

서울

What Agentic AI asks of open source strategy

A new layer for OSPOs



Ana Jiménez Santamaría, Linux
Foundation (TODO Group, PyTorch)





👋 Hi, I'm **Ana**

Sr. Project Manager, [The Linux Foundation](#)

Agentic AI Foundation Ambassador

Open source strategy, OSPO and community health

🧑 Where I contribute



Four things to cover

- 01 Questions coming from the industry
- 02 Agent = model + harness and its relationship to governance
- 03 The open pieces at three layers
- 04 How openness moves layer by layer (and what to do about it)



New Ingredients. Similar Recipe?

- NIST AI Risk Management Framework
- ISO/IEC 42001
- EU AI Act
 - Enforcement opened on August



New Ingredients. Similar Recipe?

- NIST AI Risk Management Framework
- ISO/IEC 42001
- EU AI Act
 - Enforcement opened on August



New Ingredients. Similar Recipe?

- NIST AI Risk Management Framework
- ISO/IEC 42001
- EU AI Act
 - Enforcement opened on August



The United Nations itself has framed OSPOs as an organizational mechanism for turning open source in AI infrastructure from an aspiration into operational capability

What organizations are asking

- Which protocols should we standardise on?
- What policy governs AI-assisted contributions to our repos?
- How do we know what an agent actually did?
- Can we audit it later?
- Where does human oversight sit in the agent loop?



What organizations are asking

Same four dimensions of OSPO value

- **Resilience** to preventing lock-in
- **Cost avoidance** for not paying for a migration you could have designed out
- **Risk foresight** for catching a misleading agent output before it reaches production
- **Strategic influence** for being in the room while the standards are still forming



서울

Which model?



서울

Which harness?




서울

Agent = model + harness



What means for Engineers



The engine generates power. **The harness turns it into a useful, safe vehicle.**

Image Source: [Agent Harness Where Engineers Add Value When Models Keep Getting Smarter](#), by Trista Pan

🧠 Foundation Model (LLM)


🔧 Harness

- └ Instructions & context
- └ Tools
- └ Memory & state
- └ Execution loop
- └ Permissions
- └ Policy Enforcement
- └ Observability

What means for Engineers

The Model–Harness Training Loop

How harness primitives get absorbed into model training



🧠 Foundation Model (LLM)

🔧 Harness

- └ Instructions & context
- └ Tools
- └ Memory & state
- └ Execution loop
- └ Permissions
- └ Policy Enforcement
- └ Observability

What means for Organizations

- The model is probabilistic and cannot be audited line by line
- The harness is where “deterministic guardrails” can go
- Provenance, policy and ecosystem engagement (the OSPO's existing skills) attach at the harness
 - E.g **AAIF Project goose** is an open harness: runtime and orchestration across many models and providers

"which model?" vs "which harness?"

The harness is where teams could differentiate



서울


keep the model a choice

Open Infra Stack for Agentic AI





**open harness +
open protocols**



**Agentic AI
Foundation**

The open pieces. Relation to OSPO

OPEN PIECE	WHAT IT DOES	OSPO EXAMPLE
MCP	Connects the harness to tools, resources and reusable prompts over an open client-server protocol	Standardise tool integrations on an open protocol; reduce vendor lock-in
MCP servers	Host and execute capabilities: repo metrics, database queries, internal services	With security and platform, define the organisation-approved capability set
AGENTS.md	Repository-specific instructions an agent reads	Standardise contribution policy and AI-assisted development guidance across projects
agentgateway	Shared routing, authentication, observability and policy across agents, models and MCP servers	Translate open source and AI governance policy into audit requirements
A2A	Lets agents communicate and coordinate with other agents	Preserve interoperability across multi-agent systems

— Ana Jiménez Santamaría · What Agentic AI Asks of Open Source Strategy · OSS Korea 2026



Openness moves **layer by layer**

"Is this open source or not?" gets more granular: "Is the harness open?" "is the protocol open?" "is the model swappable, or are we cornered into one vendor?"

An open harness plus an open protocol keeps the model a choice

Familiar moves, one layer up

- **Standards advocacy:** Open harness, open tool protocol, open repository guidance, interoperable infrastructure
- **Ecosystem engagement:** The project health questions don't change
 - who governs it, is development concentrated in one company, is the roadmap open, can we influence direction, will the standard stay interoperable?



Final Remarks

- 01 Questions arriving are about resilience, cost avoidance, risk foresight and strategic influence
- 02 Know and improve your harness
- 03 Know the open pieces and which layer each sits at
- 04 Keep openness granular, so the model stays a choice rather than a vendor lock-in



Group shared learnings:
<https://github.com/todogroup/working-groups/blob/main/wg-agentic-ai/meeting-notes.md>

**Building this with peers: the TODO Group
Agentic AI to Empower OSPOs Working
Group meets biweekly!**



OPEN
KOREA

SOURCE SUMMIT

THE LINUX FOUNDATION

서울

