# Quick Evaluation:
# Are You Ready for AI Agent Deployment?

**Time required:** 5–10 minutes
Answer these questions based on your current capabilities. Score 1 point for each YES answer.

---

## Foundation: Visibility and Isolation (5 questions)

### 01. Can you inventory every AI system that has tool access?

- ☐ **YES** – We maintain a complete registry of AI systems with their permissions documented.
- ☐ **NO** – We lack visibility into what AI systems exist or what they can access.

### 02. Do you know what your AI agents can actually do in your environment?

- ☐ **YES** – We've mapped AI permissions to actual system capabilities.
- ☐ **NO** – We cannot definitively answer what actions AI systems could perform.

### 03. Are you isolating MCP servers and tools today?

- ☐ **YES** – AI operations run in isolated environments with network segmentation.
- ☐ **NO** – AI systems run on general infrastructure without special isolation.

### 04. Have you identified over-privileged AI systems?

- ☐ **YES** – We've audited and flagged AI systems with excessive permissions.
- ☐ **NO** – We haven't systematically reviewed AI privilege levels.

### 05. Do you log all AI-driven actions with sufficient retention?

- ☐ **YES** – We capture comprehensive logs (90+ days) of all AI tool invocations.
- ☐ **NO** – Logging is incomplete or retention is insufficient.

Foundation Score: ___ / 5

---

# Quick Evaluation:
# Are You Ready for AI Agent Deployment?

**Time required:** 5–10 minutes
Answer these questions based on your current capabilities. Score 1 point for each YES answer.

## Implementation (5 questions)

06. Are you using authorization controls with scoped credentials?

- ☐ **YES** – AI agents use OAuth with PKCE, not API keys or static credentials.
- ☐ **NO** – We use API keys or haven't implemented OAuth.

07. Do all AI-callable tools have structured schemas with validation?

- ☐ **YES** – Every tool has defined input/output schemas that block malformed inputs.
- ☐ **NO** – Tools accept loosely defined inputs without strict validation.

08. Have you configured human-in-the-loop controls for high-risk operations?

- ☐ **YES** – AI systems pause for human input on sensitive or ambiguous operations.
- ☐ **NO** – AI systems operate autonomously without human checkpoints.

09. Do you verify the authenticity of tools in your MCP ecosystem?

- ☐ **YES** – Tools undergo security review and integrity verification before deployment.
- ☐ **NO** – Tools are deployed without formal validation.

10. Can you monitor and detect anomalous AI behavior?

- ☐ **YES** – We have real-time monitoring with alerting for unusual patterns.
- ☐ **NO** – We lack real-time visibility into AI operations.

MCP 2.0 Score: ___ / 5

# Quick Evaluation:
# Are You Ready for AI Agent Deployment?

**Time required:** 5–10 minutes
Answer these questions based on your current capabilities. Score 1 point for each YES answer.

---

## Zero-Trust Alignment (5 questions)

11. Do you use per-application service identities for AI systems?

- ☐ **YES** – Each AI system has dedicated identity with scoped permissions.
- ☐ **NO** – AI systems share credentials or use broad service accounts.

12. Have you refactored legacy tools to reduce privilege?

- ☐ **YES** – We've reviewed and right-sized permissions for AI-accessible tools.
- ☐ **NO** – Legacy tools retain their original broad permissions.

13. Are you aligning MCP usage to zero-trust principles?

- ☐ **YES** – Every AI action is authenticated, authorized, and continuously validated.
- ☐ **NO** – AI operations rely on implicit trust after initial authentication.

14. Do you have incident response procedures for AI-related events?

- ☐ **YES** – Our playbook covers AI scenarios with tested disable/containment procedures.
- ☐ **NO** – We lack specific AI incident response capabilities.

15. Can you recover from AI-related errors or compromise?

- ☐ **YES** – We have tested rollback procedures for AI operations.
- ☐ **NO** – We lack specific recovery capabilities for AI systems.

Zero-Trust Score: ___ / 5

---

**Your Total Score: ___ / 15**

# Quick Evaluation:
# Are You Ready for AI Agent Deployment?

## Scoring Guide

### 0–5 points:

Not ready for AI agent deployment. Critical gaps in basic governance. Focus on Foundation items before any deployment.

### 6–9 points:

Ready for controlled pilots. You have basics but gaps remain. Pilot with non-production systems and low-risk operations while addressing MCP 2.0 implementation.

### 10–12 points:

Ready for limited production. Strong foundation with partial MCP 2.0 implementation. Deploy for specific use cases with appropriate monitoring and controls.

### 13–15 points:

Production ready. Comprehensive governance in place. You can deploy AI agents confidently for most business operations.

## Priority Actions by Score

### If you scored 0–5:

1. Complete AI system inventory within 30 days.
2. Implement isolation and logging immediately.
3. Create governance policy before any new deployments.

### If you scored 6–9:

1. Implement authorization controls for all AI systems.
2. Define and validate tool schemas.
3. Configure human-in-the-loop workflows before production.

### If you scored 10–12:

1. Complete remaining MCP 2.0 items.
2. Enhance monitoring and anomaly detection.
3. Conduct security assessment before scaling.

### If you scored 13–15:

1. Maintain governance through continuous improvement.
2. Share your approach as industry best practice.
3. Prepare for next-generation AI security challenges.

# Quick Evaluation:
# Are You Ready for AI Agent Deployment?

## Next Steps

1. Complete this assessment with your security and IT teams.
2. Review results with CIO/CISO to prioritize actions.
3. Use the Risk Decision Guide to classify your AI deployments.
4. Reassess quarterly as you implement improvements.

Based on Werner Nel's MCP 2.0 security analysis. Part of the MCP 2.0 AI Security Risk Analysis Readiness Report.