine career paths for AI-augmented roles, and incorporate AI performance into KPIs at all levels.

Attendees

Name	Company	Designation
Anthony Bussanich	J.P. Morgan	Executive Director - AI Strategy & Enablement
Brijesh Singh	Wipro	Chief AI and Strategy Officer
Celso Yamaguchi	Safra National Bank	Chief Technology Officer
Daryl Martis	Salesforce	Director of AI Product
Felix Jorge	Hitachi Vantara	Field Chief Technology Officer
Linda Powell	BNY	Enterprise Deputy Chief Data Officer/ Head of Data Governance
Paolo Pelizzoli	The Clearing House Payments Company LLC	Head of Innovation and Platform Engineering
Puneet Bhatnagar	Blackstone	Senior Vice President, Head of IAM - Cybersecurity, Blackstone Technology & Innovations
Roger Burkhardt	Broadridge	Chief Technology Officer
Samantha Sprague	Discover Financial Services	Expert Architect - Data, Machine Learning, and Generative AI
Srini Masanam	Citi	Global Head of Surveillance Data Quality
Suzy Wassef Fahmy	Spectrum Business for Enterprise	Field CTO, Digital Transformation/ Modernization Consultant
Vijay Kukreja	Bank of America	Vice President
Wim Verleyen	Humana Inc	Associate Vice President, Data Science, Research & Development
Xiaoyang (Yvonne) Clarke	Wells Fargo	Program Manager, AVP