




IAPP PRIVACY. SECURITY. RISK. + AI GOVERNANCE GLOBAL 2026

Training 6-7 October

Workshops 7 October

Conference 8-9 October

AI GP exam 9 October

SEATTLE

#IAPPpsr26

Governance for the Age of Frontier Models

A board-ready approach to governing frontier capabilities



#IAPPpsr26

WELCOME AND INTRODUCTIONS



Nandita Das, CIPT
Senior Manager
EY Technology Consulting



Naomi Matin
US Privacy Leader
EY

#IAPPpsr26



AGENDA OUTLINE

- I. Welcome and Introductions
- II. Why governance needs to change
- III. What executives should own
- IV. How to operationalize quickly
- V. Q&A



Frontier capability leapt ahead. Governance didn't.

Frontier models now act autonomously, at machine speed, across systems you don't own.

AUTONOMY	BEFORE AI assists human experts	NOW Models discover, chain and act end-to-end without a human in the loop
SPEED	BEFORE Weeks of expert effort from discovery to impact	NOW Hours: discovery-to-exploit now runs at machine speed
REACH	BEFORE Mature, well-tested systems assumed safe	NOW Flaws hidden for decades surfaced by AI

ONE FRONTIER MODEL, THREE MONTHS

0 → 181

working exploits in the same benchmark: prior Claude model vs. Mythos Preview ¹

10,000+

high/critical vulnerabilities found by early partners in one month ²

< 1 hour

to a first working exploit for a newly disclosed vulnerability ³

Model: Claude Mythos (Anthropic).
Sources: See [Resource List](#) slide

The question isn't whether it arrives. It's whether **your decisions can keep pace.**

Frontier risk arrives from both directions

What you deploy and what's deployed against you need one governance model, not two.

OUTBOUND

WHAT YOU DEPLOY

- Agents with tool access and standing permissions
- Copilots reaching sensitive and personal data
- AI-driven decisions affecting customers and employees
- Frontier AI built into the SaaS and vendor tools you buy

GOVERN ONCE

SHARED FOUNDATIONS

- Data inventory and classification
- Identity and access, including agents
- Third- and fourth-party oversight
- Incident response and resilience

INBOUND

WHAT'S DEPLOYED AGAINST YOU

- Machine-speed discovery and exploitation
- AI-assisted data exfiltration
- Attackers exploiting flaws in vendor and open-source software you rely on
- Remediation volume that outpaces patch and change processes

Four stress tests for your current governance

1 POLICY

Does policy say who can approve an autonomous use, and what evidence is required?

2 DECISIONS

Can you approve, pause or roll back an agent in hours, not at next month's committee?

3 SYSTEMS

Do controls sit in the access and deployment path, or only in documents?

4 EVIDENCE


Could you show a regulator today what each AI system can reach?

Governance becomes real when policy, decision rights, controls and evidence connect.

#IAPPpsr26

Board oversight & management ownership

Frontier governance fails on unclear decision rights, not missing principles.




Pre-agree who can pause a frontier capability, **before you need to.**

#IAPPpsr26

Set risk appetite by autonomy and reach

Tier by what a capability can do and touch, not by model name. Oversight scales up from there.




Tier by what it can do and touch, then apply the same rules every time.

#IAPPpsr26

One control plane, one source of evidence

Build controls once and reuse the evidence for privacy, security, audit, the board and regulators.



What the board sees

Five leading indicators, each backed by control-plane evidence

- % of frontier uses tiered, with a named owner
- Agents with standing privileged access
- Overdue remediation from failed tests
- Critical vendors without AI terms or attestation
- Time to pause or roll back a high-tier use

DECISIONS IT DRIVES

Approve • invest • remediate • pause • retire

Report to drive decisions, not to fill a dashboard

#IAPPpsr26

90-day agenda: make governance usable

Start narrow, prove the operating model and scale from evidence.

0-30 DAYS

Locate + prioritize

- Inventory active and planned frontier use cases.
- Identify high-impact processes and external vendors.
- Confirm accountable business and technology owners
- Define the first risk appetite guardrails

31-60 DAYS

Design + embed

- Define decision rights and the escalation path
- Apply tiering to the inventory
- Embed controls in procurement, access and deployment paths
- Perform pilot testing for selected high-risk use cases

61-90 DAYS

Operate + evidence

- Start board reporting on the five leading indicators
- Run one incident / rollback exercise
- Review exceptions and overdue remediation
- Agree scale plan, funding and cadence

Five questions for your board

1

Where does frontier capability act on our data today?

2

Who can pause it, and how fast?

3

Which high-tier uses sit outside our risk appetite?

4

Which critical vendors can't show evidence of their controls?

5

What evidence shows our controls work?

Begin with a focused 90-day sprint and let evidence decide what scales.

#IAPPpsr26

RESOURCE LIST

- 1 Anthropic Frontier Red Team (2026, April 7). *Assessing Claude Mythos Preview's cybersecurity capabilities.*
<https://www.anthropic.com/research/mythos-preview>
- 2 Anthropic (2026, May 22). *Project Glasswing: An initial update.*
<https://www.anthropic.com/research/glasswing-initial-update>
- 3 Anthropic Frontier Red Team (2026, June 8). *Claude Mythos Turns N-Days Into N-Hours With Rapid Exploit Creation*
<https://www.securityweek.com/claude-mythos-turns-n-days-into-n-hours-with-rapid-exploit-creation/>

How Did Things Go? (We Really Want To Know)

Did you enjoy this session? Is there any way we could make it better? Let us know by filling out a speaker evaluation.

1. Open the IAPP Events app.
2. Select **IAPP P.S.R. + AIGG 2026**
3. Tap "Schedule" on the bottom navigation bar.
4. Find this session. Click "Rate this Session" within the description.
5. Once you've answered all three questions, tap "Done".

Thank you!

#IAPPpsr26

