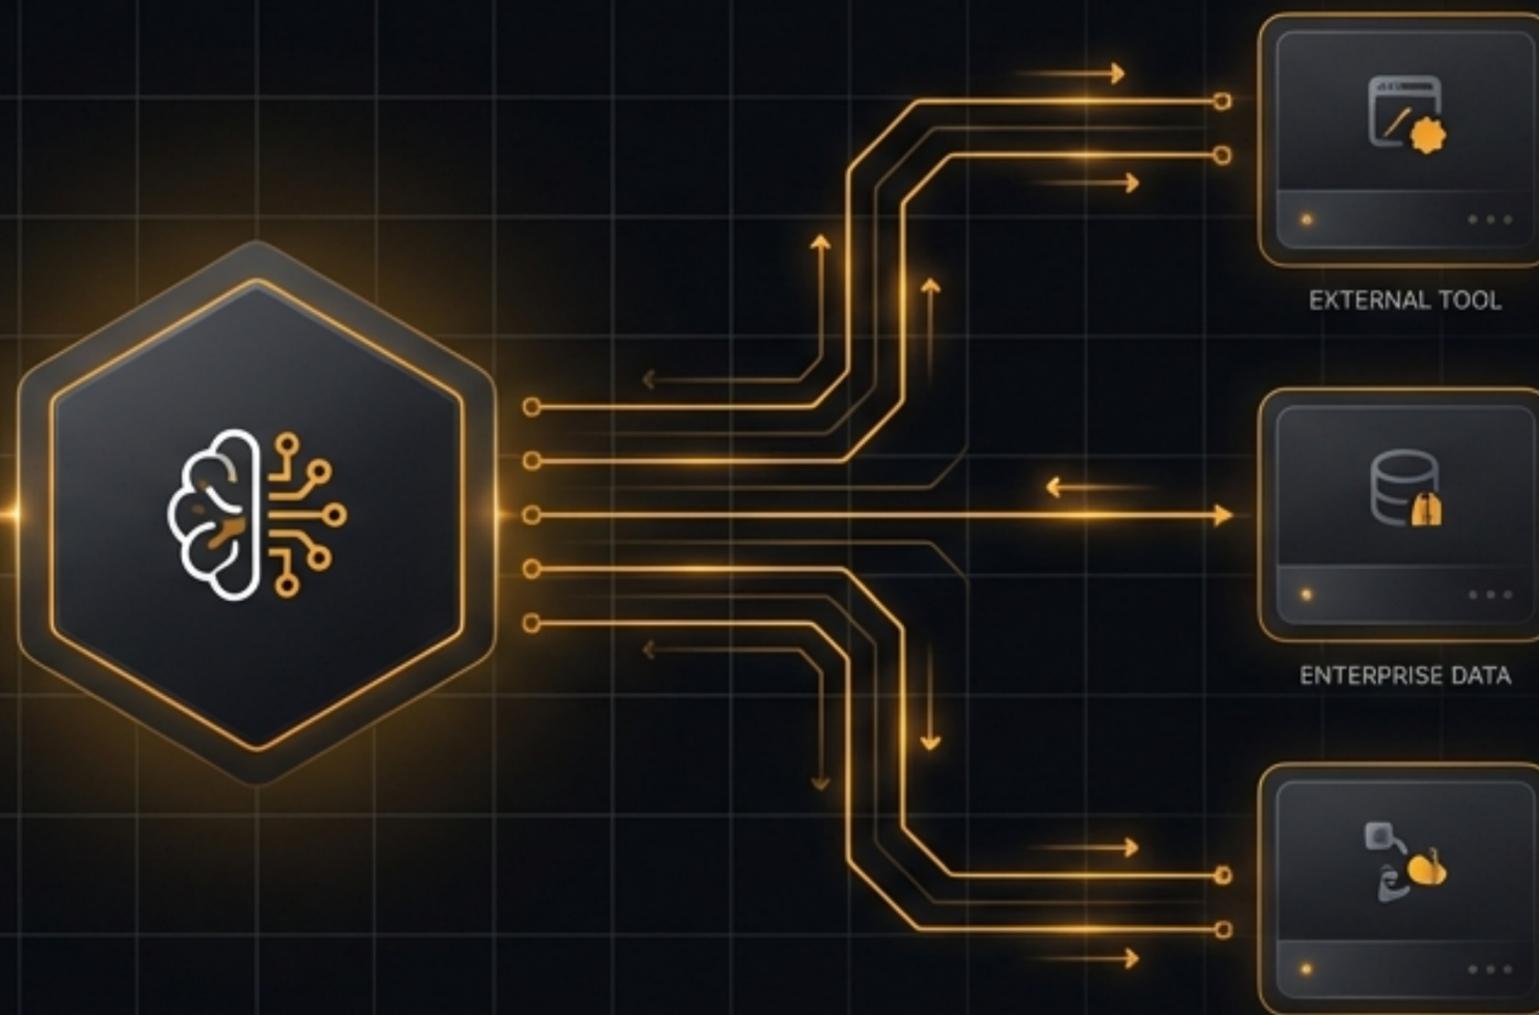


# The Model Context Protocol

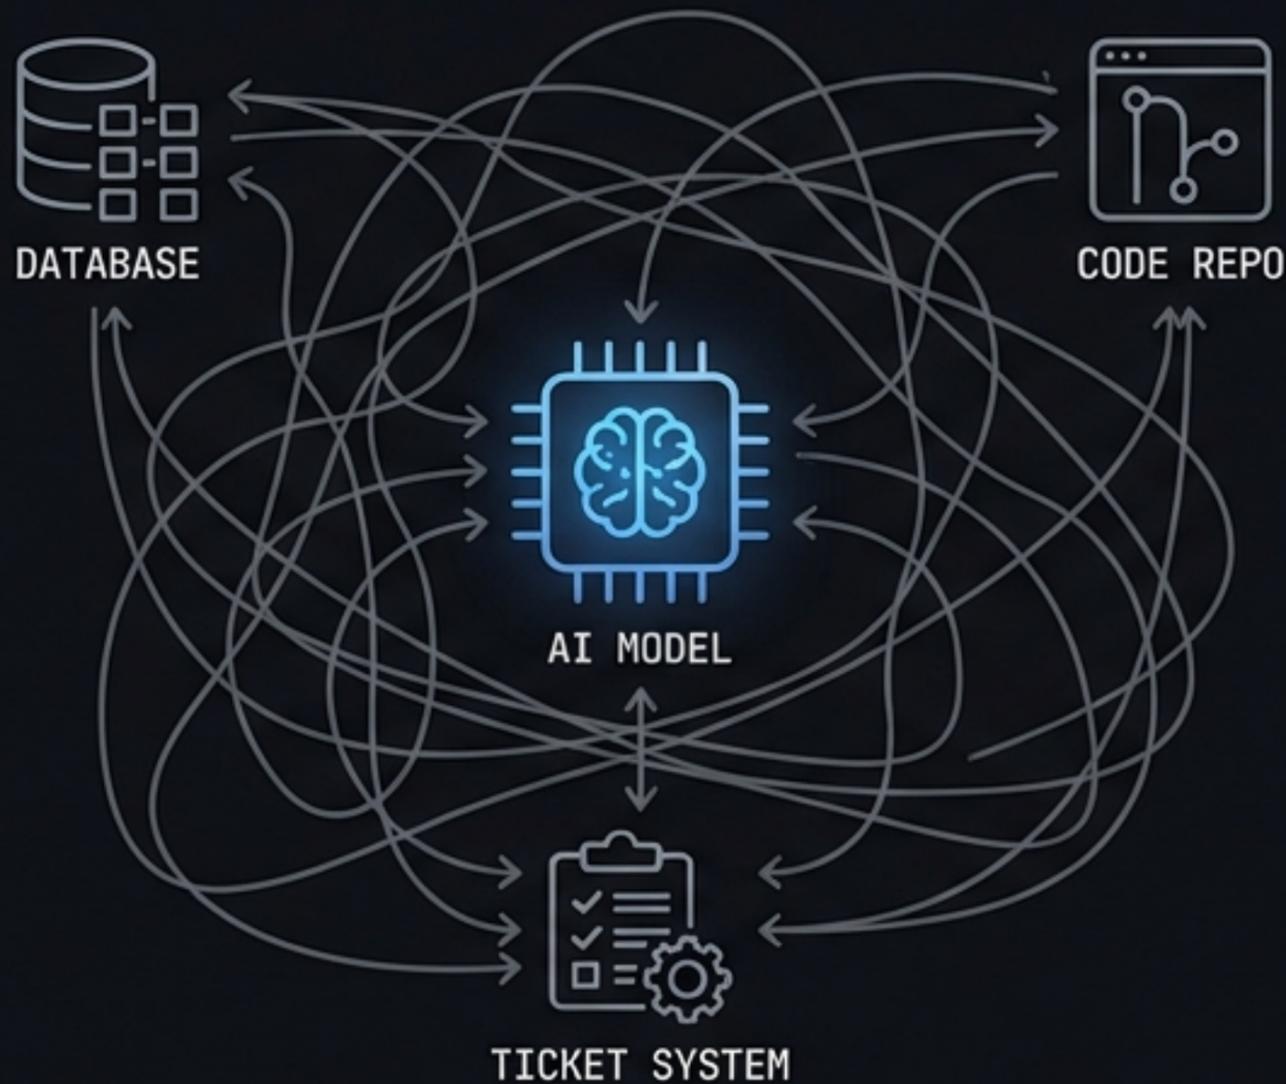
VERSION: 2026  
ARCHITECTURE GUIDE



Standardizing the connection between AI models, external tools, and enterprise data via JSON-RPC 2.0.

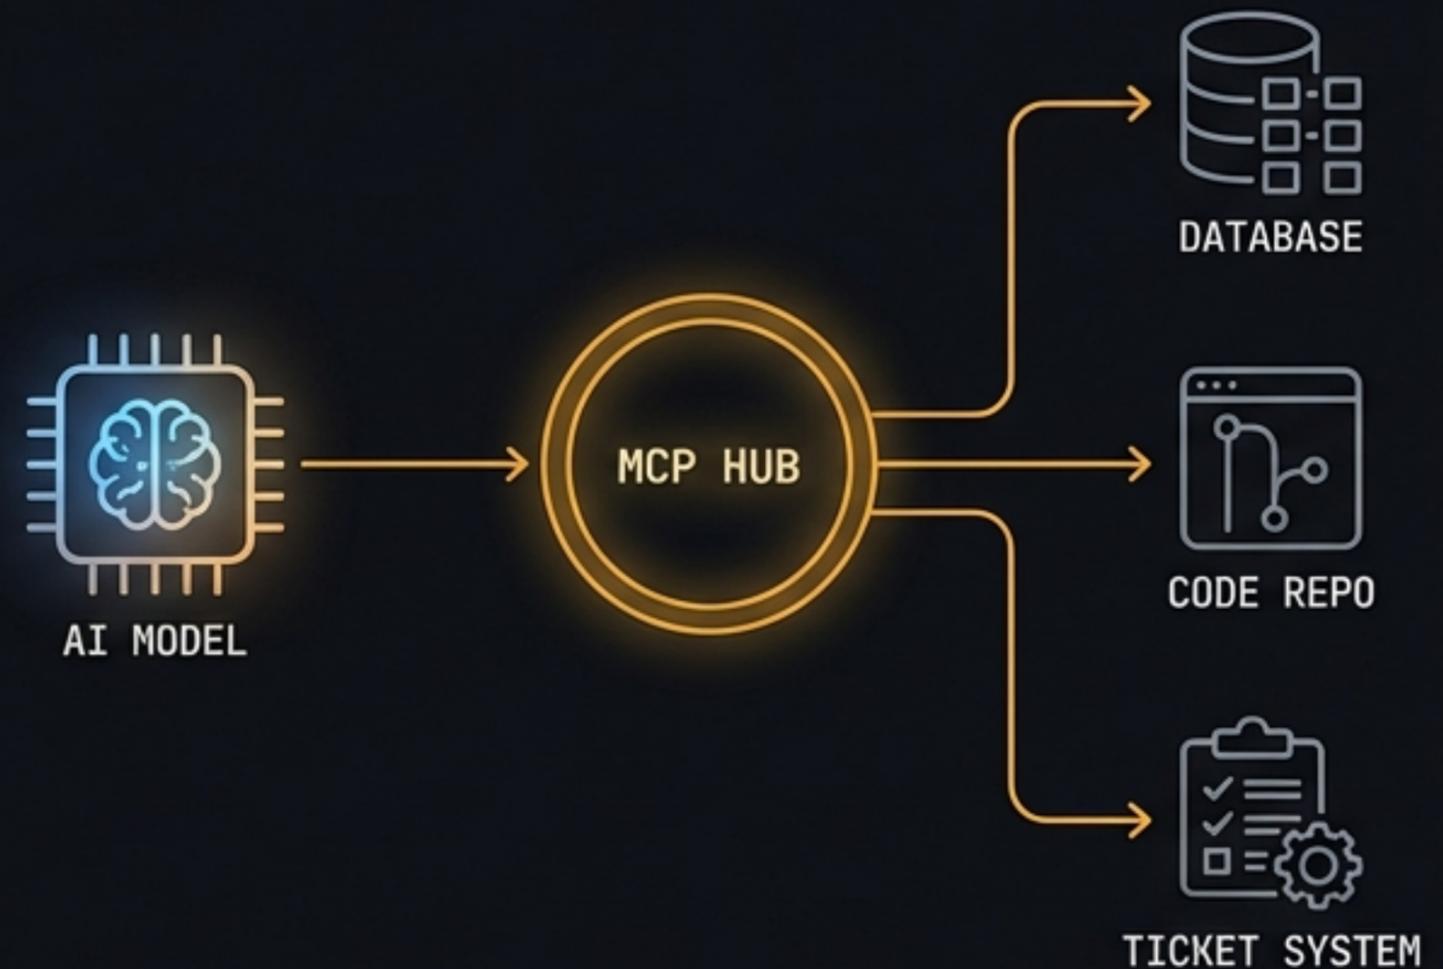
# Fragmented Integrations vs. The Universal Hub

## The Old World: N-to-N Complexity



Custom glue code, brittle mapping, unique auth per API.

## The MCP Standard: Hub-and-Spoke



**Key Insight:** MCP eliminates the need to write custom API connectors for every new tool. Write the server once; any MCP-compatible agent can discover and utilize it.

# The USB-C Adapter for Agentic AI



## The Hardware (USB-C):

Standardized port. Plugs into any device. Instantly negotiates power and data protocols.

## The Software (MCP):

Standardized interface. Plugs into any LLM Host. Instantly exposes tools, resources, and context windows.

“If function calling was ‘the model can call a tool’, MCP is ‘the model can call tools in a repeatable, portable, and vendor-agnostic way.’”

# Architectural Positioning Matrix

	Traditional APIs	Agent Skills	A2A Protocols	Model Context Protocol
Primary Function	Static data retrieval	Local static reference scripts	Agent-to-Agent dialogue	Dynamic Tool & Data Discovery
Connection Type	Hardcoded JetBrains Mono	Local markdown/scripts JetBrains Mono	Network communication JetBrains Mono	Universal JSON-RPC 2.0 JetBrains Mono
Data Freshness	Requires polling	Static/File-based	Real-time dialogue	Real-time / Stateful
Setup Complexity	High (Custom glue code)	Low (Local IDEs)	Medium	Low (Standardized)

# The Host, Client, and Server Triad

## The Host

(e.g., Claude Desktop, Cursor, n8n). Contains the LLM and the user interface.

## The Client

Embedded inside the Host. Acts as the translator and traffic controller.

The MCP Protocol

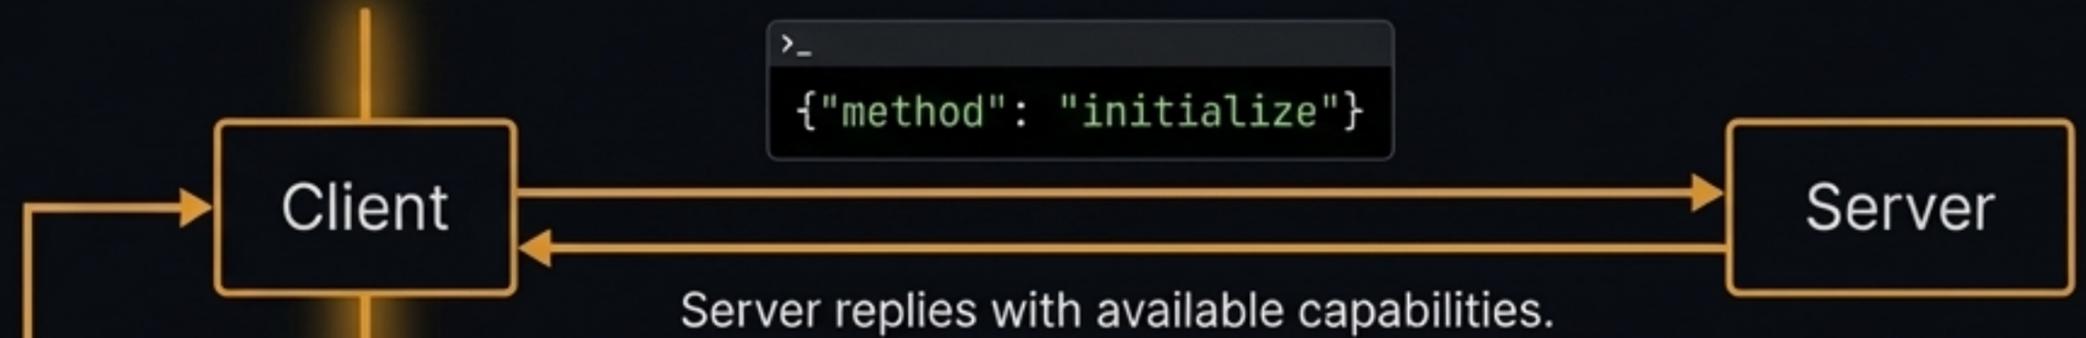
## The Server

(e.g., GitHub, PostgreSQL, local files). A lightweight program exposing capabilities.

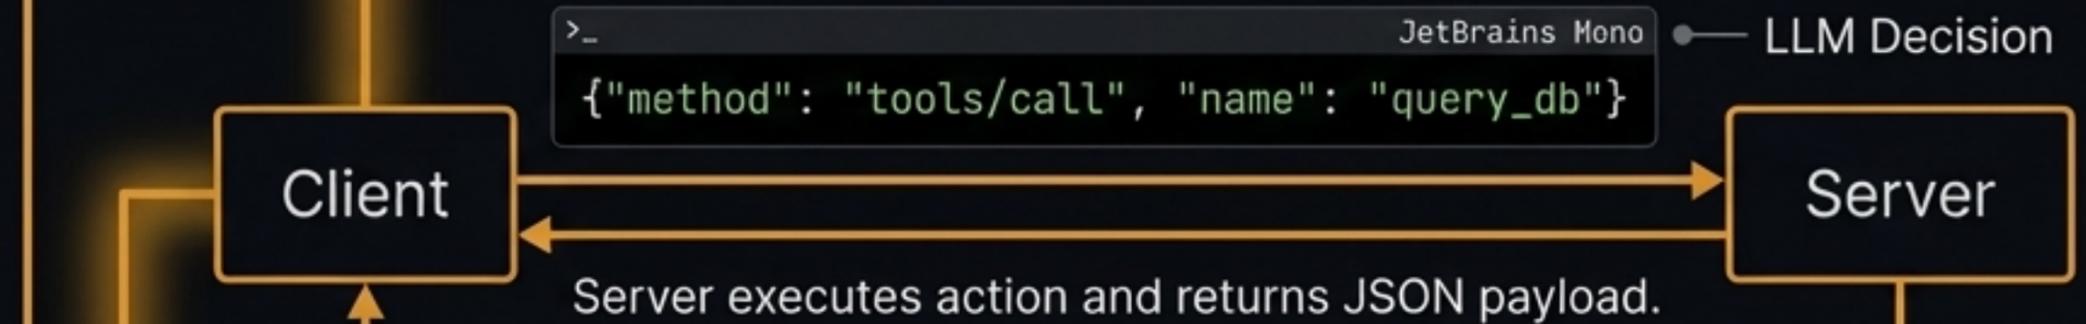
**Definition:** The Host houses the MCP Client, which invokes the protocol to maintain a 1:1 connection with the MCP Server.

# Anatomy of a Request Lifecycle

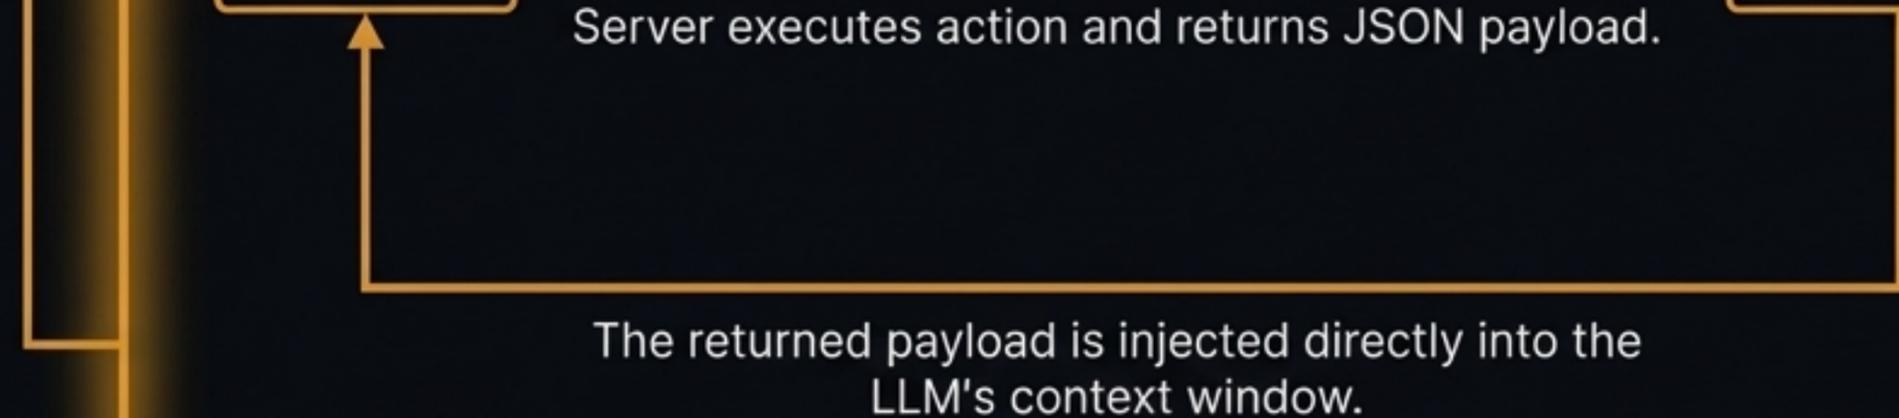
Phase 1:  
Initialization  
(Handshake)



Phase 2:  
Message Exchange



Phase 3:  
Context Update



**Architectural Note:** Everything happens in a stateful session, ensuring the AI maintains context throughout complex, multi-step workflows.

# The Transport Layer Diagnostic

## Local Transport (stdio)

- **Environment:** Laptop / Single Machine
- **Mechanics:** Standard input/output pipes. JSON-RPC messages exchanged locally.
- **State:** Inherently stateful.
- **Security:** Inherits local machine permissions.

## Remote Transport (HTTP + SSE)

- **Environment:** Cloud / Distributed Architecture
- **Mechanics:** Server-Sent Events (SSE) for server-to-client, POST requests for client-to-server.
- **State:** Streamable HTTP (supports both stateful and stateless).
- **Security:** Requires strict OAuth 2.1 or token-based authorization.

# The Tripartite Capability Model

## Pillar 1: Tools (Actions)

Functions invoked by the client.

### Examples:

- Sending a Slack message, executing a PostgreSQL query, triggering an API.

## Pillar 2: Resources (Context)

Read-only data exposed by the server.

### Examples:

- Database schemas, tracking logs, internal markdown wiki pages.

## Pillar 3: Prompts (Workflows)

Structured prompt blueprints.

**Examples:** Takes the prompt-engineering burden off the user by hardcoding complex instructions into the server.

# Dynamic Discovery and Tool Constraints

**The LLM Limit:** Current language models have cognitive caps (e.g., Copilot Studio supports a maximum of ~70 tools). Exposing an entire ERP system at once breaks the model.

```
System Log – MCP Client
Agent: "What tools do you have?"
Server: "Returning tools list (Page 1 of 5)..."
```

**Dynamic Tool Mode:** Instead of pre-loading all tools, the MCP Client can dynamically query the server to discover available objects at runtime.

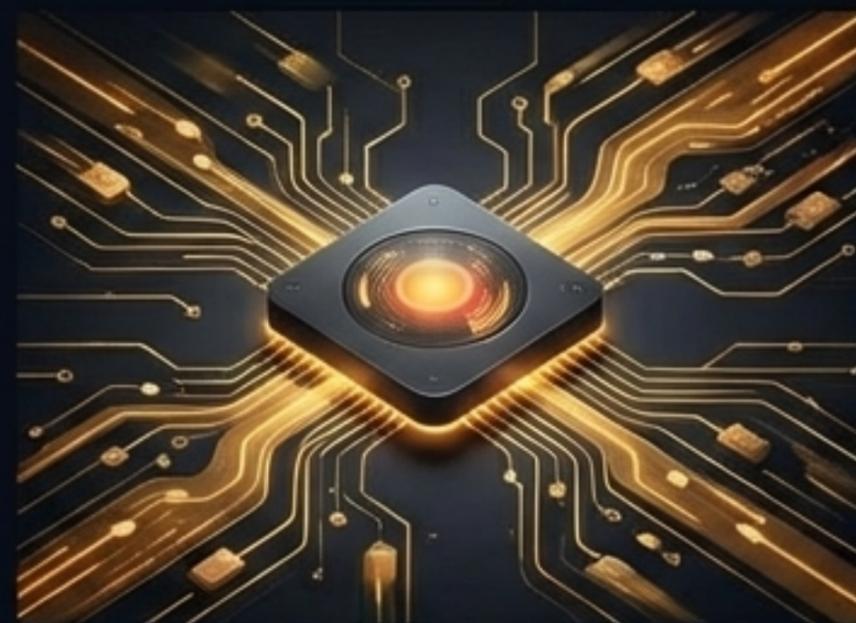
**Selective Configuration:** Administrators configure specific profiles (e.g., granting 'read-only' vs. 'CRUD' access) to selectively limit the surface area exposed to the AI.

# Differentiating MCP from Agent Skills



## Agent Skills

Local, file-based enhancements. Usually a markdown file containing specialized procedures and local scripts. Best for localized coding agents (e.g., Claude Code).

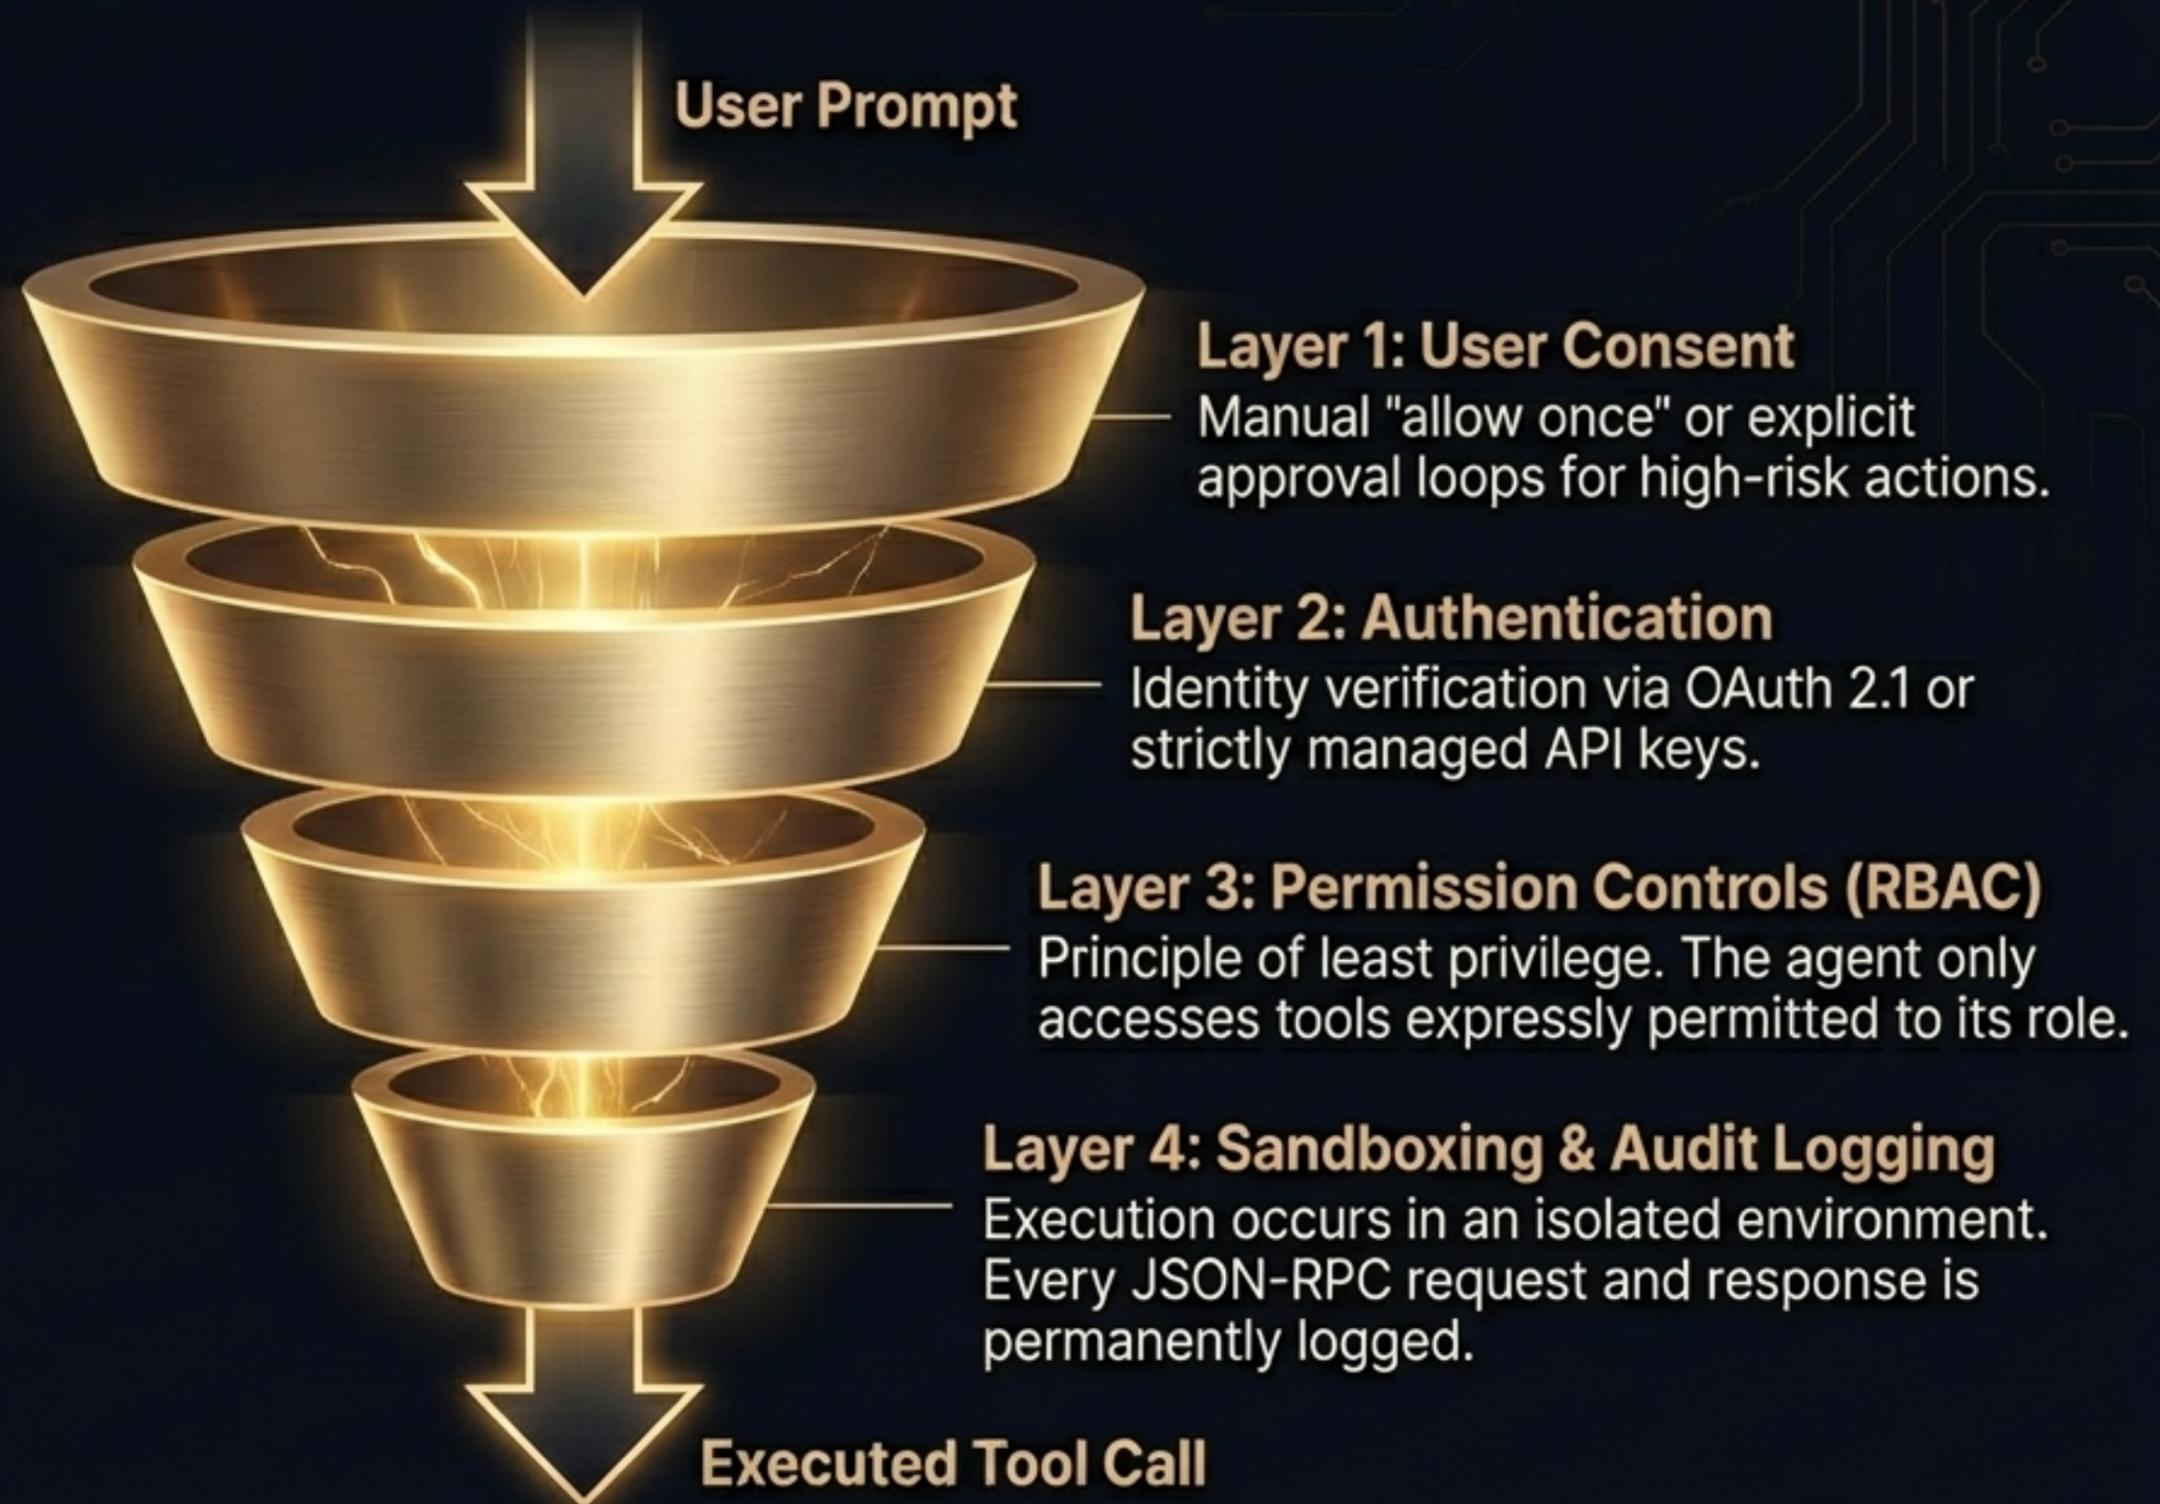


## MCP Servers

Live infrastructure. Provides a standardized interface for real-time, dynamic data (e.g., Kafka topics, live ticketing systems) and authenticated cloud actions.

**Verdict:** They are complementary. Skills enhance local reasoning; MCP connects the reasoning to the outside world.

# The Enterprise Security Funnel



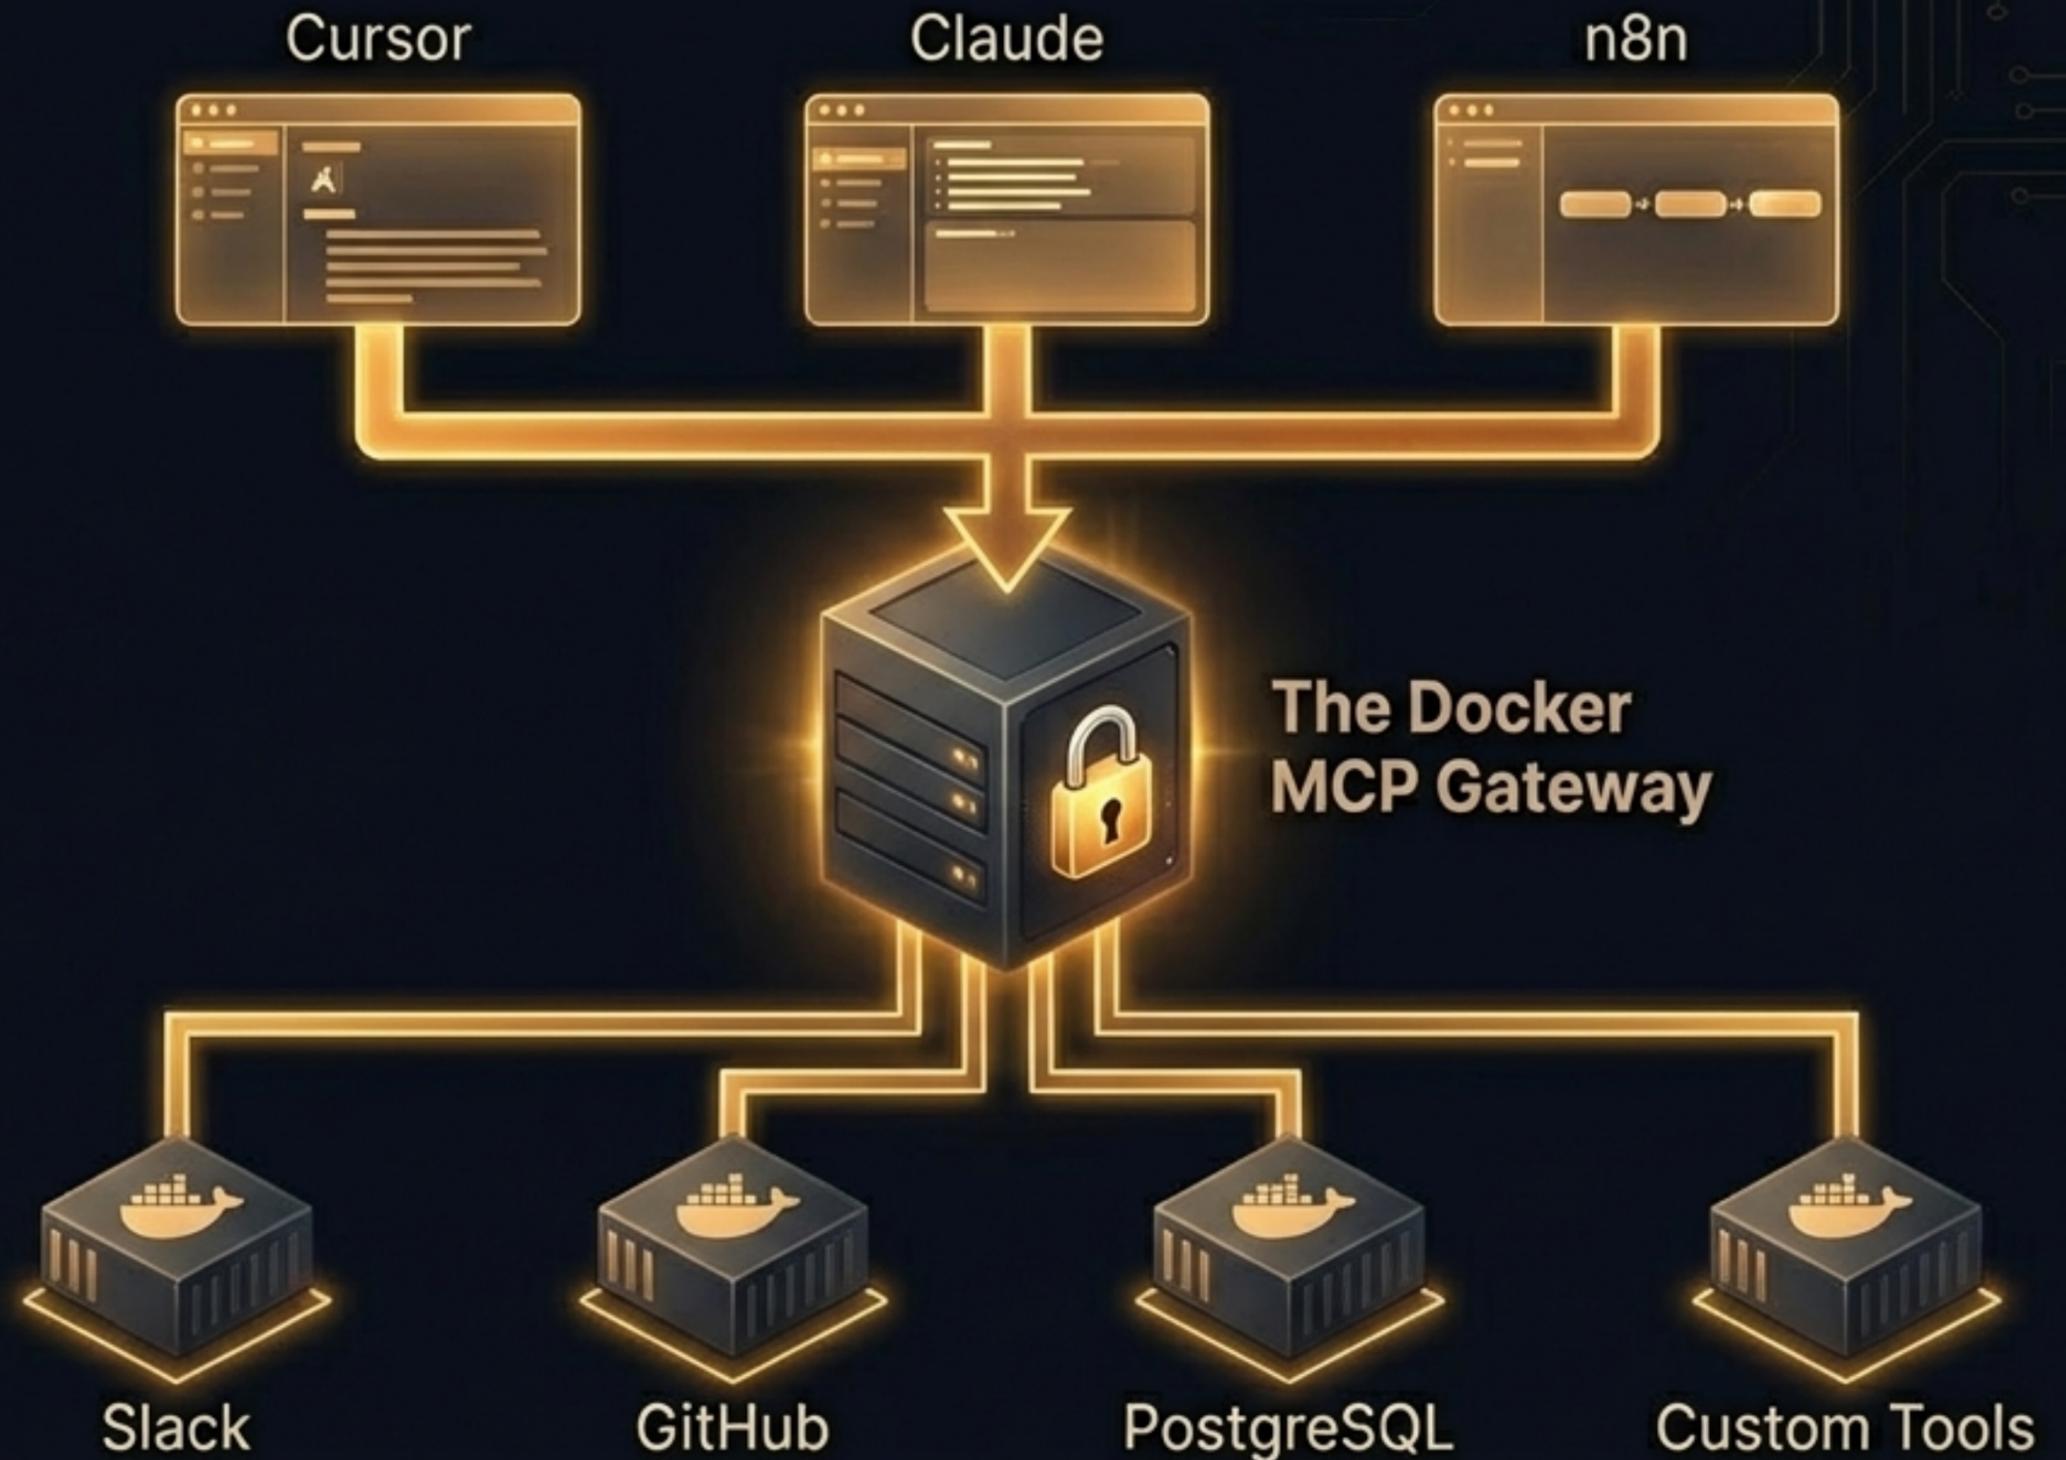
# Enterprise Governance and Lifecycle Management



**Strategic Impact:** Platform engineers maintain absolute security control while AI engineers gain instant access to certified, compliant toolsets.

# Centralized Orchestration: The Docker Gateway

Instead of configuring credentials and connections for dozens of distinct servers per client, the Gateway manages a single secure connection to the host and routes requests to containerized servers. Handles token management, catalog aggregation, and environment isolation natively.

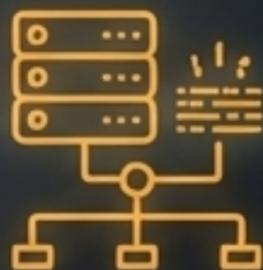


# High-Impact Enterprise Workflows



## Developer Productivity

The agent reads a **GitHub** repository, identifies a bug in the **CI/CD** logs, creates a fix, pushes a **PR**, and updates the **Confluence** documentation—all via synchronized **MCP** tools.



## IT Operations

The agent accesses a **ServiceNow** MCP server, securely reads a user's role, provisions the requested cloud infrastructure, and resolves the **Jira** ticket.



## Regulated Industries (Fintech/Pharma)

The agent executes sensitive data retrieval with mandatory human-in-the-loop approval gates, strictly adhering to **RBAC** policies enforced by the **MCP** server.

# Silent Analytics and Self-Evaluation



## The Concept:

MCP isn't just for user-triggered tools. It can run as a **background injection** via system prompts.

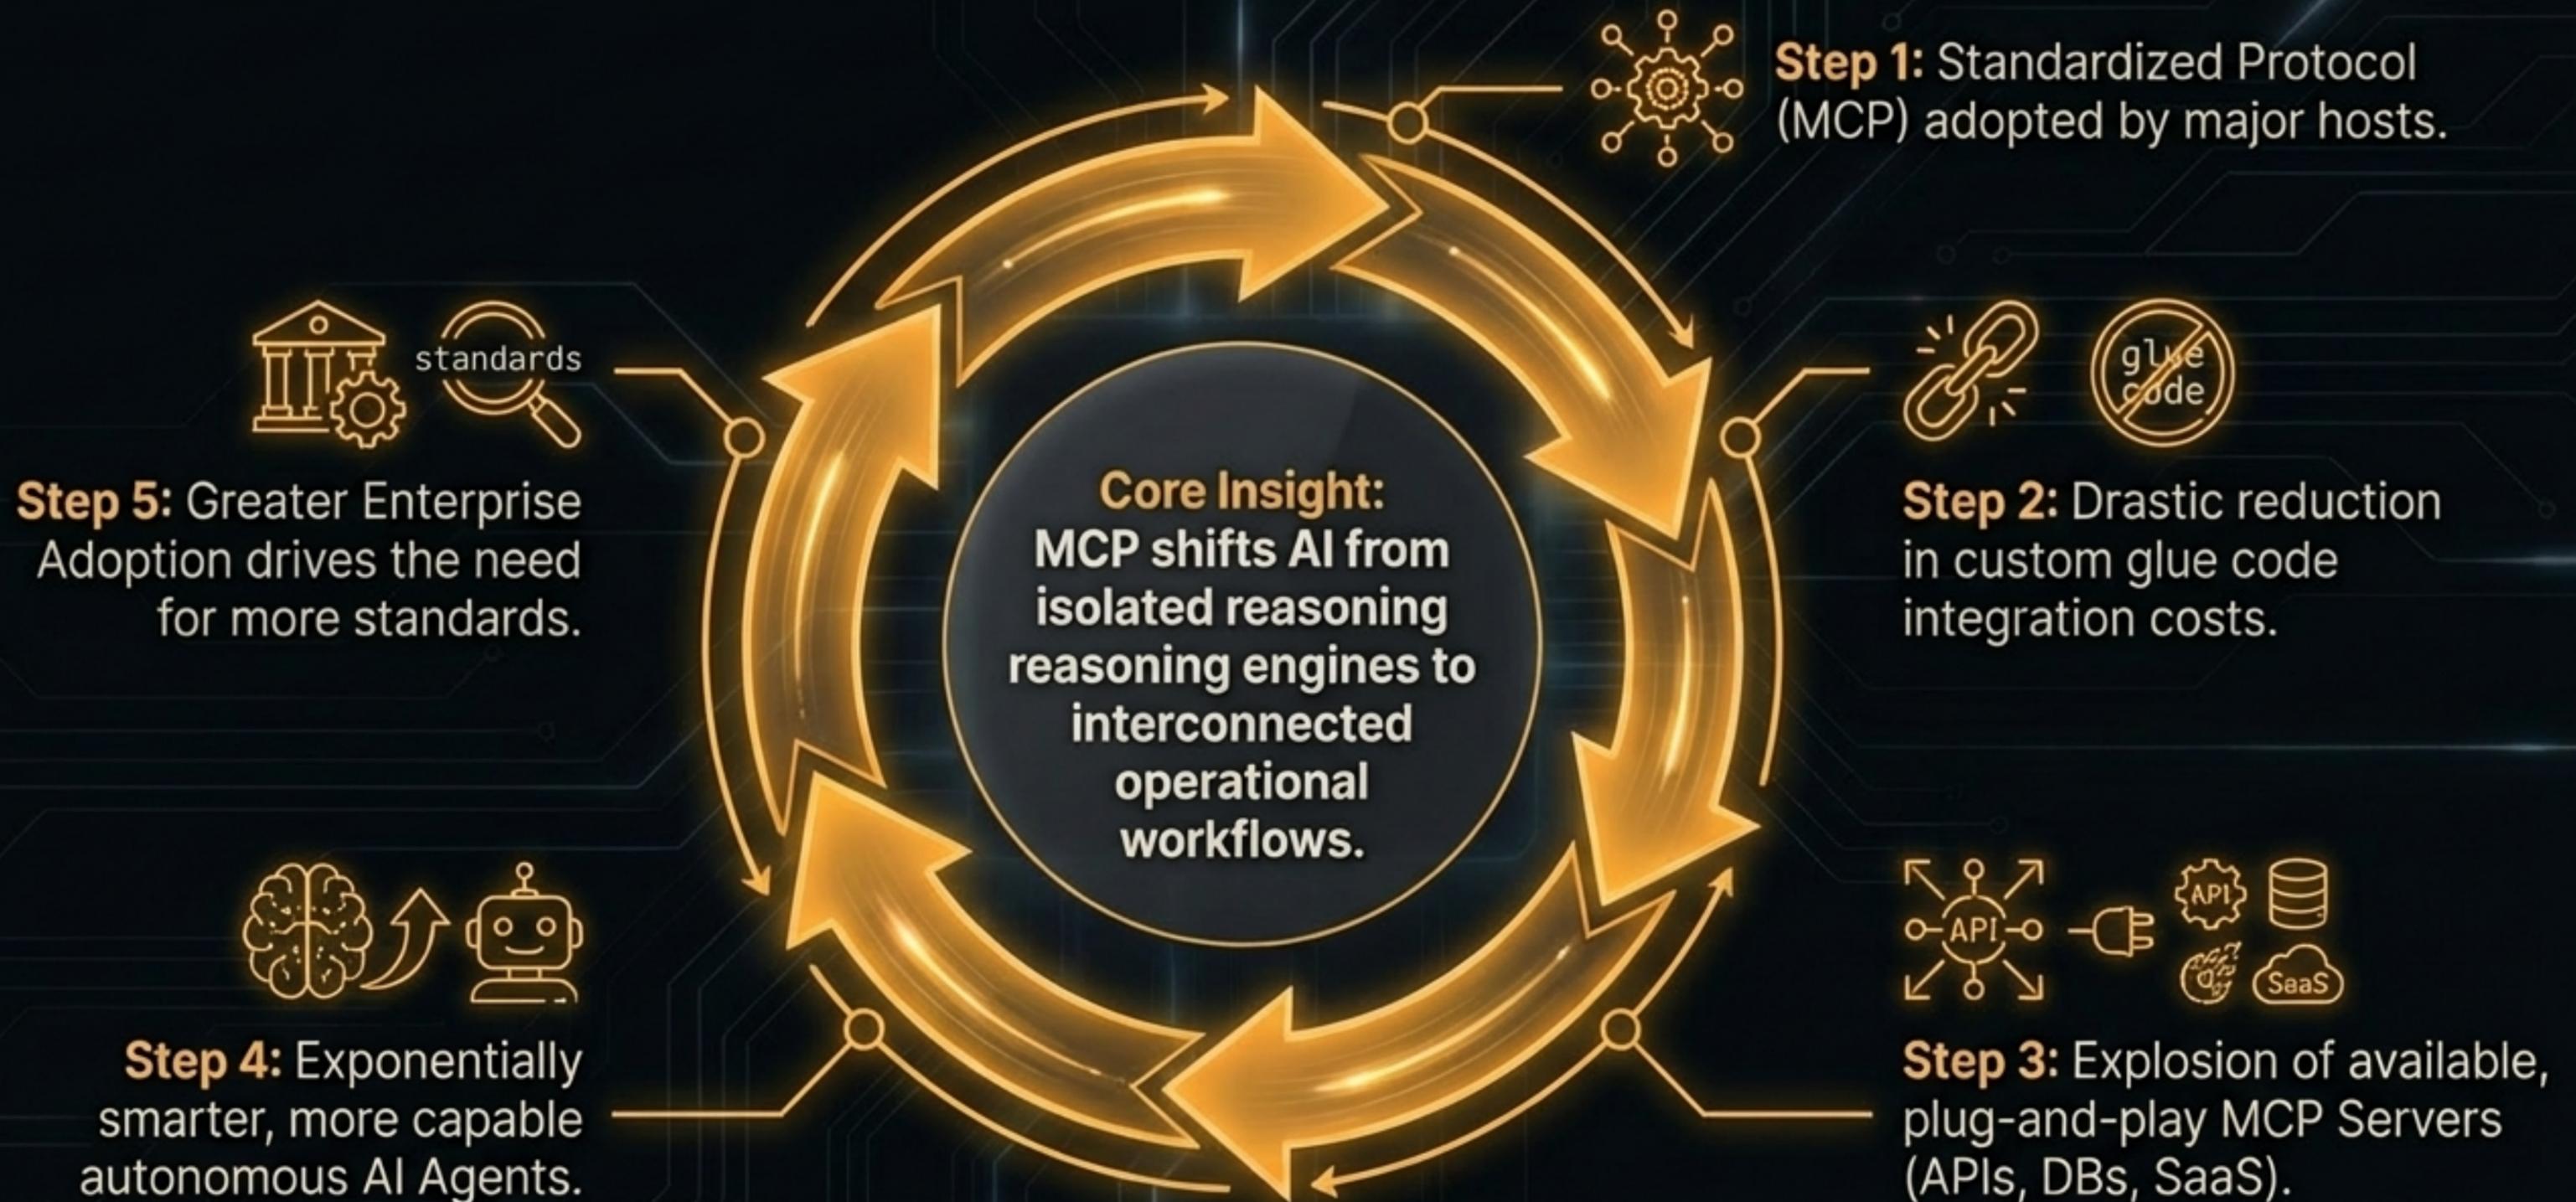
## The Execution:

Every time the AI generates a response, it is prompted to secretly call an **Analytics MCP** server, sending back telemetry on its own reasoning steps, tool usage, and evaluation metrics.

## The Result:

Unprecedented **visibility** into agentic behavior, functioning like **Google Analytics** but for autonomous AI reasoning.

# The Agentic Interoperability Flywheel



# The 2026 Horizon

## 1. Multimodal Protocol

Moving beyond text. MCP will support chunked messages and multipart streams, allowing agents to see, hear, and process images/video mid-stream.

## 2. Hierarchical Agents

Shifting from single-agent tool use to multi-agent coordination, where supervisor agents manage specialized sub-agents through shared MCP servers.

## 3. Open Governance

The protocol transitions from proprietary origins to a community-driven standard, cementing MCP as the foundational infrastructure of the next web.

**Closing Note:** AI isn't just getting smarter; it's getting connected.