



Developer Track

Building Secure AI Agents with ContextForge MCP Gateway

Shravan Raghu

Summer AI Intensive

August 2026

| The Challenge of Agent Security

Emerging Threat Classes

AI Agents operate autonomously, introducing risks like prompt injection, unauthorized tool execution, and tool-poisoning attacks.

Governance & Compliance

Scaling agents requires centralized auditing, PII redaction, and strict access controls to ensure data privacy across ecosystems.

Standard tool-calling lacks the control plane necessary for secure enterprise adoption.

| Model Context Protocol (MCP)




The "USB-C" for AI Context

The Model Context Protocol establishes an open standard for how AI agents connect to data and tools.

- **Standardized:** JSON-RPC based approach.
- **Universal:** Connects any LLM to any data source.
- **Flexible:** Supports resources, tools, and prompts.

Current Challenges with MCP



Point-to-Point Chaos

Connecting agents directly to MCP servers creates a "Wild West" environment where every interaction is unmediated.

Exposed Secrets: API keys must be managed by the agent itself.

Zero Visibility: No central audit log for tool execution across multiple agents.

Unfiltered Data: LLMs can inadvertently pass PII or malicious payloads directly to servers.


| What is ContextForge Gateway?

A Security-First Middleware

ContextForge is an open-source registry and proxy that federates MCP servers, A2A agents, and REST/gRPC APIs.

It acts as a centralized "control plane" that sits between your AI clients and external tools, providing governance, discovery, and observability.

"ContextForge" MCP Gateway aggregates your MCP Servers:



| Gateway Architecture



Federation

Unifies multiple MCP and REST services into a single clean endpoint for AI clients.



Virtualization

Composes multiple backend servers into a single logical "Virtual Server" for the agent.



Multi-Protocol

Translates between HTTP, JSON-RPC, SSE, stdio, and gRPC seamlessly.

*ContextForge is the cornerstone of our agentic strategy,
securing tool calls for over 160,000 consultants at IBM.*

— Mihai Criveti, IBM Distinguished Engineer

Q&A

 github.com/IBM/mcp-context-forge

<https://ibm.github.io/mcp-context-forge/>

DEMO

Ready to build secure AI agents?

