

From Tools to Platforms

MCP Patterns for Building Open Agent Ecosystems

Guangya Liu
JPMC / Executive Director

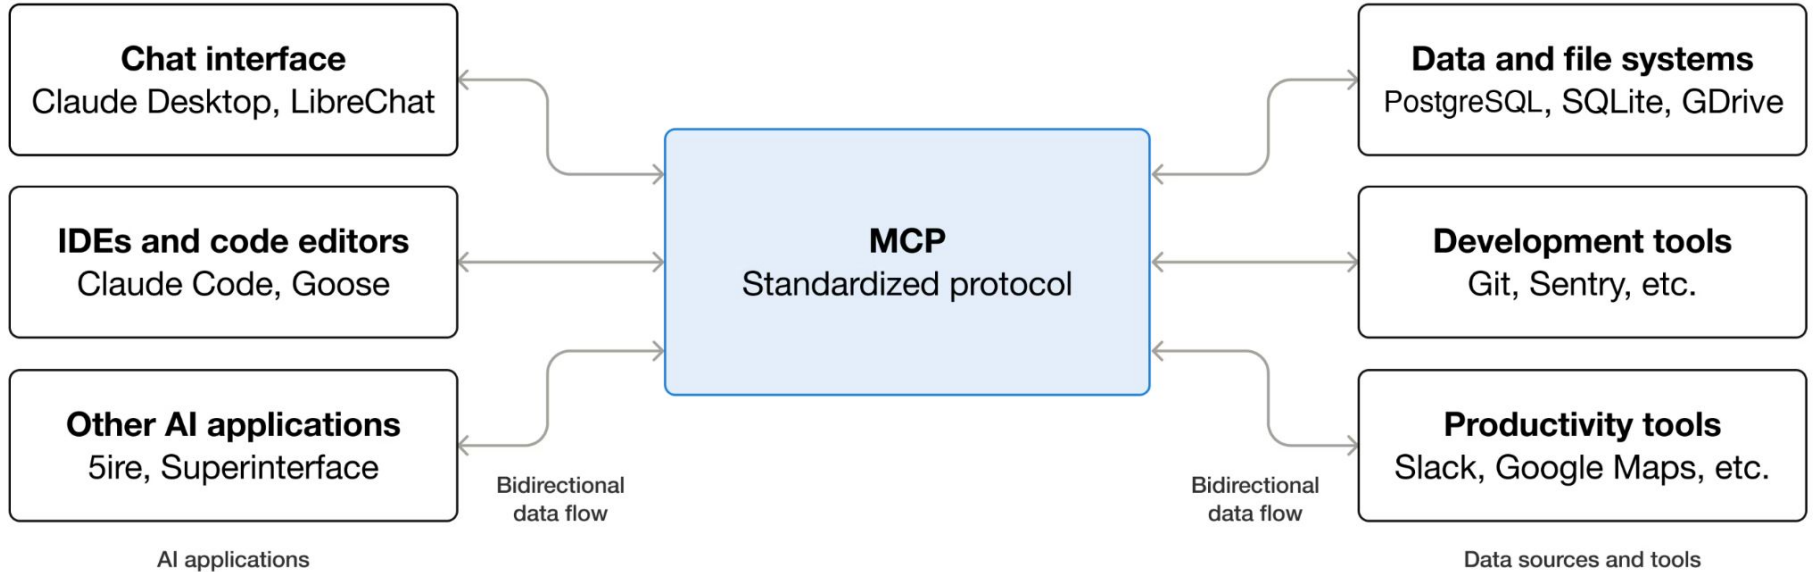
Self Intro

- Guangya Liu gyliu513@gmail.com
- Executive Director at JPMC (ex-RedHat, ex-IBM)
 - SRE Agent for JPMC
 - FullStack AI Observability with OpenTelemetry
 - GPU Optimization for AI/ML/HPC
- Open Source <https://github.com/gyliu513>
 - Contributor for llmstack(ogx), llm-d, kueue
 - Apache Mesos Committer and PMC Member
 - Former Maintainer for OpenTelemetry GenAI Semantic Convention
 - Former Maintainer for OpenStack

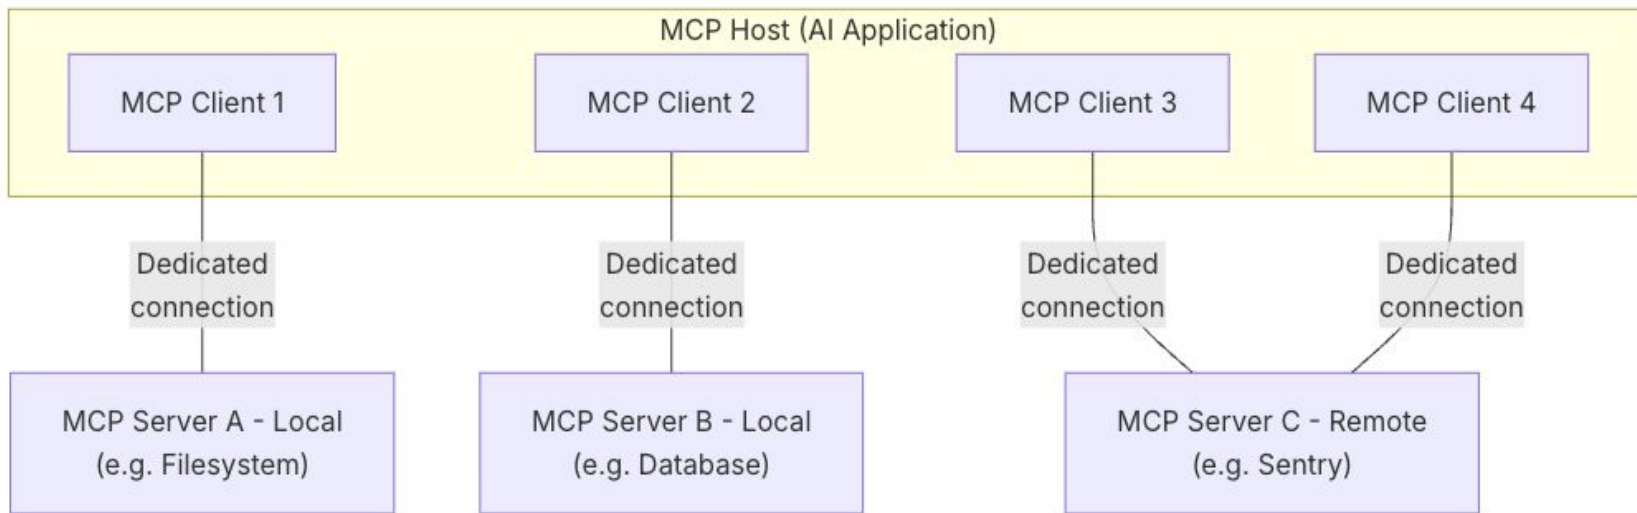
Agenda

- MCP Intro
- Single MCP Server
- Multiple MCP Servers

MCP Intro



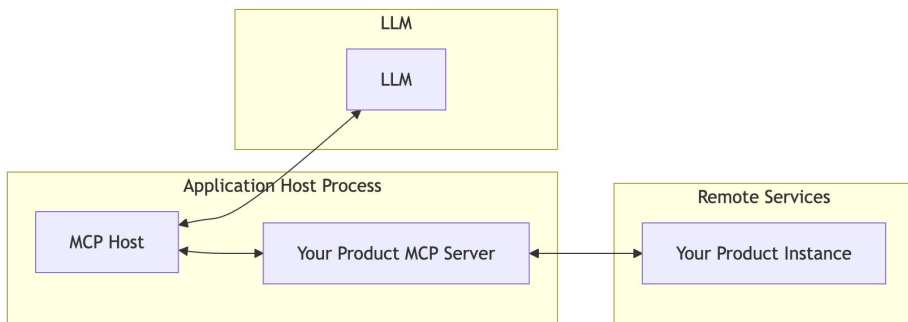
MCP Intro



Product/Platform Level MCP Server

- External MCP Server
- Internal MCP Server

Platform/Product level External MCP Server

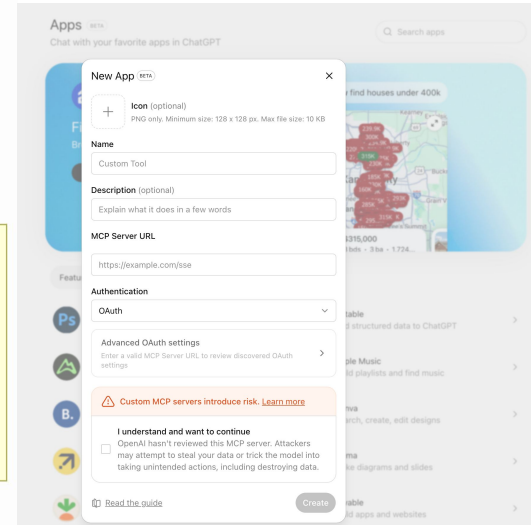
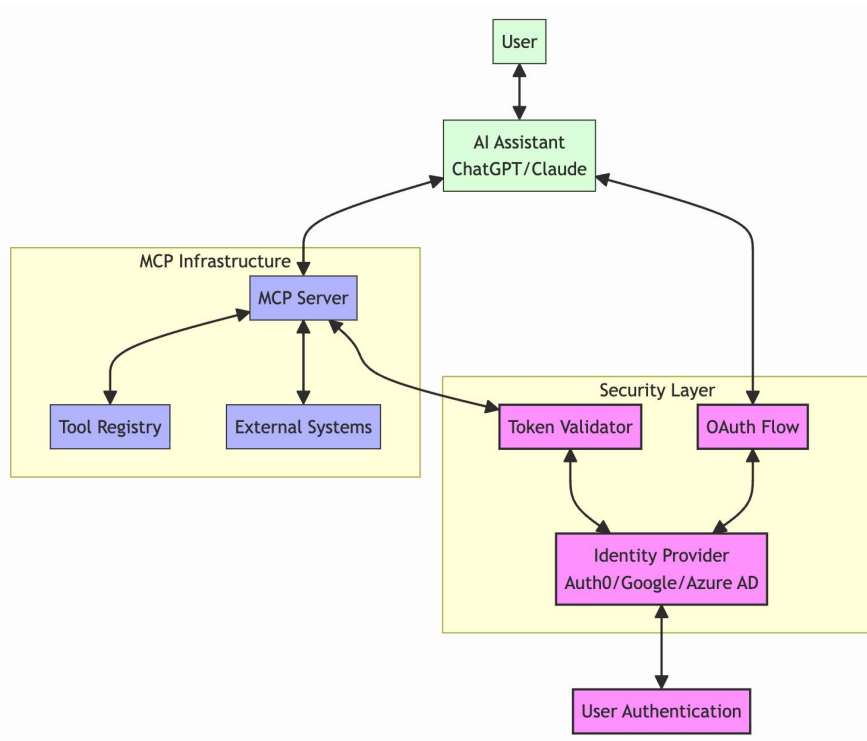


The screenshot shows a "Directory" interface with a search bar and a list of connectors. The interface is organized into sections:

- Skills**
- Connectors** (Active section):
 - Anthropic & Partners** (Filter):
 - Filesystem**: Let Claude access your filesystem to read and write files.
 - Control Chrome**: Control Google Chrome browser tabs, windows, and navigation.
 - pdf-viewer**: Read, annotate, and interact with PDF files — interactive viewer with search, navigation, annotations, form filling,...
 - PowerPoint (By Anthropic)**: Control Microsoft PowerPoint with AppleScript automation.
 - Word (By Anthropic)**: Control Microsoft Word with AppleScript automation.
 - Apify**: Extract data from any website with thousands of scrapers, crawlers, and automations on Apify Store.
 - Windows-MCP**: MCP Server that enables Claude to interact with Windows OS.
 - Fantastical**: Read events and tasks, create new items, make basic edits, and delete items from Fantastical for Mac.
- Plugins**

OAuth for External MCP Server

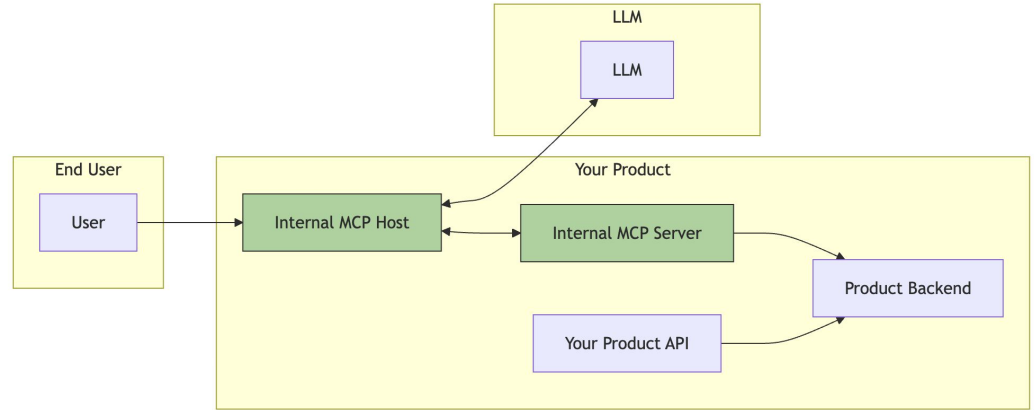
- Without OAuth
 - Copying access tokens by hand.
 - Managing token rotation and expiry.
 - Worrying about leaks in logs or environment variables.
- With OAuth
 - **Instant sign-in experience** – users click “Connect” and authorize via Google, GitHub, or SSO.
 - **No manual token management** – short-lived, auto-refreshed access tokens.
 - **Granular permissions** – only the scopes you need (`mcp.read`, `mcp.write`, etc.).



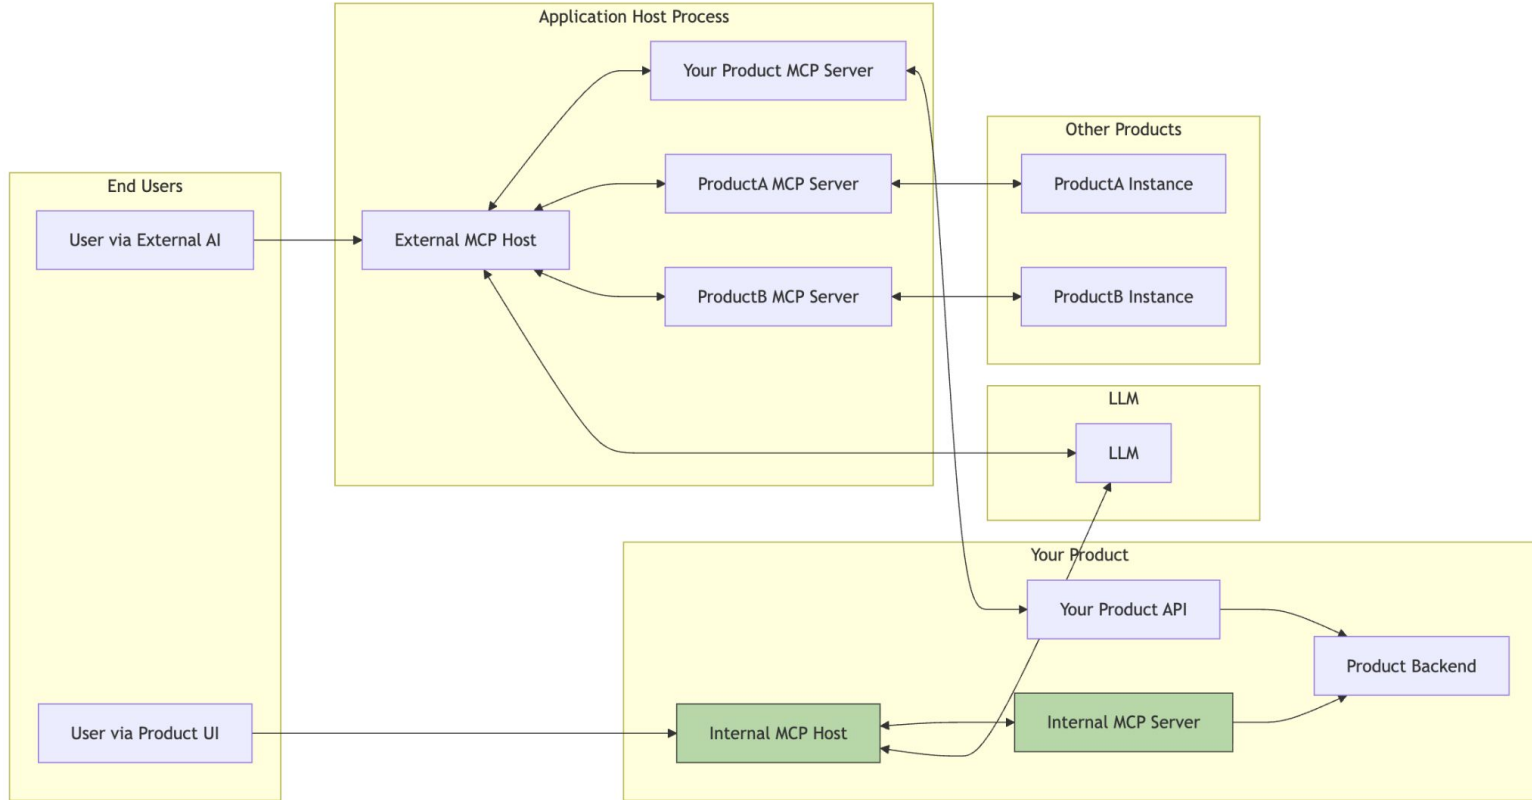
Platform/Product level Internal MCP Server

- Self Hosted MCP Host

- Your own build-in chat bot for product
- Dynatrace
- DataDog
- ...



Combining them Together

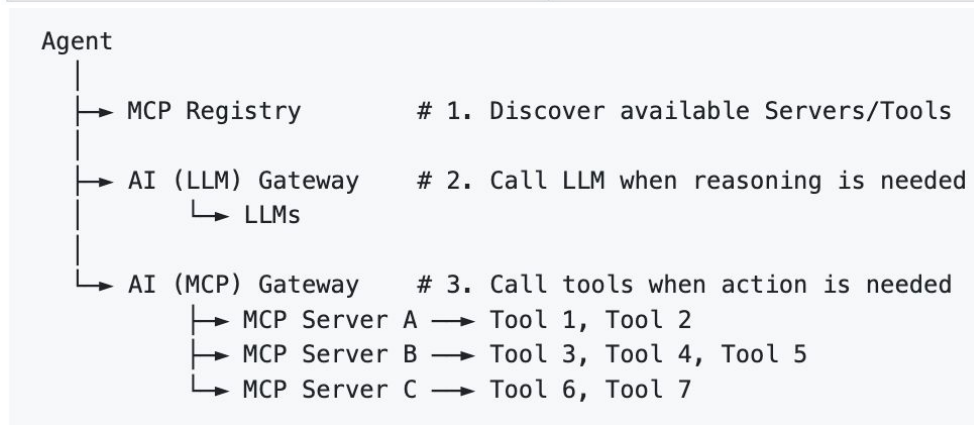


Multiple MCP Servers

- MCP Gateway
- Skill

MCP Gateway

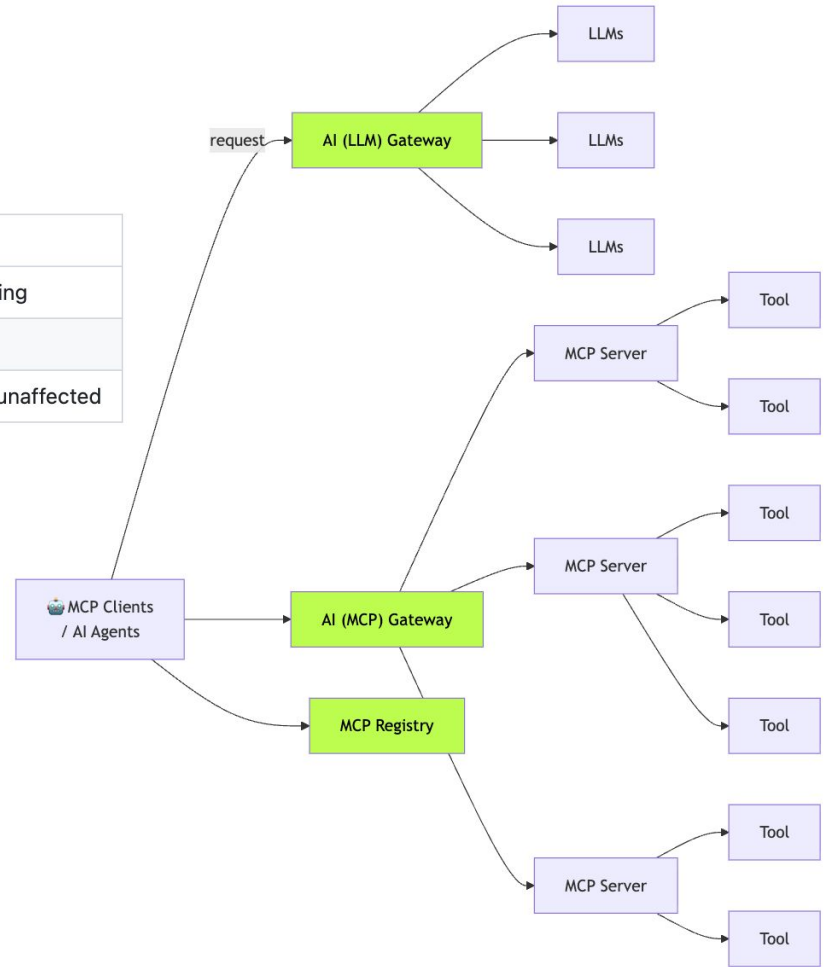
Problem	Gateway's Solution
Agent needs to integrate dozens of Tools	Connect to one Gateway; it manages everything
Tool-level auth, rate limiting, auditing, observability	Handled centrally at the Gateway layer
Servers dynamically joining/leaving	Registry + Gateway handle discovery; Agent unaffected



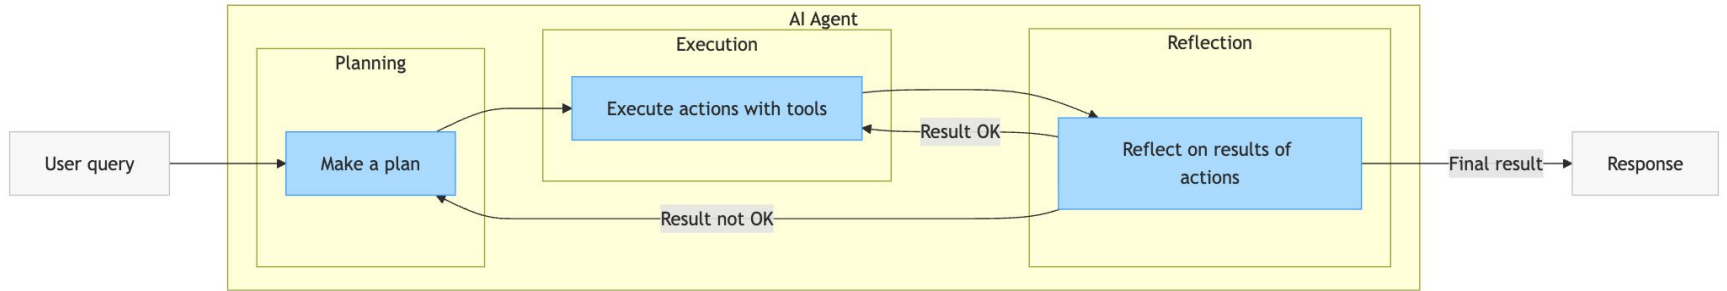
<https://github.com/kgateway-dev/kgateway>

<https://github.com/obot-platform/obot>

<https://github.com/IBM/mcp-context-forge>



ReAct Agent



Skill

```

---
name: orchestrator
description: ReAct meta-skill - Plan -> Execute -> Reflect loop until goal is met
trigger: "orchestrate {task}"
tools:
  - bash: [skill_list, skill_call, skill_call_parallel]
  - mcp: [mcp_discover, mcp_call]
permissions:
  skills: read
  mcp: all
max_iterations: 5
---

# Orchestrator Skill (ReAct Loop)

## Outer Loop - repeat until Reflect says DONE (max 5 iterations)

### Phase 1 - PLAN / REVISE
**iter = 1 (First time):**
- 'skill_list()' + 'mcp_discover()' - understand available capabilities
- Decompose task into Task DAG: `[[{id, skill, input, depends_on}]`

**iter > 1 (after a REVISE):**
- Review: what results came back? what was empty/wrong/missing?
- Revise DAG: retry failed tasks with different skill, add missing tasks, prune useless ones
- Record reason for each change (for next reflection to reference)

### Phase 2 - EXECUTE
while tasks remain in DAG:
  ready = [tasks where all depends_on are resolved]
  results = skill_call_parallel(ready)
  inject {{id.result}} into dependent task inputs

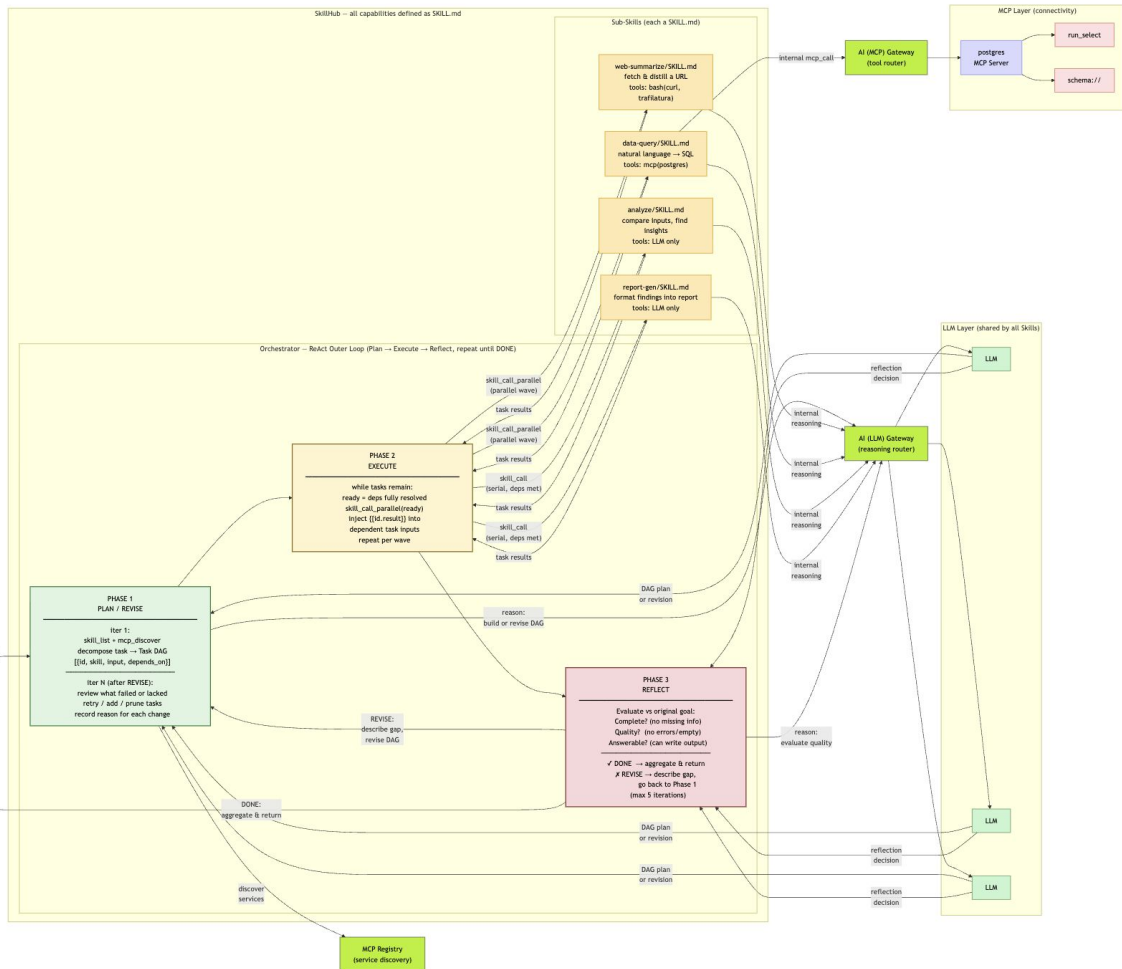
### Phase 3 - REFLECT
Evaluate results against the **original goal** (not just "did tasks finish"):

1. **Complete?** - do we have all the information needed?
2. **Quality?** - any results empty, errored, or too shallow?
3. **Answerable?** - can we write a correct, complete final answer right now?

+ All pass: **DONE** - exit loop, go to Aggregate
+ Any fail: **REVISE** - describe the gap clearly, go back to Phase 1

## Aggregate (after loop exits)
Merge all results -> return final markdown output.

```



Tools or MCP

- With Skills, should I use tools or MCP
- Tools
 - Short-lived execution
 - Agent-local capabilities
 - Lightweight interactions
- MCP
 - Long-lived integrations
 - External systems/services
 - Reusable and standardized connections

MCP Pattern



Thanks