

Enterprise AI Readiness & Governance

21-Page Action Plan

A practical framework to assess readiness, design the target state, establish responsible governance, implement production AI, and continuously measure business value.

Discover → Assess → Prioritize → Design → Govern → Build → Pilot → Validate → Deploy → Operate → Measure → Optimize

Executive Overview

Enterprise AI readiness is a business transformation discipline, not a single technology deployment. Inteltel combines strategy, infrastructure, data, cybersecurity, operating-model design, and governance into one executable program.



Inteltel principle: move from AI experimentation to governed, repeatable enterprise capability.

Program outcomes

• Align AI investment with measurable business priorities.

• Create an enterprise inventory of AI systems, models, data dependencies, and owners.

• Establish target architecture and reusable AI platform capabilities.

• Implement governance gates for risk, security, privacy, quality, and human oversight.

• Move selected use cases from pilot to production with measurable KPIs.

Three-stage transformation

Stage 1 assesses readiness and risk. Stage 2 designs the target architecture, governance model, and investment roadmap. Stage 3 builds, pilots, validates, deploys, and operates AI capabilities under continuous governance.

The Inteltel AI Transformation Framework

The framework connects business outcomes to the technical and governance capabilities required to operate AI safely at enterprise scale.



Governance is embedded across every stage rather than added after deployment.

1. Assess

Discover business priorities, inventory current AI, evaluate maturity, assess data and infrastructure, identify security gaps, and establish the initial risk register.

2. Design

Prioritize use cases, define the target AI architecture, design governance, assign accountability, create risk tiers, and establish lifecycle approval gates.

3. Implement & Govern

Build shared platform services, launch controlled pilots, validate performance and risk controls, operationalize human oversight, deploy production AI, and continuously optimize.

Phase 1 — Executive Alignment & Discovery

Begin with leadership alignment so AI investment is tied to business value, risk appetite, and enterprise priorities.



Actions

• Conduct executive and business-unit workshops.

• Define desired outcomes: growth, productivity, service quality, resilience, security, or cost optimization.

• Document strategic constraints, regulatory obligations, risk appetite, and investment horizons.

• Identify executive sponsor, transformation leader, and accountable business owners.

Decision point

Agree on the business problems AI should address and the outcomes that will determine whether the program advances.

Phase 1 — Enterprise AI Inventory

Create a controlled view of where AI is already being used and where new AI capabilities are planned.



Inventory scope

• Generative AI assistants and copilots.

• Machine-learning models and analytics.

• AI-enabled SaaS applications.

• Internally developed models and APIs.

• Third-party models, data services, and autonomous/agentic workflows.

Required metadata

Capture owner, purpose, users, data sources, model/provider, integrations, business criticality, access method, risk classification, deployment status, and monitoring requirements.

Phase 1 — AI Readiness Assessment

Assess the organization across six connected readiness domains. The objective is to identify constraints before they become production failures.



Six domains

• Business: strategy, sponsorship, value cases, funding, and KPIs.

• Data: availability, quality, ownership, lineage, retention, and privacy.

• Infrastructure: cloud, data center, compute, storage, network, GPU, and edge readiness.

• Security: identity, Zero Trust, API protection, data security, monitoring, and incident response.

• AI technology: model platforms, integration, MLOps/LLMOps, testing, and automation.

• Governance: policy, accountability, risk management, documentation, compliance, and oversight.

Output

Produce a maturity baseline, evidence-backed gaps, remediation priorities, and readiness dependencies for each priority use case.

Phase 1 — Data Readiness

AI quality and trust depend on the data foundation. Inteltel evaluates whether enterprise data can be discovered, governed, protected, and delivered to AI systems reliably.



Actions

• Identify authoritative data domains and accountable owners.

• Assess quality, completeness, freshness, lineage, classification, and access controls.

• Map sensitive, regulated, confidential, and intellectual-property data.

• Define approved retrieval, vectorization, retention, and model-training practices.

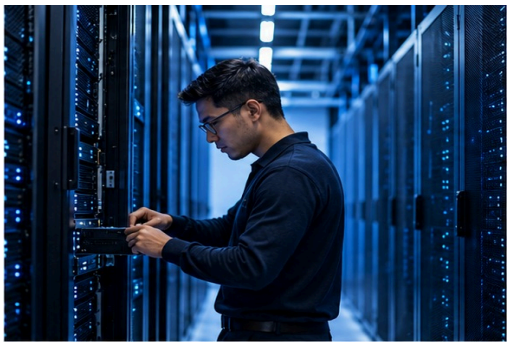
• Establish data-quality metrics and exception workflows.

Decision point

Determine which use cases have production-ready data, which require remediation, and which should be deferred.

Phase 1 — Infrastructure & Platform Readiness

Evaluate whether hybrid infrastructure can support AI workload performance, resilience, security, and economics.



Technical assessment

• Cloud and on-premises compute capacity, including accelerators where required.

• Storage performance, data locality, backup, and recovery.

• LAN/WAN/Internet/cloud connectivity and latency.

• API integration, identity federation, secrets management, and service connectivity.

• Observability, logging, telemetry, capacity planning, and cost controls.

Target outcome

A prioritized modernization backlog covering compute, data platforms, connectivity, security controls, observability, and AI platform services.

Phase 1 — Security, Risk & Governance Gap Analysis

Identify how AI changes the enterprise attack surface and where governance controls must be strengthened.



Review areas

• Identity and privileged access for users, services, models, and agents.

• Data leakage, prompt injection, insecure tool use, API abuse, and third-party risk.

• Privacy, intellectual property, retention, and approved-use requirements.

• Model reliability, harmful output, bias where relevant, and explainability needs.

• Incident response, audit evidence, monitoring, and exception management.

Deliverable

Create the initial AI risk register with owner, severity, affected systems, required treatment, due date, and residual-risk acceptance authority.

Phase 1 — Use-Case Prioritization

Convert discovery into an investable portfolio rather than a collection of disconnected experiments.



Scoring dimensions

• Business value and strategic alignment.

• Implementation feasibility and integration complexity.

• Data readiness and data sensitivity.

• Security, legal, operational, and reputational risk.

• Time-to-value, total cost, scalability, and reuse potential.

Portfolio structure

Organize opportunities as Quick Wins, Strategic Initiatives, and Transformational AI. Select a small number of pilots with clear owners, baselines, and measurable success criteria.

Phase 2 — Target Enterprise AI Architecture

Design a reusable architecture that separates user experiences, AI services, data, infrastructure, and control functions.



Logical layers

- Users & Business Applications
- AI Applications, Copilots & Agents
- AI Orchestration, Guardrails & API Services
- Enterprise Models, LLMs & ML Services
- Enterprise Data, Knowledge & Retrieval
- Cloud, Data Center & Edge Infrastructure
- Network, Identity, Cybersecurity & Observability

Cross-cutting controls

Governance, security, logging, policy enforcement, evaluation, cost management, and lifecycle evidence apply across every layer.

Phase 2 — AI Governance Operating Model

Establish decision rights so AI can scale without creating unmanaged business or technology risk.



Governance council

• Executive sponsor / business leadership

• CIO / CTO and enterprise architecture

• CISO and security operations

• Data and AI leadership

• Legal, privacy, compliance, and enterprise risk

• Business-unit owners and product leaders

• AI engineering, platform, and operations teams

Operating rhythm

Define intake, classification, architecture review, risk review, approval, exception handling, incident escalation, quarterly portfolio review, and annual policy refresh.

Phase 2 — AI Risk Tiers

Use risk tiers to match governance effort to business consequence rather than applying the same process to every AI use case.



Tier 1 — Lower risk

Internal productivity, drafting, summarization, and low-impact assistance with limited sensitive data.

Tier 2 — Moderate risk

Operational recommendations, customer assistance, workflow automation, or systems using sensitive enterprise information.

Tier 3 — High risk

Consequential decisions, regulated activities, material financial or operational impact, sensitive data, or significant autonomous action.

Control scaling

Higher tiers require stronger documentation, independent validation, security testing, human oversight, monitoring, approval authority, and evidence retention.

Phase 2 — AI Lifecycle Governance Gates

Every production AI capability should move through explicit governance checkpoints with defined evidence and accountable approvers.



A governance gate should enable informed decisions, not create an undefined approval queue.

Lifecycle

Idea → Business Case → Risk Classification → Architecture Review → Data Approval → Security Review → Model Evaluation → Governance Approval → Production → Continuous Monitoring → Retirement

Go / no-go evidence

• Named business and technical owners.

• Approved data sources and access model.

• Architecture and threat model.

• Evaluation results and acceptance thresholds.

• Human-oversight design and escalation path.

• Monitoring, incident response, rollback, and retirement plan.

Phase 2 — Responsible AI Policy & Standards

Translate governance principles into practical rules that teams can follow during design, development, procurement, and operations.



Policy areas

• Acceptable and prohibited AI use.

• Data handling, privacy, confidentiality, and intellectual property.

• Human oversight and accountability.

• Model and vendor due diligence.

• Testing, evaluation, documentation, and change control.

• Security requirements for prompts, tools, APIs, identities, and agents.

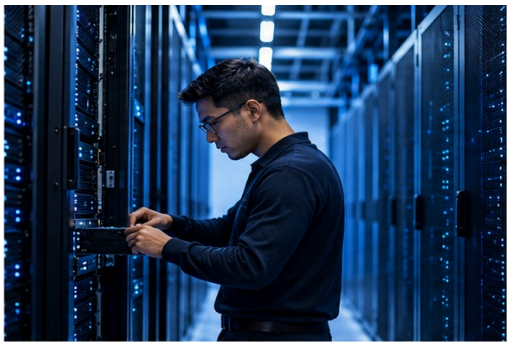
• Transparency, user communication, recordkeeping, and incident reporting.

Implementation

Publish policy with technical standards, templates, control owners, evidence requirements, exception workflow, and training.

Phase 3 — Build the Enterprise AI Foundation

Implement shared capabilities that reduce duplicated engineering and create consistent security and governance.



Foundation services

• Approved model and AI-service access layer.

• Enterprise identity, secrets, service accounts, and least-privilege access.

• Data connectors, retrieval services, vector stores, and knowledge controls.

• Prompt/model evaluation, testing, and release pipelines.

• Central logging, telemetry, policy enforcement, and cost management.

• MLOps/LLMOps for versioning, deployment, rollback, and monitoring.

Architecture objective

Create reusable patterns that business teams can consume without rebuilding governance and security for every project.

Phase 3 — Controlled Pilot Execution

Use pilots to validate business value and operating assumptions before enterprise-scale deployment.



Pilot sequence

• Define baseline performance and business KPI.

• Confirm owner, users, data, architecture, and risk tier.

• Build the minimum production-relevant workflow.

• Run security, privacy, model-quality, and operational tests.

• Conduct controlled user acceptance and capture human feedback.

• Compare results against baseline, cost, and risk thresholds.

Pilot examples

Enterprise knowledge assistant, service-desk copilot, network operations intelligence, cybersecurity analysis, document automation, or customer-service assistance.

Phase 3 — Validation, Security Testing & Human Oversight

Production approval requires evidence that the AI capability performs acceptably under normal, abnormal, and adversarial conditions.



Validation

• Accuracy and task-success metrics.

• Reliability, latency, availability, and failure behavior.

• Hallucination/error analysis and groundedness where applicable.

• Prompt-injection, data-exfiltration, unsafe tool-use, and API testing.

• Bias/fairness testing when relevant to the use case.

• Load, cost, resilience, recovery, and rollback testing.

Human oversight model

AI detects → AI analyzes → AI recommends → Human approves consequential actions → Automation executes → Platform records and monitors.

Phase 3 — Production Operations: AIOps, MLOps & LLMOps

Operationalize AI as a managed enterprise service with continuous visibility across models, infrastructure, data, security, cost, and business outcomes.



Monitor continuously

• Model quality, drift, errors, and evaluation trends.

• Infrastructure utilization, latency, availability, and capacity.

• Data quality, retrieval performance, and access anomalies.

• Security events, policy violations, and unusual agent/tool behavior.

• Token/API usage, cloud consumption, licensing, and unit economics.

• Adoption, human overrides, customer outcomes, and business KPIs.

Response model

Define alert thresholds, triage ownership, incident severity, model rollback, kill-switch authority, communications, root-cause review, and control improvements.

12-Month Enterprise AI Roadmap

Sequence the program so governance, architecture, platform engineering, and business adoption mature together.



Executive decision points: approve target state → fund foundation → approve pilots → accept production risk → authorize scale.

Months 1–2 | Assess

Executive alignment, discovery, AI inventory, maturity assessment, data/infrastructure review, risk analysis, and use-case prioritization.

Months 3–4 | Design

Target architecture, governance operating model, policies, risk tiers, lifecycle gates, security architecture, and investment roadmap.

Months 5–7 | Build

Shared AI platform, data integration, identity/security controls, observability, evaluation pipelines, and operating procedures.

Months 7–9 | Pilot

Launch priority use cases, validate value and risk, complete user acceptance, and prepare production support.

Months 9–12 | Scale & Govern

Production deployment, governance dashboard, portfolio reviews, reusable patterns, training, and expansion to additional business units.

Executive Scorecard, Deliverables & Next Steps

The program closes the loop between strategy, governance, engineering, and measurable business value.



Workstream	Core deliverable	Executive measure
Strategy	AI strategy & prioritized portfolio	Value, sponsorship, investment alignment
Readiness	Maturity scorecard & gap plan	Critical gaps closed
Governance	Policy, risk tiers & lifecycle gates	Coverage, exceptions, unresolved risk
Architecture	Target AI architecture	Reusable patterns & platform adoption
Security	AI security controls & risk register	Control effectiveness & incidents
Implementation	Pilots & production releases	KPI improvement & adoption
Operations	AI governance dashboard	Quality, availability, cost & risk trends

Recommended next steps

- Appoint the executive sponsor and cross-functional AI steering team.
- Run a 3–6 week enterprise AI readiness assessment.
- Select 2–3 priority use cases with measurable baselines and accountable owners.
- Approve the target architecture, governance model, and first-year investment roadmap.
- Launch controlled pilots and establish the AI governance dashboard before scaling.

Inteltel operating model

Discover → Assess → Prioritize → Design → Govern → Build → Pilot → Validate → Deploy → Operate → Measure → Optimize

Reference alignment: NIST AI Risk Management Framework (Govern, Map, Measure, Manage) and ISO/IEC 42001 AI management-system concepts. This document is an Inteltel consulting framework and should be tailored to each organization's legal, regulatory, industry, risk, and technology requirements.

INTELTEL | Enterprise AI Strategy • Infrastructure • Cybersecurity • Governance