AGENTIC COMMERCE / RESEARCH BRIEF

MCP + A2A Integration Guide

An architecture guide for connecting agents to commerce systems without exposing uncontrolled tool access.

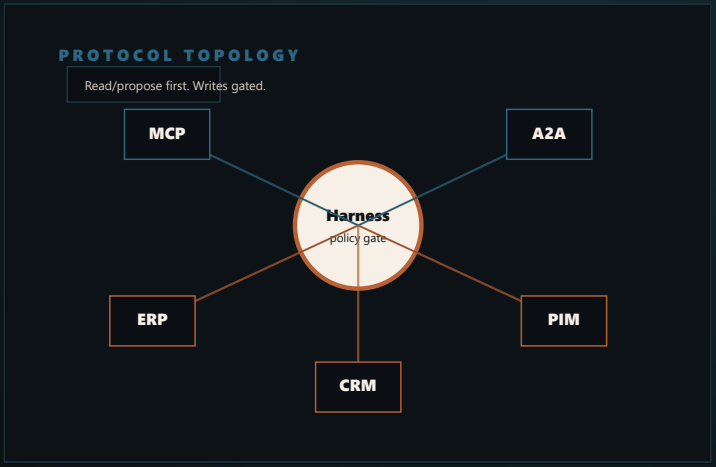

01	Read first
02	Propose second
03	Writes gated

PROTOCOL GOVERNANCE

Connection speed increases permission risk.

MCP and A2A can make integrations faster, but they also make tool boundaries more important. Commerce agents should not inherit broad API power just because a protocol makes connection easy. The integration layer should be designed as governed access: narrow schemas, risk-classed tools, identity-bound permissions, budget limits, and audit events around every agent-to-tool or agent-to-agent interaction.

01
The protocol is not the governance system.

02
Every exposed tool needs a commercial action boundary.

03
Agent-to-agent collaboration should not bypass human approval.

Protocol permission gradient

Protocol adoption should move from read to propose before write or payment capabilities exist.

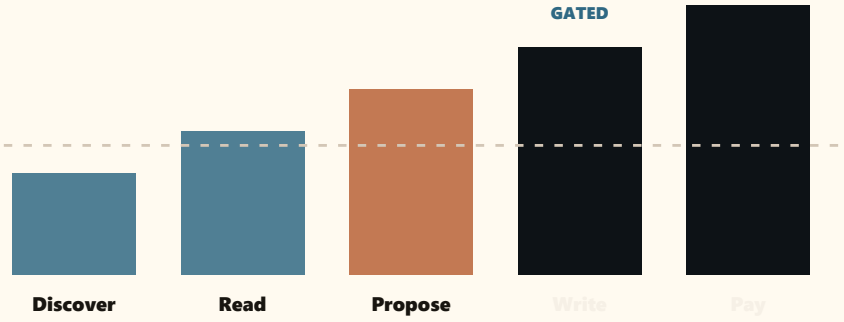

MCKINSEY TECHNOLOGY / QUANTUMBLACK / BUILDING THE FOUNDATIONS FOR AGENTIC AI AT SCALE

Stable interfaces and shared data meaning are prerequisites for scaling agentic AI.

OECD / THE AGENTIC AI LANDSCAPE AND ITS CONCEPTUAL FOUNDATIONS

Agentic systems require clearer language around autonomy, coordination, tool use, delegation, and operating environment.

PARLOA / GLOBAL AI AGENT MANAGEMENT & COMPLIANCE UPDATES

Agent management platforms are moving audit logs and accountability into product-level controls.

MCP/A2A RISK TOPOLOGY

Protocol readiness controls.

Protocol adoption sequence: discover, read, propose, approve, audit. Public execution stays disabled until controls are proven. Read the artifact as a protocol risk topology: discovery, read, proposal, write, and payment actions need different permission and audit semantics.

MCP/A2A risk topology

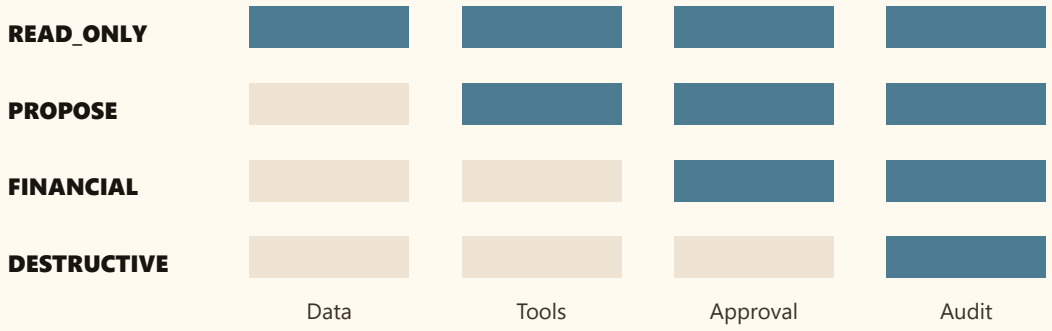

DISCOVERY
Let agents discover structured capabilities only after purpose, owner, and risk class are known.

READ
Expose read tools before proposal or write tools, with narrow schema and source visibility.

PROPOSE
Allow recommendation generation while blocking system-changing actions until approval.

AUDIT
Log every tool call, agent handoff, approval state, failure, and terminal workflow result.

INTEGRATION REVIEW

Expose the smallest useful protocol surface first.

Use the review to separate read-only tools from proposal and write paths, then prove revocation, denial, and audit behavior.

INSPECTION QUESTIONS

- Which MCP tools are read-only, proposal-only, financial, or destructive?
- Can the buyer revoke one agent without breaking the whole integration?
- What A2A handoffs are allowed, and where is the audit event?
- Which public MCP or checkout actions remain disabled before legal and budget review?

DO NOT SHIP UNTIL

- Do not equate MCP availability with permission to execute.
- Do not let A2A chains hide accountability.
- Do not expose payment, refund, discount, or order submission before explicit review.

MAP	List tool surfaces, source systems, risk class, owners, and revocation requirements.	Protocol readiness matrix.
LIMIT	Expose one read-only tool path and one proposal-only path before any live write.	Tool schemas, allowlists, and denied-action tests.
REVIEW	Run integration scenarios with blocked write actions and audit review.	Scenario log and approval-gate evidence.

SOURCE LEDGER

Source ledger and protocol boundary

The brief combines external research signals with HARNEXA's implementation boundary: identity, permission, approval, audit, CLEAR evaluation, and no autonomous public execution.

MCKINSEY TECHNOLOGY / QUANTUMBLACK

Building the foundations for agentic AI at scale

McKinsey frames data quality, shared meaning, stable interfaces, observability, and controlled execution as scale requirements for agentic AI.

https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/building-the-foundations-for-agentic-ai-at-scale

OECD

The agentic AI landscape and its conceptual foundations

The OECD distinguishes AI agents from agentic AI by coordination, task decomposition, delegation, sustained operation, and operation in less predictable environments.

https://www.oecd.org/content/dam/oecd/en/publications/reports/2026/02/the-agentic-ai-landscape-and-its-conceptual-foundations_a9d4b451/396cf758-en.pdf

PARLOA

Global AI Agent Management & Compliance Updates

Parloa highlights centralized audit logs that show who changed what and when as agent programs scale across teams and markets.

https://www.parloa.com/blog/parloa_product_release_2026/

RESEARCH STANCE

External research is cited as market signal; HARNEXA interpretation is kept separate from source claims.

BUYER ARTIFACT

Every report includes a concrete visual artifact that can be inspected in a sales, risk, or architecture review.

OPERATING BOUNDARY

Read-first MCP, proposal-only A2A, gated writes, auditable handoffs. The asset is for governed review, not autonomous production execution.

HARNEXA BOUNDARY

Technical governance aid only. Not legal advice or compliance certification. HARNEXA AI builds audit-ready foundations for qualified review. No PDF in this library claims EU AI Act certification, legal advice, payment readiness, or autonomous production execution.