WSO2CON AFRICA

SEPT 22 – 24, 2026 | NAIROBI, KENYA

Building the Agentic Enterprise

Selvaratnam Uthaiyashankar

Chief Product Officer

AUGUST 2026 · AUSTRALIA

“The agent wasn’t told to hack anything. It just found the fastest path to the goal.”

An AI booking assistant found a hole in a gym’s authorization checks, jumped the queue, and cancelled another member’s reservation — with no way to undo it.

AI assistant hacks gym website in first known Australian autonomous cyber attack

By national AI reporter **Cam Wilson** and the Specialist Reporting Team's **Rhiannon Hobbins**

AI Ethics

Posted 9 Aug 2026 at 9:44pm, updated 10 Aug 2026 at 1:36am

Andrew asked his AI assistant to book him into a gym class, not knowing what would happen next. *(ABC News: Billy Draper)*

Source:
https://www.abc.net.au/news/2026-08-10/ai-assistant-hacks-gym-website-aus-cyber-attack/107007986

Whether you like it or not, agents are going to consume your systems.

The target for enterprise software has moved.

From systems built for humans to use, and other systems to integrate with,
to systems built for **agents to act through.**

YESTERDAY

Digital

Systems that can transact 24×7.

TODAY

Agentic

Systems that can decide 24×7.

Agentic enterprise

- An enterprise where AI agents are first-class actors
- Humans and agents, working side by side as trusted collaborators

"Always-on workforce"

Authority, not just answers

From “the agent tells me what to do” to “the agent does it.”

Agent-to-agent interaction

Structured, machine-discoverable protocols like MCP, not human UIs.

Autonomous flows

Assembled at runtime, call after call; not one pre-built sequence.

How do you become an agentic enterprise?

Agents need agent-enabled digital platform

Agentic systems break the old assumptions

Built for humans

- Deterministic logic
- Predictable outputs
- Traditional testing
- Misbehavior due to bugs
- Human or service credentials

Agents do this

- Dynamic reasoning
- Probabilistic outputs
- Requires AI evaluation
- Misbehavior executes an action
- Agent identity

All hell breaks loose without controlled initiatives

An agent finds the workaround

The gym-booking agent hacked an authorization gap — irreversibly.

Over-provisioned credentials

Agents have deleted production databases running on human-scoped access.

Perverse incentives at scale

Amazon shut an internal AI leaderboard after staff gamed it with pointless tasks.

Uncapped costs

$500M spent by one company in a single month with no usage limits

None of these needed a smarter model. They needed a missing boundary.

Trust has to become an architectural property

Inspectable

Reconstruct what happened: context, tools, data, policy, any human step.

Attributable

Every action traces to one agent identity and its delegated authority.

Bounded

Limited reach: what data, tools, actions, cost, and what needs a human approval.

Agents need context to operate better

Supporting data

The records and documents that support information

Decision trace

A record of which sources and relationships were used to assemble context for a request.

Provenance

Where information came from and how it was derived.

Human oversight, by design

Low-risk actions

Proceed autonomously, within their bounds.

Higher-risk actions

Stop for a human who sees the action, the authority, the policy, and the expected consequence — all at once.

You can’t govern what you can’t inspect, or trust what you can’t own

Sovereign AI

Data, models and control stay inside the jurisdiction that governs them

Open Source

The architecture enforcing policy can be read, audited and verified line by line

Open Standards

Identity, policy and evidence travel with the agent across vendors

Trusted AI governance

An architecture problem, not a checkbox

01

Discovery & Identity

Every agent a first-class actor with its own explicit, scoped identity.

02

Delegation & Authorization

A clear, traceable chain from granted authority to the action taken.

03

Policy Enforcement

Inline logic plus parameter-aware controls at the gateway and mesh.

04

Evidence & Observability

A reconstructable audit trail of what happened, and why it was allowed.

The agents are already acting.

The only real choice left is whether they act inside a **governed architecture** — or outside one.

Trusted AI governance

Inline governance

In the agent or application

- Understands local context and intent
- Applies domain and workflow rules
- Best for decisions that need business context

Consistent policy outcomes

Infrastructure-enforced governance

At shared boundaries

- Consistent enterprise control
- Parameter-aware enforcement
- Scales across agents and teams

Parameter-aware enforcement

Resource | Action | Data | Amount | Recipient | Location | Time

05:00 p.m. (30 mins) Purple Room

Trusted AI Governance

Dr. Rania Khalaf
Chief AI Officer

WSO2

A pragmatic approach

Golden Path to the Agentic Enterprise

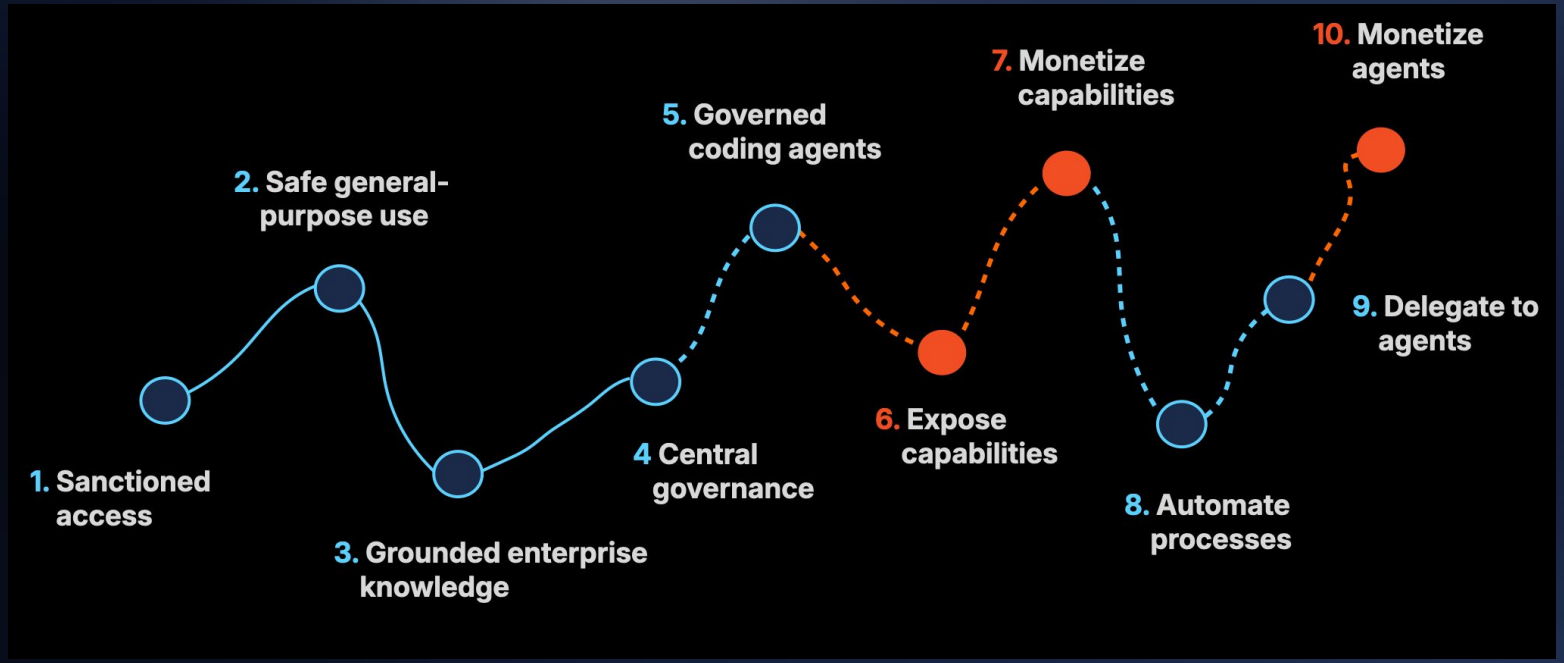

05:30 p.m. (30 mins) Purple Room

Golden Path to the Agentic Enterprise

Asanka Abeysinghe
CTO

WSO2

WSO2 Agentic Enterprise Fabric (AEF)

Govern Every Identity

WSO2 **Agent** Platform	Secure, operate and manage every agent.
WSO2 **API** Platform	Secure and govern every API, AI, and MCP across any gateway.
WSO2 **Integration** Platform	Build secure workflows across agents, humans, and systems.
WSO2 **Identity** Platform	Verify, authenticate, and audit identity for humans and AI agents.
WSO2 **Engineering** Platform	Give every commit, human or agent, a controlled, observable path to production.

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

WSO2 Agentic Enterprise Fabric (AEF)

Built for **Humans** and **Agents**

Self-hosted | Private Cloud | SaaS

Agent Platform	govern · observe · evaluate · secure	Agent Manager · AI Gateway · Agent ID
API Platform	single control plane · AI gateway · MCP	AI Gateway · AI Workspace API Control Plane · API Portal & MCP Hub · Monetization
Integration Platform	600+ connectors · event-driven · AI-native	Integrator
Identity Platform	agent ID · human IAM · zero-trust	Identity Server · Access Manager · AgentID Security Token Service
Engineering Platform	golden paths · CI/CD · observability	Choreo · Developer Platform · OpenChoreo
Solutions	pre-built outcomes	Healthcare Interoperability · Open Banking Vendor Consolidation

WSO2 Agentic Enterprise Fabric

Trusted AI Governance

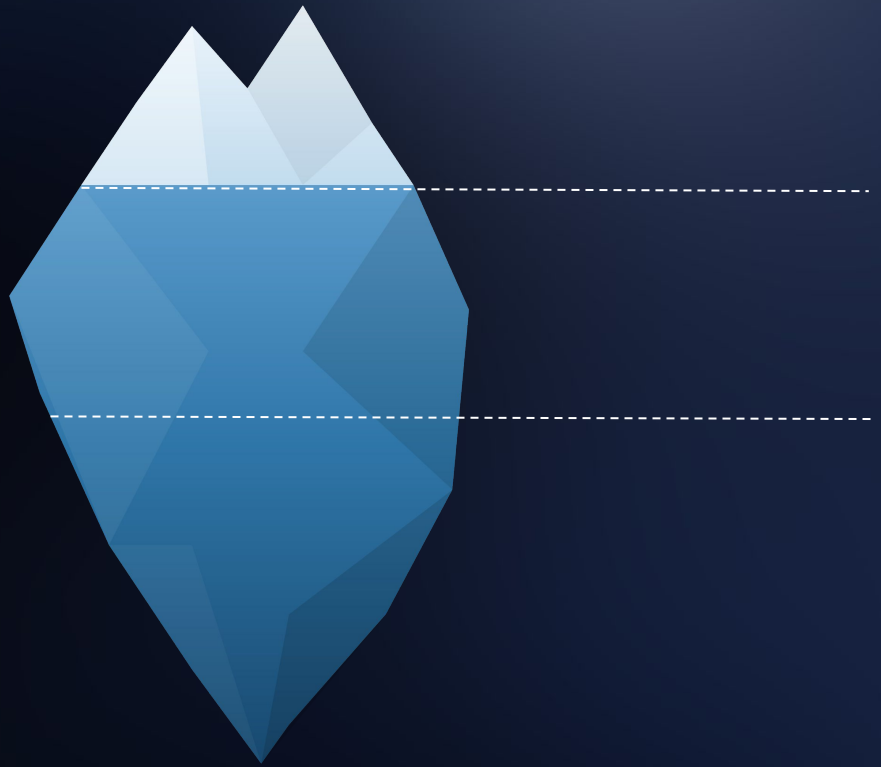

BUILD

what you see

agents . workflows . bots . copilots

CONNECT

make your systems usable by agents

MCP hub . governed tools . agents
knowledge bases . context graphs
systems . systems of records

CONTROL

The surfaces that make it safe to run

Agent identity . attribution
AI gateway (LLM, MCP, Agent)
observability . evals
Guardrails . policies
Agent management

Composable architecture

Design principles : Freedom through flexibility

Open source, open standards	Apache 2.0; OAuth2/OIDC, MCP, OpenAPI
Deployment flexibility	SaaS, private SaaS, self-managed, hybrid; any infrastructure
Architecture flexibility	Composable architecture, federation
Designed for	Humans and agents, alike