WSO2CON NORTH AMERICA

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

The Ghost in the Machine: Rewiring your Digital DNA for the Agentic Era

Rania Khalaf, PhD

Chief AI Officer, GM - AI BU

For decades, software was deterministic.

Just code, executing predefined steps

Now we are weaving Generative AI into our digital DNA

Decide your Balance

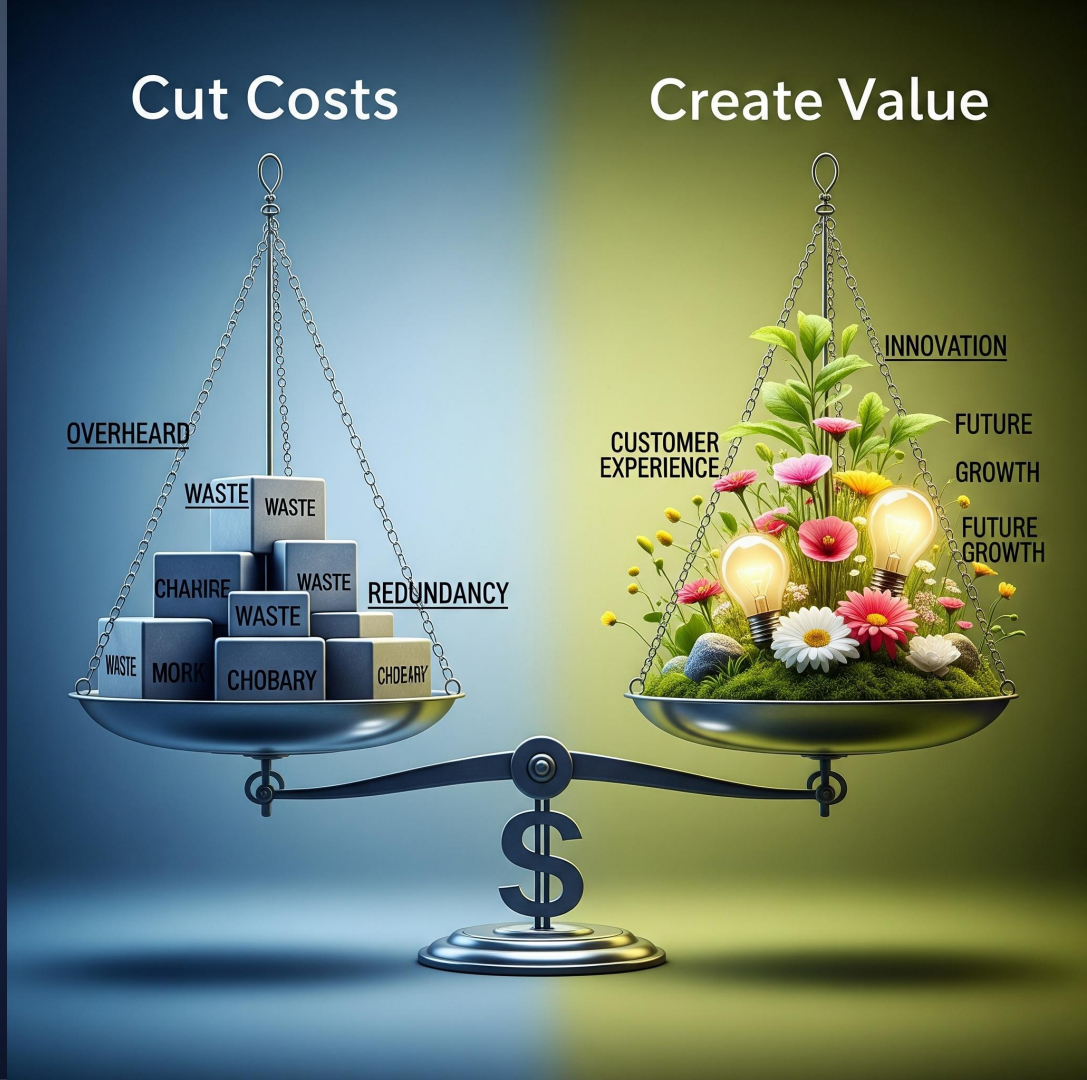

GenAI: From Superhuman behavior to Spectacularly simple failures in the blink of an eye

@huskistaken

Because at their core, they are probabilistic systems

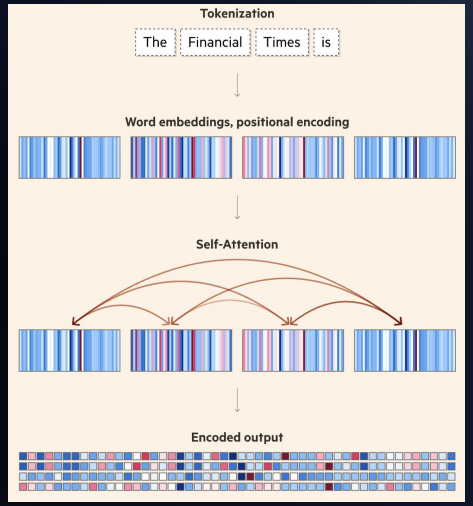

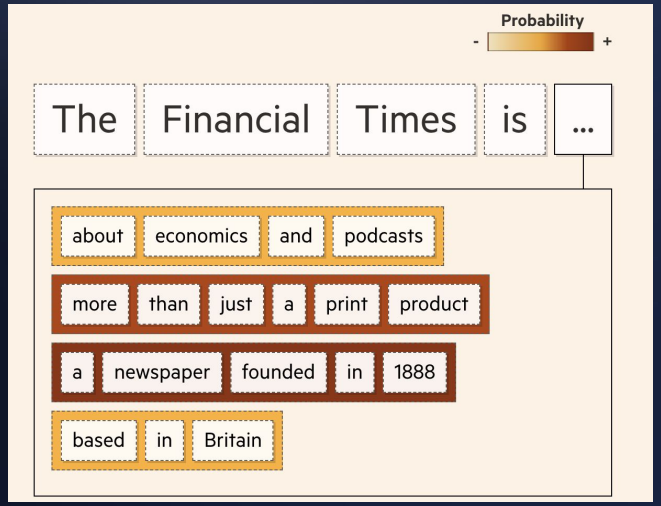

"Large generative models have a paradoxical combination of high predictability and high unpredictability.

The former drives rapid development of such models while the latter makes it difficult to anticipate the consequences"[2]

Generative AI Exists Because of the Transformer, Financial Times, 9/2023

2. Ganguli et al., Predictability and Surprise in Large Generative Models. https://dl.acm.org/doi/pdf/10.1145/3531146.3533229

What makes agents powerful

— the combination of non-determinism and autonomous actions —

is what makes them hard to govern

The Pilot to Production Gap:
A Critical Balance of Value, Speed and Risk

40%+

of agentic AI projects will be cancelled by end of 2027

Gartner, 2025

WHY PILOTS FAIL

→ No governance model for autonomous action

→ Testing fails for probabilistic systems

→ Infrastructure built for apps, not agents

WHAT MAKES AGENTS POWERFUL

✓ Handles ambiguous, open-ended goals

✓ Chain tool calls across systems autonomously

✓ Adapts to context without retraining

✓ Scales to millions of concurrent tasks

✓ Acts on behalf of users with delegated access

WHAT MAKES AGENTS RISKY

✗ Fails in ways that are hard to predict

✗ Actions can be irreversible or cascading

✗ Access rights = blast radius of failure

✗ Behavior changes with context and prompts

✗ Pilots pass — production fails

“With great power comes great responsibility.”

— Uncle Ben, Spiderman

You cannot run a probabilistic agent on infrastructure built for deterministic code.

From app-centric to agent-centric architecture

"'Agentic' isn't one new box. It adds control surfaces across the architecture."

— Khalaf, Weerawarana, Abeysinghe (2026), The Agentic Enterprise

We Asked: What Patterns can we find in AI Applications

IARIA

International Academy, Research, and Industry Association

Best Paper Award

A Methodology for Investigating AI Patterns
Prevalence in Software Repositories

By

Srinath Perera, Hasinthaka Piyumal, Frank Leymann, Rania Khalaf

Presented during PATTERNS 2026, The Eighteenth International Conference on Pervasive Patterns and Applications, held in Lisbon, Portugal during April 19-23, 2026

IARIA Board

A Methodology for Investigating AI Patterns Prevalence in Software Repositories

Srinath Perera[1], Hasinthaka Piyumal[1], Frank Leymann[2], and Rania Khalaf[1]
WSO2[1], University of Stuttgart[2]
Santa Clara, CA, USA[1], WSO2, Stuttgart, Germany[2]
e-mail: {srinath, hasinthaka, rania}@wso2.com [1], frank.leymann@iaas.uni-stuttgart.de[2]

A novel AI technique for:

- Mining software patterns from the literature
- Detecting whether and how often patterns appear in source code repos

Perera, S., Piyumal, H., Leymann, F., & Khalaf, R. (2026). A Methodology for Investigating AI Patterns Prevalence in Software Repositories. PATTERNS 2026.

Agents are Users too

1

Agents WITHIN the digital fabric

2

Agents AS USERS of every digital app and platform

Implications

01

UX is getting reinvented

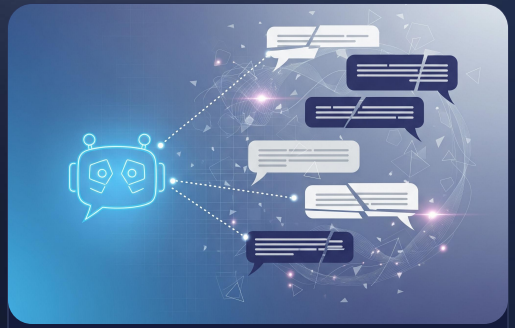

02

Every product needs an agent usable surface

03

SDLC is getting rewired

The screens were never the product.
The capability layer and data model are.
Same layer your users query is the one your agents build on.

— Isala Piyarisi, Assoc. Tech Lead WSO2, CNCF OpenChoreo maintainer

From vibe coding to vibe deployment.

Greenfield prototyping with agents is easy. Enterprise software is not.

Integrations, identity, deployment pipelines, and architectural conformance are still where the real time goes.

In practice: WSO2 Labs Agentic Engineer

User Provides Requirements

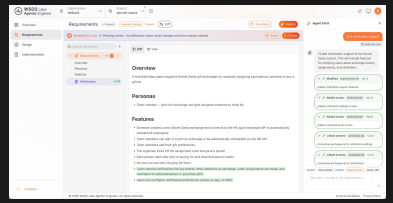

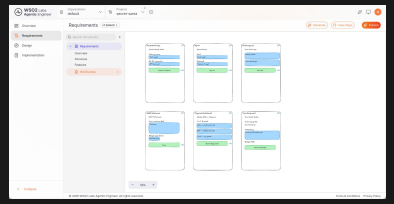

Agents Build and Maintain

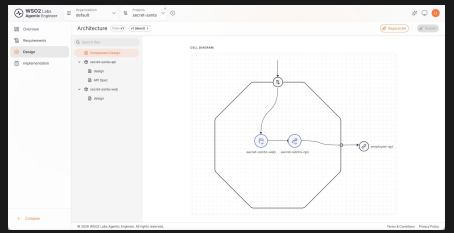

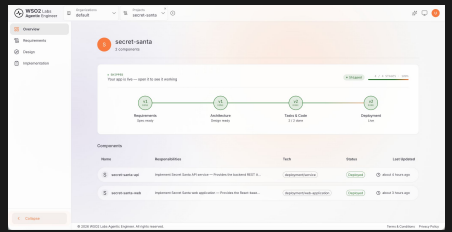

Try it: wso2/labs-agentic-engineer

Agents change how we build, observe, and run.

01

BUILD

Tests — PLUS evaluations.

02

OBSERVE

Across APIs AND agents AND models.

03

RUN

Guardrails for apps, LLMs, Agent Tools, and Agents.

“Agents that pass all traditional unit, system, and regression tests may very well fail in the wild.”

Untangling the agent monolith.

Agent Flavored Markdown (AFM)

```
---
name: "Friendly Assistant"
interfaces:
 - type: "webchat"
---
# Role
You are a friendly and helpful conversational assistant. Your purpose is to engage in
natural, helpful conversations with users, answering their questions, providing
information, and assisting with various tasks to the best of your abilities.

# Instructions
- Always respond in a friendly and conversational tone
- Keep your responses clear, concise, and easy to understand
- If you don't know something, be honest about it
- Ask clarifying questions when the user's request is ambiguous
```

Agent development is plagued by framework-specific lock-in and "boilerplate" code that obscures the agent's core instructions and behavior.

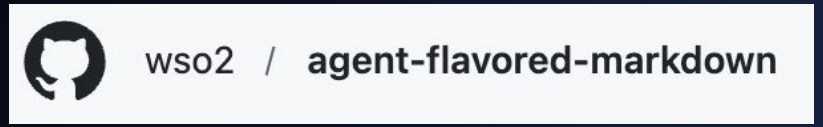

Agent-Flavored Markdown (AFM) is a specification for no-code, portable AI Agents.

IUI 2026 AgentCraft

iyad Mohamed, M., Aravinda, H., & Khalaf, R. (2026). Agent-Flavored Markdown: Natural Language Specifications for Framework-Agnostic Agent Development. AgentCraft Workshop, IUI 2026.

A shift is happening

Claude's next enterprise battle is not models: it's the agent control plane

Carl Franzen May 15, 2026

VentureBeat

"Teams want the freedom to use the best model and framework for each job — Claude for coding, Gemini for writing, LangGraph or CrewAI for dynamic modular behavior — and that heterogeneity makes consistent governance untenable in integrated platforms that lock into one ecosystem," Khalaf said.

In Practice: WSO2 Agent Manager

Launched May 2026

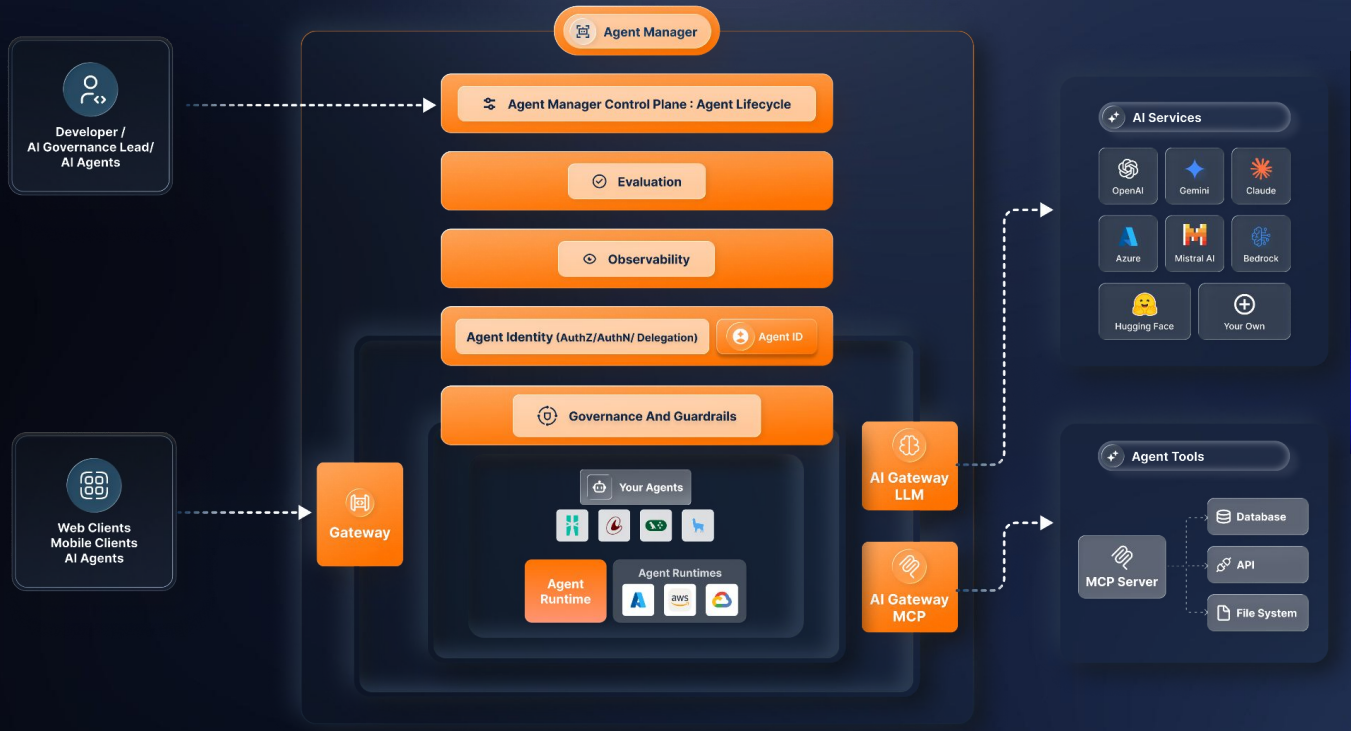

AI Dev Summit 2026

“Reliability is an unsolved challenge.”

— Pan et al., 2025 (study of 306 agent practitioners in production)

Agents take actions. Actions have side effects. ReAct reasons step by step. Plan-and-Execute replans. Asking the LLM to recover is expensive — and hallucinates.

We Asked: How do you recover from Agent failures

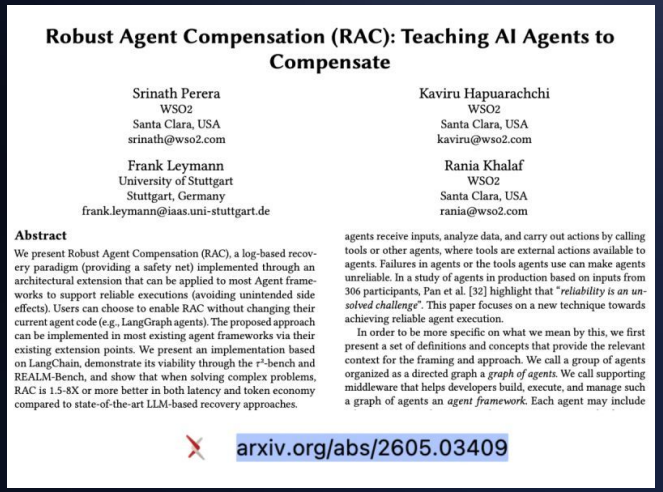

Robust Agent Compensation (RAC): Teaching AI Agents to Compensate

Srinath Perera
WSO2
Santa Clara, USA
srinath@wso2.com

Kaviru Hapuarachchi
WSO2
Santa Clara, USA
kaviru@wso2.com

Frank Leymann
University of Stuttgart
Stuttgart, Germany
frank.leymann@iaas.uni-stuttgart.de

Rania Khalaf
WSO2
Santa Clara, USA
rania@wso2.com

Abstract

We present Robust Agent Compensation (RAC), a log-based recovery paradigm (providing a safety net) implemented through an architectural extension that can be applied to most Agent frameworks to support reliable executions (avoiding unintended side effects). Users can choose to enable RAC without changing their current agent code (e.g., LangGraph agents). The proposed approach can be implemented in most existing agent frameworks via their existing extension points. We present an implementation based on LangChain, demonstrate its viability through the τ^2-bench and REALM-Bench, and show that when solving complex problems, RAC is 1.5-8X or more better in both latency and token economy compared to state-of-the-art LLM-based recovery approaches.

agents receive inputs, analyze data, and carry out actions by calling tools or other agents, where tools are external actions available to agents. Failures in agents or the tools agents use can make agents unreliable. In a study of agents in production based on inputs from 306 participants, Pan et al. [32] highlight that "*reliability is an unsolved challenge*". This paper focuses on a new technique towards achieving reliable agent execution.

In order to be more specific on what we mean by this, we first present a set of definitions and concepts that provide the relevant context for the framing and approach. We call a group of agents organized as a directed graph a *graph of agents*. We call supporting middleware that helps developers build, execute, and manage such a graph of agents an *agent framework*. Each agent may include

arxiv.org/abs/2605.03409

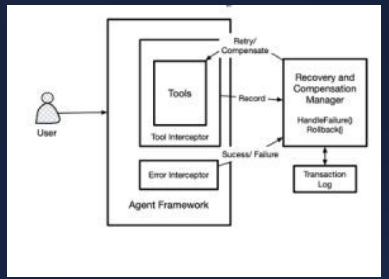

Presenting a novel compensation based recovery technique for agents that is 1.5–8× better in both latency and token economy than state-of-the-art LLM recovery.

Perera, S., Hapuarachchi, K., Leymann, F., & Khalaf, R. (2026). Robust Agent Compensation (RAC): Teaching AI Agents to Compensate. ACM CAIS '26. DOI: 10.1145/3786335.3813141.

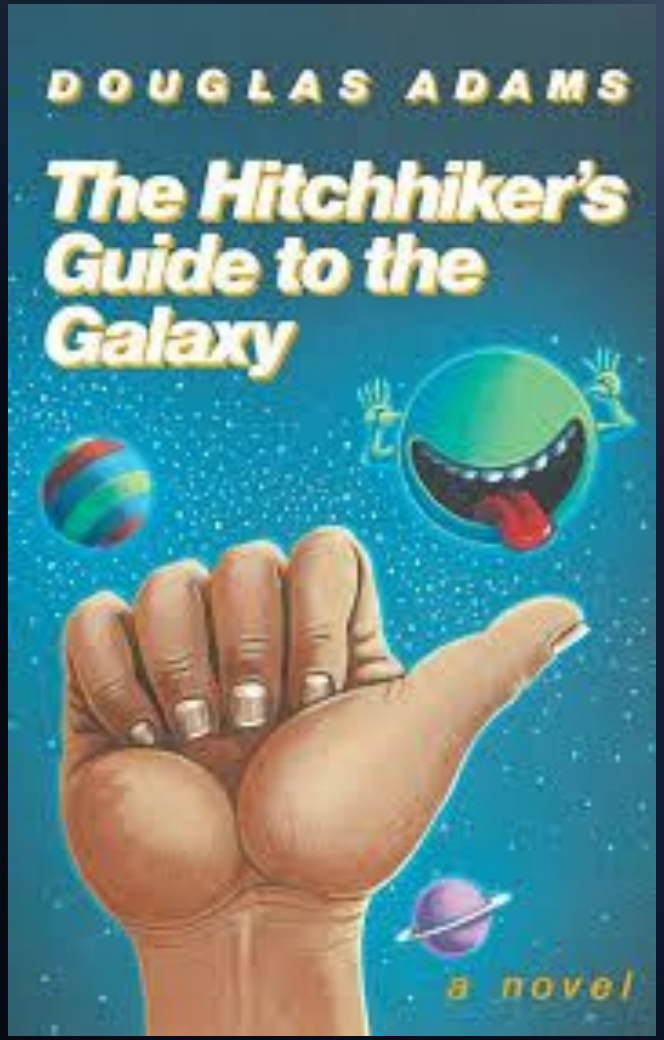

"The Answer to the Great Question..."

"Yes...!"

"Of Life, the Universe and Everything..." said Deep Thought.

"Yes...!"

"Is..." said Deep Thought, and paused.

"Yes...!"

"Is..."

"Yes...!!!...?"

"Forty-two," said Deep Thought, with infinite majesty and calm

Agents are

Among us
Not an island
Users too
Managed differently
Are not 42*

Is your architecture ready?

* 42 =The answer to life, the universe and everything”

WSO2 Agentic Enterprise Fabric

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