



## SEVEN DAYS IN ORGANIZATIONAL INTELLIGENCE

How organizations capture, govern and reuse what they know.

# The National Archive Has Decided What AI Data May Be Deleted

The National Archives and Records Administration told federal records officers that the definition of a federal record reaches the inputs, outputs, data and audit trails involved in the use of AI, and that none of it may be disposed of without an approved schedule.

## 02 IN THIS ISSUE

# Schedules, Tickets, Polished Drafts

In a cluster analysis of hundreds of articles, about hundreds of organizations, these are the key developments and themes emerging.

Powered by [Hawkeye](#); analyzed by Claude, checked with ChatGPT, and human editorial oversight.

An organization's memory gets decided twice: once by what anyone bothers to type up, and once by what it is required to keep. This week a cloud provider went looking through resolved support tickets for the first, and the National Archives ruled on the second.

COVER IMAGE: CARD CATALOGS IN THE BASEMENT OF THE LIBRARY OF CONGRESS, WASHINGTON DC - WIKIMEDIA COMMONS (CC BY-SA)

1

NEWS BRIEFS

WHAT GOES UN-CHECKED

## Schedules, Tickets, Polished Drafts

A study on polish, a capture pattern, a legal graph, and a directory entry.

2

FORESIGHT

## One Good Bet, One Wild Card

Two scenarios that could unfold in this market.

3

COVER STORY

WHAT YOU MUST KEEP

## The National Archives Says AI Material Is a Federal Record and Cannot Be Deleted on Anyone's Say-So

An official memo last Friday and the archivists who have been arguing over it.

4

FEATURE 1

THE KNOWLEDGE YOU ALREADY PAID FOR

## AWS Goes Looking for Knowledge in the Resolved-Ticket Graveyard

Thousands of problems solved a month, and almost none of it written down.

---

5

FEATURE 2

WHAT MEMORY COSTS

## Glean Starts Billing for Memory While a Benchmark Says Knowledge Work Is Mostly Unsolved

Enterprise memory became a metered line item this month.

---

6

MOVERS AND SHAKERS

## This Week in the News

Nine people quoted on issues of organizational memory.

---

7

FUTURE

## The Future of Organizational Intelligence

Three important dates coming up in this field.

03

BROUGHT TO YOU BY

BROUGHT TO YOU BY

# Humanity Labs

---

Humanity Labs gives wealth management firms an AI Workforce that makes growth profitable. In a people-based business, serving more clients means hiring more people, so cost scales with revenue. The AI Workforce breaks that trap.

It works inside the systems your team already runs and delivers done work across the front, middle, and back office, while your people stay focused on the decisions that matter.

Every job it runs writes into one shared memory, building the firm's Organizational General Intelligence (OGI) so it gets smarter as it grows.

The result is the Firm of the Future: a firm that serves more clients and grows without adding cost at the same rate.

Humanity Labs works with some of the world's leading wealth management firms, which together oversee more than \$750 billion in client assets, and is backed by leading venture firms including Google Ventures and Vestigo Ventures.

Learn more at [HumanityLabs.ai](https://HumanityLabs.ai)

## IN THIS ISSUE

# Schedules, Tickets, Polished Drafts

“Data alone is not the same as knowledge.”

- CORY ST. PIERRE, KLEINFELDER, AUG 25

Organizational intelligence is the practice of keeping what a company learns from doing its own work. Most of that learning never gets written down, and the part that does now arrives looking more finished than it is. Here are four news items this week on what an organization records, and who checks it.

## MOST NOTABLE

## Polished Drafts Stop People From Spotting What Is Missing

Atlassian’s Teamwork Lab gave 903 knowledge workers the same proposal for reducing IT support tickets. The proposal contained two deliberate flaws: no adoption strategy and the wrong success metrics. Readers who got the AI-polished version were 22 percent less likely to catch the first flaw and 15 percent less likely to catch the second. They spent 62 percent more time on it, were 18 percent less willing to offer feedback, and were 2.5 times more likely to judge it finished work. Labelling the same document “Status: Early Draft” and asking for feedback nearly erased the entire effect. Managers and Gen X readers, though, spotted the flaws regardless.

ATLASSIAN TEAMWORK LAB

ALSO NOTABLE

## Amazon Publishes a Pattern for Capturing What Retiring Technicians Know

Amazon Web Services published a reference build on Monday for a knowledge system aimed at the oldest problem in this field: the expertise that leaves when a person does. The write-up describes capturing production procedures and maintenance protocols before experienced technicians retire, delivered through a voice-first interface, with the material held in Amazon S3 and retrieved through Amazon Bedrock Knowledge Bases. The company names manufacturing, healthcare, financial services, energy and government agencies as settings for it. The deployment is scripted to stand up in hours.

AWS MACHINE LEARNING BLOG

---

ALSO NOTABLE

## Law Firms' Own Files Become the Thing the Software Answers From

Google Cloud took Gemini Enterprise for Legal into preview on Tuesday with Cleary, Freshfields, Weil and Williams & Connolly as launch firms, connecting through Model Context Protocol into iManage, NetDocuments, RelativityOne, Everlaw, Harvey and Thomson Reuters. Google says a firm's ethical walls are inherited automatically and that client files and negotiated positions stay inside the firm's perimeter. Wolters Kluwer announced the publisher's version the day before, linking laws, rulings and commentary into a knowledge graph; Philipp Mueller, its legal and regulatory chief product officer, called it the biggest change to the company's content since it went digital. Early access opens in the coming months.

GOOGLE CLOUD

## Okta Gives an Agent a Directory Entry and an Owner

Okta made AI agents first-class identities in its Universal Directory on Monday, issuing short-lived tokens in place of stored credentials, included in core single sign-on at no extra cost. The most useful part of the announcement is the company's description of what it is fixing: most agents reach enterprise data through static keys and one-off grants and, in Okta's words, "operate as anonymous traffic with no owner, no policy, and no audit trail."

## 05 FORESIGHT

## FORESIGHT

# One Good Bet, One Wild Card

Two possible scenarios in the future of organizational memory.

## GOOD BET

RENEWAL CYCLES THROUGH 2027

Buyers will start asking vendors which named person each entry in the record resolves to, and will treat a shared service account as a failed answer.

## EVIDENCE

- 01** Identity provider Okta's survey of 292 executives and 492 knowledge workers found 96 percent were confident in their identity systems for software that is not a person. But only 34 percent apply the same controls to agents that they apply to employees. A gap that wide between belief and practice is likely to come up in procurement questionnaires.  
[Okta, May 27](#)
- 02** The AI AGENT Act, S.5051, introduced July 21, 2026, defines a custodial user agent as one authorized to act in a documented, limited and revocable manner, requires real-time records of the actions it takes, and directs NIST to develop standards for verifying that a user delegated the authority. Procurement language often tracks draft standards before they pass. [NIST, Feb 2026](#)

A federal agency will be asked to produce an agent's working record and will find it was already deleted under no approved schedule.

---

#### EVIDENCE

- 01** The US National Archives and Records Administration told federal agency records officers on August 21, 2026 that the definition of a federal record applies to the inputs, outputs, data and audit trails involved in the use of AI, and that agencies may only dispose of those records under a NARA-approved schedule. Few agencies have a schedule that names AI materials. [NARA, Aug 21](#)
- 02** Related underlying conduct is already going through the courts. Three days before that NARA memo, on August 18, a former senior official at the National Institute of Allergy and Infectious Diseases pleaded guilty to a conspiracy count connected to concealing federal records and evading records law during the pandemic. He faces up to five years. [US DOJ, Aug 18](#)

## COVER STORY

# The National Archives Says AI Material Is a Federal Record and Cannot Be Deleted on Anyone's Say-So

One memo on Friday moved the decision about what an agency may forget from the engineering team to the records officer. Corporate archivists have been making the same argument for years without a statute behind them.

## SOURCES:

NATIONAL ARCHIVES

LUCIDEA

HELP NET SECURITY

US DOJ

## QUICK TAKE

- NARA told federal agency records officers Friday that the definition of a federal record reaches AI inputs, outputs, data and audit trails (- National Archives, Aug 21)
- Agencies may dispose of that material only under a NARA-approved records schedule, and few schedules name AI materials
- Corporate archives face the same question without a statute: budget pressure, record volume, and a preservation mandate that makes records management significantly harder (- Lucidea, Aug 24)

- Europe's Article 12 logging duty has been live since August 2, with a six-month minimum retention for most high-risk systems
- 

The National Archives and Records Administration told federal agency records officers on Friday, August 21, how to apply the definition of a federal record to “inputs, outputs, data, audit trails, software, and other materials involved in the use of AI.” Memorandum AC 11.2026 does not ask agencies to collect anything new. It tells them what they may not throw away. ([National Archives, Aug 21](#))

**Why it matters:** the memo regulates deletion rather than collection. Agencies may dispose of AI-related federal records only under a NARA-approved records schedule. An agency that cannot point to a schedule naming its AI materials has not been given permission to tidy them away. The records officer, not the engineering team, now decides what the organization is allowed to forget.

**Between the lines:** NARA drew its boundary narrowly, and the exclusions are worth reading. The guidance “does not establish policy related to AI governance, e-discovery, privacy, security, or ethical use.” It answers one question, about custody, and leaves untouched many other questions.

**Meanwhile, in the corporate market:** corporate archivists have been arguing the same case all year without a memo to point at. Margot Note, an archivist and consultant writing on Monday, named the four fights: proving the archive's worth when it is read as a support function rather than a revenue line, managing the volume that email and chat now produce, protecting material that has to stay both discoverable and confidential, and reconciling records management with preservation. ([Lucidea, Aug 24](#))

**Yes, but:** those last two disciplines pull in opposite directions. Records management exists to destroy material once its retention period expires. Archiving exists to keep what has enduring value. Both are memory policy, and an organization that funds neither is making the decision by default, in favour of deletion.

---

1

Approved schedules an agency needs  
before it may delete any AI material

4

Fights a corporate archive is having at  
once, per Margot Note

---

CONTINUED ON NEXT PAGE →

## SEVEN DAYS IN ORGANIZATIONAL INTELLIGENCE

SUBSCRIBE

COVER STORY

ISSUE 3 • AUG 26 2026

## 07 COVER STORY

← CONTINUED FROM PREVIOUS PAGE

## WHAT THE RULE DOES NOT REACH

The duty to keep a record of what a system did is older than this month, and in Europe it is already enforceable. The record-keeping obligations in Article 12 of the European Union's AI Act took effect on August 2, three weeks before the NARA memo. High-risk systems must automatically record events in a form regulators read as tamper evident, with a minimum retention of six months for most systems and 24 months for biometric identification and law enforcement uses. The duty falls on deployers as well as providers, so an organization running someone else's system must keep its own logs. In Germany the Bundesnetzagentur took up supervision the same day. ([Help Net Security, Aug 3](#))

**The precedent:** concealing federal records is a charged offence, not a theoretical one. On August 18, three days before the memo, a former senior official at the National Institute of Allergy and Infectious Diseases pleaded guilty to a conspiracy count connected to concealing federal records and evading records law during the pandemic. He faces up to five years. ([CBS News](#))

**“Discoverable and useful for business purposes while also being protected from unauthorized disclosure.”**

- MARGOT NOTE ON CORPORATE ARCHIVES, AUG 24



## 08 FEATURE

THE KNOWLEDGE YOU ALREADY PAID FOR

# AWS Goes Looking for Knowledge in the Resolved-Ticket Graveyard

Enterprise support teams solve thousands of problems a month and turn almost none of it into anything another person can find. A reference build published Wednesday reads the closed tickets and drafts the articles nobody wrote.

## SOURCES:

AWS MACHINE LEARNING BLOG

KLEINFELDER

LUCIDEA

## QUICK TAKE

- AWS published KnowledgeForge on Wednesday, a reference build that mines resolved IT service management tickets for potential knowledge base articles (- AWS, Aug 19)
- It scores drafts against ten weighted quality dimensions and retires near-duplicates above 0.95 similarity
- A knowledge manager reviews every article before it publishes; the system drafts, it does not decide

- Kleinfelder, writing about utilities the same week: “Data alone is not the same as knowledge” (- Kleinfelder, Aug 25)
- 

Every resolved support ticket is a small, complete record of a problem somebody solved. Too often, almost none of them become anything a colleague can find six months later. AWS published a reference build on Wednesday, called KnowledgeForge, that treats that pile as the raw material it is: it clusters resolved IT service tickets, drafts knowledge base articles grounded in the documentation that already exists, and curates the existing library at the same time. ([AWS Machine Learning Blog, Aug 19](#))

**Why it matters:** this is knowledge the organization already paid for. The engineer’s time was spent, the fix was found, the customer was answered. What was not funded in many cases is sufficient time for "after call work" to write it down in a meaningful form the next person can use, which is why knowledge bases fill with duplicates and stale pages while the same incident gets solved a fourth time.

**How it works:** the build scores each draft against ten weighted quality dimensions and detects near-duplicates at 95 percent similarity or better. Then it stops. A knowledge manager reviews every proposed article before publication. The system drafts and ranks; a person still decides what the organization is going to claim it knows.

**The same problem, different industry:** utilities have it in physical form. Cory St. Pierre, a program manager at Kleinfelder, described asset information scattered across “records management platforms, GIS systems, work management applications, engineering drawings, spreadsheets, shared drives, and email archives,” with naming conventions that drifted over decades. His question is the one every knowledge base eventually faces: “When information differs between sources, which one should be believed?” ([Kleinfelder, Aug 25](#))

**Yes, but:** KnowledgeForge is a reference architecture, not a product with published results. AWS names no customer, reports no reduction in repeat tickets, and publishes no before-and-after on article quality. The mechanism is documented in detail but the outcome is not documented at all.

**What's next:** the discipline is older than the tooling and still mostly unstaffed. Stan Garfield, answering the question of what a new knowledge manager should do in a first 90 days, put listening and finding the existing communities ahead of any system. The tools now draft faster than an organization can decide what is worth keeping, which puts the bottleneck back where it was: on the person who reviews. ([Lucidea, Aug 20](#))

---

## 09 FEATURE

## WHAT MEMORY COSTS

# Glean Starts Billing for Memory While a Benchmark Says Knowledge Work Is Mostly Unsolved

Enterprise memory moved from a bundled feature to a metered line item this month. The public test of whether it works has not been passed by anyone.

## SOURCES:

[GLEAN RELEASE NOTES](#)[SIERRA](#)[FORRESTER](#)[PR NEWSWIRE](#)

## QUICK TAKE

- Glean began billing Engram Memory usage on Core Suite plans on August 18; Enterprise Flex plans start consuming credits on September 16 (- Glean, Aug 18)
- The best score on Sierra's knowledge benchmark is 37.4 percent of tasks passed on a first attempt (- Sierra, May 13)
- Sierra's own reading: the leading model still fails roughly 60 percent of these tasks at maximum reasoning effort
- Pinecone says its knowledge layer beat that model by 1.0 percentage point,

Glean put a price on remembering. In release notes this month, the enterprise search company said Engram Memory usage is now billed on its Core Suite plans, with customers on Enterprise Flex credit plans free until September 15 and consuming credits from September 16. The company was careful to say the change “does not change the underlying memory experience.” ([Glean Release Notes, Aug 18](#))

**Why it matters:** what a company remembers about its own work has been sold as a property of the product being used. Metering it makes organizational memory a quantity somebody has to approve. The same notes told administrators to review usage before the billing date if they rely on persistent memory, which is a statement with a budget owner attached.

**Catch up quick:** the pitch for paying is that the knowledge layer, not the model, is now where accuracy and cost are decided. Pinecone made that same argument when it took its product Nexus generally available on August 6, selling a governed knowledge layer that runs inside the customer's own cloud. Ash Ashutosh, the company's chief executive, framed the stakes as ownership: “every model call risks handing proprietary knowledge to a system that can turn around and compete with you.” ([PR Newswire, Aug 6](#))

**Yes, but:** read the benchmark the claim rests on. Sierra's tau-knowledge test gives an agent 698 documents across 21 product categories and tasks that take an average of 18.6 documents and 9.5 tool calls each. Pinecone reported 47.4 percent success, against 46.4 percent for GPT-5.5, a margin of one point. On Sierra's own public leaderboard the leading configuration sits at 37.4 percent of tasks passed first time, and Sierra's assessment is that the benchmark is “nowhere near saturated: even the leading model fails roughly 60% of these tasks at maximum reasoning effort.” ([Sierra, May 13](#))

**The category forming around it:** Forrester named the thing on August 20, and the definition is worth reading closely because it is doing category work. Boris

Evelson and Indranil Bandyopadhyay call a context layer “the next evolution of semantic layers and knowledge graphs,” and they are precise about what it takes from each: the business semantics and governance of a semantic layer, and the ontological modeling of a knowledge graph. The third ingredient is what makes it separate. A context layer “continuously incorporates runtime context such as events, decisions, actions, and outcomes,” producing what they call “a living model of the enterprise that enables AI reasoning, automation, and decision intelligence.” A semantic layer settles what the company means by a customer. A knowledge graph holds how customers connect to contracts and territories. A context layer adds what happened this morning. ([Forrester, Aug 20](#))

**Why the naming is the news:** Forrester has scheduled a landscape report on the category for the end of the fourth quarter of 2026, with a Wave evaluation to follow it. That sequence is how a category acquires a buying process: the landscape establishes who is in the market, the Wave scores them against published criteria, and procurement gets a shortlist it can defend. Organizational memory stops being a feature folded into somebody’s search product and starts being a line item with vendors ranked underneath it. The criteria do not exist yet, which is the part worth watching, because whoever shapes them defines what counts as memory infrastructure.

**The big picture:** Glean and Pinecone are pricing the knowledge layer as infrastructure while the public evidence for it is a single-digit margin on a test most attempts fail. A buyer can reasonably pay for both. The price is the part that has been settled.

|                                                                                               |                                                                                       |
|-----------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| <div>37.4%</div> <div>Best first-attempt score on Sierra's public knowledge leaderboard</div> | <div>1.0 pt</div> <div>Margin Pinecone reports over GPT-5.5, in its own testing</div> |
|-----------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|



## 10 MOVERS AND SHAKERS

## MOVERS AND SHAKERS

# This Week in the News

Nine people who made the news this week regarding Organizational Intelligence.

GROWTH • 50 TO 999 PEOPLE

3 PEOPLE

## Margot Note

Archivist and consultant, author for Lucidea Press

[LINKEDIN →](#)

Named the four fights a corporate archive is having at once: budget, volume, confidentiality, and preservation pulling against records management.

## Stan Garfield

Knowledge management author and consultant

[LINKEDIN →](#)

Puts listening and finding the existing communities ahead of any system, for a knowledge manager's first 90 days.

## Ash Ashutosh

Chief Executive, Pinecone

[LINKEDIN →](#)

"Argued that every model call risks handing proprietary knowledge to a system that can compete with you." [P.9 →](#)

SCALE • 1,000 TO 9,999 PEOPLE

3 PEOPLE

Cory St. Pierre

Program Manager, Kleinfelder

LINKEDIN →

“When information differs between sources, which one should be believed?” P.8 →

Boris Evelson

Vice President and Principal Analyst, Forrester

LINKEDIN →

Co-authored the definition of a context layer as a living model of the enterprise, published August 20.

Ric Smith

President of Products and Technology, Okta

LINKEDIN →

“Shouldn’t require trading away security or visibility.” P.4 →

ENTERPRISE • 10,000 TO 99,999 PEOPLE 2 PEOPLE

Molly Sands

Head of Teamwork Lab, Atlassian

LINKEDIN →

Ran the experiment showing a polished draft reads as finished work and stops people checking it.

Philipp Mueller

Chief Product Officer, Wolters Kluwer Legal & Regulatory

LINKEDIN →

“The biggest change to our content since we took it digital.” P.4 →

HYPERSCALE • 100,000 OR MORE PEOPLE 1 PERSON

# Thomas Kurian

Chief Executive, Google Cloud



Said legal research outputs are grounded in primary authority rather than in what a model absorbed during training.

## SEVEN DAYS IN ORGANIZATIONAL INTELLIGENCE

SUBSCRIBE

FUTURE

ISSUE 3 • AUG 26 2026

11

THE FUTURE OF ORGANIZATIONAL INTELLIGENCE

# Subscribe

A weekly read on how organizations capture, govern and reuse what they know. Free, and it arrives before the vendor briefings do.

SUBSCRIBE

NO SPAM. UNSUBSCRIBE ANYTIME.

UPCOMING

NEAR-TERM

NOV 16 2026

## KMWorld 2026, JW Marriott, Washington DC

The field's own conference meets a month after the first full quarter under NARA's guidance, with records officers and knowledge managers in the same building.

MORE →

## Forrester publishes its context layer landscape

The report is planned for the end of the fourth quarter and is the first time the category is mapped rather than described, naming who Forrester counts as being in this market. A Wave evaluation, which scores those vendors against published criteria, follows it.

[MORE →](#)

## Wolters Kluwer's Libra rollout reaches its final phase

The phased release running through the fourth quarter of 2026 and the first quarter of 2027 is the first test of whether a publisher's curated graph outperforms a firm's own files.

[MORE →](#)

### PREVIOUS EDITIONS

Aug 21 2026 **Retention, Tickets, Code Context**

[→](#)

Aug 10 2026 **Private By Default**

[→](#)

## 12

HOW THIS MAGAZINE WAS CREATED

## THE METHOD

# How This Magazine Was Created

This magazine, and all the magazines in the Humanity Labs AI Magazine series, were created using [Hawkeye](#), an AI system built for people who work in fields too large and too fast-moving to keep on top of alone.

Every week Hawkeye reads thousands of items across mainstream press, trade publications, government records, vendor announcements and company newsletters, then clusters them to find the themes that are actually emerging rather than the ones already being discussed.

For this issue the consideration set ran to 112 candidate items across seven source classes. A first draft was written with Claude Opus, checked against primary sources, and edited by Marshall Kirkpatrick.

Where a claim rests on a company's account of its own product, this magazine says so on the page.