MCP:
Model Context Protocol

Roger Lam

An open standard for connecting models and data

MCP addresses this challenge. It provides a universal, open standard for connecting AI systems with data sources, replacing fragmented integrations with a single protocol. The result is a simpler, more reliable way to give AI systems access to the data they need.

What is MCP?

The Model Context Protocol (MCP) lets you build servers that expose data and functionality to LLM applications in a secure, standardized way. Think of it like a web API, but specifically designed for LLM interactions. MCP servers can:

- Expose data through **Resources** (think of these sort of like GET endpoints; they are used to load information into the LLM's context)
- Provide functionality through **Tools** (sort of like POST endpoints; they are used to execute code or otherwise produce a side effect)
- Define interaction patterns through **Prompts** (reusable templates for LLM interactions)
- And more!

- Descriptions from their blog post and their Python SDK
- Side note: I'm very impressed by the polish on every released piece of work

Connect to not only external services but local files

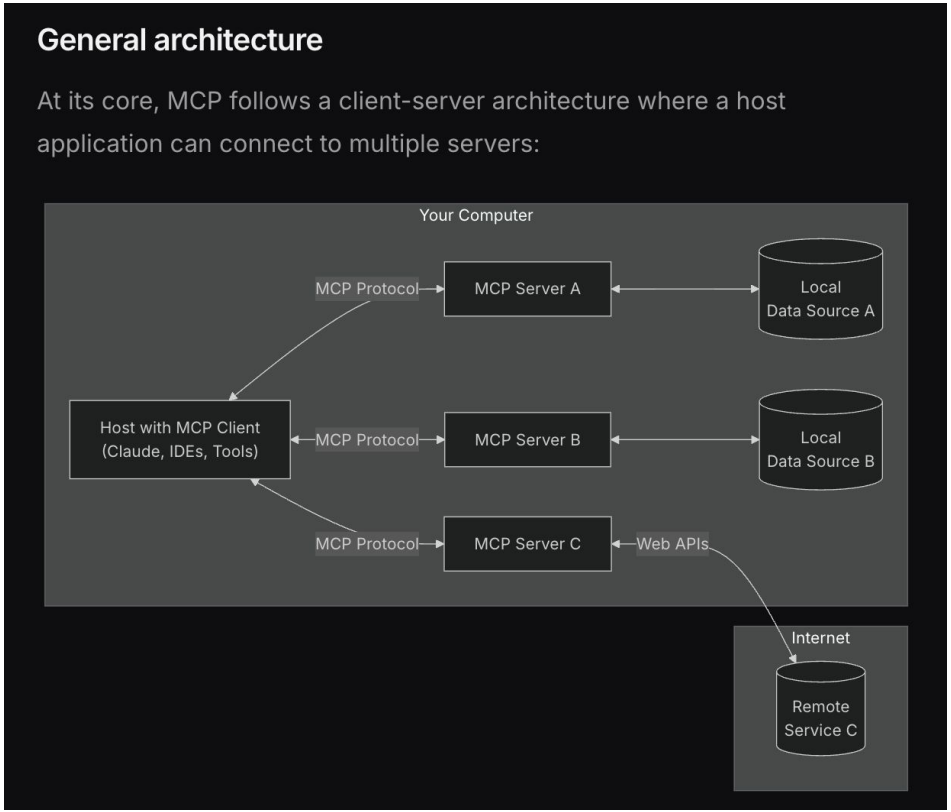

This is a stateful connection… for now

The Model Context Protocol (MCP) defines a rigorous lifecycle for client-server connections that ensures proper capability negotiation and state management.

1. **Initialization**: Capability negotiation and protocol version agreement
2. **Operation**: Normal protocol communication
3. **Shutdown**: Graceful termination of the connection

There's an ongoing discussion about statefulness and a proposal to change it announced today.

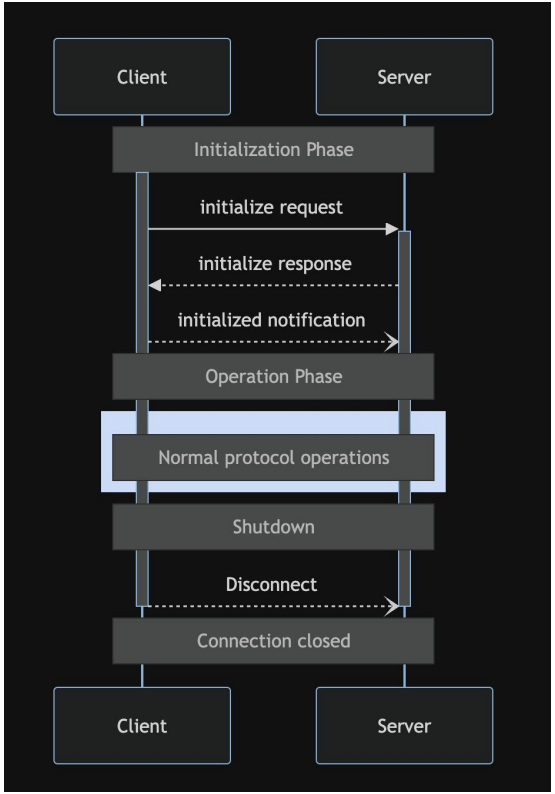

Reminds me of GraphQL

- A translation layer
- I thought AGI was supposed to figure this out for us
- Will agents just be more code under-the-hood?

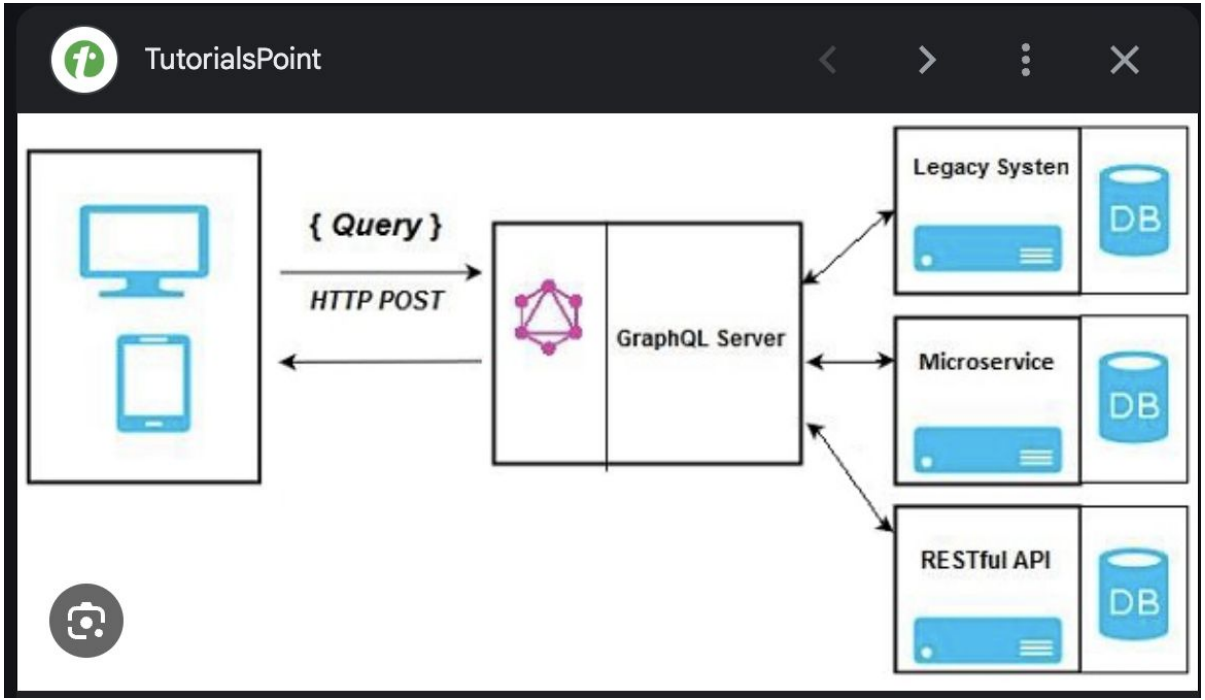

Security is still a thing. And it’s scary.

From the tools section in the MCP docs.

Remember, these servers will need access to credentials to access services. Think booking flights.

Each of these MCP servers will need to be hosted and trusted. So will each Airline host their own MCP Server?

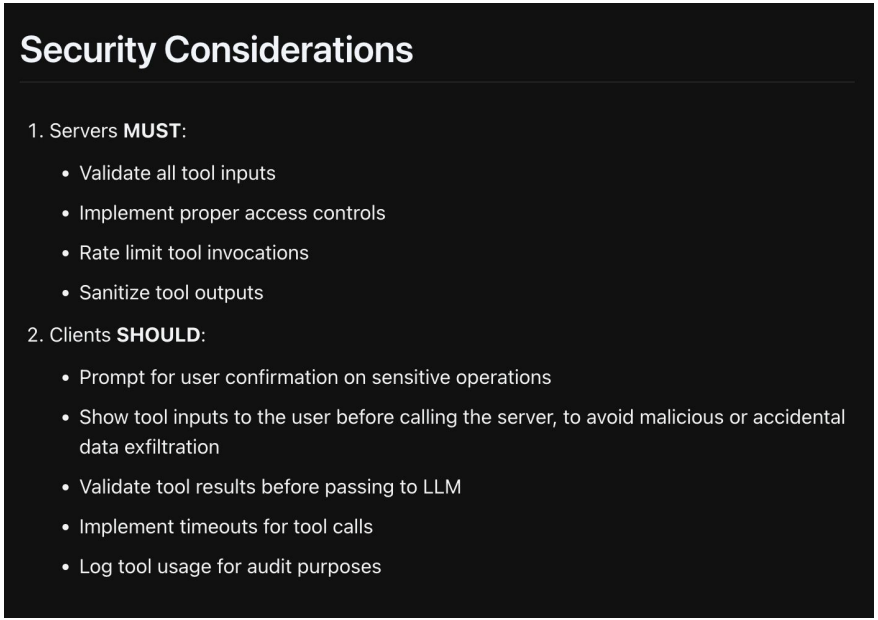

Security Considerations

1. Servers **MUST**:
   - Validate all tool inputs
   - Implement proper access controls
   - Rate limit tool invocations
   - Sanitize tool outputs
2. Clients **SHOULD**:
   - Prompt for user confirmation on sensitive operations
   - Show tool inputs to the user before calling the server, to avoid malicious or accidental data exfiltration
   - Validate tool results before passing to LLM
   - Implement timeouts for tool calls
   - Log tool usage for audit purposes

Alternatives: llms.txt

Jeremy Howard proposed a llms.txt file for website to host and for LLMs to scrape or access.

Much easier to maintain.

Here's an example of `llms.txt`, in this case a cut down version of the file used for the FastHTML project (see also the full version:

```
# FastHTML

> FastHTML is a python library which brings together Starlette, Uvicorn, HTMX,
and fastcore's `FT` "FastTags" into a library for creating server-rendered
hypermedia applications.

Important notes:

- Although parts of its API are inspired by FastAPI, it is *not* compatible with
FastAPI syntax and is not targeted at creating API services
- FastHTML is compatible with JS-native web components and any vanilla JS
library, but not with React, Vue, or Svelte.

## Docs

- [FastHTML quick start]
(https://answerdotai.github.io/fasthtml/tutorials/quickstart_for_web_devs.html.md
): A brief overview of many FastHTML features
```

We propose adding a `/llms.txt` file to websites that are designed for reading by language models, not just humans. `llms.txt` is a file that outlines the information that a model may want to retrieve (with links) when assembling context for LLM prompts relevant to a website. Here's an example llms.txt file.

Alternatives: Models that can understand screens

It seemed like Apple was close to this breakthrough - and many others chasing (Rabbit Action Model).

Recent delays are suggesting otherwise.

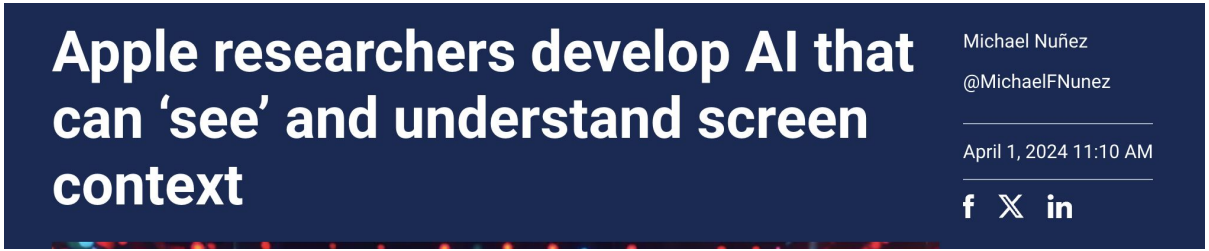

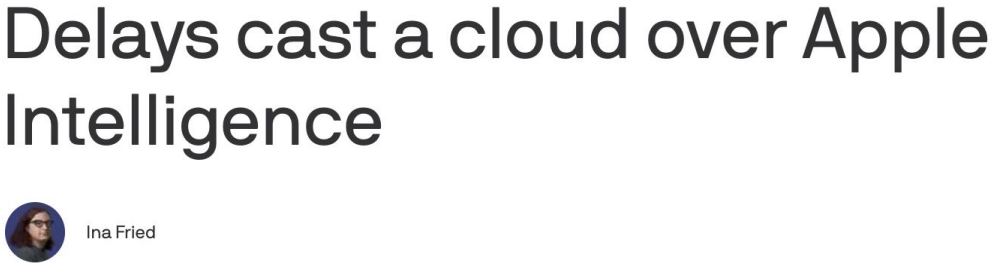

Thanks! Hope this helped.

Follow for more.

lamroger.com

linkedin.com/in/lam-roger