Agentic Trust Requirements

March 2026

Lead Author: Damian Glover

DIF

Project Background & Goals

The Trusted AI Agents Working Group is keen to anchor its work in concrete use cases, in order to enable real-world solutions built on DIF specifications. To support these efforts, DIF commissioned a study to identify compelling Enterprise use cases for cross-boundary Agent-to-Agent communication, and gather Enterprise perspectives on what's needed to bring solutions into production.

Project Goals

1. Understand the conversations happening within Enterprises around Agentic AI

2. Identify current, emerging and potential Enterprise use cases where AI agents need to operate across trust boundaries

3. **Identify key identity and trust requirements**

Research methodology

Literature review to understand overall maturity & industry adoption of Agentic AI

3x industry-specific deep dives, including in-depth interviews with subject matter experts

DIF

Source: DIF interviews with industry & technology experts, October 2025 - January 2026

Enterprise AI Adoption: Current Reality vs. Future 'Agentic Web'

Current Enterprise Reality

AI-augmented Workflows & LLMs

In Production: LLM-based chatbots and AI-augmented workflows are deployed.

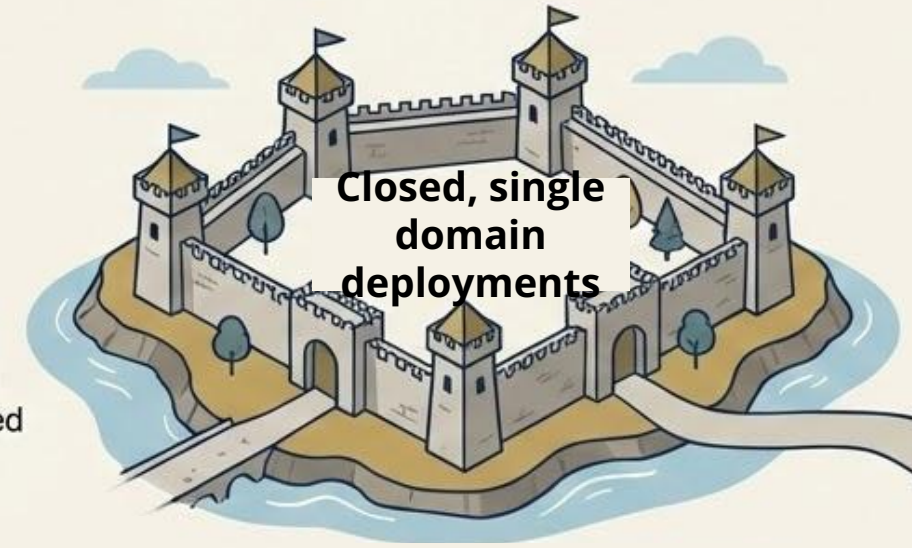

Platform ecosystems

Walled-Garden SaaS: Interaction via AI-augmented workflows within platforms like SAP.

Loosely coupled API orchestration

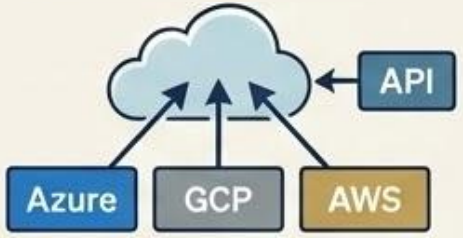

Cloud & API-based: Integration within cloud environments via APIs, not open protocols.

Future Outlook & Barriers

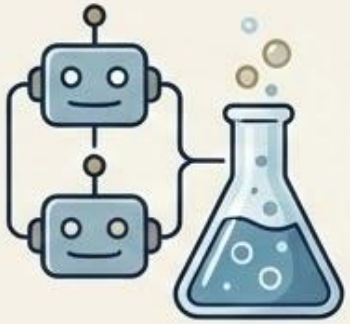

Autonomous Agents & Multi-Agent Systems

Experimental: Usage is still experimental; interactions confined within enterprise boundaries.

API vs. Open Protocols (e.g., A2A)

Open, distributed agent choreography

A LONG way off!: The fully realized "agentic web" is still ~7 to 10 years away

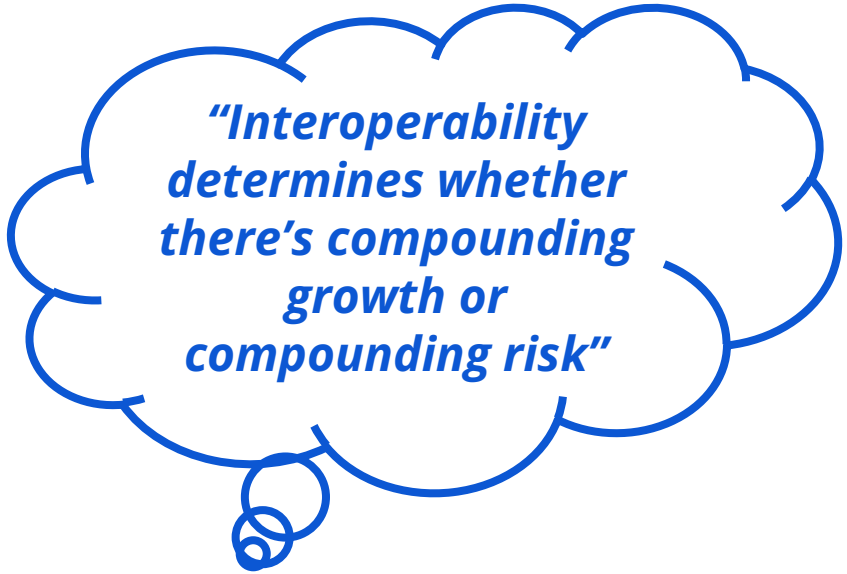

Source: Patchwork: Synching on AI Standards, 10 December 2025

DIF

Insurance sector - context

- Legacy tech and data silos
- High cost of claims processing & fraud
- Current agentic initiatives are internal RPA-style projects

Source: DIF interviews with insurance industry participants

Insurance sector - agentic timeline

Simple claims processing	Agentic commerce	End-to-end claims automation
Already happening, e.g. AI Jim (Lemonade), Insurance Co-Pilot (Allianz), Qover Also co-pilots to assist underwriting decisions	Within 1-3 years Key drivers: enable agentic product discovery, drive top-line growth	5 year program, major players will start planning in 2027 Key drivers: efficiency, fraud reduction

Source: DIF interviews with insurance industry participants

Claims Automation - use case overview

Description

- AI agents orchestrate the insurance claims lifecycle across multiple external participants — insurers, brokers, adjusters, repair shops, healthcare providers, and policyholders.

Why cross-boundary trust is needed

- Agents must autonomously extract, verify, and transfer key data (incident details, documents, photos, repair estimates, medical bills) between systems, ensuring policy compliance and regulatory transparency while reducing manual handling.

Feasibility considerations

- Selective disclosure: Only claim-relevant data (not full customer records) should be exchanged, to protect privacy
- Trust through verification: Agent actions must be logged, attested, and audited, ensuring accountability across parties

Source: DIF interviews with insurance industry participants

DIF

Claims Automation - example workflow (automotive)

Submits claim with verified credentials (Proof of policy, photos, police report)

Verifies credentials, checks policy coverage, determines claim type, routes to appropriate specialist agent

Requests estimates from network of verified repair shops

Each shop's agent has VCs proving licensing, insurance, reputation score

(for injury claims)

Submits bills with verifiable credentials (licensed provider, treatment codes)

Selective disclosure: Only relevant medical info shared

(for large claims):

Automatically notified if claim exceeds threshold

Verifies primary insurer's agent authority

Confirms coverage participation

Executes payment after all verifications complete

Creates immutable record for audit trail

Source: DIF Interviews with industry participants

DIF

The Trust Gap in Agent Interactions

Understanding where trust breaks down in the AI agent ecosystem reveals why current approaches fall short of enabling safe, scalable automation.

User ↔️ Agent

❌ No Strong Linkage

Weak or non-existent binding between the user and their agent

Agent ↔️ Business

❌ No Trust Layer

No cryptographic proof of authorization or identity

Consent & Authorization

❌ Missing Framework

No standardized consent or delegation mechanism

Additional Gaps for Claims Automation

Policy-bound agent authorisation (rather than broad, bearer-based access)

End to end auditability

Clear accountability & liability frameworks for agents acting on behalf of organizations or individuals

Industry-specific schemas

Source: DIF Hospitality & Travel SIG, December 2025, Industry Expert Interviews

DIF

Platforms & payments infra can bridge some of the gaps

01	**Core Systems** (Guidewire, Duck Creek, Applied Epic)	• 70% of insurers run on these platforms for policy admin, claims, billing • Users can enable agents to submit quotes for approval, endorse, cancel, and renew policies, etc
02	**AI/ML Platforms** (Reserv, Shift Technology, Tractable)	• Provide AI for specific insurance use cases (claims triage, fraud detection, damage assessment) • Already building agent-like functionality but without identity / interoperability layer.
03	**Payment / Agent Infrastructure** (Stripe, Plaid, fintechs, Open Finance rails)	• Stripe + others are building payment agent infrastructure • E.g. Stripe enables merchants and service providers to share product data with AI agents, use payment tokens & make machine payments (e.g. agents paying for API calls / data access)

Source: DIF Interviews with industry participants

DIF

Recommendations

- Enterprises: start by piloting verifiable credentials for agents (internal) and adopt attestation for critical agents so you can later interoperate. Define liability/playbooks.
- SaaS/platform providers: offer bridges (exportable credentials, signed audit trails, clear SLAs) to increase interoperability while preserving platform controls; participate in standards.
- **Standards bodies: prioritize delegation & selective disclosure**
- Researchers: focus on scalable revocation, attestation at ML/model level, sybil-resistant identity schemes, and human-interpretable provenance.
- Integrators / developers: design using capability-scoped tokens and fallback safe modes (human in loop) for cross-domain flows; prefer composable, auditable messages.
- Industry bodies & Regulators: develop data schemas, facilitate industry thinking on liability frameworks for agent actions

Source: DIF research

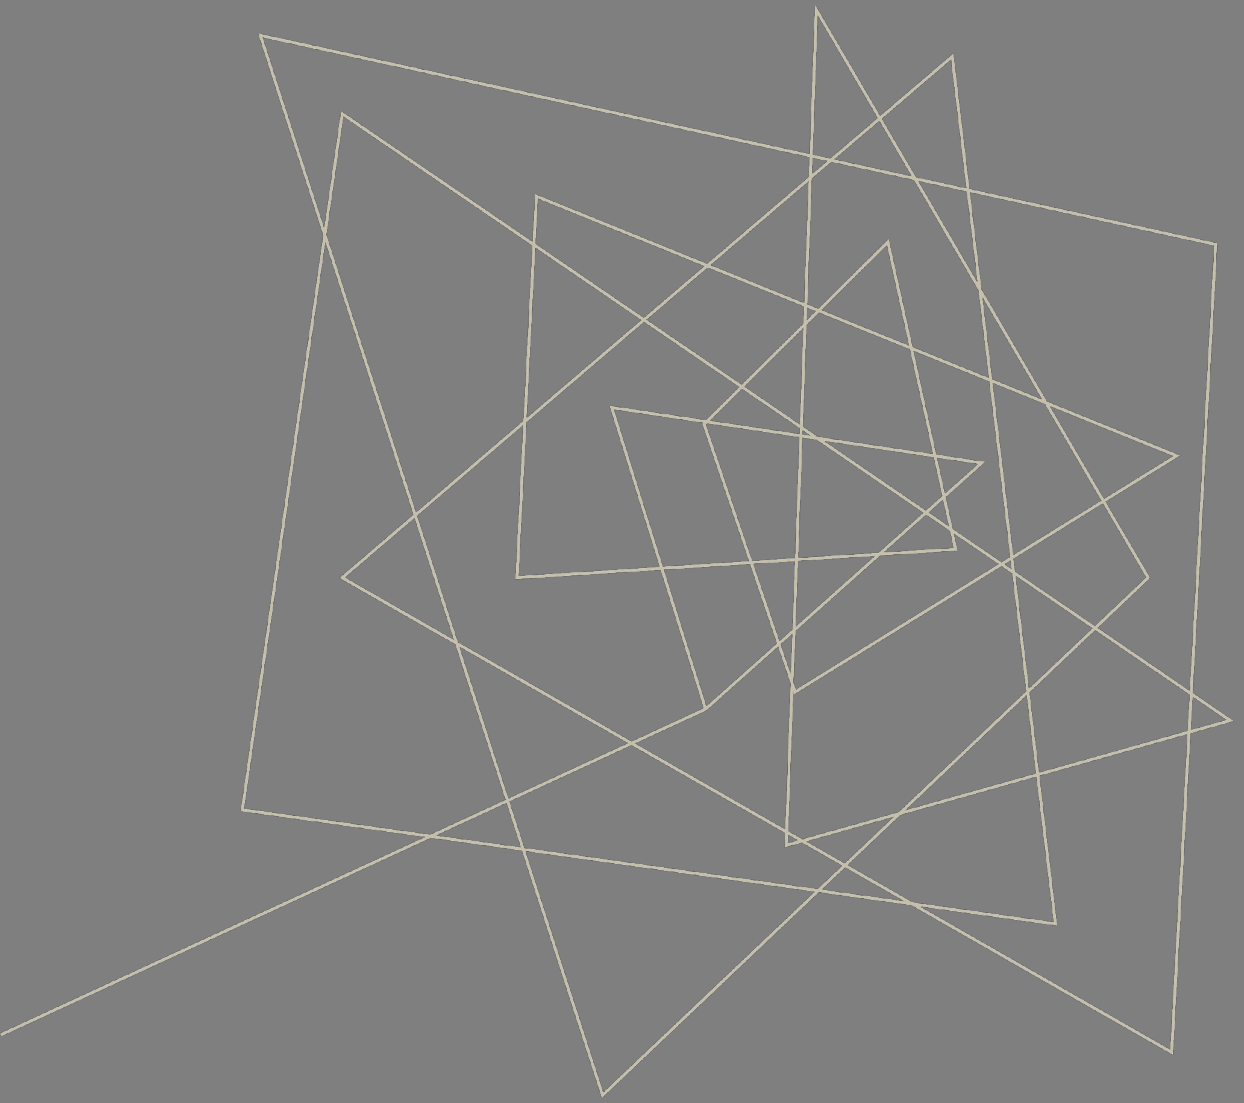

Thank you for your attention!