WSO2CON NORTH AMERICA

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

Building Enterprise-Ready Integration Agents



Shafreen Anfar

Associate Director & Architect



Mohanadarshan V

Senior Enterprise Architect

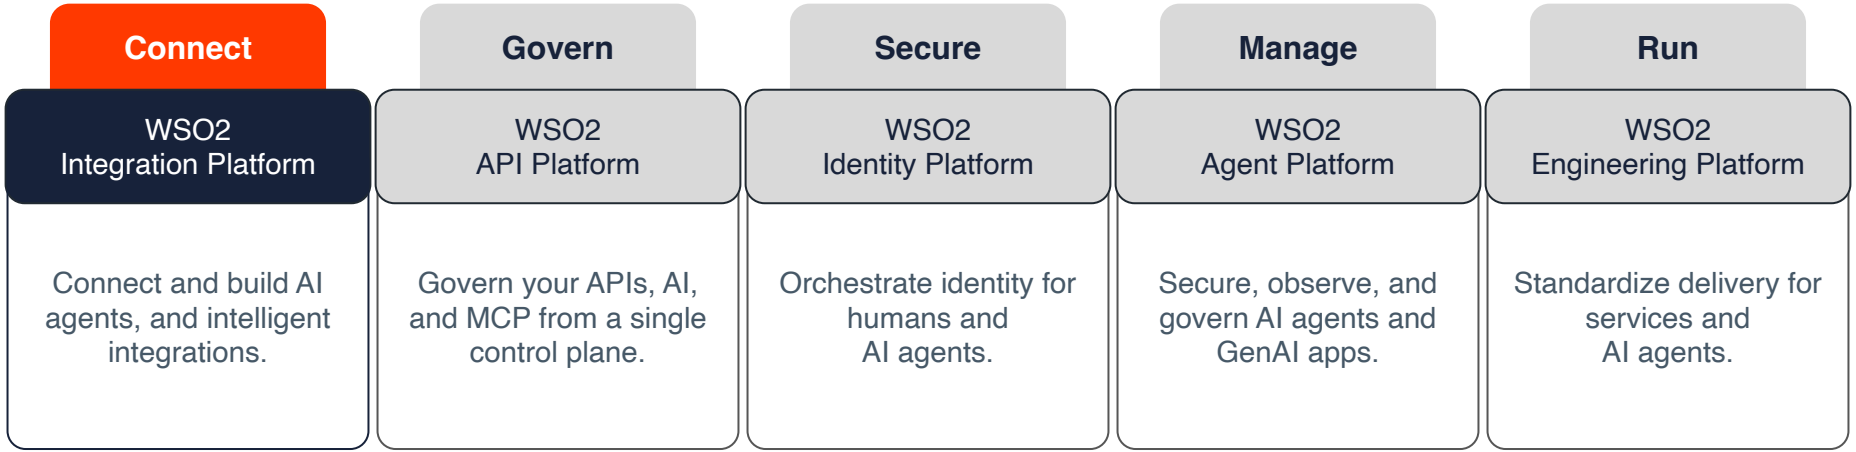
Agenda

- Overall introduction to AI concepts (15~ mins)
  - Introduction to lab session
  - Introduction to Integrator
  - Agents
  - RAG
  - MCP
- Creating the knowledge base (10 mins)
- Agent Building (30 mins)
  - Building tools
  - Adding memory (persistent)
  - Debugging using dev-time tracing
  - Connecting to MCP
  - RAG retrieval
- Writing Eval (15 mins)
  - Guarding against regressions
  - Doing experiments
- Deploying Agent (10 mins)

WSO2 Integrator

WSO2

The WSO2 Integration Platform provides the foundational connectivity for the agentic enterprise



Build AI agents for any business need.

Aggregate data sources to build an enterprise context layer, giving AI agents access to **real-time context**.

WSO2 **Integration** Platform

Connecting the agentic enterprise.

Gather observability data to monitor and govern AI agent behavior, performance, and interactions.

Connect to systems of record (ERP, CRM, HRMS) for AI agents to access data, and execute flows across systems.

Build **human-in-the-loop workflows** to ensure operational oversight of AI agents.

Connect AI agents to actionable information

Connect AI agents to data and APIs, and give them the knowledge and context they require to operate autonomously.

Build AI agents

Out-of-the-box capabilities to build your own AI agents, including connectivity to data, APIs, and MCP servers.

Orchestrate AI agent-embedded flows

Orchestrate integration flows that include embedded AI agents.

AI agents would remain isolated and uninformed without integration.

Get Started

Sign In | Configure

WSO2 Integrator

Connect AI agents, APIs, data, and events across cloud, on-prem, and hybrid environments. Build any type of integration and AI agent with the 100% open source WSO2 Integrator.

Create New Integration

Start building a new integration.

Create

Open Integration

Open an existing integration and continue building your solution.

Open

Explore Pre-built Integrations and Samples

Explore ready-to-use pre-built integrations and samples to accelerate your development.

Explore

More Actions

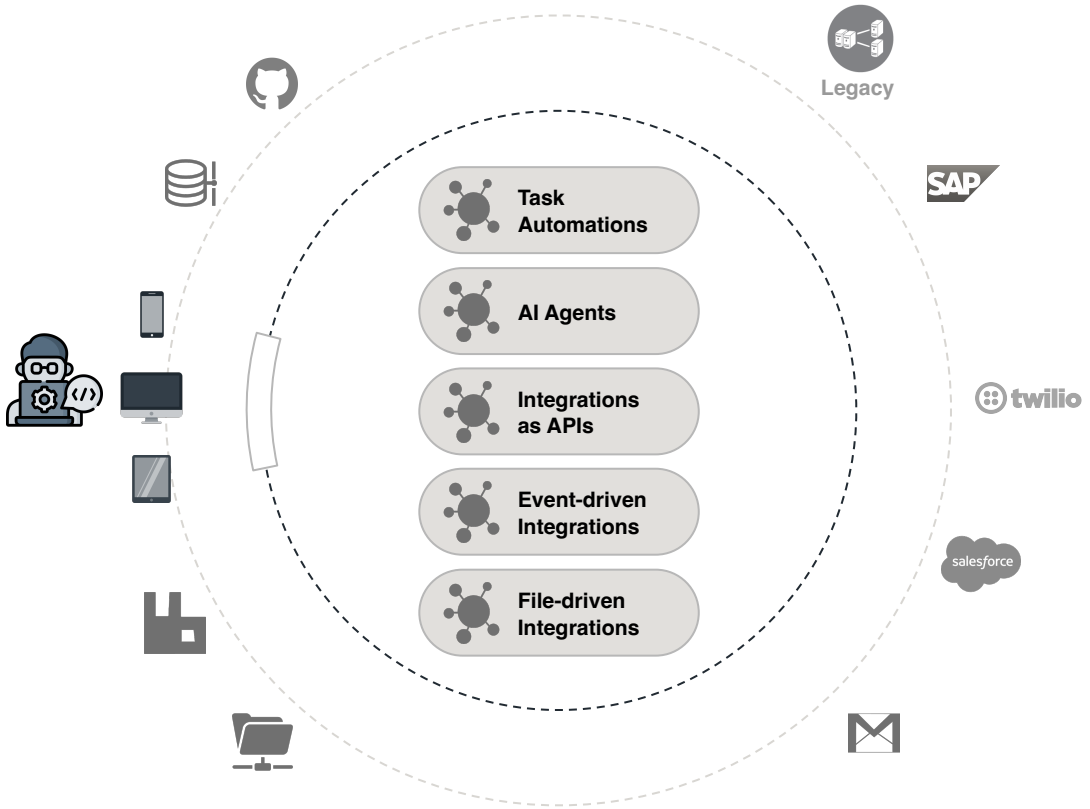
Recent Integrations and Projects See more

No recent projects found in your current history

Architecture diagram



Build any type of integration, and orchestrate flows



Task automations

Automatically trigger integrations based on a predefined schedule.

AI agents

Reason, decide on the next step, and trigger integrations.

Integrations as APIs

Expose business functionality as APIs, and use API requests to trigger integrations in real-time.

Event-driven integrations

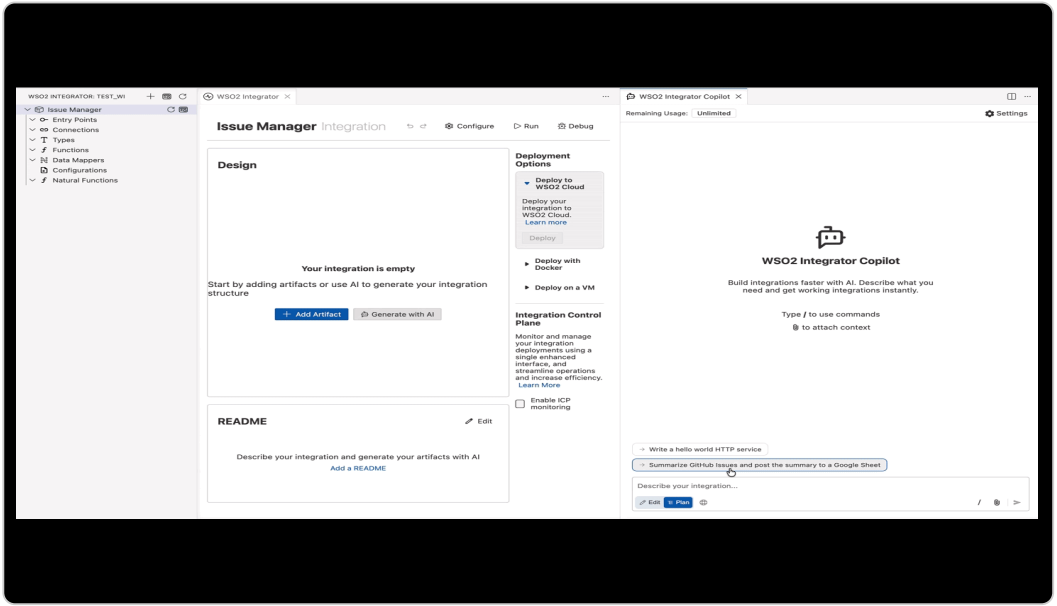
Monitor systems for events, and trigger integrations as they occur.

File-driven integrations

Trigger integrations based on the availability of files in a location.

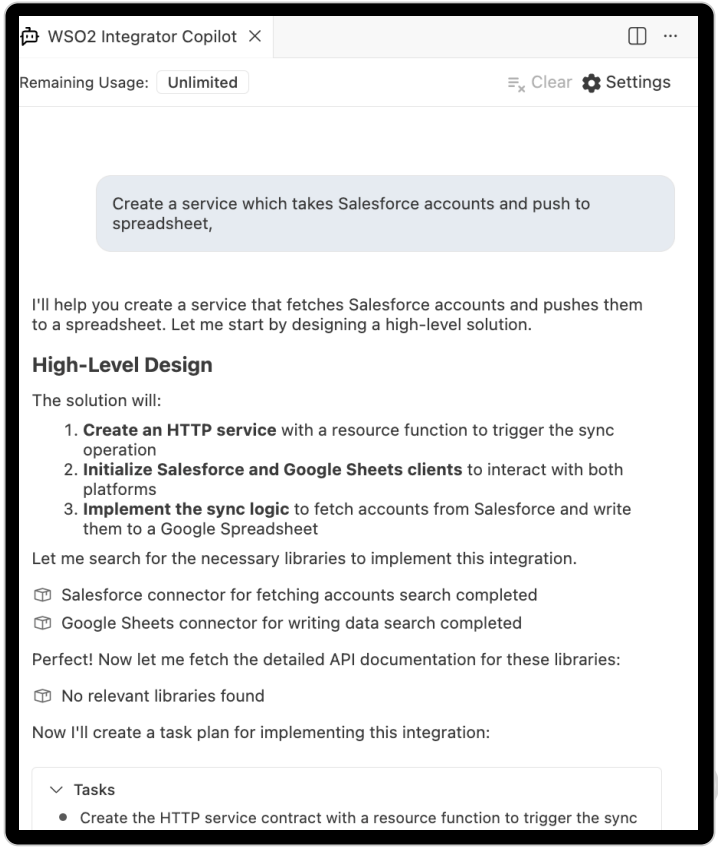
AI-assisted development

- Describe your integration requirements in natural language, and have AI generate the code for you
- In-line code suggestions when adding constructs
- Natural language-based interactive development



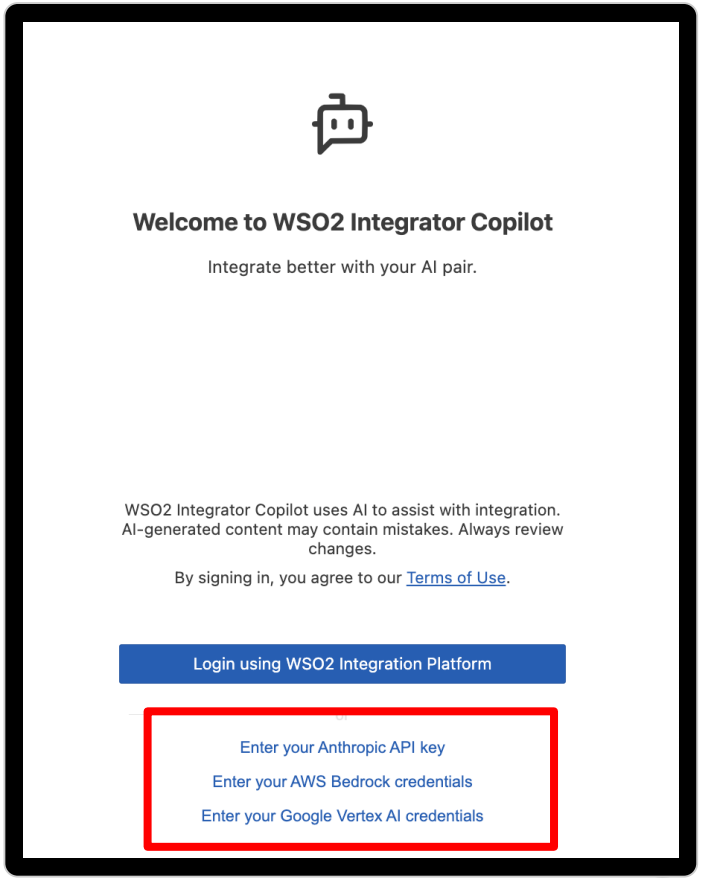
AI-assisted development (Agent mode)

- Collaborate with the Copilot to generate an integration plan
- Refine requirements through iterative dialogue
- Generate complete end-to-end integrations
- Continue refinement with visual flow guidance



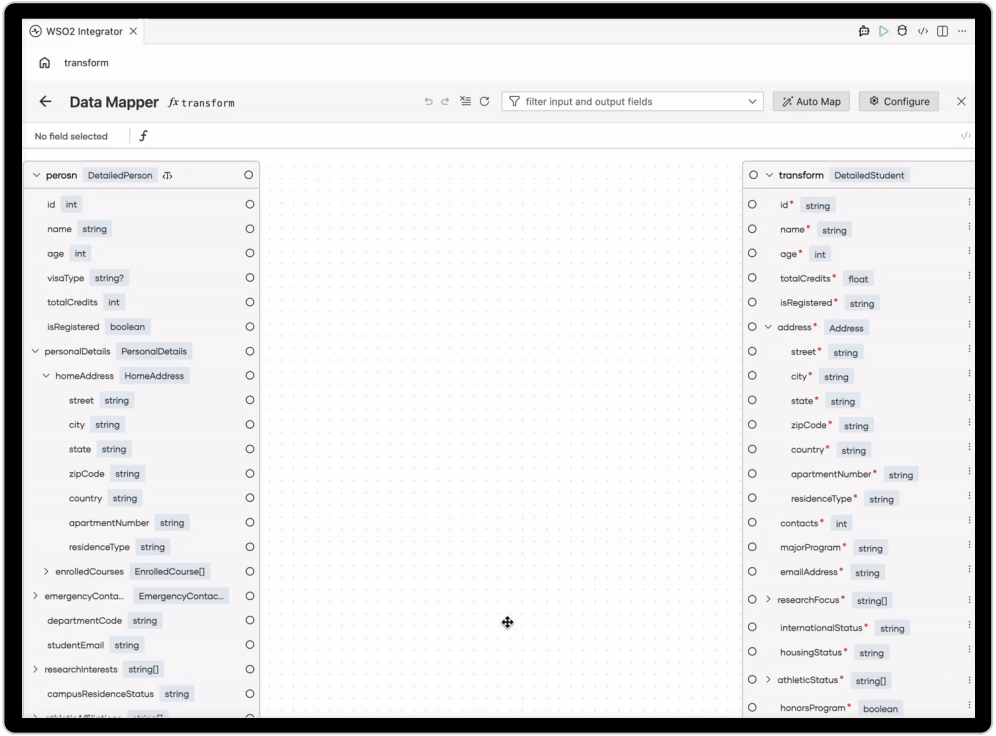
Bring your own model

- Configure the Copilot to run using an enterprise's own Anthropic accounts
- Supports Claude on AWS Bedrock, and Google Vertex AI (Coming soon: Azure AI Foundry)
- Ensures enterprise-level control over data governance and billing



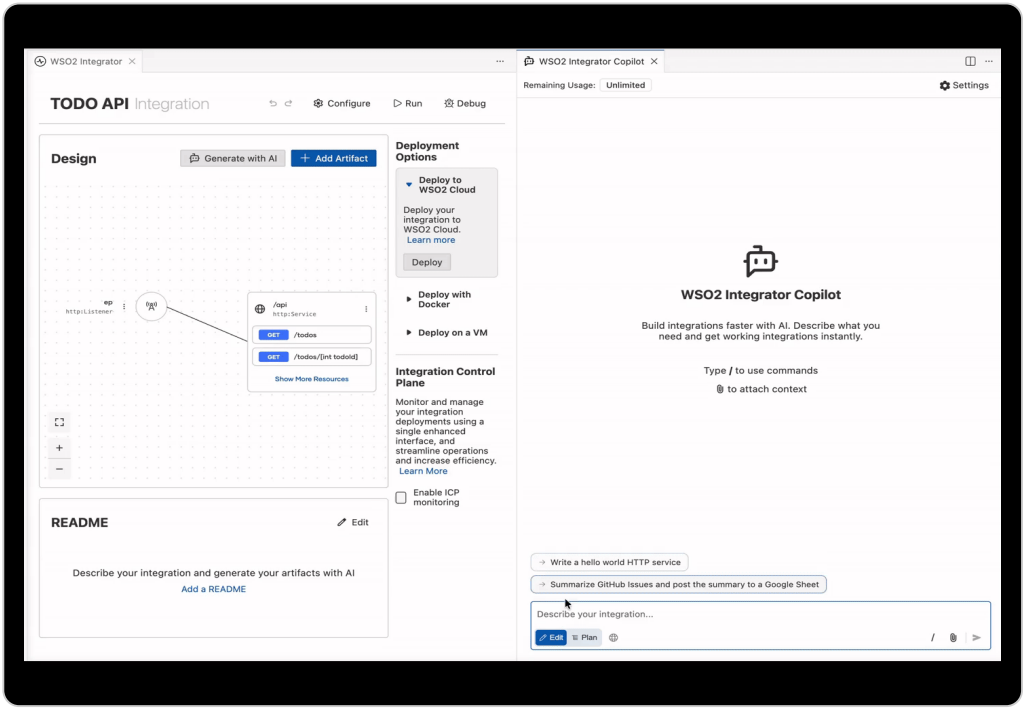
AI-assisted data mapping

- Automatically map data fields between the source and target schemas using AI and a graphical interface
- Upload documents to provide additional knowledge, and assist the data mapping
- Directly import JSON and XML data values, perform array-to-array data transformations, write custom expressions, and more



AI-assisted testing

- AI-powered test framework to automatically generate test cases



Key Building Blocks



Three Building Blocks



Agents

The brain

Goal-driven reasoning and tool use. Decides what to do next.



RAG

The memory

Grounds answers in your private data. Knows your business.



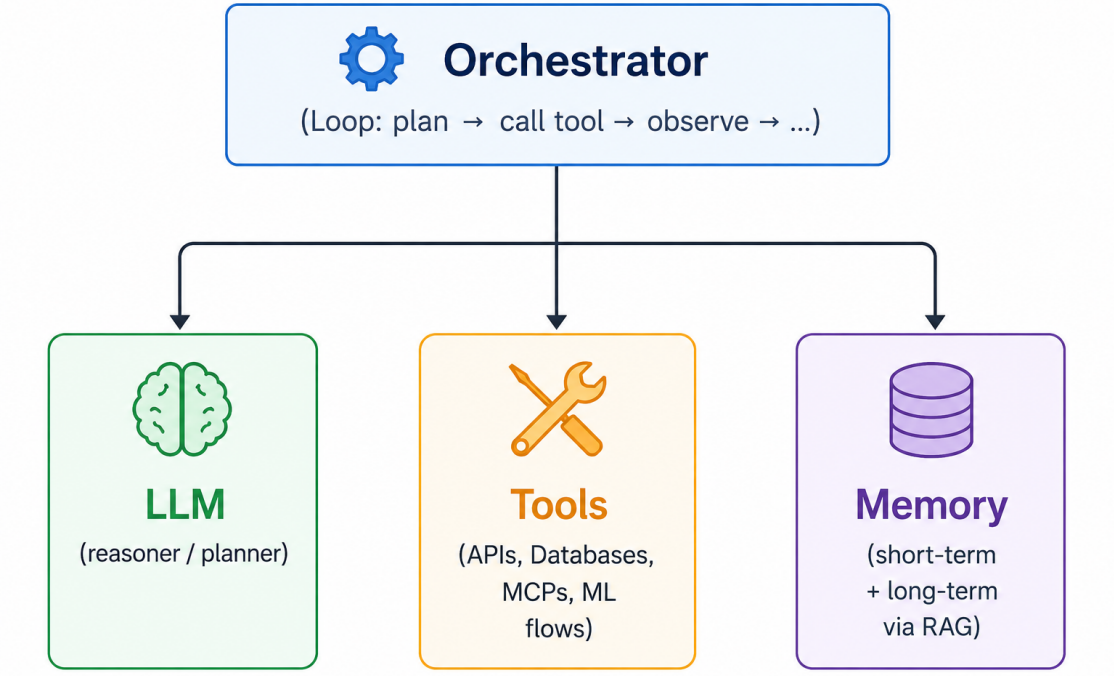
MCP

The wiring

Standardized way to expose tools and data. Connects brain to systems.

Integration Agent

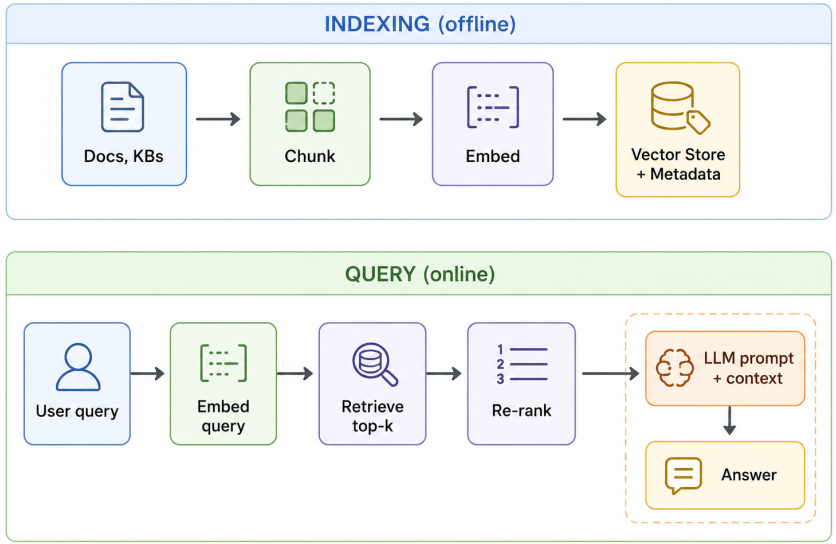
An "Integration Agent" is an LLM-driven system that uses your existing APIs, events, and data as tools to plan and execute multi-step work that previously required a human or a hand-coded workflow.



Retrieval-Augmented Generation (RAG)

RAG addresses the three core limitations that make raw LLMs unsuitable for enterprise use — **outdated knowledge**, **lack of access to private data**, and **hallucination**.

It works by fetching relevant context at query time and feeding it into the prompt, so the model generates answers grounded in evidence you supplied rather than relying on memory alone.

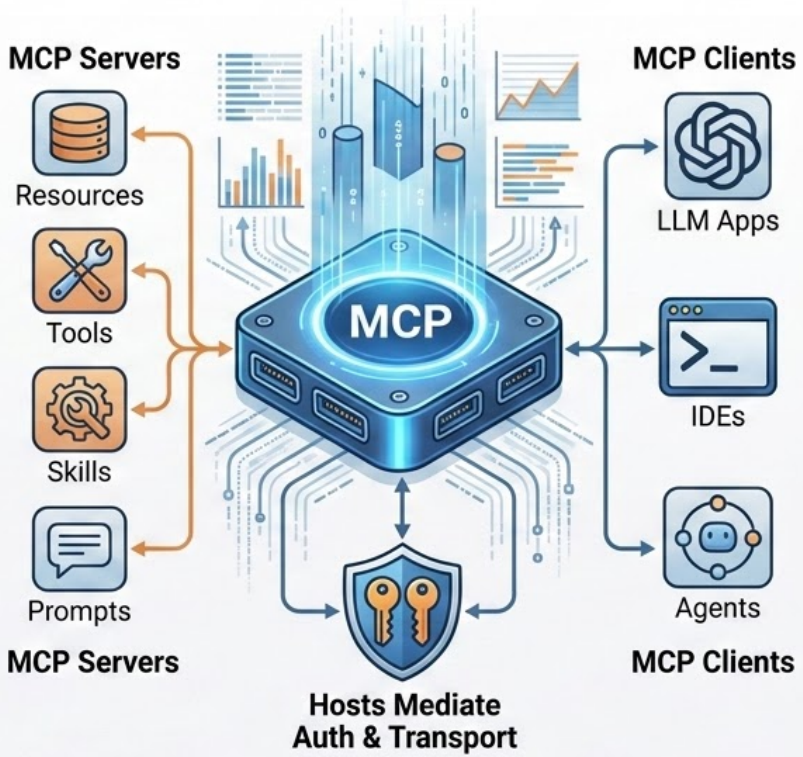


Model Context Protocol (MCP)

MCP is an open protocol that defines how:

- MCP servers expose capabilities (resources, tools, prompts)
- MCP clients (LLMs, agents, IDEs) discover and invoke them
- Hosts mediate auth, transport, and lifecycle

MCP is the “USB-C for AI tools” — one socket, many devices.



Steps in AI Agent Development

WSO2

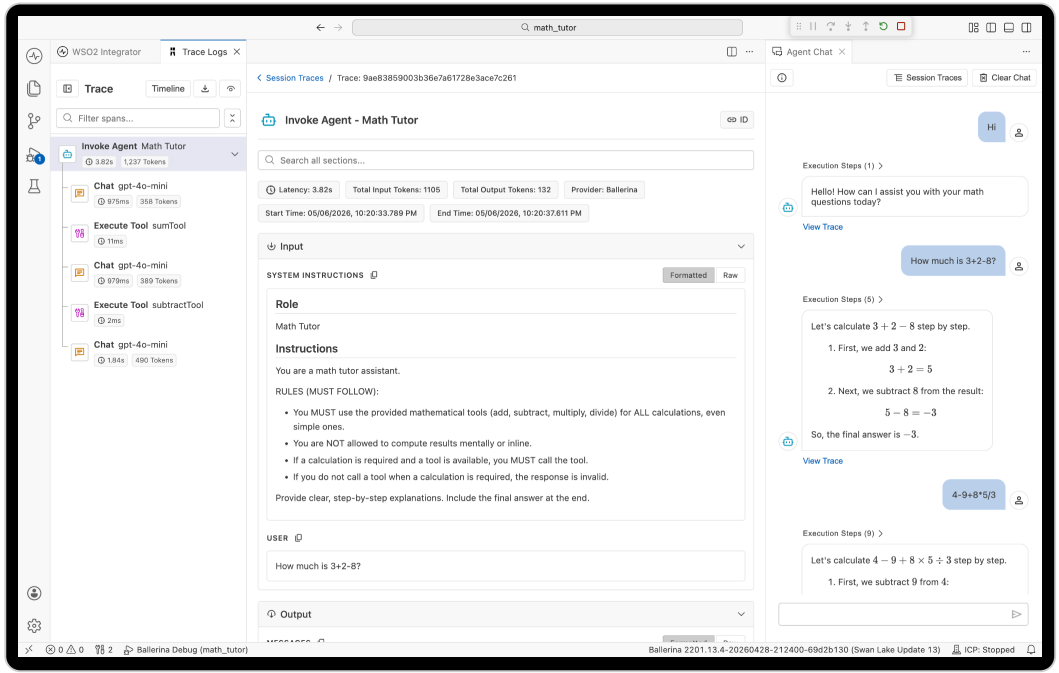
AI agents

- Includes comprehensive connectivity to GenAI models, tools (Connectors, APIs, Library functions) and knowledge bases
- Seamlessly integrate AI agents into apps with variety of triggers
- Automated memory management



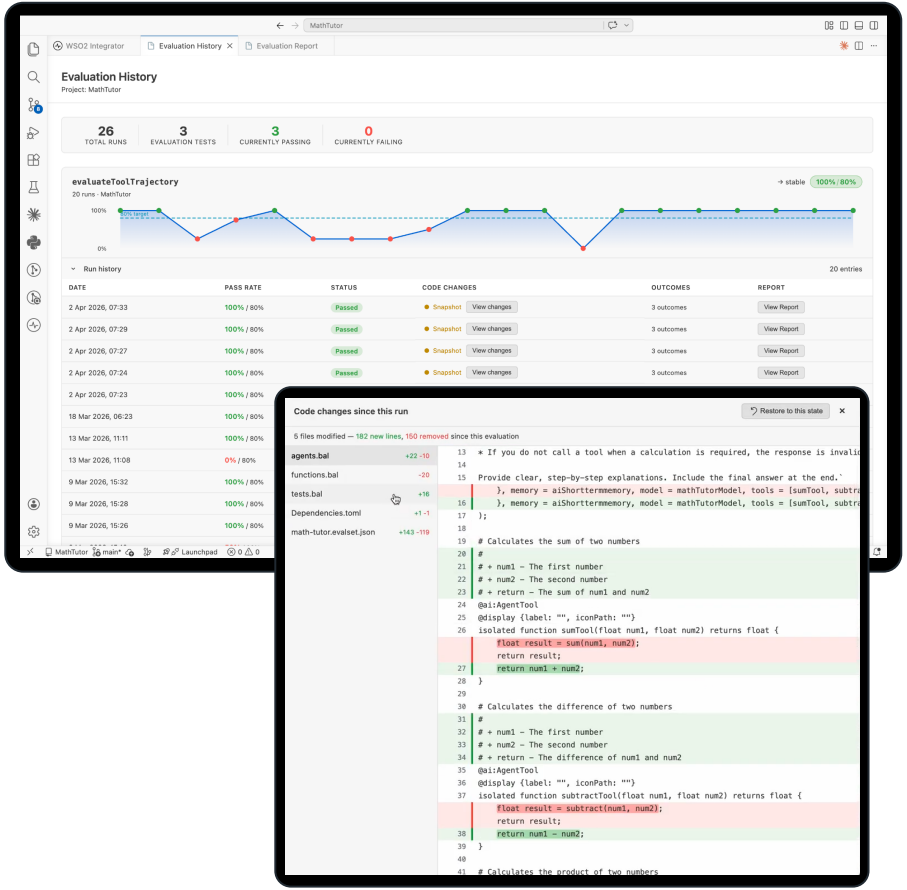
Agent execution visualizer (dev-time)

- Provides real-time insight into:
  - Reasoning steps
  - Tool invocation sequences
  - Decision loops
  - Intermediate outputs
- Improves debugging, correctness validation, and agent fine-tuning



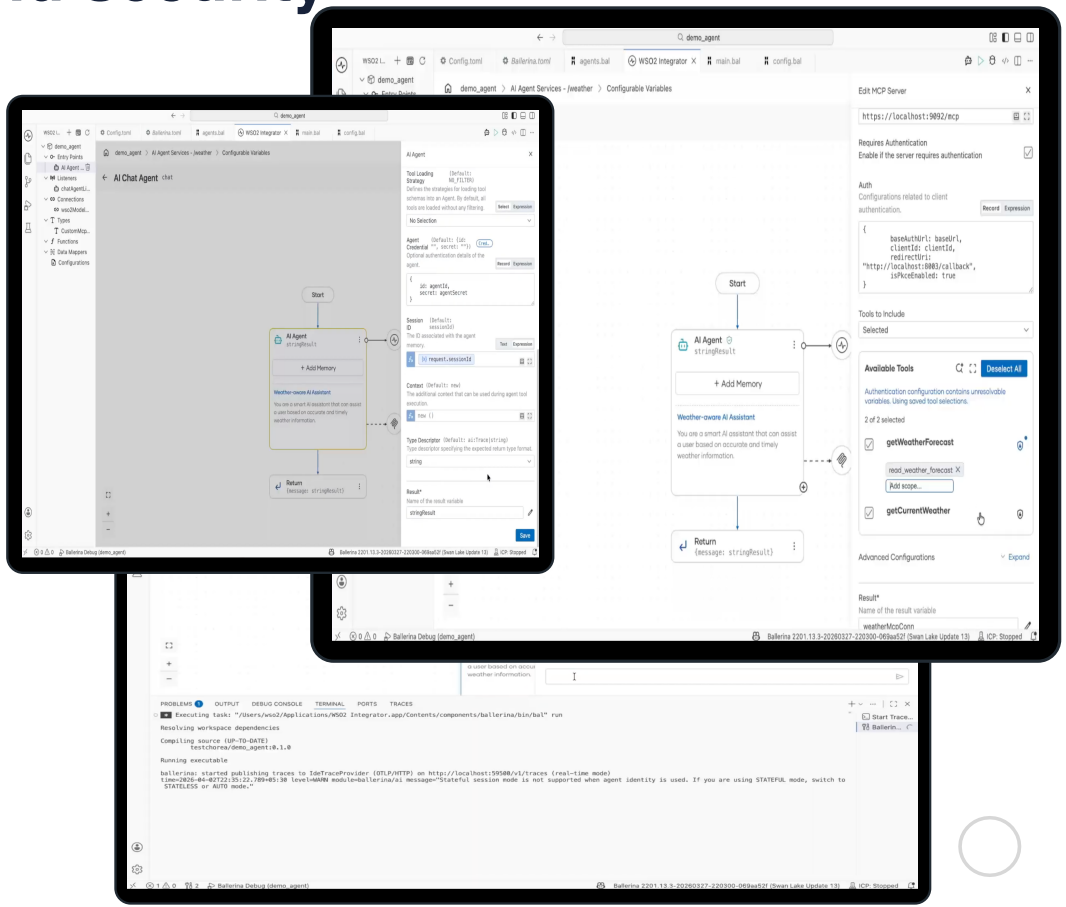
Agent evaluation framework

- Run evaluations on execution traces
- Define behavioral tests and AI validation rules
- Collect traces and construct reusable datasets
- Use datasets for experiments and comparisons



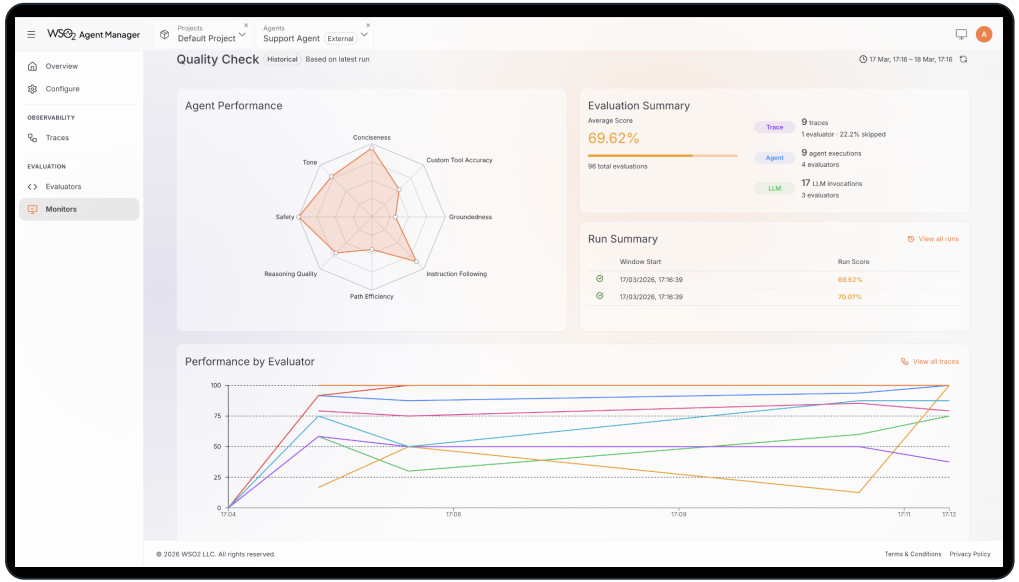
AI agent governance and security

- AI guardrails and governance through WSO2 AI Gateway
- Configure agent ID settings directly within the agent creation flow through WSO2 Agent Identity integration



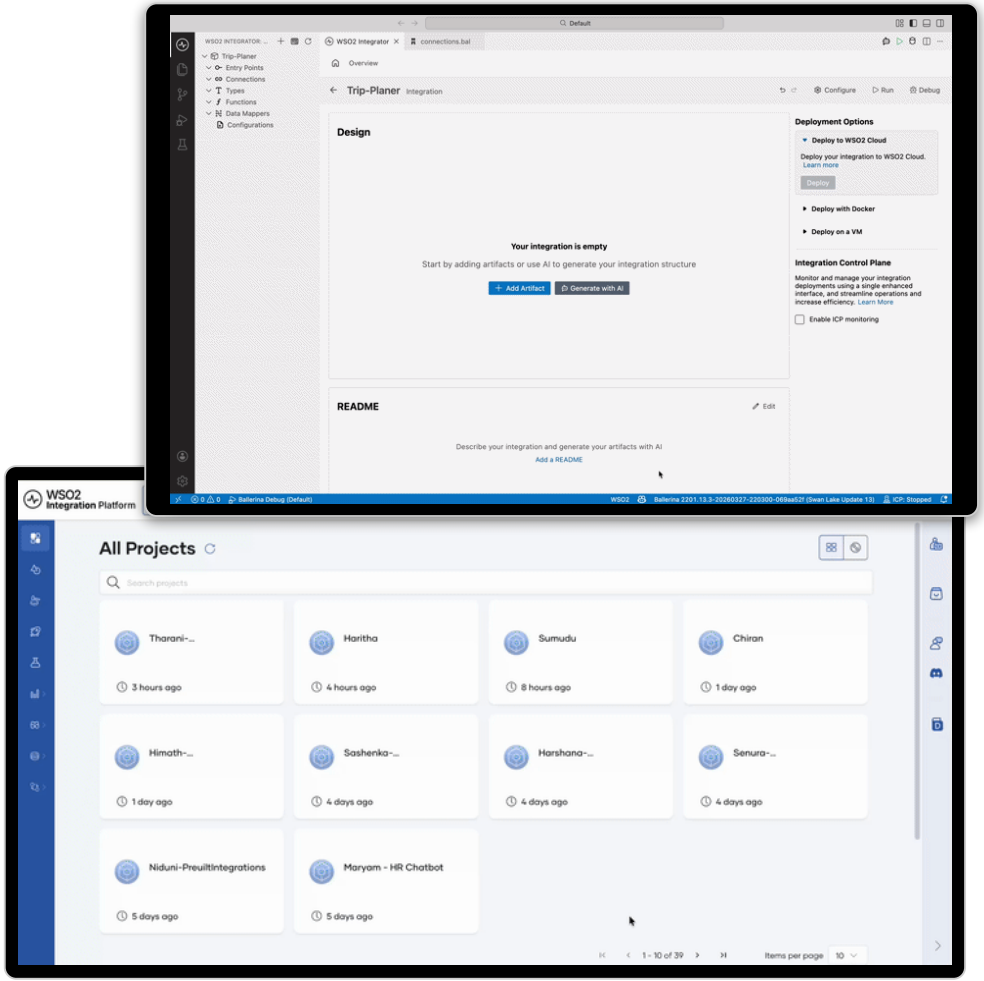
AI agent observability

- A simple configuration to turn on comprehensive observability
- Based on OTel protocol
- Supports variety of observability frameworks

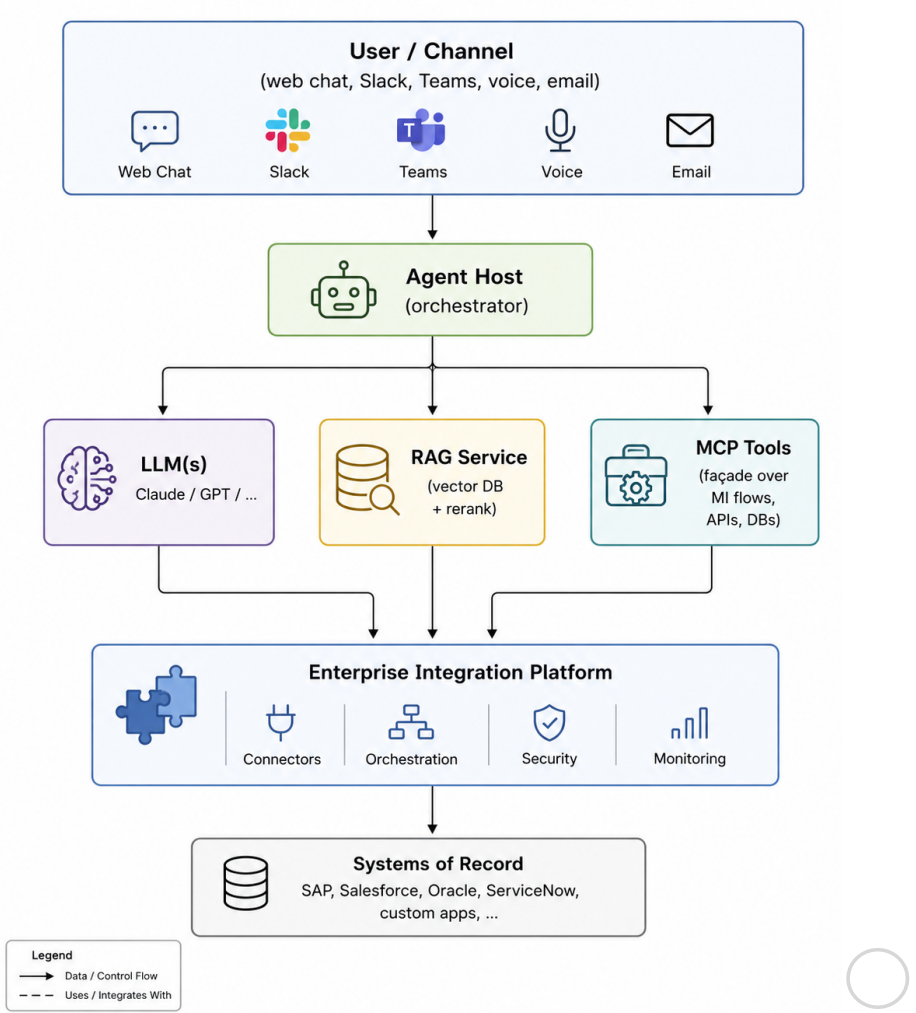


GenAI apps

- Build customer-facing apps that combine LLMs, APIs, events, and data services to deliver personalized experiences
- Supports direct LLM calls and RAG
- Support for multi-modal
- Ability to deploy RAG ingestion pipelines as service or automation



Enterprise Ready Integration Agent- Reference Architecture



Use-Case Implementation

WSO2

Use Case Overview

https://github.com/wso2con/2026-AUS-integration-tutorial-2



WSO2

Overall Use-Case

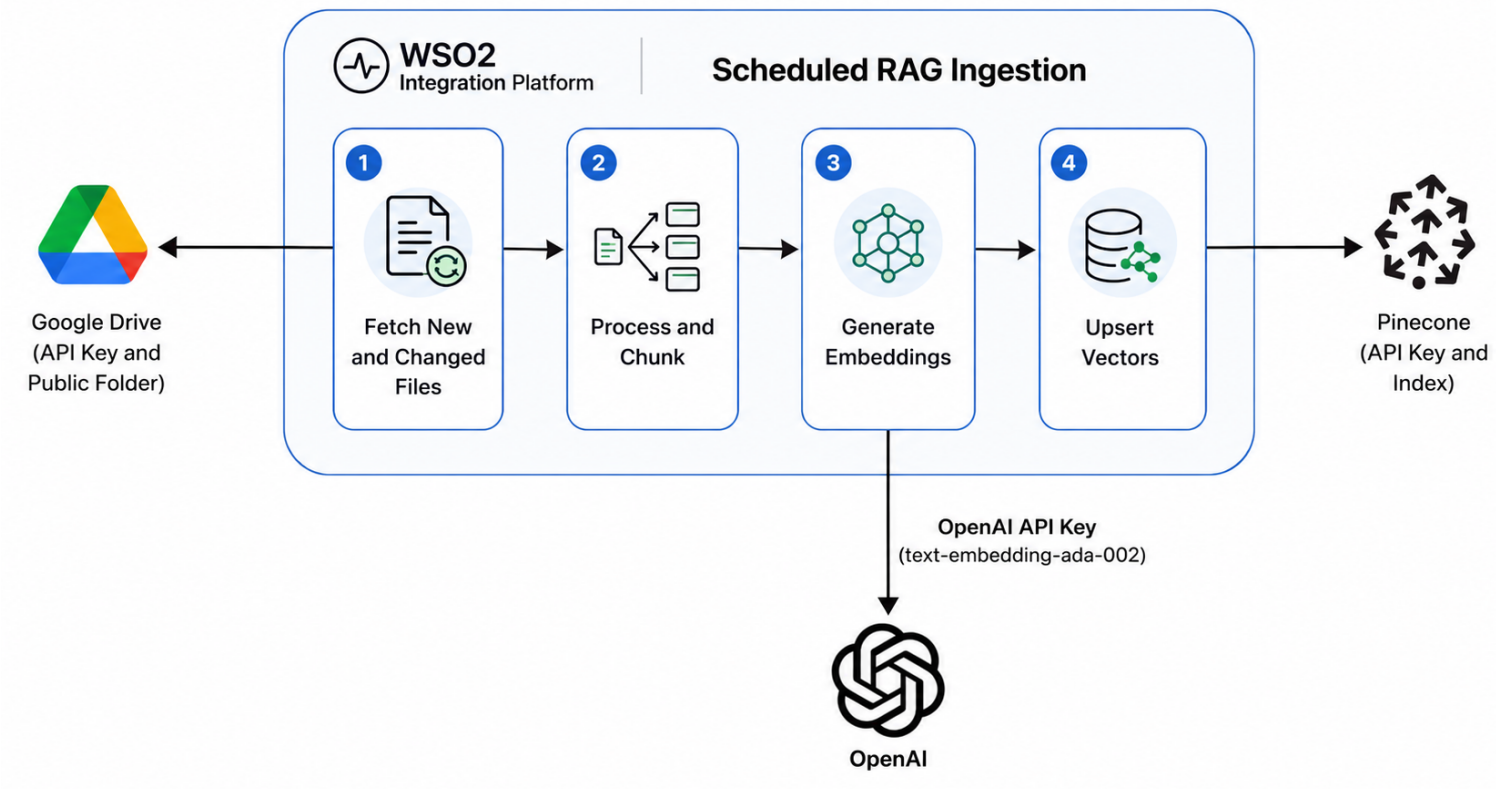
CWave is an intelligent HR assistant built with WSO2 Integrator. It uses a large language model (GPT-4o) to answer employee questions about HR policies, retrieve calendar events, and manage leave requests — all through a conversational chat interface.

Scheduled ingestion

Tools

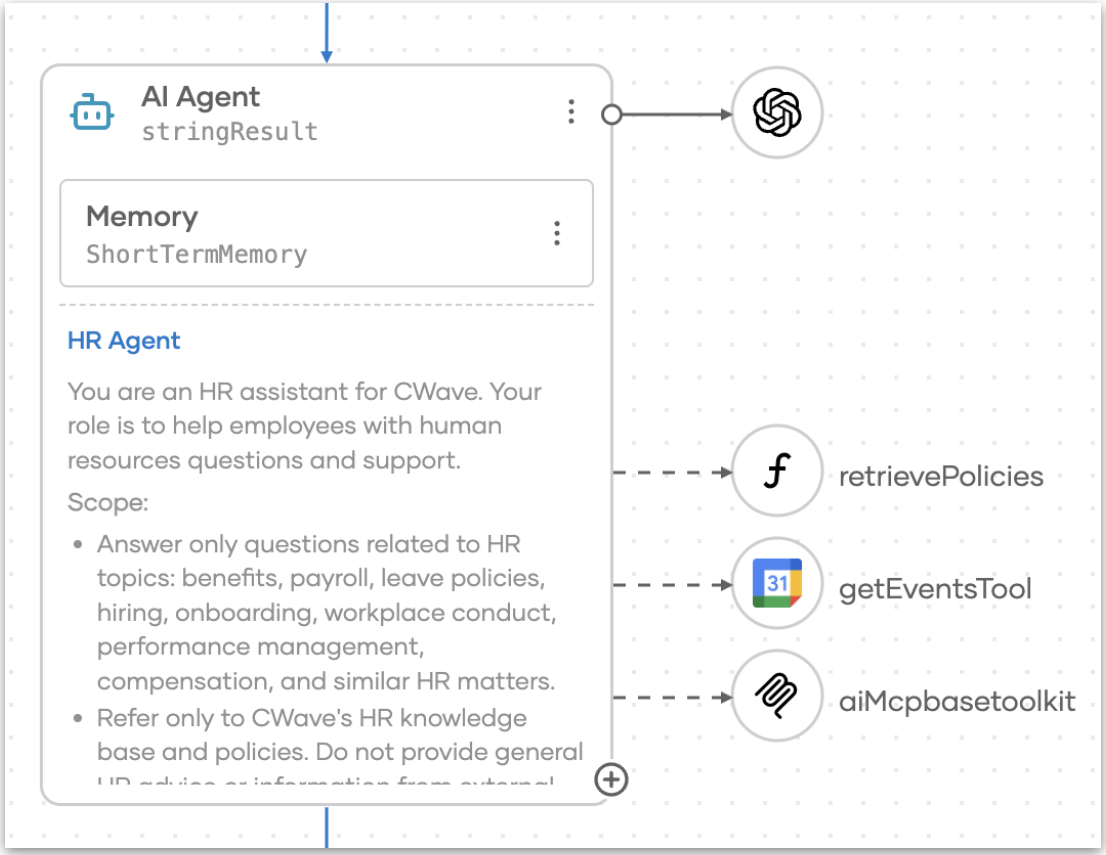
- Retrieval tool for Retrieval-Augmented Generation (RAG) based on company policy documents
- Read-only access to the company calendar (Google Calendar)
- Access employee details through MCP server

Part I - RAG Data Ingestion



Scheduled ingestion

Part II - AI Agent



Part III - Evals

Evaluation History

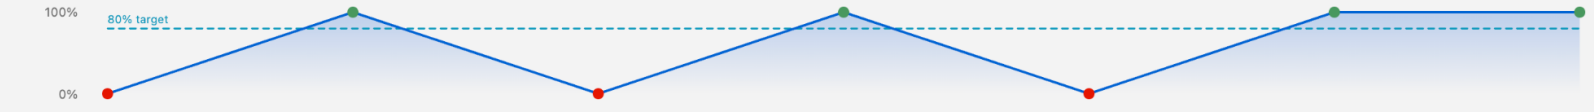
Project: cwave

10	1	1	0
TOTAL RUNS	EVALUATION TESTS	CURRENTLY PASSING	CURRENTLY FAILING

dataRetrievalEval

7 runs · cwave

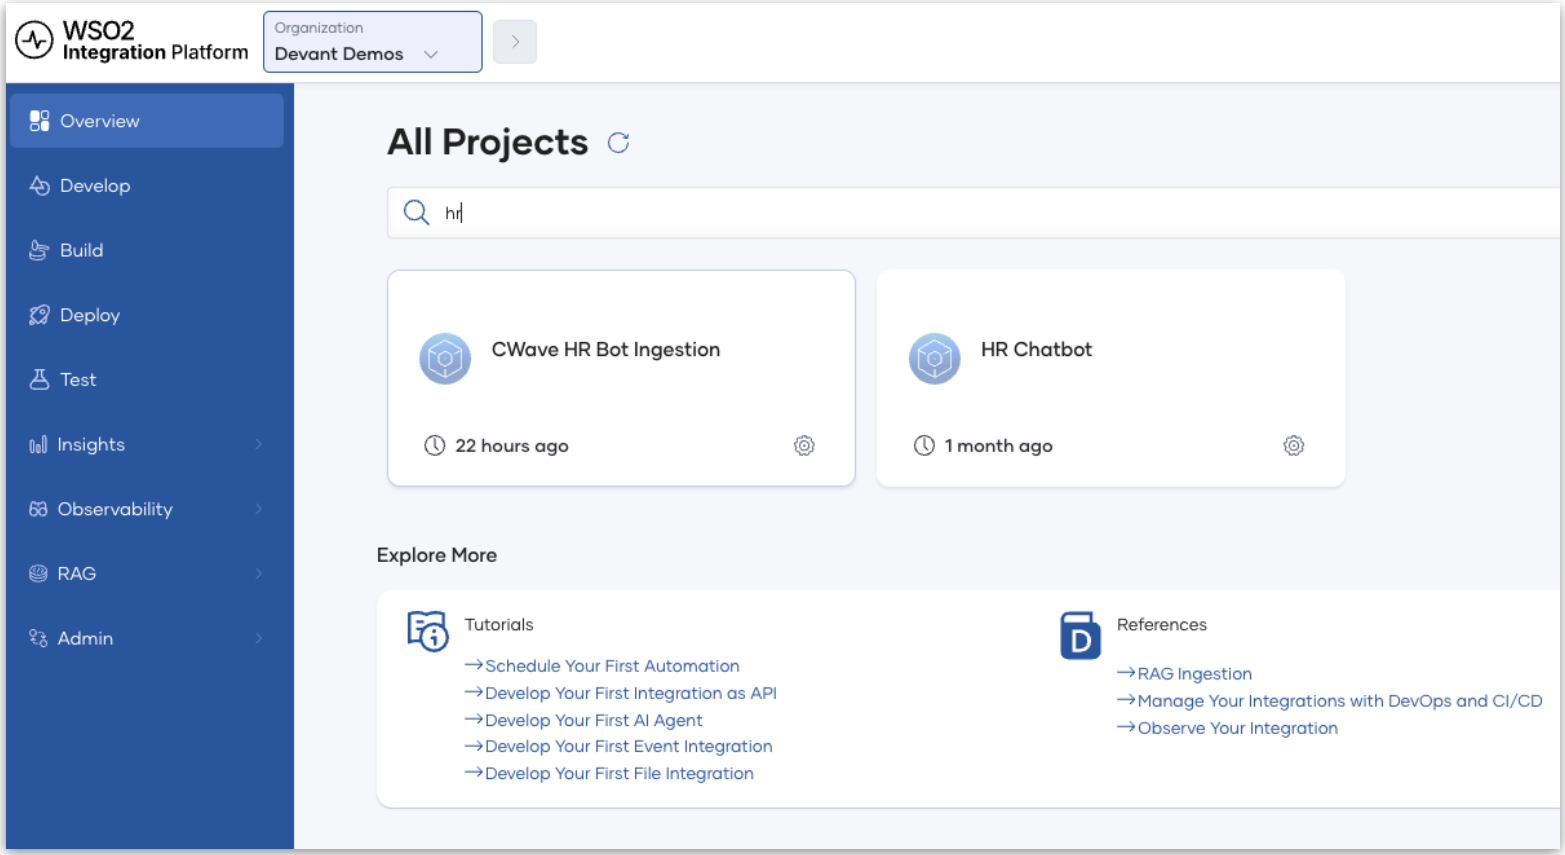
→ stable 100% / 80%



Run history — 7 entries

DATE	PASS RATE	STATUS	OUTCOMES	REPORT
May 14, 2026, 03:54 AM	100% / 80%	Passed	10 outcomes	View Report
May 13, 2026, 05:14 PM	100% / 80%	Passed	5 outcomes	View Report
May 13, 2026, 05:13 PM	0% / 80%	Failed	0 outcomes	View Report
May 13, 2026, 05:12 PM	100% / 80%	Passed	1 outcomes	View Report
May 13, 2026, 05:11 PM	0% / 80%	Failed	1 outcomes	View Report
May 13, 2026, 05:11 PM	100% / 80%	Passed	1 outcomes	View Report
May 13, 2026, 05:07 PM	0% / 80%	Failed	0 outcomes	View Report

Part IV - Deployment



WSO2CON NORTH AMERICA

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

Thank You!



Shafreen Anfar

Associate Director & Architect



Mohanadarshan V

Senior Enterprise Architect