

Enterprise AI Readiness & Governance Guidance

Foundational Considerations Before and During Enterprise AI Adoption

A practical baseline for public- and private-sector organizations

AI inherits the environment exactly as it exists today.

Prepared by Peter Gallinari

Technology, Security, Privacy, Data Governance, and AI Readiness Advisor

Purpose and Scope

Artificial intelligence is rapidly becoming part of everyday enterprise operations. Organizations are exploring AI to improve productivity, automate processes, support decision-making, accelerate research, enhance customer and citizen services, and create new products and capabilities.

Those opportunities are significant. However, responsible AI adoption requires more than selecting models, writing policy, training users, or implementing technical guardrails. AI operates within - and inherits - the organization's existing data, identity, access, privacy, security, vendor, and operational environment.

This guidance provides a generalized baseline for organizations considering, piloting, procuring, or already deploying AI. It is intended to complement existing AI policies, risk frameworks, security standards, privacy programs, procurement practices, and industry-specific requirements.

The central readiness question is not simply: 'Are we ready to use AI?' It is: 'Do we understand and govern the environment AI is being allowed to inherit?'

The Foundational AI Readiness Principle

AI does not independently correct weaknesses in the environment to which it is connected. Existing conditions may instead become easier to discover, correlate, automate, or scale.

- Existing permissions
- Existing data sprawl and duplication
- Unclear data ownership
- Incomplete or inconsistent classification
- Stale or excessive access
- Third-party and SaaS exposure
- Shadow IT and Shadow AI
- Weak lifecycle and retention practices
- Security and privacy gaps
- Operational inconsistencies

AI can accelerate both mature governance and immature governance. The underlying environment strongly influences which outcome an organization experiences.

Recommended AI Readiness and Governance Areas

1. Enterprise Data Discovery and Visibility

Organizations should establish sufficient visibility into the data AI may access, process, retrieve, summarize, correlate, or act upon.

Recommended considerations include:

- What data exists across the enterprise
- Where structured and unstructured data resides

- Who owns or is accountable for the data
- Who has access to it
- How it is shared, copied, or duplicated
- Whether unknown or unmanaged repositories exist
- What data resides in collaboration platforms, file shares, cloud services, endpoints, email, and non-production environments

Key Principle: You cannot govern, secure, or responsibly enable AI against data you do not understand.

2. Data Classification and Sensitivity Awareness

Data classification provides context for determining how information should be accessed, shared, protected, retained, and used by AI.

Recommended considerations include:

- Enterprise classification standards
- Public, internal, confidential, sensitive, regulated, and restricted data
- Metadata and labeling practices
- Automated classification capabilities
- Classification accuracy and validation
- AI handling rules based on classification

Key Principle: AI accelerates access to data. Classification helps determine whether that access is appropriate.

3. Business Ownership and Accountability

Technology teams may operate infrastructure and platforms, but business functions should remain accountable for the data and business processes they create and use.

Recommended considerations include:

- Defined data owners and stewards
- Business accountability for classification
- Data custodian versus data owner responsibilities
- Governance councils and escalation paths
- Executive accountability for AI-enabled processes

Key Principle: IT manages technology. The business owns and is accountable for its data.

4. Identity, Access Management, and Permission Drift

AI systems frequently rely on the permissions already granted to users, applications, service accounts, repositories, and connected platforms. Existing access weaknesses can therefore become AI weaknesses.

Recommended considerations include:

- Least privilege
- Identity lifecycle management

- Access recertification
- Permission inheritance
- Stale groups and accounts
- Collaboration-platform exposure
- AI access boundaries
- Runtime authorization and continuous verification

Key Principle: AI does not create every access problem. It can expose and amplify access problems that already exist.

5. Privacy and Data Protection

AI adoption should be evaluated against existing privacy obligations and organizational commitments regarding the collection, use, sharing, retention, and protection of personal and sensitive information.

Recommended considerations include:

- Privacy impact and risk assessments
- Purpose limitation and data minimization
- Sensitive and regulated information
- Retention and deletion
- Transparency and notice
- Automated decision impacts
- Cross-border considerations
- Individual rights and applicable legal obligations

Key Principle: AI governance should not operate separately from the organization's privacy program.

6. Security Readiness

AI introduces new attack surfaces and can also interact with longstanding weaknesses in enterprise systems and data environments.

Recommended considerations include:

- Secure architecture and configuration
- Zero Trust principles
- Secrets and credential protection
- Prompt injection and data exfiltration risks
- Logging and monitoring
- Security testing
- Incident response integration
- Model, application, API, and connector security

Key Principle: Security controls are strongest when the organization first knows what it is protecting and where it resides.

7. Third-Party, Vendor, and Supply-Chain AI Risk

Organizations increasingly consume AI through SaaS products, cloud platforms, embedded features, APIs, vendors, and business partners.

Recommended considerations include:

- AI supplier due diligence
- Provider and deployer responsibilities
- Data-use and model-training terms
- Subprocessors and downstream providers
- Contractual controls
- Logging and auditability
- Exit and portability provisions
- Ongoing vendor monitoring
- Embedded AI feature evaluation

Key Principle: Organizations remain accountable for their data even when third parties process it.

8. AI Inventory and Use-Case Governance

Organizations should know where AI is being used and for what purpose. Governance becomes difficult when AI capabilities are introduced without a maintained inventory and defined ownership.

Recommended considerations include:

- AI system and model inventory
- Business owner and technical owner
- Approved use cases
- Data sources and connectors
- Risk classification
- Decision authority
- Human oversight requirements
- Lifecycle status and retirement

Key Principle: You cannot govern AI capabilities you do not know exist.

9. Shadow AI and Unauthorized Use

Employees and business units can adopt consumer or embedded AI capabilities faster than traditional procurement and governance processes can respond.

Recommended considerations include:

- Shadow AI discovery
- Unauthorized data sharing
- Consumer AI usage
- Acceptable-use standards
- Approved tools and alternatives
- Employee awareness

- Monitoring and enforcement
- Exception processes

Key Principle: Shadow AI is the next evolution of Shadow IT.

10. Human Oversight and Operational Accountability

Human oversight should be designed according to the risk, autonomy, and consequences of the AI-enabled process rather than treated as a generic checkbox.

Recommended considerations include:

- Human-in-the-loop and human-on-the-loop models
- Approval and escalation boundaries
- Decision accountability
- Exception handling
- Auditability
- Risk-based review
- Separation of recommendation from authorization

Key Principle: Automation does not remove accountability.

11. Continuous Governance, Monitoring, and Change Management

AI environments evolve continuously. Models, prompts, data sources, connectors, permissions, business processes, and vendor capabilities can all change after initial approval.

Recommended considerations include:

- Continuous monitoring
- Model and output evaluation
- AI telemetry
- Drift detection
- Access and configuration change monitoring
- Periodic reassessment
- Policy and control updates
- Continuous compliance validation

Key Principle: Point-in-time approval is not continuous governance.

12. Execution-Layer Governance and Runtime Controls

Policies and documentation establish expectations, but higher-risk AI environments increasingly require controls that can enforce those expectations during execution.

Recommended considerations include:

- Pre-execution authorization
- Runtime policy enforcement
- AI decision gating
- Tool and action permissions

- Continuous verification
- Transaction and action logging
- Trust and risk signals
- Kill switches and containment mechanisms where appropriate

Key Principle: Effective AI governance must ultimately connect policy to operational enforcement.

13. Non-Production and Development Data

Development, testing, analytics, and sandbox environments deserve explicit attention because they can contain copied production information, broad access, or weaker controls.

Recommended considerations include:

- Prohibition or strict control of live production data in non-production environments
- Synthetic or masked data alternatives
- Test-data governance
- Developer access
- Temporary copies and exports
- Retention and disposal
- AI experimentation environments

Key Principle: AI experimentation should not become a reason to weaken established data-protection practices.

14. Regulatory, Sector, and Public-Trust Requirements

AI governance should incorporate the legal, regulatory, contractual, ethical, and trust requirements relevant to the organization's sector and jurisdiction.

Recommended considerations include:

- Government and citizen data
- Healthcare and patient information
- Financial and consumer information
- Education records
- Critical infrastructure
- Records management and transparency
- Industry-specific regulation
- Public accountability and explainability

Key Principle: Responsible AI must reflect the obligations and trust expectations of the environment in which it operates.

A Practical Pre-Deployment AI Readiness Check

Before approving a material AI deployment, leadership should be able to obtain evidence-based answers to questions such as:

- Do we know what data the AI can access - directly and through connected systems?
- Do we know where that data resides and who owns it?
- Has the relevant data been classified and labeled appropriately?
- Have access permissions been reviewed and validated?
- Do we understand data duplication, sharing, and non-production exposure?
- Have privacy, security, legal, records, and sector-specific obligations been assessed?
- Do we know which model, provider, connectors, tools, and subprocessors are involved?
- Is the AI use case inventoried, owned, risk-assessed, and approved?
- Are human accountability and escalation boundaries defined?
- Can policy requirements be monitored or enforced operationally?
- Can the organization log, investigate, contain, and stop the AI capability if necessary?
- Is there a process for continuous reassessment as the AI capability and underlying environment change?

Readiness should be demonstrated with evidence - not assumed because a policy, framework, or vendor assurance exists.

Leadership Perspective

AI governance should enable innovation, not simply restrict it. But sustainable innovation requires an environment that leadership understands well enough to accept, mitigate, or avoid the risks AI may introduce or amplify.

Organizations should resist treating AI readiness as a technology-only initiative. It is an enterprise operating-model issue involving business leadership, data governance, privacy, cybersecurity, architecture, procurement, legal, compliance, records management, risk, human resources, and operational teams.

The strongest AI environments will not necessarily be the fastest adopters. They will be the organizations that understand their environments well enough to deploy AI responsibly.

Final Perspective

Artificial intelligence is an operational transformation capability. It can accelerate productivity, innovation, insight, and service delivery. It can also accelerate poor access decisions, weak governance, data sprawl, privacy exposure, security weaknesses, and unclear accountability.

The objective is not to delay AI until an organization achieves a perfect environment. Few organizations ever will. The objective is to understand the environment well enough to make informed, risk-based deployment decisions - and to identify foundational weaknesses before AI magnifies them.

AI inherits your environment. Know the environment before you scale the AI.

This document is intended as generalized guidance and a practical starting point. Organizations should adapt it to their business, risk profile, regulatory obligations, technology environment, and AI use cases.