THE LINUX FOUNDATION
OPEN SOURCE SUMMIT
NORTH AMERICA

Embedded Linux Conference

Taming MCP Server Sprawl:
Securing and Scaling the Model Context Protocol in Production

Jeffrey Borek
Sr Program Director
Open Technologies
IBM
jborek@us.ibm.com

Olivia Buzek
Lead Developer Advocate
AI Technologies
IBM
ombuzek@us.ibm.com

The Rise of AI Agents in the Enterprise

From Pilots to Production

- Agents are rapidly moving beyond PoCs into real production workflows
- Organizations are now connecting LLMs to APIs, business systems, developer tools and operational data to automate complex tasks
- They need a ”standardized” way for models to discover and interact with external capabilities **safely and reliably**

Enter the Model Context Protocol (MCP)

- MCP is an open-source “standard” introduced by Anthropic in November 2024

- Using MCP, AI applications like Claude or ChatGPT can connect to data sources (e.g.. local files, databases), tools (e.g., search engines, calculators) and workflows (e.g., specialized prompts)—enabling them to access key information and perform tasks

Understanding the Model Context Protocol

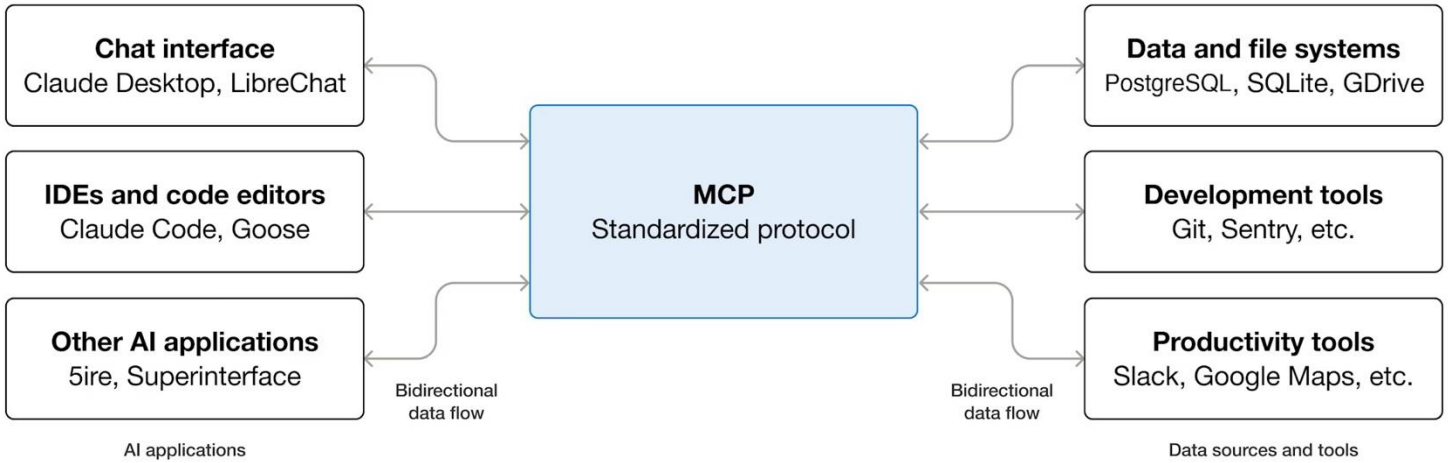

https://modelcontextprotocol.io/docs/getting-started/intro

MCP Adoption Statistics 2026

9.4K
Public MCP Servers

41%
Custom MCP Server

78%
Enterprise Teams w/ MCP

Vertical and Use-Case Adoption

MCP server distribution by category, computed across the 9,400+ registry-listed servers as of mid-April 2026:

Category	% of Registry	Representative Servers
Developer tools	32%	GitHub, GitLab, Linear, Sentry, Datadog
CRM and sales	14%	Salesforce, HubSpot, Pipedrive, Close
Data and analytics	12%	Snowflake, Databricks, BigQuery, Postgres
Documentation and wikis	11%	Notion, Confluence, Google Drive
Marketing automation	9%	HubSpot, Google Ads, Mailchimp, Customer.io
Customer service	7%	Zendesk, Intercom, Front
Search and retrieval	6%	Brave Search, Tavily, ElasticSearch
Other (cloud, finance, design, niche)	9%	AWS, GCP, Stripe, Figma, Airtable

Developer tools dominate the registry not because demand is higher there but because developer-tool teams shipped MCP support earliest (every IDE that adopted MCP in Q1 2025 motivated dozens of community servers in the same quarter). The CRM, data, and marketing categories grew fastest in late 2025 once first-party vendor servers landed.

https://www.digitalapplied.com/blog/mcp-adoption-statistics-2026-model-context-protocol#vertical-and-use-case-adoption

The Hidden Problem: MCP Server Sprawl

When Flexibility Becomes Operational Risk

- As organizations scale adoption, MCP deployments can quickly become fragmented. Teams independently deploy MCP servers, expose APIs, and connect sensitive systems without centralized visibility or governance. Over time, this creates “MCP server sprawl” — a rapidly expanding ecosystem of loosely governed endpoints and trust relationships

- The resulting risks resemble earlier eras of API and microservice sprawl, but with added AI-specific complexity. Inconsistent authentication, unclear authorization boundaries, weak auditability, and uncontrolled tool exposure can introduce security gaps and compliance concerns. Without intentional architecture, organizations may lose control of how agents access data and execute actions

The Agentic Value Chain

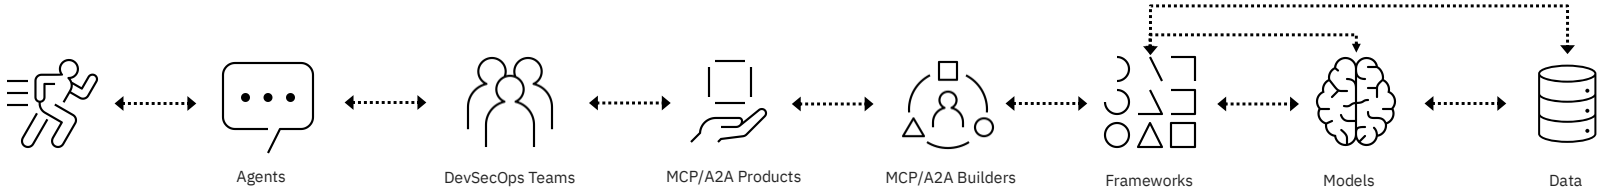

- A value chain is emerging that connects the production of agents to consumption by other systems, applications, and users
- Along this chain there are multiple interconnected processes and constituencies, each with different needs and procedures
- Agent developers are looking for open agentic tools that can help them move quickly, but this must be balanced by managing the multiple risks that these nascent workloads can create

Identity, Authentication, and Trust

- A foundational design principle for MCP environments is establishing strong identity and trust models. Every interaction between an AI agent and an MCP server should be authenticated using enterprise-grade identity controls. This includes service identity management, token validation, credential rotation, and integration with existing IAM systems

- Equally important is defining trust boundaries. Not every agent should access every tool, and not every MCP server should be universally discoverable. Organizations should apply least-privilege principles, isolate sensitive workflows, and ensure that authentication mechanisms are designed for machine-to-machine interactions at scale

Governance & Policy Enforcement

- Production MCP environments require centralized governance capabilities that extend beyond simple connectivity. Enterprises need policy enforcement mechanisms that define which tools agents can access, under what conditions, and with what level of oversight. Governance must address both security and operational consistency.

- This includes lifecycle management for MCP servers, approval workflows for new integrations, inventory tracking, compliance monitoring, and runtime policy controls. Observability is also critical: organizations must be able to audit agent actions, monitor behavior patterns, and investigate failures or misuse in real time.

In the tech industry, history doesn’t repeat itself
but occasionally it rhymes...

- Over two decades ago, SOAP services → REST APIs → API gateways → microservices

MCP plays a similar role for AI agents that APIs played for applications. Before APIs were standardized, every integration was custom and brittle

MCP is doing something similar for AI-native systems:

- Agents become the ”applications”
- MCP servers become “AI-accessible services”
- Tool invocation becomes the equivalent of API calls
- Context exchange becomes standardized vs. custom prompts

MCP/AI Gateways Available in Open Source

There are several emerging options for open source-based MCP/AI gateways:

- IBM ContextForge
- Microsoft MCP GW
- Agentic Community
- Docker Agent GW
- Agent Gateway
- Kuadrant GW

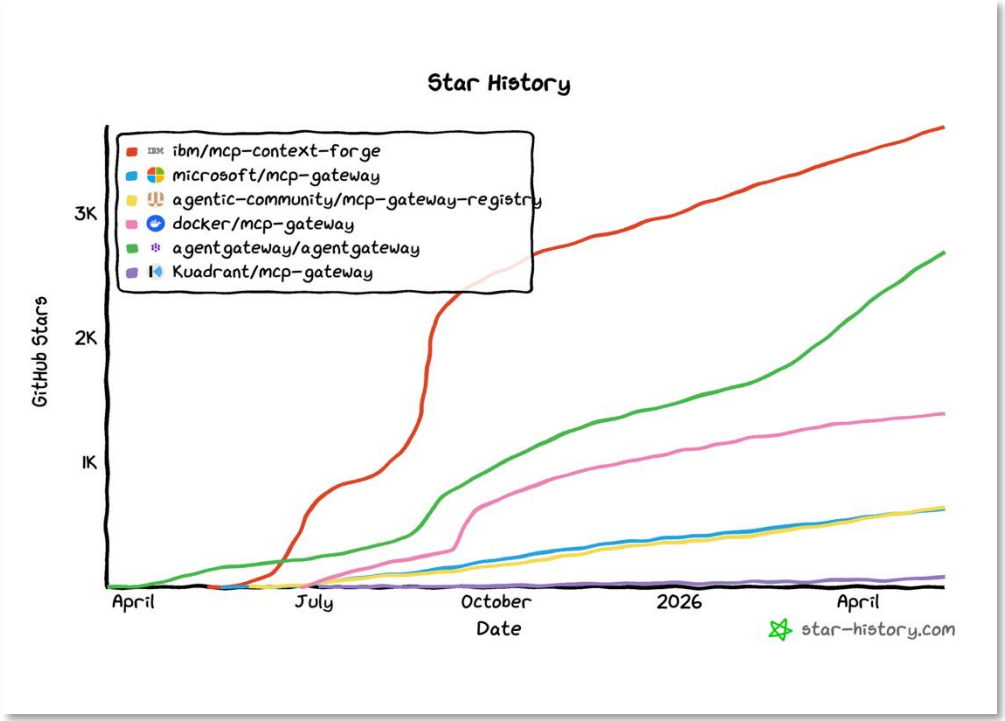

Context Forge MCP with Pluggable Security

MCP Context Forge provides a unified middleware for proxying and mediating access to MCP Servers, enforcing security policies for tools, resources, prompts, and creating virtual servers.

Security extensions implemented as LSM-style hooks to MCP Context Forge

Initial security plugins include:
- Guardrails
- Access control and sandboxing policy enforcement (OPA)
- Security labeling for data flow analysis and control flow integrity

Open source, open standards

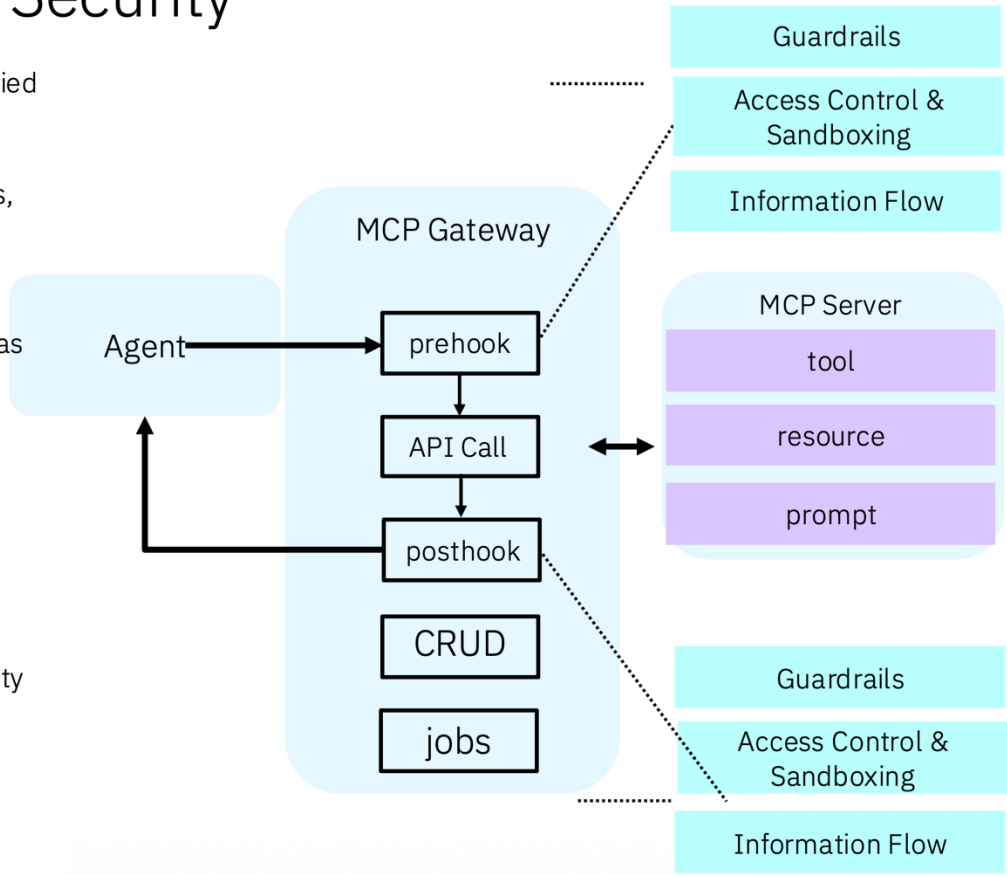

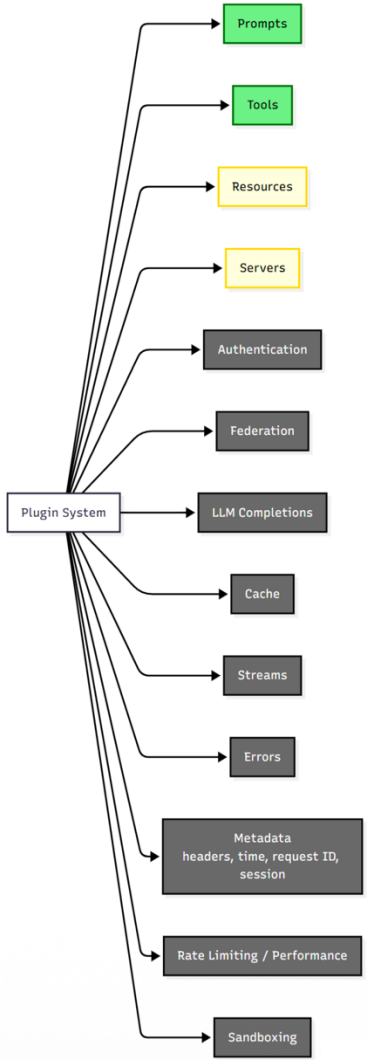

Trusted Plugins Catalog: secure or improve agent outputs

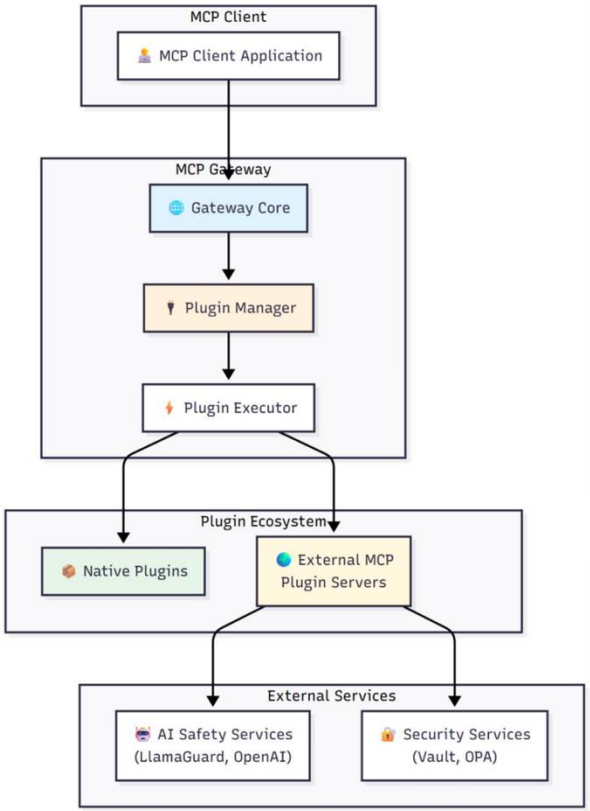

Internal Plugins

35 plugins

rate_limiter	robots_license_guard	virus_total_checker
secrets_detection	sql_sanitizer	code_safety_linter
schema_guard	citation_validator	file_type_allowlist
url_reputation	safe_html_sanitizer	markdown_cleaner
json_repair	html_to_markdown	deny_filter
regex_filter	pii_filter	harmful_content_detector
resource_filter	argument_normalizer	ai_artifacts_normalizer
code_formatter	header_injector	license_header_injector
privacy_notice_injector	output_length_guard	response_cache_by_prompt
cached_tool_result	circuit_breaker	retry_with_backoff
summarizer	timezone_translator	watchdog
resources	external	

External Plugins

4 services

- clamav_server
- opa

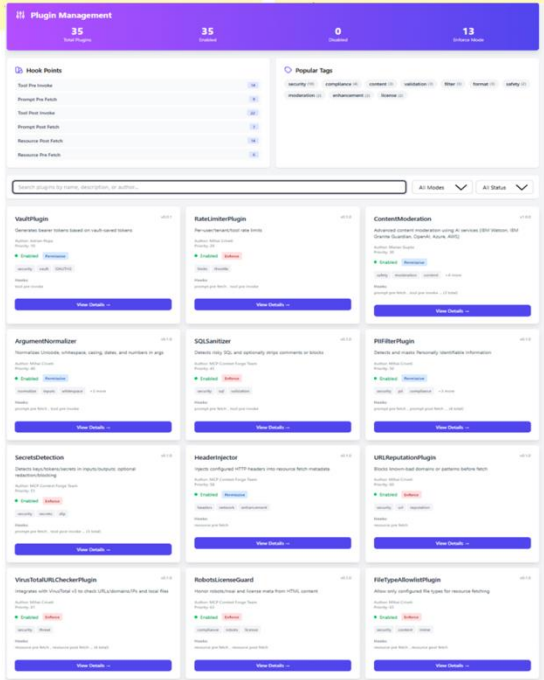

ContextForge: Enterprise AI Gateway & Security for MCP & A2A

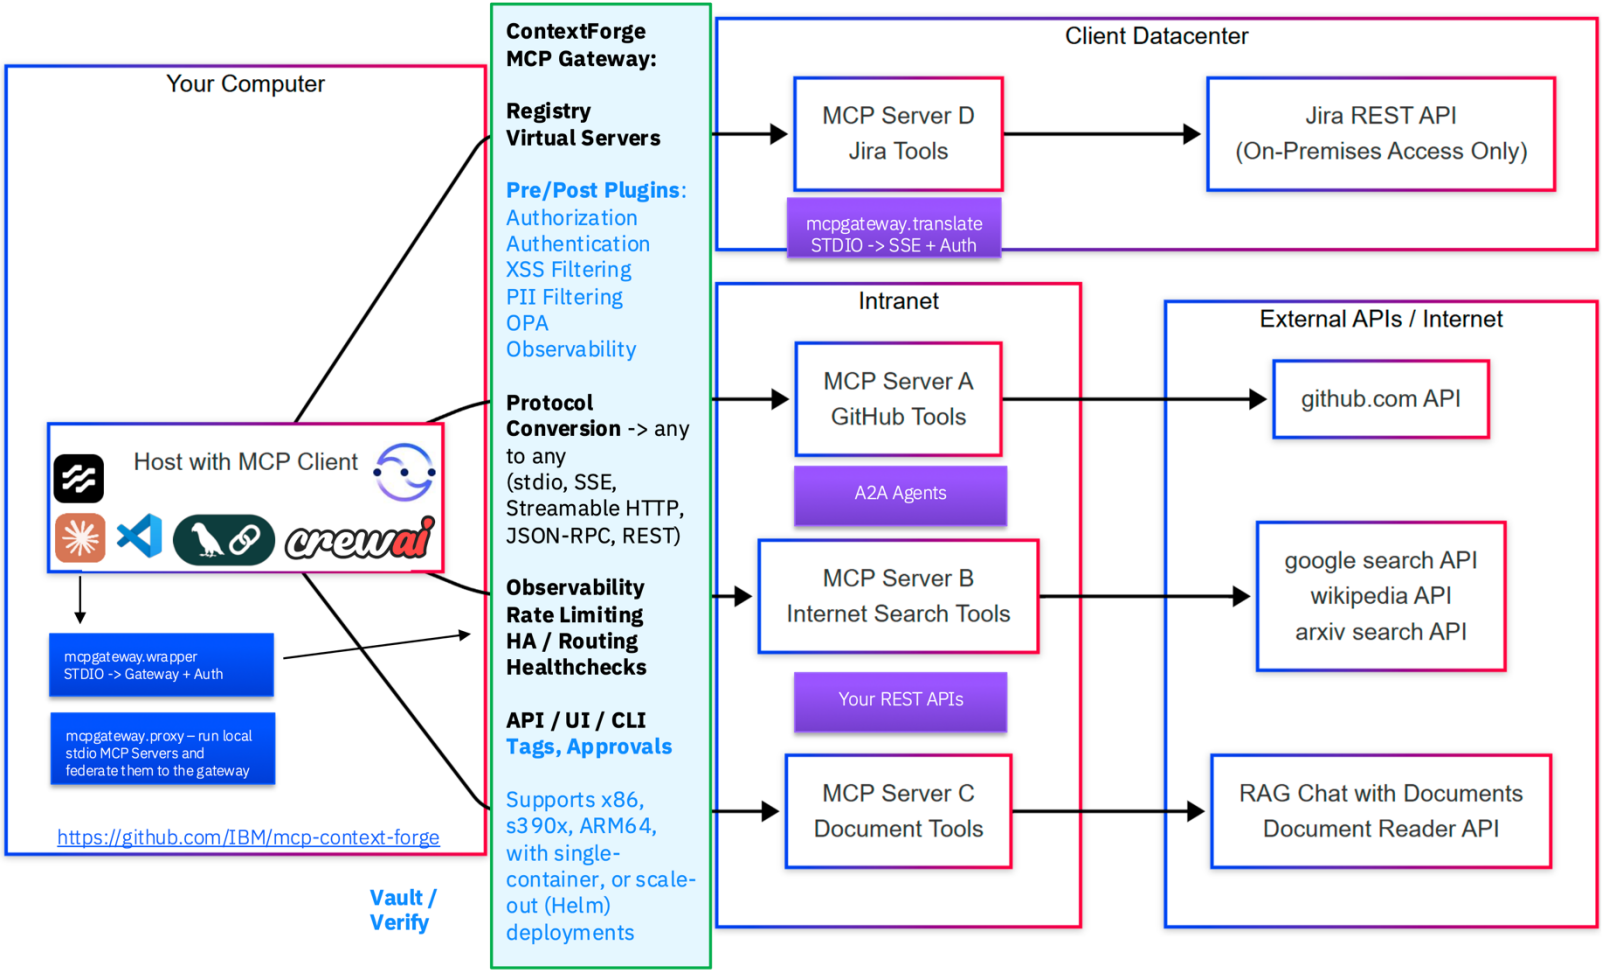

Mercy General Hospital

STAFF & DEPARTMENT ACCESS MAP

STAFF

- **GS** — Gamu Sija — Physician
- **PP** — Pafan Pogaldo — Nurse
- **RT** — Ranpu Tifon — Pharmacist
- **CT** — Cola Tatnath — Rad Tech
- **VN** — Vinfi Nuncuvir — Medical Coder

DEPARTMENTS — CONTEXTFORGE TEAMS

Clinical Docs
5 tools

GS PP RT CT VN

Orders
4 tools

GS PP RT CT

Billing
4 tools

VN

Patient Portal
3 tools

GS PP

Team Access Matrix — Who Sees What

Each department maps to a ContextForge team with scoped MCP tools

	Clinical Docs	Orders	Billing	Portal
GS **Gamu Sija** Physician	Full	Full	—	Full
PP **Pafan Pogaldo** Nurse	Read	Read	—	Full
RT **Ranpu Tifon** Pharmacist	Meds / Allergies	Meds	—	—
CT **Cola Tatnath** Rad Tech	Allergies	Imaging	—	—
VN **Vinfi Nuncuvir** Medical Coder	Codes only	—	Full	—

HEALTHCARE RBAC DEMO — CONTEXTFORGE ARCHITECTURE

Team-scoped MCP federation with virtual servers, RBAC, and plugin-based PII protection

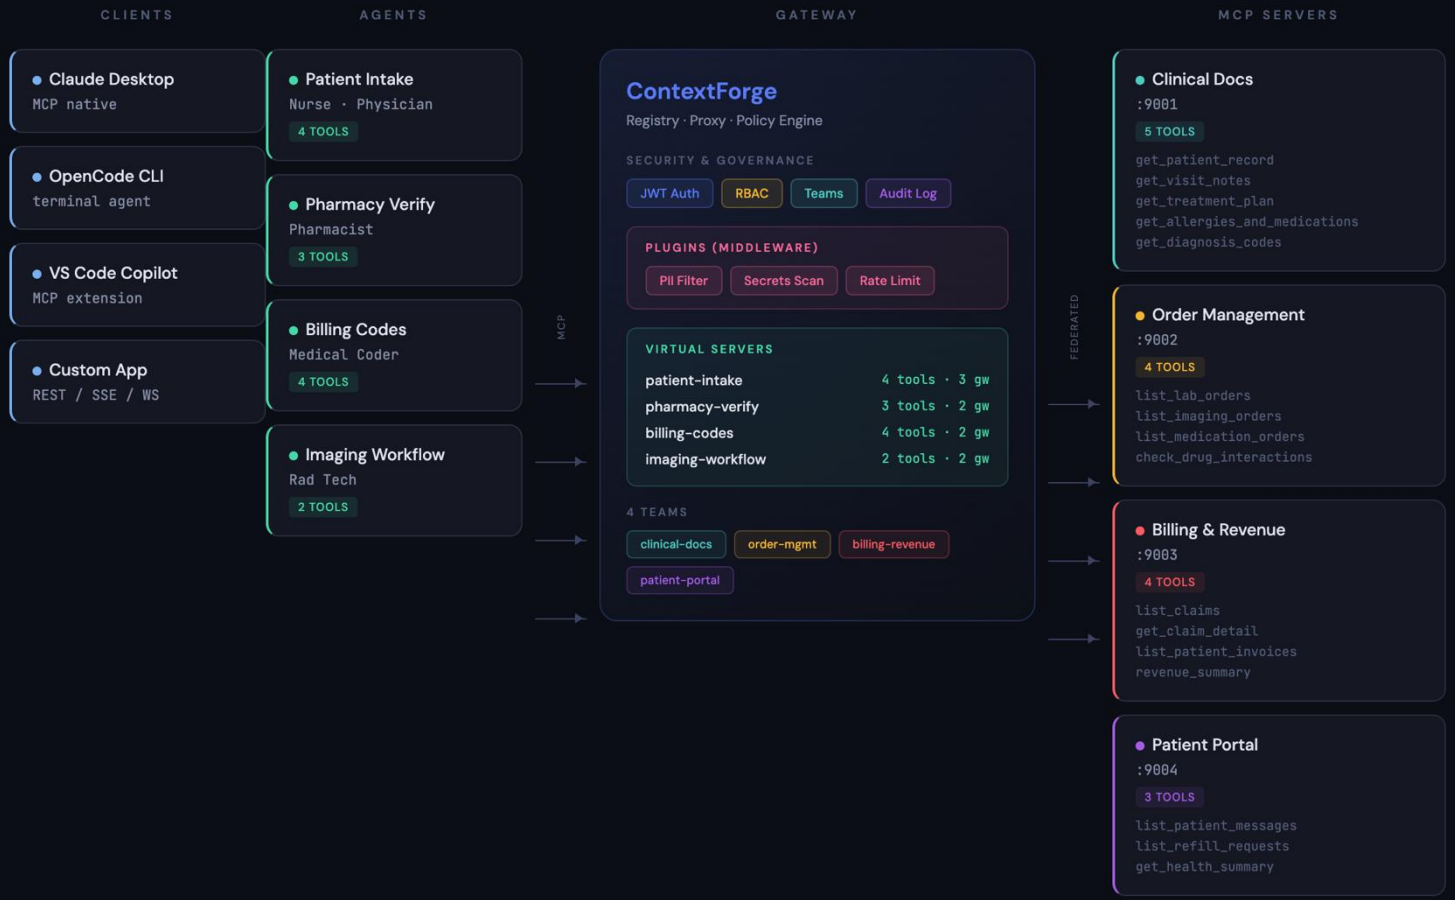

ContextForge Demo

ContextForge

Github
https://github.com/ibm/mcp-context-forge

Docs
https://ibm.github.io/mcp-context-forge/

Questions?