AI-Driven Platform Engineering with MCP

Making AI a Context-Aware Operational Partner

Platform Engineering Summit 2025

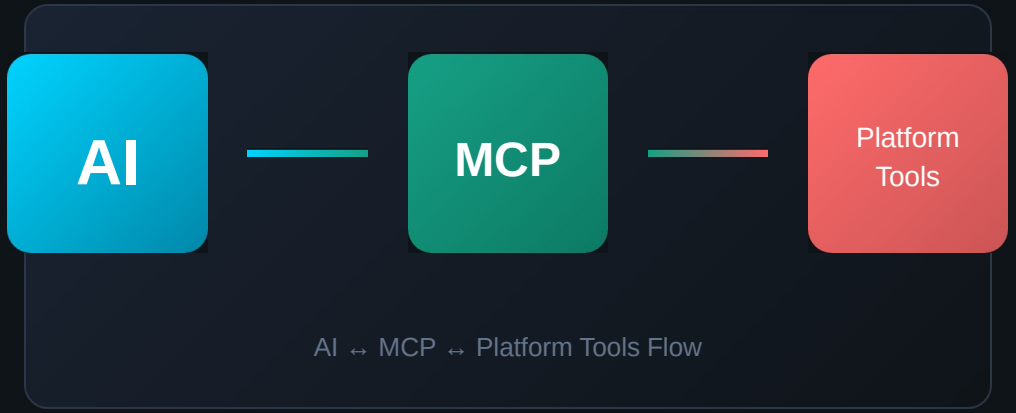
Session Overview

PURPOSE

Show how AI can transform platform engineering workflows

KEY TOPICS

1. Leveraging AI for operational efficiency
2. MCP servers across cloud-native stack
3. Building custom MCP servers



Why AI for Platform Engineering?

CHALLENGES

- ✗ Manual toil in infrastructure and observability
- ✗ Slow incident response times
- ✗ Complexity across multiple tools
- ✗ Context switching overhead

AI OPPORTUNITIES

- ✓ Natural-language queries across systems
- ✓ Context-aware automation
- ✓ Faster debugging and insights
- ✓ Unified interface for all tools

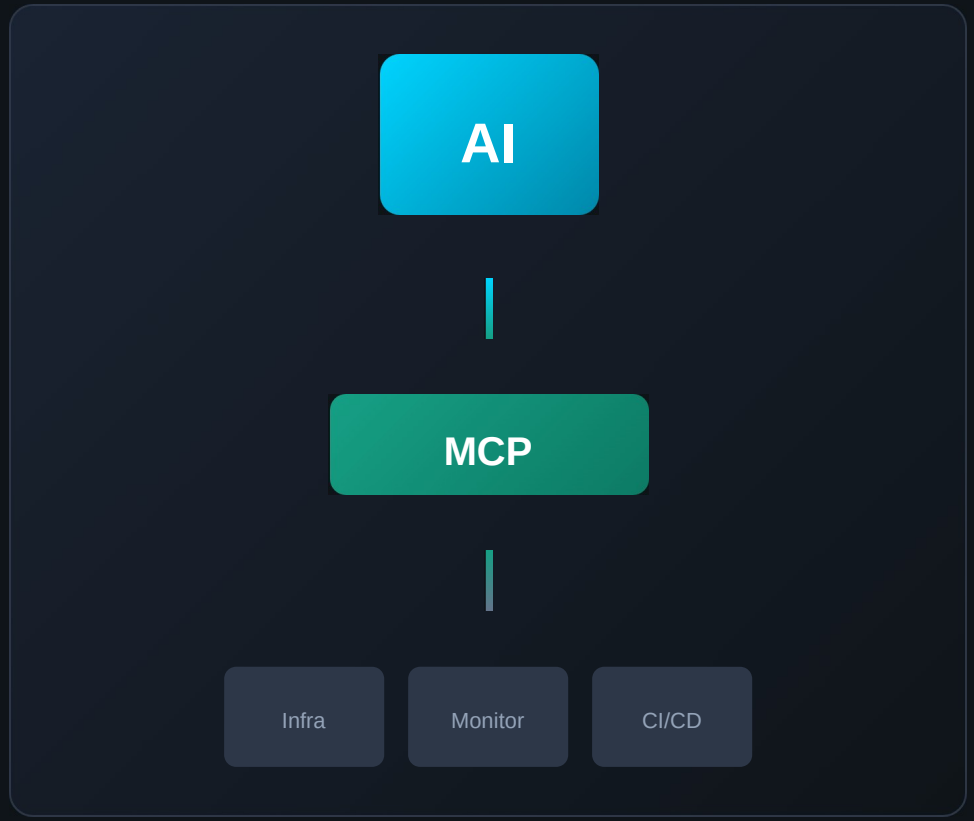
MCP – The Key Enabler

WHAT IS MCP?

Model Context Protocol – enables AI to access operational context securely and in real-time

ROLE IN PLATFORM ENGINEERING

- Exposes infrastructure, monitoring, and CI/CD data to AI
- Acts as a bridge between AI models and live systems
- Secure, standardized protocol for tool access



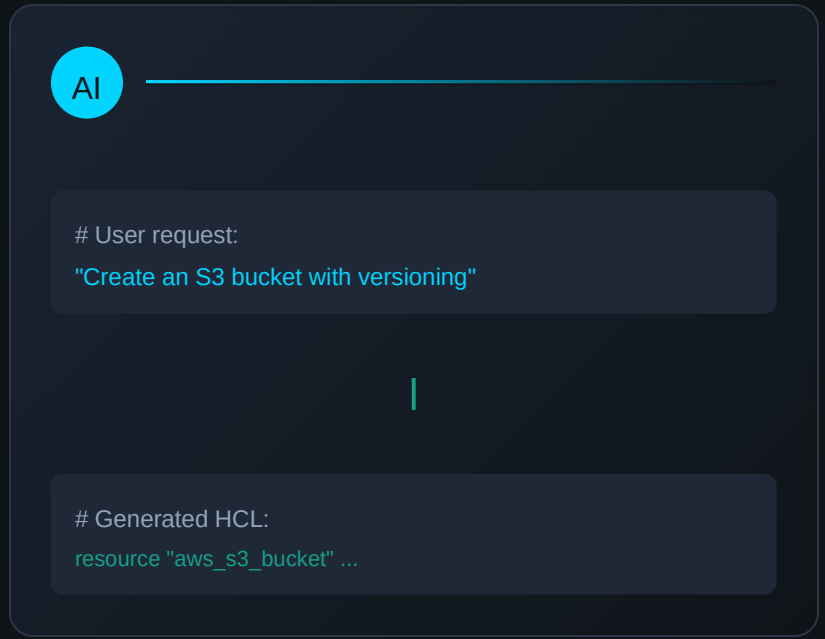
Terraform MCP Server

USE CASE

Intent-driven Infrastructure as Code

CAPABILITIES

- Generate Terraform modules from natural language
- Analyze plan output in natural language
- Detect drift and unsafe changes
- Suggest remediation strategies



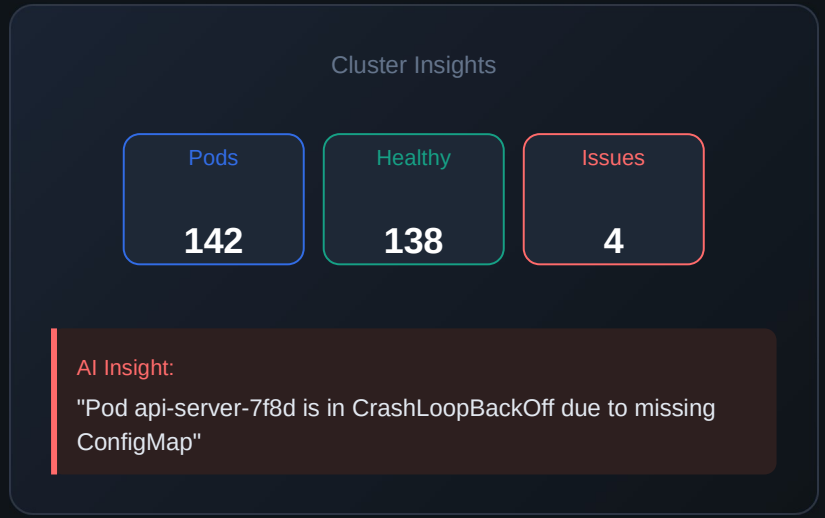
Kubernetes MCP Server

USE CASE

Debugging cluster workloads

CAPABILITIES

- Inspect pod state, logs, and events
- Diagnose CrashLoopBackOffs
- Identify configuration issues
- Suggest scaling and resource adjustments



Grafana MCP Server

USE CASE

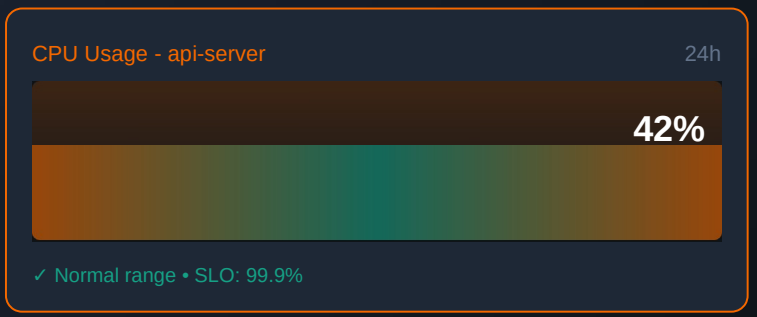
Live system monitoring

CAPABILITIES

- Query metrics using natural language
- Generate dashboards automatically
- Summarize trends and anomalies
- Track SLO compliance

User Query:

"Show me CPU usage for api-server last 24h"



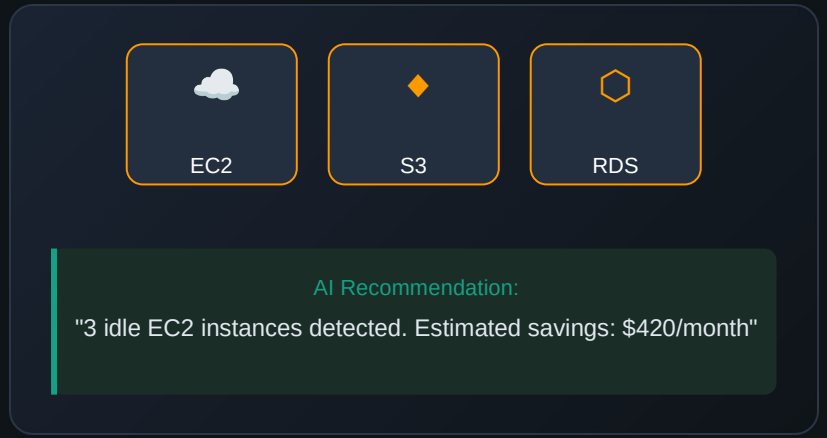
AWS MCP Servers

USE CASE

Cloud resource inspection and optimization

CAPABILITIES

- Explore EC2, S3, RDS configurations
- Identify misconfigurations
- Detect compliance gaps
- Suggest cost optimizations



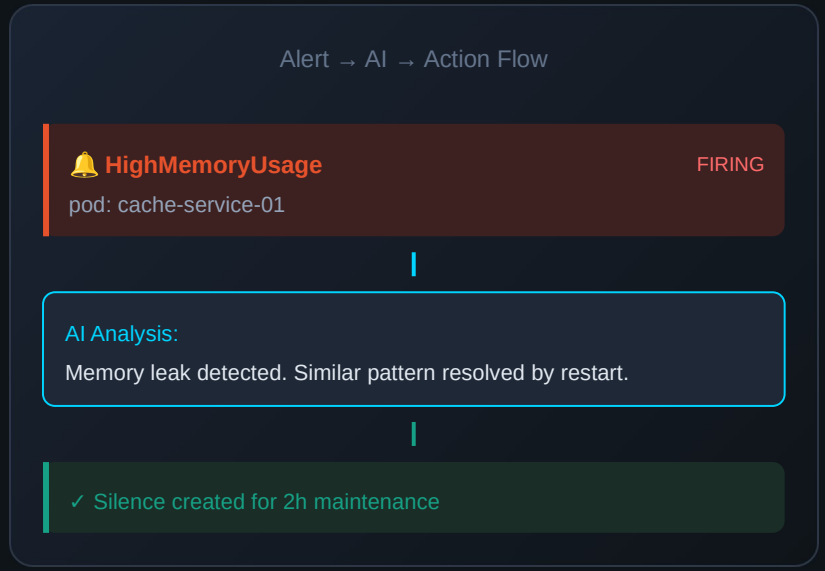
Alertmanager MCP Server

USE CASE

Intelligent alert handling

CAPABILITIES

- Inspect and triage active alerts
- Reduce noise with smart silences
- Provide contextual recommendations
- Correlate alerts across systems



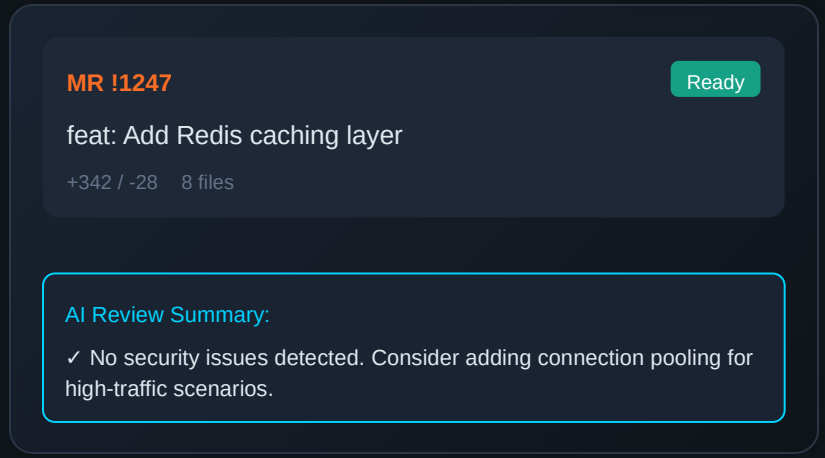
GitLab MCP Server

USE CASE

Version control and CI/CD context

CAPABILITIES

- Review merge requests with AI
- Track infrastructure changes
- Provide deployment insights
- Pipeline status and analysis



Fast

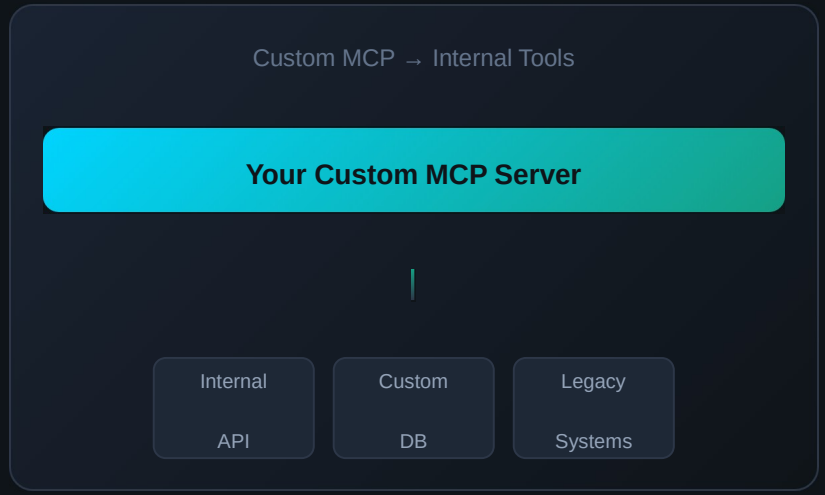
FastMCP

USE CASE

Build custom internal MCP servers

CAPABILITIES

- Extend AI access to proprietary tools
- Standardized API for internal platforms
- Built-in security and access control
- Python framework for rapid development



Key Takeaways

1 AI as Interface

AI becomes a first-class platform interface for all operations

2 Practical MCP

Real use cases across Terraform, K8s, Grafana, AWS, Alertmanager, GitLab

3 Extend It

Build custom MCP servers to extend AI capabilities internally

Session Objectives

1

Demonstrate AI Integration

Show how AI integrates with platform tools using the Model Context Protocol

2

Real-World Workflows

Present practical use cases and workflows you can apply immediately

3

Enable Custom Integrations

Empower you to design and build your own MCP-based integrations

Who Should Attend

Platform Engineers
Building and maintaining internal platforms

SREs
Managing reliability and incident response

DevOps Engineers
Automating deployment and operations

Infrastructure Developers
Curious about AI in infrastructure



No Prior Experience Required

MCP knowledge not required to attend and benefit from this session

Questions?

Let's Discuss



AI Integration Protocol Platform

Thank you for attending!