

Access control platform built for AI agents

Corporate IAM assumes human users with logins and sessions; AI agents act continuously, impersonate service accounts, and inherit far more privilege than any single task needs, with no per-action identity or revocation story.

Access control platform built for AI agents should be tested as a narrow first-win workflow for Security or platform engineering lead at a company deploying autonomous agents against internal systems.

HIGH DIFFICULTY

PER-AGENT-IDENTITY MONTHLY PRICING WITH ENTERPRISE POLICY AND COMPLIANCE TIERS.

56/100

VALIDATION VERDICT / RESEARCH

Validation is a weighted rubric, not a guarantee. Use the next validation step before building.

Confidence	55%
Lifecycle	Crowding
Timing	21/100
Rubric	INAV-VALIDATION-2026-06-04

CROWDING

Window closing

Demand signal	5.5/10
Problem severity	6.3/10
Willingness to pay	5/10
Competitive saturation	6.7/10

VERDICT

Research • 56/100

Access control platform built for AI agents should be tested as a narrow first-win workflow for Security or platform engineering lead at a company deploying autonomous agents against internal systems.

THIS WEEK'S TEST

Interview twenty security leads blocking agent deployments, then pilot the broker on one high-value workflow per design partner and measure whether it unblocks the rollout.

KILL IT IF

Fewer than five qualified buyers agree to discuss the workflow after targeted outreach.

— READER DEMAND SIGNAL

Would you build this? Would you pay?

One tap each — anonymous, one vote per question per day. Tallies update as readers weigh in.

Would you build this?

Be the first to weigh in

Would you pay for this?

Be the first to weigh in


■ ■ ■

IDENTITY AND ACCESS MANAGEMENT FOR AI AGENTS

SIGNAL MODEL

Access control platform built for AI agents

Access control platform built for AI agents should be tested as a narrow first-win workflow for Security or platform engineering lead at a company deploying autonomous agents against internal systems.



Point	Value (Relative)
1	Baseline
2	Slightly below baseline
3	Slightly above baseline
4	Below baseline
5	Above baseline
6	Slightly above baseline
7	Peak (highest point)
8	Below peak, above baseline

VALIDATION

56/100

Research

CONFIDENCE

55%

Editorial confidence

SCORE AVG

6/10

Scorecard average

PROOF

5.8/10

Proof signal average


SCORE RADAR

Decision balance



VALUE EQUATION

Offer strength



MARKET MAP

Category king candidate

CUSTOMER VALUE



High value plus high uniqueness deserves deeper research; lower uniqueness requires a clear distribution advantage.

VALIDATION FUNNEL

From pain to product.

1	Buyer pain Security or platform engineering lead at a company deploying autonomous agents ag...	5.4/10
2	Concierge proof Interview twenty security leads blocking agent deployments, then pilot the broker...	5.8/10
3	Paid wedge Concierge review or paid template	6/10
4	Repeatable product Per-agent-identity monthly pricing with enterprise policy and compliance tiers.	5.9/10

EVIDENCE HEATMAP

Signal intensity.

WHY NOW 5/10 Demand visibility	WHY NOW 4/10 Tooling readiness
WHY NOW 4/10 Budget clarity	WHY NOW 7/10 Competitive window
PAIN 5/10 Repeated workflow friction	MONEY 4/10 Budget hypothesis

URGENCY

6/10

Switching pressure

DISTRIBUTION

8/10

Reachable buyer language

Crowding (21/100): demand exists, but funded or visible competitors are compressing the window.

Deterministic stage assignment from re-check status, demand signals, complaint echo, and competitive saturation.

21/100

CROWDING

Adoption substrate is up 945.8% across matched packages.

3 matched company signals raise saturation.

Demand

68/100

Not old enough for a 30-day re-check yet.

Saturation

90/100

3 funded signals across 3 matched competitor signals.

Complaint echo

22/100

Matched adoption substrate is up 945.8%.

— VALIDATION ENGINE

Evidence-backed idea-validation score.

The score uses a versioned 2026 rubric across demand, problem severity, willingness to pay, competitive saturation, and feasibility.

56/100

Research

Research is the current validation verdict: competitive saturation is the strongest signal, while feasibility is the main evidence gap to close before scaling the build.

Rubric version: INAV-VALIDATION-2026-06-04 / generated August 12, 2026

Demand signal

5.5/10

24% WEIGHT

Demand looks thin because the report has 2 source-backed signal(s), an editorial confidence of 55/100, and a defined buyer in Identity and access management for AI agents.

- Agent frameworks authenticate with static keys and broad scopes, which security teams flag as their top blocker to wider agent rollout.
- Target buyer: Security or platform engineering lead at a company deploying autonomous agents against internal systems

Problem severity

6.3/10

22% WEIGHT

Problem severity is thin when the buyer pain, customer value, and dream-outcome scores are combined.

- Corporate IAM assumes human users with logins and sessions; AI agents act continuously, impersonate service accounts, and inherit far more privilege than any single task needs, with no per-action identity or revocation story.
- Agent frameworks authenticate with static keys and broad scopes, which security teams flag as their top blocker to wider agent rollout.

Willingness to pay

5/10

20% WEIGHT

Willingness to pay is weak; the model has a monetization hypothesis, but it must still be proven through paid pilots or explicit pricing objections.

- Per-agent-identity monthly pricing with enterprise policy and compliance tiers.
- Interview twenty security leads blocking agent deployments, then pilot the broker on one high-value workflow per design partner and measure whether it unblocks the rollout.

Competitive saturation

6.7/10

18% WEIGHT

No source-backed direct match is recorded yet, so saturation risk is treated as unknown rather than proof of novelty.

- Existing-product check has no named direct match.
- Competitive score rewards a narrow wedge, not absence of research.

Feasibility

4/10

16% WEIGHT

Feasibility is weak for a high build if the MVP is limited to the first measurable workflow.

- Interview twenty security leads blocking agent deployments, then pilot the broker on one high-value workflow per design partner and measure whether it unblocks the rollout.
- Okta, Microsoft, and cloud providers are all positioned to extend existing IAM to non-human identities.

Next validation step

Interview twenty security leads blocking agent deployments, then pilot the broker on one high-value workflow per design partner and measure whether it unblocks the rollout.

Seven days to a build / kill decision.

Derived from this report's own validation test, channels, offers, and kill criteria.

Each day has a threshold, so the week ends in a decision instead of a feeling.

DAY 1

Build the buyer list

List 50-100 named security or platform engineering lead at a company deploying autonomous agents against internal systems prospects from Community pain posts and Direct outreach — names, not categories.

Threshold: 50+ named, reachable buyers on the list.

DAY 2

Join the watering holes

Join and observe Reddit / forums, Launch communities, Review and alternative pages. Collect the exact words buyers use for this pain.

Threshold: 10+ verbatim pain quotes captured.

DAY 3

Send first outreach

Send the cold outreach template (below) to 15 buyers from the day-1 list, personalized with one detail each.

Threshold: 15 sent; 3+ replies of any kind.

DAY 4

Run buyer interviews

Hold 15-minute calls using the interview script (below). Listen for current workarounds and what they cost.

Threshold: 3+ completed interviews.

DAY 5

Run the report's validation test

Interview twenty security leads blocking agent deployments, then pilot the broker on one high-value workflow per design partner and measure whether it unblocks...

Threshold: Problem resonance: 5+ calls or 10+ detailed replies.

DAY 6

Make the smoke offer

Offer "Concierge review or paid template" at \$19-\$99 to every interviewed buyer. Manual delivery is fine — payment is the signal.

Threshold: 1+ pre-commitment (payment, signed LOI, or scheduled paid pilot).

DAY 7

Decide against the kill criteria

Score the week against this report's kill criteria, then take the stated next validation step: Interview twenty security leads blocking agent deployments, then pilot the broker on one high-value workflow per design partner and measure whether it unblocks...

Threshold: A written build / keep-testing / kill decision.

Pass signal

Pass: thresholds on days 3, 4, and 6 are met — proceed to the next validation step with real buyer language in hand.

Fail signal

Kill or rethink if the week confirms: Fewer than five qualified buyers agree to discuss the workflow after targeted outreach.

Decision scorecard.

The report is structured to force a yes, no, or test decision instead of leaving the reader with a loose brainstorm.

Opportunity PROMISING Access control platform built for AI agents has an editorial confidence score of 55/100 before live buyer validation.	6/10
Problem PROMISING Corporate IAM assumes human users with logins and sessions; AI agents act continuously, impersonate service accounts, and inherit far more privilege than any single task needs, with no per-action identity or revocation story.	5/10
Feasibility NEEDS PROOF A high build can work if the MVP stays limited to the first repeated workflow.	4/10
Why now EXCEPTIONAL 2026 is the year agents moved from demos to production workflows, and every deployment currently borrows a human’s OAuth token or a god-mode API key because agent-native IAM doesn’t exist yet.	9/10

Revenue potential

\$250K-\$2M ARR potential if the wedge proves budget urgency and becomes a recurring workflow.

Execution difficulty

Execution is high; the main constraint is staying narrow enough for a first proof loop.

Go-to-market

Start with manual concierge output, direct outreach, and community proof before paid acquisition.

Founder fit

Best for an AI-assisted solo founder who can interview the buyer and ship a focused first version quickly.

1. Lead magnet

Access Control Platform Built For Ai Agents checklist

Free

Helps Security or platform engineering lead at a company deploying autonomous agents against internal systems audit the painful workflow before buying software.

Capture qualified leads and learn the buyer’s exact language.

2. Frontend offer

Concierge review or paid template

\$19-\$99

Delivers the first useful output manually before automation is trusted.

Validate urgency, workflow fit, and willingness to pay.

3. Core offer

Access control platform built for AI agents focused SaaS

\$49-\$499/month

Turns the recurring manual workflow into a repeatable product loop.

Create the recurring revenue product after the narrow wedge survives tests.

4. Continuity

Monitoring, benchmarks, and monthly reporting

\$99-\$1,000/year add-on

Keeps the buyer engaged with ongoing proof, saved time, or reduced risk.

Increase retention and make the product part of a routine.

5. Backend offer

Done-with-you setup, agency, or team rollout

Custom

Adds implementation help, integrations, and workflow migration.

Capture higher-value accounts once the productized wedge is proven.

Price-anchored revenue scenarios.

Derived from this report's "Core offer" offer-ladder stage (\$49-\$499/month). These are price-anchored scenarios, not market-size claims.

Proof 10 CUSTOMERS Ten paying customers proves willingness to pay and funds continued validation.	\$490-\$4,990 MRR
Wedge 50 CUSTOMERS Fifty customers in one niche makes the workflow the default in that circle and feeds referrals.	\$2,450-\$24,950 MRR
Vertical leader 250 CUSTOMERS A few hundred accounts in one vertical is a real business before any horizontal expansion.	\$12,250-\$124,750 MRR
Break-even At \$49-\$499/month, 1 customers cover the stated Local-first MVP budget: \$0-\$10K before paid acquisition. budget within a month; fewer if they land at the top of the range.	
Sizing the buyer universe Size the buyer universe in one day: count security or platform engineering lead at a company deploying autonomous agents against internal systems reachable through the report’s channels (directories, associations, communities) until the list stops growing — the test only needs the first 100 names, not a TAM estimate.	
Pricing benchmark No public look-alike products were recorded in this report, so price against the manual workaround’s time cost, not against software.	

Mixed reflexivity — execution over secrecy.

Does publishing this opportunity strengthen it or crowd it? low confidence.

Publish or protect

No dominant reflexivity signal. Publish the analysis, but note that execution speed and distribution matter more than secrecy here; revisit if the space shows saturation.

Signals

- FULFILLING

 Platform play — matched "platform"

— EVIDENCE	
Why now and proof signals.	
Why now	
5/10 Demand visibility Agent frameworks authenticate with static keys and broad scopes, which security teams flag as their top blocker to wider agent rollout. Build only if the complaint repeats across interviews, posts, or existing workflow artifacts. <code>oauth.net</code>	
4/10 Tooling readiness AI-assisted product work and managed infrastructure reduce the first-version cost. The first release should automate one high-friction step rather than become a broad platform. <code>en.wikipedia.org</code>	
4/10 Budget clarity Per-agent-identity monthly pricing with enterprise policy and compliance tiers. Ask for money during validation before building the full workflow. <code>oauth.net</code>	
7/10 Competitive window The wedge is specific enough to test without claiming the whole market. Position around one buyer and one measurable first-win outcome. <code>oauth.net</code>	
Proof signals	
5/10 Pain: Repeated workflow friction Agent frameworks authenticate with static keys and broad scopes, which security teams flag as their top blocker to wider agent rollout. <code>oauth.net</code>	

4/10

Money: Budget hypothesis

Security or platform engineering lead at a company deploying autonomous agents against internal systems is the first group to test because the monetization path is: Per-agent-identity monthly pricing with enterprise policy and compliance tiers. [oauth.net](#)

6/10

Urgency: Switching pressure

Urgency becomes real only if the current workaround costs time, risk, money, or reputation every week. [en.wikipedia.org](#)

8/10

Distribution: Reachable buyer language

The first channel should be whichever source lane already contains the buyer's vocabulary.

[oauth.net](#)

— POSITIONING

Market gaps and execution plan.

Underserved segments

- Security or platform engineering lead at a company deploying autonomous agents against internal systems who still run the workflow in spreadsheets, generic docs, email, or chat threads.
- Small teams in Identity and access management for AI agents that feel the pain weekly but are too narrow for broad incumbents.
- New adopters who need guided proof before committing to a larger platform.

Feature gaps

- A narrow workflow that reaches value without configuration-heavy onboarding.
- A buyer-facing proof artifact that shows time saved, risk reduced, or communication improved.
- A handoff path from manual concierge service to repeatable software.

Differentiation levers

- Use specificity as the wedge: one buyer, one workflow, one measurable result.
- Show proof earlier than broad competitors with before-and-after examples and small pilot data.
- Keep implementation lighter than incumbent suites or generic AI assistants.

Execution snapshot

Type	SaaS product
Timeline	8-12 weeks
Budget	Local-first MVP budget: \$0-\$10K before paid acquisition.
Initial offer	Concierge review or paid template
Build only the first-win workflow for "Access control platform built for AI agents" and keep research, setup, and exceptions manual until the wedge is proven.	

Weekly

Community pain posts

Use communities and forums where Security or platform engineering lead at a company deploying autonomous agents against internal systems already describe the painful workflow.

Problem teardown, interview ask, and short demo clip / 5 qualified calls or 10 detailed replies in 7 days

Daily during validation

Direct outreach

Direct conversations are the fastest way to verify budget ownership and switching cost.

Concierge pilot offer with a manually prepared sample / 3 paid pilots, LOIs, or budget-owner follow-ups

Bi-weekly

Searchable comparison content

Alternative and comparison pages reveal objections, pricing language, and buying intent.

Before-and-after page or alternatives memo for the exact workflow / Organic clicks, booked demos, or waitlist joins from comparison intent

Once MVP is clickable

Launch directory

Launches test whether the promise is legible to people outside the first interview set.

Single-purpose demo and first-win story / 25% demo completion or 10 waitlist joins

COMPETITIVE LANDSCAPE	
Alternatives, incumbents, and whitespace.	
This section names likely workarounds and public players so the report can argue where the wedge is still open.	
Access control platform built for AI agents should be positioned against generic AI assistants, no-code workarounds, and any vertical incumbent that already owns Identity and access management for AI agents. The opening is a narrower first-win workflow for Security or platform engineering lead at a company deploying autonomous agents against internal systems.	
DIRECT	ServiceTitan Field service platform Relevant to field service, HVAC, appliance repair, contractor, and service dispatch ideas.
WORKAROUND	Airtable No-code database Competes when the first version can be modeled as a lightweight database and workflow view.
ADJACENT	HubSpot CRM and marketing platform Competes for sales, marketing, client follow-up, webinar, and service pipeline workflows.
WORKAROUND	Asana Project management Competes where the buyer can express the workflow as tasks, owners, and due dates.

WORKAROUND

Notion

Workspace and documentation

Competes when buyers can solve the pain with templates, checklists, and shared pages.

Whitespace

- A narrow workflow that reaches value without configuration-heavy onboarding.
- A buyer-facing proof artifact that shows time saved, risk reduced, or communication improved.
- A handoff path from manual concierge service to repeatable software.
- Use specificity as the wedge: one buyer, one workflow, one measurable result.
- Show proof earlier than broad competitors with before-and-after examples and small pilot data.
- Keep implementation lighter than incumbent suites or generic AI assistants.
- Own the specific buyer workflow instead of selling a broad AI assistant.

Positioning moves

- Lead with the exact buyer: Security or platform engineering lead at a company deploying autonomous agents against internal systems.
- Show a proof artifact for: Interview twenty security leads blocking agent deployments, then pilot the broker on one high-value workflow per design partner and measure whether it unblocks the rollout.
- Name the generic-assistant workaround directly and explain what it misses.
- Offer concierge setup before promising a full platform.

Public source

ServiceTitan

<https://www.servicetitan.com/>

Public source

Airtable

<https://www.airtable.com/>

Public source

HubSpot

<https://www.hubspot.com/>

Public source

Asana

<https://asana.com/>

Public source

Notion

<https://www.notion.com/>

Public source

Report source

<https://oauth.net/2/>

Public source

Report source

https://en.wikipedia.org/wiki/Access_control

Who's already moving in Business Ops

Public companies and funding signals the intelligence graph links to this vertical (related by keyword overlap — sized players, not direct competitors). Source: [/graph.json](#) .

FIELD SERVICE MANAGEMENT

\$625M

ServiceTitan

Operations software for contractors and field-service trades: scheduling, dispatch, quotes, jobs, and crew management.

IPO · 2024-12-12

RESTAURANT AND HOSPITALITY OPERATIONS

\$870M

Toast

Restaurant point-of-sale and hospitality operations including kitchen workflow, guest management, and food service.

IPO · 2021-09-22

Execution-readiness scorecard.

The score turns the report into bottlenecks, accelerators, and a dated first-month launch plan.

44/100

Research first

Access control platform built for AI agents scores 44/100 for execution readiness. The recommended next step is Interview twenty security leads blocking agent deployments, then pilot the broker on one high-value workflow per design partner and measure whether it unblocks the rollout.

Execution scorecard is generated from report validation, confidence, feasibility, founder fit, and difficulty.

Bottlenecks

- Okta, Microsoft, and cloud providers are all positioned to extend existing IAM to non-human identities.
- Standards for agent identity are unsettled; building on the wrong abstraction means rework.
- A broad AI assistant can flatten differentiation unless the wedge is painfully specific.
- The first release can become a generic dashboard if the job is not named tightly.
- Needs real buyer access, not only desk research.

First milestones

- 2026-08-12: Frame the wedge
- 2026-08-15: Interview 10 people who match the buyer persona.
- 2026-08-19: Ship a clickable demo or concierge workflow that produces the first useful artifact.
- 2026-08-26: Run one paid pilot or collect explicit pricing objections before automating the rest.

Value equation, matrix, and ACP.

Fit, roast, and kill criteria.

6/10

Founder fit

A solo or AI-assisted founder with direct access to Security or platform engineering lead at a company deploying autonomous agents against internal systems.

ADVANTAGES

- Can talk to the buyer before writing much code.
- Can ship a narrow first-win demo quickly.
- Can use local-first research artifacts to keep validation moving without a large team.

GAPS

- Needs real buyer access, not only desk research.
- Needs proof of budget or repeated urgency.
- Needs a crisp wedge before broad product work starts.

Roast

Promising enough to test, not strong enough to build broadly.

BLIND SPOTS

- Okta, Microsoft, and cloud providers are all positioned to extend existing IAM to non-human identities.
- A broad AI assistant can flatten differentiation unless the wedge is painfully specific.
- The first release can become a generic dashboard if the job is not named tightly.

HARD QUESTIONS

- Who wakes up already trying to solve this?
- What do they stop paying for or stop doing when this works?
- What proof would make a skeptical buyer trust it in one screen?
- What is the smallest paid version of this idea?

Kill criteria

- Fewer than five qualified buyers agree to discuss the workflow after targeted outreach.
- No buyer can name a current cost in time, money, risk, or reputation.
- The first demo does not produce a clear next step, paid pilot, or specific objection.

Next actions

- Write the one-sentence promise and test it in the strongest channel.
- Create the lead magnet and use it to recruit interviews.
- Build the smallest demo that proves the first win.

Move from reading to testing.

Local-first handoff cards copy prompts or structured data without requiring an account.

BUILD THIS IDEA

Copy the focused build brief for a coding agent.

COPY

ROAST

Copy the critique lens and blind spots before committing time.

COPY

LANDING PAGE

Copy a landing-page brief based on buyer, pain, and validation.

COPY

BRAND PACKAGE

Copy positioning inputs for naming, messaging, and design direction.

COPY

AD CREATIVES

Copy campaign angles for buyer-problem validation.

COPY

EXPORT DATA

Copy structured JSON for a research engine, roadmap tracker, or another agent.

COPY

FOUNDER FIT

Copy the founder-fit self-check before entering build mode.

COPY

Outreach template and interview script.

Built from this report's buyer, pain language, and channels. Personalize one detail per message — these are starting points, not spam ammunition.

Cold outreach message

QUESTION ABOUT ACCESS WORKFLOW

HOW ARE YOU HANDLING CORPORATE IAM ASSUMES HUMAN USERS WITH LOGINS AND SESSIONS;...

15 MINUTES ON A IDENTITY AND ACCESS MANAGEMENT FOR AI AGENTS WORKFLOW?

Hi {{firstName}},

I'm researching how security or platform engineering lead at a company deploying autonomous agents against internal systems handle this today: Corporate IAM assumes human users with logins and sessions; AI agents act continuously, impersonate service accounts, and inherit far more...

I'm not selling anything yet — I'm testing whether "Access control platform built for AI agents" is worth building, and I'd rather learn from people living the workflow than guess.

Would you trade 15 minutes for first access (and a say in what gets built) if it goes ahead?

{{yourName}}

COPY MESSAGE

Buyer interview script

1. Walk me through the last time this happened: Corporate IAM assumes human users with logins and sessions; AI agents act continuously, impersonate service accounts, a... What did you actually do?
2. What does that workaround cost you — in hours, money, or risk — in a normal month?
3. What have you already tried or bought to fix it, and why didn't it stick?
4. If "An identity broker that issues short-lived, task-scoped credentials to agents, maps every action to..." existed, what would have to be true for you to switch in the first week?
5. Who else feels this worse than you do — and would you introduce me?

WHERE TO SEND IT

- Community pain posts — Problem teardown, interview ask, and short demo clip
- Direct outreach — Concierge pilot offer with a manually prepared sample
- Searchable comparison content — Before-and-after page or alternatives memo for the exact workflow
- Reddit / forums — Post a problem teardown for Identity and access management for AI agents and ask how people solve it today.
- Launch communities — Ship a narrow demo and watch which promise gets clicks.

Build and review prompts.

Build prompt

Build a narrow MVP for "Access control platform built for AI agents" for Security or platform engineering lead at a company deploying autonomous agents against internal systems. Preserve the evidence, build only the first-win workflow, include source links, and treat Interview twenty security leads blocking agent deployments, then pilot the broker on one high-value workflow per design partner and measure whether it unblocks the rollout. as the first acceptance gate.

Review prompt

Review the "Access control platform built for AI agents" MVP for over-breadth, unsupported claims, weak buyer proof, privacy risk, and missing validation instrumentation. Do not approve expansion until the kill criteria and success metrics are measurable.

[documentation / oauth.net](#)

OAuth 2.0

The delegation standard agents currently stretch beyond its human-centric assumptions, motivating agent-native credentials.

[encyclopedia / en.wikipedia.org](#)

Access control - Wikipedia

Background on access-control models that all presume human principals rather than autonomous software actors.

If this exact wedge isn't yours, these are adjacent.

Derived deterministically from this report's buyers, vertical language, and business model.

Same problem, different buyer: Budget owner who feels the operational cost of the broken workflow.

The workflow pain in this report is not exclusive to security or platform engineering lead at a company deploying autonomous agents against internal systems. Budget owner who feels the operational cost of the broken workflow. faces the same friction with their own budget and urgency.

First test: Re-run day 3 of the sprint (15 outreach messages) against this buyer only, and compare reply rates before changing anything else.

Same workflow, adjacent vertical: pick the nearest regulated niche

No second vertical matched this report's language strongly, which usually means the wedge is horizontal. Horizontal wedges win by going vertical first.

First test: Pick the vertical where the pain costs the most per incident and rewrite the promise in its vocabulary.

Same wedge, alternate model: a productized service (fixed-price, done-for-you delivery)

This report monetizes via "Per-agent-identity monthly pricing with enterprise policy and compliance tiers.". Concierge delivery validates willingness to pay before any software exists and earns the workflow knowledge the product needs.

First test: Offer both versions on day 6 of the sprint and let the first pre-commitment choose the model.

Where this report sits in the intelligence graph.

Links from the ontology layer. Declared links are explicit in the research record; inferred links are keyword overlap and labeled as such. Full graph at /graph.json.

EVIDENCE INDEPENDENCE 90/100

5 source domains, 48 evidence edges. Dominant family: federalregister.gov. Audit all provenance .

Complaint evidence

- Privacy, trust, and data-control anxiety — keyword overlap (control, users)

— FULL NARRATIVE