



MCP
Dev Summit
Toronto

The Agentic Enclave: Building Secure Agent Ecosystems with MCP and Control Planes

Kice Nakeyar




MCP in 60 seconds



- MCP gives agents and hosts a common way to discover and invoke tools, resources, and services.
- For security, the key point is simple: MCP often sits on the path between model reasoning and action.
- If that path becomes standard, it becomes a practical place to attach governance and auditability.

Shared responsibility for agentic systems

Anthropic's four-layer model remains a useful way to explain where controls live.



MCP security today



Identity

who is connecting



Authorization

what is allowed



Gateway visibility

what is being called



Enterprise access

who manages access

The gap

MCP can secure the connection and improve visibility, but end-to-end accountability is still distributed. There is no common evidence model that ties objective, policy decision, tool action, and outcome across multiple MCP servers.



From access to action

Connection security asks...

Can this agent or user authenticate?

Does it have a valid token or scope?

Will the target service accept the request?

same
session


Governed execution asks...

Should this action happen now?

Is the context acceptable?

Should it be denied, modified, or approved?


Direct MCP connections fragment context



Who saw the whole transaction?

The Agentic Enclave

One governed path to many MCP capabilities.



Defense in depth across the four layers

ANTHROPIC'S FOUR LAYERS

Model

Training and model-level safeguards

Harness

Instructions, approvals and orchestration

Tools

MCP servers, APIs and service permissions

Environment

Runtime, data, filesystem and network

Complementary safeguard

MCP-CP does not replace model training.

Verified task context

Bind identity and permissions to the run.

Gateway enforcement

Action policy, approvals, credential isolation.

Runtime containment

Restrict bypass routes and access to secrets.

Govern the action. Constrain bypass. Retain the evidence.

Boundaries and non-goals

What it does

Provides one governance point for MCP workflows.

Policy, approvals, scoped credentials, and evidence can be applied once at the agent boundary.

What it does not do

Does not make external tools trustworthy or fully observable.

Third-party MCP servers and SaaS still own their internal execution, telemetry, and authorization decisions.

Why that matters

This is a pattern, not a protocol replacement.

Direct MCP remains valid. The extra layer matters when you need consistent governance across many agents or boundaries.

Goal: governed interoperability, not ownership of every component.

Call to action: standardize the evidence

What MCP standardized

- A common way to connect agents to tools and services.
- A stable path where gateways and policy can begin to attach.

What is still missing

- A vendor-neutral evidence model for delegated actions.
- Portable telemetry that can flow to a SIEM, dashboard, or observability stack.

Interaction Standard



Evidence Standard

Objective • tool action • policy decision • approval •
credential context • outcome



MCP
Dev Summit
Toronto

If you enjoyed my presentation (or didn't 😊), please take a minute to provide some feedback. Thank you!



The Agentic Enclave: Building Secure Agent Ecosystems with
MCP and Control Planes