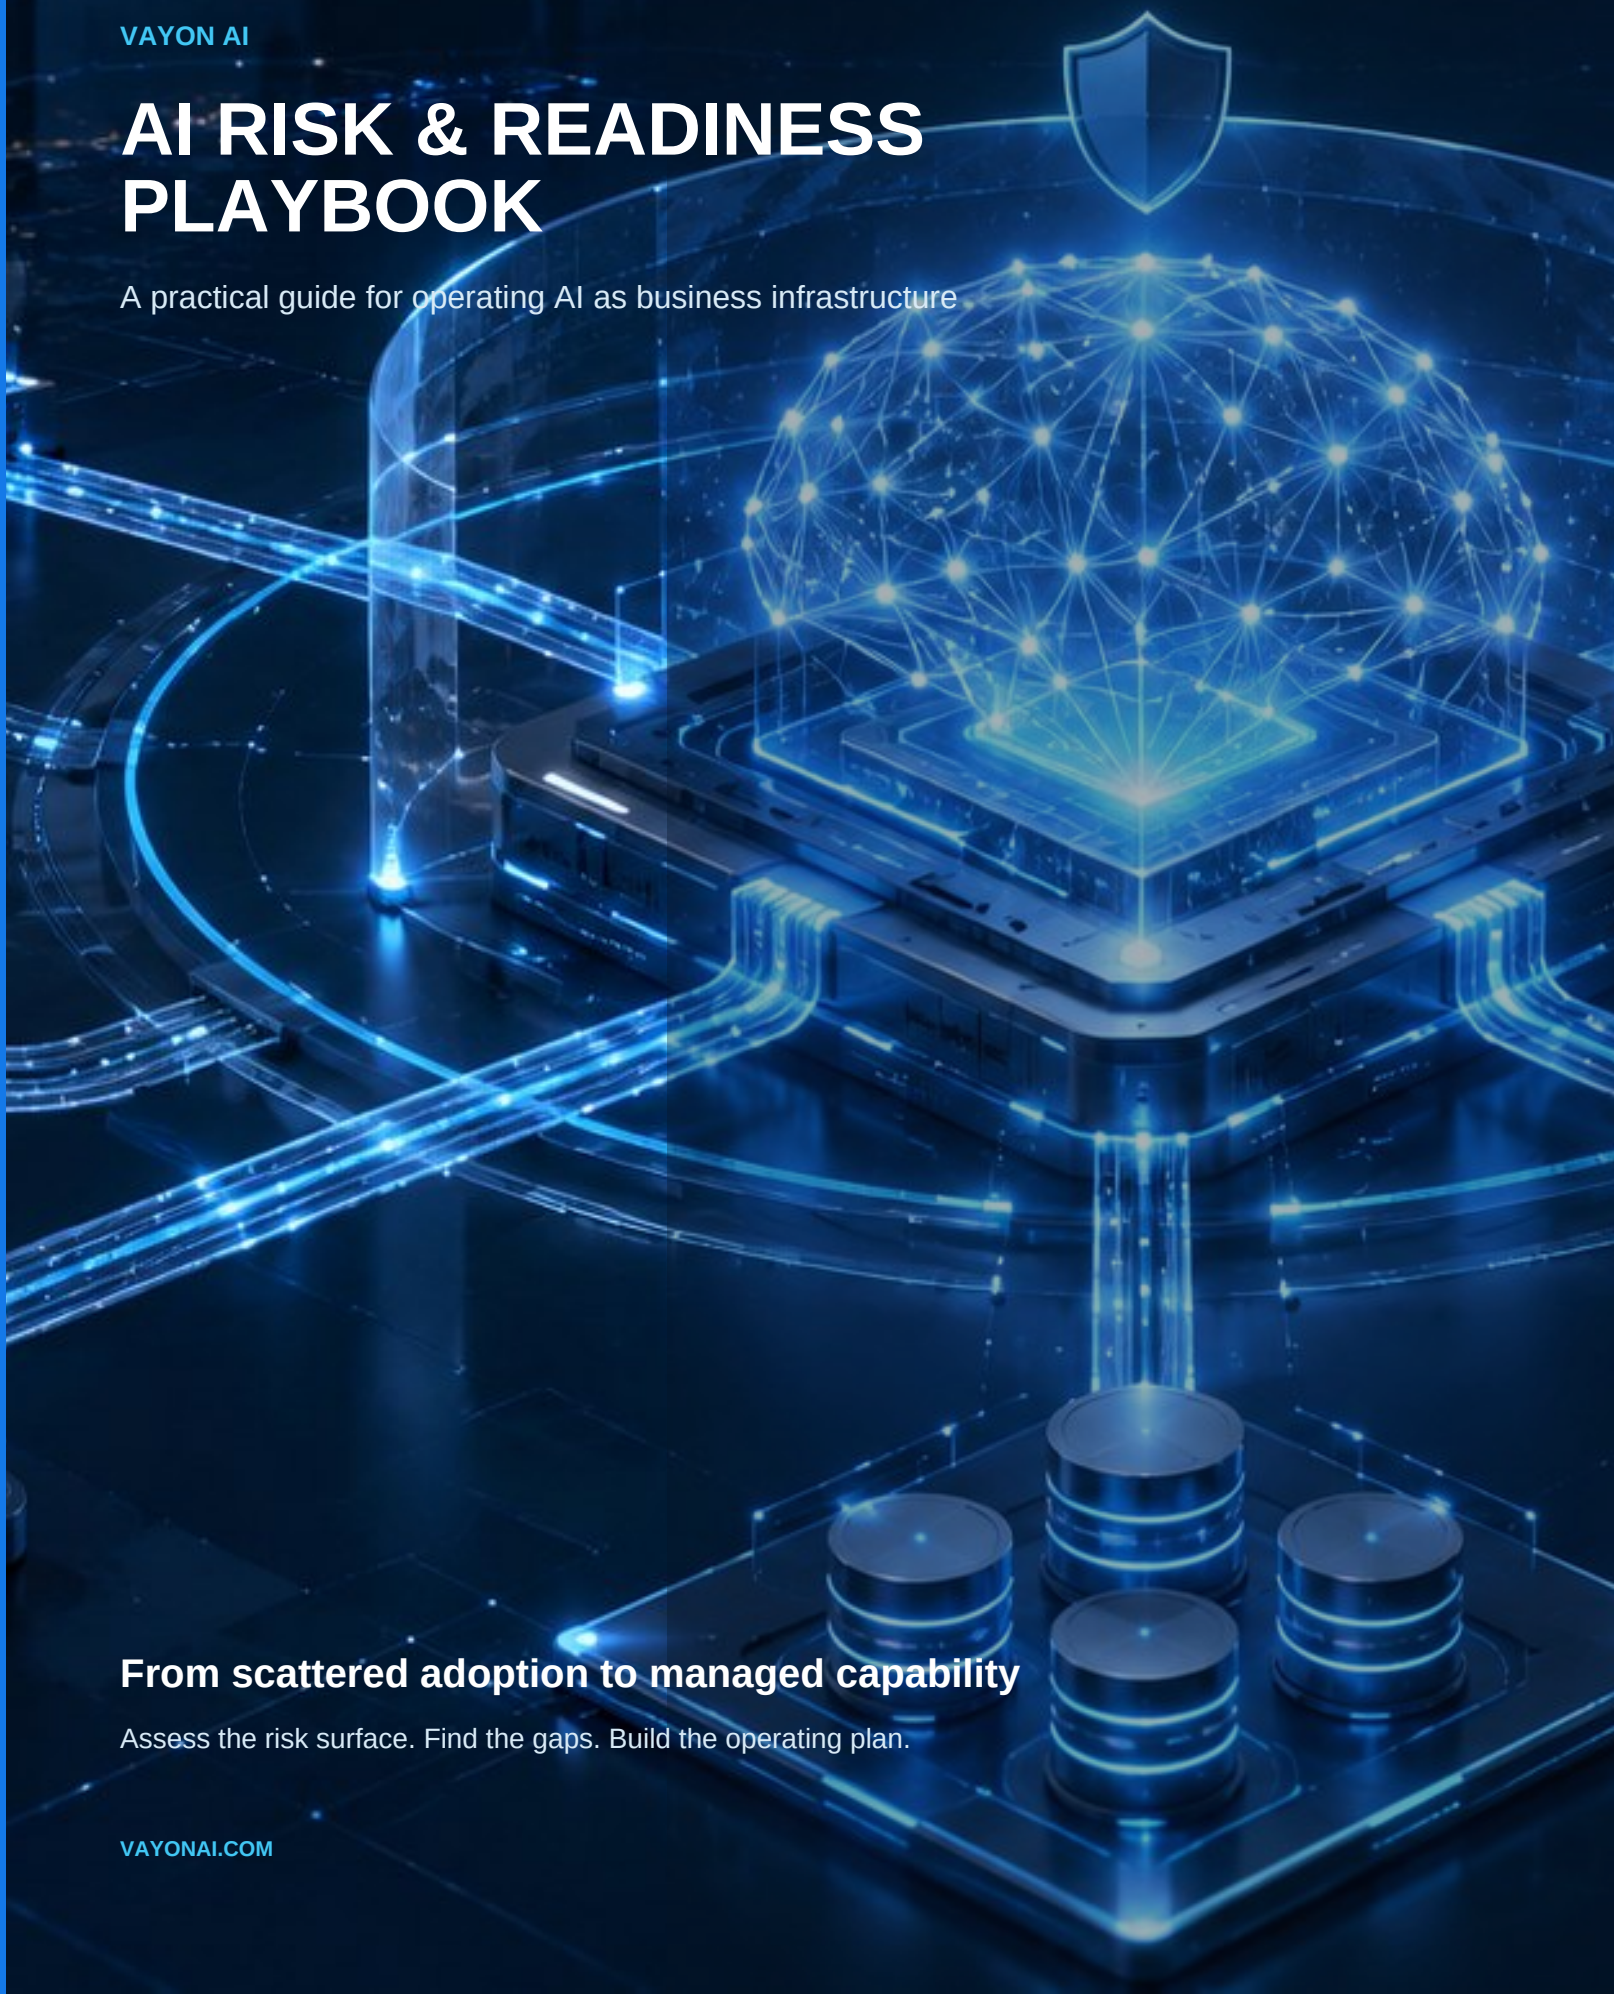
VAYON AI

AI RISK & READINESS PLAYBOOK

A practical guide for operating AI as business infrastructure



From scattered adoption to managed capability

Assess the risk surface. Find the gaps. Build the operating plan.

VAYONAI.COM

Executive brief

AI is moving from an experimental tool to a business dependency. The risk is no longer limited to what employees type into public tools. It now includes the models, providers, APIs, agents, applications, and workflows that increasingly shape how work gets done.

The operating question

Can the organization see, govern, route, secure, and recover AI-enabled work across models, providers, environments, users, systems, and agents?

A useful AI program connects five outcomes:

Visibility - know where AI is used and who owns it.

Control - enforce policy at the point of consumption.

Resilience - continue critical work when a model or provider changes or fails.

Economics - match workloads to appropriate capability and cost.

Accountability - assign owners and prove controls work.

What this playbook delivers

- A risk surface that leaders can use to structure discussion.

- A four-level maturity model for identifying the current state.

- A five-question assessment with evidence prompts and scoring.

- A 30/60/90-day operating plan and a leadership decision worksheet.

The central principle

Do not block AI to stay safe, and do not adopt AI without control. Operate it as managed infrastructure - securely, efficiently, and continuously.

The AI infrastructure model

AI becomes infrastructure when it is embedded in repeatable work. Use the model below to map what is connected, what is controlled, and what can fail.

LAYER	WHAT IT INCLUDES	CONTROL QUESTION
PEOPLE	Employees, developers, operators, business owners, and decision makers	Who can use AI, for what purpose, and with what accountability?
APPLICATIONS	Business systems, copilots, embedded AI features, workflows, and integrations	Which processes now depend on AI-enabled functionality?
AGENTS	Automated assistants, tools, actions, and recurring autonomous workflows	What can an agent do, and how is it monitored or stopped?
MODELS & PROVIDERS	Foundation models, APIs, local models, vendors, contracts, and changes	Can we replace a model or provider without disrupting work?
BUSINESS VALUE	Decisions, service delivery, productivity, revenue, and operational continuity	What happens to the business if the AI layer is unavailable?

Why this matters

A formal AI program may account for only part of actual usage. The rest may already be creating dependencies that are invisible to security, procurement, continuity, finance, or business leadership.

The first step is not more policy. It is a reliable map of where AI is present and what it supports.

The enterprise AI risk surface

Risk expands as AI moves deeper into the operating environment. Review each dimension for both direct exposure and dependency risk.

DIMENSION	TYPICAL FAILURE MODE	WHAT GOOD LOOKS LIKE
Data & security	Sensitive data reaches an unapproved route, or outputs are used without appropriate controls.	Data classification, identity, protection, logging, and tested response.
Governance	A policy exists, but no control changes what happens at consumption time.	Policy is translated into enforceable rules, approvals, and evidence.
Provider dependence	A provider changes an API, price, contract, geography, or availability.	Dependencies are known, alternatives are assessed, and portability is tested.
Agents & automation	An agent takes an unsafe action or becomes part of a process nobody owns.	Scopes, tools, approvals, monitoring, and stop mechanisms are defined.
Cost & architecture	Every workload uses the same expensive path, regardless of need.	Workloads are routed by data, capability, availability, and cost.
Continuity	A critical process stops when a model, API, or provider is unavailable.	Fallbacks, recovery steps, human workarounds, and ownership are tested.

A practical test

For every AI use case, ask: What data is involved? Which model and provider are used? Who owns the outcome? What policy applies? What happens if the service is unavailable? What evidence proves the control is working?

AI operating maturity

Maturity is not about using more AI. It is about increasing visibility, control, resilience, and repeatability as adoption grows.

LEVEL	CHARACTERISTICS	NEXT MOVE
1. Ad hoc	AI use is scattered, ownership is unclear, and risks are managed case by case.	Create visibility. Inventory use, owners, models, providers, and critical workflows.
2. Aware	Leadership knows AI is material, but controls are mostly documentation and manual review.	Translate policy into enforceable controls and define minimum evidence.
3. Managed	AI use is visible, policies affect consumption, and key dependencies have owners.	Add routing, portability, cost controls, and tested continuity paths.
4. Resilient	AI is operated as a governed capability with measurable outcomes, alternatives, and recovery.	Continuously improve based on evidence, incidents, cost, and business value.

Readiness indicators

You are moving from awareness to management when leaders can answer these questions with evidence: Where is AI used? Which processes depend on it? Which policies are enforced? What does it cost? Which alternatives exist? Who owns continuity?

The maturity goal is not perfect control. It is controlled progress with fewer surprises.

Five-question assessment scorecard

Score each question from 0 to 3. Use evidence, not estimates. A score of 0 or 1 should become an action in the operating plan.

#	QUESTION	EVIDENCE TO LOOK FOR	0-3
1	**Do we know where AI is already being used?**	Inventory, owner, application, model, provider, agent, and embedded capability evidence.	
2	**Can we enforce policy when AI is consumed?**	Identity, route, data controls, approvals, logging, exceptions, and response evidence.	
3	**Which business processes depend on models or providers?**	Critical workflow map, concentration, portability, contract, and change-impact evidence.	
4	**Can we route workloads by security, capability, availability, and cost?**	Classification, routing rules, model choice, cost visibility, and fallback evidence.	
5	**Could critical AI-enabled processes continue if a provider failed?**	Alternate route, recovery procedure, human workaround, and tested continuity evidence.	

Score interpretation

0-5: foundational visibility and control work is urgent. 6-10: important controls exist but are inconsistent or untested. 11-15: the organization has a managed baseline and should focus on resilience and optimization.

Overall score: _____ / 15 Priority gap:

_____

A 30/60/90-day operating plan

Use the sequence below to move from uncertainty to managed execution. Adapt it to the organization's risk, regulatory, and operating context.

TIMEFRAME	OUTCOMES	DELIVERABLES
FIRST 30 DAYS	Make AI visible and assign accountability.	AI inventory; critical-use-case list; owners; provider and model map; initial policy gaps.
DAYS 31-60	Make controls operational.	Minimum enforcement rules; data and route controls; logging; exception process; cost baseline.
DAYS 61-90	Test resilience and scale the operating model.	Fallback test; continuity runbook; portability assessment; executive scorecard; next-quarter roadmap.

Five operating principles

- Inventory before optimizing.
- Enforce at consumption time, not only in documents.
- Treat provider and model dependence as continuity risk.
- Route workloads by fit, not habit.
- Measure evidence, cost, resilience, and business value together.

Leadership checkpoint

At day 90, leadership should be able to decide which AI capabilities can scale, which dependencies require mitigation, and which controls must become non-negotiable.

Operating plan worksheet

Turn the assessment into named work. Keep each action concrete, owned, dated, and verifiable.

GAP / RISK	ACTION	OWNER	DATE	EVIDENCE

Evidence examples

- Approved route or policy rule with an audit record.
- Inventory record linked to a named owner and business process.
- Test result showing a fallback or recovery path works.
- Usage and cost report connected to workload decisions.
- Review record showing an exception was assessed and closed.

A control is not complete because it was designed. It is complete when the organization can show that it works.

Make AI operable

AI governance is becoming part of how the organization operates. The practical objective is not to slow responsible adoption. It is to make AI usable at scale without losing visibility, control, resilience, or economic discipline.

Use this playbook to start the conversation, identify the first material gaps, and create a plan that business and technology leaders can own together.

The next leadership decision

Which AI-enabled business process will we map, govern, and resilience-test first?

_____

Executive sponsor: _____ Date: _____

VAYON AI

Practical guidance for a more resilient AI future.

For more practical guidance on enterprise AI governance and operations, visit VayonAI.com.