



AI Text, Vision, Voice & Beyond

Governance for Multimodal AI



Welcome to Leading AI Governance – Session 6

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

PRECISION ABOUT WHERE AI GOVERNANCE DOES ITS WORK

What You Will Leave With

Operational Clarity

AI governance is the discipline that operates at the moment of decision

Structure

The approach for pulling data governance, privacy, compliance, and ethics into the decision moment.

Action

Something you can take into tomorrow's leadership meeting.

Where This Session Fits

Five sessions with one continuous argument. Today's session extends into multimodal data for AI.

January: Cutting through the noise. Visibility, decision rights, controls.

February: Lessons from enterprise leaders. How you decide matters as much as what you decide.


March: AI vs. Data Governance. Decision stewardship as its own discipline.

April: Designing accountability. If it isn't codified, it isn't real.

May: Governance in the wild. Govern the layer, not the tool.

Today - Session 6: Governance for Multimodal AI

Govern what your AI concludes and what gets done about it





Multimodal AI Defined

A Working Definition and a Distinction That Matters

Three different ways “multimodal AI” can be interpreted

1 The Model

LLM models (GPT-4o, Gemini, Claude) natively process multiple modalities in a single forward pass.



2 The Application

Domain tools: voice analytics, document intelligence, meeting transcription. Process-specific modalities for specific use cases.



3 The Data


Unstructured content of every kind flowing into AI systems (e.g., text, audio, image, video, document), regardless of which model consumes it.

This talk focuses on the data

Multimodal Data is Not New

Unstructured data is the unfinished business of data governance.

- Emails and chat messages (Slack, Teams)
- Support tickets and customer correspondence
- Meeting transcripts and recordings
- Contracts, agreements, and policy documents
- Scanned PDFs in shared drives



An Everyday Example - the Meeting Transcript

The most prevalent multimodal AI deployment.

The transcript bot is in the meeting and it is:

- Transcribing every word, identifying every speaker
- Summarizing the discussion and extracting action items
- Characterizing tone and engagement
- Inferring sales signal, project status, sentiment

The four questions most programs cannot answer:

- What may be transcribed?
- What may be summarized or inferred?
- Who owns the artifact afterward?
- What is the retention policy?

These questions don't sit with one function. Data governance, AI governance, compliance, ethics, and privacy need a seat at the table

And this is the easy case; you know the AI is in the room.


What AI Is Actually Doing

The authorized use case: meeting transcripts

The organization authorizes a **TOOL** to 1. Record, 2. Summarize, 3. Retrieve, 4. Generate, 5. Draft

The **DATA** being analyzed is presumed to be the conversation*

The **DECISION** is being influenced by what AI ingested *and* what it inferred



* Behind the scenes AI creates additional “unseen” data that is being used to influence action

- **Extracting** structured signal from unstructured content
- **Characterizing** sentiment, topic, urgency, intent
- **Inferring** relationships, status, behavior, priorities
- **Retaining** derived characterizations in places governance doesn't see

The Invisible Inference Dataset

Signals extracted from source data become the inference dataset on which AI performs its analysis. Homegrown models have some control over dataset properties, but these configurations are IP for many vendors and therefore inaccessible.



Example – Copilot reading SharePoint

You give Copilot access; it returns answers. How it decides what's relevant, how it interprets, how it composes – all the vendor's black box IP.

Mobley v. Workday (N.D. Cal., 2023)

2026 - Court accepted vendor AI can function as agent of the employer. Workday faces direct liability. Every employer using Workday's AI is now potentially liable for what Workday's AI did.

**When you cannot govern what you cannot see -
*You must govern the resulting decision***

The Pattern Across These Deployments

None of these felt like multimodal AI when they were procured, yet all of them are.


Each of these deployments:

- Consumes inputs your governance program never inventoried
- Produces inferred data that lives where governance doesn't see
- Influenced a productivity or operational decision with no codified context

This is happening on a foundation that was not solid:

- Unstructured content was the unfinished work of data governance
- AI did not create that gap

AI industrialized it



Multimodal AI Requires a New Level of Enterprise Collaboration

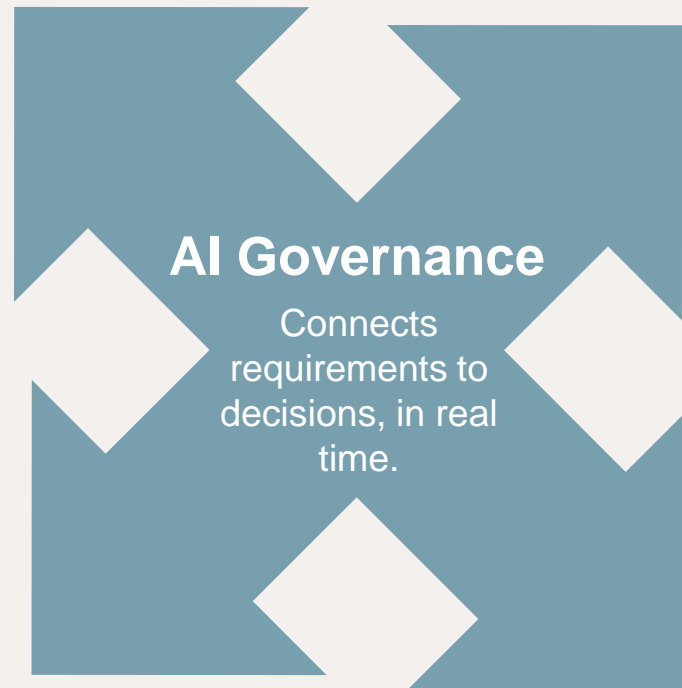
Four disciplines have new work; AI Governance is at the center

Data Governance

Must extend the disciplines it built for structured data to multimodal data.

Privacy

Must extend its discipline to reach the inferences AI is drawing about people and the actions that follow from them.



Ethics

Must move from deployment review to decision review - present each time a consequential AI decision is made.

Compliance

Must identify the regulatory events vendor AI is creating on the organization's behalf.

Examining the Decision with a Multimodal Lens

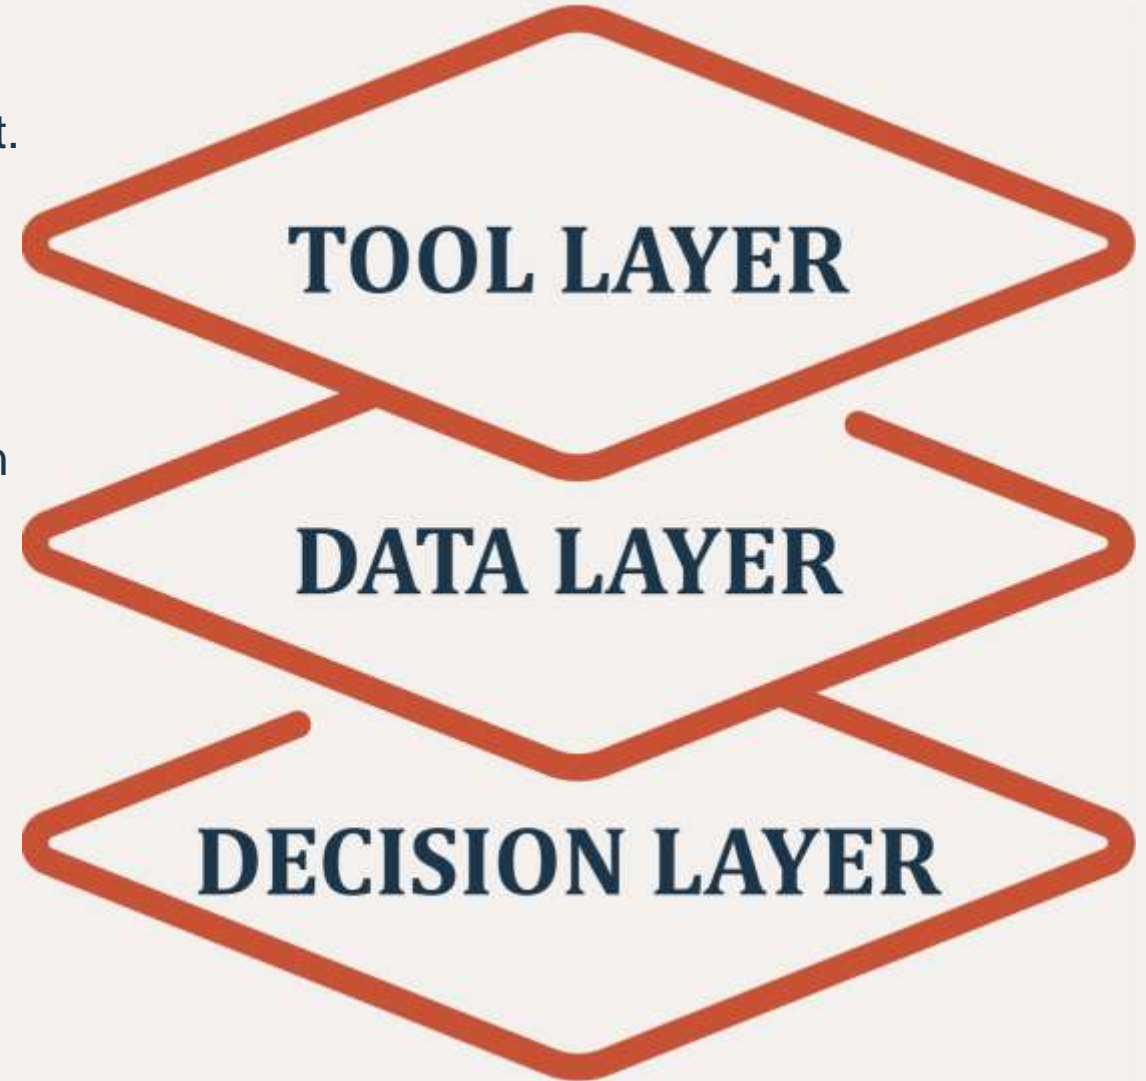
From May: One Event, Three Governance Questions

A salesperson uses ChatGPT to draft a customer commitment email.

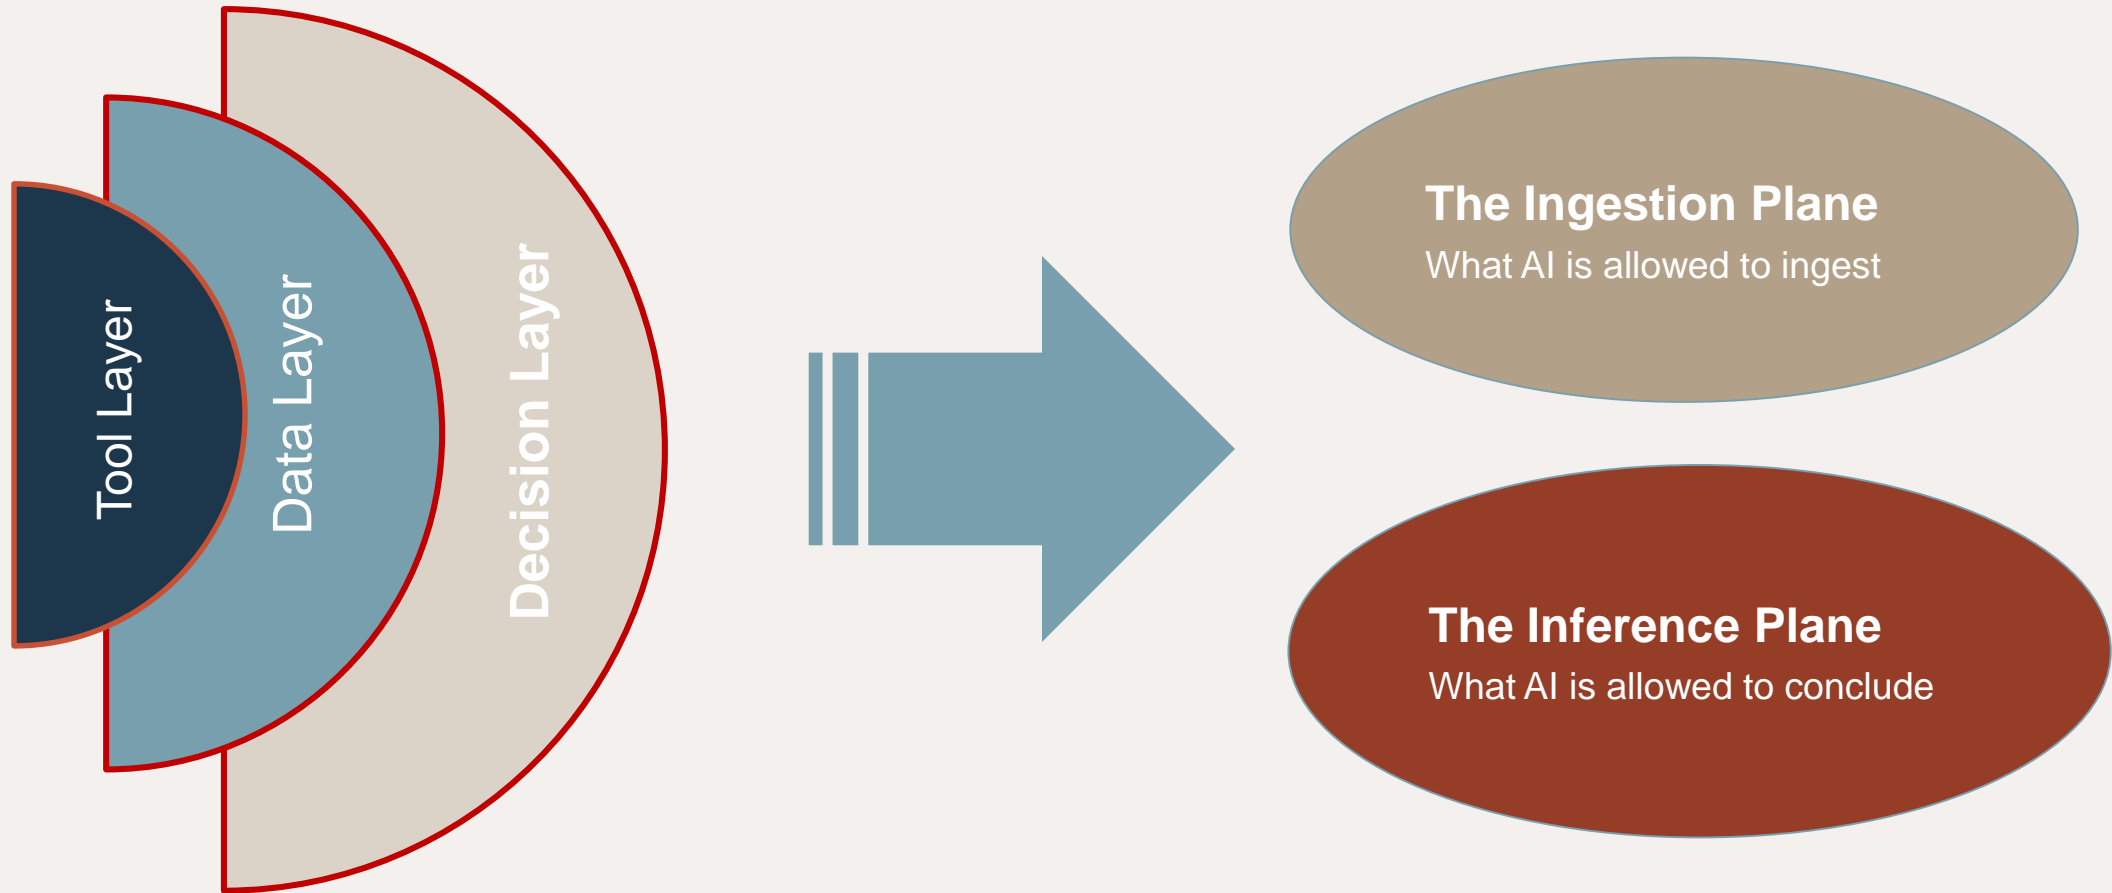
Tool layer. ChatGPT, browser, personal account.
Tool-layer detection: medium-priority generic AI use.

Data layer. Customer name, account size, contract terms, internal renewal pricing crossed an external boundary. High-classification breach.

Decision layer. A commitment was made on behalf of the company, with terms partly authored by an AI you have no audit access to. If the customer holds you to those terms, there is no record of what shaped them.



Applying the Decision Layer to Multimodal Data



One Event > Two Planes > One Decision


A meeting transcript under the microscope

The Ingestion Plane

The bot takes in: Tone.
Speakers. Words. Pace.
Pauses. Voice.

The Inference Plane

The bot concludes: Who
dominated. Who disengaged.
Sentiment. Action items.
Deal probability. Emotional state.



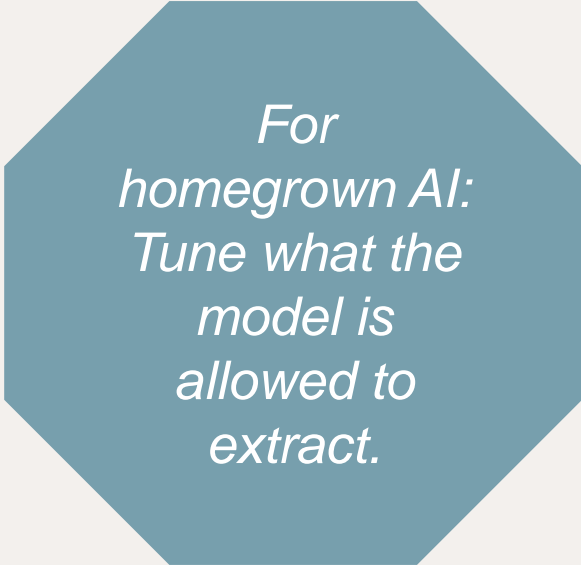
What follows: **Decisions**

Summary distributed. Scorecard filed.
Pipeline updated.
Coaching flag raised. Retention
concern escalated.

Govern the Ingestion Plane

You cannot tell Copilot how to process data, You can tell it which documents to read.

- 1. Decide what AI tools are allowed***
- 2. Decide what AI can connected to***
- 3. Decide what content is out of scope***



*For
homegrown AI:
Tune what the
model is
allowed to
extract.*

Govern The Result of the Decision – The Action

You cannot govern what AI concludes. You can govern what humans do with what AI surfaces.

1. *The Action Register*

2. *The Human Checkpoint*

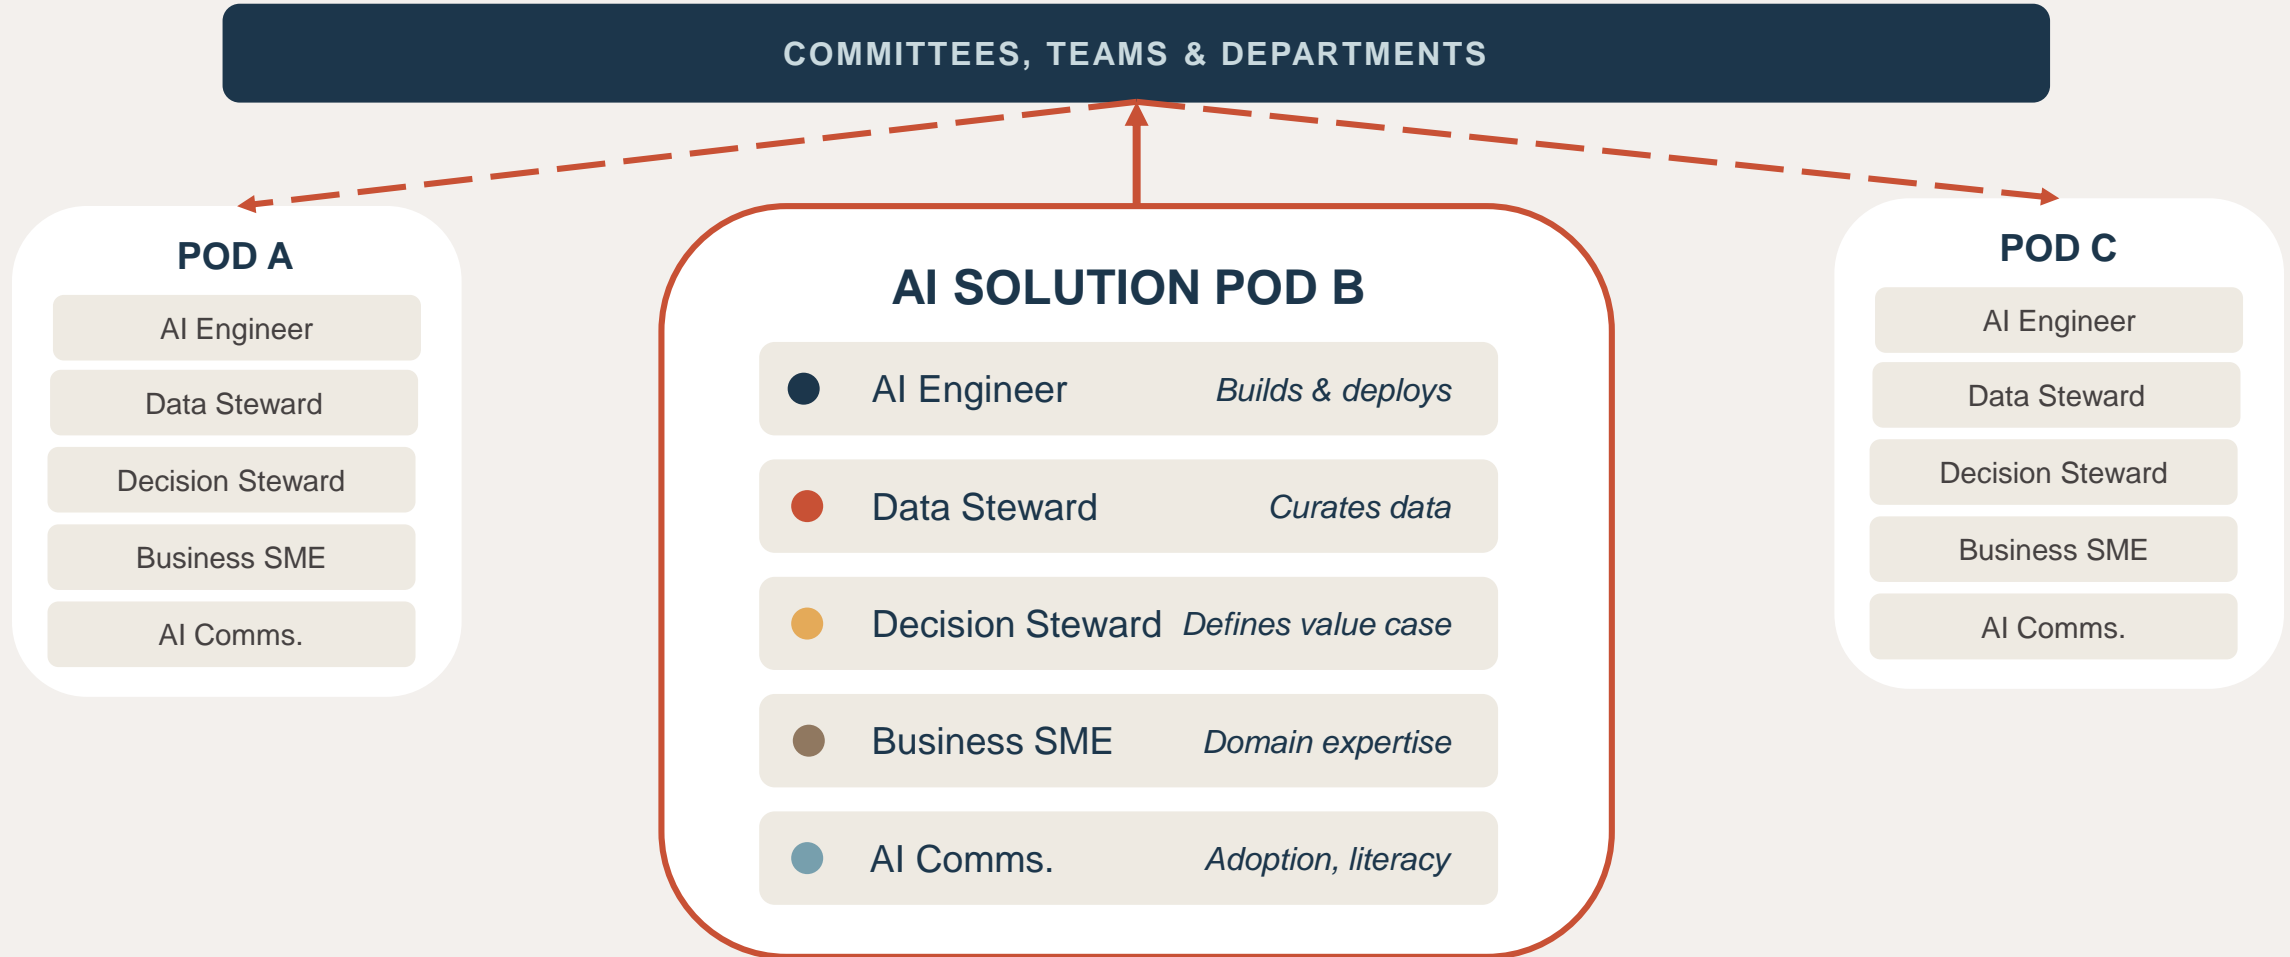
3. *The Evidentiary Trail*

*For homegrown
AI: Extend the
register
backward into
the inference
layer.*

Governing in Real Time – The Pod Revisited

From April: 5 in-the-Pod – the connective tissue for governance at scale

Each pod builds targeted AI solutions at speed, while enabling lateral thinking, upward alignment, and accountability.



5-in-the-Pod for Multimodal Decisions

Scaling decisions at the speed your business moves – from Agent to Enterprise

Solution Focus

AI Engineer *Builds & deploys*

Data Steward *Curates data*

Decision Steward *Defines value case*

Business SME *Domain expertise*

AI Comms. *Adoption, literacy*

Decision Focus

Technical Insight

*Data governance and **Compliance***

*AI Governance and **Ethics***

*Use Case and **Decision Accountability***

Privacy and Disclosure

Why this isn't just data governance

And why it doesn't take work away from Privacy or Compliance

Privacy & Compliance

Sets the requirements

- What disclosure looks like
- What lawful basis is needed
- What the regulatory obligations are

Data Governance

Makes requirements operational for data

- Ensures the data complies
- Sets up operating models
- Builds decision structures

Doesn't replace privacy or compliance, it makes them enforceable.

AI Governance

Same thing — for AI decisions

- Makes privacy, compliance, ethics, and data governance requirements operational *at the moment AI is producing decisions*

The requirements come from where they have always come from. The implementation is what's new.

Newer Modalities

Voice and Video — New Data Modalities

How these modalities are reshaping the inference plane

Voice analytics on customer interactions

Perception: tone, pace, breathing, vocal tremor

Inference: satisfaction, stress, escalation intent, demographic guesses, fraud risk

Action: routing, priority, retention flags, offer triggers

Video and image inference

Perception: facial expressions, gaze, posture, facial geometry

Inference: emotion, attention, engagement, identity, demographic categorization

Action: hiring (and firing) decisions, security response, access control



Voice and Video - New Considerations

Voice and video aren't just new modalities. They carry additional legal and ethical weight

Biometric data classification Voice prints and facial geometry are special-category data under most privacy regimes. Different consent bar, different retention requirements, different private-right-of-action exposure. Most practitioners blur this with regular voice and video processing.

Protected-characteristic inference happens by default Gender, age range, ethnicity, disability. Voice and video draw these as side effects of doing the authorized work. The EU AI Act Article 5 prohibition is structured around this fact: the inference itself is the issue, not just how it's used.

The subject often doesn't know the AI is observing Meeting transcript bots announce themselves. Voice analytics on the same call doesn't. Video at retail entrances doesn't. Emotion inference in interview platforms doesn't. The "you know the AI is in the room" comfort from Slide 8 disappears, and notice/consent frameworks become structurally harder.

In Summary

AI governance is the discipline that makes the requirements operational at the moment a decision gets made.



Data Governance

Ensures the inputs are accurate



Privacy

Defines the right to know and refuse



Compliance

Names the regulatory obligations



Ethics

Names the values choices



AI Governance

Makes all of it operational at the decision

*The pod is where it happens. Multi-discipline roles. One pod. Every decision.
The decision is what differentiates AI Governance from Data Governance.*

The Message to Carry Forward

SAME WORK. DIFFERENT ALTITUDE.

For Your Next Board Update:

- *AI governance is the discipline that operates at the decision.*
- *Making the requirements of data governance, privacy, compliance, and ethics operational where AI is producing outputs the organization will act on.*

For your next Governance Committee:

- *Govern what AI ingests*
- *Govern what humans do with what it surfaces.*
- *Name the pod*



Let's Stay Connected
Scan the QR Code



Kelle O'Neal | Founder & CEO
Lisa Wintrick | Executive Advisor



Citations and Sources

All material referenced in this outline has been verified at primary source.

EU AI Act — prohibitions and enforcement

- Regulation (EU) 2024/1689 — the EU AI Act. Article 5(1)(f) prohibits emotion inference in workplace and education. Effective Feb 2, 2025; enforceable by Member State authorities since Aug 2, 2025.
- Official AI Act portal — <https://artificialintelligenceact.eu/article/5/>
- AI Act Service Desk (EU Commission) — <https://ai-act-service-desk.ec.europa.eu/en/ai-act/article-5>
- European Commission Guidelines on prohibited practices, C(2025) 884 final, Feb 4, 2025

Digital Omnibus — provisional delay of high-risk obligations

- Council of the EU — Provisional agreement reached May 7, 2026, deferring Annex III high-risk AI obligations to Dec 2, 2027 and embedded high-risk obligations to Aug 2, 2028. Final adoption pending.
- Council of the EU press release (May 7, 2026) — <https://www.consilium.europa.eu/en/press/press-releases/2026/05/07/artificial-intelligence-council-and-parliament-agree-to-simplify-and-streamline-rules/>
- Inside Privacy / Covington analysis (May 18, 2026) — <https://www.insideprivacy.com/artificial-intelligence/eu-ai-act-update-timeline-relief-targeted-simplification-and-new-prohibitions/>

CNIL (France) — 2026 enforcement priorities

- CNIL announced on April 2, 2026 that recruitment is its #1 GDPR enforcement priority for 2026, with focus on automated decision-making tools.
- Lexology coverage (April 3, 2026) — <https://www.lexology.com/library/detail.aspx?g=f4a54dcb-2909-4f77-8fcf-2d6829d287c8>

Mobley v. Workday — AI hiring discrimination

- Mobley v. Workday, Inc., No. 3:23-cv-00770 (N.D. Cal.) — class action alleging Workday’s AI-driven applicant screening discriminated by age, race, and disability. Plaintiff rejected from over 100 jobs through Workday’s platform, often within minutes.
- July 2024 — Motion to dismiss denied; case allowed to proceed.
- May 16, 2025 — Preliminary collective certification under the ADEA.
- July 7, 2025 — Workday’s HiredScore AI features within scope.
- February 2026 — Nationwide collective action authorized. Opt-in deadline March 7, 2026.
- March 6, 2026 — Court rejected Workday’s argument that ADEA disparate-impact doesn’t apply to applicants. Court accepted theory that vendor AI can function as agent of employer.
- Civil Rights Litigation Clearinghouse — <https://clearinghouse.net/case/44074/>
- Court notice authorization (Feb 17, 2026) — <https://www.wigginschilds.com/news/workday-case-update/>

NOTE: Case in active litigation. Lisa should refresh status the day of delivery.

Prior FSFP sessions referenced

- January 2026, Session 1 — “Cutting Through the Noise” — visibility, decision rights, controls
- February 2026, Session 2 — “Lessons from Enterprise Leaders” — how you decide matters as much as what you decide
- March 2026, Session 3 — “AI Governance vs Data Governance” — decision stewardship; data governance evolves in service to AI; Decision Stewardship Committee structure
- April 2026, Session 4 — “Designing Accountability” — if it isn’t codified, it isn’t real; the Accountability Orbit; 5-in-the-Pod team structure (Decision Steward, Data Steward, AI Engineer, Business SME, AI Communications)
- May 2026, Session 5 — “Governance in the Wild” — govern the layer, not the tool; three layers (tool / data / decision)