

THE AI INSTITUTE / RESEARCH FOR LEADERS

This Week, AI Governance Gets Specific

Three dated moments will turn broad principles into operating decisions on data, human judgement and assurance.

Boards, C-suite and senior management · Research period September 2026 – September 2026

[Download designed PDF ↓](#)

A single enterprise AI policy is no longer enough; leaders must translate principles into domain rules for data access, retained human judgement and proof before deployment.

KEY POINTS

What this paper means for leaders

-
- 01 **7 September: decide the organisation's position on data access across organisational boundaries.**
 - 02 **8–11 September: state which decisions and forms of participation must remain visibly human.**
 - 03 **9–10 September: define what must be demonstrated before AI enters a high-consequence process.**
 - 04 **Treat consultations and conferences as early design inputs, not as enacted law.**
 - 05 **By Friday, name an owner and produce one practical rule for each domain.**
-

One policy cannot answer three operating questions

A strong AI policy can state values. It cannot decide whether another organisation may act for a customer, what counts as genuine participation in an AI-mediated institution, or which tests permit AI into a safety-critical process.

Those questions move forward this week through separate UK, European and global forums. They are not one regulatory event, and most are not new law. Their value is as an early warning: governance is leaving the principles page and entering the operating model.

The board-level response is simple. Do not commission another general policy. Assign three practical decisions to the executives who already own the underlying work.

Governance becomes useful when it changes a permission, preserves a human decision or sets a deployment threshold.

Decide who may use data—and on whose behalf

The UK's consultation on data intermediaries closes on 7 September. It asks how third parties might act for individuals to exercise data rights, and where legal ambiguity or operational friction blocks that model.¹ A wider UK call for evidence on data regulation and AI closes on 9 September, covering access, quality, governance across organisations, transparency and rights.²

These are consultations, not obligations. But the business question already exists wherever AI depends on data moving between customers, platforms, suppliers and partners. Ask the chief data, legal and product leaders for one position: which permissions follow the data, who can exercise them, and how a person can challenge or withdraw that authority.

This UK process may influence other markets, but it cannot stand in for them. Privacy law, data portability, sector rules and consumer expectations differ. Multinationals need a shared decision structure with jurisdiction-specific answers.

03 8-11 SEPTEMBER

Name the judgement that must remain human

UNESCO's Digital Learning Week will address synthetic knowledge, AI tutors, autonomous systems and what it calls “ghost learners”—a challenge to assumptions about evidence, participation and presence.³ The immediate audience is education, but the question travels into every organisation that trains, assesses, credentials or promotes people.

If AI can draft the work, coach the learner and help assess the result, leaders need to say where human judgement remains visible. That may be authorship, final assessment, professional sign-off, promotion or a decision affecting a person's rights.

UNESCO guidance is influential rather than binding. Infrastructure, language and institutional capacity also change what responsible use looks like. The durable executive output is not a universal ban or permission. It is a short list of decisions that remain human, plus the record that shows they were.

04 9-10 SEPTEMBER

Define what must be proven before deployment

ASA's AI Days will connect updates to its proposed AI concept paper with ongoing aviation rulemaking, assurance, human factors and ethics-based assessment.⁴ Aviation is a specific regulated sector; its requirements should not be copied into every workflow.

The broader discipline does travel. A supplier demonstration is not assurance. Before AI enters a high-consequence job, the accountable leader should know which failure conditions were tested, who can intervene, what changes trigger retesting and which evidence is acceptable to the relevant authority.

A Council of Europe and ETH Zurich conference on 7 September reinforces the same movement from ethical and legal principles toward implementation.⁵ The useful signal will be concrete changes to supervision, procurement or deployment—not the strength of the speeches.

THREE OUTPUTS FOR FRIDAY

DATA

Permission

Who may use which data, for whom, and how authority ends.

Owner: data + legal

PEOPLE

Judgement

Which decisions remain visibly human and on the record.

Owner: workforce + risk

PROOF

Assurance

What must be demonstrated before high-consequence use.

Owner: operations + assurance

05

BY FRIDAY

Ask for three one-page rules, not another principles workshop

Regional events in Tunisia, Australia and Hungary will add investment, infrastructure and adoption perspectives this week. They matter because implementation capacity, language, public institutions and market structure change what can be done. They should not be treated as proof of regional adoption.

Use the week to collect specifics. Which permissions are changing? Which decisions remain human? Which tests become a real threshold? Ask each owner for one page that can change a contract, workflow or release decision.

If the week's outcomes remain declaratory, keep the current plan. If they produce a clearer right, assurance test or sector expectation, update the relevant domain rule. That is the standard: not more governance material, but a better operating decision.

The Friday test – can each principle now change a real permission, decision or release?

Method and limitations

Method

This Week Ahead agenda uses official UK Government, UNESCO, EASA, ETH Zurich/Council of Europe and regional event pages current at 6 September 2026. It distinguishes consultation, conference, multilateral guidance, proposed concept paper and rulemaking, and uses the latest arXiv batch only as underlying R&D radar. The public treatment selects decisions rather than summarising every event.

Limitations

Consultations and conferences do not create immediate legal obligations. Outcomes were not available at the research cutoff. EASA's work is sector-specific, UNESCO guidance is non-binding, and UK processes apply within their stated jurisdiction. Regional event programmes do not establish market-wide adoption or value.

First published 6 September 2026 · Updated 6 September 2026 · Research period September 2026 – September 2026 · Research current to 6 September 2026 · Version 1.0 · Suggested citation: The AI Institute, *This Week, AI Governance Gets Specific* (2026).

References and source notes

01	UK Government, Empowering people through data intermediaries ↗ Official open consultation, published 8 June and updated 27 August 2026; closes 7 September 2026.
02	UK Government, Data regulation in the age of AI ↗ Official call for evidence; closes 9 September 2026. It is evidence gathering, not enacted policy.
03	UNESCO, Digital Learning Week 2026 ↗ Official programme for 8–11 September 2026; multilateral and globally oriented, but not legally binding.
04	EASA, Artificial Intelligence Days 2026 ↗ Official regulator event for 9–10 September 2026, covering assurance, human factors, proposed concept paper and active rulemaking.
05	ETH Zurich and Council of Europe, Conference on Trustworthy AI ↗ Official conference page for 7 September 2026; implementation discussion rather than a binding outcome.
06	AI Forward Summit, Tunisia ↗ Official programme for 7–10 September 2026; regional infrastructure, governance and investment agenda.
07	The AI Summit Australia ↗ Official event page for 7–9 September 2026; an industry forum, not market-wide adoption evidence.
08	arXiv, recent cs.AI submissions ↗ The latest available batch at cutoff was Friday 4 September 2026 with 165 entries. Preprints are not peer-reviewed by default.