



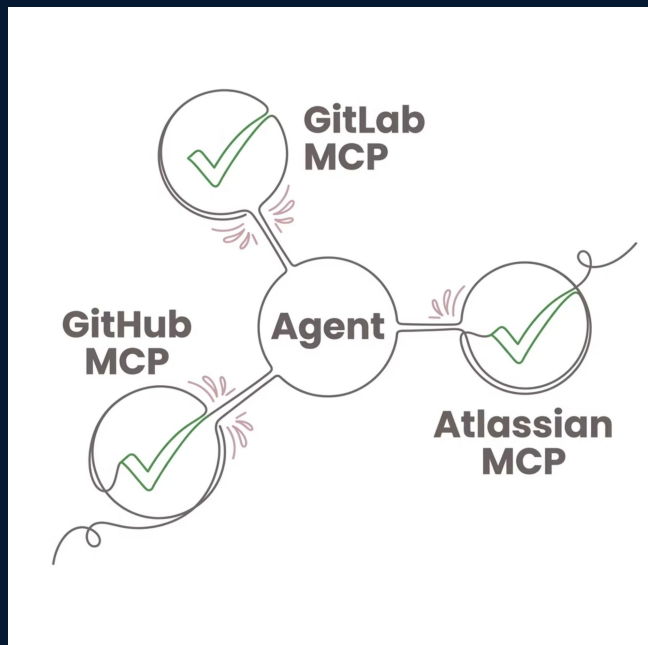
MCP  
Dev Summit  
Bengaluru

# From Alert to Revert

One MCP. 500+ tools for production triage and  
real-world DevOps

Presented by Avinash Lodhi · Architected by Steven Neisius  
Platform Engineering, Learneo

# To debug oncall issue, I **connected**



## Three MCPs. One agent. It worked.

GitHub, GitLab, and Jira integration was frictionless and promising.

### ✓ Agent connected. Schemas loaded.

- Listed issues and pull requests
- Checked repositories without friction
- First queries successful

*So naturally, I kept going...*



# Then I Connected a Dozen.

30K

Tokens burned

Before typing a single character – all from tool schema definitions

12+

MCP servers

Kubernetes, AWS, Datadog, Confluence, and more

~0

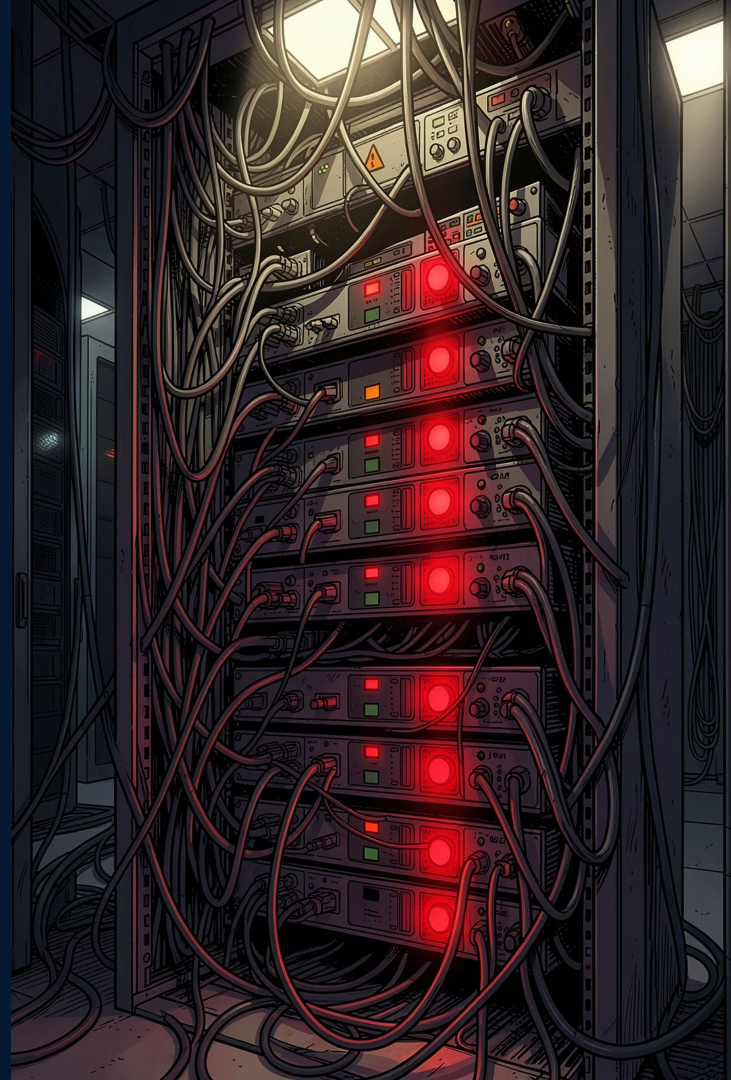
## Useful Context Remaining

Context window bloated from schemas alone – wrong tools picked, wrong answers given.

Not a model problem. A **context bloat** problem – the agent was reasoning through a **fog of schemas** it would never use.



MCP  
Dev Summit  
Bengaluru



# And Every MCP Want Its Own way to connect .

## GitHub

GITHUB\_PERSONAL\_ACCESS\_TOKEN

## GitLab

GITLAB\_PERSONAL\_ACCESS\_TOKEN

## Atlassian

JIRA\_API\_TOKEN + JIRA\_EMAIL + JIRA\_URL

## AWS


4-variable credential chain

## Kubernetes

kubeconfig mount

## groundcover

OAuth

 **Credential sprawl**, increases with every new addition, hard to manage .



# 🤖: Where should I look?

Learned: 7 product lines. Each with its own stack.

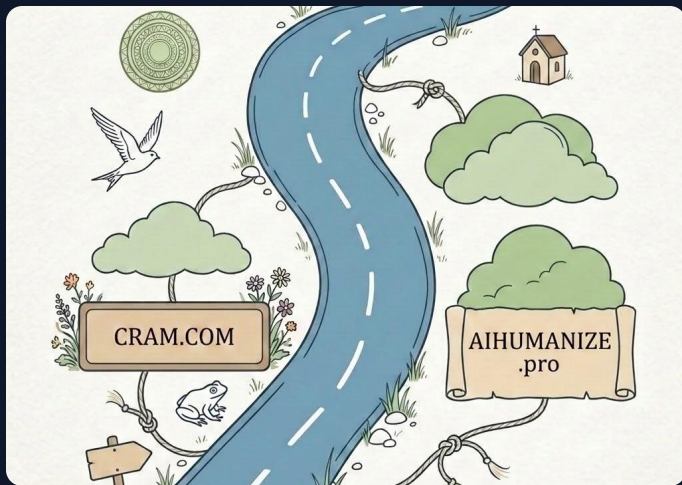
Product	Code Host	Cloud	Monitoring
Course Hero	GitLab (self-hosted)	AWS GCP	Groundcover Datadog
QuillBot	GitLab (cloud)	GCP	Datadog
Scribbr	GitHub	AWS	Cloudwatch, Datadog
LitCharts	GitHub	AWS	Cloudwatch, Datadog
StudentBrands	Bitbucket	AWS	Groundcover
LanguageTool	GitHub	AWS	Cloudwatch



**3 code hosts. 2 cloud providers. 20+ cloud accounts. 7 monitoring configs.** When an engineer asks "check the latest deployment" – which codebase? Which cluster? Which dashboard? The agent has no idea, and your prompt can't explain all of this every time.



# A Moving Target: Evolving Enterprise Context



Learnéo's product lines are not static; they're constantly evolving in the AI era.

## New Products = New Context

Additions like Cram.com or AIHumanize.pro bring unique stacks, auth, and monitoring requirements.

## Context Bloat Returns

Agents struggle without specific enterprise context, even with large windows or selective MCPs.

## The Agent's Blind Spots

Lack of context leads to incorrect assumptions, constant clarification, and hallucinations.

The agent needs an **up-to-date map** of the enterprise to act intelligently, not just more raw data.



# Too many install commands for team onboarding

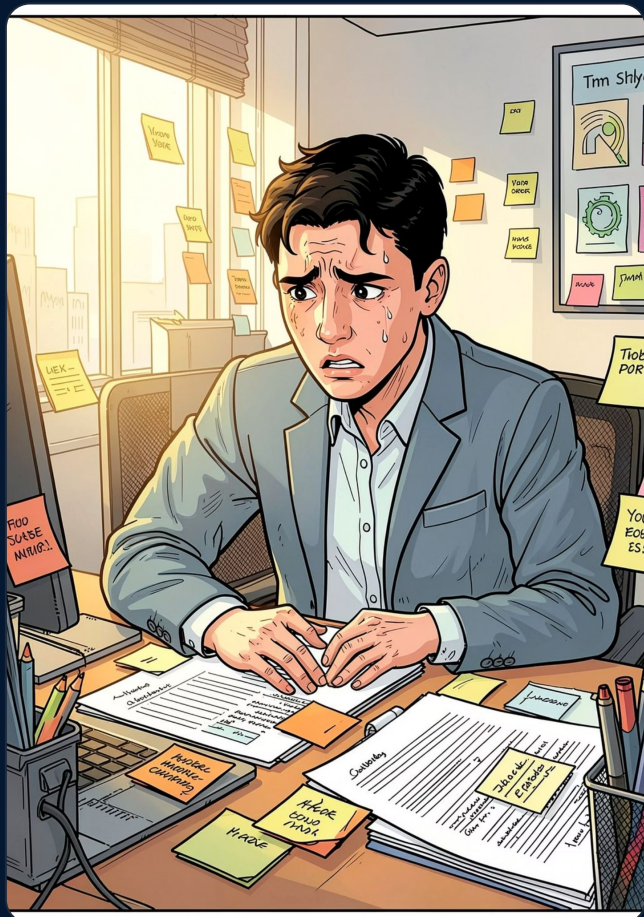
01 `npm install -g @modelcontextprotocol/server-github`

02 Download binary. Verify SHA256.

03 `pip install via uvx`

04 Set 14 environment variables across 6 different auth shapes

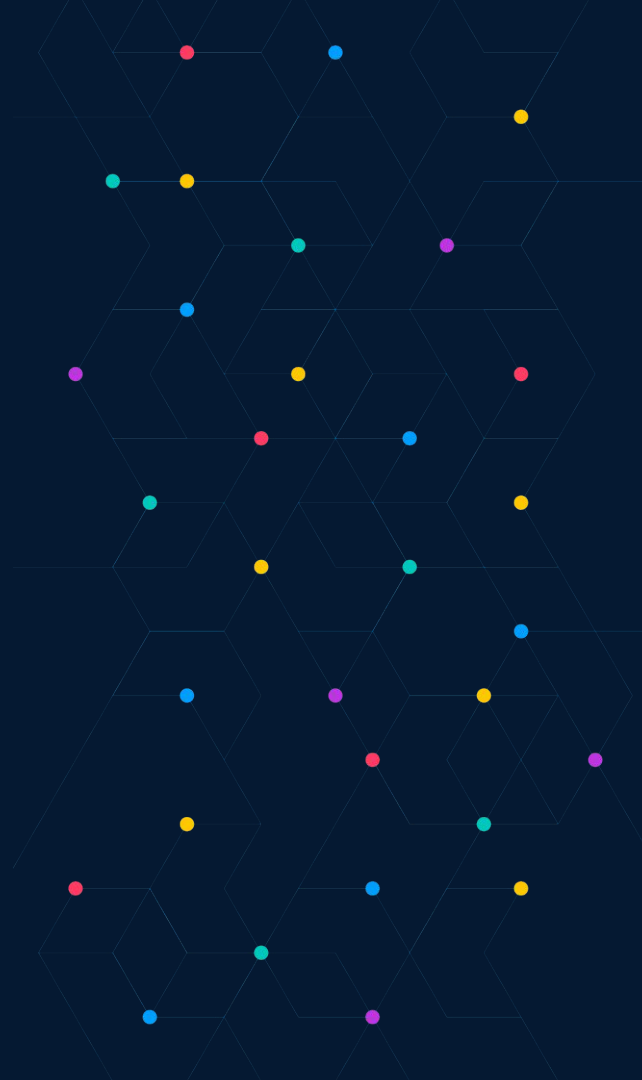
Average time to working setup: **3-4 hours**. Some engineers never finish. Every new upstream made the list longer.



# Solution

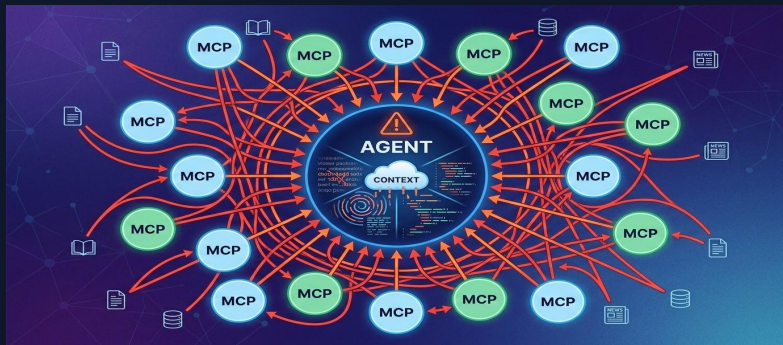


**MCP**  
Dev Summit  
Bengaluru



# One Aggregator. All the Tools. One Connection.

## BEFORE: CHAOTIC DIRECT CONNECTIONS

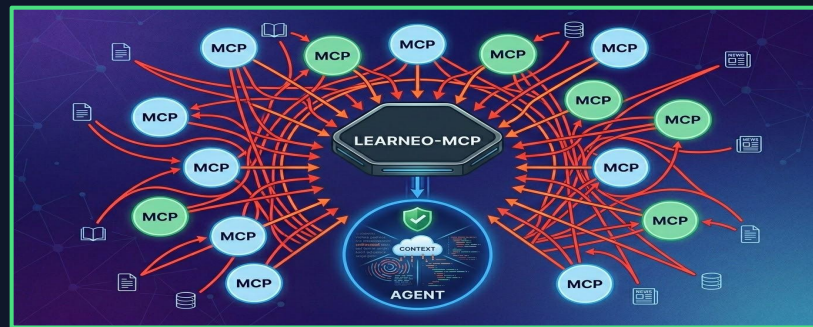


**~30,000 tokens**

Context burned at connect time

**12+ connections**

## AFTER: CLEAN AGGREGATION



**~1,500 tokens**

20× reduction in overhead

**1 connection**

Single unified mcp server to the agent



# What the Agent Actually Sees

Not 500 schemas. **Six meta-tools.** 1,500 tokens. Every upstream tool still reachable at runtime.



## `learneo_search_tools`

BM25 search over the full tool catalog



## `learneo_describe_tools`

Pull full schema for any tool, on demand



## `learneo_execute_tool`

Call any upstream tool through the aggregator



## `learneo_get_company_context`

The full enterprise tech map: all 9 product lines, their code hosts, and cloud providers.



## `learneo_setup_guide`

Self-service credential setup for engineers



## `learneo_debug_status`

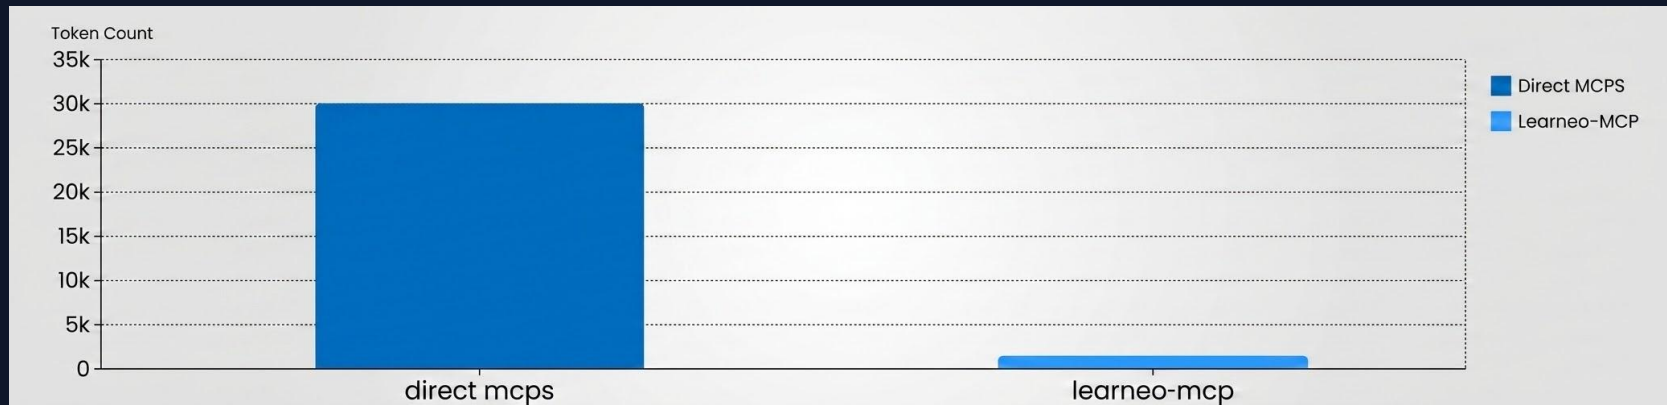
See what's connected and what's healthy



# Constant tokens. Infinite tools.

~1,500 tokens at connect – always.

Regardless of how many upstreams are configured.



## How the agent accesses 500+ tools:

```
learneo_search_tools("error rate logs groundcover")  
learneo_describe_tools(["groundcover_query_logs"])  
learneo_execute_tool("groundcover_query_logs", {...})
```

## Search is BM25:

- name (3×)
- tags (2×)
- description (capped at 150 tokens)
- upstream prefix boost



MCP  
Dev Summit  
Bengaluru

# Demo

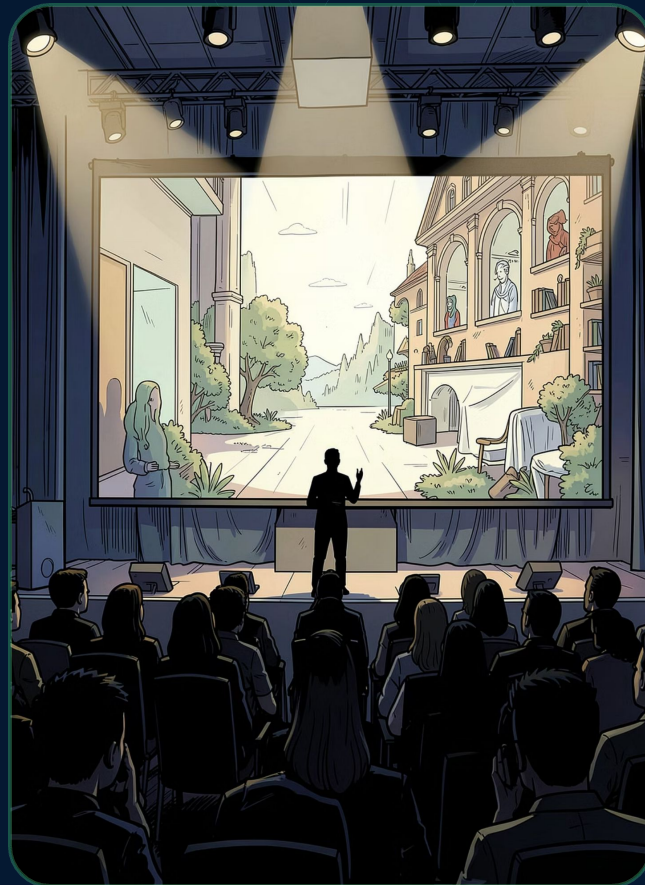
Watch the agent go from Slack alert to reverted commit.

General prompt, zero tool hints, no infrastructure briefing.

**Let's see learneo-mcp in action**



**MCP**  
Dev Summit  
Bengaluru



Home



DMs



Activity



Files



Later



More



# devops-alerts Course He...

# devops-alerts-1 Course ...

# devops-ch Course Hero S...

# devops-debug Course H...

# devops-discuss Course ...

# devops-merge-requests...

devops-team Course Her...

# devops-warn Course He...

# eng-announcements Co...

# engineering Course Hero...

# engineering-discuss Co...

# high-priority-alerts Cou...

# high-priority-response...

# high-priority-warn Cou...

# itopshelp

# learneo-updates Course ...

# lifeatlearneo Course Her...

# local-remote Course Her...

# news Course Hero Slack

# org-random Quillbot

# ph...

More unread

# platform-helpdesk



Datadog APP 3:32 PM

Triggered: AWS Personal Health Dashboard Event

@slack-devops-alerts

DD AWS Health Event Dashboard

AWS PHD Dashboard

1 event triggered this monitor, here is the last one.

---

Show more

Notified

@slack-devops-alerts



Groundcover APP 6:53 PM

Firing: Testing Mixin Job Failures in empty namespace

Silence | Investigate | See Monitor

This alert fires when the total number of failed Pods across all jobs matching testing-mixin.\* in the empty namespace are more than 0. Each failed Pod represents a container that terminated unsuccessfully within a job.

groundcover.com | Today at 6:52 PM | Added by Groundcover

B I U

Message #devops-alerts

+ Aa @

# The Agent Did All of This.

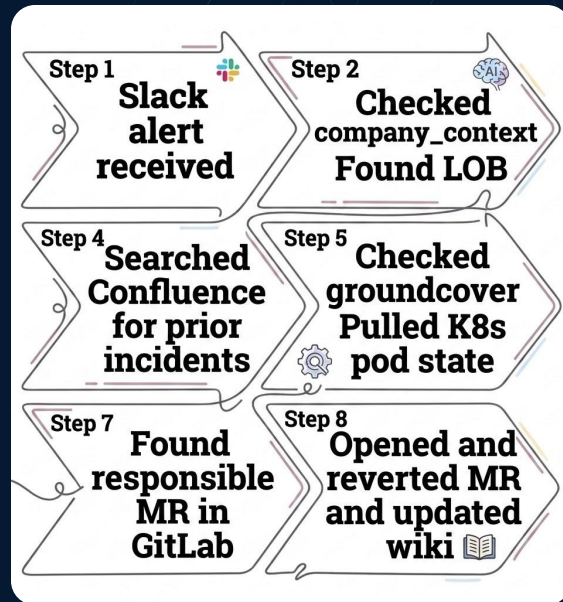
One conversation. One short prompt.

The agent navigated complex infrastructure seamlessly:

- Identified correct monitoring tools and clusters without explicit instructions
- Analyzed prior incidents to inform current resolution strategies
- Found the root cause and reverted it

## Without learneo-mcp:

custom scripts per LOB, or 14 MCPs burning your context budget, or a 500-word prompt explaining your stack on every tun



# The agent knew.

## COMPANY-OVERVIEW

**Course Hero** GitLab (self-hosted), AWS, Groundcover, Datadog

QuillBot GitLab.com GCP Datadog

Scribbr GitHub AWS Cloudwatch, Datadog

LanguageTool GitHub AWS Cloudwatch

## How it worked in the demo:

- Agent called `learneo_get_company_context`
- ID from metadata → deduced **CourseHero LOB**
- Independently knew: Groundcover, gitlab, the right EKS cluster

*I gave a one-line prompt. It didn't ask clarifying question.*



## MCP instructions field:

LOB → code host, monitoring, cluster, AWS account

# Setup in minutes, not hours.

## Central deployment

Read-only tools → credentials already there.

Connect once. Gated by Okta OAuth. Use immediately. Nothing to install.

## Write tools / personal tokens

`learneo_setup_guide`

- Agent asks which upstreams you want
- Tells you exactly which tokens to get
- Copy-paste them into one `.env` file
- Agent handles the rest

Pick and choose which upstreams to enable. No format chaos. No 14-step checklist.

Before: ~~3-4 hours~~ ~~14 env vars~~ ~~format chaos~~ **After: ~10-20 minutes · one .env file**



# Learning from past incidents

## A Confluence space owned by the agents

Central repository for all operational knowledge, accessible and modifiable by the AI agent.



### Confluence wiki

Org structure, data models, and "where-to-look" navigation pages.



### LLM Wiki

Catalog of past investigations and technical post-mortems for reference.



### Skills

Reusable playbooks for complex tasks (e.g., automated EKS upgrades).

Agent reads during triage · Writes what it learns · **Incident response gets faster every time**



# Architected by Steven Neisius

Architect and creator of Learneo MCP.

- ✓ Designed the meta-tool pattern.
- ✓ Coded the aggregator end to end.
- ✓ Set the security model for company-wide deployment.

*Everything you saw today rests on his architecture.*

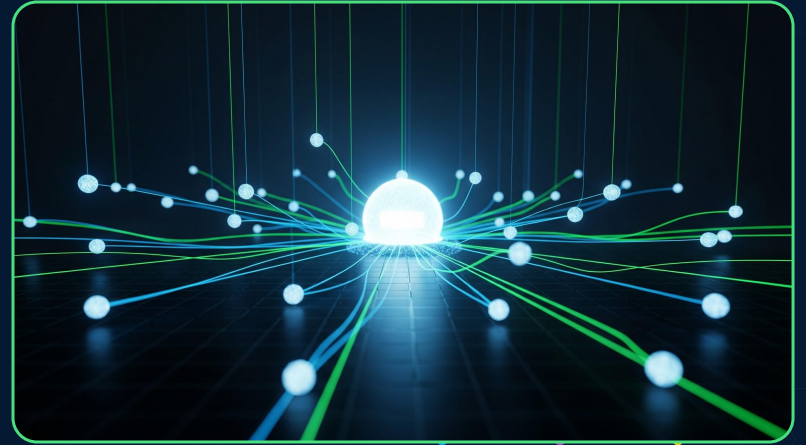


# Learneo MCP



Special thanks to:

**Adam, Eilidh, Manshubh, Rich, Jing & the whole team.**



**MCP**  
Dev Summit  
Bengaluru