



SEPT 22 - 24, 2026 | NAIROBI, KENYA

Golden Path to the Agentic Enterprise



Asanka Abeysinghe

CTO










Picture credit: <https://www.thesfmarathon.com/san-francisco-ranked-best-running-city/>

Direction without
a path is just
aspiration





"I've never been a runner"


Half marathon finisher

Every plan includes strength training · Move to the next only when you're ready

Agentic Enterprise



The **path** matters



Golden Path to the Agentic Enterprise

Specification - CC-BY 4.1

<https://github.com/wso2/reference-methodology/blob/master/golden-path-agentic-enterprise.md>

The Golden Path to Become an Agentic Enterprise

A WSO2 Methodology

Version: Summer-2026

First Published On: Summer-2026

Authors (in alphabetical order)

- Asanka Abeysinghe | Chief Technology Officer, [WSO2, Inc](#)
- Mifan Careem | Chief Solutions Officer, [WSO2, Inc](#)
- Derric Giling | GM API Platform BU, [WSO2, Inc](#)
- Rania Khalaf | Chief AI Officer, [WSO2, Inc](#)
- Ruwan Linton | Chief Customer Officer, [WSO2, Inc](#)
- Selvaratnam Uthaiyashankar | Chief Product Officer, [WSO2, Inc](#)

Contents


- [Overview](#)
- [Roles used in this document](#)
- [How to use this](#)
- [Step 1: Sanctioned access](#)
- [Step 2: Safe general-purpose use](#)
- [Step 3: Grounded enterprise knowledge](#)
- [Step 4: Central governance with distributed enforcement](#)
- [Step 5: Governed coding agents](#)



The **path** is built
one step
at a time

Before **autonomy**,
establish control

The Golden Path



Build the foundation

step 1 - 4



Step 1 **Sanctioned access**

Give every employee a **safe**, attributable way to use AI


- **Approved AI access** behind enterprise identity
- Every interaction attributable to a **person**
- Unsanctioned AI access **blocked** or **monitored**



Step 2 **Safe general-purpose use**

Make everyday AI use **safe, controlled, and economically sustainable**


- Keep enterprise data out of model training
- Apply guardrails before rollout
- Put budgets and limits around usage



Step 3 **Grounded enterprise knowledge**

Ground AI in enterprise knowledge without breaking existing access boundaries.


- Connect AI to curated enterprise knowledge
- Enforce source-level access control at query time
- Start read-only, then allow isolated low-risk actions



Step 4 **Central governance** with **distributed enforcement**

Define **governance centrally**, enforce it where AI is actually used

- Author **policy** once and apply it consistently
- Maintain a live register of models, tools, apps, and agents
- Reuse governed capabilities instead of rebuilding them everywhere



← Evidence and drift detection back to policy




Expand
what **AI** can do
step **5** - **7**

Step 5 **Governed coding agents**

Extend the same governance model to **coding agents** and **AI-assisted development**

- Keep proprietary source and secrets on private endpoints
- Scope what agents can read and what they are allowed to do
- Scan and review AI-generated code before it reaches production




..... X No direct push to main

Step 6 **Expose your capabilities to external agents**

Turn enterprise capabilities into governed tools that external agents can discover and use


- Publish capabilities through MCP, APIs and events
- Give every tool a clear scope, owner, and access policy
- Authenticate and authorize every external agent call



Step 7 **Monetize your capabilities**

Turn governed capabilities into products that can generate revenue

- Meter usage in a way you can bill against
- Enforce entitlements, quotas, and rate plans at the gateway
- Reuse governed capabilities instead of rebuilding them everywhere



Move from **assistant** to
autonomy
step 8 - 9

Step 8 **Automate entire operational processes**

Move from agent-assisted work to agent-first **execution**

- Run agents in a managed runtime with clear identity, ownership, and budgets
- Bind agents only to the tools and actions they actually need
- Put evaluation, tracing, authorization, and human approval into the execution path




.....➔ Human approval tier for high-risk, irreversible, regulated actions

Step 9 **Agents as part of your workforce**

Let agents act on behalf of people, without losing the person behind the authority

- Give agents short-lived, task-scoped delegated authority
- Preserve the full chain of who initiated the action and what was granted
- Give every agent a human sponsor and make access revocable



Turn autonomy into a
business model
step **10**


Step 10 **Monetize your agents**

Turn agents themselves into products and new revenue streams

- Package agents around tasks or outcomes, not just usage
- Isolate tenants, budgets, policies, and evidence
- Preserve identity, delegation, and auditability for every customer interaction




One **path** many **outcomes**





Architecture



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 · Agent Builder



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 on the Golden Path



A three-layer engagement



Strategy



Architecture



Technology



How it all comes Together

Building the **Agentic Enterprise** :the north star

Golden Path to the Agentic Enterprise :the journey

WSO2 **Agentic Enterprise Fabric** :the technology foundation

Trusted AI Governance :the control model

What **actually** matters

- not the **steps**

- not the **tools**

- never even the **agents**

Outcomes & Value

Asanka Abeysinghe

Gautham Pallapa

INVALUABLE

Define Value • Align Execution • Deliver Outcomes



Foreword by **Vaughn Vernon**

Leading domain-driven design practitioner and author

Becoming an
agentic enterprise is
not a **sprint**,
you need a **direction**, a **path**, the right
foundation, and **control** along the way

Pace beats speed.

