

# Model Context Protocol (MCP)

Expose enterprise capabilities as discoverable tools

## 01 / WHAT IT IS

An interoperability protocol through which an AI host uses an MCP client to discover and invoke tools, access resources, and obtain reusable prompts from an MCP server.

## 02 / WHY IT MATTERS

A useful enterprise application should not require users to leave ChatGPT or Claude, reopen another UI, and manually repeat a workflow. MCP reduces bespoke integrations while preserving controlled access to backend capabilities.

## 03 / HOW TO IMPLEMENT

Wrap approved Question Assistant APIs as typed MCP tools (for example, `list_assessments`, `get_assessment`, `start_assessment`, `review_answers`). Validate inputs, check identity and authorization at every call, enforce idempotency for writes, and return structured results and job identifiers.

## 04 / QUESTION ASSISTANT · STATUS + METRICS

The Question Assistant retained its own UI/UX, while its APIs were exposed through an MCP connector. Users could initiate questionnaires from ChatGPT and Claude and view returned results there. This integration was implemented; the exact tool names shown here are illustrative.

**Measures:** Successful tool-call rate; median launch latency; permission-denial correctness; response completeness; failed/duplicate launches.

## 05 / INDUSTRY TOOLS & PLATFORMS

MCP specification; modelcontextprotocol SDKs; Salesforce MCP client (for future reuse)

**Design caution:** An MCP transport is not an authorization system. Server-side checks and business rules remain mandatory.

# Connectors & Plugins

Turn access and workflows into a reusable installation

## 01 / WHAT IT IS

A connector configures service reachability and credentials. A plugin packages related capabilities—such as Skills and MCP configuration—so users or organizations can discover and reuse them. Exact formats and UI vary by host.

## 02 / WHY IT MATTERS

A working integration is not yet a manageable product. Packaging improves onboarding, documentation, versioning, capability discovery, ownership and workspace-level governance.

## 03 / HOW TO IMPLEMENT

Create one Question Assistant plugin with two Skills: Questionnaire Workspace and Question Bank Administration. Attach the existing MCP-backed capabilities, define connection setup and permitted actions, ship setup documentation, and test supported host surfaces separately.

## 04 / QUESTION ASSISTANT · STATUS + METRICS

The Question Assistant had its MCP connector and a user-focused Skill. Packaging them as a distributable plugin and adding the Administrator Skill is the proposed next phase—not an assertion that such a plugin was deployed at Fortress.

**Measures:** Time to connect; successful installation rate; active users; onboarding support requests; revoked/expired connections detected.

## 05 / INDUSTRY TOOLS & PLATFORMS

OpenAI Plugins; ChatGPT connected apps; Claude connectors; Salesforce Agentforce Registry

**Design caution:** A plugin is a distribution construct, not a substitute for backend authorization or a universal package format across vendors.

# Agent Skills

Give AI applications repeatable domain expertise

## 01 / WHAT IT IS

Reusable instructions, examples and optional files that teach an agent how to complete a task using available capabilities. A Skill guides the workflow; the MCP server performs live reads and writes.

## 02 / WHY IT MATTERS

Raw APIs force every model interaction to reinvent business workflow logic. Skills encode sequencing, exception handling and output expectations in a reusable, inspectable form.

## 03 / HOW TO IMPLEMENT

Skill A (Questionnaire Workspace): dashboard, list, filter, start, status, review answers. Skill B (Administration, proposed): browse question-bank library; add, edit, clone, retire questions; create or clone 250-, 150- and 75-question templates; draft and publish versions with approval.

## 04 / QUESTION ASSISTANT · STATUS + METRICS

A user-facing Question Assistant Skill was implemented for ChatGPT and Claude. The second, admin-focused Skill is a proposed expansion. It must not expose destructive operations merely because the Skill instructions mention them.

**Measures:** Task completion; correct skill selection; tool-call efficiency; admin change errors; drafts approved versus rolled back.

## 05 / INDUSTRY TOOLS & PLATFORMS

Anthropic Agent Skills; OpenAI plugin Skills; skill.md patterns

**Design caution:** Human approval, role checks and version-safe data changes belong in backend policy, not solely in prompt instructions.

# Agent2Agent Protocol (A2A)

Delegate complex tasks between independent agents

## 01 / WHAT IT IS

A standard for agent discovery and task-oriented interaction using concepts such as Agent Cards, messages, tasks, status updates and artifacts.

## 02 / WHY IT MATTERS

An independent supply-chain agent may need to delegate a complete vendor assessment—not individually orchestrate every questionnaire API call.

## 03 / HOW TO IMPLEMENT

Add an A2A-facing Question Assistant Agent with a documented Agent Card, supported tasks, authentication policy, input/output schemas and lifecycle handling. Agentforce or another agent could delegate assessments and receive progress and final artifacts.

## 04 / QUESTION ASSISTANT · STATUS + METRICS

A2A was not implemented in the original Question Assistant. The proposed demo uses an Agentforce supply-chain agent that asks a new Question Assistant Agent to perform a vendor-risk assessment and return a structured summary.

**Measures:** Task handoff reliability; delegation latency; retry and resume behavior; artifact quality; unauthorized handoff rejection.

## 05 / INDUSTRY TOOLS & PLATFORMS

A2A protocol specification; Google Agent Development Kit; compatible agent frameworks

**Design caution:** An existing MCP endpoint does not automatically provide an A2A interface. Agentforce A2A product support and deployment details must be validated for the intended environment.

# Agent Orchestration & Planning

Coordinate roles, tool calls, decisions and handoffs

## 01 / WHAT IT IS

The runtime or workflow logic that sequences or coordinates agents, tools, validation steps and human decisions; it need not rely on any particular multi-agent framework.

## 02 / WHY IT MATTERS

A reliable result requires explicit responsibilities and completion criteria; multiple agents alone do not guarantee consistent outcomes.

## 03 / HOW TO IMPLEMENT

Define a state machine: receive questionnaire → retrieve evidence → Extractor drafts answers → Validator checks claims → Evaluator scores quality → human review for exceptions → publish. Add escalation, retry and reconciliation rules.

## 04 / QUESTION ASSISTANT · STATUS + METRICS

The historical Question Assistant used distinct Extractor, Validator and Evaluator LLM roles with human oversight. Reimplementing that pattern using LangGraph, CrewAI or another general-purpose orchestration framework is an optional architectural enhancement, not a verified historical tool choice.

**Measures:** Pass-through accuracy; disagreement rate; human-review workload; total assessment time; error recovery.

## 05 / INDUSTRY TOOLS & PLATFORMS

LangGraph; CrewAI; Google ADK; OpenAI Agents SDK

**Design caution:** The roles are confirmed; the precise historical orchestration implementation has not been independently verified.

# Context Engineering & Memory

Put the right evidence and state in the right context

## 01 / WHAT IT IS

Selection, shaping and lifecycle management of the evidence, definitions, conversation state and permitted retained memory supplied to each agent step.

## 02 / WHY IT MATTERS

Long prompts with mixed-quality sources dilute reasoning, expose irrelevant data and make successive agent steps inconsistent.

## 03 / HOW TO IMPLEMENT

Use permission-scoped retrieval, source provenance, semantic definitions, structured state and explicit memory policies. Keep tenant boundaries intact, set retention periods and distinguish durable assessment records from conversational memory.

## 04 / QUESTION ASSISTANT · STATUS + METRICS

The existing Question Assistant necessarily processed questionnaire content and evidence, but a dedicated persistent-memory architecture has not been confirmed. The animation presents a proposed approach that also connects to our Beyond RAG semantic-layer article.

**Measures:** Relevant-context recall; source citation precision; cross-tenant leakage rate; token cost per validated answer; stale-context incidents.

## 05 / INDUSTRY TOOLS & PLATFORMS

LangGraph stores/checkpoints; LangChain; Cube semantic layer; Mem0 (optional)

**Design caution:** A semantic layer, RAG store, workflow state and long-term agent memory are complementary but distinct components.

# Agent-User Interaction (AG-UI)

Make agent progress, approvals and results visible

## 01 / WHAT IT IS

An event-based way for an agentic backend and user-facing application to exchange state updates, messages, tool activity and approvals.

## 02 / WHY IT MATTERS

Business users need to know which questions are running, which answers require review, what the agent has done and when they must intervene.

## 03 / HOW TO IMPLEMENT

Stream questionnaire counts, progress events and review requests to a dashboard. Add approval interrupt/resume interactions and evidence links. Use AG-UI where its protocol fits; a bespoke UI can also satisfy these needs.

## 04 / QUESTION ASSISTANT · STATUS + METRICS

The original Question Assistant had its own UI/UX. Adoption of the AG-UI protocol itself is a proposed enhancement; it is not a historical implementation claim. Host-native ChatGPT or Claude rendering may require platform-specific UI facilities.

**Measures:** Time to recognize a blocked assessment; approval completion rate; user comprehension; UI event latency; accessibility errors.

## 05 / INDUSTRY TOOLS & PLATFORMS

AG-UI protocol; CopilotKit; platform-native interactive app/extension facilities

**Design caution:** AG-UI is a possible frontend protocol, not mandatory for dashboards; clients differ in UI capabilities.

# Durable Execution & Events

Keep assessments running when things fail or pause

## 01 / WHAT IT IS

Persisted workflow execution with checkpoints, retries, event handling, asynchronous status updates and controlled recovery.

## 02 / WHY IT MATTERS

An assessment with 250 questions may take substantial time or await approvals. A timeout or model failure should not silently restart the whole process or duplicate records.

## 03 / HOW TO IMPLEMENT

Assign each assessment a stable ID; checkpoint progress at safe boundaries; retry transient failures; enforce idempotency for create/publish calls; queue slow work; support cancellation, resumption and audit-ready event histories.

## 04 / QUESTION ASSISTANT · STATUS + METRICS

Question Assistant historically supported questionnaire statuses, but the particular durable-execution framework is not confirmed. A Temporal-backed or equivalent future architecture is illustrative.

**Measures:** Workflow recovery rate; duplicates prevented; time to resume; stuck assessments; 95th-percentile end-to-end completion time.

## 05 / INDUSTRY TOOLS & PLATFORMS

Temporal; LangGraph persistence; Salesforce Flow; queue/event infrastructure

**Design caution:** “Exactly once” end-to-end is difficult across systems; use idempotent effects and explicit reconciliation rather than promising perfect delivery.

# Identity, Discovery & Authorization

Enforce least privilege across every boundary

## 01 / WHAT IT IS

Identity assertion, discovery, credential handling, policy enforcement and auditable authorization for humans, service accounts, agents, tools and delegated actions.

## 02 / WHY IT MATTERS

A convenient Skill should not allow an ordinary analyst to edit the question bank; an agent should not inherit more authority than the calling user or service.

## 03 / HOW TO IMPLEMENT

Use trusted authentication, explicit scopes and tenant binding; approve MCP tools per role; protect admin writes; require confirmation for destructive changes; log the actor, tool, purpose and outcome. For A2A, validate agent identity and delegation rights as well.

## 04 / QUESTION ASSISTANT · STATUS + METRICS

The User Skill and future Admin Skill make role separation concrete: business users may launch/review; authorized admins may draft and publish templates. No specific historical SSO/RBAC implementation is claimed here.

**Measures:** Unauthorized-action success rate (target zero); privileged-tool exposure; audit coverage; credential rotation compliance; time to revoke access.

## 05 / INDUSTRY TOOLS & PLATFORMS

OAuth/OIDC; Okta/Auth0; Open Policy Agent; Salesforce Agentforce policies

**Design caution:** A model's reasoning or a Skill description must never be the sole authority for an administrative action.