



Dependable Outsourcing
Since 1999

Agentic AI

Readiness Assessment

A Self-Audit Framework for Enterprises
Considering AI Agents



Dependable Outsourcing
Since 1999

Agentic AI promises a self-optimizing, autonomous ecosystem—but only 1 in 10 enterprises is ready to deploy it at scale. Having worked with that 1 (and 9 others), we know what tells. Which is why we have devised a self-audit framework that helps diagnose Agentic AI readiness gaps across infrastructure, governance, and strategy—before making costly implementation mistakes.

Use it to:

- Benchmark your organization's AI maturity
- Identify where AI Agents can realistically add value
- Prioritize foundational investments for scalable adoption



Alignment Audit

Goal

Determine if Agentic AI aligns with your core business model and competitive priorities.

Evaluation Area	Key Questions	Ideal Indicators
Business Relevance	Are the identified use cases tied to measurable KPIs (e.g., cost-to-serve reduction, faster decision cycles)?	Clear ROI linkages; use cases validated through pilots.
Problem Solution Fit	Do the workflows involve multi-step decision-making or cross-system orchestration (vs. simple automation)?	At least one validated workflow that benefits from autonomy.
Stakeholder Buy-In	Is executive sponsorship in place with a defined governance structure?	Steering committee with technical & business leads.

Action Imperative:

Only proceed if at least **2 out of 3 indicators** score positively.

Data & Infrastructure Readiness

(Your Foundation for Agentic AI)

Goal

Evaluate whether your current data quality and systems can support autonomous, reasoning-driven workflows.

At this stage, the focus is not on building autonomous systems, but on **creating the right preconditions**: clean, connected data; interoperable systems; and optimized cloud environments.

Area	What to Assess	Why It Matters
Data Centralization	Are key datasets unified or still siloed?	Agents need real-time, cross-system visibility.
Data Quality & Standards	Are schemas and identifiers consistent?	Clean, standardized data enables reasoning.
APIs & Integrations	Do systems expose APIs or event streams?	Agents act through API communications, which is why closed systems can block progress.
Cloud & Scalability	Are workloads cloud-based and elastic?	Agentic models demand flexible, on-demand compute.
Observability	Is monitoring centralized?	Continuous telemetry enables learning and adaptation.
Security & Compliance	Are access, encryption, and audit controls in place?	Protects sensitive data before autonomy scales.

Action Imperative:

- Invest in AI-first data engineering
- Modernize integration layers
- Set up logging and monitoring
- Move compute workloads to elastic cloud infrastructure (AWS, Azure, GCP)
- Implement data catalogs and access control policies for governance



When you can integrate real-time, high-quality data from at least three core systems (e.g., CRM, ERP, support platform) into a unified warehouse with API accessibility, you're ready to design limited Agentic AI workflows.

Organizational Capability & Governance

Goal

Assess whether your teams and structures can support the operational complexity of AI Agents.

Focus Area	Assessment Lens	Ideal State
Talent Readiness	Do you have AI engineers, prompt designers, and human-in-the-loop QA specialists?	You should have at least 30% of the tech workforce trained on LLMOps & AI orchestration.
Governance & Ethics	Are there established protocols for AI accountability and human override?	Clear RACI model (responsibility assignment); risk committee in place.
Change Management	Is there a communication and adoption plan for non-technical teams?	Documented adoption roadmap with executive oversight.

Action Imperative:

- Identify which roles (engineering, operations, compliance) are directly impacted by the adoption of Agentic AI.
- Train existing teams in AI fundamentals, prompt engineering, and human-in-the-loop QA.
- Establish defined RACI roles.
- Form a small, cross-functional oversight body (IT, compliance, business ops).
- Integrate checkpoints for human validation in any decision-critical process.

Use Case Validation

Goal

Identify where Agentic AI can deliver measurable outcomes—and where it cannot.

Use Case Category	Agentic AI Fit?	Why / Why Not
Routine back-office automation (e.g., data entry, scheduling)	No	Overkill; rule-based scripts or process automation are more efficient.
Complex service workflows (e.g., supply chain coordination)	Yes	High inter-system dependencies; AI Agents can optimize dynamically.
Real-time decision support (e.g., fraud detection, customer escalation routing)	Yes	Contextual reasoning and adaptive responses add value.
Creative strategy or compliance reviews	Partial	Requires nuanced human judgment; partial augmentation only.

Action Imperative:

Select 1 or 2 high-impact use cases for a limited proof of concept (POC) rather than broad deployment. It is even better to get your use case validated by AI consultants.



Dependable Outsourcing
Since 1999

Risk, ROI, and Sustainability Mapping

Goal

Balance ambition with accountability.

Factor	What to Assess?	Recommended Threshold
ROI Visibility	Can you quantify success in terms of time, cost, or quality improvements? And how soon will that be measurable?	Tangible KPIs within 6–9 months.
Regulatory Risk	Are workflows compliant with GDPR, CCPA, SOC 1 and 2, HIPAA, or other local regulations?	100% alignment with internal data policies.
Sustainability	Can AI Agents continuously learn without performance drift?	Active monitoring & retraining loop established.

Action Imperative:

If more than 2 risk indicators fail to meet thresholds, postpone large-scale Agentic AI adoption.



Dependable Outsourcing
Since 1999



Rohit Bhateja

Director of Digital Engineering, SunTec India

Enterprises often pursue autonomy without fixing the fundamentals. Agentic AI only works when your data is reliable, your systems are connected, your governance is mature, and your teams know how to operate in an AI-driven environment. Without that foundation, autonomy becomes ambition without execution, and even the most advanced AI Agents will collapse under their own promise.

Autonomy is the outcome, but readiness is the fundamental prerequisite.