



# ARKOVA

— Issue Once, Verify Forever —

## The Gap in the Agentic IAM Discussion

Over the past year we have seen an explosion in organizations adopting AI and autonomous agents to streamline work and eliminate costs. From helping developers code, to managing an inbox, executives are finding ways to incorporate AI into their daily workflows. But with that adoption comes a growing risk, and while many companies are beginning to find ways to verify an agents identity, nobody is talking about the liability that comes from autonomous agents acting on regulated documents, credentials, and version control repositories. Arkova aims to tackle that issue

### The five levels of verified identity



#### Level 0 Point-in-time verification

Identity checks with no portable identity retained.

Today



#### Level 1 Bound identity

A verified human holds a portable, key-bound credential.



#### Level 2 Bound authorization

A verified identity signs a specific transaction.



#### Level 3 Bound delegation

A signed mandate authorizes an agent within defined scope.



#### Level 4 Bilateral verification

Both human and business are cryptographically verified.



#### Level 5 Networked verified trust

Any verified party can transact with any other.

The progression begins here →

## What is Arkova?

Document, credential, and licensing verification for the Agentic economy. According to a recent [Boomi](#) study, **86%** of enterprises have deployed AI agents.

Autonomous agents are being integrated into enterprise workflows to move beyond simple chat responses and actively execute multi-step business processes. Agents are extracting data from invoices, reviewing contracts, and translating multilingual files, connecting to existing enterprise data stores and enterprise identity providers, and automating employee onboarding.

Arkova is a middleware that extracts a documents unique metadata, without ever taking custody of the document itself, and anchors that metadata to an immutable cryptographic ledger verifiable anywhere, by anyone, at anytime. Even if Arkova disappears, the verification remains immutable and verifiable

## Why is Document Verification Important to the Agentic OS?


Despite **86%** of organizations surveyed reporting that they have deployed agents, only **34%** of those surveyed say they trust the actions their AI agents are taking. As agents begin to act on behalf of organizations at an increasing rate, verifying the actions of agents is going to become a job unto itself, erasing possible gains the agents themselves may provide. Poor document and credential management is already a multi-trillion dollar issue, globally. A 2024 [Deloitte and Docusign study](#) found that nearly **\$2 Trillion USD** in global economic value is destroyed annually because of poor agreement management. While an [Adobe](#) survey found **48%** of professionals struggle to locate documents quickly, another **48%** of organizations experience deteriorated customer relationships due to agreement delays.

**1. Agreement management ineffectiveness drives significant destruction in global economic value.** Current agreement management processes are often inefficient, leading to a nearly £1.4 trillion loss in annual global economic value from lost productivity and revenue opportunities.

”

“Our research shows that companies with disconnected agreement management workflows spend an extra 18% of time working on agreements. These inefficiencies can have a material impact on a company's productivity, employee morale, bottom line, and long-term business outcomes.” – Jonathan Jones, Managing Director, Deloitte Tax LLP

Things become even more staggering when you look beyond contracts and corporate documentation, and focus on credential fraud. According to the [Association of Certified Fraud Examiners](#), **88%** of individuals committing occupational fraud **passed** pre-employment background checks with no red flags. **4% had a prior fraud conviction** that wasn't caught or didn't block the hire.



The issues we are seeing with document and credential management are only going to be proliferated as agents start replacing humans. Without guardrails, your organization is at risk of becoming a cautionary case study at HBS.

## When do Organizations Need to Start Preparing?

**Now.**

Enforceable global regulations demand it. While the EU AI Act recently adjusted its strict compliance deadlines for standalone high-risk systems to December 2, 2027, and embedded systems to August 2028, other regulatory pieces like the Colorado AI Act (effective now) and initial phases of the EU AI Act are already active. [Research indicates](#) that **72%** of organizations have already faced breaches linked to non-human identities, making proactive compliance an operational necessity rather than a legal box to check.

| Regulation/Framework            | Enforcement Date | Core IAM Impact                                                                                        |
|---------------------------------|------------------|--------------------------------------------------------------------------------------------------------|
| Colorado AI Act                 | February 2026    | Mandates strict documentation of AI decision-making processes and algorithmic risk assessments.        |
| EU AI Act (General Provisions)  | August 2, 2026   | Governance, transparency, and penalty provisions go into active effect.                                |
| EU AI Act (Annex III High-Risk) | December 2, 2027 | Deadline for standalone high-risk applications (e.g., biometrics, employment screening tools).         |
| EU AI Act (Annex I High-Risk)   | August 2, 2028   | Applies to high-risk AI embedded in traditionally regulated physical products (e.g., medical devices). |

The potential legal exposure for enterprises utilizing autonomous AI agents on regulated documents and version control repositories is severe. As agent architectures pivot from content generation to independent conduct, deploying organizations bear [primary liability](#); courts explicitly reject the defense that "the AI acted autonomously".

The independently verifiable proof package Arkova delivers for anchored documents can help organizations protect themselves in case of an external audit.

## Who is Currently Tackling this Issue?

No single centralized global authority currently verifies whether the documents and instructions processed by autonomous AI agents are completely genuine. There are [5 layers](#) to verified identity. Not one of them deals with ensuring the **internal documents** your agent is working on is accurate.

While automated platforms like Scoreplex and Wisedocs parse legal filings, invoices, and medical documents by cross-referencing them against official public or corporate registries, they act specifically against official public or corporate registries. That’s great if you want to verify an SEC Edgar filing your agent acted on was accurate, but what about if you want to verify a private, internal contract that your agent is working on? Or internal financial records? Right now there is a real gap in the market.

For example, you’re a paralegal at a law firm, a client is claiming your firm breached the contract you both had, they have attached the contract and highlighted the supposed breach. Your agent, having read your email, begins to search your internal records to verify the claims your client made. A few seconds later your agent replies to the client that their version of the contract doesn’t align with theirs. Who’s right? At best you spend a few hours reviewing things, attending meetings with your

clients representation, and things resolve quickly. At worst you're dragged into a months long, expensive dispute. With Arkova we integrate with e-signature platforms like DocuSign to extract and anchor that contracts unique metadata as soon as it is executed. In the same scenario as the one above, your agent would've found the disputed contract, delivered you a link to your organizations record of it (complete with proof package) as well as that anchors public ID. You would then send that proof package + public ID + link to [search.arkova.ai](https://search.arkova.ai) + the executed version of that contract to your client.

## Where are Agents Acting? Why is this a Risk?


***Everywhere.***

And we're starting to see issues. [NHI sprawl](#) has become an alarming security risk. Existing IAM frameworks, including widely used protocols, such as OAuth 2.0, OpenID Connect (OIDC), and Security Assertion Markup Language (SAML), were designed for a more deterministic digital era. They presume predictable application behavior and a single authenticated principal, human, or static machine identity. Agentic AI violates the assumptions these frameworks make. While Digital Identifiers (DIDs) and Verified Credentials (VCs) - both of which Arkova utilizes - are a start, it's not broad enough. Organizations need to be proactive in addressing document verification issues before an incident occurs that puts them on the front page of the Wall Street Journal. While others are working to address the broader picture, Arkova is working to become that verification layer for the Agentic OS.

## What You Can Do Today

***You do not need to wait for regulatory deadlines to close this gap.***

Verify a record yourself. Every document Arkova anchors gets a public ID and a proof package that anyone can check at [search.arkova.ai](https://search.arkova.ai), without an account and without access to the document itself. More than 3.4 million records are already anchored and independently verifiable.



If you run agents on regulated documents: bring us your highest risk workflow, whether that is executed contracts, professional credentials, or a version controlled repository, and we will show you what the audit trail looks like when your agent acts on it. Connect Docusign or Google Drive and your next executed agreement is anchored automatically, with no change to how your team works.

If you invest in the agentic infrastructure stack: the identity layer is being solved. The document layer underneath it is not. We are happy to walk you through the architecture, the regulatory timeline, and where Arkova fits.

**Lets Talk:** founders@arkova.ai

***Issue Once. Verify Forever.***