GreenLake MCP servers

Roman Nersisyan, Vandewilly Silva

January, 2026

Agenda

01	The AI Knowledge Gap
02	Model Context Protocol
03	Natural Conversations, Real Answers
04	What HPE Has Built
05	Security
06	MCP Generator / Testing with LLM
07	Versioning
08	What’s Next

What If You Could Just Ask?

Get an instant, accurate answer.

- No dashboards.
- No queries.
- Just conversation.

You could just ASK:

"How many devices do we have?"

The AI Knowledge Gap

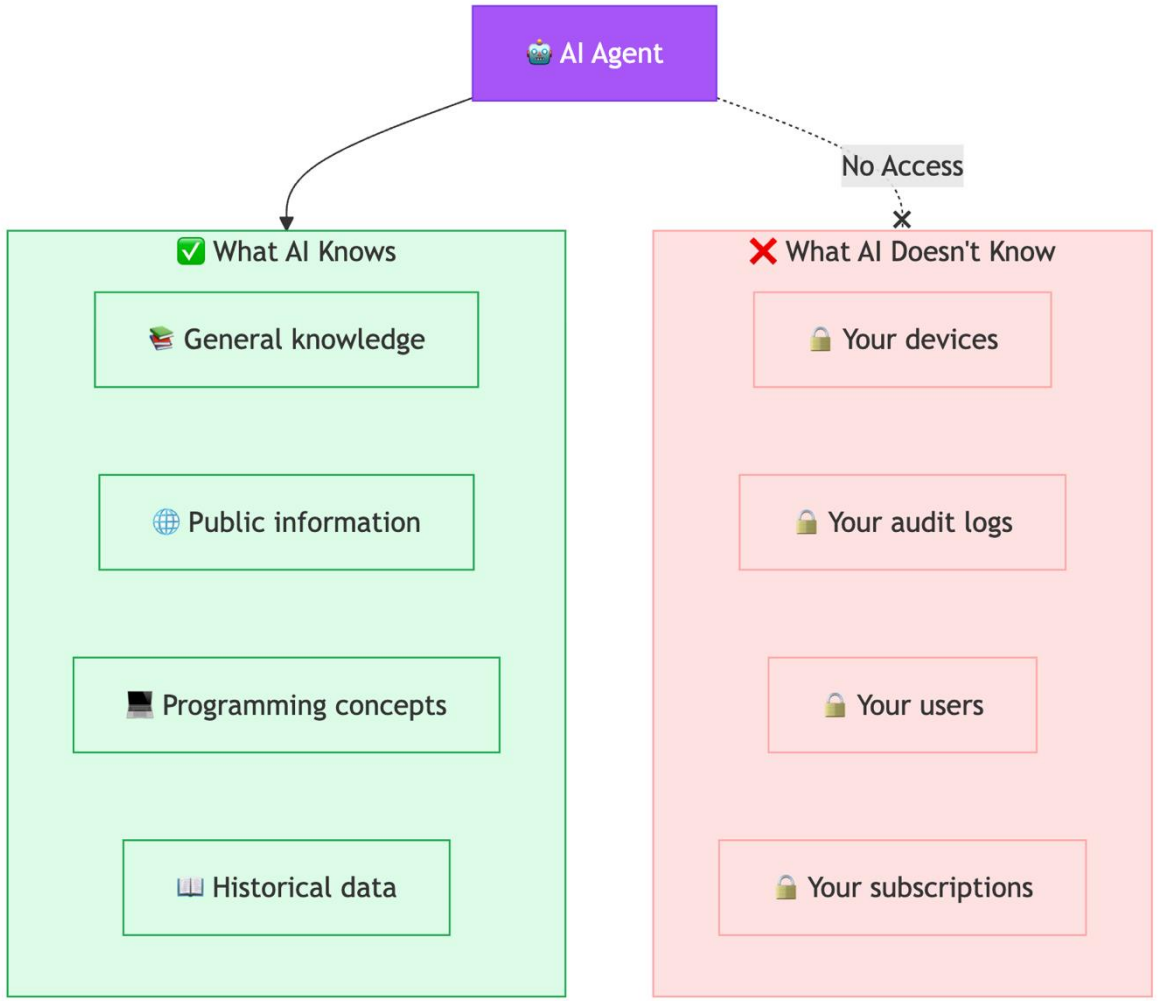

AI Models Are Trained on Static Data

- Wikipedia, books, articles, code from the internet.

AI lacks access to live, real-time enterprise data.

- What servers you have
- Who logged in yesterday
- When your subscriptions expire
- What changes happened last week

Model Context Protocol

- Before MCP no standard protocol for AI agents to connect to systems
- MCP was created by Anthropic in Nov 2024
- One of the fastest growth repo in GitHub history
- Rapid adoption by all major players
- AWS, Azure and others become part of the Core Committee
- Okta and others influenced evolution of MCP authorization specification.
- Become de facto standard in industry
- In December 2025 Anthropic donated MCP project to newly formed Agentic AI Foundation under Linux Foundation.
- Google, AWS, OpenAI, Anthropic, Cloudflare and others become founding members of the new foundation.
- Specification continues to evolve

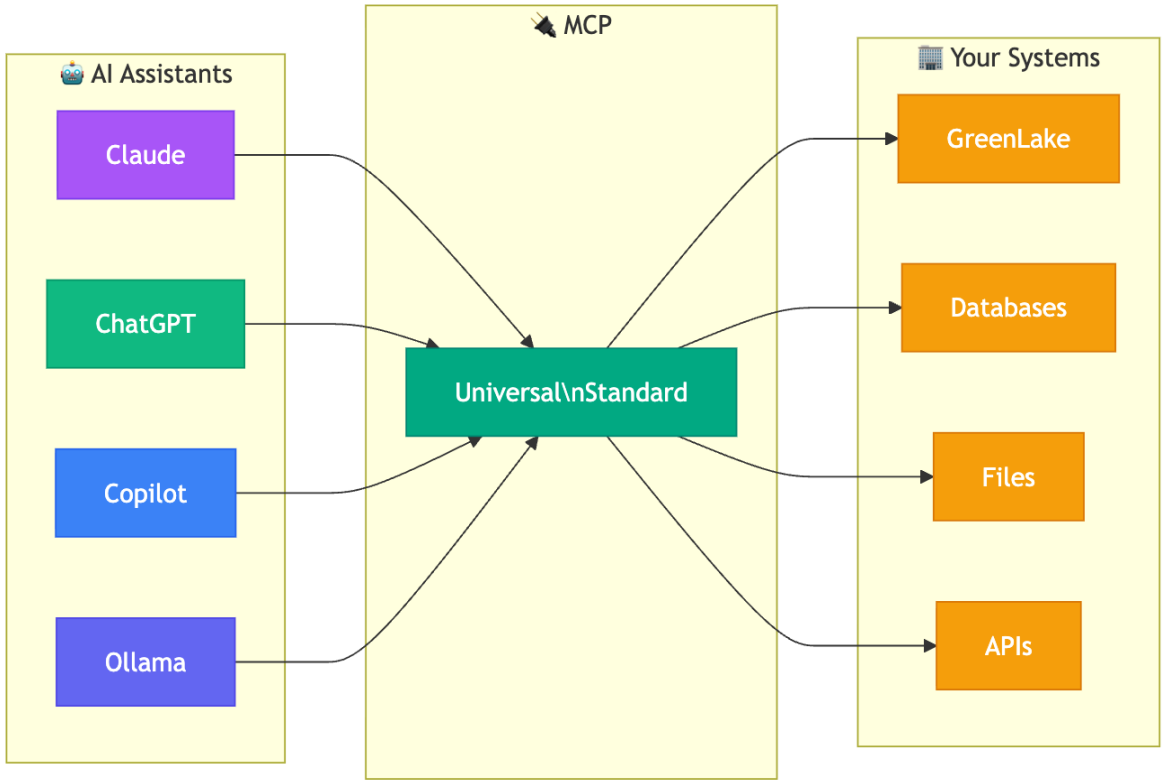

Natural Conversations, Real Answers

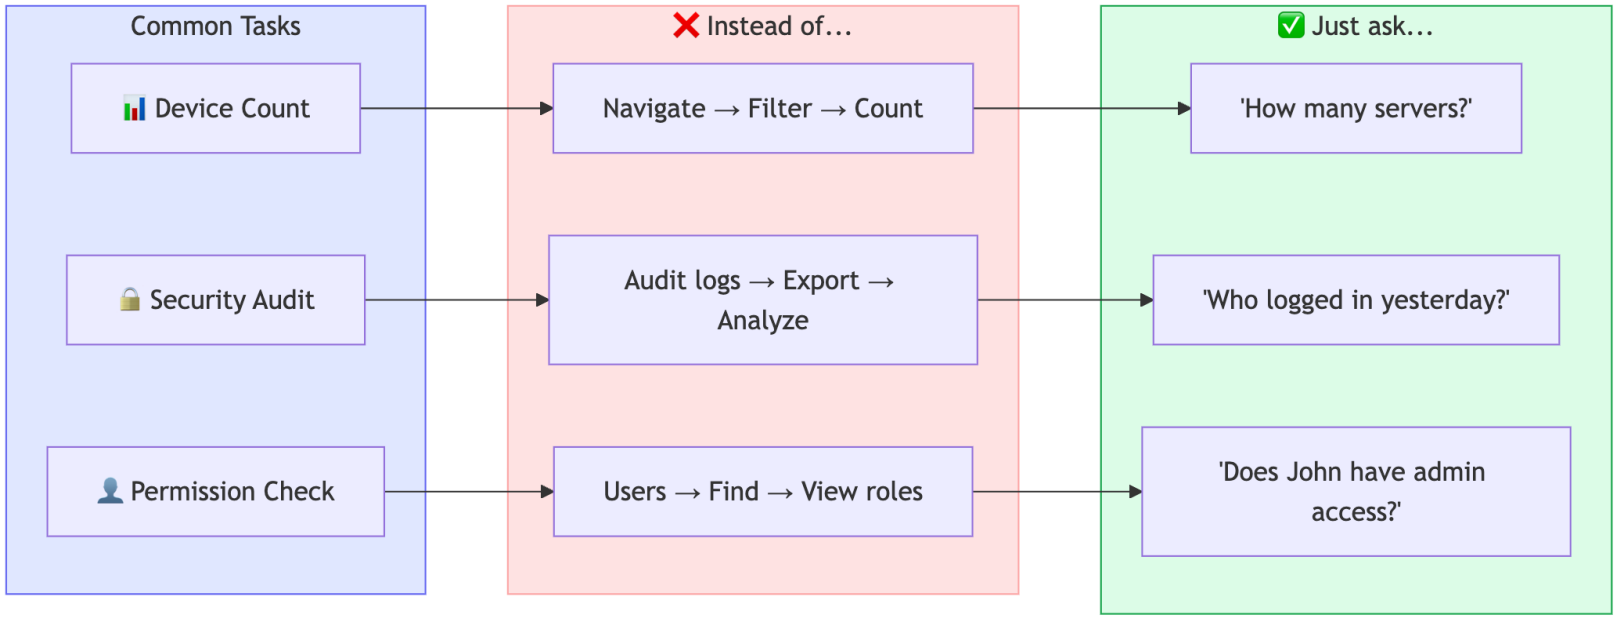

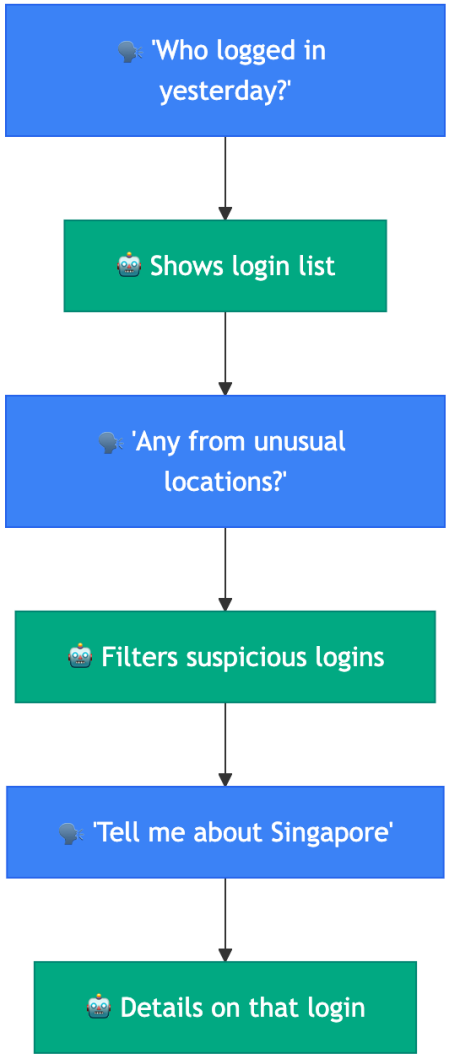

GreenLake MCP Servers

Talk to Your Infrastructure

Our Approach

The Debate - Many argue MCP servers shouldn't be built on OpenAPI specs as REST APIs aren't always suitable for MCP servers

Partial Truth - Custom MCP built on new APIs works well for systems with limited authorization granularity or such as single user accounts.

Enterprise Reality - Hyperscalers or system such as GreenLake have granular authorization (RBAC, ABAC, tag-based, etc) enforced via REST APIs/SDKs

THE RISK - Building parallel APIs for MCP is doubled effort & expanded attack surface to maintain authorization parity

GreenLake Advantage - APIs adhering to GreenLake API standards, one credential per user per workspace for services which enables uniform MCP server delivery independent from actual service

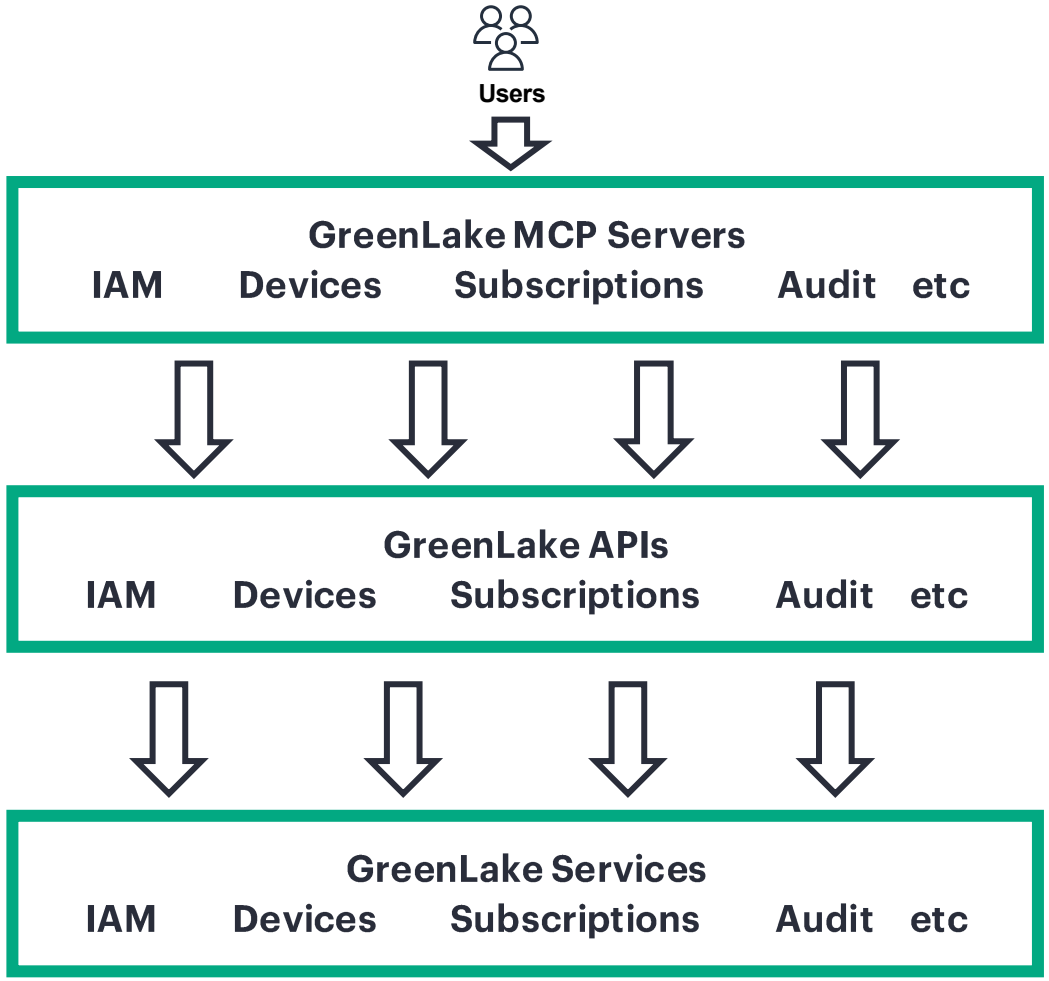

What HPE Has Build

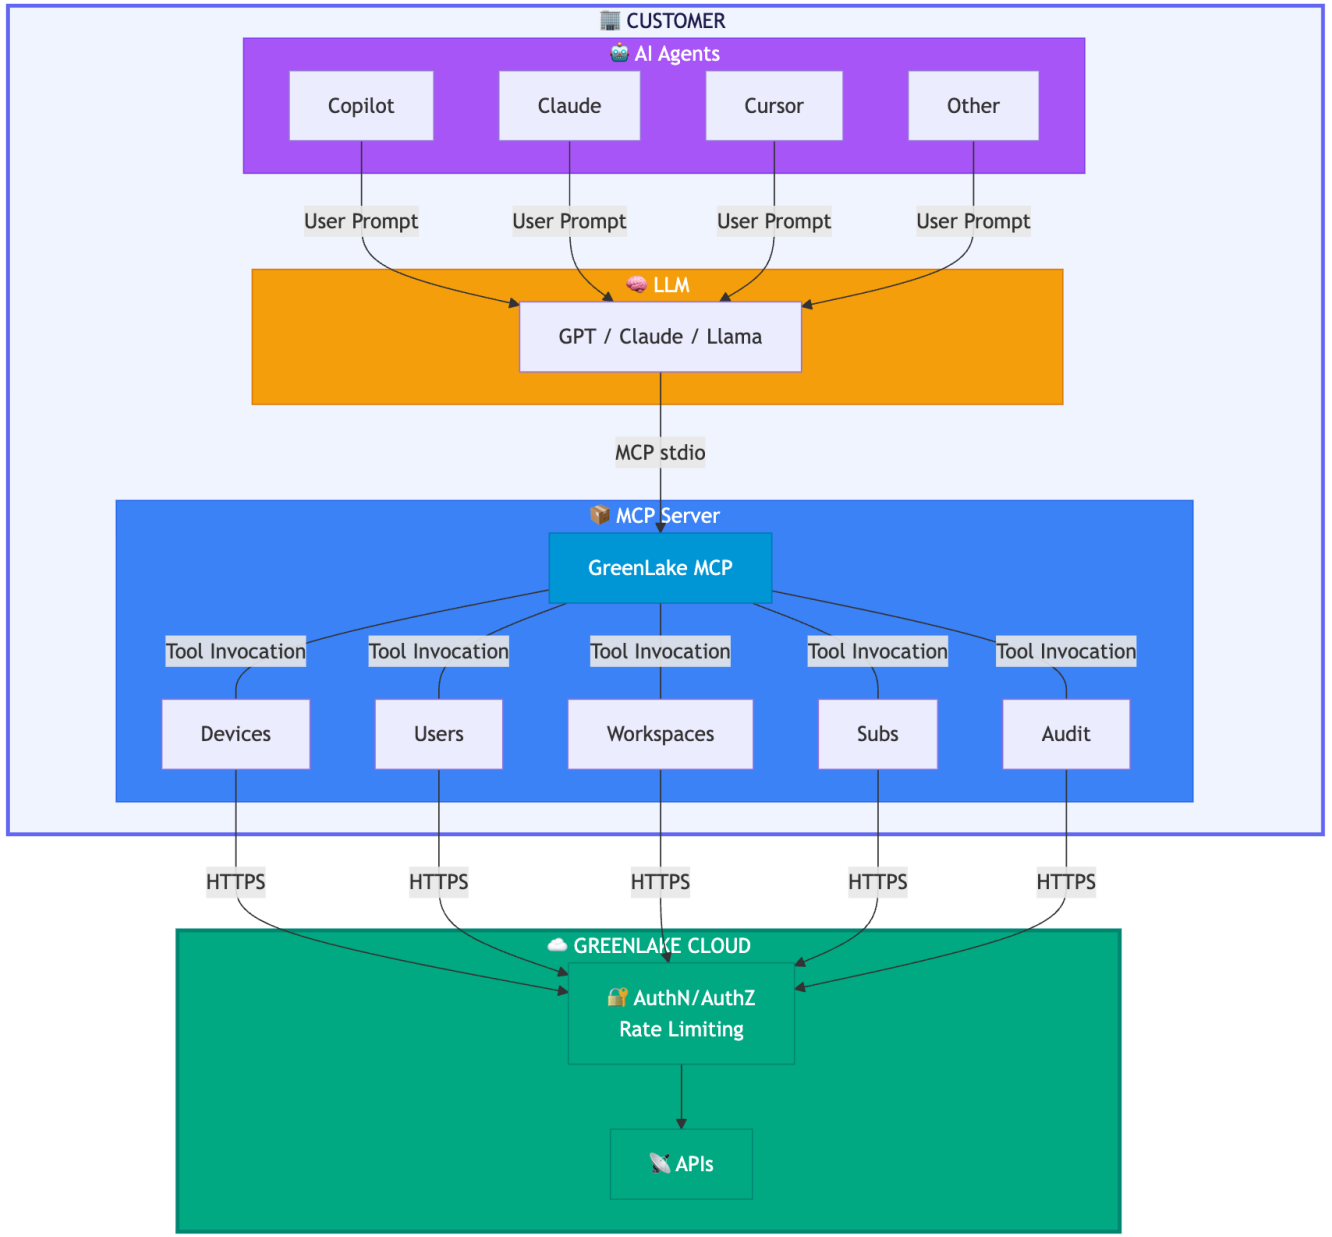

Runs in customer environment

- AI Agents: any MCP-compatible agent
- LLM Layer: AI models that understand natural language and decide which tools to use
- MCP Server: translates AI tool calls into API requests

Greenlake Cloud

- AuthN / AuthZ / Rate Limiting: Already in place in GreenLake Cloud — no new infrastructure needed, uses existing security controls
- REST APIs: Existing GreenLake APIs returning live infrastructure data

How MCP Works

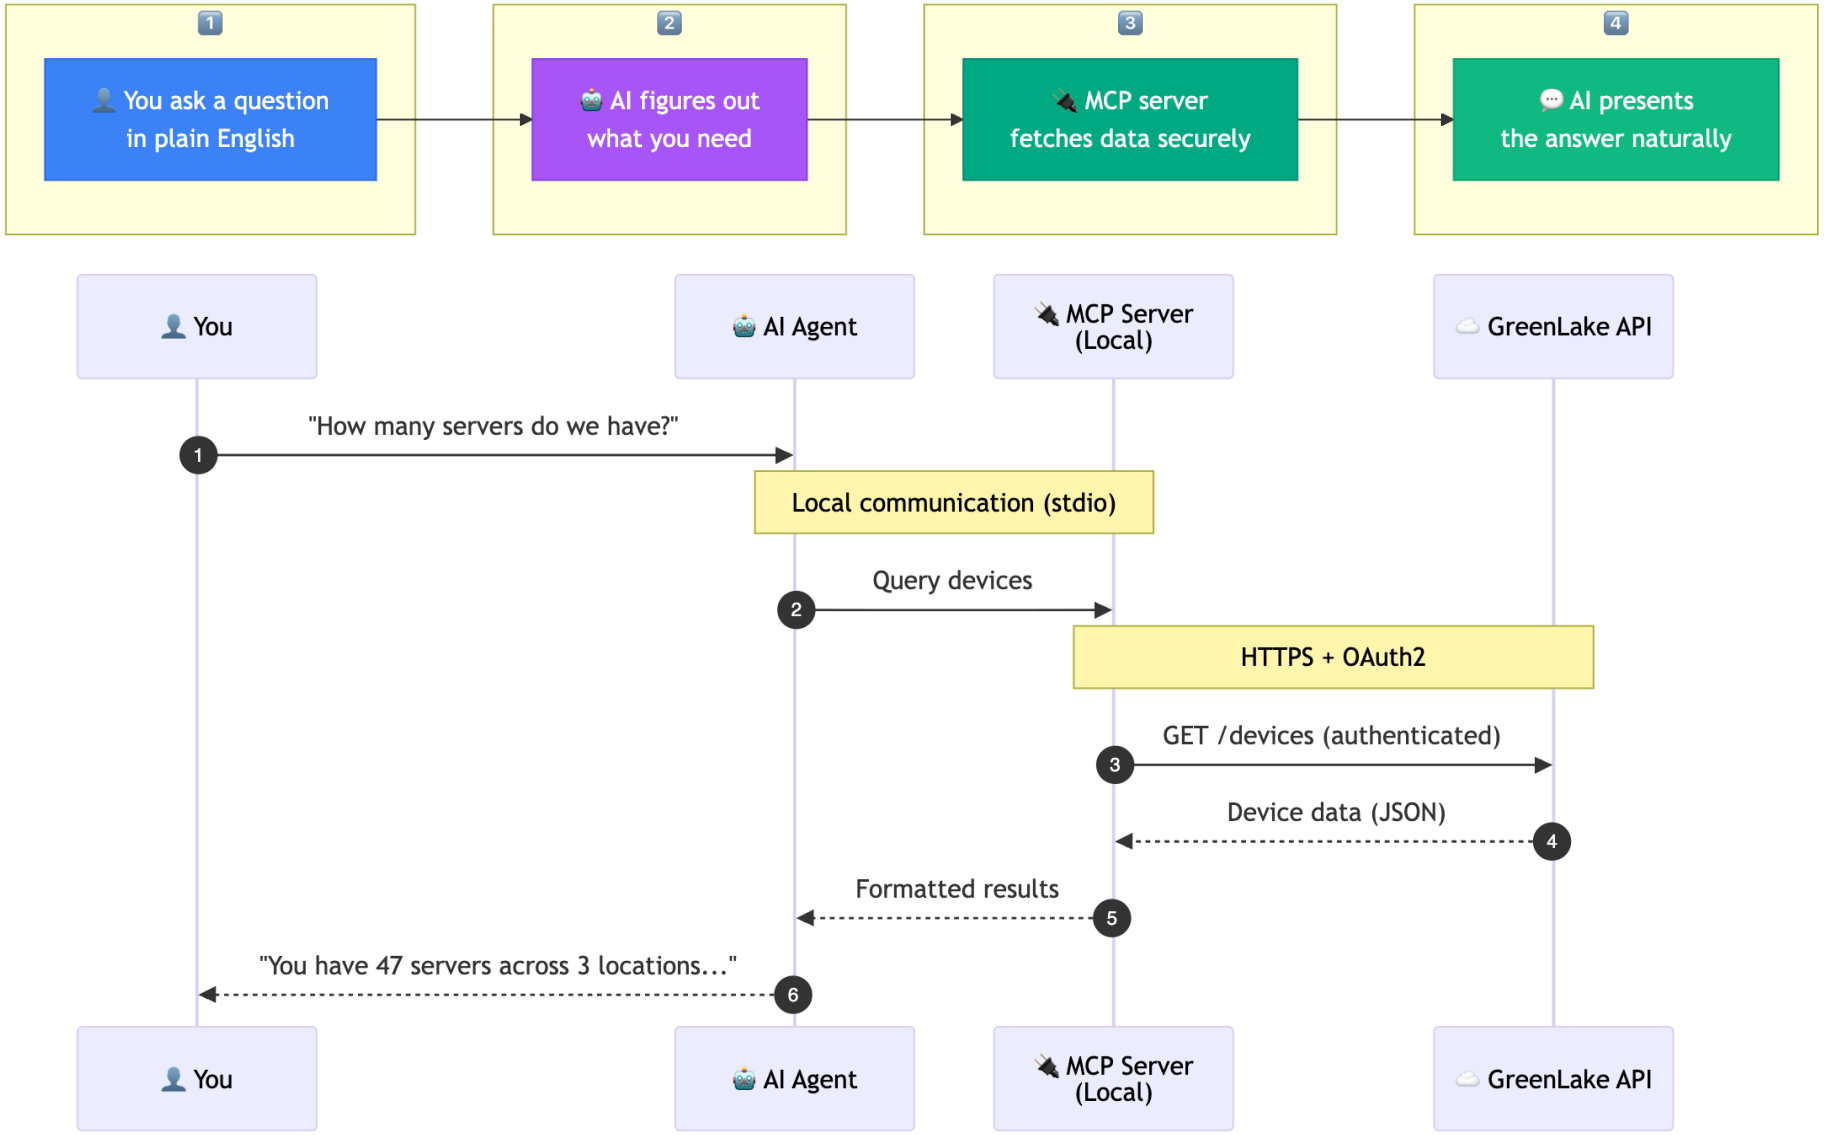

What is Available Today

DEVICES

Knows about your hardware inventory.

Servers, storage, network equipment.

USERS and WORKSPACE

Knows about people, their permissions, and organization structure.

AUDIT LOGS

Knows about who did what and when.

Your security and compliance trail.

SUBSCRIPTIONS

Knows about your licenses and renewals.

Is This Secure?

READ ONLY

AI can look at your data but cannot modify.

ENCRYPTED

All traffic uses HTTPS/TLS

YOUR PERMISSIONS

Uses your existing access

LOCAL

Runs on your machine

How We Build MCP Servers

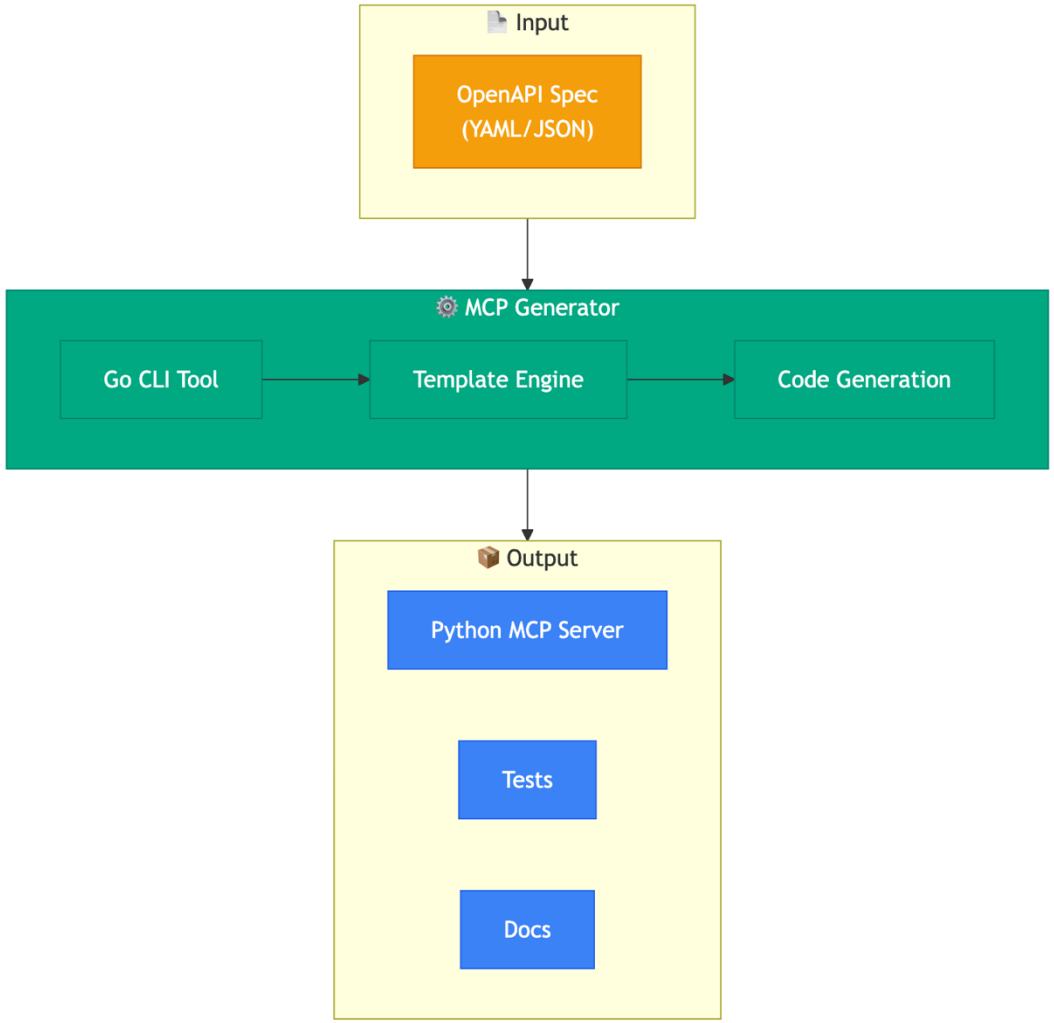

From OpenAPI spec to working MCP server in minutes.

Automatically creates:

— A complete Python MCP server

— All the tools mapped from API endpoints

— Authentication handling

— Tests and documentation

Testing with LLM in the Loop

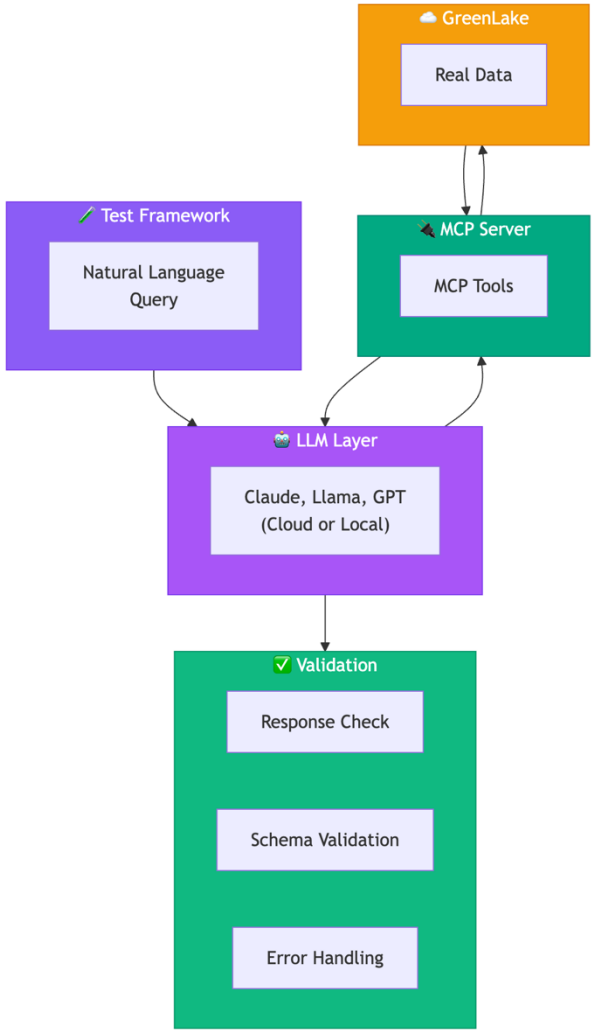

LLM-Powered Testing

Natural language tests that validate real-world usage.

- We send NATURAL LANGUAGE queries to the test framework
- The framework uses a real LLM (Claude, Llama, etc)
- The LLM calls the MCP server tools
- The server queries real GreenLake data
- We validate the responses

Why LLM Testing Matters

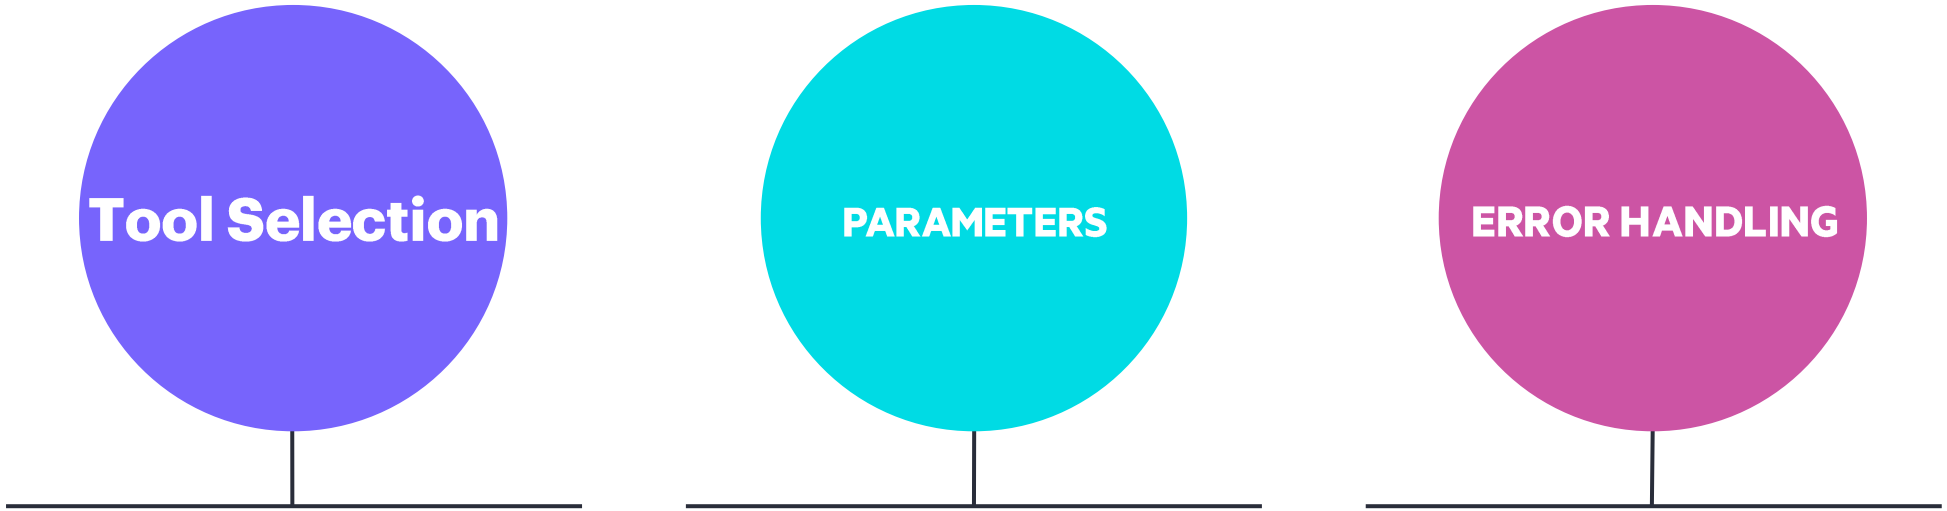

"Show me servers" should use the devices tool, not audit logs

"Servers in Dallas" should filter by location

Empty results, invalid queries, API errors

How We Keep MCP Servers Updated

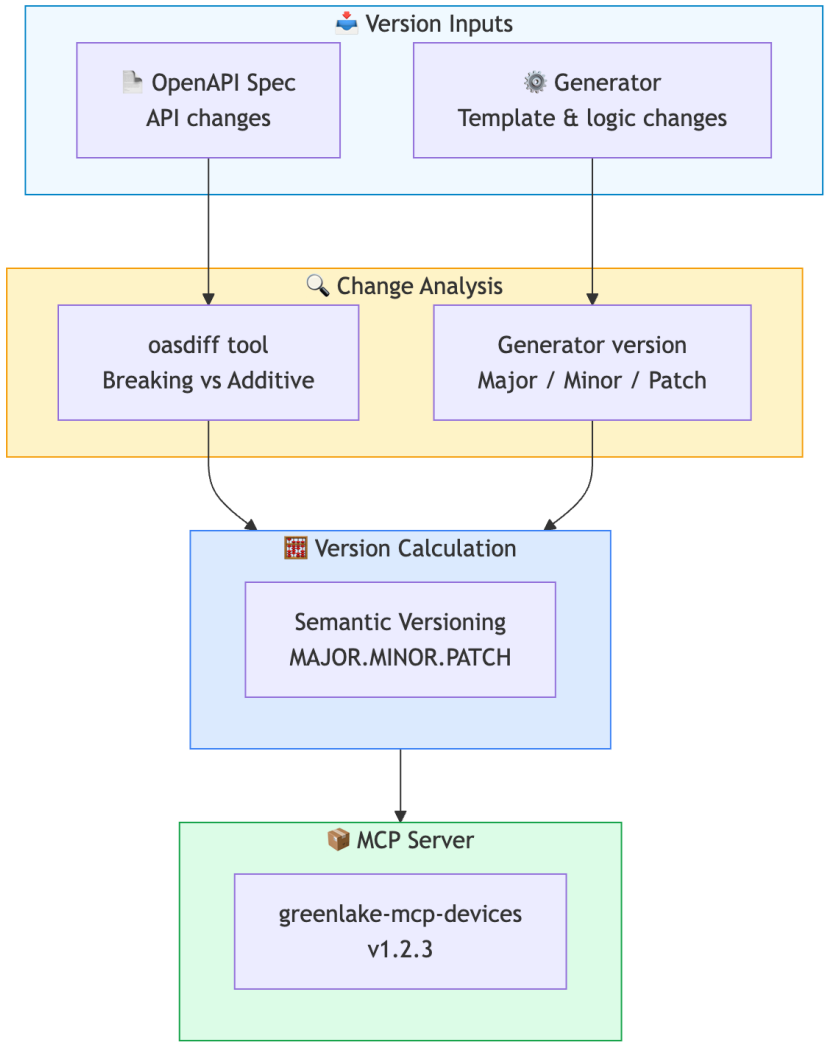

OpenAPI changes + Generator changes = Automatic version calculation

OPENAPI SPEC CHANGES

When the upstream HPE API changes, detection happens automatically

A tool called "oasdiff" analyzes what changed

It identifies: breaking change? New feature? Just a fix?

GENERATOR CHANGES

When the generator is improved (templates, auth, features)

Each generator release has its own version

The MCP server version is CALCULATED automatically:

Breaking change in API or generator → MAJOR bump (2.0.0)

New endpoint or feature → MINOR bump (1.3.0)

Bug fix or documentation → PATCH bump (1.2.4)

Live Demo

Try It Yourself

Visit GitHub, follow the guide, and talk to your infrastructure.

https://github.com/HewlettPackard/gl-mcp

https://developer.greenlake.hpe.com/docs/greenlake/mcp-server/public

GreenLake
MCP
servers

What's Next?

- Remote MCP Servers
- Write Operations
- Expand to More Services

Q/A

Thank You