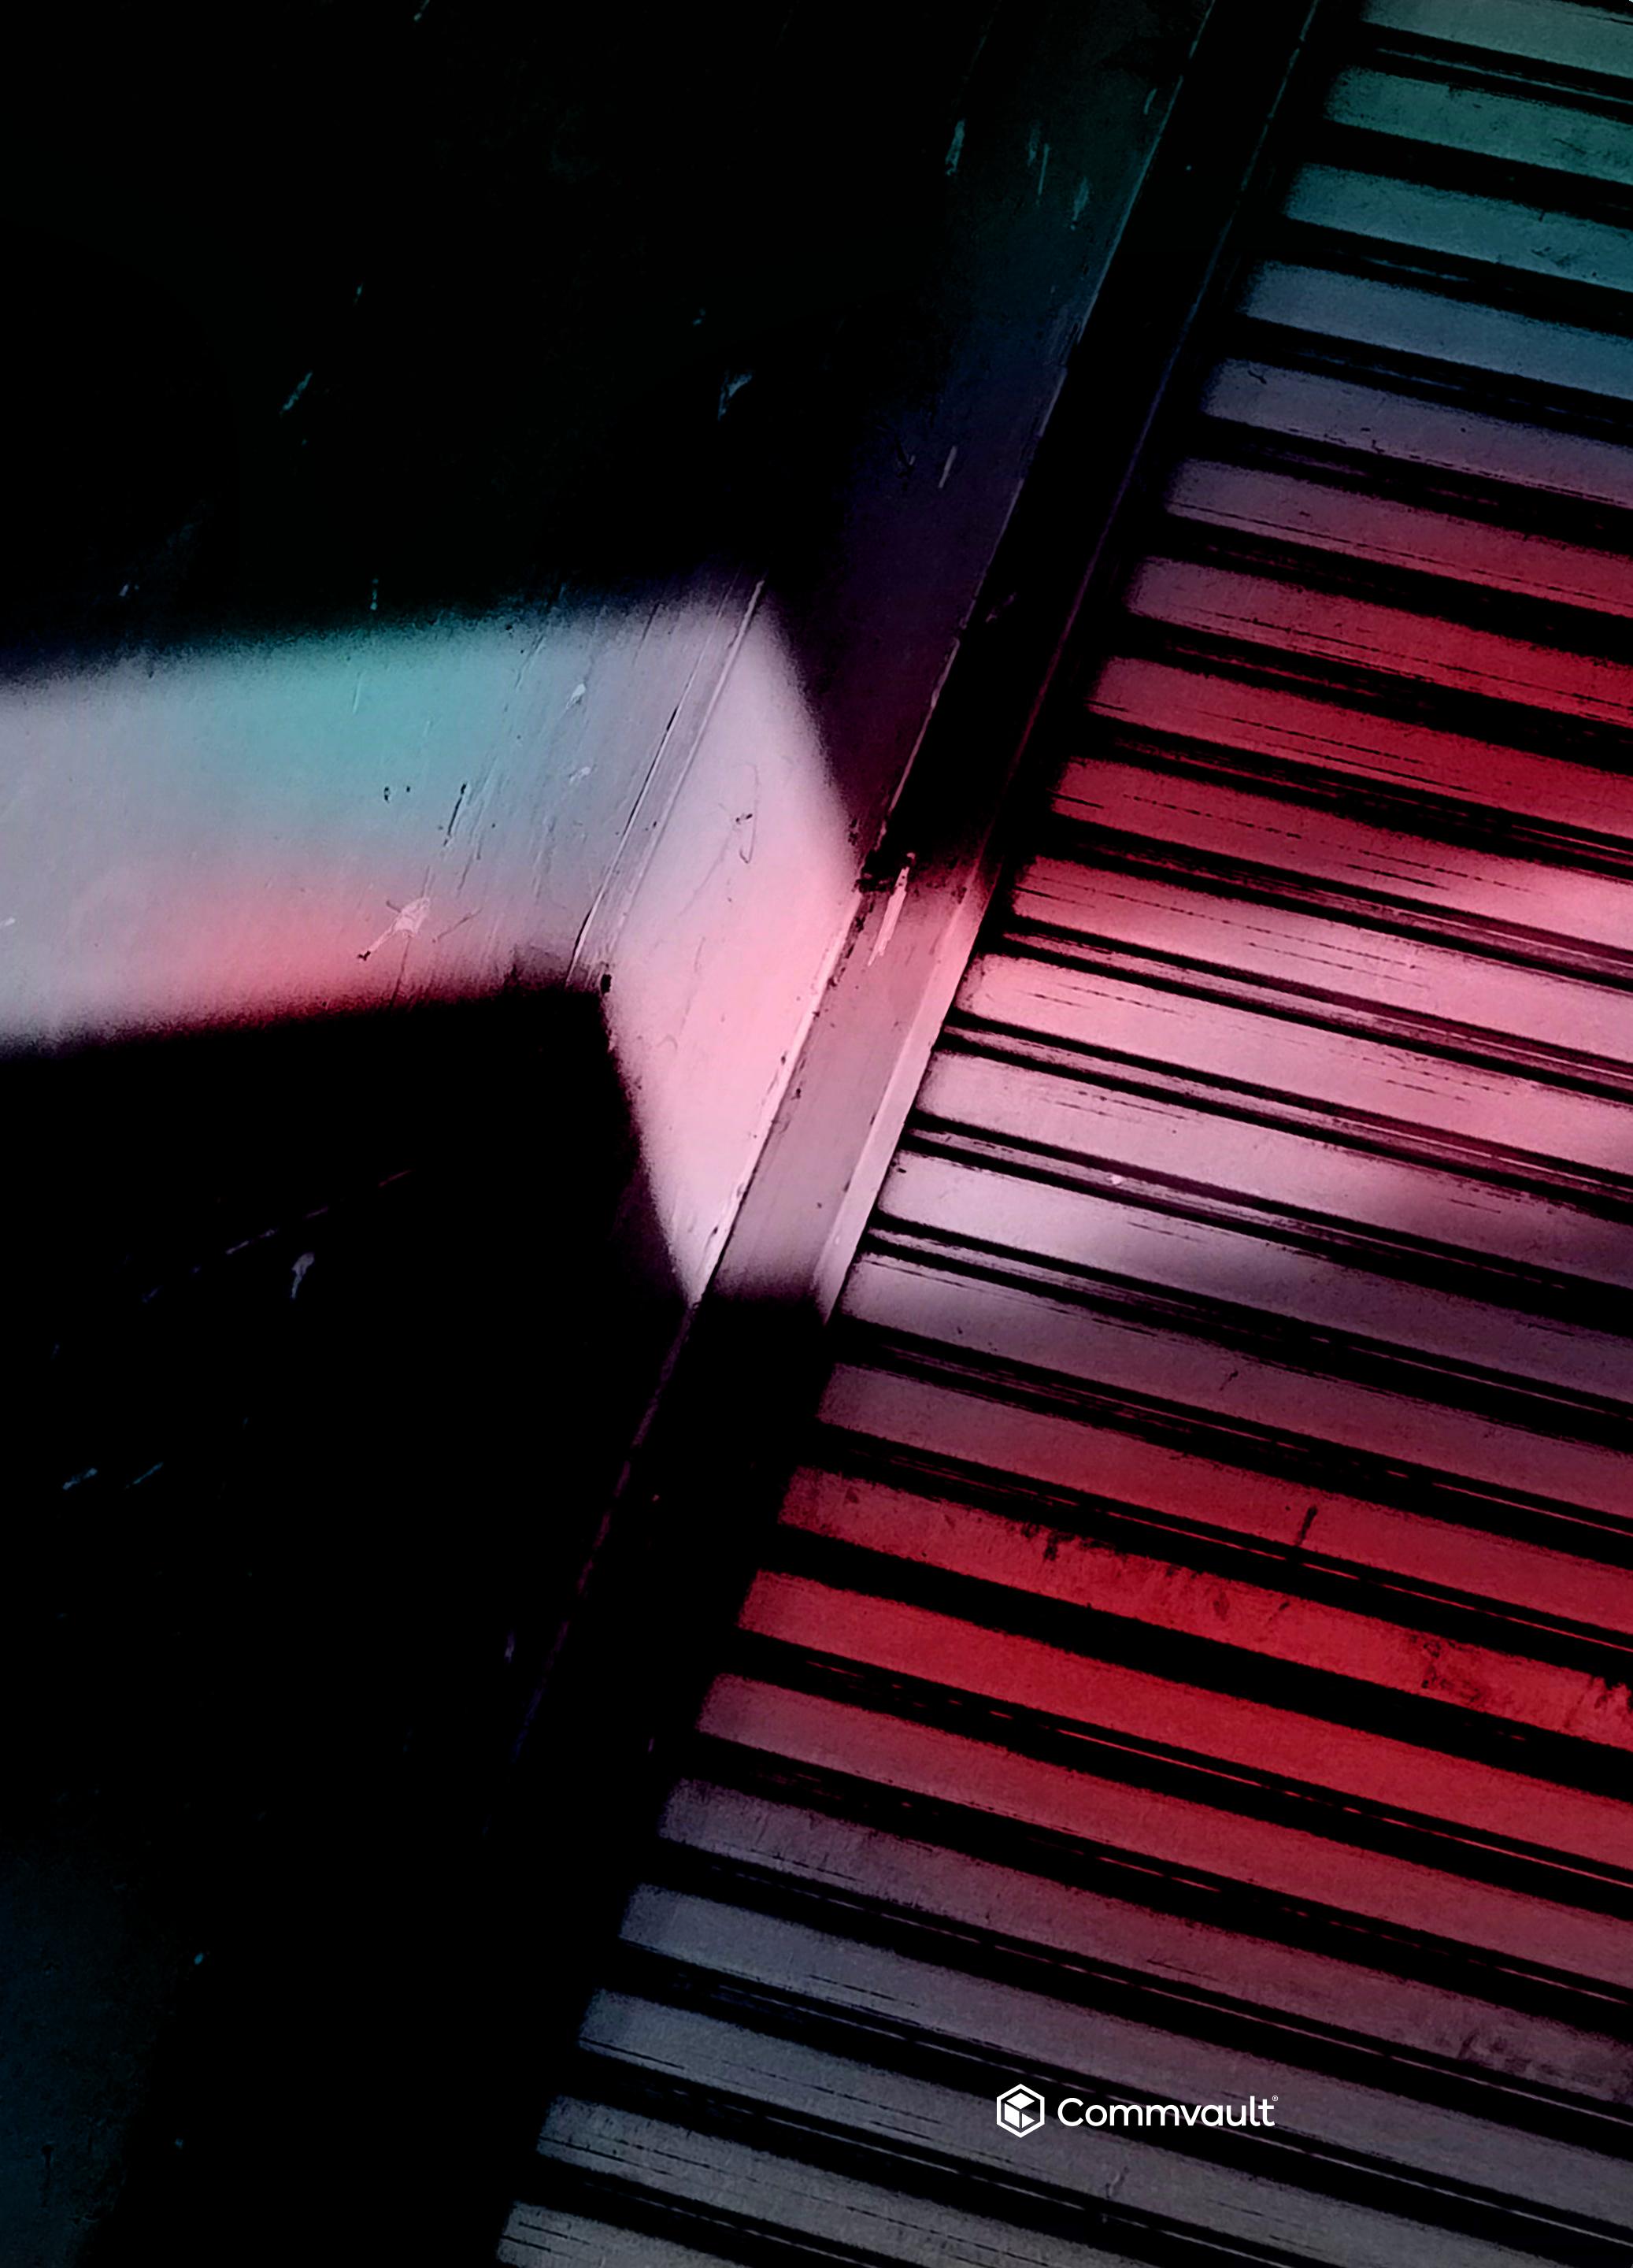


AI Security Risk Analysis:

What Changed with MCP 2.0
and What to Do Now

A Readiness Report for CIOs and CISOs

READINESS REPORT ISSUE 2 JANUARY 2026



Additional Resources

Visit readiverse.com/mcp to take our **Readiness Self-Assessment** and learn more about how to determine if MCP 2.0 controls are sufficient to protect your organization.

Join the Readiverse Community – Connect with CIOs and CISOs building resilient AI operations at readiverse.com.

This Readiness Report was prepared based on analysis of the MCP 2.0 specification and enterprise AI security requirements.

Published: January 1, 2026
Commvault Security Research Team

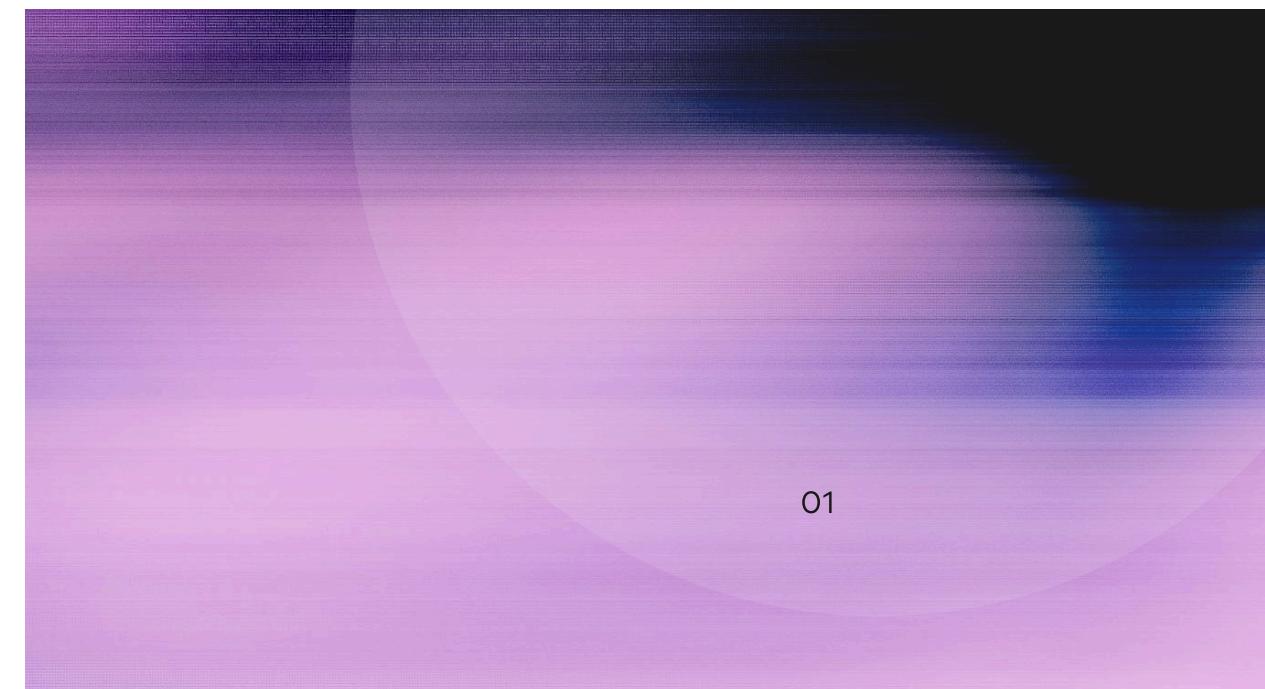
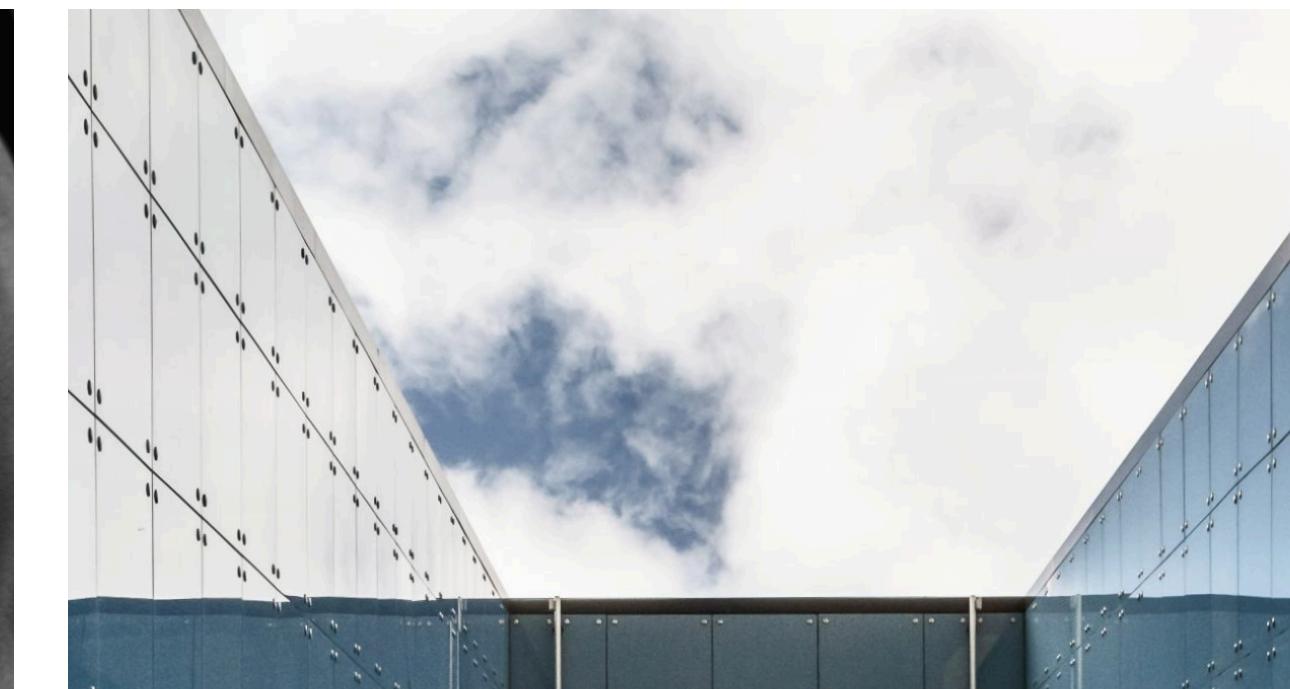
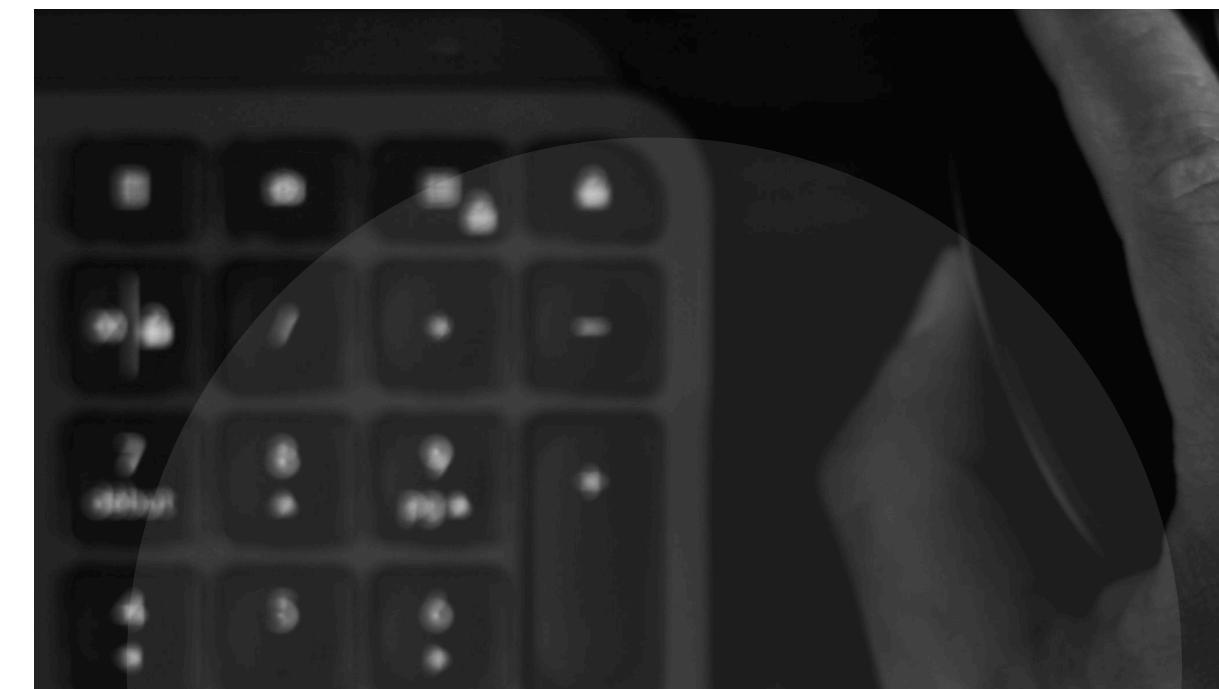
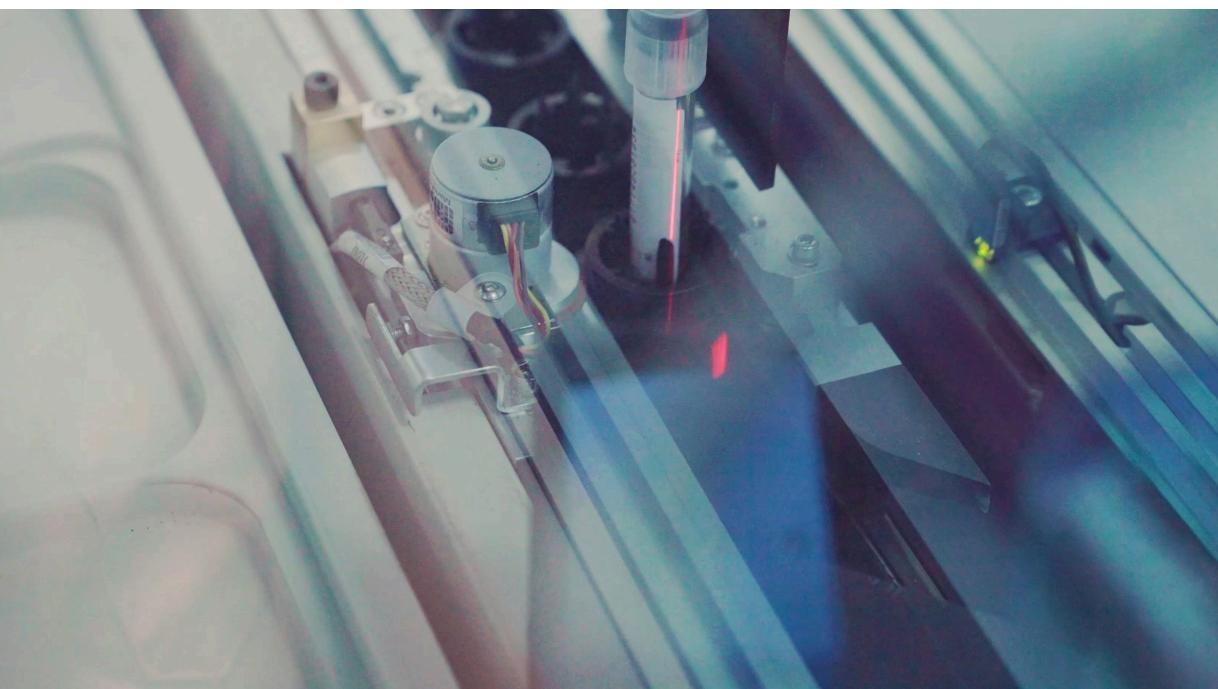
For practical tools to help you evaluate your readiness for MCP 2.0, including an assessment and a decision-making guide, visit [Readiverse.com/MCP](https://readiverse.com/MCP)

Executive Summary

Artificial intelligence has reached a critical operational milestone. AI systems now execute commands, call enterprise tools, and initiate workflows that directly affect business operations and security posture.

The Model Context Protocol (MCP) has emerged as the standard enabling these interactions, and Version 2.0 introduces the first structured governance framework for agentic AI. Three foundational controls – OAuth-based authorization, structured tool schemas, and elicitation workflows – transform AI agents from experimental systems into candidates for enterprise deployment.

These enhancements make AI actions auditable, behavior predictable, and deployment controllable. However, MCP 2.0 doesn't resolve all security challenges. Critical gaps remain in server identity, tool provenance, and runtime isolation. This report explains what changed, what remains unresolved, and how to prepare for governed AI agent deployment.



Why MCP 2.0 Matters Now

The Question Every Board Will Ask

"Are we ready to deploy AI agents that can take action on behalf of the company?"

This isn't a technology evaluation. It's a governance decision. MCP 2.0 provides the framework that makes an affirmative answer possible with appropriate controls.

The Deployment Timeline

Within 12 to 18 months, most large enterprises will have AI agents in production workflows¹: executing database queries, modifying configurations, managing workflows, and accessing sensitive information based on natural language instructions.

The critical question is whether these deployments will be governed from inception or retrofitted with security after incidents expose weaknesses.

Early movers that establish governance frameworks now will define operational standards. Late adopters will inherit technical debt that becomes increasingly expensive to remediate.

Regulatory Alignment

MCP 2.0's control mechanisms align with emerging regulatory requirements. The EU AI Act mandates documentation and auditability for high-risk AI systems. The Digital Operational Resilience Act (DORA) requires financial institutions to demonstrate control over automated operations. U.S. executive orders emphasize AI safety and transparency.

The auditability that MCP 2.0 enables (structured logging of authorization, tool invocation, and human oversight) provides the evidence base these frameworks require. Organizations implementing MCP governance now have the ability to position themselves ahead of compliance mandates.

The Competitive Dynamic

As the first high-profile AI incidents occur, market expectations will shift rapidly toward governed deployment. Organizations with mature AI governance frameworks will scale confidently while competitors pause to retrofit controls.

This isn't about being first to deploy AI agents. It's about being first to deploy them safely.

1. [AI Agents: The Final Frontier of the Enterprise](#), SnapLogic.

The Three MCP 2.0 Security Enhancements

1. Authorization Controls: Bounded Permissions

What Changed:

MCP 2.0 replaces implicit trust with explicit, scoped authorization. Each credential is bound to a specific system and cannot be reused across services. Credential sharing between systems is prohibited.

Why It Matters:

Before MCP 2.0, AI systems operated with broad, loosely defined permissions. A compromised credential could provide access across multiple systems. MCP 2.0 creates enforceable boundaries. Credentials for email access cannot query databases or modify cloud infrastructure.

Real-World Impact:

Before: An AI assistant's leaked token could access any connected system.
After: Credentials are scoped to specific services; a leaked credential has no database access.

Business Value:

Credential exposure becomes contained rather than catastrophic. When incidents occur, blast radius is designed to be limited to specific authorization scope.

The Three MCP 2.0 Security Enhancements

2. Structured Tool Schemas: Predictable Behavior

What Changed:

Every tool must define precise input and output schemas. The server validates all arguments before execution. AI models cannot generate free-form commands or invent parameters.

Why It Matters:

Schema enforcement eliminates injection vulnerabilities. Previous approaches allowed dynamic command construction, creating opportunities for malformed inputs and attacks. Structured validation makes tool behavior deterministic and testable.

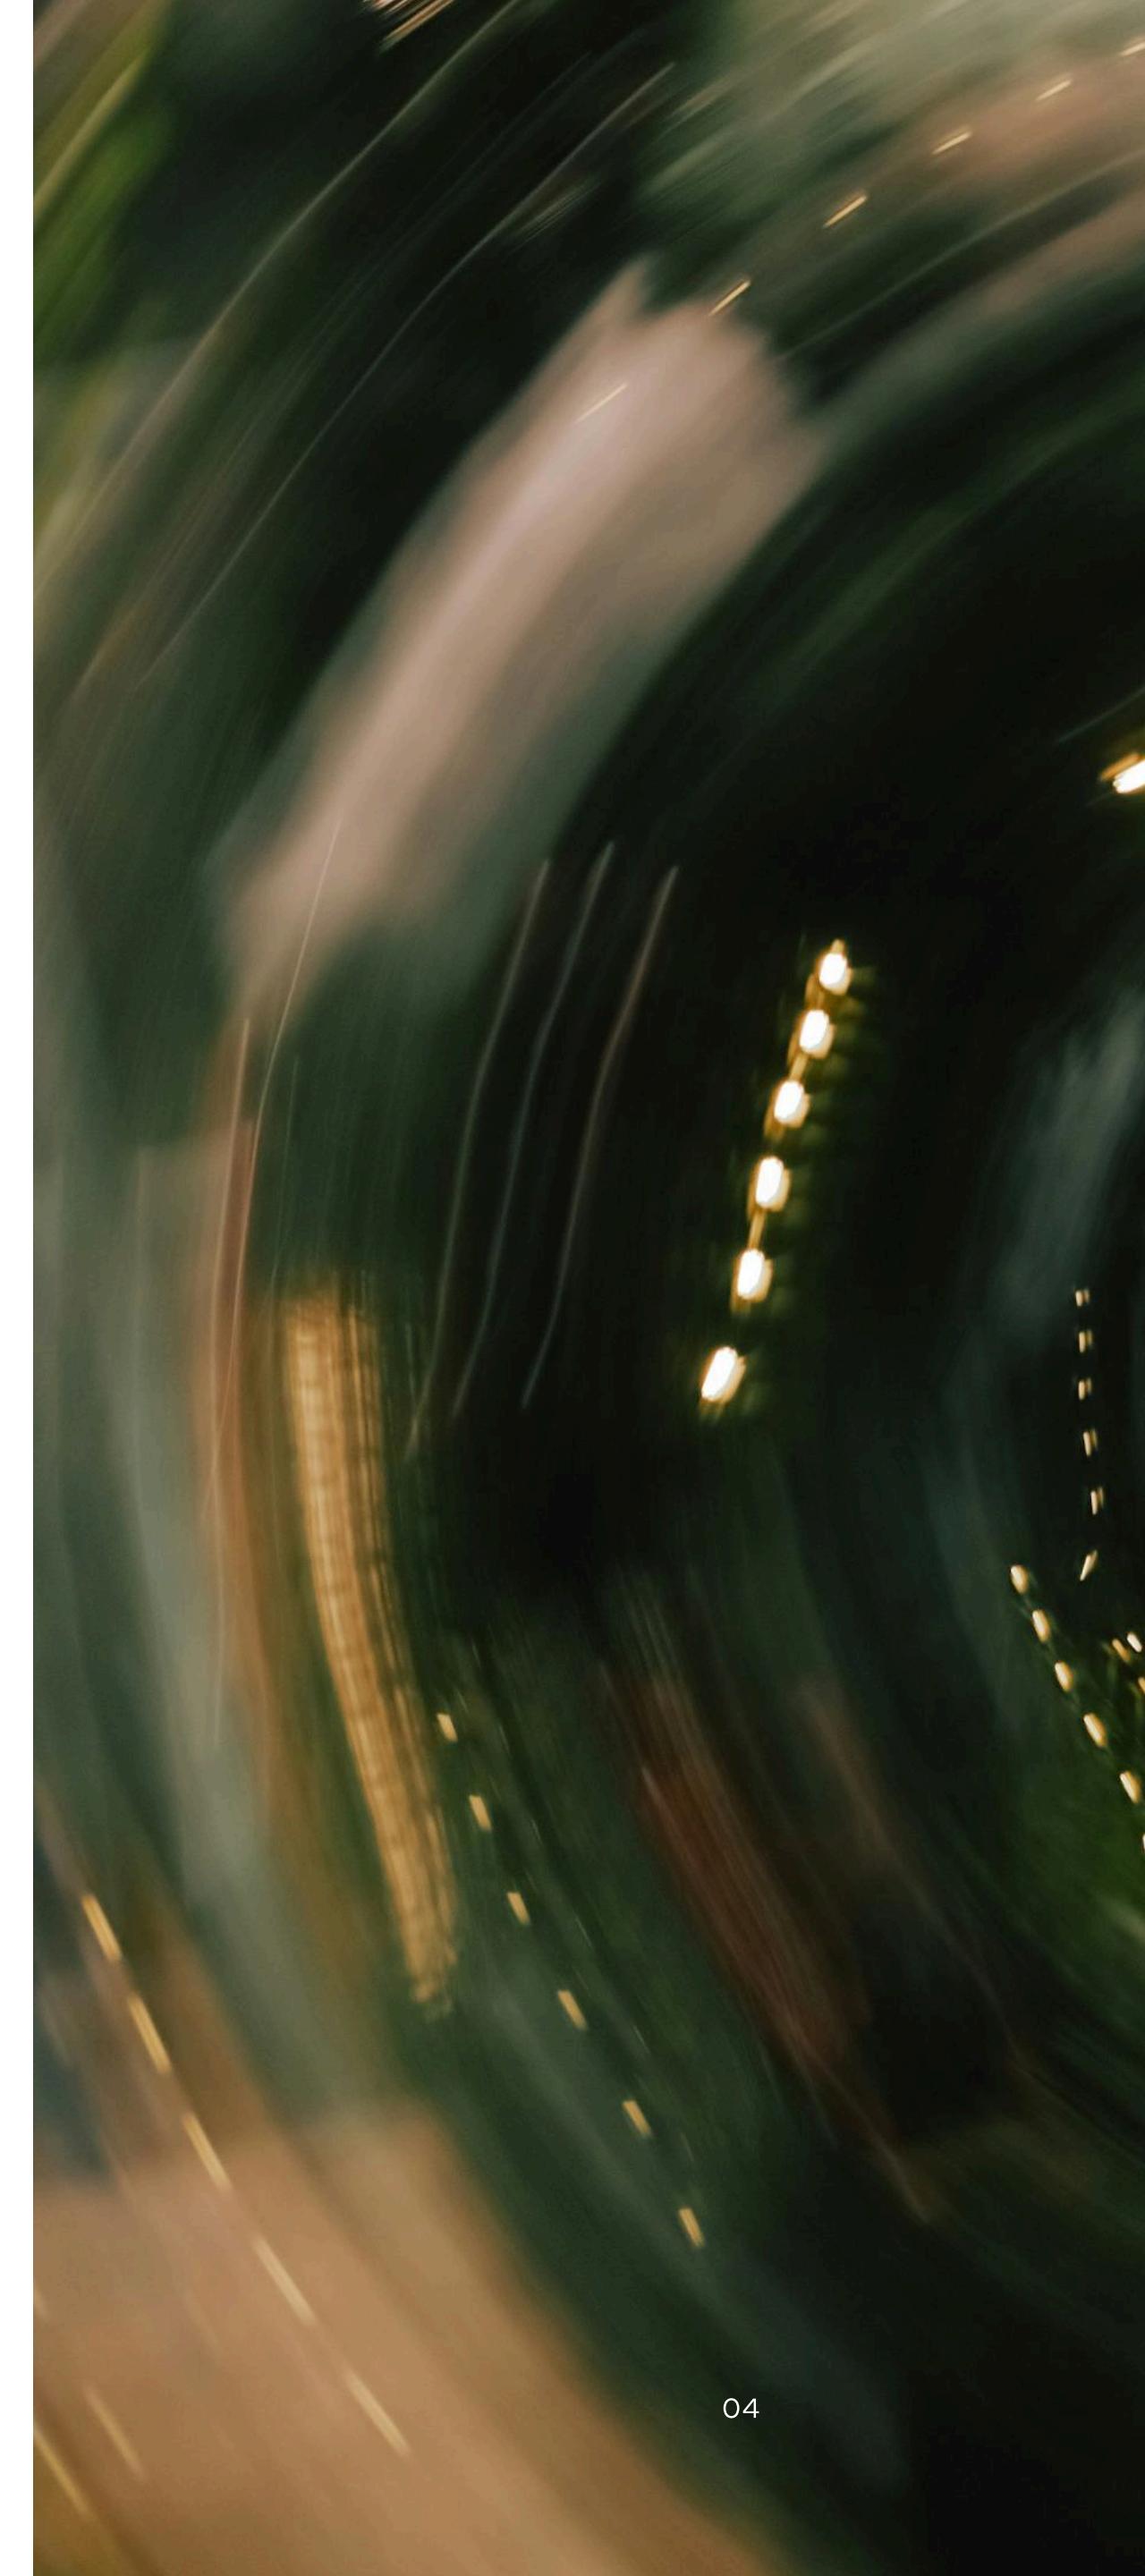
Real-World Impact:

Before: AI told to "delete old files" might execute broad deletion commands with unpredictable scope and irreversible results.

After: Schema requires explicit parameters: action type, specific directory path, retention period in days, and confirmation of dry-run mode before execution.

Business Value:

AI operations become auditable and repeatable. You can test behavior before production and demonstrate to auditors that AI follows defined, validated processes.



The Three MCP 2.0 Security Enhancements

3. Human-in-the-Loop Controls: Supervised Autonomy

What Changed:

MCP 2.0 pauses AI operations to request human input when information is missing, ambiguous, or requires approval. The system asks for clarification instead of guessing.

Why It Matters:

This transforms AI from autonomous systems that guess into supervised systems that seek validation. Organizations maintain human oversight for high-stakes decisions, while gaining efficiency for routine tasks.

Real-World Impact:

Before: AI asked to "prepare Q4 report" might hallucinate missing data or make incorrect assumptions.

After: System asks: "Which quarter (Q4 2024 or Q1 2025)? Include subsidiary data? Intended for board or investors?"

Business Value:

Efficiency gains with built-in checkpoints that help prevent costly errors. Users retain control over sensitive decisions while delegating routine execution. Creates audit trails showing explicit human authorization for all significant action.

What Remains Unsolved

MCP 2.0 represents significant progress, but understanding remaining governance gaps is essential for deployment planning. These limitations require compensating controls:

01. Server Identity Verification
The Gap: No mechanism to verify MCP server authenticity.
Impact: Exposure to impersonation attacks and credential interception.

What You Should Do: Deploy servers only on verified, controlled infrastructure. Implement network segmentation to limit exposure.

02. Tool Provenance and Supply Chain Security
The Gap: No built-in mechanism to verify tool authenticity or detect unauthorized modifications.
Impact: The MCP ecosystem operates like an app store without security vetting. Compromised tools can spread undetected across your environment.

What You Should Do: Maintain internal tool registries with manual validation. Require security review before any tool deployment. Implement integrity checking.

03. Runtime Isolation
The Gap: No built-in execution boundaries or privilege controls. Tools run with whatever access the host environment permits.

Impact: Compromised tools can access sensitive systems, move laterally across your network, or exfiltrate data.

What You Should Do: Run MCP servers in isolated environments. Restrict network access to only required services. Apply minimum necessary permission.

04. Prompt Manipulation and Logic Attacks
The Gap: While structured schemas help prevent direct command injection, attackers can manipulate AI decision-making through carefully crafted prompts or metadata.

Impact: Attackers can convince AI to execute unintended actions, access unauthorized data, or bypass security controls through seemingly legitimate requests.

What You Should Do: Validate all inputs before AI processing. Review AI-generated actions before execution. Monitor for unusual patterns in AI behavior.

05.

Legacy Tool Over-Privilege

The Gap: Existing tools retain broad permissions. MCP 2.0 doesn't force tool redesign or privilege reduction.

Impact: Legacy tools with excessive permissions remain dangerous when accessible to AI.

What You Should Do: Audit all tools for privilege scope. Refactor high-risk tools. Create tool-specific service accounts with minimal permissions.

06.

Multi-Agent Coordination

The Gap: MCP 2.0 governs individual tool calls but provides no guardrails for interactions between multiple AI agents.

Impact: Multiple AI agents can amplify each other's actions, create feedback loops, or produce unpredictable combined behaviors.

What You Should Do: Limit action frequency per agent. Implement automatic shutoffs for runaway processes. Monitor for unusual interaction patterns between agents.

07.

Observability and Anomaly Detection

The Gap: MCP 2.0 enables logging but provides no built-in mechanisms to detect unusual behavior or enforce policies in real time.

Impact: You can investigate incidents after they occur, but cannot proactively identify problems as they develop.

What You Should Do: Implement comprehensive logging with security monitoring integration. Establish baseline behavior profiles. Deploy anomaly detection.

Recommended Next Steps

Your immediate actions depend on where you are in AI adoption.

If You're Still Experimenting with AI

Your Advantage: Build governance into your foundation rather than retrofitting later.

Immediate Actions:

1. Adopt MCP 2.0 as standard from Day One.
2. Establish governance policies now: Define risk tiers, approval workflows, elicitation triggers.
3. Create an AI tool registry even if it contains few tools initially.

What to Avoid: Don't defer governance until "production scale." Early patterns become organizational defaults.

Timeline: Complete governance framework and pilot one MCP 2.0 deployment within 60 days.

If You're Moving AI to Production

Your Challenge: Need governance that doesn't break existing deployments while helping prevent future risk.

Immediate Actions:

1. Conduct AI capability audit: Document every system, access, and actions.
2. Classify systems using risk framework: Privilege × data sensitivity.
3. Address over-privilege immediately: Reduce to minimum necessary.
4. Migrate new deployments to MCP 2.0 as standard.
5. Plan staged migration for existing high-risk systems.

What to Avoid: Don't assume existing systems are "good enough" because incidents haven't occurred yet.

Timeline: Complete risk assessment within 30 days. Migrate high-risk systems within 90 days.

If You Already Have AI Agents Deployed

Your Reality: You're operating with AI governance debt. Remediation is urgent but must avoid operational disruption.

Immediate Actions:

1. Perform comprehensive security assessment across all deployments.
2. Implement compensating controls for systems that cannot immediately migrate.
3. Integrate AI incident response plan into company incident response procedures.
4. Prioritize migration ruthlessly – focus on admin privileges and regulated data access.
5. Establish continuous monitoring to help with anomaly detection.

What to Avoid: Don't let "it's working now" delay action. Gaps compound as you scale.

Timeline: Emergency assessment within 14 days. Compensating controls within 30 days. Migration plan with 90-day target for high-risk systems.

Conclusion

MCP 2.0 establishes the first structured governance framework for AI agents operating in enterprise environments. Its three core controls – authorization boundaries, structured schemas, and human-in-the-loop workflows – create the accountability and predictability required for production deployment.

These capabilities represent meaningful progress, but implementation requires recognizing both what MCP 2.0 solves and what remains unaddressed. The structural gaps demand compensating controls, careful planning, and ongoing vigilance.

Organizations that establish strong AI governance frameworks now, while AI agent deployment is still early, will define operational standards for their industries. Those that defer will risk accumulating technical debt and compliance risk that becomes expensive to remediate.

The opportunity is clear: Build the architecture that makes AI agents an asset rather than a liability. The timeline is compressed: Most enterprises will have AI agents in production within 18 months. Governance must precede deployment, not follow it.

MCP 2.0 creates the foundation. Your implementation determines the outcome.

Additional Resources

Visit readiverse.com/mcp to take our **Readiness Self-Assessment** and learn more about how to determine if MCP 2.0 controls are sufficient to protect your organization.

Join the Readiverse Community – Connect with CIOs and CISOs building resilient AI operations at readiverse.com.

This Readiness Report was prepared based on analysis of the MCP 2.0 specification and enterprise AI security requirements.

Published: January 1, 2026
Commvault Security Research Team