F5 regional CXO roundtable series

Hong Kong edition

Architecting the AI-enabled enterprise

Key takeaways | July 16, 2025

Lessons from the Hong Kong CXO roundtable

An actionable path for scaling AI for business outcomes

Executive summary

The Hong Kong edition of the F5 CXO Roundtable brought together senior leaders across banking, telecom, education, cybersecurity, and public sector organisations to explore how enterprises can move beyond AI experimentation and deliver meaningful business outcomes at scale.

While AI has entered the strategic agenda for most organizations, many remain in the early phases, limited to proof of concepts and internal trials. The roundtable emphasized that scaling AI is not just a technological upgrade. It is a transformation that demands cross-functional leadership, deliberate governance, and enterprise-wide capability building.

Three cross-cutting themes emerged

- Governance must evolve from reactive to embedded. Most risk frameworks remain manual and disconnected from development workflows, increasing enterprise exposure as AI use accelerates.
- People, not platforms, will determine success. Cultural resistance, unclear ownership, and lack of enablement are stalling momentum. Trust and capability must be cultivated across all layers of the organization.
- Architectural design is lagging ambition. Fragmented tooling, inconsistent metrics, and siloed pilots are slowing progress toward unified, scalable deployment.

To address these challenges, the discussion focused on five key areas.

1. **Strategic imperatives**: Embedding AI into the fabric of business operations through steering committees, centralized funds, and market-aligned capital plans.
2. **Critical challenges:** Overcoming barriers such as regulatory compliance, shadow AI, fragmented ownership, and cultural pushback through structured enablement and clear mandates.
3. **Implementation architecture:** Building modular, secure, and interoperable AI stacks that are resilient by design, accessible across environments, and shaped by evolving geopolitical realities.

4. **Success metrics:** Shifting focus from model accuracy to enterprise value by tracking adoption, use case impact, value creation, governance maturity, and scaling readiness.
5. **Next steps:** Institutionalizing learnings and scaling trusted pilots by starting with a centralized AI model and gradually transitioning to a federated structure as governance frameworks mature, with tailored capability programs across functions.

The call to action is clear: Organizations that lead with trust, unify leadership, and invest in scalable architecture will not just adopt AI. They will redefine their industries with it.

1. Strategic imperatives for enterprise-scale AI adoption

To unlock enterprise-wide value from AI, organizations must move beyond experimentation and lay the foundations for scale. The following imperatives outline the leadership, governance, and investment shifts required to embed AI into the core of business operations.

1.1 Drive change through citizen innovation

Insight: AI capabilities are advancing faster than enterprise readiness. The core challenge is not technology, but internal resistance and a lack of structured enablement for business users.

Recommendation

Foster a culture where AI adoption is democratized across functions through hands-on access, guided experimentation, and visible leadership support.

Actions

- Deploy citizen AI enablement programs targeted at core business users (e.g. HR, finance, ops). Start with curated, low-risk tools such as copilots, no-code builders, and workflow assistants.
- Integrate AI fluency into leadership development and performance management. Treat digital proficiency as a promotable skill.
- Create structured sandboxes with real datasets, model templates, and support channels to encourage safe experimentation.

1.2 Operationalize AI governance as code

Insight: Most governance approaches today are reactive, manual, and disconnected from development workflows. Risk teams are underpowered, and governance lacks real-time enforcement.

Recommendation

Transition from manual governance models to embedded, real-time controls aligned to model risk and business exposure.

Actions

- Define a three-tiered risk classification model based on data sensitivity and model autonomy.
- Embed red-teaming, audit trails, and explainability into the development pipeline.
- Implement governance as code through CI/CD integration and automated policy enforcement.

1.3 Set clear and central AI budget ownership

Insight: AI-related costs are often spread across IT, business units, and innovation teams. In some cases, small teams are self-funding POCs, while strategic investments face delays due to unclear ownership and ROI expectations.

Recommendation

Establish a dedicated capital allocation strategy that aligns funding with enterprise goals and creates transparency across initiatives.

Actions

- Establish a centralized AI innovation fund with cross-functional governance
- Standardize investment templates for use-case proposals with defined success metrics.
- Tie future funding to demonstrated outcomes through a structured reinvestment loop.

1.4 Establish a cross-functional AI steering committee

Insight: AI must be treated as a strategic and transformative imperative. Yet in many enterprises, siloed decision-making results in fragmented execution. Business units compete for resources, IT becomes a bottleneck, and AI initiatives are often assessed prematurely through a narrow ROI lens.

Recommendation

Create a central AI governance forum to align on enterprise-wide priorities, assess initiatives through strategic and long-term value creation criteria, and accelerate safe experimentation. The goal is to enable rapid learning while building foundational AI readiness and mitigating competitive risk.

Actions

- Launch an AI steering committee with representation from the CEO, CIO, CDO, CRO, and key business heads.
- Empower the committee to approve initiatives based on strategic fit, not just near-term ROI.
- Prioritize low-cost pilot projects that demonstrate proof of value and promote organizational learning.
- Use the forum to allocate data access, manage vendor relationships, and track AI maturity, regulatory readiness, and investor expectations.

1.5 Federate innovation, centralize guardrails

Insight: Employees are already adopting AI tools without formal oversight. Blocking access creates shadow usage, while ungoverned adoption introduces security and compliance risks.

Recommendation
Enable innovation at the edge while enforcing control through centralized policy, monitoring, and tooling.

Actions
- Adopt a federated AI operating model with business-led experimentation and central oversight.
- Publish pre-approved toolkits including models, prompt templates, and usage boundaries.
- Monitor usage with integrated telemetry, audit logs, and opt-in feedback loops.

1.6 Align AI initiatives to investor and market expectations

Insight: Boards and investors increasingly demand transparency on AI strategy, value creation, and risk management. Most organizations lack the structure to report this effectively.

Recommendation
Treat AI as a core component of corporate strategy and embed it within enterprise reporting frameworks.

Actions
- Define and publish an “AI Capital Plan” linked to strategic outcomes, cost savings, and new value creation.
- Establish quarterly reporting to the board covering pipeline, budget utilization, and risk metrics.
- Integrate responsible AI and performance tracking into ESG and enterprise risk disclosures.

2. Critical challenges

While strategic imperatives guide how to embed AI, progress often stalls due to organizational, cultural, and structural barriers. This section outlines the most critical challenges, and the actions leaders must take to overcome them.

2.1 Unclear AI starting points stall enterprise momentum

Mitigation: Define structured entry points for AI across business domains, supported by a central framework and guided experimentation.

Action: Run facilitated use-case discovery workshops aligned to business goals, supported by reusable pilot frameworks.

2.2 Shadow AI exposes the enterprise to unmanaged risk

Mitigation: Offer secure, sanctioned AI tools to eliminate the need for unapproved alternatives.

Action: Deploy enterprise-approved AI platforms with built-in monitoring and enforceable usage policies.

2.3 Fragmented budget ownership delays execution

Mitigation: Centralize early AI investment decisions under a unified governance model.

Action: Establish a cross-functional AI innovation fund with shared accountability between business, IT, and finance.

2.4 Inconsistent success metrics limit cross-team learning

Mitigation: Adopt standardized KPIs to evaluate and replicate what works across pilots.

Action: Create a central reporting template that tracks pilot impact using common business metrics.

2.5 Cultural resistance slows adoption and trust in AI

Mitigation: Position AI as a tool for empowerment, not displacement, and build confidence through targeted education.

Action: Launch opt-in, function specific AI training paired with internal success stories to shift mindsets.

2.6 Compliance ambiguity limits AI scalability

Mitigation: Define applicable regulatory requirements early and embed responsible AI practices to reduce implementation friction.

Action: Establish a joint evaluation framework with legal, risk, and IT teams. Prioritize compliant use cases to enable safe and auditable experimentation.

3. Implementation Plan: Scalable AI architecture

To transition from isolated pilots to enterprise-scale AI, organizations need a modular, secure, and observable architecture. The Hong Kong roundtable shared practical design principles and implementation steps to help build resilient, governed, and scalable AI platforms.

3.1 Design principles

Modularity: Enterprise AI must be built for agility and reuse. A modular, API-first architecture enables rapid experimentation without disrupting core systems.

- Containerized design allows easy replacement of models and scaling across teams.
- Vendor neutrality preserves flexibility and cost control by supporting multi-cloud and on-prem environments.
- Geo-political adaptability ensures AI resilience to regional infrastructure limits, data sovereignty laws, and regulatory shifts.

Security: AI systems must be secure by default, especially when dealing with regulated or sensitive data.

- Privacy by design with risk-tiering ensures that workloads are protected according to their sensitivity.
- Secure deployment options such as private cloud and on-prem infrastructure are essential for industries like finance and healthcare.

Observability: Visibility is critical for building trust in AI systems. Enterprises must know what models are doing, where data flows, and how outputs are generated.

- Observability provides real-time insights into system behaviour for governance and audit.
- Explainability ensures decisions made by AI are understandable, especially in critical functions like underwriting or hiring.

Interoperability: AI must integrate into the tools and workflows employees already use. Adoption depends on relevance at the point of work.

- Seamless integration with CRM, ERP, and productivity platforms ensures AI supports real operations.
- Flexible architecture allows AI to work across varied systems without dependency on specific vendors.

3.2 Core architecture components

Enterprise AI requires an architecture that is modular, governed, and adaptable across use cases. The following layers form the technical foundation for scalable deployment.

Data layer
Ingest and process data in real time while maintaining structured warehousing for lineage, auditability, and governance. This layer ensures data is both actionable and compliant.

Model layer
Orchestrate a mix of in-house models, open-source LLMs, and third-party APIs. Selection should be based on use case sensitivity, performance needs, and cost considerations.

Governance layer
Centralize control through a unified model registry, role-based access management, logging, and policy enforcement. This ensures oversight is embedded across the development lifecycle.

Interface layer
Deploy task-specific AI copilots directly into business systems. Embedding intelligence at the point of work drives adoption and delivers measurable impact.

3.3 Implementation steps

- Launch **low-risk internal pilots** to build familiarity and uncover operational friction.
- Identify **high-value use cases** by assessing data readiness and business impact.
- **Upgrade infrastructure** to support modular, hybrid AI deployment across environments.
- **Embed governance** early through risk-tiering, access control, and usage monitoring.
- **Integrate observability and explainability** into all AI workflows for transparency.
- **Scale through federated teams** using shared infrastructure and approved templates.

4. Success metrics

AI progress should be measured using pragmatic metrics that reflect business value, organizational readiness, and responsible growth.

- **Cost-to-value efficiency:** Demonstrated savings or productivity gains that justify investment beyond pilot stages.
- **Organizational adoption:** Uptake of AI tools across business units, workflows, and roles.
- **Governance maturity:** Availability of audit trails, explainability mechanisms, and adherence to policies.

- **Scaling readiness:** Funding of follow-on initiatives and strong user feedback that fuels expansion.
- **Use case value creation:** Business outcomes and measurable impact derived from implemented use cases.

Action: Define and monitor metrics tied to business impact, adoption, and compliance to ensure responsible growth and sustained executive sponsorship.

5. Next steps

With foundational pilots in place, the focus must now shift to scaling AI responsibly and embedding it into the fabric of the organization. Leaders must act decisively to expand usage, foster trust, and build lasting capability across teams.

- **Institutionalize AI in business functions** by identifying repeatable, high-impact workflows in HR, finance, and operations that benefit from automation or augmentation.
- **Shape a narrative of augmentation** by actively positioning AI as a tool that enhances employee performance rather than threatens it.
- **Operationalize peer-driven momentum** by converting early pilot wins into case studies and embedding them into onboarding, training, and townhalls.
- **Build capability at the edge** through immersive, role-specific learning journeys that combine real datasets, tool exposure, and certification-linked outcomes.
- **Evolve the AI operating model** from a centralized structure to a federated one as governance frameworks mature, enabling greater autonomy and execution.

The true measure of AI maturity lies not in experimentation but in execution at scale. As reinforced by the Hong Kong roundtable, the path forward is not just about deploying smarter tools. It is about building smarter enterprises. This requires bold leadership, unified governance, and a sharp focus on real business outcomes. Organizations that move with clarity, embed trust, and empower their people will not just adopt AI. They will lead with it.

Attendees

Name	Company	Designation
Alex Wu	CLP	Head of Digital Engineering
Billy Leung	Bank of China(Hong Kong) Limited	Head of Enterprise Architect (Business Solution)
Brian Lee	HKEX	Managing Director, Group Head of Technology Risk
Cara Yick	Fubon Bank	Chief Information Officer
Damian Lum	MTR Corporation Limited	Head of Data & Analytics, Enterprise Integration
Darron Sun	Hong Kong Housing Authority	Head of Information Technology
Hackker Wong	TransUnion	Head of Application Delivery
Horace Ma	HKIRC	Head of Public Mission - Cybersecurity
Kylie Fung	Vocational Training Council	Director, Information Technology Office
Marcos Chow	HKT	Group Chief Information Officer
Patrick Chiu	Gammon Construction Limited	Technical Manager
Roger Wong	The Education University of Hong Kong	Chief Information Officer
Samuel Hui	HKBN Enterprise Solutions	Co-Owner and Chief Operating Officer
Wilson Tang	HKBN Group/HKCNSA	CISO