ITIL®

ITIL® AI Governance

Global Best Practice

AI NATIVE

For all organisations and people aiming to use AI confidently and responsibly!

PeopleCert

White Paper

ITIL®
AI Governance

Global Best Practice

For all organisations and people aiming to use AI confidently and responsibly!

Published by PeopleCert International Limited
Publication printed in Greece or reproduced electronically in Greece

Version 1.0 (November 2025)

V1.0 | November 2025

CONTENTS

List of figures

List of tables

Authoring Team

James Finister

James Finister has been an ITSM Thought Leader for over 30 years. With a background in outsourcing, operating models, corporate governance, statistics and IT security, he has a wide-ranging experience of real-world issues across geographies and sectors. He was instrumental in the adoption of the earliest version of ITIL and remains an ITIL author. He also contributes to many international standards on AI and ITSM.

Dmitry Isaychenko

Portfolio Director at PeopleCert, responsible for partnership programmes. An experienced IT consultant, solution architect, ITIL 4 co-author, and trainer with over 25 years in the IT industry, spanning network administration, systems engineering, software development, and IT service management. Passionate about measurable, data-driven approaches to IT management. Co-author of Metrics-Based Service Management.

Kaimar Karu

Kaimar bridges ITSM, DevOps, and Agile with over 20 years of international experience, from hands-on IT operations and coding at startups to leading Estonia's digital transformation. At AXELOS (acquired by PeopleCert), he led the initiative to reshape ITIL, helping it embrace DevOps principles. Kaimar has co-authored key ITIL publications, including *ITIL Practitioner* and *Sustainability in Digital and IT*, and contributed to several others. His pragmatic, forward-thinking advice helps organizations embrace agility while navigating complexity. A frequent keynote speaker on service and product management, digital transformation, enabling governance, Cynefin, and AI ethics, Kaimar advocates that frameworks are useful scaffolding — never the goal itself.

Stephen Mann

Stephen is Principal Analyst and Content Director at the ITSM-focused industry analyst firm ITSM.tools, an independent IT and IT service management marketing content creator, and a frequent blogger, writer, and presenter on the challenges and opportunities for ITSM professionals. He has previously held positions in IT research and analysis (at IT industry analyst firms Ovum and Forrester and the UK Post Office), ITSM consultancy, enterprise IT service desk and ITSM, IT asset management, innovation and creativity facilitation, project management, finance consultancy, internal audit, and product marketing for a SaaS ITSM technology vendor.

Simone Jo Moore

Simone is a catalyst for human-led transformation in the age of AI fusing emotional intelligence with digital strategy. Shaping ethical, people-centred approaches through Service Management, HR, and Organizational Change. As Editorial Director of The Youth Rise in Power and The Era of HumanisingIT docuseries she amplifies emerging voices and future-focused narratives. A Top 25 Thought Leader and industry award winner, Simone is a global consultant, master trainer, and speaker. Her work—including contributions to ITIL 4: High Velocity IT, VeriSM, SIAM, DevOps, and ADapT— blends behavioural science with business acumen. Whether shaping policy or igniting mindset shifts, Simone brings rigour, heart, and a touch of poetic rebellion to every conversation.

Roman Zhuravlev

Roman has been responsible for continual development of ITIL content since joining the ITIL team in 2016. He is ITIL 4 Master, ITIL v3 Expert, and ITIL v2 Service Manager with over 25 years' experience in digital product and service management. Roman is a Senior ITIL Architect.

Executive Summary

Artificial Intelligence (AI) is transforming the way organizations operate, make decisions, and deliver value. Yet, without effective governance, its adoption can create significant risks, from biased decisions and data misuse to compliance failures and reputational damage. This paper provides a practical playbook for embedding **effective AI governance** into corporate structures, enabling organizations to capture the benefits of AI responsibly and sustainably.

The core message is that **traditional IT governance is insufficient** for AI. AI introduces unique challenges: autonomy in decision-making, lack of transparency, evolving ethical dilemmas, and fast-changing regulatory requirements. To address these, organizations must extend governance across four perspectives: **Decision Authority & Risk Management, Ethical Principles & Responsible AI, Data Governance & Performance Management, and Regulatory Compliance & Operational Standards**.

The paper applies the **PeopleCert 6C model** to classify AI capabilities (Creation, Curation, Clarification, Cognition, Communication, Coordination) and map them against AI-specific risks. It outlines a structured, four-step approach for adapting governance: stress-testing current frameworks, defining AI-specific requirements, designing governance adjustments, and operating governance dynamically.

Finally, the paper stresses that governance should evolve into **stewardship**. This means moving beyond control to care: safeguarding human dignity, embedding ethics, and creating trust. Effective AI governance is not a barrier to innovation but the enabler of safe, ethical, and value-driven adoption of AI. Organizations that embrace this mindset will be best positioned to leverage AI's transformative potential while maintaining compliance, resilience, and stakeholder confidence.

1 Why should we care?

The subject of 'governance' or 'IT governance' might not usually be on your IT Service Management (ITSM) radar, but it should be, because the corporate use of artificial intelligence (AI)-based capabilities 'blows traditional IT governance out of the water'. This might sound dramatic, but the risks of ungoverned adoption of AI must not be underestimated.

This guidance is designed to help your organization establish effective AI governance. Later sections describe how to do this. However, this introduction first explains the need and the context.

1.1 Recognizing the corporate need for AI governance

The opportunities for corporate AI use are great, both at the individual and corporate levels, but so are the risks, including legislative and regulatory compliance risks. However, the corporate need for AI governance is not solely compliance-based, with the drivers for AI governance ranging from externally imposed to internally motivated ones. Importantly, AI governance, when done well, is not about control; it is about care. It is the scaffolding that holds up trust, the quiet pulse behind ethical transformation, and the choreography that lets humans and machines move in sync, not just in service of efficiency, but of dignity.

It is unlikely that the requirement for effective AI governance is unknown in your organization. AI governance was a hot topic in ITSM in 2024 and remains so now. The annual ITSM.tools content poll for 2025[1] found that 'governance (including AI governance)' was the highest priority 'learning' area for ITSM professionals. However, there is often a gap between ambition and action. Hence, AI governance within your organization may not be as effective as it could or should be. Recent research revealed that:

- While 90% of organizations utilize AI in daily operations, only 18% have a fully implemented AI governance framework. (Legalfly)[2]
- 82% of organizations already use AI agents, but only 44% have policies in place to secure them. (Sailpoint)[3]

These surveys highlight the need for AI governance. However, numerous real-world examples of where AI-based capabilities have adversely affected corporate operations, compliance, and reputation are likely more valuable in convincing ITSM and business professionals of the need for effective AI governance.

1.2 Do not assume your IT governance has AI governance covered

Consider these two questions:

- How does corporate AI adoption affect your IT organization in terms of risks?
- Or, more importantly, how does use of AI (corporate and non-corporate) affect your wider organization's risk profile?

Your response might be, "Oh, our corporate governance team and practices have this covered, along with the sterling work they do with traditional IT governance". However, do you (and they)

[1] https://itsm.tools/itsm-trends-for-2025/
[2] https://www.legalfly.com/report-overviews/ai-governance-gap-key-findings
[3] https://www.sailpoint.com/press-releases/sailpoint-ai-agent-adoption-report

truly understand how introducing AI affects your organization's risk and compliance posture? Additionally, can you collectively manage the additional risks AI adoption presents?

Importantly, these questions relate to both corporate AI capabilities and what could be called 'Shadow AI', where individuals and business functions independently use free or paid-for AI capabilities without the involvement of your IT organization. After all, no matter how great your corporate AI capabilities are, your organization's employees will still likely use free AI tools. Ultimately, AI influences corporate decision-making, regardless of the corporate stance on its use.

It is important to recognize that simply adding AI-focused "extensions" to your existing corporate IT governance practices is likely insufficient. Instead, work might be needed to fix the foundations of the existing governance practices before building the required additional AI governance capabilities.

This need can relate to AI and IT governance. For example:

- AI knowledge might be lacking, that is how AI models work and where the common risks arise.
- Understanding of AI use cases might be limited, that is how different AI types and models can be applied to different ITSM opportunities and issues.
- Existing IT governance practices could have tensions, weak spots, or omissions.
- The current approach to governance of technology may be too slow and rigid to keep up with the speed of AI.

This AI governance guidance offers practical tools to assist with this need to fix the foundations, including:

- governance pattern assessment and modelling, to assess and adjust the corporate governance of technology
- the PeopleCert 6C Model for AI, to assess the use of AI and identify risks and relevant controls
- an assessment and improvement approach, to make the corporate governance of technology AI-ready and AI-relevant.

2 Where are we now? Corporate governance of technology

2.1 Introduction

Every organization requires four enabling capabilities: the core building blocks, which must be well-designed and well-executed for successful business operations and at the times of transformations. While closely intertwined and dependent on each other, the development of each of these capabilities should be done consciously.

Leadership sets the vision and fosters the organization's culture, inspiring people toward agreed goals. Leaders maintain organizational direction, resolve conflicts, and sustain motivation across teams. They champion change, communicate the vision for the future state, and guide the organization through uncertainty.

Governance sets rules and framework for decision-making and accountability. It ensures consistent policy application, risk management, and regulatory compliance across operations. Governance establishes change controls, approval and escalation mechanisms, maintains oversight of the changes, and ensures transformation activities align with organizational objectives.

Strategy determines long-term direction and competitive positioning for the organization. It guides resource allocation, market positioning, and capability development to maintain competitive advantage. It identifies new opportunities, defines target operating models, and reshapes the organization's approach to markets and operations.

Management handles day-to-day execution and operational delivery. Managers coordinate teams, optimize processes, and deliver products and services to meet established targets and customer expectations.

The approach taken to any of these four must consider the organization's surrounding environment and the specifics of how the organization is expected to operate.

Adoption of AI is likely to impact governance and management of organizations. It may also require a transformation of the organization at many levels, in which case a relevant approach to governance and management of transformation should be identified and adopted.

This section focuses on understanding existing governance patterns and the resulting constraints that need to be considered for the AI adoption. The following sections explore the impact of AI on governance of technology and offer guidance on adjusting the organization's governance system to address this impact. Equipped with this information, organizations can address the challenges described in the Introduction and in Section 3.

2.2 Governance patterns

Organizations' approach to governance vary shaped by the organizational values, culture, business context, and other factors. Organizations typically adopt one of four dominant governance patterns along the axes of authority (centralized ↔ distributed) and assurance (structured ↔ emergent), illustrated in Figure 2.1 and described through the lens of governance activities in Table 2.1.

- **Directive:** top-down control through formal hierarchies, strict procedures, and standardized processes.

- **Guided:** central vision and strategic direction with local freedom in execution and implementation.
- **Federated:** multiple distributed units with delegated authority coordinating through formal structures.
- **Autonomous:** self-organizing teams making decisions through peer collaboration and influence.

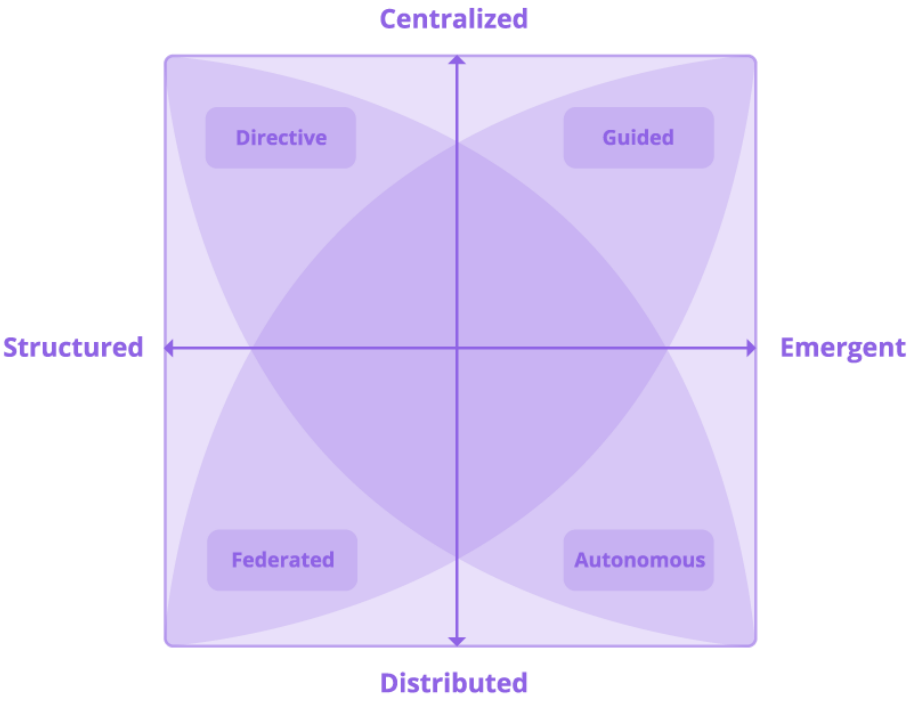

Figure 2.1 Governance patterns defined by authority and assurance

To help with identifying the closest-matching governance pattern, the frequently observable characteristics of the four governance activities as described in Table 2.1 can be used. The activities used for this comparison are based on the *ISO/IEC 38500:2024 Information technology – Governance of IT for the organization* standard:

- Engage stakeholders
- Evaluate (the use of IT in the organization)
- Direct (the use of IT in the organization)
- Monitor (the use of IT in the organization)

Please keep in mind that these descriptions represent the extreme version (that is, what would be the case for an organization in the far corner of the quadrant) of the patterns and can manifest themselves in more balanced approaches within the quadrant in the context of the specific organization. It is highly unlikely for an organization to manifest pure extreme characteristics across all four governance activities; every organization is somewhere between the four extreme corners.

Table 2.1 Characteristics of governance patterns

	Directive	Guided	Federated	Autonomous
Engage stakeholders	Top-down information flow; stakeholders receive decisions, not consulted.	Leadership storytelling and vision-casting; idea harvesting from below.	Protocol-based engagement per unit; inter-unit coordination forums.	No designated engager; anyone can convene anyone.
Evaluate	Performance against predetermined standards; deviation requires justification.	Success redefined as context evolves; leading indicators valued over lagging.	Unit-specific interpretation of enterprise metrics; systematic comparison across units.	No imposed metrics; groups define what matters to them.
Direct	Permission required for decisions; escalation for any deviation.	Direction through 'commander's intent'; how is flexible if why is clear.	Documented decision rights matrix; structured escalation paths between units.	No formal directing; influence through expertise not position.
Monitor	Compliance verification; variance reports trigger intervention.	Pattern sensing over rule-checking; insights valued over compliance.	Standardized exception handling; portfolio-level integration reviews.	No formal monitoring; peer visibility creates natural accountability.
Typical signals of ineffectiveness	Initiative paralysis without approval; blind spots from filtered information.	Different interpretations of intent cause divergence; vision without execution.	Units optimize locally, suboptimize globally; coordination consumes excessive resources.	Lack of coherent direction; critical dependencies missed; uneven capability development.

To assess the current governance pattern and identify the pattern most suitable for the current and anticipated context of the organization, the following ten characteristics of governance can be used:

1. Approach to change
2. Success metrics
3. Integration requirements
4. Risk tolerance
5. Environmental stability
6. Decision velocity
7. Compliance requirements
8. Stakeholder diversity
9. Governance scope
10. Governance capabilities.

The final position may fall clearly within one quadrant, indicating consistent governance approach, or in boundary zones between quadrants, suggesting hybrid governance pattern.

Boundary positions are not weaknesses; they may indicate sophisticated balancing of competing needs. Both the average position and the spread of characteristics should be documented, as high variance suggests either organizational flexibility or potential governance conflicts requiring attention.

For example (see figure 2.2): six out of ten governance characteristics have been placed in the Directive quadrant, and the other four in the Guided quadrant. This positions the organization's governance pattern close to the border between Directive and Guided, and rather high on the centralized/distributed axis.

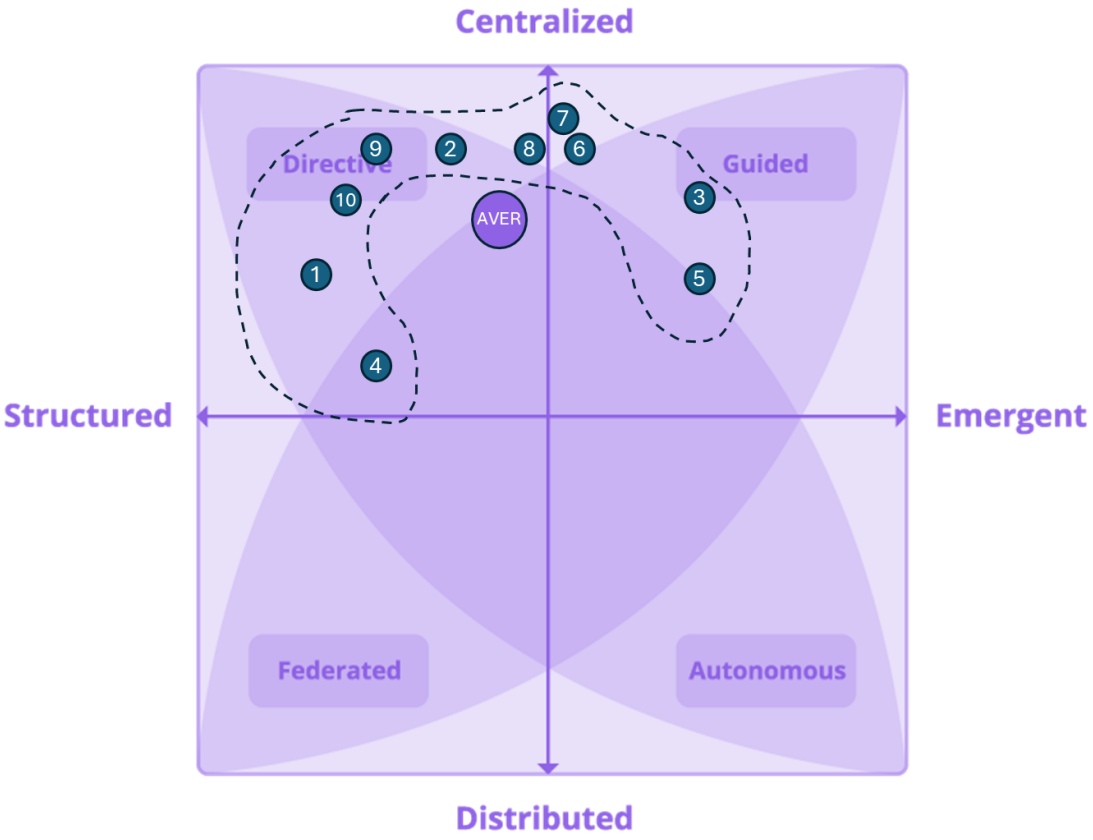

Figure 2.2 Example of the governance characteristic on the 2x2 governance matrix

The clusters and outliers should be analysed carefully: for example, if seven characteristics cluster in 'directive' but three fall in 'autonomous', this tension reveals important dynamics. Consideration should be made whether outliers represent organizational strengths to leverage or conflicts creating friction.

By understanding the current governance pattern, organizations can assess the potential impact of AI adoption and the AI readiness of the governance system. The current governance pattern serves as a baseline for any adjustments required to address the challenges brought by AI and ensure sustainable and effective governance of technology in the organization.

---

3 What is the challenge?

The rapid adoption of Artificial Intelligence across organizations introduces fundamental tensions into existing governance structures. Unlike traditional technology implementations that typically enhance or automate existing processes, AI systems introduce potentially autonomous decision-making capabilities that challenge established patterns of authority, accountability, and control. These systems learn, adapt, and make decisions at speeds that often exceed human oversight capacity, whilst simultaneously requiring careful stewardship to ensure they operate within acceptable boundaries.

The governance implications extend far beyond technical considerations. AI systems can perpetuate or amplify bias, make decisions that affect individuals' lives and livelihoods, and operate in ways that may not be fully explainable even to their creators. They consume vast amounts of data, including personal and commercially sensitive information, whilst generating insights and decisions that can reshape business operations and customer experiences. The regulatory landscape surrounding AI continues to evolve rapidly, with new requirements emerging across jurisdictions that organizations must navigate whilst maintaining operational effectiveness.

Traditional directive governance frameworks, designed for predictable and controllable systems, often prove inadequate when applied to AI solutions. The challenge lies not in replacing existing governance structures, but in understanding how AI specifically impacts each governance perspective and adapting accordingly. This requires a systematic approach that examines AI's unique characteristics (its autonomy, opacity, learning capability, and broad applicability) against established governance practices.

The four AI impact perspectives presented here provide a structured lens through which organizations can analyse how Artificial Intelligence affects their technology governance. Rather than creating entirely new governance frameworks, these perspectives enable systematic assessment of AI's impact on existing governance characteristics, facilitating the design of appropriate oversight mechanisms that balance innovation with responsibility. This approach proves particularly valuable when organizations apply the same diagnostic methodology used for other technology transformations, ensuring consistency with established governance assessment practices whilst addressing AI's distinctive challenges.

The perspectives have been designed to cover the critical governance considerations without overwhelming complexity. Each perspective captures related governance challenges that naturally cluster together, enabling focused analysis whilst ensuring comprehensive coverage. This structure supports both initial assessment and ongoing governance adaptation as AI implementations evolve and mature.

Note that these four perspectives are never isolated; they are intertwined and interdependent in the real-world governance landscape. Same applies to the risk described in each of the following sections: although primarily manifested in one of the dimensions, they may originate from and affect the others. Organizations should use this description of the AI challenges and risks as a starting point for their own analysis, adapting this guidance to the evolving business and technology context.

3.1 Decision authority & risk management

Also known as 'structural governance', this perspective encompasses the fundamental questions of who holds decision-making power in AI-enabled environments and how risks are identified, assessed, and controlled through governance mechanisms. The combination reflects the reality that in AI-enabled management systems, decision authority and risk management are inextricably linked. Those with the authority to deploy, configure, or modify AI systems must also bear responsibility for managing the risks these systems create.

AI challenges traditional decision-making hierarchies by operating at speeds and scales that often preclude human intervention. An AI system processing thousands of loan applications per hour or managing real-time traffic flows cannot pause for human approval on each decision. This reality necessitates a shift from approval-based governance to guardrails-based governance, where humans establish boundaries and principles within which AI systems operate autonomously.

The risk landscape for AI differs substantially from traditional technology risks. Model drift can cause AI solutions to gradually degrade in performance without obvious warning signs. Adversarial attacks can manipulate AI systems in ways that bypass conventional security measures. Training data can become outdated or biased, leading to decisions that seemed reasonable during development but prove problematic in production. These risks require new forms of monitoring, assessment, and response that must be embedded within existing risk management frameworks.

Accountability structures must evolve to address the distributed nature of AI decision-making. When an AI system makes an incorrect decision, establishing accountability requires tracing through model training, data preparation, algorithm selection, deployment decisions, and ongoing monitoring, often involving multiple teams and timeframes. Governance structures must clearly define roles and responsibilities across this extended chain whilst ensuring that accountability gaps do not emerge between technical and business domains.

The integration of AI risk management with enterprise risk management requires careful consideration of risk appetite, tolerance levels, and escalation procedures. Traditional risk assessment methods may prove inadequate for evaluating AI risks that can emerge from complex interactions between data, algorithms, and operating environments. New risk management capabilities, including AI-specific monitoring tools and risk assessment methodologies, often need to be developed alongside governance structure adaptations.

The key risks primarily related to this perspective are:

Autonomy in decision-making: AI introduces a new level of autonomy where decisions can be made without direct human intervention. Risks include loss of human oversight, unclear accountability for AI-driven actions, and potential automation of flawed or unethical decisions.

Relevant countermeasures for these risks include:

- defining clear accountability and human oversight roles for AI decisions
- establishing policies for human-in-the-loop or human-on-the-loop controls
- setting predefined conditions or thresholds for when human intervention is mandatory (high level of risk, low confidence score, ethical or sensitive decisions, regulatory requirements, etc.).

Transparency and explainability: many AI models, especially deep learning models, operate as black boxes with decisions that are hard to interpret or explain. This reduces auditability,

stakeholder trust, and regulatory compliance, especially where explanations are legally or ethically required.

Relevant countermeasures for these risks include:

- using inherently interpretable models, to reduce complexity and increase transparency
- ensuring that AI models are designed and implemented to support post-hoc explainability techniques when complex models are used
- defining explainability expectations based on the system's risk profile and requiring model cards to detail purpose, limitations, and interpretability approach.

3.2 Ethical principles & responsible AI

Sometimes called 'relational governance', this perspective addresses the moral and societal implications of AI systems, encompassing both foundational ethical principles and their practical implementation in AI development and deployment. The pairing recognizes that ethical considerations cannot remain abstract philosophical concepts but must be translated into concrete practices, policies, and oversight mechanisms.

Ethical principles provide the foundation for responsible AI governance, drawing from established frameworks including beneficence (doing good), non-maleficence (avoiding harm), justice (fairness and non-discrimination), autonomy (respecting human agency), and transparency (explainability and openness). These principles, well-established in fields such as medical ethics and research, require careful adaptation to the AI context where decisions may affect thousands of individuals simultaneously and where the decision-making process may be opaque even to the system's creators.

The challenge lies in operationalizing these principles within practical governance structures. Fairness, for instance, must be defined in measurable terms that can be assessed during development and monitored in production. Different stakeholders may have different perspectives on what constitutes fair treatment, requiring governance processes that can balance competing interests whilst maintaining ethical standards. Transparency presents similar challenges. Whilst complete explainability may be technically impossible for complex AI systems, governance must ensure sufficient transparency to enable appropriate oversight and accountability.

Responsible AI implementation requires embedding ethical considerations throughout the AI lifecycle, from initial conception through deployment and ongoing operation. This includes processes for ethical review during development, mechanisms for detecting and correcting bias, procedures for handling ethical concerns raised during operation, and frameworks for stakeholder engagement on ethical issues. The governance challenge involves ensuring these processes are robust enough to address genuine ethical concerns whilst remaining practical enough to enable innovation and operational effectiveness.

The intersection between ethical principles and business objectives requires careful navigation. Ethical AI is not merely about compliance or risk mitigation. It encompasses fundamental questions about the role of AI in society and the responsibilities of organizations deploying these systems. Governance structures must facilitate meaningful consideration of these broader implications whilst enabling practical decision-making about specific AI implementations.

The key risks related to this perspective are:

Bias and fairness risks: AI can embed and amplify social, racial, gender, or cultural biases present in training data or model design. This results in unfair or discriminatory outcomes, loss of customer trust, legal challenges, and ethical concerns.

Relevant countermeasures for these risks include:

- using diverse and representative training datasets
- involving diverse stakeholders in model review
- performing bias audits on training data and AI models
- establishing remediation plans for identified biases.

Hallucination and mistakes in generated outputs: AI, particularly generative AI, can produce inaccurate, fabricated, or misleading information ("hallucinations"). This can cause operational errors, misinform users, and potentially lead to legal or reputational damage if relied upon without human validation.

Relevant countermeasures for these risks include:

- establishing human review and validation processes for critical AI outputs
- continuously monitoring AI outputs for errors and inconsistencies
- training staff and users to critically assess and verify AI-generated outputs before acting on them, especially in high-risk or decision-making contexts
- developing escalation protocols for detected mistakes.

3.3 Data governance & performance management

Sometimes referred to simply as 'data governance', this perspective combines the stewardship of data resources that fuel AI systems, with the measurement and management of AI system performance and effectiveness. The connection reflects the fundamental dependency of AI performance on data quality, relevance, and appropriateness, whilst recognizing that effective performance management requires sophisticated understanding of data flows and dependencies.

Data governance for AI extends beyond traditional data management to encompass the entire data lifecycle as it relates to AI systems. Training data quality directly impacts AI performance, but assessing this quality requires understanding how data will be used by specific algorithms and in particular operating contexts. Data lineage becomes critical when AI systems exhibit unexpected behaviour and investigation requires tracing back through training datasets, preprocessing steps, and data sources. Privacy and consent management must address not only current data use but also potential future applications as AI capabilities evolve.

The dynamic nature of AI systems creates ongoing data governance challenges. Unlike traditional systems that typically process data in predictable ways, AI systems may identify new patterns or relationships that raise fresh questions about data appropriateness or consent. Model retraining may require new data or different uses of existing data. The governance framework must be sufficiently flexible to address these evolving requirements whilst maintaining appropriate controls and protections.

Performance management for AI systems requires new metrics and monitoring approaches that go beyond traditional system performance indicators. AI system accuracy, bias, drift, and robustness must be continuously monitored, but these metrics only have meaning within specific business contexts and use cases. An AI system that maintains high technical accuracy but fails to

deliver business value, or one that performs well on average but exhibits bias against particular groups, requires governance intervention despite meeting some performance criteria.

The feedback loops between performance monitoring and data governance create complex dependencies that must be managed through integrated governance approaches. Performance issues may indicate data quality problems, whilst data governance decisions can significantly impact system performance. The governance framework must enable coordinated decision-making across these interdependent domains whilst maintaining clear accountability for both data stewardship and performance outcomes.

The key risks related to this perspective are:

Access control misalignment between source data and AI outputs: AI models are trained on data with different access controls but often cannot enforce these controls when generating outputs. This can cause unauthorized disclosure of sensitive information or force the organization to create multiple segregated AI models, increasing costs and risks of insufficient data.

Relevant countermeasures for these risks include:

- using policy-aware or context-aware Retrieval-Augmented Generation (RAG) to constrain the information retrieved and used in AI-generated outputs based on user roles or access rights.
- using model segregation or data tagging to restrict output access
- auditing AI outputs regularly for unauthorized data disclosures
- avoiding application of AI models to highly sensitive data unless effective security enforcement mechanisms are in place.

Data availability and quality for AI training: many organizations, especially SMBs, face insufficient volume or quality of data needed for effective AI training. Poor data leads to overfitting, biased models, or poor generalisation, reducing AI effectiveness and increasing risk of errors.

Relevant countermeasures for these risks include:

- assessing data sufficiency and quality before model training when possible
- implementing data cleansing and augmentation practices
- monitoring model performance for signs of data-related issues.

3.4 Regulatory compliance & operational standards

Also known as ‘procedural governance’, this perspective encompasses both external regulatory requirements that organizations must meet and internal operational standards that organizations establish to ensure consistent, high-quality AI governance across their operations. The combination recognizes that compliance extends beyond meeting minimum regulatory requirements to encompass the operational excellence necessary for sustainable AI deployment.

Regulatory compliance for AI presents unique challenges due to the rapidly evolving nature of both AI technology and the regulatory landscape. The EU AI Act, algorithmic accountability requirements, sector-specific regulations, and evolving privacy laws create a complex compliance environment that organizations must navigate whilst maintaining operational effectiveness. Unlike compliance with established regulations where requirements are well-understood, AI compliance often requires interpretation of new requirements in the context of rapidly evolving technology capabilities.

The global nature of many organizations and AI systems creates additional complexity as different jurisdictions develop different regulatory approaches. An AI system developed in one country, trained on data from multiple countries, and deployed globally may be subject to multiple regulatory frameworks with potentially conflicting requirements. Governance structures must enable navigation of this complexity whilst ensuring that compliance requirements do not inadvertently compromise system effectiveness or create conflicting obligations.

Operational standards provide the internal framework within which AI systems are developed, deployed, and operated. These standards translate regulatory requirements and organizational policies into practical procedures, quality gates, and operational practices. They encompass technical standards for AI development, operational procedures for AI deployment and monitoring, and organizational standards for roles, responsibilities, and decision-making authority.

The development of operational standards for AI requires balancing multiple considerations including regulatory compliance, risk management, operational efficiency, and innovation enablement. Standards must be sufficiently detailed to ensure consistent practice whilst remaining flexible enough to accommodate the diversity of AI applications and the rapid pace of technological change. They must integrate with existing operational standards and procedures whilst addressing the unique characteristics of AI systems.

The relationship between regulatory compliance and operational standards creates opportunities for organizations to exceed minimum compliance requirements by developing operational excellence in AI governance. Organizations that establish robust operational standards often find regulatory compliance becomes a natural outcome of good governance practice rather than a separate compliance burden. This approach enables organizations to anticipate regulatory developments and adapt more readily to changing requirements whilst maintaining operational effectiveness.

The key risks related to this perspective are:

Legal, regulatory, and compliance risks: AI systems must comply with data protection laws (for example, GDPR), intellectual property rights, and emerging AI-specific regulations (for example, EU AI Act). Challenges include demonstrating audit trails, handling liability, and adapting to evolving regulatory landscapes.

Relevant countermeasures for these risks include:

- maintaining up-to-date knowledge of applicable AI and data regulations and standards
- integrating regulatory requirements into the design, procurement, and deployment of AI solutions
- documenting compliance measures and maintaining audit trails for AI system operations.
- training staff regularly on relevant legal and ethical requirements
- establishing robust processes to monitor and respond to changes in AI and data regulations.

Operational and lifecycle risks: AI models may experience drift and degradation over time, reducing accuracy and fairness. Managing the AI lifecycle, including updates, retraining, version control, and retirement, is complex and requires adapted processes.

Relevant countermeasures for these risks include:

- implementing continuous monitoring for model drift and performance degradation
- developing robust model update, retraining, and retirement procedures
- establishing AI-specific incident response plans

- maintaining version control and documentation for AI models.

Extended supply chain and third-party AI governance risks: AI governance risks extend beyond organizational boundaries when sharing sensitive data with suppliers or AI service providers. Lack of transparency or governance over third-party AI usage can lead to data misuse, non-compliance, or unintended exposure.

Relevant countermeasures for these risks include:

- conducting due diligence on AI governance practices of suppliers
- including AI governance requirements in contracts
- monitoring and auditing third-party AI usage regularly
- establishing incident reporting agreements
- defining and enforcing policies specifying data types or sensitivity levels that are prohibited from being processed or stored in third-party or cloud environments.

Organizational readiness and skills: successful AI governance requires AI literacy and cultural readiness among leaders and employees. Lack of understanding or resistance to AI adoption can cause misaligned expectations, governance gaps, and insufficient risk mitigation.

Relevant countermeasures for these risks include:

- providing AI literacy and ethics training
- building organizational awareness, clear accountability, open communication, and leadership commitment to ensure responsible and ethical AI use
- aligning AI initiatives with business strategy
- establishing clear communication channels for AI governance.

Note that although the list of the key AI-related risks used in this guidance is thought to be valid at the time of writing and represents a good starting point for AI risk assessment, every organization should continually identify and assess risks in the context of its unique business and technology context. This guidance offers an effective approach to risk assessment that is supposed to be adopted and adapted to the organization's needs and circumstances.

3.5 What are we using AI for?

Of course, impact of the AI risks vary depending on the organization's context, current governance and management approaches, management capabilities, and other organization-specific factors.

In addition to those, one important factor to be considered is the AI capabilities adopted by the organization. ITIL describes six key AI capabilities found (in different combinations) in AI solutions: Creation, Curation, Clarification, Cognition, Communication, Coordination. Together they form the 6C model providing a functional classification of AI solutions. The model helps understand and communicate the range of possible applications of AI. It can also enhance AI governance by helping organizations tailor risk profiles, controls, and countermeasures to the specific functions of AI solutions, especially beyond it is intended use.

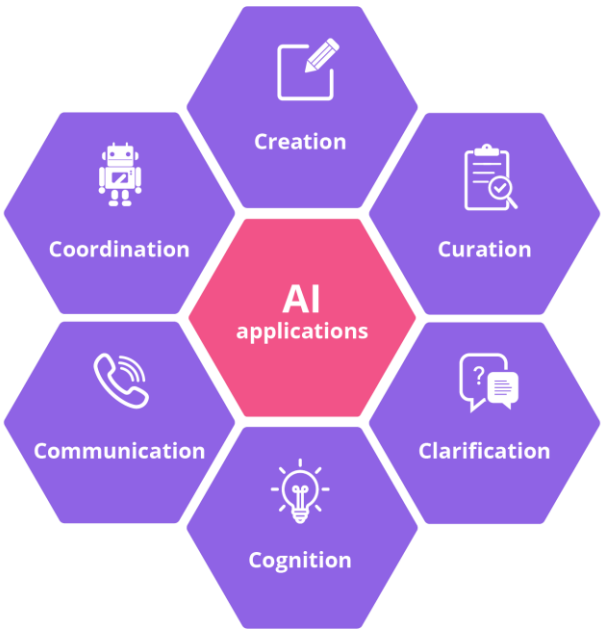

Figure 3.1 The 6C model for functional AI capabilities

These capabilities are not mutually exclusive; most real-world use cases combine them. Each of the six capabilities is described in the Table 3.1 below.

Table 3.1 The 6C model: functional AI capabilities

AI Capability	Description
1. Creation	AI generates net-new outputs in response to prompts or triggers. This includes producing content, code, documentation, or any other artifacts that did not exist before.
2. Curation	AI improves the quality, organization, and relevance of existing data or knowledge by identifying redundancies, outdated information, inconsistencies, or incompliance.
3. Clarification	AI helps users find, understand, navigate, or improve existing content by summarising, rephrasing, restructuring, or translating it.
4. Cognition	AI identifies patterns, anomalies, or hidden insights across data, enabling proactive detection, forecasting, and analysis.
5. Communication	AI acts as a communicative interface helping users interact with services and systems naturally.
6. Coordination	AI autonomously executes, orchestrates, or triggers actions across systems, often in response to events, requests, or patterns.

Typically, the key AI risks vary depending on the AI application as described in table 3.2.

Table 3.2 AI Risks level typical for the 6C AI capabilities

Governance perspectives	AI risks	Creation	Curation	Clarification	Cognition	Communication	Coordination
Decision Authority & Risk Management	**Autonomy in decision-making**	Low	Low	Low	Medium	Medium	High
	Transparency and explainability	Medium	Medium	Medium	High	Medium	High
Ethical Principles & Responsible AI	**Bias and fairness risks**	High	Medium	High	High	High	Medium
	Hallucination and mistakes in generated outputs	High	Low	Medium	Medium	High	Medium
Data Governance & Performance Management	**Access control misalignment between source data and AI outputs**	High	Medium	Medium	Medium	High	Low
	Data availability and quality for AI training	High	High	Medium	High	Medium	High
Regulatory compliance & Operational standards	**Legal, regulatory, and compliance risks**	High	Medium	Medium	High	High	High
	Operational and lifecycle risks	Medium	Medium	Medium	High	Medium	High
	Extended supply chain and third-party AI governance risks	Medium	Medium	Medium	Medium	Medium	Medium
	Organizational readiness and skills	Medium	Low	Low	Medium	Medium	Medium

The typical risk levels indicated in table 3.2 represent a starting point for risk assessment conducted by organizations in their business and technology context. These levels may vary from solution to solution and from organization to organization and evolve with time. Like other guidance in this publication, this table is supposed to be adapted by organizations to their evolving context.

Understanding the current pattern of the corporate governance of technology and of the current and anticipated functional capabilities of the AI solutions adopted by the organization is very important. However, this is not sufficient: it helps to know 'where we are' and to some extent 'where we want to be'. The following sections will help define effective AI governance and navigate the journey towards it.

4 Where do we want to be? Defining effective AI governance

A governance assessment (described in Section 2) positions organization's governance system along the axes of authority (centralized to distributed) and assurance (structured to emergent), revealing whether an organization operates with Directive, Guided, Federated, or Autonomous governance patterns. This and following sections help to understand what constitutes effective AI governance and how to assess an organization's readiness for adopting AI solutions. The approach assumes that most organizations will not completely restructure their governance simply to accommodate AI. Instead, the focus is on identifying where existing governance can be adapted through bridging mechanisms, where AI solutions must be scoped to fit governance constraints, and where governance evolution becomes necessary for organizational success.

Effective AI governance extends beyond traditional IT oversight by recognizing that Artificial Intelligence systems operate within inherent limitations that must be understood and respected. The foundation of effective governance lies in understanding what questions can reasonably be asked of different AI solutions and ensuring they are deployed only for purposes where they can provide meaningful, reliable outputs.

The stakes of AI governance extend far beyond technical performance metrics. When AI systems analyse data and make or influence decisions, they impact business operations, individual lives, and the broader society. These cascading impacts mean that effective AI governance must consider the full spectrum of consequences when designing oversight mechanisms, ensuring that enthusiasm for AI's capabilities does not overshadow careful consideration of its appropriate application boundaries.

Drawing from the pillars of Ethical AI (Context, Capability, Culture, Conduct, Change, and Consequence) AI governance should:

- recognize the context of AI use, including emotional and social contexts
- match its capability to human needs, not just key performance indicators (KPIs)
- be aligned with the cultural context
- enforce ethical conduct through accountability
- adapt to ongoing change
- recognize and own consequences of AI use, even the unintended ones.

AI governance is not limited to defining rules; it shapes the human-AI relationships, creating safe, intelligent, human-aligned futures.

AI governance is not an optional layer or a separate silo; it is an integral part of your organization. It must be woven into the broader tapestry of technology and business governance. AI is more than just another system. It is an actor in your service ecosystem, making decisions, shaping experience, and sometimes even learning what we unconsciously reward or tolerate.

Effective AI governance demonstrates four essential characteristics that distinguish it from more traditional IT oversight.

- First, governance must be **fit-for-purpose**, meaning the intensity of oversight matches both the AI system's inherent risks and its operational boundaries.
- Second, governance must remain **adaptive, capable of evolving** as AI systems learn and as organizational understanding of their limitations deepens.
- Third, governance must **integrate with existing organizational structures** rather than creating parallel oversight systems that complicate decision-making.

- Finally, governance must **address all four AI impact perspectives** simultaneously, recognizing that decisions about AI authority, ethical implications, data usage, and regulatory compliance cannot be separated without creating governance gaps.

These are, in essence, the characteristics of any efficient governance system, but experience from the industry has demonstrated that the underlying principles have not always been followed well.

5 How do we get there? Adapting governance for AI

The proposed approach for assessing and implementing AI governance recognizes that effective governance creates value through inclusive design rather than merely ensuring compliance. It acknowledges that AI solutions affect diverse stakeholder groups in different ways, requiring assessment processes that capture varied perspectives and types of expertise. This also provides a systematic method for evaluating current governance capabilities against AI requirements while ensuring that human needs remain central to governance design.

The AI governance assessment and implementation approach follows four interconnected steps that build systematically toward inclusive governance design:

- **Step 1:** Using the knowledge of the current governance pattern and the adopted (or expected) AI capabilities, **stress-test existing governance** against extreme AI scenarios, drawing on diverse stakeholder experiences to identify potential breaking points.
- **Step 2: Determine governance requirements** by understanding how different groups interact with and are affected by specific AI solutions.
- **Step 3: Design and implement adjustments** that respect both organizational constraints and stakeholder needs.
- **Step 4: Ensure that governance remains responsive** to evolving stakeholder experiences as AI solutions learn and adapt.

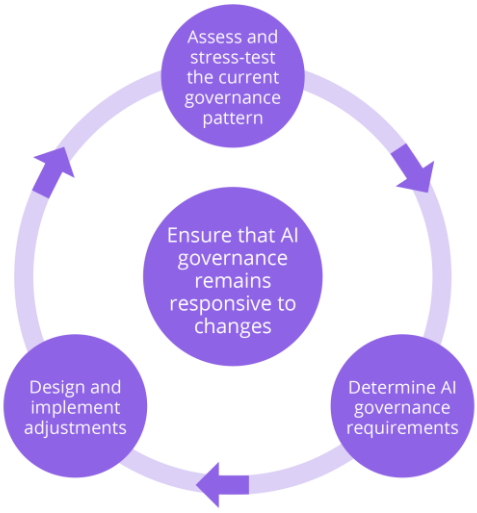

Figure 5.1 AI governance assessment and implementation approach

---

Each step deliberately incorporates multiple perspectives. Stress-testing includes not just technical failure modes but human impact scenarios drawn from frontline experience. Requirements gathering values experiential knowledge alongside technical specifications. Adaptation design involves affected communities in shaping oversight mechanisms. Operational governance creates feedback loops that capture ongoing stakeholder experiences and concerns.

The approach recognizes that effective governance emerges through collaboration rather than top-down design. For example, when warehouse workers help design governance for an AI scheduling system, the resulting oversight addresses practical realities that distant governance

committees might miss. This collaborative approach produces governance that stakeholders understand and support because they helped create it.

5.1 Key principles for AI governance assessment and implementation

Four principles ensure the assessment and implementation process serves both organizational needs and human values.

- First, assessment must **include diverse voices from the start, not as an afterthought**. This means actively seeking input from groups often excluded from governance discussions: frontline workers who will use AI systems daily, customers from varied backgrounds who will experience AI decisions, employees with different cognitive styles who may interact with AI differently, and community members who bear indirect consequences of AI deployment.
- Second, assessment must **value different types of knowledge equally**. Technical expertise matters, but so does experiential understanding, cultural awareness, and emotional intelligence. A neurodiverse team member might identify AI interface issues that others miss. A customer service veteran understands conversation patterns that inform Communication AI governance. These insights deserve equal weight with technical risk assessments and compliance requirements.
- Third, assessment must **reframe risk to include human and social impacts** alongside technical and financial concerns. Traditional governance often defines risk in terms of system failures, data breaches, or regulatory penalties. AI governance must additionally consider risks to human dignity, community trust, and social cohesion. When a Cognition AI analyzes neighborhood patterns for city planning, the risk assessment must include potential impacts on community character and resident wellbeing, not just technical accuracy metrics.
- Fourth, AI governance should be **designed for stewardship, rather than control and compliance only**. Technology professionals are already deeply familiar with enabling outcomes, supporting humans, and balancing optimization with stability. Stewardship simply expands this lens to include the broader emotional, ethical, and societal impacts of AI technologies.

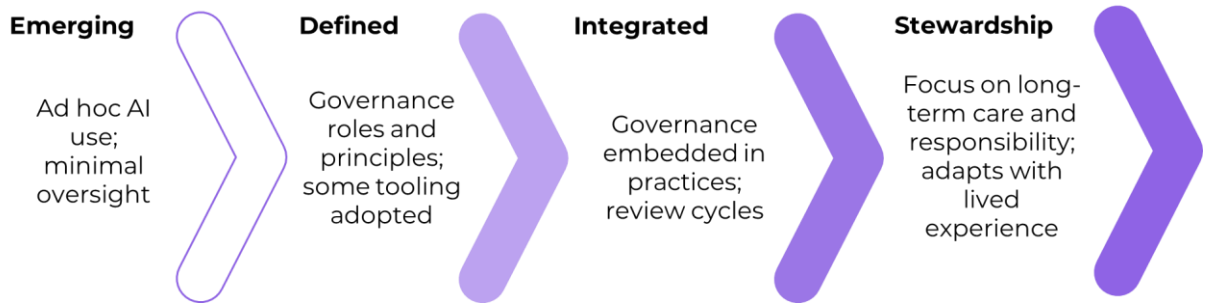

Figure 5.2 AI governance maturity indicators

Although these principles should be considered when building any kind of human-impacting digital services and products, they are often ignored due to lack of understanding or lack of interest. When it comes to AI solutions, ignoring these aspects will bring much bigger risks than before.

5.2 Implementation pathways and sustained effectiveness

The journey from initial AI governance assessment to mature operational oversight requires practical pathways that respect organizational realities.

5.2.1 Pattern-specific adaptation strategies

Each governance pattern requires different adaptation strategies to accommodate AI while maintaining core organizational characteristics.

Directive governance adaptations focus on maintaining central control while enabling the operational speed that AI requires. This is achieved through pre-approval frameworks that clearly define the boundaries within which AI can operate autonomously. Automated reporting systems provide real-time visibility to central authorities without creating bottlenecks in AI operations. Organizations can implement staged delegation, gradually expanding AI decision-making authority as confidence builds through demonstrated success. For example, a government agency might begin with AI systems that only make recommendations requiring human approval, then progress to AI making decisions with human audit rights, eventually reaching autonomous operation within carefully defined parameters.

Guided governance adaptations establish clear principles that enable flexible AI implementation across the organization. Strategic boundaries define what outcomes AI should achieve without prescribing exact methods, allowing teams to innovate within acceptable limits. Monitoring focuses on outcomes rather than process compliance, measuring whether AI delivers intended value while respecting core principles. Organizations can create designated experimentation zones where teams can explore AI capabilities within established safety parameters. A retail organization, for instance, might establish customer satisfaction and fairness principles at the corporate level, then allow individual stores to adapt their Communication AI implementations to serve their specific community needs and cultural contexts.

Federated governance adaptations create coordination mechanisms that respect unit autonomy while ensuring coherent AI deployment. Cross-unit AI councils bring together representatives to share learnings and align approaches without imposing uniform solutions. Common AI platforms provide technical consistency and shared capabilities while still allowing functional customization for unit-specific needs. Rotating leadership of AI initiatives prevents any single unit from dominating governance decisions and ensures diverse perspectives shape AI evolution. A multinational manufacturing company might deploy Cognition AI for quality control across all factories, with each site adapting the system to detect defects specific to their product lines and local manufacturing conditions, while all sites share insights about effective detection patterns and contribute to improving the core AI models that benefit the entire network.

Autonomous governance adaptations rely on peer accountability and collective wisdom rather than hierarchical control. Communities of practice emerge where practitioners share experiences and develop collective norms for responsible AI use. Transparent AI operations enable peer review and constructive feedback, creating natural quality control through professional accountability. Reputation systems recognize and reward responsible AI innovation, encouraging positive behaviours without formal oversight. A network of independent journalists might use Creation AI for investigative research assistance, with each journalist maintaining editorial independence while the community develops shared standards for verifying AI-generated leads, peer-reviews controversial uses, and collectively maintains a repository of effective prompts and fact-checking procedures that all members can access and improve.

5.2.2 Building the evolution case

When existing governance proves incompatible with essential AI capabilities, compelling cases for evolution must be built. Missed opportunities should be documented by tracking where beneficial AI deployment was prevented by constraints, with impacts quantified in terms meaningful to stakeholders. Examples should be gathered of peer organizations achieving benefits that current governance structures prevent. The accumulating costs of maintaining incompatible governance should be calculated and presented.

Evolution should be framed as enhancement rather than replacement. Documentation should demonstrate how core values are maintained while new capabilities are enabled through adjusted governance. Pilot programs should be proposed to test new approaches in contained environments. Coalitions should be formed among stakeholders who have experienced governance constraints firsthand. Most importantly, governance evolution should be connected directly to organizational mission and core objectives, showing how enhanced governance enables better achievement of fundamental goals.

Conclusion

Artificial Intelligence is no longer a peripheral technology; it is an active participant in decision-making, service delivery, and organizational transformation. This playbook has shown that while the opportunities of AI are immense, so too are the risks – ranging from ethical and compliance challenges to systemic failures that can undermine trust, reputation, and value creation.

Effective AI governance must therefore be understood not as an optional add-on, but as a core capability of corporate governance. It requires more than technical "guardrails" or isolated compliance exercises. Governance must extend across four interdependent perspectives **(decision authority and risk management, ethical principles, data governance, and regulatory compliance)** while also being responsive to context, culture, and human impact.

The **ITIL 6C model** provides organizations with a practical lens for understanding functional AI capabilities, mapping their associated risks, and designing appropriate countermeasures. By stress-testing current governance, defining new requirements, and embedding stewardship principles, organizations can adapt governance frameworks to keep pace with AI's speed, scale, and complexity. Crucially, this evolution must be **dynamic**, recognizing that AI systems learn and change, and so governance must also operate in cycles of continual monitoring, feedback, and improvement.

The journey from governance to stewardship highlights a fundamental shift: from controlling technology to **caring for its human, organizational, and societal consequences**. AI governance is not simply about preventing harm; it is about enabling responsible innovation, creating trust, and aligning machine intelligence with human values.

Success will depend on organizational readiness, cultural maturity, and inclusive participation. Leaders, practitioners, regulators, and communities must all have a voice in shaping how AI is used and governed. When approached in this way, AI governance becomes not a barrier to innovation but the scaffolding that sustains it; a framework that allows organizations to harness AI's transformative potential responsibly, ethically, and effectively.

In sum, the path forward is clear: **govern AI not as a tool to be restrained, but as a partner to be stewarded.** This mindset enables organizations to capture the benefits of AI while safeguarding people, principles, and purpose.

ITIL®

THE LANGUAGE OF GROWTH

We would love to hear your thoughts!

Scan the QR code, share your feedback by 16th November 2025, and be recognised for your contributions!

ESG Commitment

PeopleCert has been accredited by Lloyd's Register, UK (now LRQA), in accordance with ISO 14001 for Environmental Management since 2006. Recognized through numerous awards, we remain committed to ESG leadership and the preservation of our planet.