



SEVEN DAYS IN ORGANIZATIONAL INTELLIGENCE

How organizations capture, govern and reuse what they know.

AI Can Rewrite Code, Not Intent

Forrester's Biswajeet Mahapatra says AI can scan a legacy system in hours but still can't say why any rule survived inside it. The same gap just surfaced in a federal records memo and a court fight over deleted OpenAI logs.

02 IN THIS ISSUE

Appraisal, Interviews, Substrate

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.

Organizational memory has two separate tests: what an organization is allowed to keep, and what it can still explain once the person who understood it is gone. This week, a Forrester analyst named that second test for private-sector code, a federal court tested it against OpenAI's own logs, and a government records agency admitted it hasn't solved it either.

COVER IMAGE: BOXES OF DOCUMENTS ON REPOSITORY SHELVING, THE NATIONAL ARCHIVES, KEW - WIKIMEDIA COMMONS (CC BY 3.0)

1

NEWS BRIEFS

WHO KEEPS THE RECORD

Appraisal, Interviews, Substrate

A new Archivist, a sanctions fight over deleted chat logs, and a retirement wave nobody's caught up to.

2

FORESIGHT

One Good Bet, One Wild Card

Two scenarios for how the record-keeping question resolves.

Forrester Says AI Cannot Recover Why A Legacy Rule Exists

A modernization warning for private-sector systems, tested this week in a court and a federal agency.

A Company Archivist Is Interviewing Engineers Before They Retire

Burns & McDonnell's oral-history method, and the retirement wave it's racing.

Enterprise Memory Becomes Infrastructure, Sold Separately From the Model

Elastic, Graphwise and two earlier moves toward owning the facts, not the reasoning.

This Week in the News

Six people on the record-keeping and memory questions this issue covers.

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

IN THIS ISSUE

Appraisal, Interviews, Substrate

"The threat of the Great Retirement is more than a demographic shift. It's a wakeup call for organizations and C-suite leaders to act now."

- LYNDA BRAKSIEK, APQC, 2026

Organizational intelligence is about keeping what an organization learns from its own work, but it's also about being able to prove it later. Two events in the news this week put the emphasis on that second part: a court that may treat a technically-preserved but unreadable record as no record at all, and the question of how far behind most employers still are on the oldest version of the problem, workers leaving before anyone captured what they knew.

MOST NOTABLE

OpenAI Faces Sanctions Over Chat Logs a Court Ordered Preserved

A federal sanctions fight over OpenAI's deleted ChatGPT logs, and the 20-million-record production a judge called unusable, is this issue's main story, starting on the next page.

SUBJECT TO INQUIRY

ALSO NOTABLE

The Senate Confirms a New Archivist of the United States

The Senate confirmed Bradford P. Wilson as the 12th Archivist of the United States on August 7, after he testified before the Homeland Security and Governmental Affairs Committee and drew bipartisan support. Wilson was nominated in March; he previously led the James Madison Program in American Ideals and Institutions at Princeton University before retiring from that post in 2024. The Archivist runs the National Archives and Records Administration, the agency that told federal records officers on August 21 which uses of AI create a record the government may not simply delete, and that now owns the same appraisal gap Forrester describes for private-sector systems in this issue's main story.

NATIONAL ARCHIVES

ALSO NOTABLE

A Nonprofit Research Group Puts a Number on the Retirement Wave

APQC's 'Great Retirement' research estimates that 3.6 million baby boomers are leaving the workforce every year, cresting through 2025, for a cumulative 61 million gone by 2030, most without ever documenting the tacit expertise they built up. 'The threat of the Great Retirement is more than a demographic shift,' said Lynda Braksiek, APQC's principal research lead for knowledge management. 'It's a wakeup call for organizations and C-suite leaders to act now.' The survey found the resulting knowledge exodus is a board-level concern most organizations still lack the process, technology or culture to manage.

APQC

FORESIGHT

One Good Bet, One Wild Card

Two ways the record-keeping question could resolve over the next two years.

GOOD BET

NEXT RECORDS-SCHEDULE CYCLE, 2027

Agencies and large litigants will start being judged on whether an AI-related record can still explain a specific decision after the fact, not just on whether it was technically kept.

EVIDENCE

- 01** Andrew Potter's August 31 critique documents that NARA's appraisal policy has not been revised since 2007, despite the agency's own 2020 Cognitive Technologies white paper recommending exactly that review. A named gap in the standard-setter's own guidance is usually the first thing the next version of that guidance closes. [Meta Archivist, Aug 31](#)
- 02** A federal court has already rejected "technically preserved" as sufficient in practice: OpenAI's 20-million-record production came back so heavily redacted the judge called it unusable. That is a live ruling on reconstructability, not a policy paper about it. [Subject to Inquiry, Aug 2026](#)

A court, rather than NARA, will end up setting the first binding standard for how much AI-decision evidence counts as enough to reconstruct what happened.

EVIDENCE

- 01** NARA's own guidance disclaims the harder question, stating it "does not establish policy related to AI governance, e-discovery, privacy, security, or ethical use." The agency that could set the standard has said, in writing, that it isn't going to. [National Archives, Aug 21](#)
- 02** The OpenAI spoliation motion is actively testing what "enough" looks like right now, in front of a judge who has already shown willingness to reject an inadequate production. A courtroom, not a rulemaking docket, is where this question is currently being litigated. [Subject to Inquiry, Aug 2026](#)

COVER STORY

Forrester Says AI Cannot Recover Why A Legacy Rule Exists

AI can scan a system's code, documents and telemetry faster than any human team, Forrester's Biswajeet Mahapatra argued this week. But it still can't say why any given rule survived, or whether it belongs in the future. A federal records memo and a court fight over deleted OpenAI logs show the same gap elsewhere.

SOURCES:

FORRESTER

SUBJECT TO INQUIRY

NATIONAL ARCHIVES

META ARCHIVIST

FEDSCOOP

QUICK TAKE

- Biswajeet Mahapatra (Forrester) argued Sept 2 that AI can analyze a legacy system's code, docs and telemetry fast, but 'only people can explain' why a rule survived and whether it belongs in the future state
- A federal court is already testing what happens when nobody answered that question in time: a 20-million-record sample OpenAI produced under a preservation order came back so redacted the court called it unusable

- The US government hit the identical wall from its own side, Aug 21: the National Archives and Records Administration's new AI-records memo settles who may discard what, not what must be provable later, on top of an appraisal policy unrevised since 2007
 - Three institutions, one company, one court, one federal agency, are independently running into the same missing layer: a record of why, not just what
-

AI can now read a decade of code, tickets and change logs faster than any engineering team ever could. What it still can't do, Forrester analyst Biswajeet Mahapatra argued this week, is tell you why any of it is still there. 'A rule found in code may represent a valid business requirement, an obsolete policy, a workaround for a retired system, or a defect that has persisted unnoticed for years,' the article says. 'Only people can explain why it exists and whether it belongs in the future-state application.'

Why it matters: Mahapatra's argument is aimed at software modernization budgets, and it names a failure mode most of them fund only half of. AI can accelerate behavioral discovery, mapping what a system actually does. But it cannot substitute for intent discovery, interviewing the people who know which of those behaviors are load-bearing and which are debris. 'AI can accelerate each of those activities,' the article says. 'It cannot decide which parts of the past deserve a place in the future.' A company that skips the second kind of discovery doesn't find out until the rule it deleted turns out to have been a regulatory requirement, or the one it kept turns out to have been a workaround nobody still needed.

The stakes when nobody answers it in time: a federal court is already running the experiment. A judge in the Southern District of New York ordered OpenAI in May 2025 to preserve ChatGPT output logs that would otherwise be deleted after 30 days. News publishers led by The New York Times filed a sanctions motion in July 2026 alleging OpenAI kept deleting logs anyway and

had, for more than two years, told the court it could not search data it had already built the capability to search. A 20-million-record sample OpenAI did produce came back so heavily redacted the court called it unusable, according to an August 2026 legal analysis. The company's own system did not preserve, in a form anyone could later explain, why it had done what it did.

The government hit the same wall from its own side: the National Archives told federal records officers on August 21 which uses of AI create a record the government is not allowed to simply delete. A records-management critique published nine days later, by Andrew Potter, a records-management writer who publishes as Meta Archivist, argued the memo answers the easy version of Mahapatra's question, custody, and skips the harder one, appraisal: what must survive so an AI-mediated decision can later be checked. NARA's own appraisal policy is unrevised since 2007.

20M Record sample OpenAI produced in the spoliation case, called unusable by the court	2007 Year NARA's appraisal policy was last revised, per Andrew Potter
--	---

Video player configuration error

Error 153

Watch video on YouTube

Learn more

Error 153

Inside the National Archives' AI Powered Museum Demo (Axios Live, Mar 30, 2026) NARA's own attempt to make its record holdings searchable, filmed before this week's guidance and its critics; one government case study of the same discovery problem Forrester describes for private-sector systems. Embedded in full; no clip boundary set.

CONTINUED ON NEXT PAGE →

SEVEN DAYS IN ORGANIZATIONAL INTELLIGENCE

SUBSCRIBE

COVER STORY

ISSUE 4 • SEP 3 2026

07 COVER STORY

← CONTINUED FROM PREVIOUS PAGE

WHAT TWO KINDS OF DISCOVERY BUY YOU

Mahapatra's fix isn't a better scanner. It's a second, human step most modernization programs never fund. Behavioral discovery, what AI does well, tells you what a system does. Intent discovery, what still requires a person, tells you why, and it means finding the people who ran the system and asking them directly which rules are load-bearing before those people, or their access, or their memory of the reasoning, is gone. Programs that fund only the first kind get through the project faster and then spend years finding out what they deleted.

Where the same fix would help NARA: the National Archives says its memo draws its boundary narrowly on purpose. It states plainly that it 'does not establish policy related to AI governance, e-discovery, privacy, security, or ethical use.' That leaves the harder evidentiary question, what an agency must be able to reconstruct, to whoever writes the next policy, a job that now sits with Bradford P. Wilson, confirmed by the Senate on August 7 as the 12th Archivist of the United States.

What happens if nobody does either kind in time: the OpenAI case is what it looks like when neither behavioral nor intent discovery ran before a court asked. A preservation order that only says 'Keep it' does not, on its own, guarantee anyone can later explain what a system did and why. Proving that took a sanctions motion, and, as of this writing, still isn't resolved. That is the cost a private-sector operator is actually pricing when a modernization budget skips the interviews to save time.

"AI can accelerate each of those activities. It cannot decide which parts of the past deserve a place in the future."

- BISWAJEET MAHAPATRA, FORRESTER, SEPT 2



08 FEATURE

THE KNOWLEDGE THAT LEAVES WHEN SOMEONE DOES

A Company Archivist Is Interviewing Engineers Before They Retire

Burns & McDonnell is recording its own engineers on tape while they can still explain their own decisions. Nationally, the workforce is retiring faster than most employers have a plan to catch what they know.

SOURCES:

ARCHIVAL OUTLOOK

APQC

ENBOARDER

QUICK TAKE

- Burns & McDonnell's corporate archivist records structured oral-history interviews with veteran engineers before they retire (- Archival Outlook, Jul/Aug 2026)
- APQC's Great Retirement research names the wave of departing baby boomers a board-level risk most organizations aren't set up to manage
- A 2025 HR survey found institutional knowledge loss is the single most common offboarding failure, ahead of security risk or rehiring cost

- The interview method requires no AI; it's a documentation choice being made under a shorter deadline
-

Lance Warren's job at Burns & McDonnell is to sit an engineer down before they retire and get them to explain, on tape, decisions nobody wrote down the first time. Warren, the employee-owned engineering firm's corporate research librarian and archivist, described the practice in the July/August 2026 issue of *Archival Outlook*, the trade magazine of the Society of American Archivists.

Why it matters: the method is old and the deadline is not. APQC's Great Retirement research puts the current pace at 3.6 million baby boomers leaving the workforce a year, with 61 million gone by 2030, most of them without ever documenting the tacit expertise it took decades to build. 'The threat of the Great Retirement is more than a demographic shift,' said Lynda Braksiek, APQC's principal research lead for knowledge management. 'It's a wakeup call for organizations and C-suite leaders to act now.'

By the numbers: a 2025 survey of HR leaders found institutional knowledge loss was the single most common offboarding challenge, cited by 47 percent of respondents; 76.6 percent said they were concerned about it. Roughly 42 percent estimated that inconsistent offboarding, including knowledge loss, security gaps and rehiring costs, costs their company up to \$500,000 a year. (- Enboarder, 2025)

What the interview method buys: an oral history is slower than a wiki page, and it captures something a form can't: the reasoning an engineer would only give if someone asked the right follow-up question. Burns & McDonnell keeps the recordings and transcripts in its corporate archive as institutional record, alongside the company chronologies and executive histories Warren also maintains. Nothing in the method requires AI. The AI-era pressure is what shortens the deadline, not what makes the interviews useful.

61M

Baby boomers APQC projects will have left the workforce by 2030

47%

HR leaders who name institutional knowledge loss their top offboarding challenge (- Enboarder, 2025)

Video player configuration error

Error 153

Watch video on YouTube

Learn more

Error 153

Knowledge Management (Backlog, Apr 6, 2026) A webinar on KM fundamentals and their application to archives, museums and cultural heritage organizations, the same institutional setting Warren's interview method sits in. Embedded in full; no clip boundary set.



09 FEATURE

WHAT THE SYSTEM OF RECORD BECOMES

Enterprise Memory Becomes Infrastructure, Sold Separately From the Model

Elastic says it is 'the data substrate beneath frontier models,' not a competitor to them. But three weeks earlier, a private-equity firm paid for control of a knowledge-graph company, betting the layer that remembers outlasts whichever model sits on top of it.

SOURCES:

ELASTIC Q1 FY2027

OAKLEY CAPITAL

REDIS

SNOWFLAKE

QUICK TAKE

- Elastic reported strong Q1 FY2027 growth on Sept 3 and told investors it wants to be 'the data substrate beneath frontier models,' retrieving an organization's own data while OpenAI's models reason over it
- Oakley Capital, a private equity firm, bought majority control of Graphwise, a knowledge-graph platform serving 200-plus enterprise customers, in a deal announced Aug 19

- Redis (May) and Snowflake (June) made similar moves earlier this year, building a context layer that sits between an agent and the business data it needs
 - None of these products is the model doing the reasoning; all of them are where an organization's own facts get kept and governed
-

Elastic says it does not want to compete with OpenAI, Anthropic or Google. It wants to sit underneath all of them. In its first-quarter fiscal 2027 results, reported September 3, the company described itself as "the data substrate beneath frontier models" and said an expanded OpenAI collaboration would split the work: Elastic retrieves an organization's own documents, logs and communications, and OpenAI's models reason over what it finds.

Why it matters: that division of labor, one company owns the facts, another owns the reasoning, is becoming the default shape of enterprise AI, and it means the facts now have a market of their own. Three weeks before Elastic's results, a private equity firm called Oakley Capital paid up to £20 million for majority control of Graphwise, a knowledge-graph platform serving more than 200 enterprise customers with annual recurring revenue growing more than 30 percent. Oakley isn't buying a model. It's buying the layer that tells a model what a customer's data actually means.

The pattern started earlier this year: Redis launched Iris, a context and memory platform, in May, describing it as sitting between an agent and the business data it needs to act. Snowflake followed in June with Horizon Context and Cortex Sense, giving customers a shared definition of business logic that multiple AI tools can draw from instead of each reinventing it. Both address the same failure: an agent that reasons well but works from stale or ungoverned data isn't useful, whichever model powers it.

What this means for the people who run the system: an organization that buys a context layer is making a bet about which part of its AI stack it wants to own outright. The model can be swapped; a contract change can put a different

one in the loop. But the record of what the business actually knows, who's allowed to see it, and where it came from is not something a vendor can swap out from under a company. That's the part organizations are now being asked to buy, govern and keep, on purpose, as its own line item.

£20M Oakley Capital's stake in knowledge-graph platform Graphwise, announced Aug 19	30%+ Graphwise's annual recurring revenue growth, per Oakley Capital
---	--



10 MOVERS AND SHAKERS

MOVERS AND SHAKERS

This Week in the News

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

GRASSROOTS • UNDER 50 PEOPLE

1 PERSON

Andrew Potter

Records-management writer, publishes as Meta Archivist

LINKEDIN →

"What must survive so that an AI-mediated government action can later be understood, trusted, and held accountable?" [P. 6](#) →

NETWORK • 50 TO 499 PEOPLE

3 PEOPLE

Lynda Braksiek

Principal Research Lead, Knowledge Management, APQC

LINKEDIN →

"It's a wakeup call for organizations and C-suite leaders to act now."

Atanas Kiryakov

President and Co-founder, Graphwise

LINKEDIN →

Took Oakley Capital's majority investment, calling it a partnership to expand Graphwise's platform and support wider AI adoption.

Faith Charlton

Archivist; co-editor, "User Experience and Archives," American Archivist

LINKEDIN →

Co-edited the journal's Spring/Summer 2026 special section on whether the people archives are built for can actually find what's in them.

Lance Warren

LINKEDIN →

Corporate Research Librarian and Archivist, Burns & McDonnell

Described the firm's practice of recording oral-history interviews with engineers before they retire.

Bradford P. Wilson

LINKEDIN →

12th Archivist of the United States

Confirmed by the Senate on August 7; now oversees the federal appraisal question this issue's cover story raises.

SEVEN DAYS IN ORGANIZATIONAL INTELLIGENCE

FUTURE

ISSUE 4 • SEP 3 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 under a new Archivist of the United States and an unresolved appraisal question, with records officers and knowledge managers in the same building.

MORE →

NARA's FY2027 plan to deploy AI for federal records access review

NARA's FY2027 performance plan commits the agency to using AI to speed up access review of the records it holds, a test of whether the agency practices on itself what it just told other agencies to keep.

[MORE →](#)

A ruling in the OpenAI ChatGPT-log spoliation fight

Whether the court imposes sanctions, or an adverse-inference instruction, over OpenAI's deleted and heavily redacted chat logs will be the first real test of how much reconstructable evidence a preservation order actually requires.

[MORE →](#)

PREVIOUS EDITIONS

Aug 26 2026 **Schedules, Tickets, Polished Drafts**

[→](#)

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

[→](#)

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 roughly 30 candidate items across five source classes: journalism, government and institutional records, vendor and company announcements, professional and trade publications, and litigation records. 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.