Agents @ Scale

The Platform mandate

The Core Dilemma

The Agentic Sprawl

* Illustrative samples

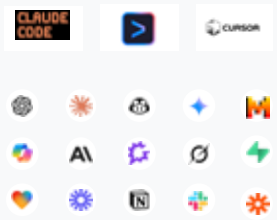

Frameworks and platforms*

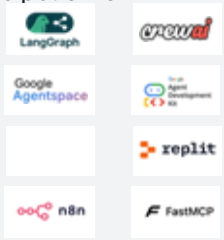

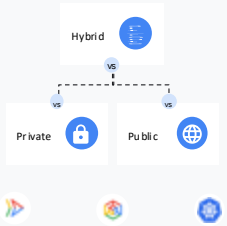

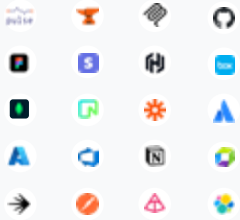

Business Outcomes

Quality

Integrity

Trust

Innovation

Google Open Source

Pillar 1 & 2: Build & Scale Seamlessly

Build

- Agent Development Kit
- 3P Agent Framework
- Agent Studio
- Agent Garden

Gemini API and models

- Gemini Models
- 3P and Open Models
- Model Inference
- Managed Training

Tools, data, and other agents

- A2A
- Grounding
- RAG
- MCP
- Search
- APIs and connectors
- A2UI
- AP2 and UCP
- Cloud Marketplace

Scale

- Agent Runtime
- Agent Sessions
- Agent Sandbox
- Agent Memory Bank

The Core: Remote, Managed MCP

The unified, fully managed MCP platform for Google Cloud

- Unified interface and simplified discoverability with Registry
- Fully managed: No infrastructure to provision. Google handles hosting, scaling, and security.
- Easy access your favorite client: Gemini CLI, Claude, ADK Agents, and more...

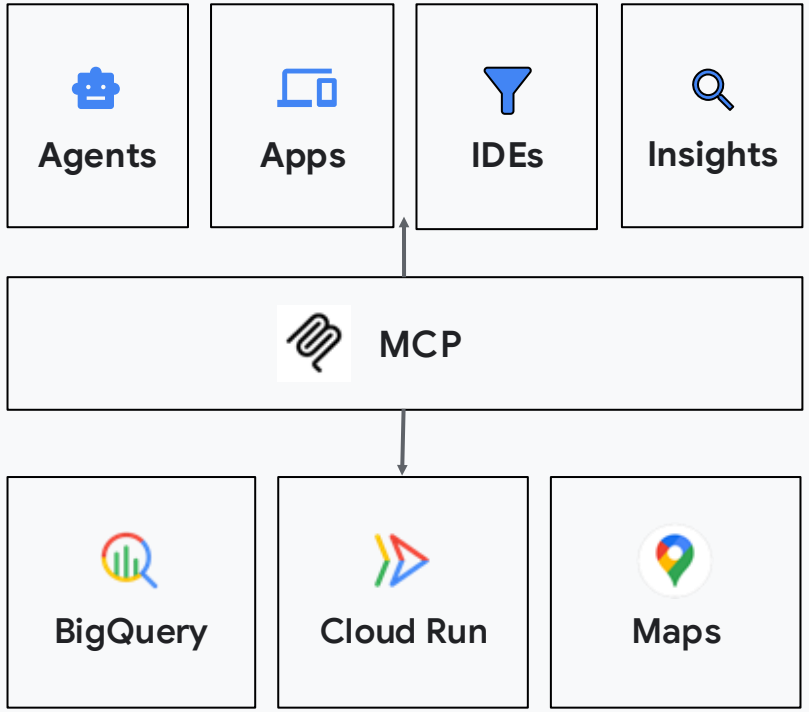

Safeguards at every agent interaction

01 **Authentication and authorization**

Who is the user & agent? What is it allowed to do?
Which actions need human authorization?

02 **Auditing**

You can’t secure what you can’t see.

03 **Agent Data & Context security**

Sanitizing what agents see and say.

Pillar 3 : Centralized Governance & Identity

Agent Identity

Automatically assigned, cryptographically-attested SPIFFE identities. Standards based least privilege approach

Agent Gateway

Single point to define, apply and enforce granular access policies across all agent tools. Built-in Model Armor.

Agent Registry

Centralized catalog for all agents, MCP Servers, Endpoints and skills. Promotes reuse and best practices.

Trusted Compliance

Ensure non-repudiable auditing, total compliance visibility and efficiency through reuse. Accelerates trusted development.

Pillar 4: Continuous Optimization

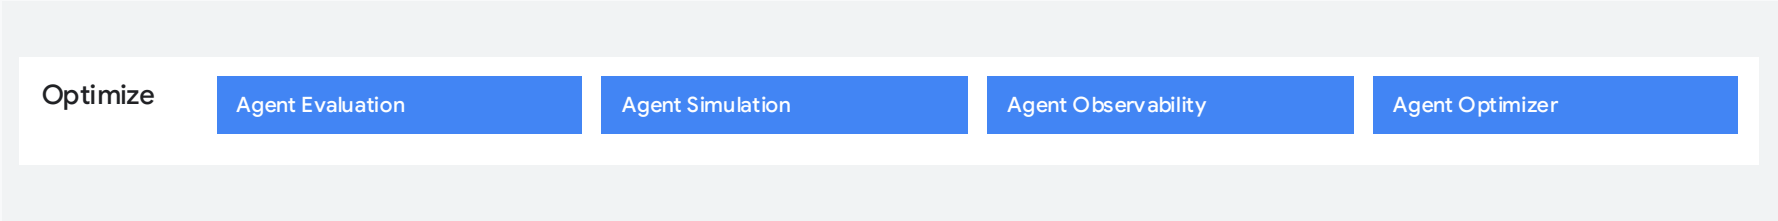

Multi-turn autoraters

Task success scores, logical conversation trajectory, and quality of tool execution

Powerful simulations

Generate realistic, multi-turn user interactions and safely test sensitive actions

Online evaluations

Monitor and observe the performance of multi-agent systems in production

Automated optimization

Automatically fix and refine agents based on real-world failures and evaluation results

Multi-agent traceability

Turnkey dashboards for for multi-agent system traceability

Platforms matter even more

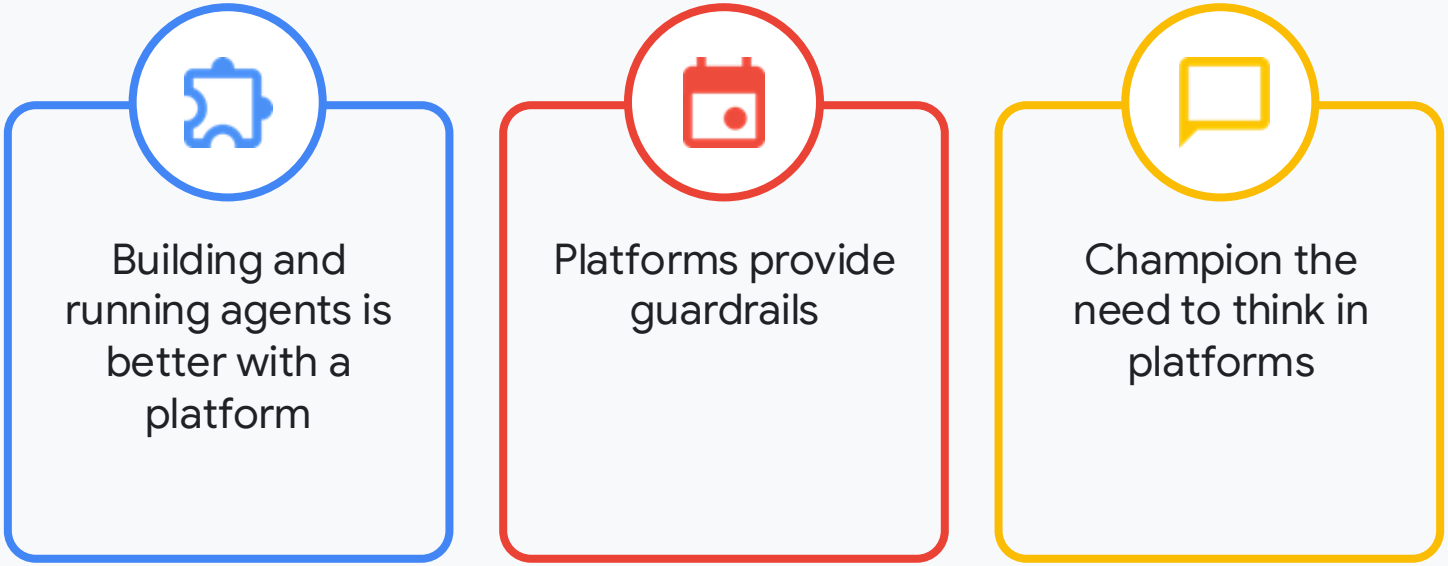

Thank you.

Romin Irani
romin@google.com