



Inter-Parliamentary Union
For democracy. For everyone.



IPU Centre for Innovation in Parliament

Parliamentary Data Science Hub in-person meeting

13–15 May 2026 | Chamber of Deputies of Italy, Rome

Implementing AI in parliaments: Taking stock and planning the next steps together

The 2026 in-person meeting of the [IPU Centre for Innovation in Parliament's \(CIP\)](#) Parliamentary Data Science (PDS) Hub took place at the Chamber of Deputies of Italy in Rome on 13–15 May 2026.

Co-organized by the CIP, the Italian Chamber of Deputies and the Chamber of Deputies of Brazil, the meeting was attended by senior officials from 23 parliaments from Africa, the Americas, Europe and the Middle East.

The meeting was convened at a pivotal moment: according to preliminary data from the *World e-Parliament Report 2026*, some 88% per cent of parliaments now report using AI, yet governance, capacity and culture are struggling to keep pace.

The focus of the session was on moving beyond theoretical discussion and towards the practical implementation challenges parliaments face when adopting AI, building on experience gained since the previous PDS Hub meeting in The Hague in June 2025.

The principal objectives were as follows:

- Reflect honestly on progress on AI in parliaments over the preceding 12 months, including obstacles and failures as well as successes.
- Shift the focus of community discussion from strategy and theory to practical implementation.
- Address the organizational transformation that AI requires, including its cultural and human dimensions.
- Promote peer learning across a community with diverse levels of digital maturity.
- Agree actionable future deliverables, including updates to shared guidelines, and new case studies.

Key findings

The three days of discussion were structured around thematic sessions. The findings below synthesize the main messages to emerge from those sessions.

AI adoption is outpacing governance

The most consistent thread across the entire meeting was the widening gap between AI adoption and institutional governance. Much AI use in parliaments is informal, undeclared and ungoverned. The strategic question facing institutions is no longer whether to adopt AI, but whether they can govern, resource and absorb it quickly enough.

Shadow AI – whereby staff and MPs use personal or unsanctioned AI tools and accounts – was identified as a pervasive and growing challenge, but the discussion also reframed it as a potential signal rather than simply a problem. Where shadow AI is widespread, it can suggest that officially sanctioned tools are not meeting users' needs. Shadow AI need not be a negative: it can prove to be a useful space for experimentation and innovation, but care is required. One institution described a practical response: rather than attempting outright prohibition, it built a governed application programming interface (API) library that shadow AI systems can connect to, bringing them within tested infrastructure while acknowledging the reality that staff will use available tools. Indeed, prohibition alone rarely works; governance that addresses the supply side is more likely to succeed.

A related concern was MPs purchasing AI tools using parliamentary allowances, outside any institutional contract or data-processing framework. Several parliaments described being forced to assume procurement responsibilities they had not anticipated.

The capacity gap is the primary barrier

Across the community, and particularly for parliaments in developing and smaller nations, the principal barriers to AI implementation are capacity first and funding second. Among parliaments at intermediate levels of digital maturity, capacity consistently emerges as the leading constraint. The least digitally mature parliaments currently have near-zero AI governance structures.

Only a few parliaments have established dedicated AI staff positions; the majority are managing AI governance on top of existing roles, creating real pressure. Discussion of dedicated AI roles pointed to a competency-led approach: rather than asking “Do we need AI staff?”, the more productive question is this: “What competencies need to exist and where should they live in the organization?”.

The discussion identified a set of functions that typically need to be explicitly assigned rather than left to be absorbed informally into existing responsibilities: strategic decision-making authority over AI; organization-wide coordination; product ownership for specific AI deployments; the creation of enabling conditions for front-line staff; and day-to-day AI use. Where these functions have no named owner, they fall through the cracks.

Alongside AI-specific roles, the importance of data stewardship was emphasized: clear data owners and data stewards are a prerequisite for any serious AI programme. Without defined accountability for data quality, provenance and access, AI systems cannot be

governed reliably. Legal advisory, procurement, AI literacy and process re-engineering were also identified as cross-cutting functions that need explicit ownership rather than being treated as someone else's responsibility.

Culture and people matter more than technology

A consistent view to emerge from the meeting was that the majority of successful AI adoption is about process, education and culture rather than technology. Parliaments that treat AI primarily as a technology challenge tend to underperform those that invest in people and change management first.

Managers were identified as a critical but often overlooked pressure point: they are expected to drive adoption from above while managing anxiety from below. Social learning – experience-sharing, peer networks and AI champions – was found to be more effective than top-down implementation. One parliament's internal innovation initiative, which attracted a large volume of use-case submissions from staff across the institution, was offered as a concrete illustration of the social learning model in action.

The meeting also surfaced a concern about the longer-term human cost of AI adoption: as AI accelerates and automates tasks that previously formed the backbone of junior staff development, institutions need to think carefully about how new parliamentary staff will acquire the institutional knowledge and procedural fluency that have traditionally accumulated through years of hands-on practice.

Data governance is the foundational layer

Strong AI governance depends on strong data governance – and the two should be designed together rather than treated as separate workstreams. Where data quality, provenance and access controls are poorly defined, AI systems built on that data will inherit and amplify those weaknesses. Conversely, where institutions have invested in clear data ownership, formal data stewardship roles and defined processes for data quality, they are significantly better placed to deploy and govern AI systems reliably.

The meeting identified challenges around privileged information – material accessible to certain users but subject to publication restrictions. The risks here were illustrated through discussion of major productivity AI tools: a deployment with access to unpublished committee material could inadvertently incorporate it into generated content, creating a real risk of parliamentary rule violations. This is not a hypothetical concern; the potential for AI tools to draw on sensitive material without explicit instruction to the contrary is a genuine governance gap. Using document labelling to instruct AI tools to ignore sensitive documents was proposed as a practical control, although this approach depends on robust and consistent labelling practices being in place first.

A consistent observation was that institutions are legally responsible for data processing but often lack effective control over the data itself. The most persistent governance gap relates to unstructured data such as emails, working documents and files in transit. Yet it is precisely this data that open productivity AI tools draw upon most heavily.

The recommended approach is use-case-driven and incremental: start with a limited, well-governed data set for a specific purpose and get the controls right for that use case before moving to the next. It is also important to recognize that data changes over time and that it needs to be actively managed, not treated as something that is collected once and then left untouched. A further point was that AI governance policy must explicitly encompass commercial and externally procured AI systems, not only those built in-house: data governance requirements, risk assessment, human oversight and bias management apply equally to tools acquired from third-party providers.

Sovereignty means managing dependency, not achieving independence

A discussion on AI sovereignty framed the concept as something that operates across four dimensions simultaneously: ethical, political, legal and technical. This framing is useful precisely because it resists reduction to a purely technical question: decisions about dependency and control are also decisions about values, accountability and institutional identity. The practical conclusion that emerged was that absolute technological independence is neither realistic nor desirable for most parliaments. Instead, the operative concept is the management of dependency: understanding reliance on external providers, reducing it where feasible and having a credible plan for when access is disrupted.

A hybrid approach was endorsed as the most practicable path: using commercial models where appropriate, while selectively applying stronger sovereignty requirements to sensitive data and high-impact decisions. Participants discussed digital nationalism and agreed that such an approach has drawbacks, since isolated national models are likely to reduce quality and increase cost without delivering genuine control. The goal should be pragmatic autonomy, not ideological purity.

Several participants described a qualitative shift in geopolitical relations, noting that this was now colouring how European institutions in particular think about their dependency on external technology providers.

An illustration of model sovereignty risk came from one parliament's experience: a provider had switched off its model two weeks before the meeting, forcing the institution to find, test and validate a replacement at short notice. While this was primarily a supplier reliability issue, the lesson that dependency is easy to underestimate when adoption decisions are made resonated widely with the parliaments present.

Agentic AI represents the next governance frontier

Agentic AI systems were seen as the next governance frontier for parliaments. Unlike conventional AI tools, which respond to prompts, agents are capable of acting autonomously, executing multi-step tasks and making decisions without human intervention at each stage. The session addressed both conceptual governance challenges and concrete emerging implementations, with real-world agentic AI use cases in parliamentary contexts shared by participating institutions. Those examples illustrated both the potential of agentic systems and the governance requirements they demand. The core challenge is qualitatively different from that of conventional AI: agentic systems create

accountability and control questions that standard governance frameworks are not designed to handle.

Three specific risk types were identified: prompt injection, whereby malicious instructions are embedded in the content the agent processes; tool abuse, whereby agents use capabilities beyond their intended scope; and data leakage, whereby information accessed during a task is inadvertently exposed.

Three principles were proposed for managing these risks: human-in-the-loop control at critical decision points as a structural feature; objective oversight mechanisms that are independent of the agent itself; and isolation of AI operations to prevent unintended cascade. On the first of these, a clear distinction was drawn between stating a commitment in policy and actually embedding human oversight in working processes. In other words, key questions – who will supervise the agent, how that will happen, and at what points in the workflow – must be answered before deployment, not left as a policy aspiration.

An important structural point was noted: introducing AI agents into existing processes without first redesigning those processes is likely to amplify existing weaknesses rather than resolve them. Process redesign before AI integration is not a preliminary step but a governance prerequisite.

The agentic AI discussion also surfaced an important shift in what procurement means in this context. With conventional AI tools, institutions tend to specify each solution in detail. With agentic systems, the more effective approach is to define goals, boundaries and security requirements and allow agents to operate within those parameters, rather than trying to anticipate and specify every individual action in advance. This has significant implications for how institutions write contracts, set accountability frameworks and evaluate vendors.

Disinformation and content integrity pose specific risks to parliaments

AI has materially lowered the cost and raised the quality of synthetic and manipulated media, making content authentication an increasingly urgent concern. Parliamentary content is widely shared on social media and is a particular target for selective editing and fabrication.

The Coalition for Content Provenance and Authenticity (C2PA) standard was presented as a practical mechanism for creating a cryptographic chain of provenance for media content, with at least one parliament planning a pilot to explore applying the standard to parliamentary videos. Participants were, however, candid about the limitations of this approach: digital signatures should be positioned as a tool for helping audiences make more informed decisions, not as a “truth engine”. Consistent adoption and sustained public awareness – analogous to the campaign that established HTTPS as a trust signal on the web – will be needed to make the standard meaningful.

Key takeaways and recommendations

The recommendations below are relevant to parliaments at all stages of the AI journey. They are ordered roughly from foundational to advanced, but all apply regardless of maturity level.

Treat AI as a governed institutional capability, not a collection of tools

Effective parliaments are moving away from ad hoc adoption and towards treating AI as a managed institutional function. This requires cross-functional governance structures and coherent oversight, with authority delegated from senior leadership.

Build data governance before scaling AI

Data governance is the prerequisite for reliable AI – and the two should be designed as an integrated system rather than as separate workstreams. Parliaments should establish clear data ownership and data stewardship roles, and aim for unified data, records and knowledge management systems before scaling AI tools. Limited, well-governed data sets should be treated as the starting point, as should a case-by-case approach to proving value. This is more likely to succeed than an attempt at comprehensive deployment.

Invest early in people, process and culture

People, process and culture – i.e. not just technology – are significant factors in successful AI adoption. Parliaments need to invest in AI literacy training at all levels, create dedicated forums and peer networks for business users, engage with MPs directly, co-develop guidance with party leadership and address replacement fears openly. Early failures should be treated as expected and valuable learning opportunities, not as reasons to halt adoption.

Address shadow AI through supply, not prohibition

Attempts to prohibit shadow AI tend to fail because they do not address the underlying demand: where appropriate and officially sanctioned tools are not available, staff will use whatever is to hand. A more effective response is to provide trusted, governed alternatives, to create guidelines and trusted processes, and to provide AI literacy training and user support – all of which minimizes risk and ensures institutional oversight.

Frame sovereignty as dependency management

Complete technological independence is not achievable for most parliaments. A more helpful question is this: “Do we understand our dependency on external providers, are we reducing it where feasible and do we have a credible plan if access is disrupted?”. Parliaments need to map their technology dependencies by country and level of reliance, develop scenarios for disruption and build workload

portability as a core strategic capability. The goal is pragmatic autonomy, not ideological purity.

Prepare governance for agentic AI now

Agentic AI systems are already available and being trialled in parliaments. Institutions that wish to go down this path must develop governance frameworks before deployment, not after. Key requirements include: process redesign before agent integration; human-in-the-loop controls embedded as structural features rather than policy commitments; objective oversight mechanisms that are independent of the agent; and isolation of AI operations. Shared infrastructure models for agentic AI should be explored, given the efficiency and quality gains they offer.

Use the CIP community as a benchmarking and learning resource

Peer learning within the CIP community was consistently identified as one of the most valuable resources available to participating institutions. Parliaments at all levels of maturity should actively use the community and its resources – such as the [Maturity Framework for AI in Parliaments](#), the [Guidelines for AI in parliaments](#) and the [Use cases for AI in parliaments](#) – all of which and more can be found on the IPU's [AI resource page](#).

Looking ahead

The meeting identified several themes for future work, including the following:

1. Further development of the AI literacy and training agenda, including how to build parliamentary capacity incrementally across institutions with different starting points
2. Continued exploration of dedicated AI roles and competency frameworks, including data stewardship roles, building on “competency atlas” approaches that map what capabilities AI work requires and where they should live in the organization
3. The sharing of approaches to AI sovereignty and dependency management, including collaborative development models that allow institutions to share and reuse solutions rather than each building its own
4. The development and sharing of practical guidance on agentic AI governance, including process redesign requirements and risk management frameworks, as well as the sharing of parliamentary use cases that are already in development or deployment

About the CIP and the PDS Hub

The [CIP](#) harnesses the potential for innovation across parliaments by bringing them together, supporting collaboration on solutions and sharing knowledge with the wider parliamentary community.

The Parliamentary Data Science Hub is a community of parliamentary practitioners convened by the CIP. It brings together a group of advanced parliaments to share experience, develop good practice and build collective capacity around digital transformation and AI.

For more information, contact innovation@ipu.org.