



SEPT 22 - 24, 2026 | NAIROBI, KENYA

Operating Models for the Agentic Era



Amila Mahaarachchi

Senior Vice President &
Head of Engineering

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.



WSO2 as an Agentic Enterprise

Where it is in action

- R&D (Engineering)
- Customer Reliability Engineering
- Site Reliability Engineering
- Marketing Operations
- Revenue Operations
- WSO2 Digital



Towards an Agentic Engineering Team

WSO2 Engineering Team

- 300+ Engineers
- 5 Products (owned by 5 Business Units)
 - On-prem and SaaS
- 2 Solutions
- 40+ product versions used in production by customers



Prior to 2026

Fragmented AI Adoption

Fragmented Tooling

Engineers utilized a mixed environment of GitHub Copilot, Claude Code, and OpenAI Codex.

Inconsistent Outcomes

Productivity gains were heavily dependent on individual skill levels regarding prompts and context management.

Missing Synergy

While common engineering challenges existed across teams, the lack of a unified toolset prevented organization-wide AI initiatives.



Standardizing

Provided a unified engineering assistant for all team members.

- **Unified Platform:** We selected Claude as the official AI tool to replace the fragmented mix of assistants.
- **Skill Equalization:** This standardization leveled the playing field, equipping every engineer with the exact same baseline capabilities.
- **Shared Foundation:** Establishing a single official tool enabled us to build consistent, team-wide agentic workflows and solve common problems collaboratively.



Software Engineering Agent

- Claude Code is the agent
- Product and WSO2 specific skills defined
- Every engineer has access to the agent and skills
- These skills made way to almost everything we do at R&D



Agent-friendly code repositories

AI-Optimized Environments

Code repositories for both mature existing products and new innovations were explicitly prepared for coding agents.

- Structures that support seamless navigation for artificial intelligence tools.

Refer: <https://github.com/wso2/api-platform>

Embedded Intelligence & Execution

Embedded Intelligence: Integrated specific skills and rules directly into the repositories to guide the AI.

Seamless Implementation: This preparation allowed coding agents to immediately understand our architecture and coding standards, etc.



Codifying Bug-Fixing Steps

Scale

We actively support nearly 1,000 customers utilizing over 40 different versions of our products.

Standardized Execution

Defined and codified the bug-fixing step into custom AI skills, including: reproduce, plan-fix, assess-risk, implement, verify-fix, create-tests, and submit-fix.

Consistent Quality

Any engineer can now execute maintenance tasks with consistency, allowing us to fix and port updates across versions faster than before.



Automated Documentation Issue Fixing Workflow

Automated Integration

- We built a custom GitHub Actions workflow powered by Anthropic APIs.

End-to-End Resolution

- Once a Github issue is created, the workflow automatically identifies problem, generates the necessary fixes, and submits Pull Requests independently.

Time Savings

- This automation successfully cut down the manual effort required to fix a documentation issue from 30 minutes to just 5 minutes.

See below for examples

<https://github.com/wso2/docs-apim/issues?q=is%3Aissue%20state%3Aopen%20label%3AAI-Agent%2FFixed>

30 Min

Time to fix a documentation issue (Before)

5 Min

Automated Resolution Time (After)



Proactive Vulnerability Detection

- **Proactive Scanning:** We utilize Claude to actively scan our entire codebase to identify security issues proactively.
- **Automated Report Analysis:** Custom skills were written specifically to parse and analyze complex security reports efficiently.
- **Risk Mitigation:** This proactive approach helps us stay ahead of the curve, mitigating the looming risk of AI-discovered security vulnerabilities in the wild.



Agents in Action



e19166 5d ago

Fix results — apim 4.6.0 — local run by @e19166

setting	value
command	/fix
duration	19m 20s

Verdict: dev test passed — labeling swe-agent/fixes .

Fix implemented and dev-tested at runtime on APIM 4.6.0 (update level 46). The gateway MCP schema sanitiser now strips only known parameter location prefixes (query_, header_, path_, cookie_, formData_) instead of splitting on the first underscore, so backend snake_case names survive. Gateway tools/list now returns agreement_number/source_id/operate and tools/call reaches the backend with those names. Verified no regression for API-derived MCP servers: cookie_/header_/path_/query_ prefixes are still stripped and tool calls still route parameters correctly.

Run locally with the swe-agent plugin v0.1.15. Analysis and evidence stay on the operator machine — ask @e19166 for detail.



☐	☒	[Docs][4.7.0] Link to federating to Envoy gateway is not working	AI-Agent/Fixed	Type/Bug	Type/Docs	1 1	2
#11863 · Akila94 opened 2w ago							
☐	☒	[APIM] [4.7.0] Configuring access logs for the HTTP Servlet transport	API-M-4.7.0	Docs/type	documentation	Priority/High	1 1
#11859 · chandimajayawickrama opened 2w ago							
☐	☒	[Doc Fix] Fix invalid hyperlink formatting for p2:string parameter in JMX monitoring docs	AI-Agent/Cannot-Fix				1 1 2
#11851 · Sine-D opened 3w ago							
☐	☒	Fix stale monetization docs	Type/Improvement				2
#11847 · dgilling opened 3w ago							
☐	☒	Opensearch Artifacts link it's broken	AI-Agent/Fixed	Type/Bug			2
#11826 · jpvadel opened 1mo ago							
☐	☒	[APIM 4.7.0] Multi Tenant Authenticator and /select-tenant/ endpoint missing from Docker image	Type/Bug				
#11822 · Lengkak-Hok opened 1mo ago							
☐	☒	Remove outdated Enabling Role-Based Access Control Using XACML documentation.	AI-Agent/Fixed				5
#11812 · tharindu1st opened 1mo ago							
☐	☒	Outdated OSGi website link in Admin Services documentation	AI-Agent/Fixed	Type/Bug			2
#11790 · MenuraAndrahennedi opened on Aug 17							
☐	☒	Remove api governance cli tool documentation	AI-Agent/Fixed	Type/Bug			4
#11720 · tharikaGitHub opened on Aug 4							



Code Quality

- **Repository Rules:** Code repositories enforce specific rules for coding agents to prevent the introduction of bugs or security flaws.
- **Mandatory Human Oversight:** Coding agents are utilized strictly for assistance; we do not rely on them 100%, and manual PR reviews remain mandatory.
- **Automated AI Reviews:** We integrated CodeRabbit to provide an additional layer of automated, AI-driven Pull Request reviews.

Note: Different teams could use the coding assistants in different capacity



Customer Reliability Engineering

Novera - Agentic Support Experience

Problem → Resolution Latency & Effort

- Substantial time and friction required for customers to resolve technical issues through traditional support tickets.

Approach → Agentic Self-Service Assistant

- Conversational AI chat integrated with autonomous tool-calling for Knowledge Base and live Web search.

Target Outcome → 30% Support Deflection

- Deflect nearly one-third of incoming customer support queries through instant self-service resolution.

Current Status → ~5% Cases Deflected

- Initial rollout active with early adoption successfully eliminating ~5% of ticket volume, scaling towards 30%.



CRE Agentic Resolution Experience

Problem → Support Resolution Effort

- Time and effort it takes for a support engineer to resolve a support issue.

Approach → Claude Code Agent

- Locally running Claude Code based Agent with several MCPs and skills.

Target Outcome → 10 Days Case MTTR

- Bring down case MTTR to 10 days.

Current Status → MTTR Decreasing

- MTTR shows a visible downward trend.



Other Teams

Other Initiatives

- **[SRE]** Smart incident alerting agent
- **[WSO2 Digital]** Novera assistant for WSO2 employees
- **[Marketing]** Ideal customer profile for marketing team

2 min

Give me a summary of leave I have taken. Mask the number of days because I want to use this in a slide. I don't want the audience to know the number of leaves I have taken.

Novera App 2 min • Edited



Here's a summary of your leave taken in 2026 (number of days hidden):

1. Feb 19 — Half day (afternoon) | Casual | Approved
2. Feb 20 — Single day | Casual | Approved
3. Apr 2 – Apr 10 — Multiple days | Casual | Approved
4. May 27 — Single day | Casual | Approved
5. Jul 24 — Half day (morning) | Casual | Approved
6. Jul 30 — Half day (afternoon) | Casual | Approved
7. Jul 31 — Single day | Casual | Approved
8. Aug 14 — Single day | Casual | Approved
9. Aug 20 – Aug 27 — Multiple days | Casual | Approved

All leave records are approved casual leave. The number of days has been masked as requested.
Good luck with your presentation! 😊



Execution

- Formed a virtual team which consists of engineers representing major product teams.
- Identified the main problems.
- Weekly / Bi-weekly milestones achieved.
- Developer experience was the priority.
- Used the engineers in the virtual team to champion the practices in their respective teams.
 - Incorporated the feedback to improve the solutions
- After considerable adoption, focused on measuring the impact too.



Outcome

- **Consistent Output:** AI output quality is no longer dependent on individual engineer prompting skills.
- **Accelerated Maintenance:** Porting updates across our 40+ product versions is significantly faster and more consistent.
- **Enhanced Security:** Proactive codebase scanning has successfully identified and neutralized numerous security issues.
- **Operational Efficiency:** Routine tasks, like documentation fixes, are now resolved with minimal manual effort.



Cost, Governance, and Insights

No cost surprises.

Predictable Costs

We selected the Claude Team plan over the Enterprise plan to maintain full control over expenses with zero cost surprises.

Rules

Rules defined at repository level guarantees same quality for the code generated by every engineer with AI assistance.

WSO2 AI Gateway

WSO2 AI Gateway can provide better insights than the Claude Team plan.



WSO2 AI Gateway

Insights

Explore usage analytics and traffic insights powered by Moesif.

< Analytics / Overview / LLM ▾

Production Day Last Month No Refresh C

API Search or enter a value

Unique Consumers

▼ 25.00%

6

Total Requests

▼ 43.15%

3,569

Average Error Rate

▲ 152.63%

7.57 %

Token Usage

▼ 46.73%


3.3M

Estimated Cost


▼ 18.26%

491.173


Cost Trend per Provider




Traffic Share by Provider and Model



Request Intensity (Day vs Hour)



Latency Trend



Next Steps

- Using Agent Platform to monitor the agents
- Using AI Gateway for better governance and insights
- Connecting agents to deliver end to end agentic experience
 - Bug fixing agent to create the update as well
- Spec driven development



Recommendations

Define your objectives

- Highest impact generating
- Achievable within 6 months
- Focus on improving quality while achieving efficiency

Strategize delivery and key results

- Identify key results for each of the above objectives
- Decide on the initial technologies (they will evolve)

Outcome lead execution

- Create small teams to deliver the above outcomes
- 2 weeks sprints and push to production at end of each sprint





SEPT 22 - 24, 2026 | NAIROBI, KENYA

Thank You!



Amila Mahaarachchi

Senior Vice President &
Head of Engineering