KEYNOTE TOPIC SHEET

Signature keynotes

Ten signature sessions on enterprise AI, agentic AI, AI security, data strategy, board governance and transformation economics.

01

Operationalizing Agentic AI

How enterprises move from generative AI experimentation to governed autonomous workflows.

Most organizations have proven that generative AI can produce output. Far fewer have proven it can be trusted to act. This session sets out what changes when AI stops answering questions and starts completing work: the data and identity prerequisites, the emerging interoperability protocols, the human oversight model, the security perimeter around tools and actions, and the operating model redesign that has to happen alongside the technology.

What the audience leaves with

- A staged path from assisted work to governed autonomy, with the control gates at each stage
- The data, context and identity foundations agents require before they can be trusted
- An oversight model that satisfies risk and audit without stalling deployment
- Where agentic workflows create value first, and where they reliably do not

Best audiences: Boards, CEOs, CIOs, CDOs, CAIOs, CISOs

02

The AI Operating Model

A practical blueprint for organizing strategy, governance, platforms, talent and funding around enterprise AI.

AI strategy documents are plentiful. Operating models that actually deliver are rare. Drawing on enterprise-wide deployments inside regulated institutions, this session lays out how to structure ownership, funding, platform decisions, product accountability, risk partnership and adoption mechanics so that AI becomes a durable capability rather than a portfolio of pilots.

What the audience leaves with

- A decision rights model covering strategy, platform, risk and product ownership
- Funding approaches that survive the move from pilot budget to run-rate investment
- Adoption mechanics that move usage from launch spike to sustained behavior
- The small number of centralized functions that genuinely need to be centralized

Best audiences: C-suite, Transformation leaders, Operating committees

03

From Data Strategy to AI Advantage

Why AI outcomes are constrained by data quality, semantics, governance and access, and how to build an AI-ready foundation.

AI performance is a downstream symptom of data condition. This session works through the specific data characteristics that determine whether an AI program scales: semantic consistency, lineage, quality at the point of capture, governed access, privacy controls and the contextual metadata that models depend on. It draws on the consolidation of more than forty data warehouses and a twenty-petabyte estate into a single modern platform.

What the audience leaves with

- The data prerequisites that actually gate AI outcomes, separated from the ones that do not
- How to sequence platform consolidation without stopping the business
- A governance approach that scales across thousands of data elements
- How to make the economic case for foundational data work

Best audiences: CDOs, CIOs, CAIOs, Data and analytics leaders

04

AI Security Is the New Enterprise Control Plane

How to secure enterprise AI, agents, data, identities, models and autonomous actions while preserving speed.

AI introduces a risk surface that legacy controls were never designed to manage: prompts, tools, model behavior, non-human identities, agent-to-agent traffic and autonomous action. This session frames AI security as a continuous control plane rather than a periodic review, and sets out what boards, CISOs and CIOs should require before autonomous systems touch sensitive data.

What the audience leaves with

- A map of the AI risk surface, including the categories legacy controls miss
- How to govern non-human identity and agent permissions at scale
- Controls that operate continuously rather than at review checkpoints
- How to preserve innovation velocity while tightening the control environment

Best audiences: CISOs, CIOs, CAIOs, Boards, Risk leaders

05

The Board Guide to AI Value and Risk

What directors should be asking management about AI investment, productivity, model risk and accountability.

Boards are being asked to approve significant AI investment while the evidence base is still forming. Built from monthly board reporting experience inside a large regulated institution, this session gives directors a concrete question set covering capital allocation, value measurement, model risk, cyber and data exposure, regulatory posture, and who is accountable when an autonomous system gets it wrong.

What the audience leaves with

- A director question set for AI investment and value reporting
- The metrics that indicate real adoption versus reported activity
- How to test whether accountability for AI decisions is genuinely assigned
- Warning signs that an AI program is generating risk faster than value

Best audiences: Boards, Directors, Audit and risk committees

06

Reengineering the Enterprise for AI

How to redesign sales, service, operations, finance and risk around AI-driven productivity rather than layering AI onto legacy processes.

The largest source of disappointing AI returns is applying new capability to unchanged processes. This session covers function-by-function redesign: what work gets eliminated rather than accelerated, how roles change, where decision rights move, and how to manage the workforce and change dimension honestly. It draws on an enterprise-wide AI deployment that reached 88 percent weekly active usage alongside process redesign.

What the audience leaves with

- A method for identifying processes that should be removed rather than automated
- How role design and incentives have to change for adoption to hold
- Sequencing that delivers early proof without destabilizing operations
- How to measure productivity credibly enough to withstand finance scrutiny

Best audiences: CEOs, COOs, Functional executives

07

AI in Financial Services: From Experimentation to Production

Lessons from banking and payments on scaling AI inside regulated environments.

Regulated institutions face a harder version of the AI problem: the same ambition, with model risk management, supervisory expectations, privacy obligations and audit evidence attached. This session covers what it actually takes to get AI into production in a bank, including governance design, model risk partnership, client experience, fraud and financial crime, capital considerations and the remediation discipline that keeps programs credible with examiners.

What the audience leaves with

- A production path for AI that risk, audit and supervisors can accept
- How to design evidence and documentation into delivery from the start
- Where AI has produced measurable outcomes in banking and payments
- How to remediate at scale without halting modernization

Best audiences: Financial services executives, Risk and compliance leaders, Boards

08

The Economics of Transformation

How to connect AI and digital investment to revenue, cost, capital efficiency, risk reduction and shareholder value.

Transformation programs fail commercially long before they fail technically. This session sets out how to build an investment case that survives contact with the CFO: value drivers that can be measured, the difference between cost avoidance and cost reduction, capital efficiency as an underused lever, and how to report progress in terms a board and an investor recognize.

What the audience leaves with

- A value framework spanning revenue, cost, capital and risk
- How to make capital efficiency a first-class transformation outcome
- Reporting that holds up under finance and investor scrutiny
- How to retire investment that is not producing returns

Best audiences: CEOs, CFOs, Boards, Strategy leaders

09

Building the AI-Native Company

The organizational, cultural, platform and talent changes required to operate as an AI-native enterprise.

There is a meaningful difference between an enterprise that uses AI and an enterprise that is organized around it. This session examines that gap across structure, culture, platform strategy, talent model, governance and capital allocation, and what incumbents can realistically adopt from companies that were built this way from the start.

What the audience leaves with

- The structural differences between AI-using and AI-native organizations
- Talent and capability decisions that compound over time
- Platform choices that either enable or constrain the next five years
- What incumbents can adopt, and what they cannot

Best audiences: Enterprise leadership teams, Innovation forums

10

What Comes After GenAI

A forward view on agents, enterprise memory, model routing, autonomous workflows and the next operating model shift.

A perspective session for audiences planning beyond the current cycle. It covers where agentic systems, multimodal interfaces, enterprise memory, model routing and autonomous workflows are heading, which shifts are likely to be durable, which are likely to be absorbed into platforms, and what leaders should be positioning for now.

What the audience leaves with

- A view of the next platform shift and its operating model consequences
- Which capabilities to build versus buy versus wait out
- The security and governance implications arriving with autonomy
- Signals worth tracking over the next several quarters

Best audiences: Innovation conferences, Technology summits, Board retreats

To discuss a session, email thomasmazzaferro@evolveconsultinggroup.co or use thomasmazzaferro.com/book