

Agentic AI Stack — Master Architecture

From goal to real-world impact — securely, modularly, and at enterprise scale

1. USER & APPLICATION LAYER

Where people and systems set goals and see results



End Users
(Web, Mobile, Chat)



Agentic Applications
(Copilots, Assistants, Automation)



Business Systems
(CRM, ERP, ITSM, etc.)

2. ORCHESTRATION LAYER

Coordinates agents, tools, data, and workflows to achieve goals



Agent Orchestrator
(Planning, Routing, Memory, Guardrails)



Multi-Agent Collaboration
(Specialized Agents, Handoffs, Consensus)

3. TOOL & CONNECTOR LAYER

Gives agents safe, standardized access to external capabilities



MCP Connectors
(Tools, Data, Systems)



APIs & SDKs
(REST, GraphQL, SDKs)



Prebuilt Tools
(Search, Browser, Code, etc.)

4. DATA & KNOWLEDGE LAYER

Provides enterprise data, context, and memory for informed actions



Enterprise Data Sources
(Documents, Databases, Business Systems)



Knowledge Retrieval
(RAG, Vector Search, Semantic Indexing)



Agent Memory
(Short-term, Long-term, User Preferences)

5. SECURITY & GOVERNANCE LAYER

Ensures safety, compliance, and responsible AI use



Identity & Access
(SSO, RBAC, Permissions)



Policy & Guardrails
(Safety, Compliance, DLP)



Monitoring & Observability
(Logs, Traces, Analytics)



Audit & Reporting
(Usage, Compliance, Risk)

6. INFRASTRUCTURE LAYER

Scalable foundation for enterprise deployment



Compute
(Cloud, On-Prem, Hybrid)



Model Layer
(Commercial, Open Source, Fine-tuned)



Storage
(Data, Vectors, Logs)



Network & Runtime
(Containers, Kubernetes, Scaling)

MCP vs. Direct API — What's the Difference?

MCP (Model Context Protocol)
Standardized agent-to-tool integration

Direct API (REST, GraphQL, etc.)
Traditional point-to-point integration



Standard protocol for AI agents to discover and use tools

Custom integration for each service



Built-in metadata (tools, schemas, permissions, context)

You define and maintain the interface



Works across multiple LLMs (ChatGPT, Claude, etc.)

Client-specific development and maintenance



Dynamic tool discovery and context sharing

Static endpoints and manual configuration



Supports governed access, scopes, and auditability

Security, permissions, and logging must be built separately



Ideal for agents and dynamic workflows

Great for simple or custom integrations

MCP doesn't replace APIs — it standardizes how agents find, understand, and safely use them. The underlying services still use APIs.

Agent Prompt Administration

BUILT • Question Assistant Admin UI

Administrators can view and modify the master prompts for all three agents:

- **Extractor**
- **Validator / QA**
- **Evaluator**

Prompts are managed outside agent code through the application's admin interface.

RECOMMENDED • Additional safeguards

Version history · approvals · regression tests
Rollback · audit trails · access restrictions



Trusted



Connected



Intelligent



Real-World Impact