Rocket software

Are Your Mission-Critical Systems Ready for Agentic AI?

A governance-first readiness checklist for enterprise operations

Agentic AI introduces execution into enterprise systems, not just insight. This checklist helps enterprise leaders assess whether their core systems, operating model, and controls are ready to support agentic AI safely, compliantly, and at scale.

This is not a tool evaluation.

It is an operating-readiness assessment.

How to use this checklist

For each question, answer **Yes**, **Partially**, or **No**.

Mostly Yes: Your environment appears structurally ready for governed agentic AI.

Mixed responses: Agentic AI should remain limited to advisory or tightly constrained use cases.

Mostly No: Introducing agentic AI today may create material operational and compliance risk.

Areas with Partially or No responses may indicate material risk when introducing agentic AI into core systems.

Section 1 | Governance and policy control

Foundational controls should be in place before moving forward.

Do you enforce policy-as-code for automated actions, rather than relying only on written policies?

◯ Yes ◯ Partially ◯ No

Can policies allow, constrain, or deny agentic actions in real time?

◯ Yes ◯ Partially ◯ No

Are execution permissions role-based and identity-aware?

◯ Yes ◯ Partially ◯ No

Can you dynamically restrict actions based on risk, context, or system state?

◯ Yes ◯ Partially ◯ No

Why this matters: Agentic AI without embedded governance can create uncontrolled automation risk.

Section 2 | Execution safety and determinism

The difference between automation and unintended action.

Are all agentic actions validated before execution, rather than only monitored after the fact?

◯ Yes ◯ Partially ◯ No

Is there a rule-based control layer preventing out-of-scope actions?

◯ Yes ◯ Partially ◯ No

Can high-risk actions require human approval?

◯ Yes ◯ Partially ◯ No

Can the system stop execution when confidence or context is insufficient?

◯ Yes ◯ Partially ◯ No

Why this matters: Speed without validation can amplify risk faster than value.

Section 3 | Auditability, traceability, and reversibility

Critical requirements in regulated environments.

After an action, can you clearly answer what happened, why it happened, which policy allowed it, and which identity executed it?

◯ Yes ◯ Partially ◯ No

Are all actions logged in a tamper-resistant audit trail?

◯ Yes ◯ Partially ◯ No

Can you confidently roll back or remediate actions and outcomes?

◯ Yes ◯ Partially ◯ No

Would this evidence stand up to a regulatory or internal audit review?

◯ Yes ◯ Partially ◯ No

Why this matters: If you cannot explain how AI acted, it should not operate in production without tighter controls.

Section 4 | Diagnostic depth and operational context

Action without understanding creates risk.

Can AI trace issues from symptom to transaction to subsystem to job to code?

◯ Yes ◯ Partially ◯ No

Does it understand system dependencies, rather than only reading metrics or logs?

◯ Yes ◯ Partially ◯ No

Can it correlate data across applications, infrastructure, and business services?

◯ Yes ◯ Partially ◯ No

Is it grounded in your operational domain, rather than returning generic model output?

◯ Yes ◯ Partially ◯ No

Why this matters: Automated action without deep diagnosis can lead to inaccurate remediation.

Section 5 | Hybrid and core-system readiness

Where many platforms become difficult to scale.

Can AI operate where data lives, without forced data movement?

◯ Yes ◯ Partially ◯ No

Are existing mainframe security and governance controls preserved?

◯ Yes ◯ Partially ◯ No

Can the same operating model apply across mainframe, distributed, and cloud environments?

◯ Yes ◯ Partially ◯ No

Does your approach avoid too many disconnected tools and manual stitching?

◯ Yes ◯ Partially ◯ No

Why this matters: Core systems modernize more successfully through integration than replacement.

Section 6 | Skills gap and operational resilience

AI should reduce dependency on scarce experts.

Can junior staff diagnose issues that previously required deep subject matter experts?

◯ Yes ◯ Partially ◯ No

Does AI explain why something happened, not just what happened?

◯ Yes ◯ Partially ◯ No

Is institutional knowledge captured, reused, and improved over time?

◯ Yes ◯ Partially ◯ No

Does agentic AI demonstrably reduce mean time to resolution in real incidents?

◯ Yes ◯ Partially ◯ No

Why this matters: AI should strengthen operational resilience, not create new bottlenecks.

Interpreting your results

Mostly Yes:
Your environment appears structurally ready for governed agentic AI.

Mixed responses:
Agentic AI should remain limited to advisory or tightly constrained use cases.

Mostly No:
Introducing agentic AI today may create material operational and compliance risk.

Closing thought

Successful enterprises do not start agentic AI with models — they start with operating discipline. Platforms built for governed, operational-grade agentic AI, such as Rocket® EVA™, are designed specifically to help enterprises move safely from experimentation to trusted operation.

If you'd like to discuss your checklist results or learn how Rocket can help you get your core systems AI-ready, connect with one of our experts. We're here to help you move forward with confidence.

Rocket® software

Modernization. **Without Disruption.**™

Talk to an expert

Visit RocketSoftware.com

MAR-18237_EVA_MidFunnelChecklist_V3