



MCP
Dev Summit
Toronto

OCTOBER 5, 2026

Beyond the Protocol

The agentic stack enterprises say is still missing

Hillary Curran

Head of Customer Innovation & AI Enablement



GURU



HEAD OF CUSTOMER INNOVATION &
AI ENABLEMENT



Hi, I'm Hillary

10+ years helping enterprise teams put AI into production



Philly based



Mom of two



Former healthcare
researcher



Customer Success
Implementation
Technical Support
Solution Engineering

QUICK PRIMER



Guru is the governed knowledge layer for enterprise AI

Est. 2013

FOUNDED BY THE TEAM
BEHIND BOOMI

3,000+

ORGANIZATIONS

\$70M+

RAISED

A FEW OF THE TEAMS WE WORK WITH



WHY I DID THIS

The model isn't the bottleneck. Everything around it is.



**10 years of
"here's why it isn't
working"**



**It was never
the model**




**But they were all
my customers**

SO WE ASKED 25 PEOPLE WHO WEREN'T



Hilary Carter



Anni Lai



THE PEOPLE BEHIND THE FINDINGS

25 people. Real names, real companies.



Richard Bian
INCLUSIONAI



**Guillaume
Blaqui re**
CARREFOUR



Vincent Caldeira
RED HAT




Emily Chen
KAIYUANSHE



Grace Francisco
M2I



Thomas Garbay
SAVOIR-FAIRE
LINUX



Mazin Gilbert
AAIF



Arnaud Giuliani
KOTZILLA



Diego Gosmar
TESISQUARE



Angie Jones
AAIF



Hao Li
HAOEASY AI



Zach Lloyd
WARP



Christian Maitre
CASPERA.LAB



Fran ois M ro
PROBABL



Ania Musial
BLOOMBERG



Rick Nucci
GURU



J r me Oufella
SAVOIR-FAIRE
LINUX



Christine Rimer
INDEPENDENT



**Jonathan
Rochelle**
LUTELY




Xavier Shay
UP BANK



Micaela Stump
ADYEN



Lin Sun
SOLO.IO



Brad Topol
IBM




Andreas Traczyk
SAVOIR-FAIRE
LINUX



Matt White
LINUX
FOUNDATION
(FORMER)


THE PEOPLE BEHIND THE FINDINGS

Four vantage points, three regions



THE FINDING THAT ORGANIZES THE STUDY

The question has changed



WHAT WE'LL COVER · FOUR PARTS

1 The agent as coordinating layer

● CONTESTED

2 The trust plumbing

● CONSENSUS

3 Accountability

● OPEN FRONTIER

4 What this means for you

● CONSENSUS

PART 1 OF 4

1 The agent as coordinating layer

● CONTESTED

2 The trust plumbing

● CONSENSUS

3 Accountability

● OPEN FRONTIER

4 What this means for you

● CONSENSUS

PART 1 • THE AGENT AS COORDINATING LAYER

Is the agent the new enterprise kernel?

THE STRONG CLAIM

“The agent is becoming the new enterprise kernel.”



Matt White

Former Global CTO of AI, The Linux Foundation

WHAT MOST SAID

*“The true enterprise shift is **architectural**, not technological... in a similar way to what we did with **cloud native**.”*




Xavier Shay

CEO, Up Bank; Chief Digital Officer, Bendigo Bank

“MCP is the USB port of the agentic world... But USB alone didn’t make the internet.”

Mazin Gilbert, Executive Director, Agentic AI Foundation



PART 1 · THE AGENT AS COORDINATING LAYER

MCP is one layer, not the whole stack



**Necessary.
Not sufficient.**

Two layers are still missing. They decide whether agents from different vendors can work together.

PART 1 · THE AGENT AS COORDINATING LAYER

Match the abstraction to the task

Skills

WHERE THEY SUFFICE

Mostly a markdown file.

Flat learning curve, easy to maintain.

MCP tools

WHERE THE TASK NEEDS THEM

JSON-RPC and real code.

State, security, two-way data flow.



Lin Sun

VP of Open Source, solo.io

“Writing a skill is so much easier than writing an MCP tool.”

PART 2 OF 4

1 The agent as coordinating layer

● CONTESTED

2 The trust plumbing

● CONSENSUS

3 Accountability

● OPEN FRONTIER

4 What this means for you

● CONSENSUS

PART 2 · THE TRUST PLUMBING

A gateway between agents and everything they call

AGENTS

Coding agents

Support agents

Analytics agents



Gateway

🔒 Checks policy before acting

🔑 Who it acts for

📊 Spend by team & agent

📖 Audit trail



MODELS


Frontier, for hard reasoning
Small & local, for volume

TOOLS & DATA

MCP servers
APIs & knowledge stores

PART 2 · THE TRUST PLUMBING

Guardrails go beneath the model



Emily Chen
Co-founder, Kaiyuanshe

“If foundation models are the brain of an agent, guardrails are its safety system.”

PART 3 OF 4

1 The agent as coordinating layer

● CONTESTED

2 The trust plumbing

● CONSENSUS

3 **Accountability**

● OPEN FRONTIER

4 What this means for you

● CONSENSUS

PART 3 · ACCOUNTABILITY

Accountability stays with humans

</> PULL REQUEST #4127

Refactor billing reconciliation job

YOU Authored by **you** · drafted with an agent

⚠ Your **byline**. Your review. Your responsibility.

*"Your byline is still attached, so you're **accountable** for it, and you must review it before it's merged."*



Angie Jones

VP of Developer Relations, Agentic AI Foundation

*"You're all managers now. A manager is **accountable** for what their team produces, whether that output came from a person or an agent."*




Rick Nucci

Co-founder & CEO, Guru

PART 3 · ACCOUNTABILITY

Who's accountable when agents hand off?




Christian Maitre
President, Caspera.lab

"...the combination of legitimate decisions produced an illegitimate result."

PART 3 · ACCOUNTABILITY

How tools go missing: “shadow amnesia”



VISIBLE, OWNED, REVIEWABLE

INVISIBLE, UNGOVERNED, STILL EXECUTING

SHADOW AI

Tools nobody approved.

SHADOW AMNESIA

Tools everybody forgot.

PART 3 · ACCOUNTABILITY

The fix is discovery and identity



A sanctioned place to build

A shared platform, so experimentation doesn't route around you.



A registry with named owners

Every tool has an entry and a human accountable for it.



Credentials that expire

Lapse unless renewed, so forgotten tools fail safely.

PART 4 OF 4

1 The agent as coordinating layer

● CONTESTED

2 The trust plumbing

● CONSENSUS

3 Accountability


● OPEN FRONTIER

4 What this means for you

● CONSENSUS

PART 4 · WHAT THIS MEANS FOR YOU

The answer depends on the seat



"The slowest move is not technology, it's human change management."



Guillaume Blaquière

Group Data and Agentic Architect, Carrefour

PART 4 · WHAT THIS MEANS FOR YOU

If you're building on Wednesday



1

Match the abstraction to the task

Skills where they're enough. MCP where you need it.



2

Build for the gateway

Make identity, attribution and metering easy to plug in.



3

Guardrails in the protocol path

Enforce at runtime, not in the prompt.



4

Make discovery and identity first-class

Registries, named owners, credentials that expire.



5

Instrument agent handoffs now

Before anyone agrees on the answer.



6

Do it in the open

Neutral venues are where these layers get defined.

PART 4 · WHAT THIS MEANS FOR YOU

And if you run AI inside an enterprise



1

Start with the bottleneck that's actually yours

Trust, cost, data, change, or review capacity. Diagnose before you invest.



2

Put the control layer in early

A gateway between agents and everything they call, from day one.



3

Name a human owner before an agent ships

Bounded authority, a named person, and instrumented handoffs.



4

Measure what the agent completes

Completion and resolution rates against the human baseline, not tokens.



5

Treat knowledge and context as infrastructure

An agent is only as good as what it can be told about the business.



6


Budget for integration, not just change management

Embedding agents in existing systems is what gets a pilot to production.

Thank you.

Check out the full study 🙌

Hillary Curran
[linkedin.com/in/hellohillary](https://www.linkedin.com/in/hellohillary)



Free and openly licensed.



[SCAN TO READ IT](#)