

The Biggest AI Assumption: Why Governance Begins Before AI

Conference Speaker Package

Peter Gallinari

Author, *Rockin Data Privacy: The Truth About Your Data*

Former Chief Data Privacy Officer – State of Tennessee

Former Chief Security Officer – GE Capital

Introduction

After more than 52 years in technology, cybersecurity, privacy, governance, healthcare, financial services, and government, I have reached a conclusion that seems increasingly absent from today's AI discussions: we are focusing on the wrong starting point.

Organizations are investing heavily in Responsible AI, AI governance, guardrails, oversight, registries, frameworks, and accountability models. Yet many of these efforts begin with a critical assumption—that organizations already understand the environments AI is being introduced into.

This session challenges that assumption and introduces a practical framework for evaluating whether foundational governance is mature enough to support AI at scale.

Session Abstract

As organizations race to deploy Generative AI, Agentic AI, copilots, and intelligent automation, most governance discussions focus on AI frameworks, guardrails, oversight committees, policies, and regulatory compliance.

But what if we are starting in the wrong place?

The greatest AI risk may not be the model, the prompt, or the agent. It may be the assumption that organizations already understand the data environments AI is being introduced into.

Many organizations continue to struggle with incomplete data discovery, inconsistent classification, unclear ownership, excessive permissions, shadow data, fragmented accountability, and limited visibility into how data moves across the enterprise.

This session explores why AI governance cannot compensate for foundational governance weaknesses and why AI often amplifies existing operational gaps through speed, scale, automation, and exposure.

This is not an anti-AI presentation. It is a reality check on why governance begins long before the first prompt is ever entered.

Key Message Graphic



WE'RE NOT FAILING AT AI GOVERNANCE.

WE'RE DISCOVERING THAT WE NEVER FULLY UNDERSTOOD OUR DATA ENVIRONMENTS IN THE FIRST PLACE.

After **52 years** in technology, cybersecurity, privacy, governance, healthcare, financial services, and government, I have reached a conclusion that seems increasingly absent from today's AI discussions.

WE ARE FOCUSING ON THE WRONG STARTING POINT.

EVERY DAY I SEE CONVERSATIONS ABOUT:

- AI governance
- Responsible AI
- AI guardrails
- AI oversight
- AI registries
- AI accountability
- AI frameworks
- AI committees

Most of these conversations are well intentioned.

BUT THEY ALL SEEM TO BEGIN FROM THE SAME ASSUMPTION:



That organizations already understand the environment AI is being introduced into.

I AM NO LONGER CONVINCED THAT ASSUMPTION IS TRUE.

THE QUESTION I RARELY HEAR ASKED IS:



Do organizations actually know where their data resides today?

- Not from a policy.
- Not from a spreadsheet.
- Not from a governance presentation.
- Operationally.

DO THEY UNDERSTAND:

- Data Discovery
- Data Classification
- Ownership
- Stewardship
- Privacy Obligations
- Permissions
- Shadow Data
- Data Movement
- Operational Accountability

Because if they don't, then AI governance may be attempting to govern an environment that is not yet fully understood.

- AI **did not create** these problems.
- AI **inherited** them.

And in many cases, it **amplifies** them through scale, speed, automation, and exposure.

THAT IS WHY I CONTINUE TO CHALLENGE THE INDUSTRY ON ONE SIMPLE POINT:

DATA DISCOVERY → DATA CLASSIFICATION → DATA OWNERSHIP → DATA PRIVACY → DATA SECURITY

THESE ARE NOT AI GOVERNANCE ACTIVITIES. THEY ARE PREREQUISITES TO AI GOVERNANCE.

THE BIGGEST AI RISK MAY NOT BE:

- THE MODEL
- THE PROMPT
- THE AGENT

IT MAY BE THE ASSUMPTION



That we understand our data environments far better than we actually do.

UNTIL WE ADDRESS THAT REALITY,

We will continue to circle the problem without ever landing it.



AND THAT MAY BE THE CONVERSATION THE INDUSTRY NEEDS MOST RIGHT NOW.

Peter Gallinari
Author | Strategic Advisor | Rockin Data Privacy

Learning Objectives

- Why AI governance may be starting too late.
- The biggest assumptions organizations make about AI readiness.
- Why internal AI does not automatically reduce governance risk.
- How AI amplifies existing governance weaknesses.
- A practical readiness model for AI adoption.

Core Framework

Data Discovery → Data Classification → Data Ownership → Data Privacy → Data Security → AI Governance

These are not AI governance activities. They are prerequisites to AI governance.

Target Audience

- CIOs and CISOs
- Chief Data Officers
- Privacy Officers
- AI Governance Leaders
- Data Governance Professionals
- Security Architects
- Risk and Compliance Leaders
- Executive Leadership Teams
- Boards and Risk Committees

Why This Session Is Different

Most AI presentations focus on governance frameworks, Responsible AI, guardrails, ethics, controls, and compliance.

What if the greatest AI risk is the assumption that organizations understand their data environments far better than they actually do?

Rather than discussing how to govern AI, this session examines whether organizations have sufficiently governed the environment AI depends upon.

Speaker Biography

Peter Gallinari is a technology, cybersecurity, privacy, and governance executive with more than 52 years of experience spanning healthcare, financial services, government, and private industry.

His leadership roles have included Chief Data Privacy Officer for the State of Tennessee, Chief Security Officer for GE Capital, Privacy Program Manager, and executive technology and governance leader.

Peter is the author of *Rockin Data Privacy: The Truth About Your Data* and host of the Rockin Data Privacy podcast.

His work focuses on helping organizations identify the gap between governance on paper and governance in practice, with a particular emphasis on data visibility, accountability, privacy, and operational readiness in the age of AI.

"We're not failing at AI governance. We're discovering that we never fully understood our data environments in the first place."