

Enterprise AI Readiness Implementation Framework

Applying the Enterprise AI Readiness & Governance Guidance Through DMAIC, Evidence, and Formal Tollgates

A practical execution framework for public- and private-sector organizations

DISCOVERY FIRST.

Evidence before assumption. Accountability before acceleration.

Prepared by Peter Gallinari

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

Purpose and Relationship to the Enterprise Guidance

The Enterprise AI Readiness & Governance Guidance defines the foundational areas an organization should understand and govern before and during enterprise AI adoption. It establishes the **WHAT**: the readiness domains, control considerations, accountability expectations, and operating principles that should be addressed.

This Implementation Framework defines the **HOW**. It provides a disciplined way to apply that guidance as an enterprise program using DMAIC thinking, evidence-based decision making, and formal tollgates.

The framework is not a replacement for the Enterprise AI Readiness & Governance Guidance. It is the execution model for putting that guidance into practice.

Key Principle: The guidance defines what must be addressed. This framework establishes how an organization proves readiness, makes decisions, documents risk, and authorizes movement to the next stage.

Why a Framework Is Needed

AI programs are often encouraged to move quickly, start with a use case, pilot, learn, and improve while building. Iterative delivery can be appropriate for AI solution development. It should not eliminate the discipline required to understand the environment into which AI is being introduced.

An organization may have legacy infrastructure, unknown repositories, duplicate information, stale permissions, shadow IT, unmanaged SaaS services, non-production copies, unclear ownership, and years of accumulated data. In such an environment, moving directly to implementation can mean building on conditions that have never been established or measured.

This framework separates two activities that are often blended together:

- AI solution development — which may be iterative, Agile, experimental, and use-case driven.
- Enterprise AI readiness — which requires sufficient discovery, evidence, analysis, accountability, and authorization before material access or deployment.

Key Principle: AI development may be iterative. Readiness must still be demonstrated.

DMAIC as the Program Discipline

DMAIC — Define, Measure, Analyze, Improve, Control — provides a useful discipline for enterprise AI readiness because it prevents organizations from jumping directly to implementation before establishing the current state.

DMAIC Phase	AI Readiness Application	Primary Outcome
DEFINE	Discover and scope the environment, stakeholders, business objective, data boundaries, systems, owners, and decision criteria.	Known scope and accountable ownership
MEASURE	Establish evidence of the current state: data locations, volumes, classifications, access, duplication, retention, control coverage, and known gaps.	Evidence-based baseline
ANALYZE	Evaluate gaps, dependencies, exposure, root causes, residual risk, and whether the environment is sufficiently ready for the proposed AI use.	Readiness and risk determination

IMPROVE	Remediate gaps, implement controls, configure the AI solution, test the use case, and address deficiencies identified through analysis.	Controlled implementation
CONTROL	Monitor, reassess, recertify, log, test, and govern changes to data, access, models, connectors, vendors, and business processes.	Sustained readiness and governance

DEFINE is where discovery begins. An organization cannot credibly define the boundaries of an AI initiative if it has not established what data, systems, repositories, identities, integrations, and owners exist within those boundaries.

Key Principle: Do not jump to Improve before Define, Measure, and Analyze have produced sufficient evidence to support the decision.

The Enterprise AI Readiness Tollgate Framework

The framework converts the Enterprise AI Readiness & Governance Guidance into a gated program. Each tollgate has defined evidence, accountable participants, decision criteria, and a formal disposition.

A tollgate is not a checklist item. It is a decision point. The organization must be able to demonstrate that the required work has been completed to an agreed level, that material gaps are known, and that someone with appropriate authority has accepted the decision to proceed.

Standard Tollgate Decision

- **PASS** — Required evidence is sufficient and the program may proceed.
- **PASS WITH CONDITIONS** — Known gaps remain, but documented compensating controls, remediation commitments, owners, and dates support proceeding.
- **HOLD** — Evidence is incomplete or risk exceeds the approved threshold. Additional work is required before proceeding.
- **EXCEPTION / RISK ACCEPTANCE** — Leadership elects to proceed despite an unmet requirement. The exception, rationale, residual risk, accountable risk owner, approval, and expiration/review date must be documented.

Key Principle: A green status without evidence is not a passed tollgate.

Tollgate 1 — Enterprise Data Discovery and Visibility

Objective: Establish sufficient visibility into the enterprise data environment relevant to the proposed AI capability before defining what AI should be allowed to access, retrieve, process, correlate, or act upon.

Required evidence should address:

- Known structured and unstructured data repositories within scope
- Collaboration platforms, file shares, databases, cloud services, SaaS platforms, email, endpoints, archives, exports, and non-production environments
- Known data flows, copies, duplication, and downstream locations
- Unknown, unmanaged, abandoned, or legacy repositories requiring investigation
- Business owners or accountable parties for discovered data
- Existing access paths and material permission exposure
- Discovery coverage, exclusions, limitations, and confidence level

Minimum tollgate question: Can leadership demonstrate, with evidence, that it understands the data environment AI is being allowed to inherit?

Suggested sign-off: Business/Data Owner, Data Governance, Technology/Architecture, Privacy, Security, AI Program Owner.

Key Principle: Discovery is first because every downstream control depends on visibility.

Tollgate 2 — Data Classification and Business Ownership

Objective: Determine what the discovered data is, how it should be handled, and who is accountable for the classification and permitted business use.

- Enterprise classification standard and handling rules
- Classification coverage for data within the AI scope

- Business validation of classifications
- Identification of sensitive, regulated, confidential, restricted, and public/internal data
- Unclassified or ambiguously classified data and remediation plan
- Authoritative sources versus obsolete, duplicate, or uncontrolled copies
- Named business ownership and stewardship responsibilities

Minimum tollgate question: Do we know what the discovered data is, who owns it, and whether its classification supports the proposed AI use?

Key Principle: Technology may manage platforms; the business remains accountable for its data.

Tollgate 3 — Privacy, Legal, Records, and Permitted Use

Objective: Establish whether the proposed AI processing is permitted and whether privacy, legal, contractual, records-management, retention, and sector obligations have been addressed.

- Purpose limitation and permitted-use analysis
- Privacy impact/risk assessment where appropriate
- Data minimization and retention requirements
- Legal, contractual, regulatory, and records-management obligations
- Notice, transparency, individual rights, and automated-decision considerations
- Cross-border or jurisdictional requirements where applicable
- Third-party data-use restrictions and downstream obligations

Minimum tollgate question: Are we authorized to use this data for this AI purpose, under the conditions being proposed?

Key Principle: Availability of data does not establish permission to use it.

Tollgate 4 — Security, Identity, Access, and Technical Boundaries

Objective: Validate that identities, permissions, architecture, integrations, and security controls appropriately constrain what the AI capability and its users can access or do.

- Least-privilege and entitlement review
- User, service, application, and non-human identity controls
- Permission inheritance, stale groups/accounts, and excessive access
- Runtime authorization and action boundaries
- API, connector, plugin, model, and integration security
- Secrets and credential protection
- Logging, monitoring, incident response, containment, and kill/disable mechanisms where appropriate
- Non-production data controls, including prohibition or strict control of live production data

Minimum tollgate question: Does actual technical reach match approved business authority?

Key Principle: Designed authority and actual technical reach must be reconciled before deployment.

Tollgate 5 — AI Use-Case and Model Readiness

Objective: Only after the foundational environment has been sufficiently established should the organization authorize the AI use case, model, agent, RAG capability, or other AI solution.

- Defined business problem, owner, intended outcome, and success measure
- Approved AI use case and risk classification
- Model/provider selection and documented rationale
- Approved data sources and connectors
- Prompt, retrieval, output, and model-risk controls
- Human oversight and decision-accountability model
- Agent/tool permissions and action boundaries where applicable
- Testing criteria, acceptance thresholds, fallback procedures, and escalation paths

Minimum tollgate question: Is the proposed AI solution appropriate for the business purpose and bounded by the environment and controls already validated?

Key Principle: First understand the environment. Then define the use case. Then choose and configure the technology.

Tollgate 6 — Deployment Authorization

Objective: Make the production decision explicit. Confirm that required prior tollgates have passed, conditional items are understood, residual risks are accepted by the correct authorities, and operating controls are ready.

- Prior tollgate status and unresolved conditions
- Testing and validation results
- Production architecture and configuration approval
- Training and acceptable-use readiness
- Operational support, incident response, and escalation readiness
- Vendor and contractual readiness
- Residual-risk statement and accountable risk owner
- Formal authorization to deploy

Minimum tollgate question: Who is authorizing this capability to enter production, based on what evidence, and with what accepted residual risk?

Tollgate 7 — Continuous Control and Reauthorization

Objective: Ensure that readiness is sustained after launch. AI environments change continuously; approval at one point in time is not permanent authorization.

- Continuous monitoring and AI telemetry
- Access and entitlement recertification
- Model, prompt, output, retrieval, and agent-behavior evaluation
- Changes to data sources, connectors, APIs, vendors, and subprocessors
- Drift, security testing, red teaming, and incident trends
- Policy, regulatory, contractual, and business-process changes
- Periodic reassessment and formal reauthorization based on risk

Minimum tollgate question: Has anything materially changed that invalidates the evidence or assumptions on which the prior approval was based?

Key Principle: Readiness is a maintained condition, not a one-time declaration.

Evidence Standard

The framework is evidence-based. Statements such as 'we believe,' 'the team is comfortable,' 'security approved it,' or 'the vendor says it is protected' should not substitute for demonstrable evidence.

Examples of acceptable evidence may include:

- Discovery scan results and repository inventories
- Data maps, lineage, inventories, and ownership records
- Classification reports and business validation
- Access-review and recertification results
- Privacy, security, legal, and vendor assessments
- Architecture and data-flow diagrams
- Configuration evidence and control test results
- AI inventory and use-case records
- Model/system evaluation results
- Logs, monitoring evidence, incident-response tests, and exception records

Key Principle: The burden is not to prove perfection. The burden is to prove that leadership understands enough of the current state to make an informed, accountable decision.

Formal Exception and Risk Acceptance

The framework does not assume every gap must be fully remediated before any AI activity can proceed. Organizations may make risk-based decisions. What should not occur is silent bypass of a failed prerequisite.

When proceeding through an unmet tollgate, the record should include:

- Requirement or control not satisfied
- Evidence showing the gap
- Business and technical impact
- Residual risk
- Compensating controls
- Remediation owner and target date
- Accountable risk owner
- Approving authority
- Expiration or mandatory review date

Standard Tollgate Sign-Off Record

Field	Required Record
Program / Use Case	Name and unique identifier
Tollgate	Gate number and title
Decision	PASS / PASS WITH CONDITIONS / HOLD / EXCEPTION

ENTERPRISE AI READINESS IMPLEMENTATION FRAMEWORK

Evidence Reviewed	Artifacts, assessments, test results, inventories, reports
Open Gaps	Known deficiencies and limitations
Residual Risk	Documented remaining risk
Required Actions	Owner, action, target date
Approvers	Named accountable roles / individuals
Decision Date	Date of authorization
Review / Expiration	Required reassessment date or triggering condition

Accountability Model

AI readiness is an enterprise responsibility. No single function should be expected to certify the entire environment.

Executive Sponsor / Risk Owner — Owns the business decision, accepts residual risk within delegated authority, and ensures unresolved issues are visible.

Business / Data Owner — Owns business purpose, data accountability, classification validation, permitted use, and acceptance of business-process impact.

AI Program / Product Owner — Coordinates the use case, evidence package, dependencies, testing, conditions, and tollgate readiness.

Data Governance — Establishes discovery, inventory, classification, ownership, lineage, quality, and lifecycle expectations.

Privacy / Legal / Records — Evaluates privacy, legal, contractual, records, retention, transparency, and permitted-use obligations.

Security / IAM — Validates identity, access, architecture, technical controls, monitoring, incident response, and runtime boundaries.

Technology / Architecture — Documents systems, integrations, data flows, environments, technical dependencies, and operational readiness.

Procurement / Third-Party Risk — Evaluates vendors, terms, subprocessors, data use, supply-chain dependencies, and exit requirements.

Internal Audit / Independent Assurance — Where appropriate, tests whether evidence, controls, and tollgate decisions are operating as represented.

Applying DMAIC Across the Tollgates

Phase	Primary Tollgates	Core Question	Exit Evidence
DEFINE	1–2	What environment, data, owners, scope, and business objective are we dealing with?	Discovery scope, inventory, ownership, classification approach
MEASURE	1–4	What is the actual current state and where are the measurable gaps?	Coverage metrics, access results, classification status, privacy/security baseline
ANALYZE	2–5	What do the gaps mean for this AI use case and what risk remains?	Risk analysis, permitted-use determination, architecture/access analysis, readiness decision
IMPROVE	4–6	What must be remediated, configured, tested, or constrained before deployment?	Remediation evidence, test results, approved configuration, deployment authorization
CONTROL	7	How will readiness remain valid as the environment changes?	Monitoring, recertification, reassessment, incident and change controls

Program Metrics

Metrics should demonstrate readiness progress and expose unknowns rather than create artificial confidence. Examples include:

- Percentage of in-scope repositories discovered and validated
- Percentage of discovered data with accountable business ownership
- Percentage of in-scope data classified and business-validated
- Percentage of privileged or material access reviewed and recertified
- Number and age of unknown/unmanaged repositories
- Number of open tollgate conditions and overdue remediation actions
- Percentage of AI use cases with complete inventory, ownership, and risk classification
- Number of approved exceptions and expired risk acceptances
- Percentage of material AI changes receiving reassessment before release

Executive Readiness Questions

- Can we demonstrate where the enterprise data relevant to this AI capability resides?
- What do we know remains undiscovered or outside current coverage?
- Who owns and classifies the data?
- Do actual permissions match intended business access?
- Are we permitted to use the data for this AI purpose?
- What can the AI or agent technically reach or do beyond the intended workflow?
- What evidence supports each answer?
- Which gaps remain, who owns them, and when will they be remediated?
- Who is accepting the residual risk?
- Who is signing the decision to proceed?
- What change would require us to stop, reassess, or reauthorize?

Final Perspective

This framework is intended to make the Enterprise AI Readiness & Governance Guidance operational. It does not advocate waiting for a perfect enterprise environment. It advocates knowing enough about the environment to make an informed decision, documenting what remains unknown or unresolved, and assigning accountability for proceeding.

AI programs can still use Agile methods, pilots, sprints, experimentation, and iterative improvement. Those methods operate inside the boundaries established by readiness governance; they do not replace it.

The discipline is straightforward: Define before Improve. Discover before Classify. Establish evidence before assumption. Require accountable authorization before acceleration.

Key Principle: DISCOVERY FIRST. You cannot classify what you have not discovered. You cannot protect what you do not know exists. You cannot govern what you cannot see.

This document is generalized guidance and a practical implementation framework. Organizations should adapt tollgates, evidence requirements, approval authorities, risk thresholds, and review frequency to their business, regulatory obligations, technology environment, and risk profile.