FIRST SAN FRANCISCO PARTNERS

What AI Governance Leaders Really Need to Know

Cutting Through the Noise

AIGOV LEADING AI Governance

DATAVERSITY®

Welcome to Leading AI Governance – Session 1

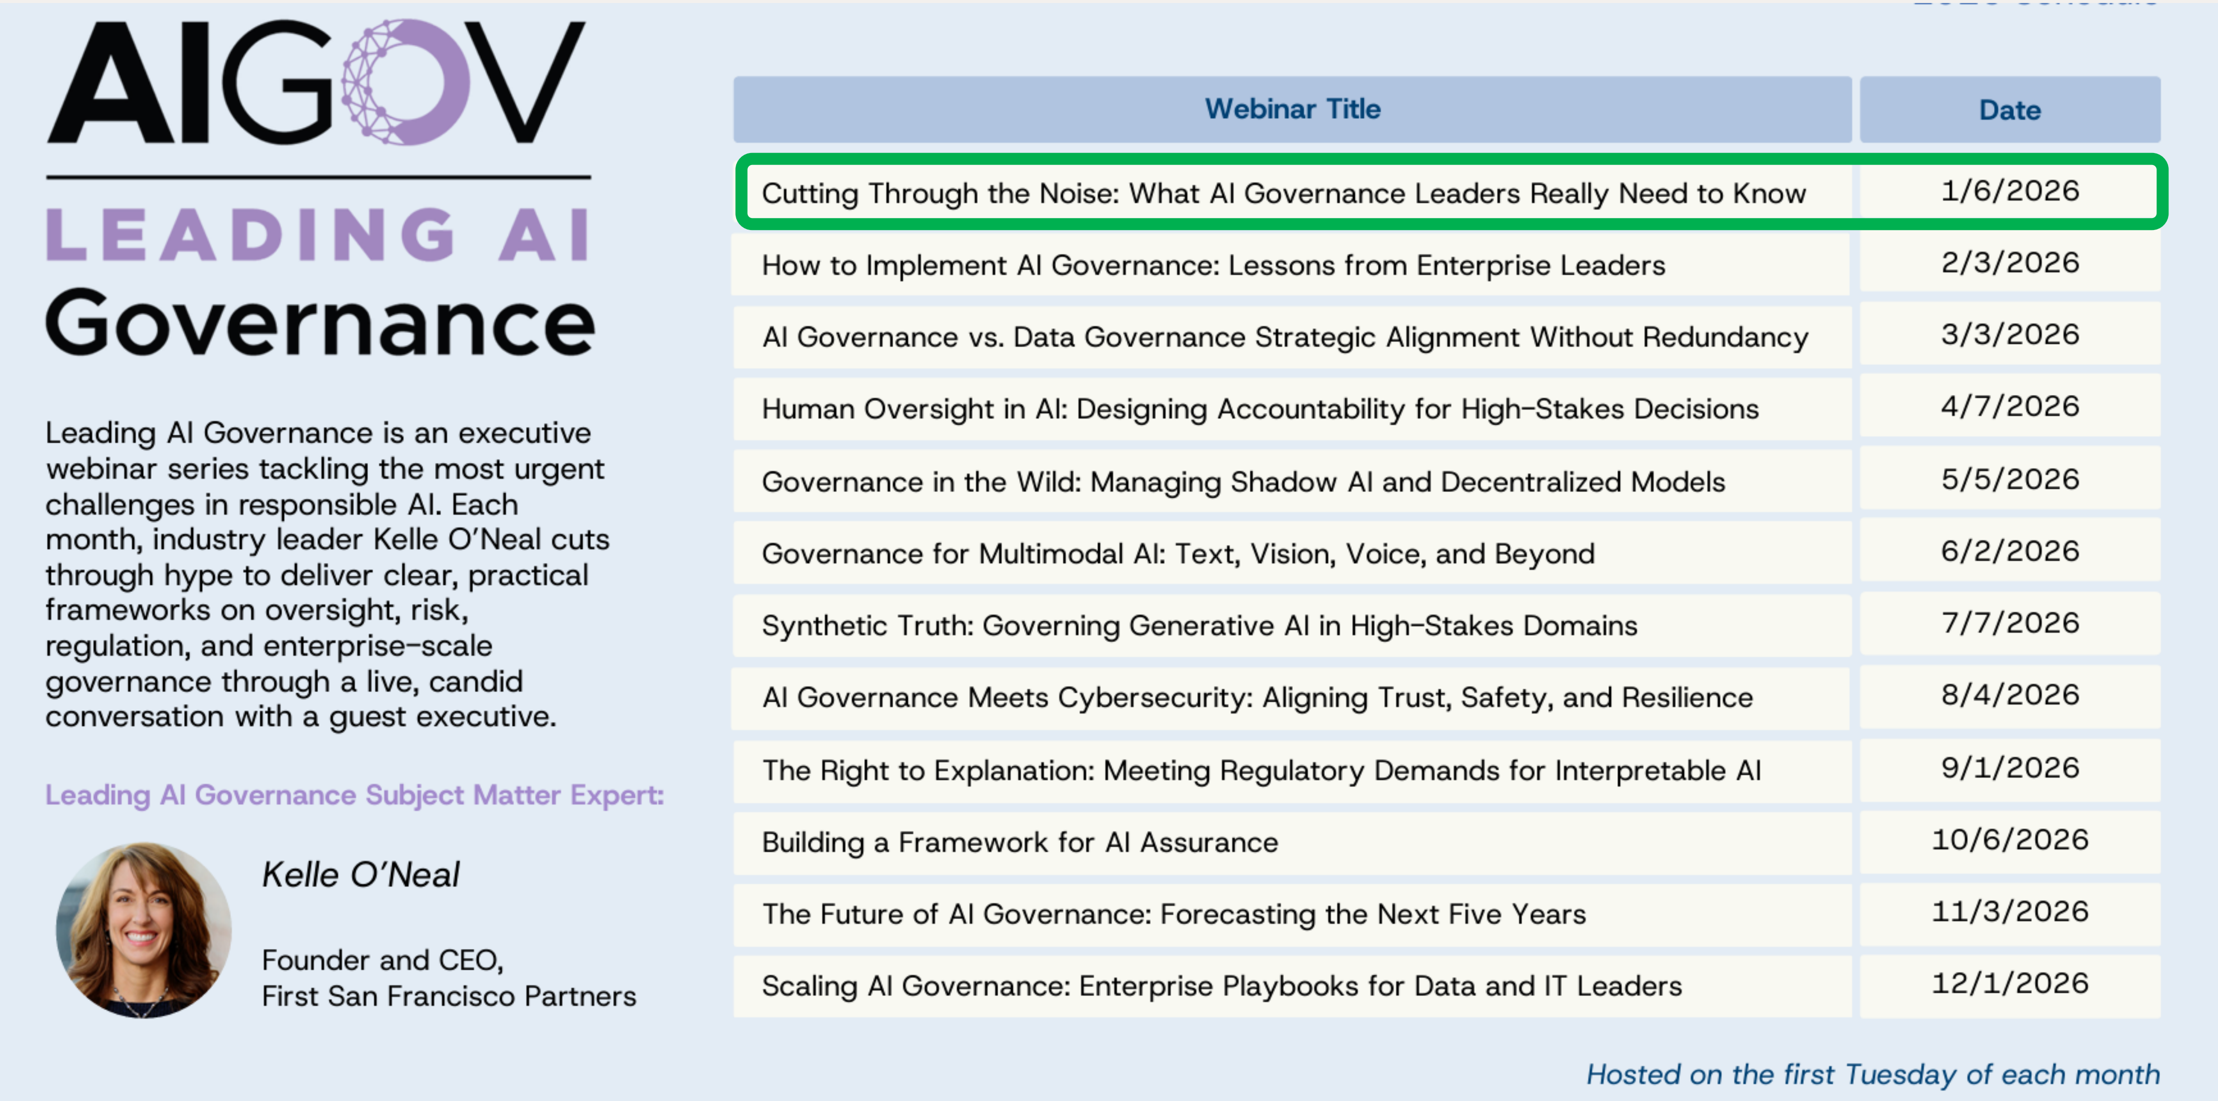

AIGOV

LEADING AI
Governance

Leading AI Governance is an executive webinar series tackling the most urgent challenges in responsible AI. Each month, industry leader Kelle O'Neal cuts through hype to deliver clear, practical frameworks on oversight, risk, regulation, and enterprise-scale governance through a live, candid conversation with a guest executive.

Leading AI Governance Subject Matter Expert:

Kelle O'Neal

Founder and CEO,
First San Francisco Partners

Webinar Title	Date
Cutting Through the Noise: What AI Governance Leaders Really Need to Know	1/6/2026
How to Implement AI Governance: Lessons from Enterprise Leaders	2/3/2026
AI Governance vs. Data Governance Strategic Alignment Without Redundancy	3/3/2026
Human Oversight in AI: Designing Accountability for High-Stakes Decisions	4/7/2026
Governance in the Wild: Managing Shadow AI and Decentralized Models	5/5/2026
Governance for Multimodal AI: Text, Vision, Voice, and Beyond	6/2/2026
Synthetic Truth: Governing Generative AI in High-Stakes Domains	7/7/2026
AI Governance Meets Cybersecurity: Aligning Trust, Safety, and Resilience	8/4/2026
The Right to Explanation: Meeting Regulatory Demands for Interpretable AI	9/1/2026
Building a Framework for AI Assurance	10/6/2026
The Future of AI Governance: Forecasting the Next Five Years	11/3/2026
Scaling AI Governance: Enterprise Playbooks for Data and IT Leaders	12/1/2026

Hosted on the first Tuesday of each month

Today’s Topic

Cutting through the Noise

At the end of this session, participants will:

- ✓ Have a working definition of AI Governance
- ✓ Understand why the AI Governance space is so heavy on buzzwords and light on action
- ✓ Test your organizational readiness with 20 key questions

AIGOV | LEADING AI Governance

Why Is There So Much Clutter?

Why AI Governance feels harder than it should

AI Governance feels cluttered because:

- Multiple functions need to contribute, and each bring their own lens
- Vendors, regulators, and internal teams use the same words differently
- AI evolved faster than enterprise operating models

As a result of the confusion:

- Leaders hear lots of words, but gain less clarity
- AI Governance is ignored because it’s too difficult, too slow and too rigid – we default to tools as solutions
- Unintentional mis-use leaves the organization exposed

Defining AI Governance in 2026

What it is, what it covers, and who's accountable

AI Governance is:

- AI Governance is the organizing framework for establishing the strategy, people, and processes for responsible creation and management of AI solutions in support of organizational goals.
- A combination of visibility, decision rights, and controls that ensures AI aligns with business, risk, and regulatory expectations.

AI Governance covers:

- Decision making rights
- Use-case visibility and ownership
- Risk tiering and approval paths
- Ongoing monitoring and escalation
- Evidence and accountability (not just intent)

AI governance turns AI ambition into accountable, repeatable, actionable decision-making.

Why AI Governance is Non-Negotiable

Research from ServiceNow and Oxford Economics' AI Maturity Index reveals that pacesetter organizations that are achieving measurable AI benefits have established cross-functional governance councils with genuine executive authority, not technical committees relegated to advisory roles.

Organizations with mature responsible AI frameworks achieve 42% efficiency gains, according to McKinsey, demonstrating that governance enables innovation rather than constraining it — provided the governance operates as an architectural principle rather than a compliance afterthought.

Why trust is the new currency in the agentic era — and what it's worth | CIO

AIGOV | LEADING AI Governance

Reactive AI Governance is Costly

Post-Incident Framework Adoption

Organizations tend to create AI governance frameworks as a reaction to major incidents or regulatory changes rather than as proactive measures.

Imbalance in Investment

Preventative action demands less funding compared to corrective measures, revealing an imbalance in risk management priorities.

Crisis-Driven Engagement

Without responsibilities and accountabilities clearly defined, organizations scramble in times of crisis, instead of taking quick, measured mitigation steps.

Delicate Balance

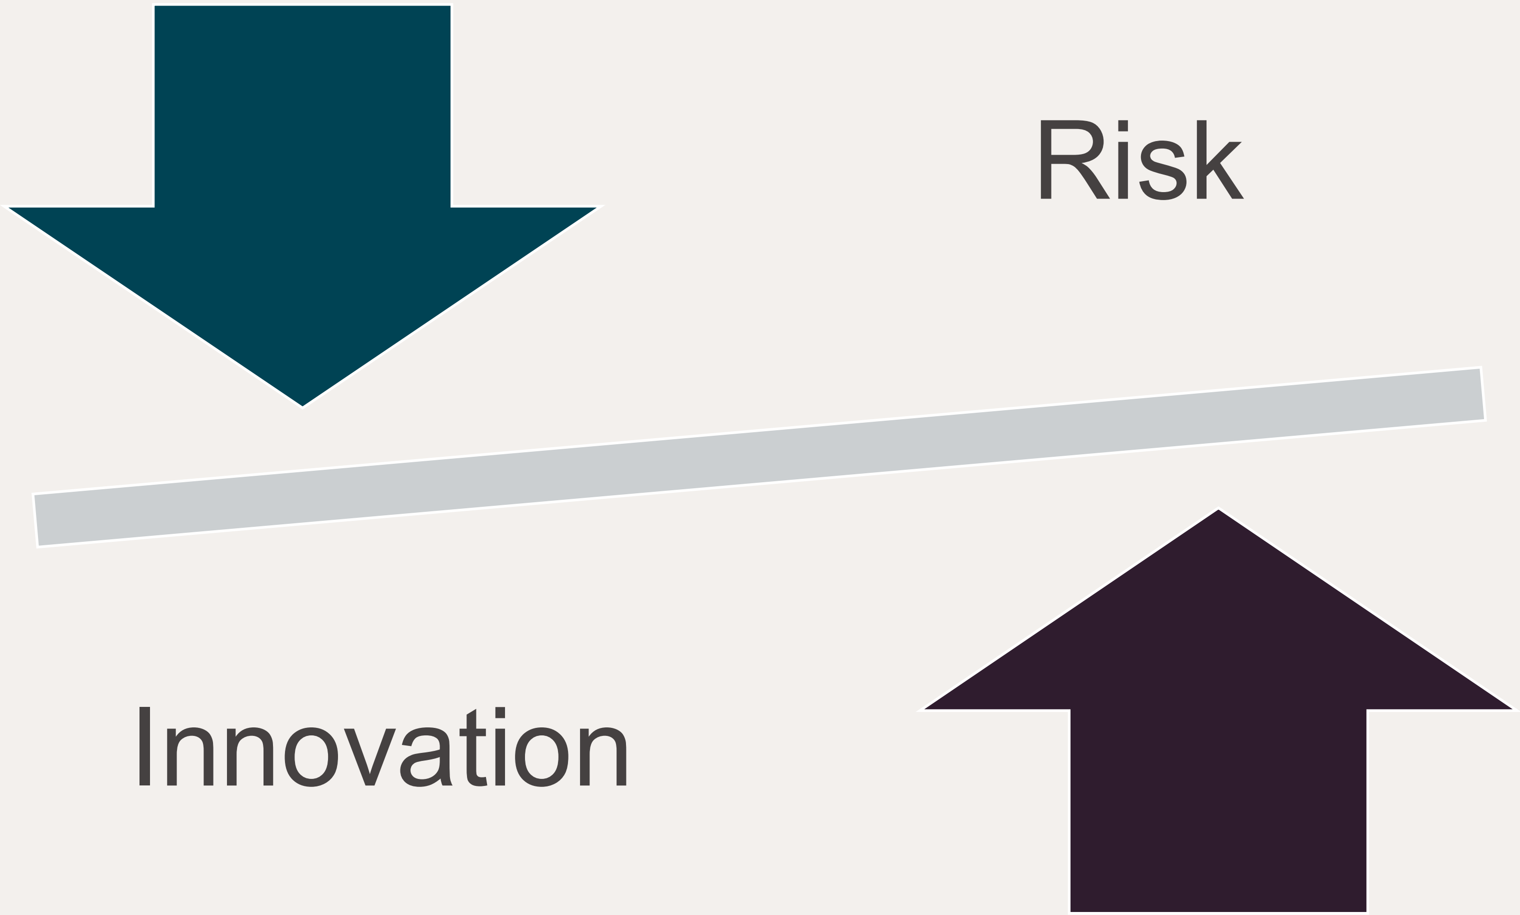

AI Governance Considerations

A rigid framework leaves no room for what matters most to you; leveraging what already works in your organization is a great place to begin.

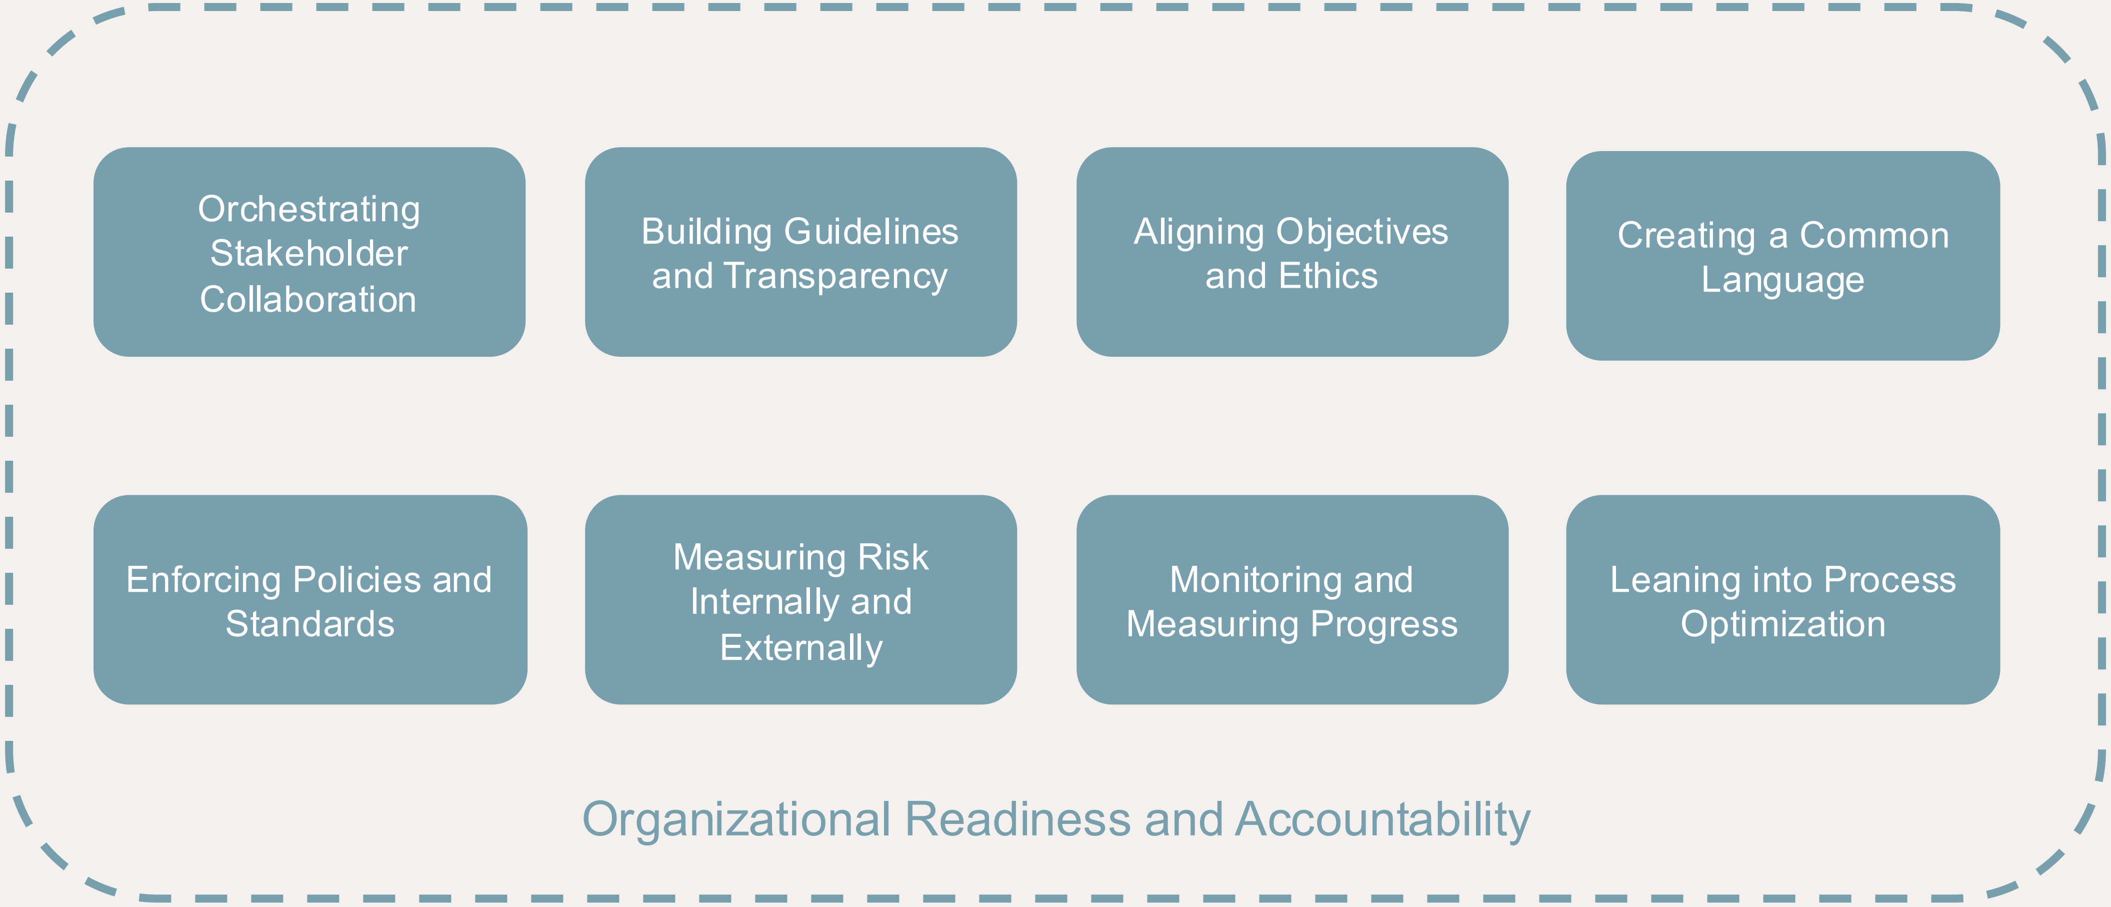

Orchestrating Stakeholder Collaboration

Building Guidelines and Transparency

Aligning Objectives and Ethics

Creating a Common Language

Enforcing Policies and Standards

Measuring Risk Internally and Externally

Monitoring and Measuring Progress

Leaning into Process Optimization

Organizational Readiness and Accountability

Progress without Clarity

AI as a rapidly evolving capability means that frameworks are adapting and evolving equally as quickly.

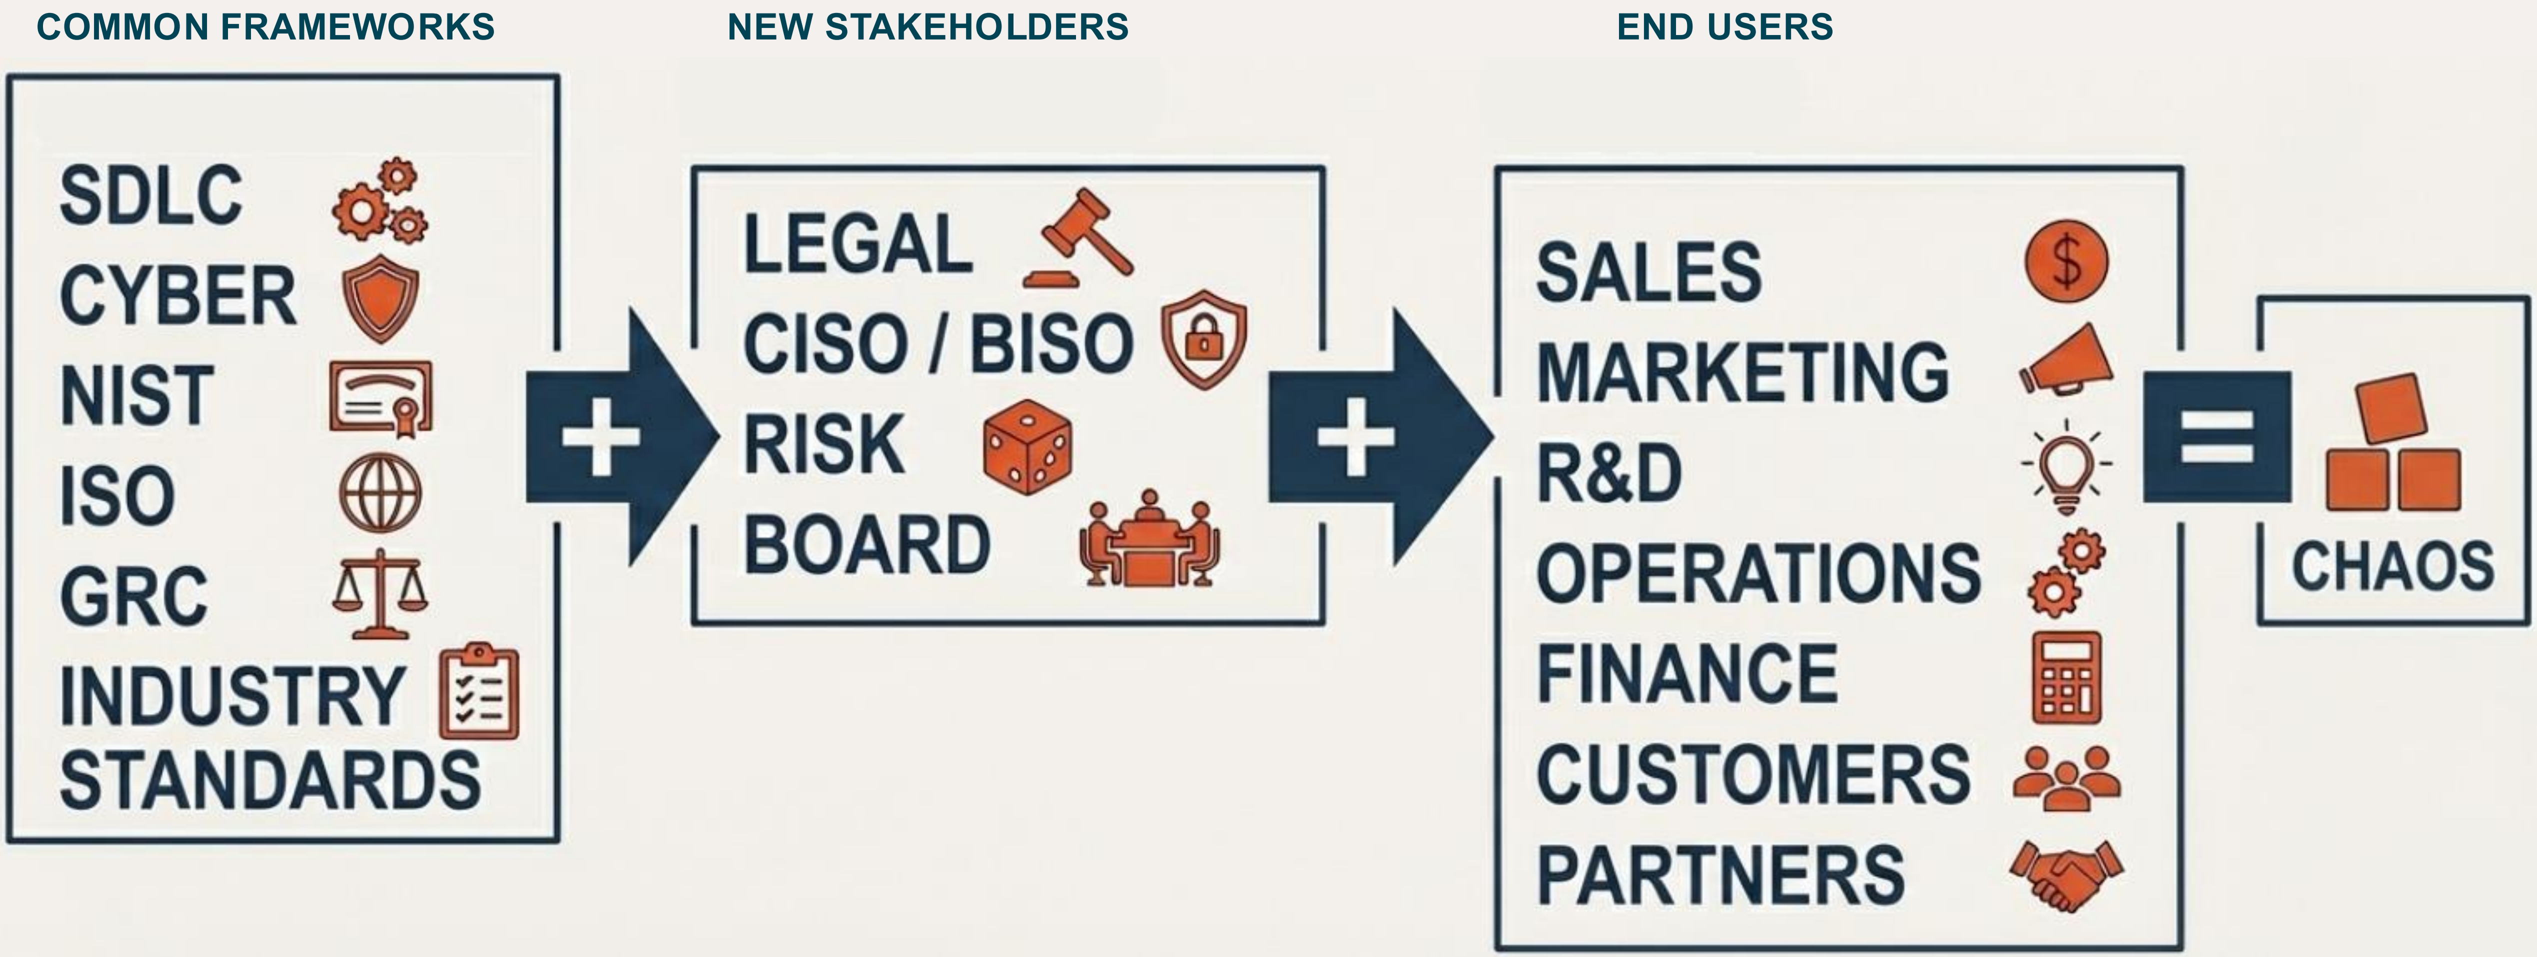

COMMON FRAMEWORKS

SDLC
CYBER
NIST
ISO
GRC
INDUSTRY STANDARDS

+

NEW STAKEHOLDERS

LEGAL
CISO / BISO
RISK
BOARD

+

END USERS

SALES
MARKETING
R&D
OPERATIONS
FINANCE
CUSTOMERS
PARTNERS

=

CHAOS

Success by Design

Creating an Orchestration Layer that cuts through the clutter, gives everyone a voice and aligns effort to value is the single most effective way to manage chaos.

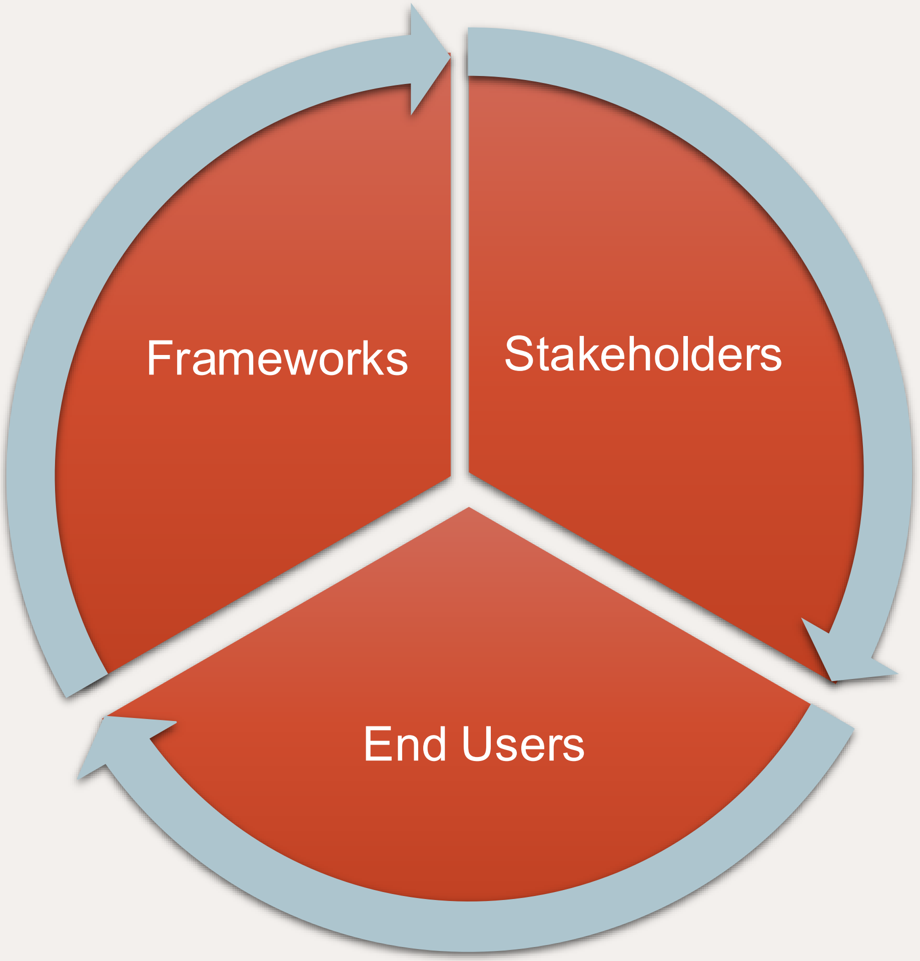

The Orchestration Layer = Decision-Making and Accountability

Accountability drives action and momentum builds rapidly when there is a clear operating model

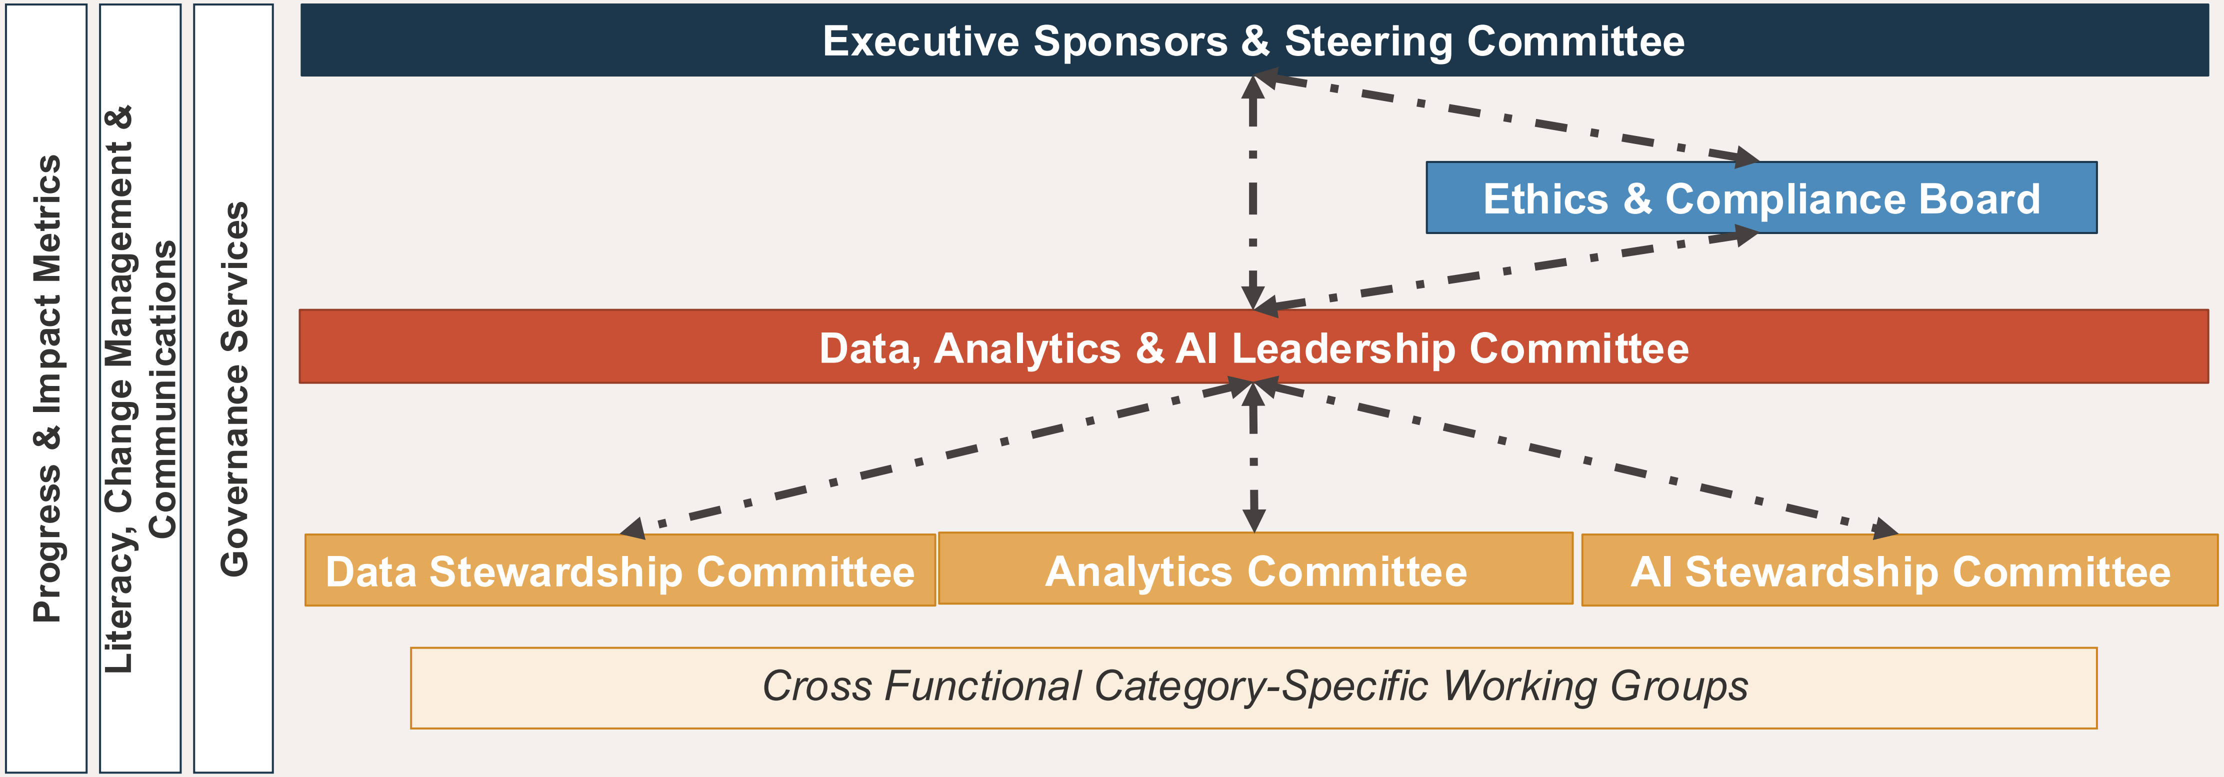
Progress & Impact Metrics
Literacy, Change Management & Communications
Governance Services
Executive Sponsors & Steering Committee
Ethics & Compliance Board
Data, Analytics & AI Leadership Committee
Data Stewardship Committee
Analytics Committee
AI Stewardship Committee
Cross Functional Category-Specific Working Groups

Governance Orchestrates AI Implementation Success

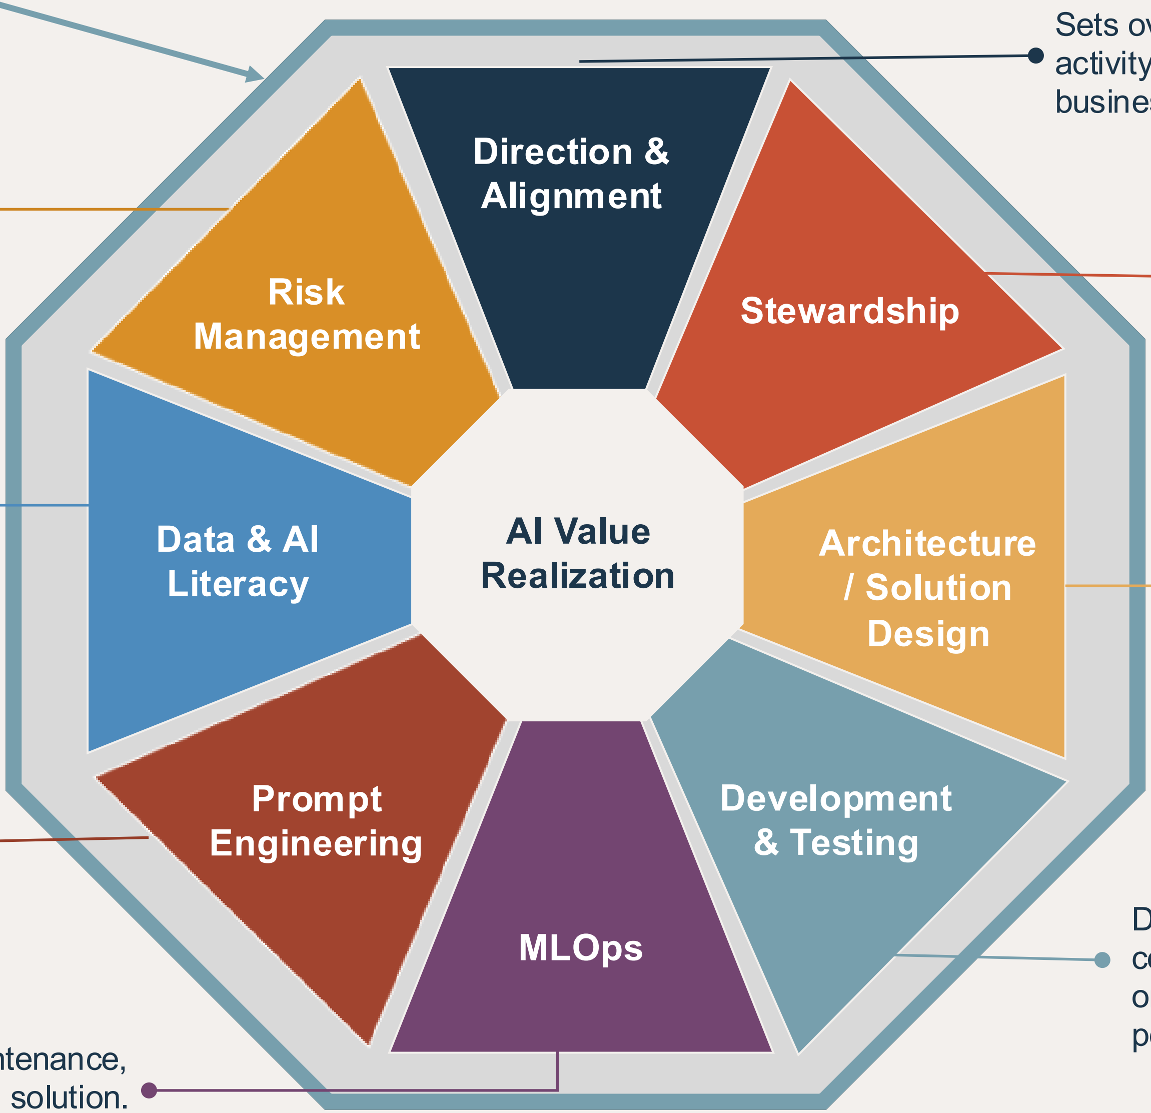

Sets overall direction, strategy, approaches, and activity prioritization ensuring alignment with the business value case.

Legal, Privacy, Compliance, Security, and Ethics risk management and impact assessments.

Collects requirements and ensures development of AI solution occurs in accordance with requirements and governance standards.

Assists with design, development, and delivery of materials that support, train, and educate the business for optimization of solution outcomes.

Ensures solution approaches and designs are created in accordance with future state design and are aligned with enterprise architecture.

Provides guidance by use case to improve human-AI collaboration outcomes by facilitating best practices for AI prompting.

Develops solution(s) and technical controls for remediation and perform ongoing testing to ensure quality, performance, and usability.

Ensures ongoing maintenance, monitoring, and continuity of AI solution.

Investment in Data slowing, as investment in AI is growing

Digital Trust and Cybersecurity as defined by McKinsey:

> "Digital trust and cybersecurity covers technologies and practices designed to ensure secure, transparent, and trustworthy digital interactions. This includes identity verification, data protection, encryption, threat detection, and blockchain-based trust systems."

https://www.mckinsey.com/capabilities/tech-and-ai/our-insights/the-top-trends-in-tech

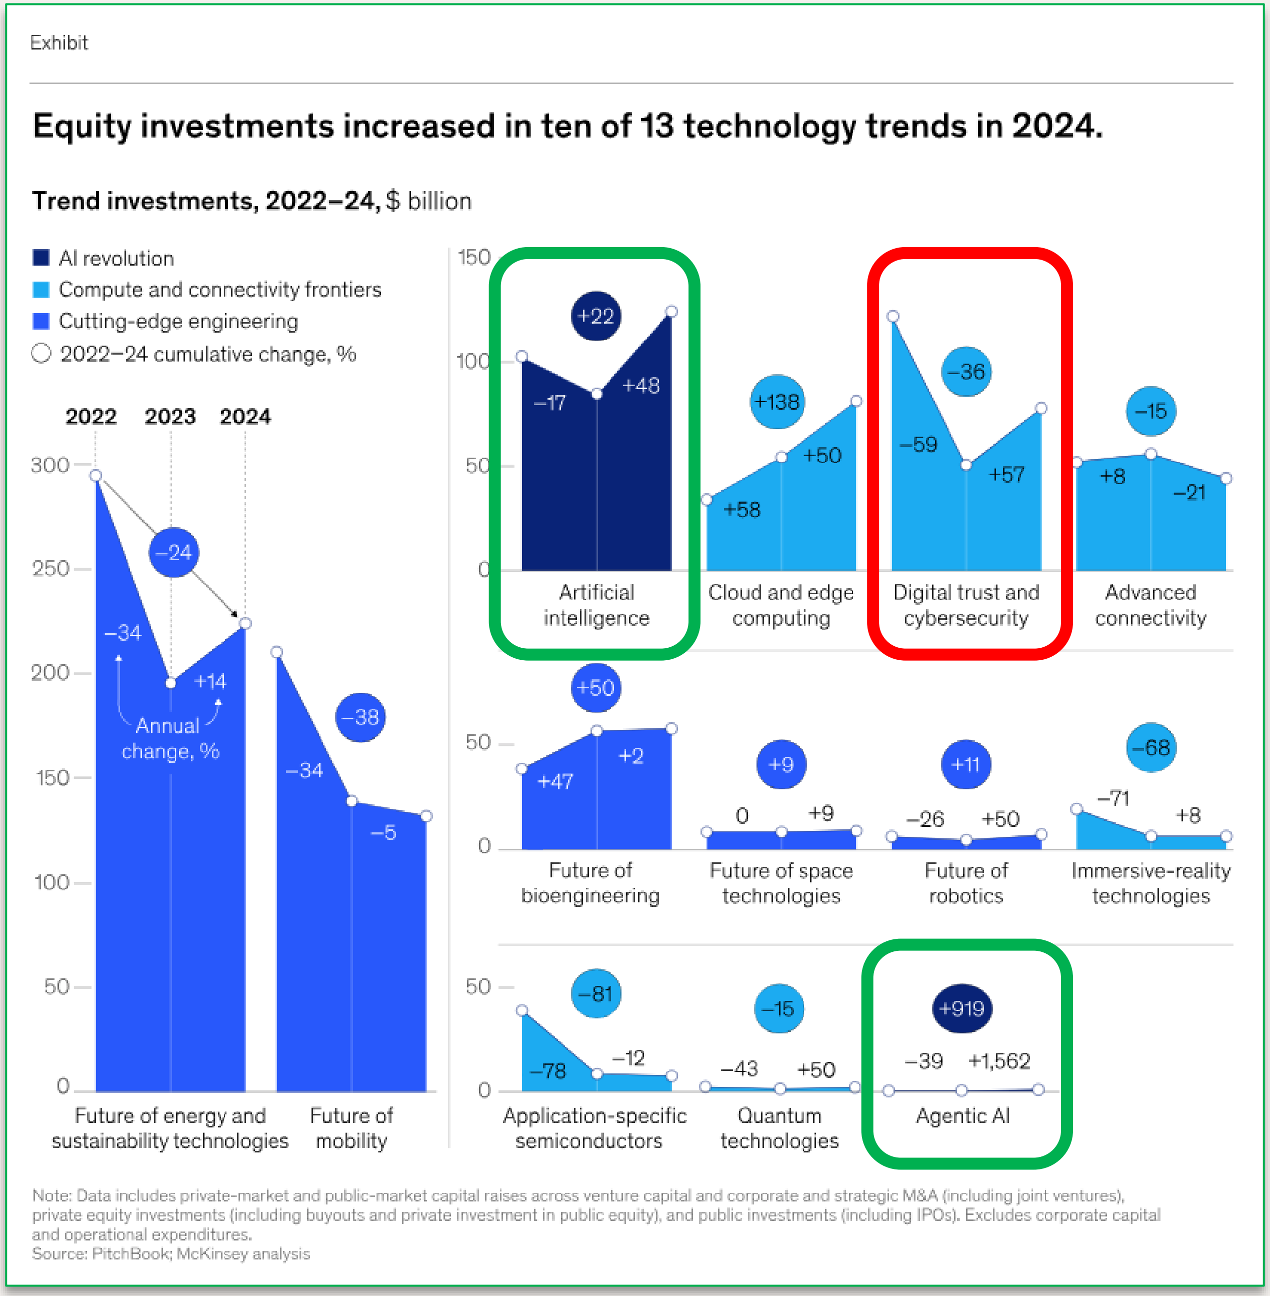

Exhibit

Equity investments increased in ten of 13 technology trends in 2024.

Trend investments, 2022–24, $ billion

Note: Data includes private-market and public-market capital raises across venture capital and corporate and strategic M&A (including joint ventures), private equity investments (including buyouts and private investment in public equity), and public investments (including IPOs). Excludes corporate capital and operational expenditures.
Source: PitchBook; McKinsey analysis

Balancing 2026 AI Spend

Less experimentation. More foundation. Real returns.

Tool Sprawl & One-Off AI Experiments

Much of last year's AI spend went to disconnected tools and pilots that didn't scale. Continuing to fund more point solutions this year only reinforces sunk costs, increases complexity, and adds data and technical debt without improving outcomes.

A Strong Data & AI Foundation

Redirecting spend to data and AI governance creates the foundation for trust and scale. Clear ownership, consistent definitions, and governed data flows reduc rework, protect prior investments, and allow AI initiatives to deliver lasting value.

Governance brings consistency and responsibility

Starting with people and process ensures that AI works for the business, and not the other way around

Be Intentional
Communicate Purpose, Assign Accountability

Avoid Silos
Share Capabilities, Differentiate Value

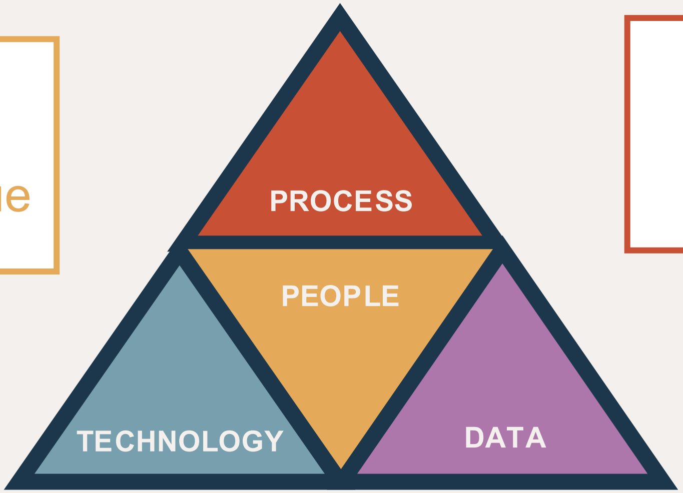

Embrace Incrementality
Don't Overengineer – Start with Essentials

Change is Constant…
and Tech is Constantly Changing

Extend Focus
Consider Users, Customers, & Stakeholders / Shareholders

AIGOV | LEADING AI Governance

What are your “20 Questions”?

The journey forward starts with a few thought-provoking questions that require purposeful answers and spark meaningful action.

How will this model behave as it scales?

How are we reskilling employees for an AI-augmented future?

How would we explain this decision to …?

What decisions are we delegating to AI?

How do we know those decisions align to our values?

What does AI failure look like?

How are we communicating our intentions clearly and honestly?

Can management assure us this system is safe?

How do you know?

Are you prepared?

What is consent?

Inspired by: AI and the Quest for Governance – Why Waiting Is No Longer an Option | LinkedIn

FIRST SAN FRANCISCO PARTNERS

Thank you for your time!

Kelle O’Neal

kelle@firstsanfranciscopartners.com

Lisa Wintrick

lisa@firstsanfranciscopartners.com

AIGOV LEADING AI Governance

Key Terms Today (2026)

What these terms mean in practice today

Responsible AI

Responsible AI is the operating model that makes AI use defensible over time.

What this means:

- Responsible AI shows up as repeatable practices: knowing what AI is in use, who owns it, how it's approved, and how it's monitored once live.
- As AI systems become more autonomous and adaptive, responsibility shifts from one-time reviews to continuous oversight and escalation paths.

What it's not:

A statement of intent or values without operational follow-through

Explainable AI (XAI)

Explainable AI is the ability to reconstruct and justify how an outcome was produced.

What this means:

- Explainability increasingly means auditability: tracing inputs, processing, and outputs in a way that supports review and accountability.
- Explanations must be appropriate to the audience, especially for regulated or high-impact decisions.

What it's not:

Surface-level transparency that doesn't support scrutiny or accountability.

Key Terms Today (2026)

What these terms mean in practice today

Synthetic Content Labeling

Synthetic content labeling is the disclosure of AI-generated or manipulated content to signal its origin.

What this means:

- Organizations are expected to disclose or mark AI-generated or manipulated content, particularly in public or high-impact contexts.
- Technical approaches focus on machine-readable markers and provenance, with known limitations.

What it's not:

A guarantee that synthetic content will be prevented or universally recognized.

Modality Risk

Modality risk reflects how AI risk changes across text, image, audio, and video.

What this means:

- Different modalities introduce distinct security, deception, and privacy risks.
- Multimodal and agentic systems expand the ways inputs can influence behavior.

What it's not:

A single, uniform risk that can be managed the same way across all AI systems.

Key Terms Today (2026)

What these terms mean in practice today

Model Risk Management (MRM)

Model risk management is the discipline of approving, controlling, and monitoring AI models as ongoing risk-bearing systems.

What this means:

- Clear ownership, approval gates, and periodic re-validation—especially as models drift, get retrained, or are updated.
- Governance expands from "model performance" to include misuse risk, data risk, and operational risk (access, logs, incident handling).

What it's not:

A one-time model validation or a static spreadsheet inventory.

AI Risk Tiering

AI risk tiering classifies AI use cases by impact and applies governance controls proportionate to that risk.

What this means:

- Higher-impact use cases require stricter documentation, testing, human oversight, and monitoring.
- Risk tiering becomes the backbone for deciding what needs reviews, what needs audits, and what can move faster.

What it's not:

Treating every AI project the same or relying on "common sense" for controls.

Key Terms Today (2026)

What these terms mean in practice today

Human Oversight

Human oversight defines where people must review, intervene, or take accountability for AI-driven decisions.

What this means:

- Clear escalation paths and stop/rollback authority when outputs are wrong, risky, or unexpected.
- Oversight increasingly shifts from "approve before launch" to "supervise during operation," especially for agents.

What it's not:

A vague statement that "humans are responsible" with no defined checkpoints.

AI Lifecycle Governance

AI lifecycle governance manages AI systems from intake through deployment, change, and retirement.

What this means:

- AI systems are governed continuously, with defined controls for updates, retraining, prompt changes, and decommissioning—not just at launch.
- Ownership, approvals, and monitoring expectations evolve as the system changes and scales.

What it's not:

A one-time review or approval that ends once the AI system is deployed.

Key Terms Today (2026)

What these terms mean in practice today

Adversarial Testing

Adversarial testing systematically probes how AI systems can fail, be misused, or be manipulated.

What this means:

- Tests for prompt injection, data leakage, harmful content, tool misuse, and agent overreach—then tracks fixes and re-tests.
- Mature governance treats this like a continuous program, not a one-off exercise.

What it's not:

Random "try to break it" testing with no remediation tracking.

Technical Guardrails

Technical guardrails are enforceable constraints that prevent unsafe or non-compliant AI behavior in real workflows.

What this means:

- Policies become "real" when embedded into access controls, tool permissions, data redaction, content filters, and safe output constraints.
- Guardrails are expected both at build time (release gates) and runtime (live protections + alerts)

What it's not:

A policy doc, a training deck, or a "please be safe" prompt.

AIGOV | LEADING AI Governance