

NASS 2026 Summer Conference

Agentic AI and Public Records: Building a Verifiable Trail for the States

A policy framework for accountable
AI in state filing processes

A new kind of participant has entered state filing workflows. It is not a new agency or a wave of hires, but AI agents, software that acts on a filer's behalf, pulling data from past filings, completing forms, and routing work on its own. For Secretaries of State, this shifts the work before a public record onto the constituent's side of the desk. The opportunity is real, since AI agents can cut the friction behind missed deadlines and rejected filings. The risk is equal. If a dispute surfaces years later and no one can reconstruct what an agent did, accountability erodes. This paper makes one core argument. States should begin moving toward a defensible, machine-readable audit trail that captures an agent's inputs, outputs, and human decision points, without reaching into proprietary model internals.

From hand-keyed inputs to agent-mediated workflows

Most public-facing state processes run on a manual loop. A person keys in data, signs, and pays, and staff catch errors and stand behind those actions if a court or auditor later asks. In the private sector, AI agents already draft, scan, and flag risks before a manager opens a file. People still hold final authority, but the agent sets the starting point.

Government offices face the same pressure to automate, but under heavier public scrutiny. Americans spend an estimated 6.5 billion hours a year on tax compliance alone.¹ The burden falls hardest on those with the fewest resources. 16% of U.S. adults have no home internet beyond a smartphone, more among lower-income residents.² For these residents and for small business owners filing on their own, an unclear form turns a routine task into a missed deadline or rejected filing. Efficiency is the easy part. Accountability is harder. If a dispute surfaces three years from now and no one can reconstruct what an agent did, the problem has not vanished, it has moved downstream. **Constituents have a basic right to know who decided what, and on the basis of which facts.**

1 6.5 Billion Hours, \$260 Billion: What Tax Complexity Costs Americans (2023), National Taxpayers Union Foundation.

2 Mobile Fact Sheet (2025), Pew Research Center.



Hard lessons from private sector agreement management

Modern agreement management offers a useful blueprint for governing automation. Studies estimate that fragmented agreement workflows, where information is scattered across disconnected tools, cost the global economy on the order of \$2 trillion a year in delays and administrative waste.³ Teams using agent-based workflows inside connected, end-to-end platforms report nearly 30 percent higher return on investment. Across organizations, 72 percent report improved agreement accuracy.⁴ Broader analysis from EY suggests that generative and agentic AI could unlock trillions of dollars in global productivity, provided it is deployed with proper guardrails.⁵

States do not need to copy commercial software wholesale, but the underlying rule applies. When data moves through a continuous, governed loop from intake to archive, human error drops. When it gets trapped in scattered email chains and stray PDFs, institutional memory erodes. For Secretaries of State, this maps onto a familiar mandate, modernizing government services and improving efficiency without surrendering accountability. A governed, end-to-end data flow advances both goals at once.

The technology: agents, systems of record, and MCP

The underlying technology is maturing quickly. The Model Context Protocol (MCP), open-sourced in late 2024, sets a secure, universal standard for AI models to talk directly to external data sources, replacing a fragile custom API per application with a single interface.⁶ Adoption is accelerating sharply. In the year after MCP's release, available MCP tools grew from roughly 5,000 to 177,000, and downloads climbed from 80,000 to 14 million, about half from the United States.⁷ In the near future, a citizen could realistically rely on an AI agent to query a state's corporate registry, pull historical data, and file a compliance form straight from the software that a business already uses — all without ever logging into the state portal.

For state registries, the upside is cleaner data and lower call volumes. But if an agent acts directly against a system of record, that system has to keep a reliable log. Auditors and courts are not interested in how a neural network works. They want to know what information entered the system, how it changed, and which person finalized the transaction. **A durable audit trail is a legal necessity.**

3 Deloitte and Docusign Research Uncovers Costly Problems in the Agreement Process (2024), Deloitte and Docusign.

4 Capitalizing on AI: How Automated Agreement Workflows Drive ROI (2026), Deloitte and Docusign.

5 The Economic Impact of Generative AI: A New Productivity Frontier (2024), EY.

6 Introducing the Model Context Protocol (2024), Anthropic.

7 How Are AI Agents Used? Evidence from 177,000 AI Agent Tools (2026), AI Security Institute.



Anatomy of a usable agent audit trail

What could a defensible agent audit trail look like? Embedding it in the model is not impossible, but external, independently validated custody is the architecture that actually works. It can be designed as a separate, machine-readable file kept alongside the filing, readable in plain language to a judge or auditor. The trail is scoped to inputs, outputs, and human decisions, while model internals and proprietary weights stay private. That keeps the requirement achievable for vendors and state CIOs and leaves room for different technical implementations to evolve over time. To stay legally useful, the trail could capture five layers of context (Figure 1):




Figure 1: The five layers of a defensible agent audit trail, from the rules the agent followed to the final record written to the state registry.

Picture a business owner who logs in one morning to find that an AI agent has swapped her company's registered agent and officers, without her sign-off. With a real audit trail, staff can pull up what the agent relied on, what it filed, and who approved it. Without one, they reconstruct the whole thing from memory while a legitimate business waits. International frameworks point the same way. The OECD AI Principles⁸ and the NIST AI Risk Management Framework⁹ both treat transparency and traceability as non-negotiable. State archivists can treat agent logs as a distinct class of record with their own retention schedules, not an IT afterthought.

8 OECD AI Principles (2019, updated 2024), Organisation for Economic Co-operation and Development.

9 AI Risk Management Framework (AI RMF 1.0) (2023), National Institute of Standards and Technology.

Practical action items for state leaders

Elected officials remain legally accountable for their agencies. AI does not change that. It adds a layer of automated processing on top of it. The early wins come from automating repetitive back-office work, not from replacing human judgment on high-stakes decisions like corporate certifications or fund routing. Five steps can help states get ahead of the shift:

- 1. Set intake ground rules.** State plainly that AI-assisted submissions are legal, but the human filer keeps full legal liability for the content.
- 2. Mandate tangible evidence.** For high-risk or fee-bearing filings, require a tamper-evident log of the agent's inputs and human interventions. You need not capture every mouse click, only enough to reconstruct the transaction in an audit. These expectations can be built into future procurements and vendor RFPs so that tools evolve toward them over time.
- 3. Adopt established frameworks.** Do not build state AI policy from scratch. Write the NIST AI Risk Management Framework into procurement contracts and vendor RFPs.
- 4. Target low-risk bottlenecks.** Apply AI where risk is low and operational payoff is high, such as routing paperwork by geography or handling status queries.
- 5. Lock down the log lifecycle.** Treat agent activity logs with the same security as primary records, with clear retention windows, redaction protocols for records requests, and tamper protection.

Conclusion

AI agents are reshaping state registries from the outside in, driven by tools businesses already use. Secretaries of State do not need to predict technology's next move to protect their systems today. A durable audit trail is not red tape. It is a legal necessity. By building toward that standard as tools mature, Secretaries of State can safeguard public records and preserve civic trust.



About DocuSign

DocuSign brings agreements to life. Nearly 1.9 million customers and more than a billion people in over 180 countries use DocuSign solutions to accelerate the process of doing business and simplify people's lives. With intelligent agreement management, DocuSign unleashes business-critical data that is trapped inside of documents. Until now, these were disconnected from business systems of record, costing businesses time, money, and opportunity. Using the DocuSign AI-native IAM platform, companies can create, commit, and manage agreements with solutions created by the #1 company in e-signature and contract lifecycle management (CLM). For more information visit <http://www.docuSign.com>.

DocuSign, Inc.
221 Main Street, Suite 1550
San Francisco, CA 94105
docuSign.com

Jacob Dale
Government Affairs Manager
jacob.dale@docuSign.com