WSO2CON NORTH AMERICA

May 20 - 22, 2026 | Austin, Texas, USA

Building Blocks of the Agentic Enterprise

Nadheesh Jihan

Senior Technical Lead

Agents: The New Digital Workforce

"When you're talking about an agent... think of it almost like an **employee or an intern** *that you would hire."*

— Andrew Kapathi

"...a **fresh college graduate** *or an* **entry-level worker***."*

— Dario Amodei (CEO of Anthropic)

Shifting the mental model from "Tools" to "Teammates"

From Deterministic to Probabilistic

Scripted (Traditional)

Deterministic Logic

Developers define every step and condition through code. That code **becomes the behavior**.

Unscripted (Agentic)

Autonomous Execution

Users define the goal. Agents decide the steps at runtime. Developers define the **boundaries and guidelines** — tools, prompts, and guardrails.

"When behavior is decided at runtime, you can't debug your way to reliability. You need a different kind of infrastructure."

Are you ready for this transformation?

Are you prepared to hire your first agent?

The Infrastructure Gap

settinggest

TRADITIONAL STACK

(You Have)

User Auth & IAM
REST APIs & Integrations
DBs & Document Stores
API Gateway & WAF
Container Orchestration
Logs and Metrics

VS

psychology

AGENTIC STACK

(Agents Also Need)

Agent Identity & Credentials
MCP & Skill Registry
Agent Memory Layer
LLM Governance & Guardrails
Durable, Sandboxed Execution
Traces of Reasoning Chains

"Existing infrastructure isn't broken—it simply wasn't designed for non-deterministic behavior."

The 10 Building Blocks

Identity

Safety & Guardrails

Memory & Context

MCP & Skill Registry

Observability

LLM Governance

Runtime

Lifecycle

Agent Interface

Eval & Feedback

"Most enterprises are building agents.
Few are building the foundation."

Agentic Identity

TRADITIONAL STACK

Human-Centric IDP

Shared, long-lived service accounts. No concept of a non-human autonomous principal.

AGENTIC STACK

First-Class Agent IDs

IDP-issued unique agent identities. Short-lived machine tokens with full per-agent audit trails.

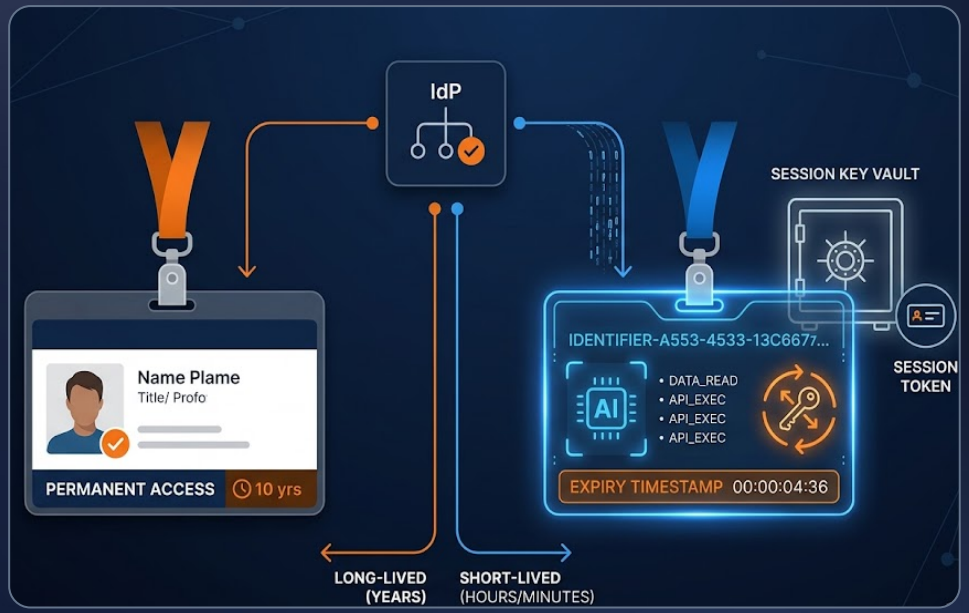

Connecting Agents to Enterprise

TRADITIONAL STACK

Traditional Integrations

Predictable system-to-system calls with documented, static flows. Built for human-facing apps where developers hardcode every logic branch.

AGENTIC STACK

Agentic Orchestration

Agents discover capabilities dynamically through **MCPs**. MCP bridges the agent to enterprise barriers.

--- OR ---

Agent is pointed to existing enterprise interfaces, and given **Skills** to manage the complexities.

MCP: The Agent-Tool Interface

Self-Describing Tools

Model Context Protocol (MCP) wraps existing APIs as tools. The agent receives the **Intent, Inputs,** and **Expected Returns**.

TRADITIONAL API

"Here is /users, figure it out."

MCP SERVER

"I can manage users. Give me an ID, I'll return the profile."

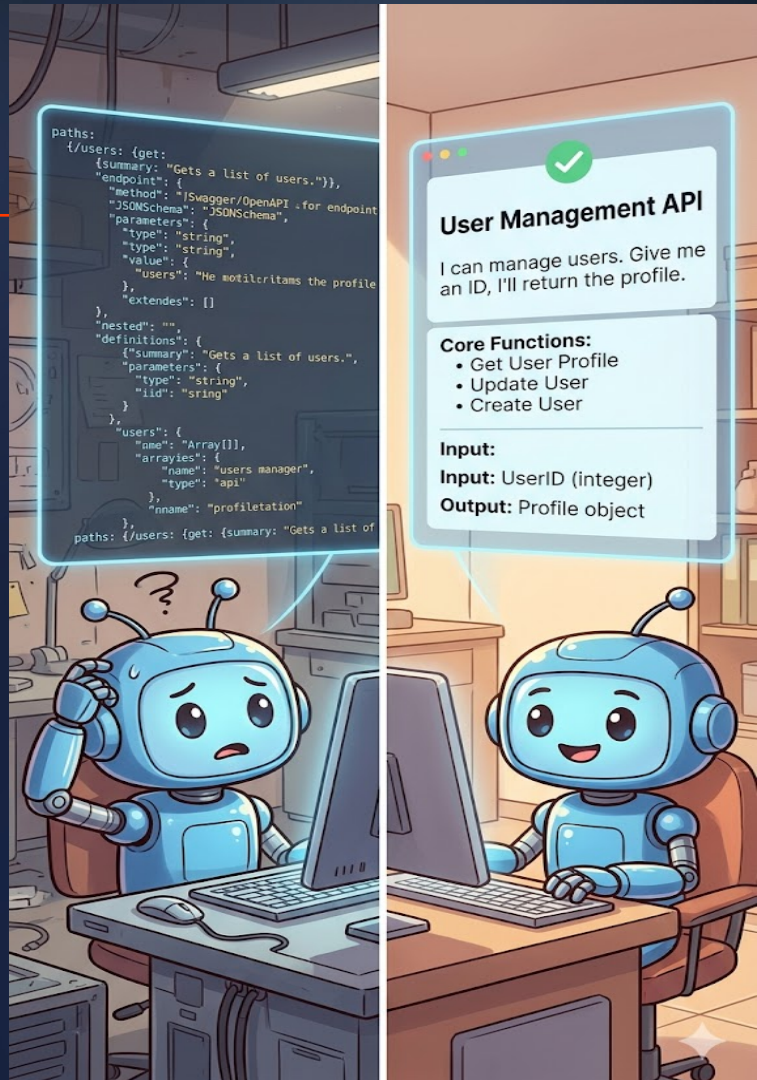

Skills for Agents

Skills bring expertise to agents:

Skills are how agents know what to do — and how to do it right.

INTERACTION

Enterprise Connectivity

How to talk to enterprise APIs and databases.

EXECUTION

Operational Compliance

How to operate within your enterprise — policies, workflows, and constraints.

INTELLIGENCE

Institutional Memory

Accumulated knowledge and learned lessons.

To manage skills at scale, you need a Skills Registry.

LLM Governance

TRADITIONAL STACK

Generic APIM

- Request-based rate limits
- No visibility into token costs or LLM-specific authorization per caller

AGENTIC STACK

AI Gateway for Agents

- Token-based cost budgets per agent identity
- Model routing, caching and fallback policies
- IDP-enforced authentication

Safety & Guardrails

Semantic Guardrails

- Detection of attack (e.g. Prompt injections)
- Avoid data leakage (e.g. PII detection)
- Control LLMs behavior (e.g. decorators) and Moderations
- Human approval
- ...and many more

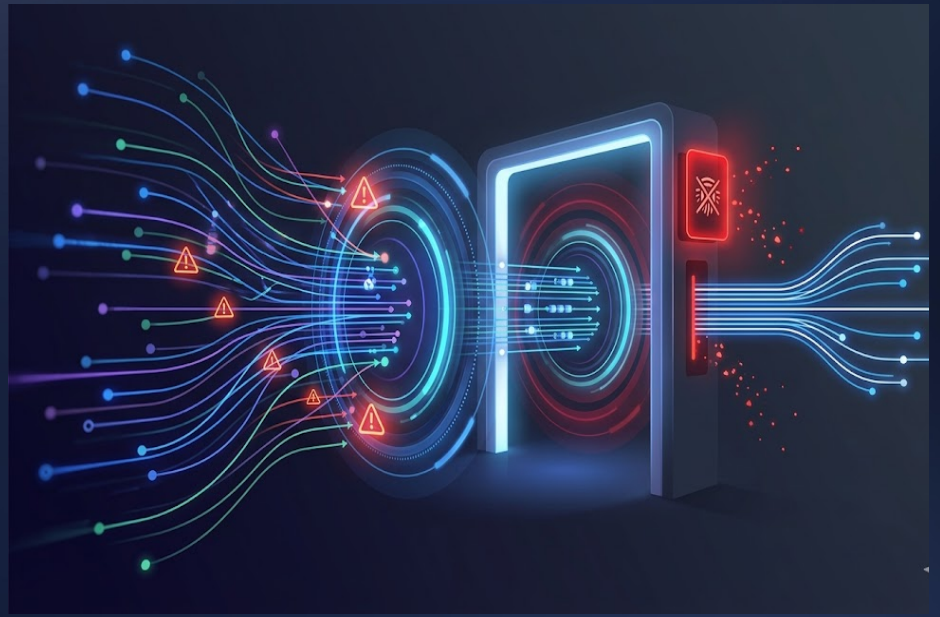

Agent Runtime

Traditional Stack

Bounded Execution

- Stateless execution model
- Failure means restart from scratch
- Zero trust at deployment

Agentic Stack

Stateful Execution

- Long-running, state-persistent execution
- Zero trust per task (Sandboxed environments)
- Granular per-agent resource quotas

Context & Memory Fabric

TRADITIONAL STACK

Query-based Access

- Data in SQL/S3 exists for human query
- Queries written into application logic
- Exists but designed for human or application access — not agent-queryable.

AGENTIC STACK

Agent Consumable Data

- RAG pipeline on unstructured data. Text-to-SQL wrappers for structured data.
- Persistent long-term memory via semantic wrappers
- Indexed for easy discovery by agents

Observability and Evaluations

TRADITIONAL STACK

Infrastructure Metrics

Latency and CPU usage.

Application logs to debug functional errors.

AGENTIC STACK

Reasoning Trace Spans

LLM-specific OTel spans for reasoning steps and tool calls.

Distributed tracing across agent swarms.

Monitoring & Evaluation

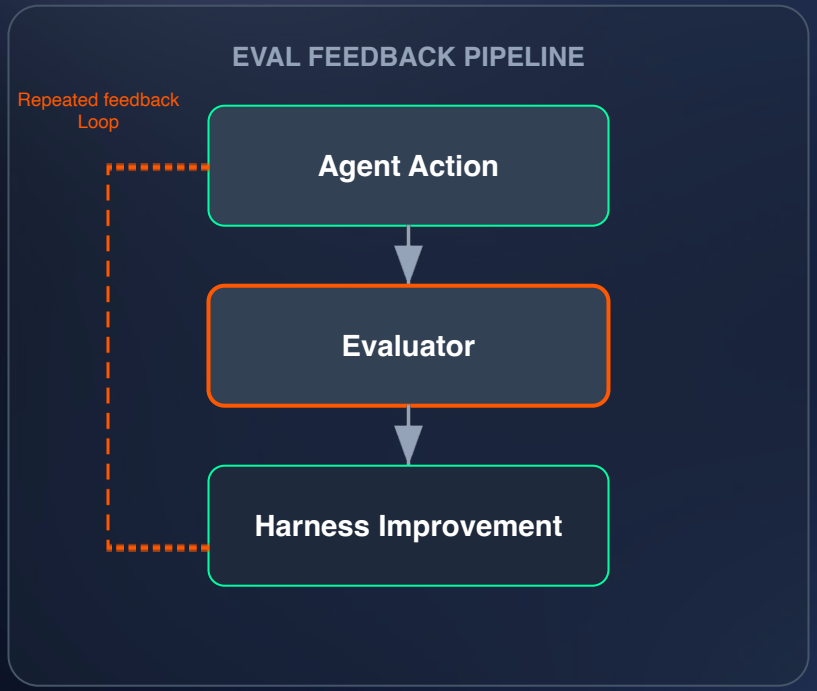

Standardizing Agent Interface

Agent logic and interface are independent

The interface is a platform concern, not an agent concern.

Chat Agent-to-human — conversational, on-demand

A2A Protocol Agent-to-agent — discovery and task delegation

Event Trigger Ambient agents — background, event-driven execution

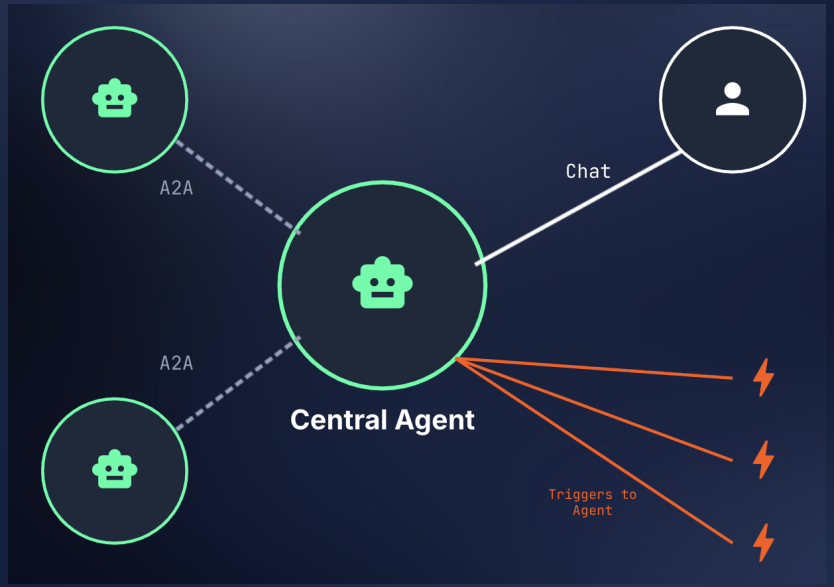

Agent Lifecycle Management

PHASE 01

Define

description

- Goals
- Responsibilities

PHASE 02

Design

architecture

- Agents
- Orchestration
- Tools/Memory
- Guardrails

PHASE 03

Develop

code

- Workflows
- Integrations
- Context

PHASE 04

Evaluate

fact_check

- Safety
- Reliability
- Accuracy

PHASE 05

Deploy

rocket_launch

- Environments
- Configuration
- Rollouts

PHASE 06

Operate & Govern

settings_suggest

- Observability
- Policies
- Drift
- Optimization

Agents are inevitable.

The only question is whether you have the **right control plane** to close the gap.

WSO2 Agent Manager

UNIFIED PLATFORM

The open control plane for the Agentic Enterprise, securely deploy, manage, and govern AI agents at scale — for ***any agent***, ***any model***, ***any runtime***.

Lifecycle
Full CI/CD with automated evaluations and managed runtime.

Observe
Distributed Semantic Tracing & Eval Loops.

Govern
Model-agnostic AI Gateway & Token Budgets.

Secure
First-class Agent Identity & Scoped Guardrails.

WSO2CON NORTH AMERICA

May 20 - 22, 2026 | Austin, Texas, USA

Thank You!