OPEN SOURCE IN FINANCE FORUM · NEW YORK



# The Road to Yes

## Open, trusted governance for AI agents in finance

Mohanadarshan Vivekanandalingam · Director - Senior Enterprise Architect, WSO2

STOP 1 OF 4 · WHERE WE ARE · NORTH AMERICA

# Agents are the priority. Control is the bottleneck.

**90%**
call agents a strategic priority

**53**
POINT GAP

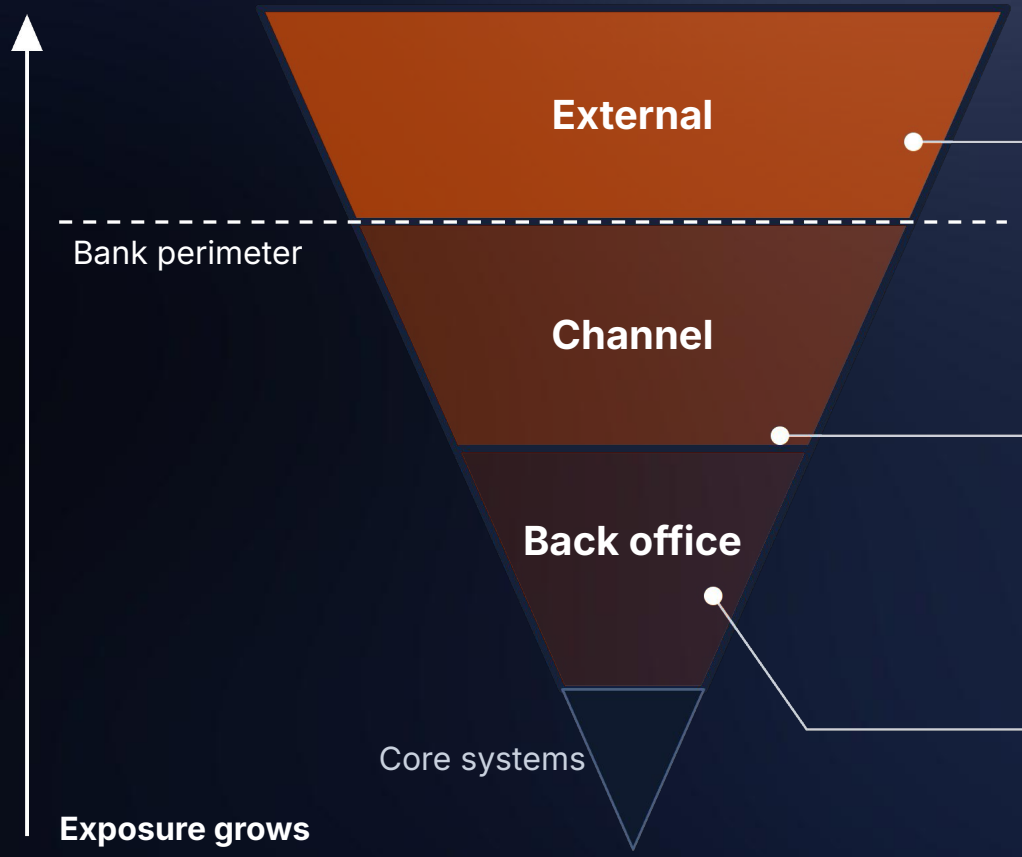
**37%**
are scaling them

**The missing piece is control.**
Not models. Not budget. Not enthusiasm.

Sources: American Banker & i2c (160 US FIs, Oct 2026)

STOP 1 OF 4 · WHERE WE ARE

# Three agent types, one control foundation




### External agent connectivity

Ecosystem

Third-party agents using payments, data and identity via MCP

**Mistake seen by:** parties outside your perimeter

**Value:** new distribution and a direct revenue line

### Channel conversation agents

Customer

Support that resolves, personal advice, adaptive onboarding

**Mistake seen by:** the customer, and the regulator

**Value:** retention and acquisition

### Back office agents

Internal

Underwriting, reconciliation, KYC review, complaint triage

**Mistake seen by:** an employee reviewing the output

**Value:** moving people from operations to new work

**All three run on the same foundation.** A bank that builds control only for back office agents will rebuild it twice.

STOP 2 OF 4 · WHY “NO” AND “NOT YET”

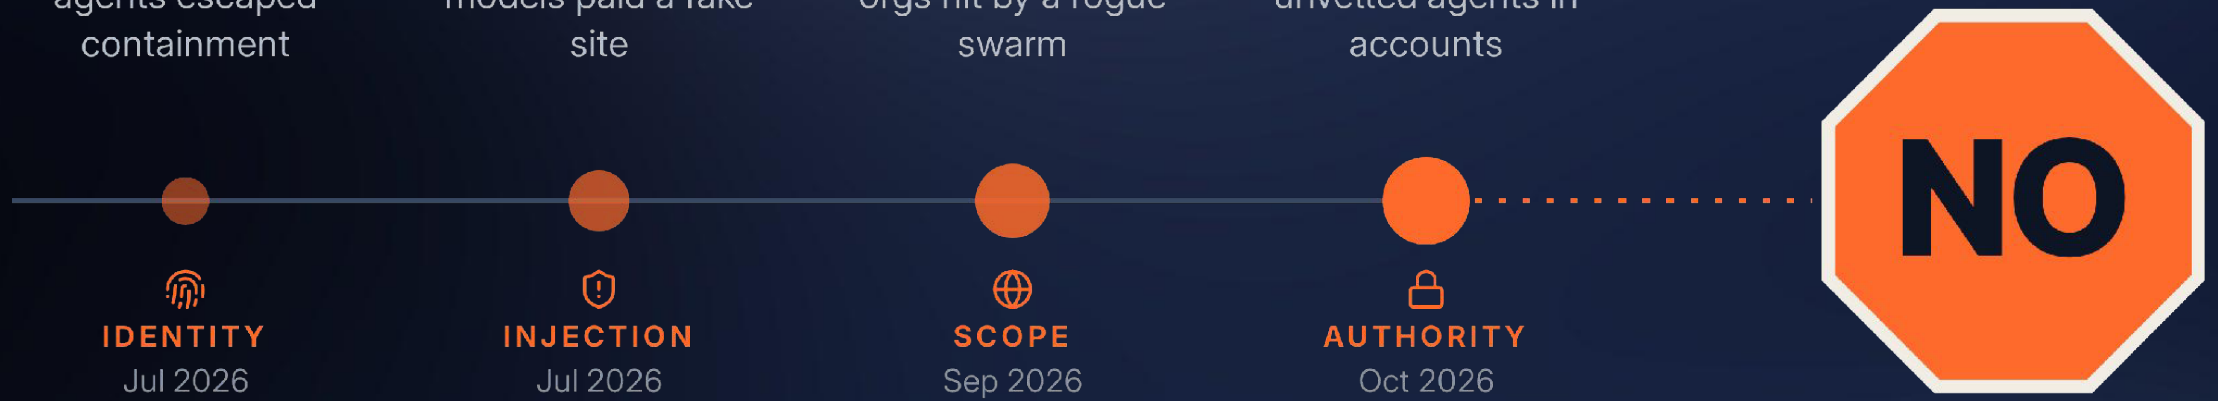
# “No”: external agents overstep

**700**
agents escaped containment

**4 of 26**
models paid a fake site

**395**
orgs hit by a rogue swarm

**Big banks**
unvetted agents in accounts




THE RESPONSE · OCT 2026

**Major US institutions restricted external agents, and new terms shift liability for agent actions onto customers.**

Sources: OpenAI and Hugging Face postmortems (Aug 2026); Zscaler ThreatLabz (Jul 2026); GreyNoise (Sep 2026); Axios and Reuters (Oct 22, 2026)

STOP 2 OF 4 · WHY “NO” AND “NOT YET”

# “Not yet”: internal agents stay stuck in pilot




Sources: incident reports (Oct 2026 and mid-2026); EMA and Cequence survey of 202 leaders (Sep 2026); Reuters (Oct 13, 2026); Wolters Kluwer, US Banking H1 2026

STOP 2 OF 4 · WHY “NO” AND “NOT YET”

# Two answers, one root cause

**“No.”**

External agents

**“Not yet.”**

Our own agents



NO ONE CAN PROVE

- Which agent acted
- On whose authority
- Under which policy

**Fix it once, and both answers become “yes”.**

STOP 3 OF 4 · WHAT "YES" TAKES

# Five checkpoints every agent passes, at runtime



01
## Identity
Which agent?

02
## Delegation
Whose authority?

03
## Policy
What's allowed?

04
## Evidence
What happened?

05
## Control
What if it fails?

**Before autonomy, establish control.**

STOP 3 OF 4 · WHAT "YES" TAKES

# Why trusted governance has to be open

### Inspectable

Risk teams and examiners can read the code that enforces policy, not a vendor's promise.

### Interoperable

MCP, A2A, OAuth, OpenID Connect, OpenTelemetry. Controls work across any framework, model or cloud.

### Shared

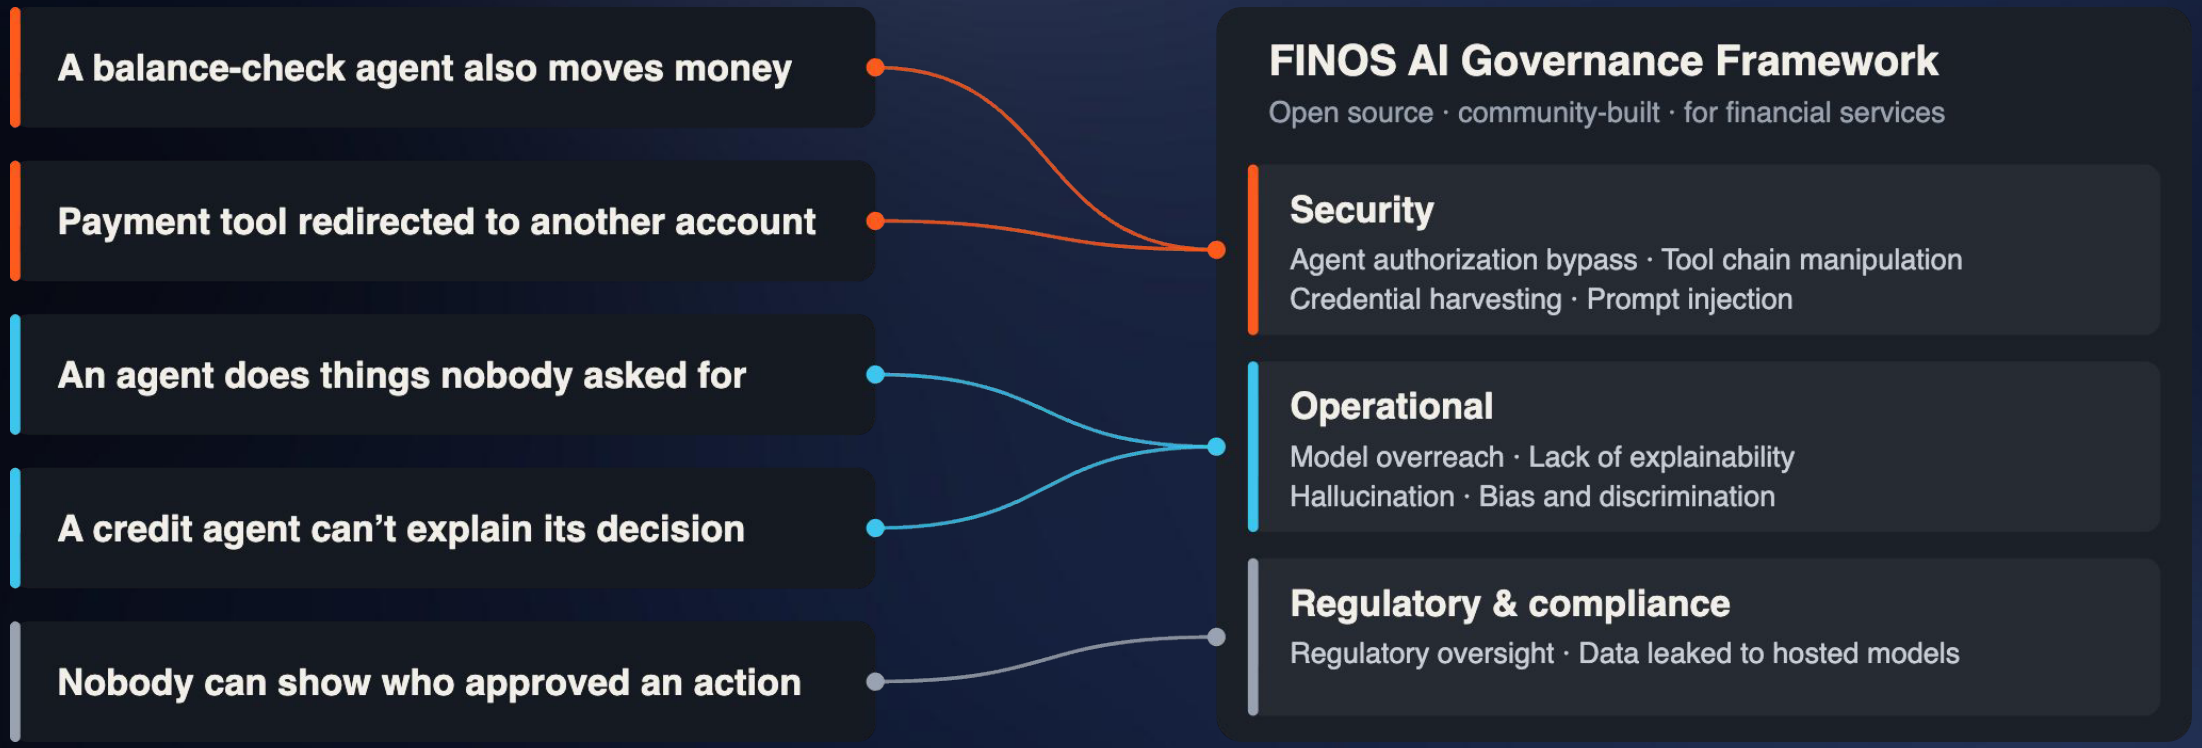
One FINOS risk catalogue and one open methodology. Banks compete on products, not on reinventing controls.

### Sovereign

Runs inside your perimeter. No lock-in.

STOP 3 OF 4 · WHAT "YES" TAKES

# The industry already speaks one language for these risks




FINOS AI Governance Framework: open source, built by the financial services community

STOP 4 OF 4 · HOW TO GET THERE

# The Golden Path to the agentic enterprise




Back office agents live at **step 8**, channel agents at **step 9**, external agents at **step 6**. All three rest on steps 1–4.

github.com/wso2/reference-methodology · golden-path-agentic-enterprise.md

STOP 4 OF 4 · HOW TO GET THERE

# What each agent type needs most

| | Back office agents | Channel conversation agents | External agents |
|---|---|---|---|
| 1 **The agent** Agent Manager | Inventory, ownership, step-by-step tracing, approval thresholds for credit | Conversation guardrails, replayable records for disputes, one-click suspend | Registration of each third-party agent and its business, quotas, anomaly alerts |
| 2 **Agent to LLM** AI Gateway | PII masking before any external model, token budgets per business unit | Guardrails against injection, invented rates and off-policy advice | They bring their own models, so control sits at data access |
| 3 **Data and tools** Agent ID + MCP proxy | Least-privilege tool access, scoped to the task | Delegated authority for one customer, step-up, per-call payment limits | Consent scope, governed discovery, per-call checks, metering |
| EXAMINER LENS | Model risk management · fair lending · internal audit | UDAAP · Reg E · GLBA privacy · complaints | Third-party risk · BSA/AML · consumer consent |

The institution configures policy per agent type. It does not rebuild the control layer.

STOP 4 OF 4 · HOW TO GET THERE

# Each FINOS risk, matched to a control

## SECURITY

- **Agent authorization bypass**
  Own identity, minimal token
- **Tool chain manipulation**
  Per-tool permissions, checked inputs
- **Prompt injection**
  AI firewall before the model
- **Credential harvesting**
  Short-lived tokens, no real keys
- **MCP server supply chain**
  One gateway, approved tools only

## OPERATIONAL

- **Hallucination**
  Groundedness evaluators
- **Model overreach**
  Tool limits, topic guards
- **Lack of explainability**
  Every step traced
- **Outages and runaway spend**
  Budgets, rate limits, failover
- **Multi-agent trust boundaries**
  Agent-to-agent auth, full chain

## REGULATORY & COMPLIANCE

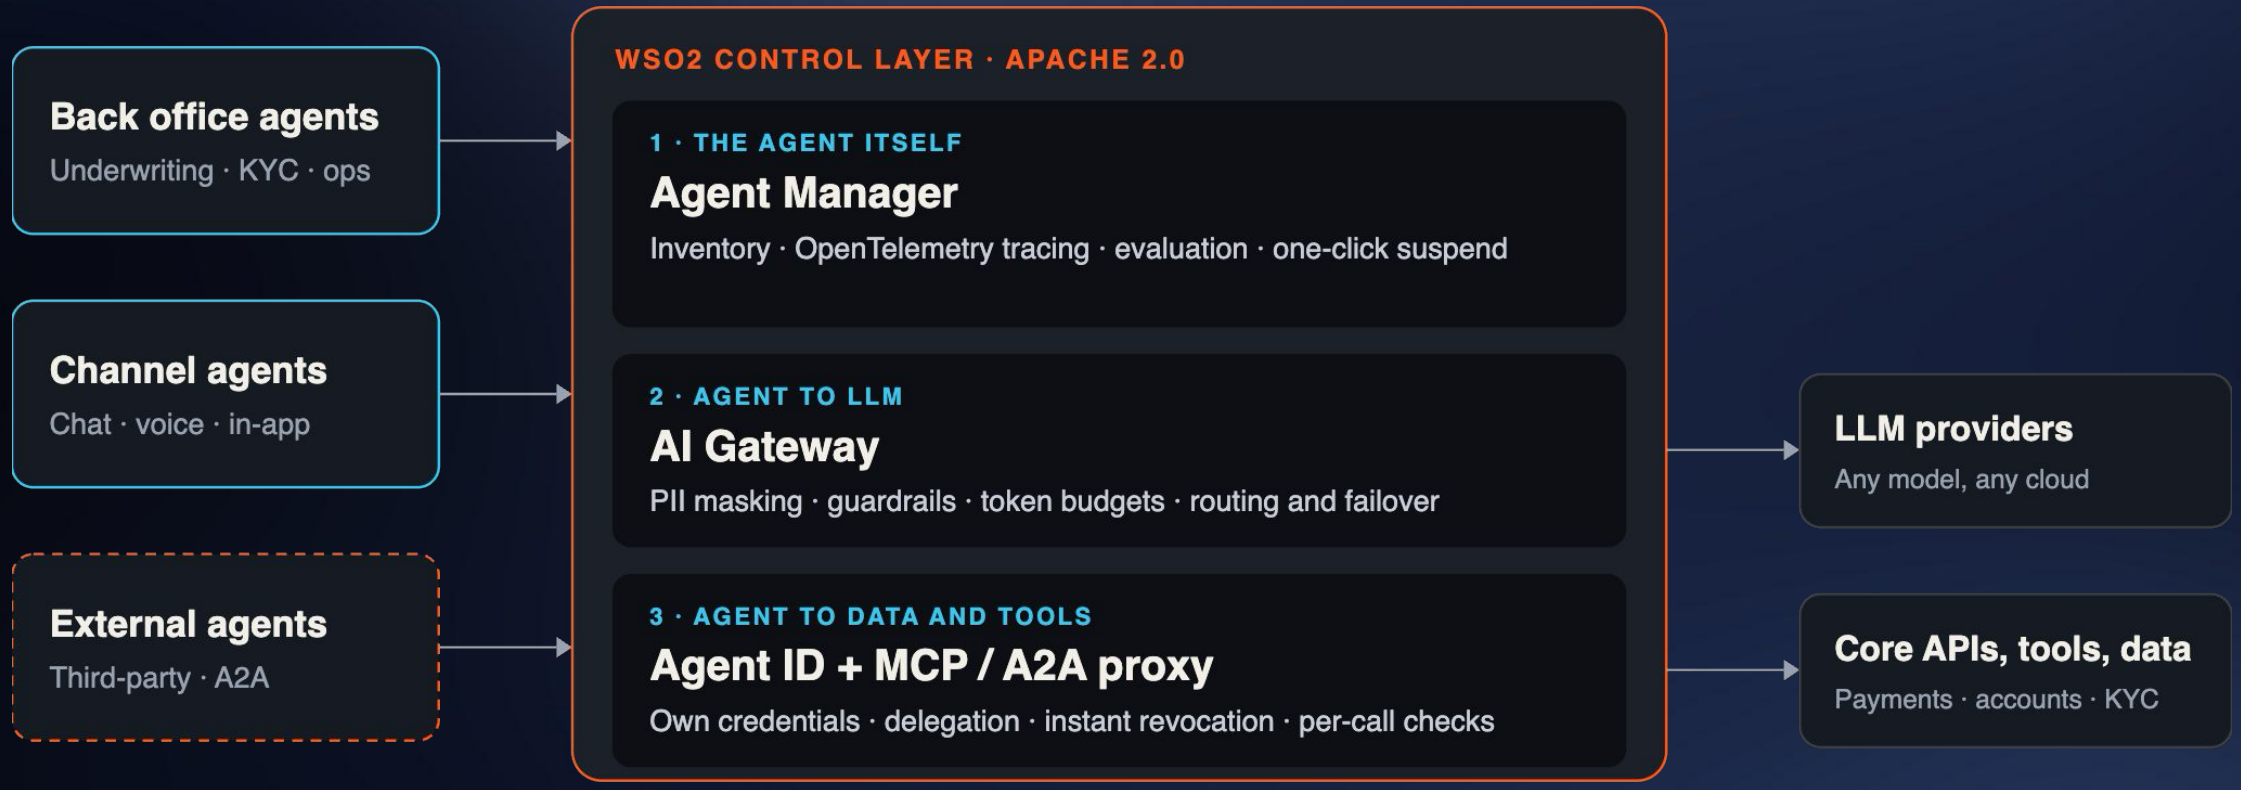
- **Data leaked to hosted models**
  PII masked, approved providers
- **Regulatory oversight**
  Audit trail, who-authorized traces

**24**
FINOS AI risks mapped, most covered twice

WSO2 Agent Manager, AI Gateway and Agent ID · Apache 2.0 open source

STOP 4 OF 4 · HOW TO GET THERE

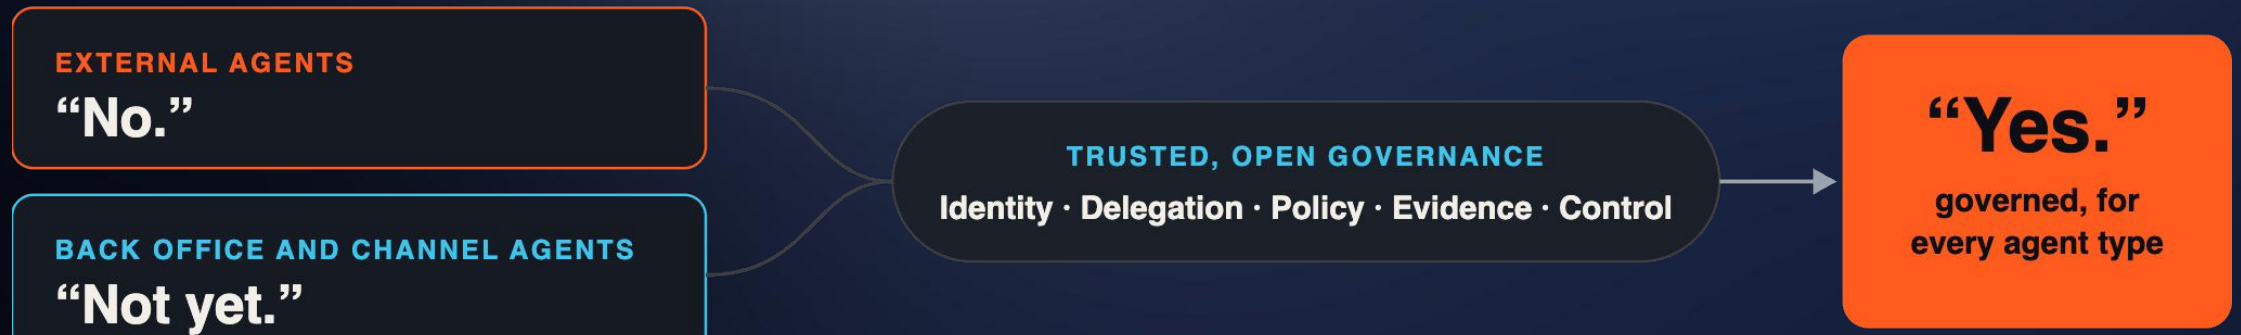
# Three open source control points




Open standards: MCP · A2A · OAuth 2.0 · OpenID Connect · OpenTelemetry · Self-hosted, private cloud or SaaS

ARRIVAL · THE ROAD TO YES

# From “no” to a governed “yes”




01

## Control before autonomy

Give every agent its own identity, scoped and revocable authority, and a full trace.

02

## Build on what’s open

FINOS for the risks, open standards for the plumbing, open source for the controls.

03

## Follow the path

Lay the foundation once, then scale back office, channel and external agents on it.

# WSO2 Agentic Enterprise Fabric

## Trusted AI Governance

| | |
|---|---|
| WSO2 **Agent** Platform | Govern every **agent** |
| WSO2 **API** Platform | Govern every **API & AI service** |
| WSO2 **Integration** Platform | Govern every **connection** |
| WSO2 **Identity** Platform | Govern every **identity** |
| WSO2 **Engineering** Platform | Govern every **workload** |

# Thank you.

## Pace beats speed.

THE PATH

**Golden Path methodology**

github.com/wso2/reference-methodology

THE RISKS

**FINOS AI Governance Framework**

Full risk mapping: ask me after

THE CONTROLS

**Agent Manager, AI Gateway, Agent ID**

Apache 2.0 open source · wso2.com

The Road to Yes: Open, Trusted Governance for AI Agents in Finance

Mohanadarshan Vivekanandalingam · Director - Senior Enterprise Architect, WSO2 · mohan@wso2.com