

Agentic AI Problem-First Design Checklist

Use this checklist to guide the design, validation, and deployment of agentic AI systems built around real business problems rather than tools or models.

Problem Definition

Have you clearly defined the business process or bottleneck the AI agent will address?

Is the problem measurable with clear success metrics?

Does this problem require autonomous decision-making or simple automation?

Scope and Boundaries

Are the agent's permissions and action limits clearly documented?

Have escalation paths been defined for edge cases or failures?

Is there a human-in-the-loop where required?

Data and Context Readiness

Are data sources reliable, current, and accessible via APIs or secure integrations?

Is sensitive or regulated data properly protected?

Have you validated data quality and consistency?

Architecture and Framework Selection

Does the selected AI agent framework support multi-step workflows?

Can the system scale horizontally?

Is observability built into the design?

Testing and Validation

Have you tested the agent across realistic scenarios?

Are failure states documented and monitored?

Is bias and behavior drift being tracked?

Security and Compliance

Are authentication and authorization mechanisms in place?

Is system activity logged for audits?

Does the system meet regulatory requirements?

Deployment and Monitoring

Are performance KPIs defined and tracked?

Is real-time monitoring enabled?

Do you have a rollback strategy?

Continuous Improvement

Are feedback loops in place for learning and optimization?

Is there a roadmap for expanding agent capabilities?

Are stakeholders reviewing outcomes regularly?