A DATA-BASED RESPONSE · PENNSYLVANIA RIGHT-TO-KNOW LAW

The Missing Half of the Story

More appeals than ever. That much is true. It is the other half that decides the question.

Bath Borough attributes the rise in right-to-know appeals to artificial intelligence, and has asked the General Assembly to act on it. The Office of Open Records' own case records do not describe an AI problem. **They describe an agency problem.**

PREPARED FOR

Office of Senator Nick Miller

PENNSYLVANIA SENATE DISTRICT 14

By Michael Long · August 20, 2026
Northampton County, Pennsylvania

Every figure traces to the Pennsylvania Office of Open Records — its published annual reports, its public docket sheets, and the complete population of its decision and docket records: 44,784 rows, 2009 through August 19, 2026.

THE QUESTION BEFORE THIS OFFICE · WHAT WAS CLAIMED · WHAT THE RECORD SHOWS

Bath Borough says AI is driving right-to-know appeals. Three numbers say otherwise.

> *"Until Pennsylvania acknowledges that* **AI is causing more problems under the RTKL…"**
>
> Borough of Bath, October 15, 2025. The same page states the Borough "**suspected**" AI was being used, and names three citations that do not exist — real findings, conceded in Part III.[7]

35.6%

Of appeals decided on their merits in 2026 ended with the agency handing over the record

Up from 25.5% in 2022, and higher in **each of the last four years** while volume rose. If AI were flooding the system with weak appeals, this would be falling.[14]

Of records ordered released in 2026 came in appeals where the agency filed nothing at all

40%

Against 24.7% in 2022. It refused the record, the requester appealed, and it then made **no submission whatsoever**. That is a finding about refusals, not filings.[15]

603

Records ordered released by August 19 — more than any complete year in the OOR's history

The previous high was 569, in 2025, with four months still to run. **A wave of meritless appeals does not produce a record number of released records.**[14][15]

WHAT THE THREE ESTABLISH TOGETHER

Busiest year on record — and **the highest share of warranted appeals ever recorded.** A flood of frivolous filings would move all three the other way. **This is the signature of records being refused that should have been released.**

CONTENTS WHERE EACH FINDING IS SET OUT

Contents

Parts I and II are the statewide record. Part III is the specific exchange this document responds to. Part IV is the part addressed to the General Assembly.

OPENING The finding in brief

The problem is not that people are appealing. It is what they find when they do. 4

Ten findings from the Office of Open Records' own case records 5

PART I What happens after someone appeals

1 Three in four release orders say the agency could not justify the denial — or never tried 6

2 One in three: what the 2025 decision corpus shows 7

3 More than one in four appeals never reached the question of whether the record was public 8

4 Relief grew faster than the caseload — which is the opposite of a frivolity surge 9

5 2026 has already set a record — not for appeals filed, but for records released 10

PART II What the appeal count cannot tell you

6 More appeals does not mean more frivolous appeals 11

7 The increase came overwhelmingly from citizens and local-agency disputes 12

8 Every kind of local government saw more appeals — not one agency, not one dispute 13

9 The outcomes moved in the opposite direction from a "frivolous surge" 14

10 The most common agency position was that the record did not exist or was not theirs 15

PART III The record in Bath

11 The Borough identified a real weakness. The weakness was fixed. The fix became the new accusation. 16

12 A resident should not need the state's foremost open-records attorney on retainer to see a bank reconciliation 17

13 What these three things actually are — and what none of them is 18

14 Why an evidence system with this much control exists at all 19

15 Where the work entered the process — what an appeal count omits 20

16 The claimed cost was not allocated by request or by stage 21

17 Three agency court dockets do not fit inside one appeal count 22

18 A withdrawal can follow production rather than futility 23

19 AP 2025-1943 tested the premise on a different record 24

20 A missed deadline and continued review are procedural work 25

21 What the public narrative leaves out 26

PART IV For the legislature

22 An agency that denies a record it should have released pays nothing for being wrong 27

23 Who advises the agency may matter more than who files the request 28

24 Six questions the Commonwealth cannot currently answer about its own open-records system 29

25 Reforms that address the denial side, not only the filing side 30

APPENDICES Series, sources and method

Appendix A · Every year in the corpus, with both denominators 31

Appendix B · Primary sources used in this report 32

Appendix B continued · Standards and published literature 33

Appendix C · How every figure in this report was produced 34

Appendix D · Bath Borough statements being answered 35

EXECUTIVE SUMMARY WHAT THE NUMBERS ACTUALLY SHOW

The problem is not that people are appealing. It is what they find when they do.

Bath Borough is right that appeals rose in 2025. But a count of appeals cannot tell you whether those appeals were justified. The Office of Open Records' own case records can, and they point somewhere the Borough's account never goes.

35.6%	2 in 5	81.1%
Of 2026 appeals that anyone actually decided, the share that ended with the agency handing over records — up from 25.5% in 2022, higher in each of the last four years[14]	Of the 603 records OOR ordered released in 2026, the share where the agency **filed nothing at all** in the appeal — 40.0%, against 24.7% in 2022[15]	Of 2026 release orders rest on an agency failure — a share that has risen in every one of the last five years: 69.6, 73.3, 76.2, 77.7, 81.1[15]

Read those three together

Bath Borough is right that more people are appealing. What follows is the part its account never reaches: **a larger share of those appeals is succeeding, not a smaller one** — and in a growing number of them the agency does not defend the refusal at all. It refuses the record, the requester appeals, and the agency then files nothing.

All three move the same way, and a flood of meritless appeals would move all three the other way. If filing had become easier and the new appeals were weaker, the success rate would fall as volume climbed; it rose in each of the last four years. If the appeals were the problem, agencies would be defending their refusals and winning more of them; instead a rising share of them file nothing.

The third number answers the obvious objection to it. An order releasing a record will often say the agency failed its burden, because that is the legal test being applied — so the level alone proves little. **The legal test has not changed since 2022. The share has risen twelve points.** That movement is not a feature of how orders are written. It is a change in what agencies are doing before the order is written.

WHAT THE DATA SUPPORT — AND WHAT THEY DO NOT

These numbers do not identify a single cause for rising appeal volume, and this report does not claim one. What they establish is narrower and harder to answer: **the appeals being filed are getting more successful, not less** — which is the opposite of what the frivolity account predicts. Any explanation of the increase that examines only who is filing, and never how requests are being answered, is incomplete.

What this response does and does not claim

NOT CLAIMED

That every agency denial is improper, that every appeal has merit, or that AI has no effect on filings at all.

CLAIMED, AND SHOWN

That the number of appeals cannot establish any of the things it is being used to establish — cause, merit, or abuse.

THE QUESTION FOR THE LEGISLATURE

Not "why are more people appealing?" but **"why do so many appeals end with the record being released?"**

EXECUTIVE SUMMARY · CONTINUED THE FINDINGS, IN ORDER OF WEIGHT

Ten findings from the Office of Open Records' own case records

Each is stated with the figure it rests on and the source it comes from. The first three carry the argument; the rest are corroboration or context.

1. **The success rate is rising, not falling.** Of appeals that reached a decision, the share ending with records released went 25.5% (2022), 29.2%, 32.6%, 33.3%, **35.6%** (2026 to date) — up in each of the last four years, while volume rose throughout.[14]
2. **Agencies increasingly do not defend the refusal.** Of orders releasing a record, the share in which the agency filed *nothing at all* rose from 24.7% in 2022 to **40.0%** in 2026. Across the five years, 713 of 2,587 release orders record no agency submission.[15]
3. **Relief outgrew volume on every available measure.** From 2022 to 2025 appeals resolved rose 23.5% in OOR's case records while appeals ending in records released rose **48.2%**. On OOR's published count of appeals *filed*, volume rose 38.0% — still below the growth in relief. [3][5][14]
4. **2026 has already broken the record — for records released.** Through August 19, OOR had ordered **603** records released. The highest figure for any complete year is 569, set in 2025. Four months of 2026 remain.[14]
5. **Three in four release orders rest on an agency failure.** Of 2,587 orders releasing records between 2022 and August 2026, 1,311 state the agency failed to carry its burden of proof and 713 record that it filed nothing at all — 1,966 orders, 76.0%.[15]
6. **More than one in four appeals is never decided.** Across 2009–2026, 12,412 of 44,784 appeals — 27.7% — ended on a procedural defect with no ruling on whether the record was public.[14]
7. **Denials fell as volume rose.** Local-agency appeals grew 29% in 2025 while the denied share dropped from 28% to 20%.[3][4]
8. **The growth is ordinary people.** 702 of the 743 additional appeals — 94.5% — came from OOR's citizen category, not companies, media, or professional filers.[3][4]
9. **The most common agency position was "we don't have it."** In 839 appeals the agency asserted it did not possess or control the record, or that the item was not an agency record at all — the largest single category of the year.[3]
10. **Volume is not unprecedented.** Appeals rose 44.3% in 2011 and 45.1% in 2015, both larger than 2025, and OOR's caseload has fallen year over year four times since 2010.[5]

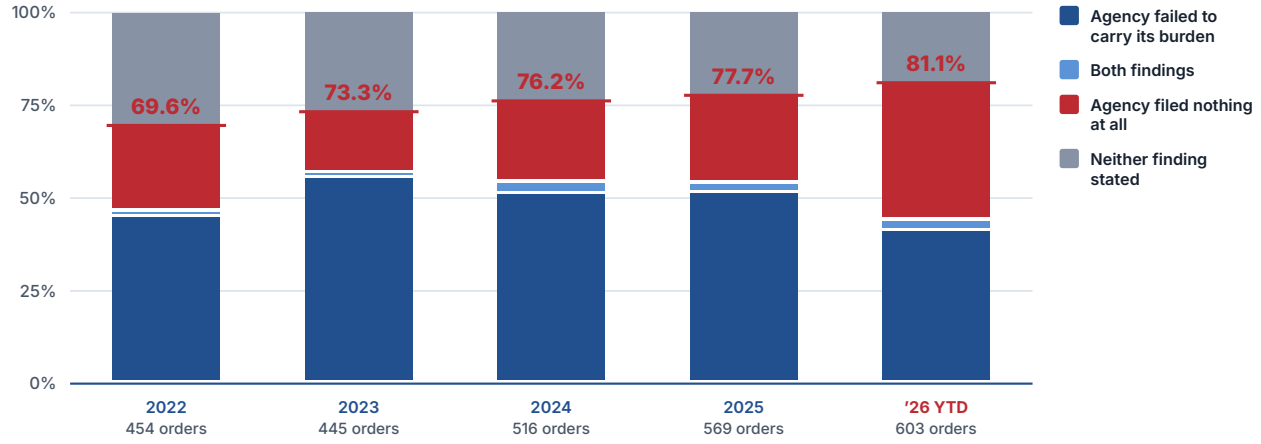
PART I · 1 THE CENTRAL FINDING

1 Three in four release orders say the agency could not justify the denial — or never tried

Every order OOR issued releasing a record between 2022 and August 19, 2026 was read and classified — 2,587 orders drawn from 16,483 appeals.[15] Under the Right-to-Know Law the **agency** bears the burden of proving that an exemption applies.[1] This is what the orders say about whether it carried that burden.

WHY OOR ORDERED RECORDS RELEASED

Every release order in the corpus, read and classified. The shaded portion of each bar is the share whose stated reason is a failure by the agency itself.



2,587 release orders reviewed across 16,483 appeals. 2026 is year-to-date through August 19.

1,311	713	76.0%
Release orders stating the agency failed to carry its burden of proof — 50.7%[1][15]	Release orders where the agency filed nothing at all on appeal — 27.6%[15]	Contain one or both findings — 1,966 of 2,587[15]

In three out of every four cases where a Pennsylvania agency was ordered to hand over a public record, the reason printed on the face of the decision is not that the requester found a clever argument. It is that **the agency could not justify its own denial, or did not show up to defend it.**

WHAT "FAILED TO CARRY ITS BURDEN" MEANS

The agency denied the record, the requester appealed, and the agency then could not produce the evidence or legal argument the statute requires of it. The record was releasable the whole time.

WHAT "FILED NOTHING" MEANS

The agency denied the record and then made *no submission whatsoever* to the appeals officer. In 2026 that has happened in **40%** of release orders — the highest share in the period.

THE WORKLOAD ARGUMENT TURNED AROUND

Agencies say appeals consume their resources. In 713 of these cases the agency spent no resources on the appeal at all — it simply never responded. **The appeal was not the burden. The denial was the decision that created it.**

PART I · 2 WHAT APPEALS ACTUALLY PRODUCE

2 One in three: what the 2025 decision corpus shows

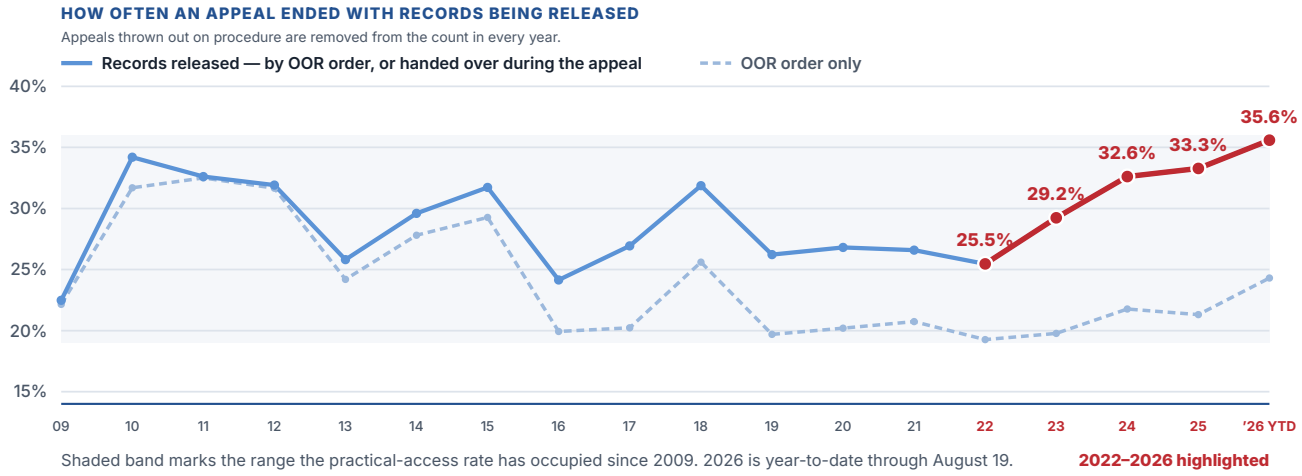
OOR's own case records is a decision/docket dataset, not the official count of appeals filed. For outcome analysis it allows a conservative screen that removes threshold dispositions — untimeliness, prematurity, jurisdiction, noncompliance, and insufficient appeals — leaving only matters that reached a substantive posture.

3,748 cases OOR closed in 2025 → **−1,077** thrown out on a technicality — nobody ruled → **2,671** left, which someone actually decided → **889** ended with records handed over

RECORDS HANDED OVER
An OOR order to release, or a docket entry showing the agency produced records during the appeal.

889 ÷ 2,671 = 33.3%

After that screen, **2,671** matters from 2025 remain. In **889** of them the docket shows either an OOR order releasing the record in whole or in part, or documented production or access during the appeal. That is **33.3%** — roughly one in three.



Applied to every year in the corpus, the same screen shows that this is not a 2025 anomaly. The release rate has stayed inside a **22%–36% band for eighteen years**, and it has risen in each of the last four. The full year-by-year series is set out in Appendix A.

PLAIN ENGLISH

Roughly one in three 2025 appeals that reached a decision produced a formal or documented access-related result. That is not a system malfunctioning — it is a review system changing outcomes.

WHY THE SCREEN MATTERS

Counting all 3,748 rows would understate the rate; counting only formal grants would understate it too. The screened denominator with both outcome types is the fairest measure available from public dockets.

STATED CONSERVATIVELY

"Practical access" does not mean every agency acted unlawfully, and it does not prove a record would have been permanently withheld without appeal. It means the docket shows access-related relief after the dispute reached OOR.

PART I · 3 THE APPEALS NOBODY DECIDES

3 More than one in four appeals never reached the question of whether the record was public

The Right-to-Know Law exists so that ordinary people can obtain public records quickly and without counsel. The corpus shows how often the process itself is what stops them.

12,412	27.7%	16.0%
Appeals thrown out on procedure, 2009–2026[14]	Of all 44,784 OOR case records[14]	Attributable to requester-side technical failure[14]

These matters ended on prematurity, untimeliness, jurisdiction, noncompliance, insufficiency under §1101, or improvident docketing. In none of them did anyone decide whether the record was public. For **2025 specifically**, 527 matters — **14.1%** of the year — can fairly be attributed to requester-side technical failure:

REQUESTER-SIDE TECHNICAL FAILURE, 2025	COUNT	WHAT IT MEANS IN PRACTICE
Untimely	247	The appeal was filed outside the window.
Noncompliance	142	A procedural requirement of the appeal was not met.
Insufficient under §1101	91	The appeal did not state what the law requires it to state.
Specificity, form and related	~47	Mostly how the underlying request was worded.

Reported openly because it is the honest counterweight: some appeals do fail for reasons that are the requester's own. The measurable figure is roughly one in seven, not the whole caseload.

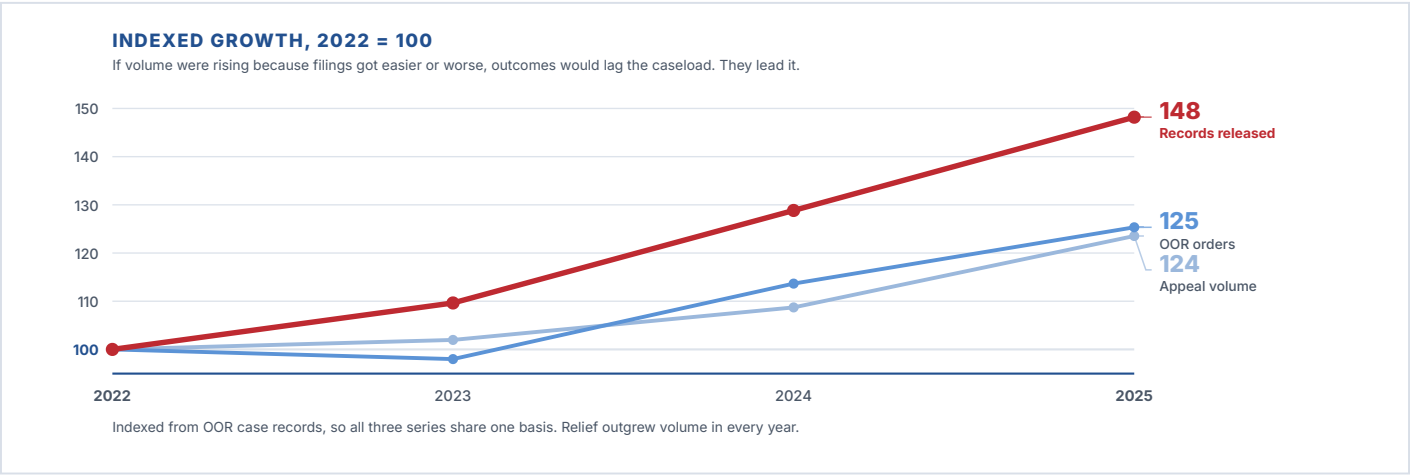
THIS IS THE PART THE VOLUME DEBATE NEVER REACHES

These 12,412 appeals ended without any ruling on the merits. The available data do not establish what happened next in any of them — whether the records were later produced, re-requested, taken to another forum, found not to exist, or never obtained at all. That is the point: **nobody knows, because nobody counts.** Thousands of appeals each year end not because a record was exempt, but because a deadline was missed, an appeal was filed too early, or it did not contain what the law requires.

PART I · 4 VOLUME VERSUS RELIEF

4 Relief grew faster than the caseload — which is the opposite of a frivolity surge

This is the test the raw-volume argument never runs. If appeals were multiplying because filing got easier or the filings got worse, the success rate should **fall** as volume climbs — the new appeals would dilute the meritorious ones. Relief has instead outgrown the caseload every year since 2022.



+23.5%	+48.2%	25.5% → 33.3%
Growth in corpus appeal rows, 2022–2025	Growth in appeal that ended in records being releaseds over the same period	Release rate after the threshold screen

The same-period comparison for 2026

Because 2026 is incomplete, the only fair comparison is the identical calendar window in each year — January 1 through August 19.

THROUGH AUGUST 19	APPEALS	ORDERED RELEASED	PRACTICAL ACCESS	RATE
2024	2,171	353	513	23.6%
2025	2,227	317	518	23.3%
2026	**3,308**	**603**	**883**	**26.7%**
Change '25→'26	**+48.5%**	**+90.2%**	**+70.5%**	**+3.4 pts**

Rate shown on the all-rows basis for period comparability. On the screened basis the 2026 year-to-date rate is 35.6%. Recently docketed 2026 matters may still be pending, which depresses rather than inflates these outcome counts.

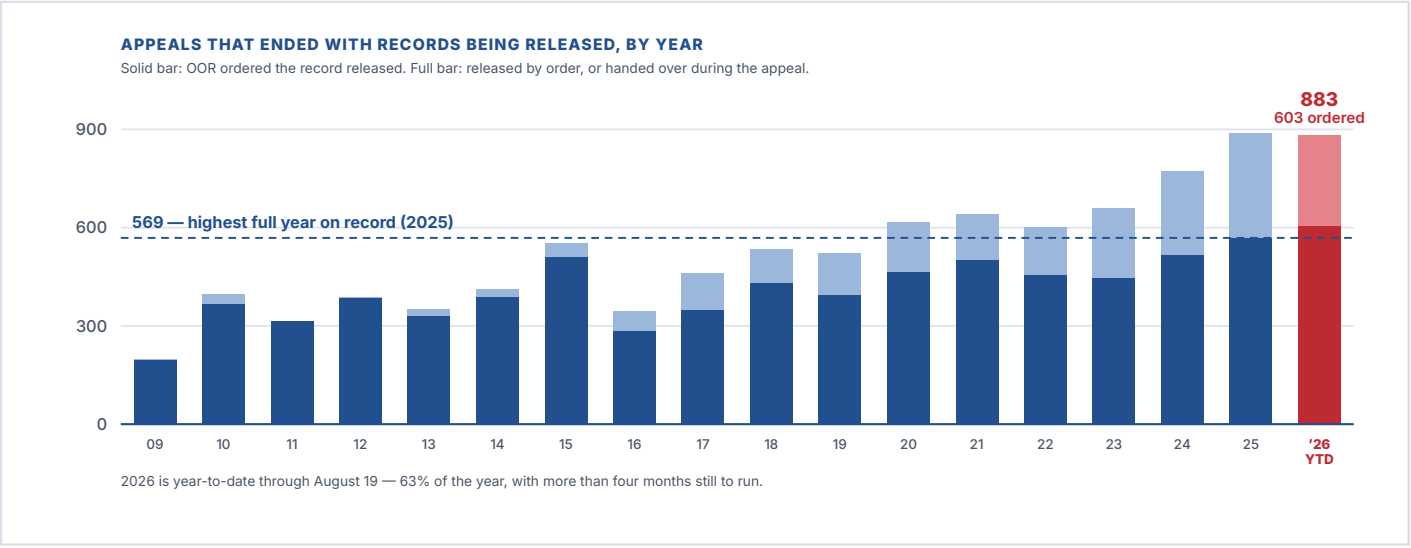
WHAT THIS DOES AND DOES NOT ESTABLISH

It does not prove that any particular denial was unlawful, and it does not establish intent. What it does establish is narrower and harder to answer: **agency responses are themselves a material driver of OOR's caseload.** An explanation of rising appeal volume that examines requester conduct while ignoring the rising share of appeals that produce records is incomplete.

PART I · 5 THE RECORD YEAR

5 2026 has already set a record — not for appeals filed, but for records released

If the rise in appeals were a rise in bad appeals, the number of records coming out the other end would flatten or fall. It has done the opposite, and 2026 is not close.



603	569	63%
Records ordered released, Jan 1 – Aug 19, 2026[14][15]	Previous high for a complete year, set in 2025[14]	Of 2026 elapsed when that record was passed

Counting appeals that ended with records handed over during the process as well, 2026 stands at **883** against 2025's full-year **889** — effectively tied, with more than four months still to run.

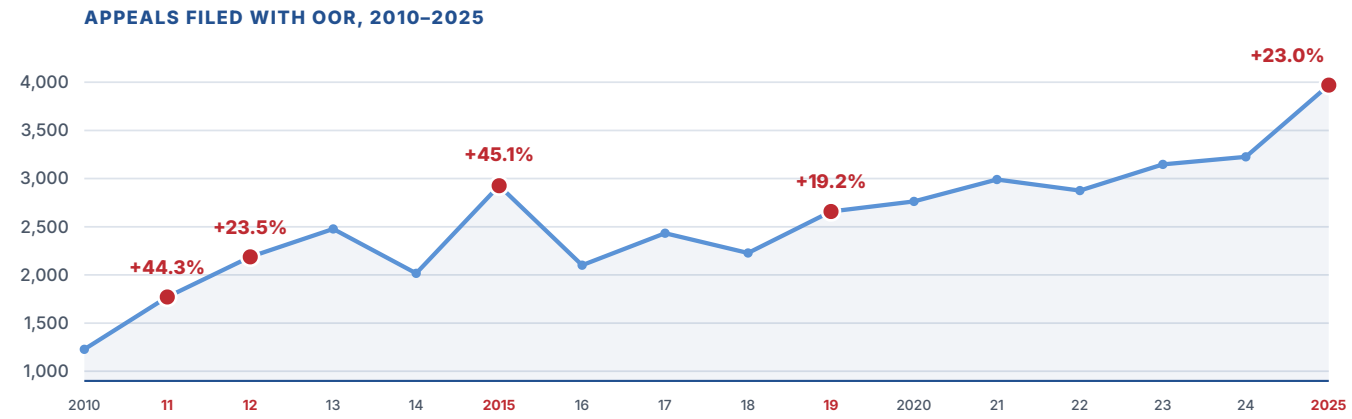
WHAT THIS RULES OUT

A wave of meritless appeals does not produce a record number of records. Whatever is driving people to the Office of Open Records, **it is ending more often than ever before with an agency releasing something it had refused to release**. That is the finding the volume figure conceals.

PART II · 6 GROWTH IN CONTEXT

6 More appeals does not mean more frivolous appeals

OOR reported 3,970 appeals in 2025, up 23.0% from 2024. That is a record number. It is not a historically unique growth rate.[3][4]



Official OOR annual-report totals. OOR's own comparable series begins at 2010; 2009 and early-2010 counts were reported on a different basis.

Pennsylvania saw **44.3% growth in 2011** (1,228 to 1,772), **23.5% in 2012**, and **45.1% in 2015** (2,017 to 2,926). Large swings in appeal volume existed long before ChatGPT or any modern generative AI tool was available to the public — and in the same period OOR's caseload *fell* year over year five times.

WHAT THE CHART PROVES — AND WHAT IT DOES NOT

"Appeals went up" cannot, by itself, establish what caused the increase; volume has moved this way before under conditions unrelated to AI. Equally, the chart does not show that AI had *no* effect in 2025. A causal claim in either direction needs evidence about the additional appeals themselves — who filed them, about what, and with what result.

RANK	YEAR	APPEALS FILED	YEAR-OVER-YEAR	WHAT WAS HAPPENING
1st	**2015**	2,926	**+45.1%**	The largest single-year increase in OOR's history.
2nd	**2011**	1,772	**+44.3%**	Early surge in the RTKL's first years of operation.
3rd	**2012**	2,188	**+23.5%**	A second consecutive double-digit increase.
4th	**2025**	3,970	**+23.0%**	A record *number* of appeals — the fourth-largest *rate* of increase.
5th	**2019**	2,658	**+19.2%**	Pre-pandemic growth, with no generative AI tools in public use.

Every figure above is taken from OOR's own annual reports. 2025 set a record in absolute volume while ranking behind 2015, 2011, and 2012 in rate of growth — three years no one attributes to AI.

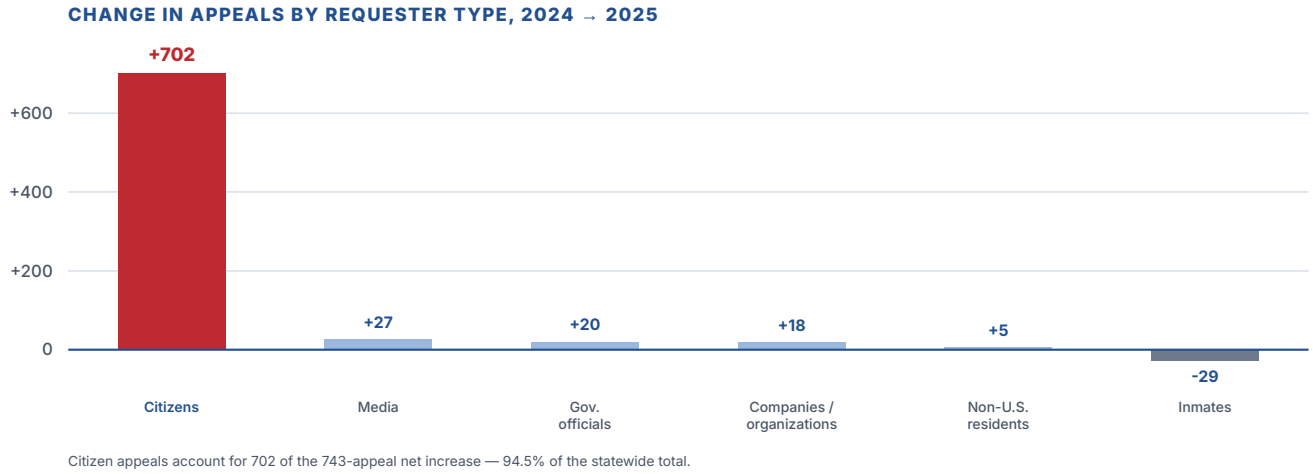
THE INFERENCE GAP

Bath's argument moves from a true premise (appeals rose) to a conclusion about cause (AI drove it) and a conclusion about quality (they are frivolous). Neither conclusion follows from the premise, and the sections that follow test both against OOR's own published data.

PART II · 7 WHO DROVE THE INCREASE

7 The increase came overwhelmingly from citizens and local-agency disputes

OOR's own requester categories show where the 743-appeal net increase came from. Citizen appeals increased by 702. Company and media filings changed only modestly, and inmate appeals fell.[3][4]



+702	94.5%	+723
Increase in citizen appeals	Share of the statewide net increase	Increase in local-agency appeals

Local-agency appeals rose from **2,490 to 3,213** — an increase of 723, or 29.0%. That single increase equals **97.3% of the statewide net rise**. The growth was not diffuse across agency types; it concentrated almost entirely in disputes with municipalities, authorities, school districts, and other local bodies.

WHAT THESE CATEGORIES CAN AND CANNOT TELL US

OOR does not report what tools each requester used. These categories therefore cannot establish how many filings involved AI. What they do show is that the growth was **broad public activity concentrated in local-agency disputes** — not, on its face, a small group automating frivolous filings.

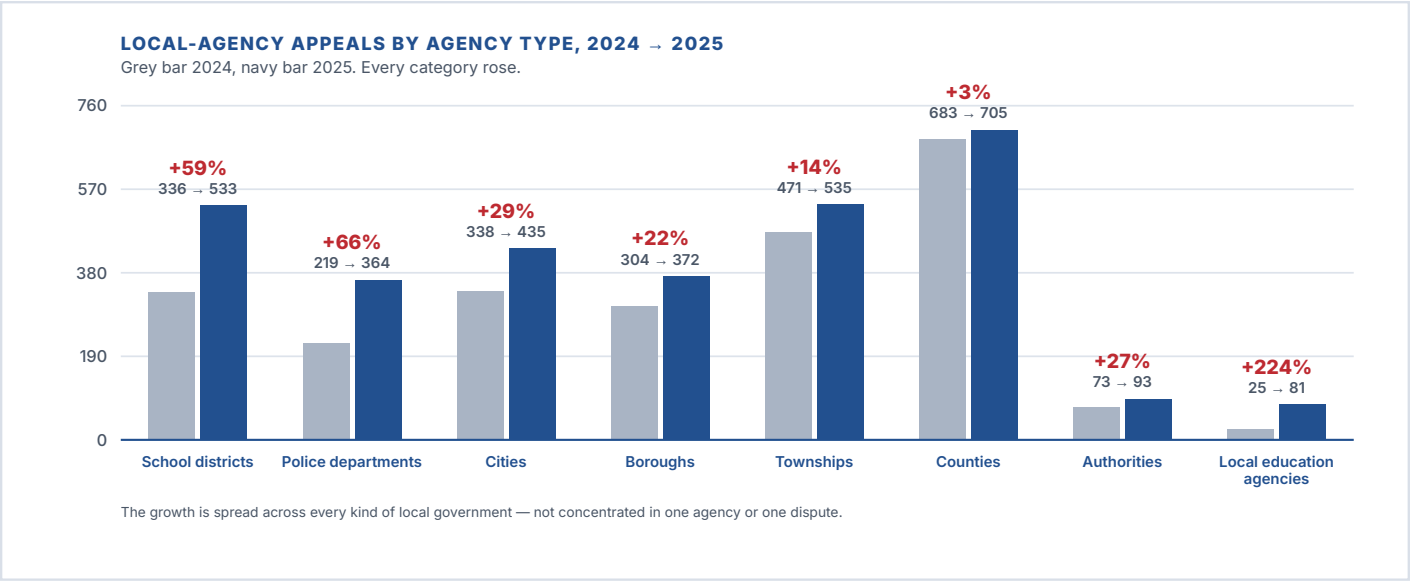
WHY CONCENTRATION MATTERS

If the increase were driven by an automation tool available to anyone, one would expect growth spread across requester types and agency levels. Instead it is concentrated in the one category — citizens appealing local agencies — where response quality is least standardized and where records officers most often serve multiple roles at once.

PART II · 8 WHERE THE INCREASE LANDED

8 Every kind of local government saw more appeals — not one agency, not one dispute

If a single prolific requester or a single controversy were driving the numbers, the growth would cluster. Broken out by agency type, it does the opposite.[3][4]



+197	***+145***	***+97***	***+68***
School districts, 336 → 533	Police departments, 219 → 364	Cities, 338 → 435	Boroughs, 304 → 372

Townships rose 64, counties 22, authorities 20, and local education agencies 56. Boroughs — the category Bath belongs to — account for **68 of the 743-appeal statewide increase**, or roughly **9%**. The remaining 91% happened somewhere else, to agencies with no connection to this dispute.

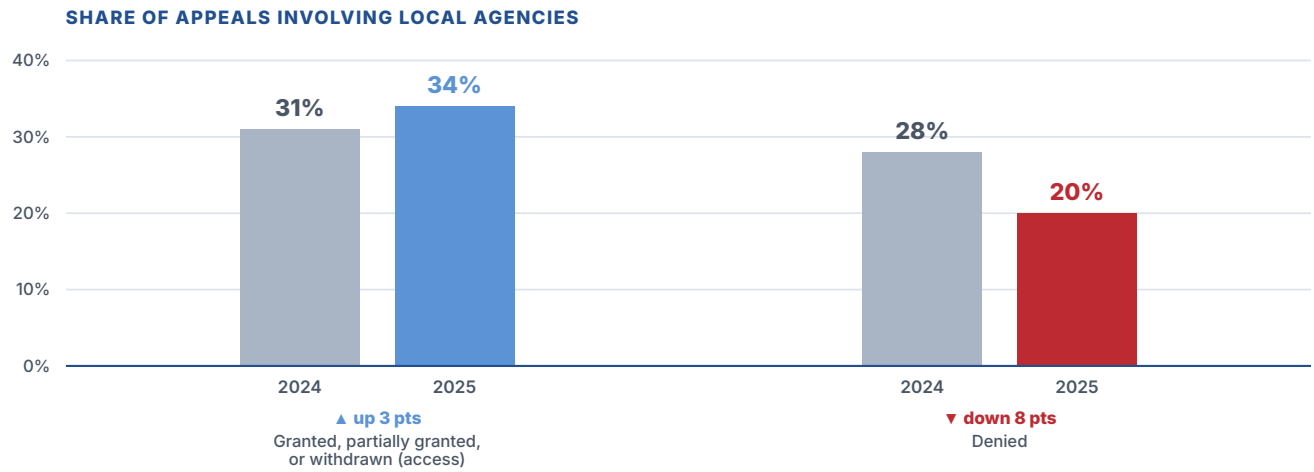
WHY THE DISTRIBUTION MATTERS

A statewide pattern this evenly spread across school districts, police departments, cities, boroughs, townships and counties is not well explained by any one requester, any one tool, or any one municipality's experience. It is the profile of **broad public demand meeting inconsistent agency response practice**.

PART II · 9 WHAT HAPPENED TO THOSE APPEALS

9 The outcomes moved in the opposite direction from a "frivolous surge"

If the 2025 increase were mainly a flood of meritless local-agency appeals, one reasonable expectation would be that favorable outcomes would be diluted. OOR's published percentages moved the other way.[3][4]



31% → 34%	28% → 20%	78%
Granted, partially granted, or withdrawn	Denied	Of completed OOR mediations successfully resolved in 2025

OOR itself explains that requesters obtain access when appeals are granted *or withdrawn*. That distinction matters here: a withdrawal frequently occurs **after the agency produces the records** during the appeal. Counting withdrawals as neutral, or as evidence that the appeal was unnecessary, inverts what usually happened.

WHY "WITHDRAWN" IS NOT NEUTRAL

An appeal is commonly withdrawn once the agency hands over what was requested. Treating every withdrawal as an abandoned or unnecessary appeal counts a requester's success as if it were a failure.

WHY MEDIATION BELONGS IN THE COUNT

Mediation resolves disputes without a final determination. A 78% resolution rate means a large block of matters ended in agreed access — work that never appears in a grant/deny tally.

WHAT A FRIVOLITY THESIS PREDICTS

If low-merit filings were driving volume, the denied share should rise and the favorable share should fall. Both moved the other way in the same year Bath describes as the problem.

THE POINT

This does not prove agency wrongdoing. It does show why raw appeal volume is a poor stand-in for "frivolous filing." The outcome data show substantial corrective value at the same time volume increased — the opposite of what a frivolity thesis predicts.

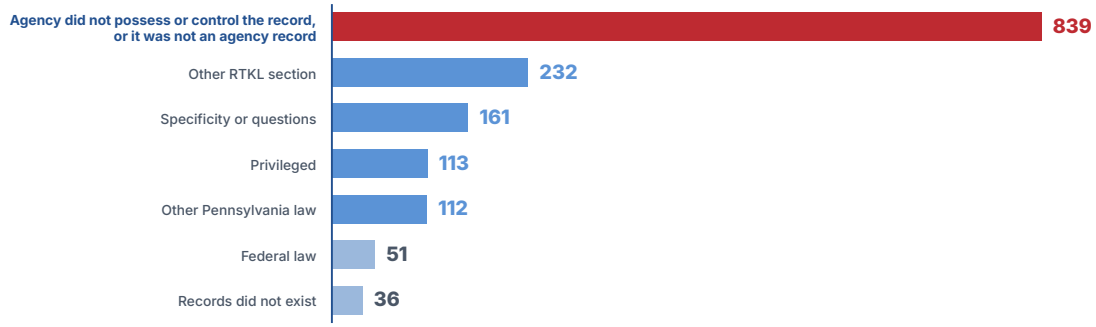
PART II · 10 WHAT AGENCIES ASSERTED

10 The most common agency position was that the record did not exist or was not theirs

OOR records the grounds agencies assert. In 2025 the single largest category was not a statutory exemption at all — it was a claim about possession, control, or whether the thing requested was an agency record.[3]

WHAT AGENCIES ASSERTED IN 2025 APPEALS

Grounds recorded by OOR across the year's appeals.



WHAT THAT CATEGORY INVOLVES

Whether a search was adequate. Whether a record exists. Whether the agency holds it. Whether the thing requested is a record at all. These are **factual questions about the agency's own recordkeeping** — and they were asserted 839 times in a single year.

WHY IT BELONGS IN A WORKLOAD DEBATE

Search adequacy, existence and possession disputes are resolved by the agency producing evidence about what it holds and how it looked. That work is generated by the response, not by the request — and it is a substantial share of what OOR spends its year deciding.

Add specificity (161), privilege (113), other Pennsylvania law (112), federal law (51), and the express "records did not exist" category (36), and the picture is of a docket dominated by **disputes about agency conduct and agency records** rather than by requesters raising exotic legal theories.

WHAT THIS CATEGORY MEANS FOR THE VOLUME ARGUMENT

A frivolity thesis predicts a docket full of requesters demanding records the law plainly exempts. The recorded grounds show something else: **the most common thing in dispute in 2025 was not whether a record was exempt, but whether the agency had it, held it, or looked for it properly.** Those are questions only the agency's own response can raise, and only the agency's own evidence can settle. They are counted here as appeals; they originate as responses.

PART III · 11 THE OBJECTION AND THE ANSWER

11 The Borough identified a real weakness. The weakness was fixed. The fix became the new accusation.

This is the part of the story the Borough's page does not narrate, and it is the part that matters most to anyone deciding public policy.

1

THE OBJECTION

Filings contained citations that were wrong — cases that did not exist as captioned, or real cases that did not stand for the stated proposition. That objection was legitimate.

2

THE RESPONSE

A source-linked corpus was built so that every authority travels with the actual final determination or court opinion, with later-treatment tracking and fail-closed publication. Its entire purpose is to make that error impossible to repeat.

3

THE NEW OBJECTION

The verification system is itself now offered as the problem, under the heading of AI — supported by material from 2022 to 2024 that predates it by years.

STATED PLAINLY, IN THE FIRST PERSON

Citations in my filings were wrong. Some cases did not exist as I captioned them; others were real but did not support the proposition I attached to them. I am not going to pretend otherwise, and the Borough was entitled to say so. The Office of Open Records identified **three** of them, in a single appeal — AP 2025-1943, September 2025.[13]

Those three errors among several hundred citations constitute a serious verification failure, not the whole character of the record. Public databases document numerous citation-error cases involving licensed attorneys, demonstrating that the problem is not confined to self-represented litigants; the peer-reviewed literature surveying those decisions treats citation verification as a profession-wide challenge.[18][19][20]

How it happened is ordinary and unflattering. A response is due in thirty days and I use all thirty. When the final version goes out there is often no time left to re-read it end to end, and errors introduced in earlier drafts survive into the filed copy. That is a process failure, and it is mine. The earlier workflow lacked a source-bound verification gate, which allowed a citation introduced during drafting to survive into the filed document without ever being checked against the primary authority. **That is the failure mode — not the technology as such, but the absence of a step that forces every proposition back to the document it comes from.**

THE TWO QUESTIONS HAVE TO BE JUDGED SEPARATELY

The earlier citation errors were legitimate grounds for criticism, and I do not ask anyone to set them aside. Whether the later source-verification process adequately addresses that problem is a different question, judged on the design and operation of the process itself — not settled by the errors that prompted it.

PART III · 12 WHY THE TOOL EXISTS AT ALL

12 A resident should not need the state's foremost open-records attorney on retainer to see a bank reconciliation

The Borough's page treats the existence of a research tool as the problem. It is worth asking what the alternative was.

The Right-to-Know Law was written so that a resident could obtain public records without hiring counsel. Very few attorneys practise this area of law, and those who do are expensive. The asymmetry in this dispute is a matter of public record: the Borough is represented by **J. Chadwick Schnee**, whose firm states that he "literally wrote 'the book' on the Right-to-Know Law," previously served as **Assistant Chief Counsel for the Office of Open Records**, authored four books on the RTKL and the Sunshine Act, and "written or edited over 10,000 final orders issued by the Office of Open Records."[17]

The requester is one person, self-represented, who has had to learn forensic accounting, software development and administrative law across four years. **On the underlying question — whether records had to be produced — the records were produced, the searches were found inadequate, or the Borough's position did not survive.**

WHAT A VERIFICATION TOOL IS FOR

It indexes only the places where authority is actually written — OOR final determinations and court opinions — and updates daily as new ones issue. Nothing enters the legal-rule results until it has been reviewed against its source, and superseded holdings stay visible rather than passing as current law. The point is to make the drafting **rest on retrieved primary documents instead of recollection**, which is precisely the defect that produced the citations the Borough complains about.

It exists because the public interface to Pennsylvania's open-records law is hard to search and easy to get wrong, and because **a resident should not need a former Assistant Chief Counsel of the OOR on retainer to obtain a bank reconciliation.**

What the alternatives actually were

HIRE COUNSEL

Few attorneys practise Right-to-Know work, and those who do are expensive. There is no contingency market for a records request.

LEGAL AID

The Pennsylvania Legal Aid Network intake process returned no referral options for municipal legal needs in this county.

PROFESSIONAL RESEARCH

Westlaw and LexisNexis sit behind attorney credentials or substantial annual commitments. Neither was realistically accessible on a private citizen's budget.

THE ASYMMETRY WORTH PUTTING TO THE LEGISLATURE

The professional legal-research platforms have incorporated artificial intelligence for decades, and their generative tools are marketed to the credentialed profession today. Counsel opposing a self-represented requester has access to all of it. **A proposal that restricts AI-assisted right-to-know filings does not restrict artificial intelligence in this dispute. It restricts it on one side of it** — the side without the credential, the budget, or the institutional subscription.

PART III · 13 THE PRODUCT BATH LABELED "AI LEGAL-RESEARCH SOFTWARE"

13 What these three things actually are — and what none of them is

Bath's headline calls the product "AI Legal-Research Software." **RTK Corpus is not an artificial-intelligence model and was not trained on anything.** It is a curated, source-linked collection and index of published legal documents. PA RTK Wizard is the research interface that searches it. MIKE OS is a separate evidence-investigation and case-management system. None of the three is a trained AI model, and the legislative argument the Borough's page builds rests on the other characterisation.[6][16]

COMPONENT	WHAT IT IS	WHAT IT IS NOT
RTK Corpus	A collection and index of published legal documents — OOR final determinations and Pennsylvania court opinions — each stored with its original source document, its version, date and a checksum.	A trained AI model. Nothing is trained. Documents are collected, indexed, and linked back to their sources.
PA RTK Wizard	The research and verification interface over that collection. Its published terms: "screen and route only," and it "does not decide appeal readiness."	An autonomous appeals engine. Per the same terms, "nothing is filed, sent, or delivered for you," and any transfer to an AI provider requires separate, submission-specific consent.
MIKE OS	A local-first evidence-investigation and case-management system built for the Bath investigation. It was not built from Bath data. Bath materials are handled as source evidence. The system records source identity, hashes, custody and processing events, derivative relationships, Bates mappings, accounting workpapers, review findings and release artifacts.	An AI model, or a system trained on Borough records. It does not independently determine legal conclusions or autonomously file, send or deliver anything.

Descriptions of what each product is are made from the author's own knowledge. Quoted terms are from the public product page.[16] That page also states the tool is "an independent research tool… not affiliated with the Commonwealth of Pennsylvania, the Pennsylvania Office of Open Records, a court, or any government agency," and its acceptable-use terms forbid using it to "automate abusive request volume."

The site's own terms are the answer to the citation-error concern. Authorities carry later-treatment tracking — *"reversed, limited, distinguished, and superseded holdings stay visible, but cannot masquerade as current law"* — and research is fail-closed: *"candidate passages stay out of legal-rule results until reviewed."* The user, not the tool, *"remains solely responsible for the truth, completeness, legality, scope, citations"* of every output.

WHY THE CORPUS WAS BUILT

Earlier research and citation problems showed exactly why source verification, provenance, later-treatment checks, and fail-closed publication were necessary. Its purpose is to make that error materially harder to repeat — not easier to scale. No research system operated by a human being can promise more than that.

WHY THE DISTINCTION MATTERS HERE

"AI can produce bad filings" is true. "A collection of published legal documents is an AI appeals engine" is a different claim, and it is not true. The headline collapses two different things into one label, and **a legislative proposal built on that label inherits the error.**

PART III · 14 STANDARDS AND CONTROL FRAMEWORK

14 Why an evidence system with this much control exists at all

This section answers a fair question: why would a private citizen build something this elaborate to make records requests? He would not, and did not. **The records requests are not the project.** They are how evidence is obtained for a financial investigation into the Borough's accounts — and findings of that kind are only worth making if they are built to a standard that survives examination by someone hostile to them.

THE REASON FOR THE STANDARD

An allegation about public money is worth nothing unless every document behind it can be traced to where it came from, shown to be unaltered, and tied to the analysis that relies on it. That requirement — not the volume of requests — is what MIKE OS was built to meet.

Standards and guidance used as design criteria

CHAIN OF CUSTODY AND AUTHENTICATION

Federal Rules of Evidence 901 and 902; Department of Justice and National Institute of Justice digital-evidence handling guidance.[21][22][23]

EVIDENCE PROCESSING AND PRODUCTION

DOJ conventions for producing electronically stored information — document-family relationships, Bates mapping, hash values and load-file controls — with NIST and SWGDE publications as supporting technical guidance.[24][25][26]

GOVERNMENTAL AND FORENSIC ACCOUNTING

Pennsylvania DCED guidance for municipal accounting; GASB; the GAO Yellow Book and Green Book; AICPA Statement on Standards for Forensic Services No. 1; and ACFE references.[27][28][29][30][31][32]

REPORTING AND REVIEW

Findings structured as source, criteria, condition and effect, with stated limitations, contrary evidence recorded rather than discarded, independent review before release, and fail-closed release controls.

STATED PRECISELY, BECAUSE PRECISION IS THE POINT

MIKE OS is an evidence system built for the Bath investigation and operated under a documented control set. **When a package reaches release-ready status under those controls,** original and derivative files are separately preserved and hash-verified; processing and custody events are recorded; page relationships are reconciled; unresolved gaps block release; and independent review is required. **Bath records are evidence processed under those controls — not training data — and MIKE OS is not an artificial-intelligence model.** The resulting evidence packages are structured to support authentication and qualified-person certification under Federal Rules of Evidence 902(13) and 902(14).

WHAT IS NOT CLAIMED

These authorities are **design and validation criteria**. They are not presented as certification by the Department of Justice, approval by any court, ISO accreditation, or endorsement by any of the bodies named. No government certification of this software is claimed, and none exists. The claim is narrower and checkable: **the controls were designed with reference to the rules, standards, guidance, and professional references identified here**, and the outputs are structured so that a qualified person can certify them if and when a proceeding requires it.

The substance of the financial investigation is outside the scope of this report and is not set out here. It is described only to explain why the evidence controls exist.

PART III · 15 THE COST CHAIN

15 Where the work entered the process — what an appeal count omits

Bath's appeal count shows that matters were filed. It does not identify which actor, decision, or procedural stage generated the resulting work.

PURPOSE AND LIMIT

These are **selected Bath matters, not a statewide sample.** They do not show that every requester position was correct or every appeal was granted. They show why a filing count cannot allocate workload without examining the request, the initial response, the appeal record, production, the determination, and later litigation. **The docket records neither whether artificial intelligence was used nor whether any tool increased the work required.**

MATTER	DOCUMENTED WORK AFTER THE INITIAL RESPONSE	LIMIT
AP 2023-1083	Second invoice review, less-redacted production, and 23 exemption logs.	OOR upheld the remaining redactions.
AP 2022-2675 / AP 2024-0001	Production during appeal or mediation, and settlement activity.	Settlement was not a merits ruling.
AP 2025-0522 / AP 2025-1943	Broad request denied; a narrower follow-up produced a different Section 705 ruling and later reports.	Page totals are not equivalent populations.
AP 2023-1598	Reconsideration, three Borough court dockets, and later motion practice by both parties.	Expense cannot be assigned entirely to either side.
AP 2026-2665	Missed agency deadline, an additional participation opportunity, and supplemental filings.	The merits remain pending.

THE CENTRAL POINT

Six of the seven selected matters involve financial or accounting records. Each contains **documented workload-producing events at several stages.** Counting requester appeals alone cannot allocate that work among the requester, the Borough, OOR, counsel, mediation, or court review. That is a statement about cost *allocation* — not a claim that every appeal was necessary or every expense Borough-caused.

PART III · 16 COST ACCOUNTING

16 The claimed cost was not allocated by request or by stage

AP 2023-1083 is the clearest reason not to infer causation from a total legal-expense figure.

AP 2023-1083 sought invoices, payment information, and hours attributable to individual RTKL requests for 2022 through March 30, 2023. **The original invoices disclosed dates, timekeepers, hours, rates, amounts, and invoice totals.** But the narrative descriptions within general-solicitor matter 00001 were broadly removed, and that matter combined RTKL and non-RTKL work. The appeal record described *estimated* cost figures rather than a contemporaneous allocation by individual request, and acknowledged that mixed billing could produce minor overstatements.[11]

After OOR required a sworn, invoice-by-invoice list identifying each redaction and its asserted legal basis, counsel conducted a second line-by-line review, supplied less-redacted invoices, and prepared **23 exemption logs**. Numerous restored descriptions expressly identified RTKL requests, OOR proceedings, and work on a prior Long appeal. Other descriptions remained redacted.

WHAT OOR ULTIMATELY HELD

The remaining privilege and work-product redactions were supported. A per-request breakdown that did not already exist did not have to be created under Section 705.

WHAT CHANGED DURING THE APPEAL

OOR's directive prompted another review, restored descriptions, a less-redacted production, an affidavit, and 23 invoice-specific logs.

THE DEFENSIBLE COST CONCLUSION

The record does not establish improper billing, bad faith, or improper final redactions, and it does not assign every charge to agency error. It establishes that **the available billing did not allocate cost by request, issue, party, outcome, or procedural stage.** It therefore cannot show how much expense resulted from requester conduct, initial agency review, counsel's approach, changed production during appeal, or later litigation.

The case file does not reproduce the precise earlier public statement that prompted the request. No exact quotation or amount is attributed to that statement here.

PART III · 17 COURT WORKLOAD

17 Three agency court dockets do not fit inside one appeal count

AP 2023-1598 is included only to show why post-OOR workload must be separated by actor and stage.

In the controlling OOR determination **upon reconsideration**, OOR found the Borough's interpretation reasonable, accepted its supplemented search evidence, upheld its attorney-client-privilege showing, and **declined to find bad faith**. OOR nevertheless required the Borough to re-review withheld communications and attachments for purely factual material that could be segregated without revealing privileged communications.[12]

The Borough initiated **three** Court of Common Pleas dockets. Two were dismissed as **moot** — one challenged a superseded determination, the other a nonfinal reconsideration action. The remaining Borough appeal reached the merits and was denied.

> **WHAT THE COURT DID — AND DID NOT — HOLD**
>
> The court held that purely factual information must be extracted where separation will not reveal privileged communications. **It did not adopt the requester's broader language concerning "application of law to fact" or "legal strategy,"** and it expressly allowed that a proper re-review might produce nothing.

On July 3, 2025, Mr Flynn stated that the Borough had re-reviewed 20 privileged emails and found no segregable factual material to release. **That is the Borough's compliance position, not a later judicial finding.** The subsequent docket also contains unsuccessful requester motions and a $1,000 sanctions order.

THE COST-ALLOCATION LESSON

This matter does not establish that all expense was agency-caused or requester-caused. It establishes that **an OOR appeal count cannot allocate the work generated by reconsideration, three agency-initiated court dockets, and later motion practice by both parties.**

PART III · 18 PRODUCTION AND MEDIATION

18 A withdrawal can follow production rather than futility

AP 2022-2675 and AP 2024-0001 document work and access after the initial agency response.

AP 2022-2675. Responsive records were produced during the appeal, and OOR addressed financial-record exemption positions the Borough had not fully substantiated. That does not convert every issue into a requester victory. It establishes that the original response did not end the work or the access question.[10]

AP 2024-0001. The matter ended through mediation, settlement production, and withdrawal. The Borough's own page reports 1,027 digital pages plus 93 additional pages released through settlement. A withdrawal therefore cannot be treated automatically as an appeal that accomplished nothing — and the settlement was **not** a merits determination that the earlier response was unlawful.[9]

SA-J required a source correction

The executed settlement and remittal required QuickBooks Audit Trail reports for **nine** specified General Fund lines covering January 1, 2019 through December 31, 2022. A complete combined copy is **120 pages**. The 60-page combined copy in the canonical source assignment is incomplete: it contains seven complete reports, only the first two pages of the eighth, and no ninth report.

WHAT IS NOT CLAIMED

The nine SA-J accounts do not overlap the 31 General Fund accounts named in AP 2025-1943, but different years, filters and report boundaries still prevent mechanical page addition. The 60-page copy is not cited as the complete nine-report production.

WHY IT MATTERS

Production during appeal or mediation is part of the workload and part of the outcome. A scoreboard recording only "appeal filed" or "withdrawn" omits that part of the process.

WHY THESE TWO MATTERS ARE IN THE REPORT AT ALL

Neither is offered as a win. Both are offered because a count of appeals filed records the same single event in each — one appeal — while the dockets record production during appeal in one and mediation, settlement production and more than 1,100 released pages in the other. **The disposition label and the work are two different quantities, and only one of them is published.** A legislative premise built on the label alone cannot see the other.

PART III · 19 THE NARROWER FOLLOW-UP

19 AP 2025-1943 tested the premise on a different record

The two QuickBooks determinations are related, but they are not total opposites and should not be described as such.

AP 2025-0522 was denied. OOR accepted Bath's General Fund interpretation, accepted the PDF production and the evidence concerning failed Excel export and inactive-account reports, and declined to find bad faith. But OOR also recognised that drawing existing information from a database is **not automatically the creation of a new record**. Footnote 6 expressly left open a separate request for smaller Excel exports within QuickBooks's capabilities.

AP 2025-1943 followed that path. The later request named particular accounts, requested annual reports, and permitted segmentation. Bath renewed a related Section 705 theory, arguing that selecting built-in report parameters would create new records. OOR accepted that pre-generated custom reports were not already stored, but rejected that legal theory on the later record: reports generated through QuickBooks's built-in functions were **extraction from an existing database**, and OOR ordered electronic production. OOR again did not require Excel and again declined to find bad faith.[13]

> **THE CITATION ISSUE**
>
> The inaccurate citations identified in AP 2025-1943 are conceded and are not relied upon in this report. That advocacy defect and the separate Section 705 merits ruling both belong in the record.

Why 677 and 699 are not a completeness equation

The later response listed six annual General Fund reports totalling **677 pages** for 2018–2023, produced as responsive to a request naming 31 accounts, plus **49 pages** for 2024. The earlier response described a 699-page report for active accounts 100.00–499.99 beginning January 18, 2021. Different date spans, account populations, filters, report classes and pagination make the totals non-equivalent. **They warrant controlled reconciliation; they do not prove that identified transactions were withheld.**

WHAT THE PAIR OF DETERMINATIONS ACTUALLY SHOWS

The same agency position was tested twice on two different records and came out two different ways. That is what a functioning review process looks like — and it is the reason a second, narrower appeal was filed rather than evidence that the first one should not have been. **Counting the two filings as two units of requester-generated cost treats the refinement of a request as a burden rather than as the mechanism the statute provides.**

PART III · 20 THE PENDING APPEAL

20 A missed deadline and continued review are procedural work

AP 2026-2665 remains pending. It is used only to identify where work entered after the original response.

The Borough invoked a thirty-day extension that included legal review and issued its final response on June 19. It then missed OOR's original **July 16** appeal-submission deadline. At **8:50 a.m. on July 17**, OOR notified the Borough that no agency submission had been received and supplied an additional opportunity to participate. Borough counsel filed a position statement and a Flynn attestation at **9:16 a.m.** The Borough later filed an additional packet whose matrix continued to mark categories "In Review," "Legal Review Needed," or "Legal/Bookkeeping Review Needed."[8]

NO MERITS INFERENCE

The extension, the missed deadline, the later filings and the review labels do not decide whether the Borough's underlying positions were lawful. They show additional procedural work after the statutory extension and final response. **The Borough did file a defence; it did so after OOR's notice and additional opportunity.**

The cost question the public narrative does not answer

Across the selected matters, work entered through initial interpretation and search, production after filing, invoice re-review, exemption logs, mediation, reconsideration, supplemental submissions, a missed agency deadline, court review initiated by the Borough, and motion practice by both parties. **None of those events can be priced reliably from an appeal count.**

RECOMMENDATION TO THE GENERAL ASSEMBLY

Before using appeal volume or requester technology as a legislative premise, require statewide, stage-specific evidence: initial disposition, records produced after filing, supplemental work attributable to each party, mediation, reconsideration, court review, initiating party, outcome, cost, and **whether AI use was actually recorded and causally connected to the workload.** Without those distinctions, the count measures filings — not cause or responsibility.

PART III · 21 THE PUBLIC NARRATIVE

21 What the public narrative leaves out

The problem is not that Bath states facts. Many of the facts it cites are real. The problem is **where the framing repeatedly stops**.

THE EMPHASIS	THE MISSING CONTEXT
"Appeals cost the Borough money."	So do agency reconsideration petitions, continued withholding positions, and court appeals. AP 2023-1598 is an example of the Borough itself carrying the dispute into court.
"AI and bad citations create workload."	The 2025 statewide surge was overwhelmingly citizen and local-agency growth, and large percentage surges occurred long before generative AI existed.
"The requester filed thousands of pages."	In AP 2026-2665, the principal post-July-17 set was 97.82% legal authorities or source evidence; only about 2% was authored merits, verification, or schedule material.
"Bath provided records."	Selected Bath appeals show records produced *during* appeal, additional search relief, partial grants, and more than 1,100 pages released through settlement.
"The appeal is simple."	AP 2026-2665 involves nine years of financial and accounting records, paper-versus-electronic medium, fees, exemptions, search adequacy, metadata, and an August 4 status spreadsheet still showing categories under review.

THE PATTERN ACROSS ALL FIVE

Each statement is defensible on its own terms. Each also stops at the point where the next fact would complicate it. That is a framing choice, not a factual error — which is why it has to be answered with the omitted half rather than a denial.

WHY IT MATTERS FOR POLICY

A statewide legislative argument built on a one-sided ledger will produce reforms aimed at requesters alone, leaving the response-quality problem that generates the appeals entirely untouched.

A FAIRER WAY TO STATE THE COST QUESTION

The public should be told not only what appeals cost, but what the underlying agency decisions were, what changed after review, what records were produced, and which side chose to keep litigating. Anything less is a one-sided ledger.

PART IV · 22 WHY THE PATTERN PERSISTS

22 An agency that denies a record it should have released pays nothing for being wrong

The preceding sections describe what happens. This section offers an explanation for *why* it keeps happening. It is presented as an argument about incentives, not as a finding from the data.

What the law actually does when a denial is overturned

The Office of Open Records decides roughly 3,700 appeals a year.[3] When it finds that a record was improperly withheld, it orders the record released. That is the entire remedy. **OOR cannot award the requester costs, cannot impose a penalty, and cannot require the agency to explain the denial.** Attorney fees and civil penalties under the Right-to-Know Law are available only from a court, and only where the requester can prove the agency acted in bad faith or that its conduct was willful or wanton — a demanding standard, pursued at the requester's own expense.[2]

WHAT A WRONGFUL DENIAL COSTS THE AGENCY

- The record is released — months later
- Its own legal fees, largely incurred anyway
- No penalty, no fee award, no finding on the record

WHAT IT COSTS THE REQUESTER

- Months or years of delay
- Their own time, and counsel if they can afford it
- Nothing at all, if they cannot navigate the appeal

THE ARGUMENT, STATED PLAINLY

Denying costs the agency nothing. The worst that happens is that it hands the record over later. An agency uncertain whether a record must go out faces no downside in withholding it and letting an appeals officer decide months from now. **None of this requires anyone to act in bad faith. The existing structure simply imposes fewer immediate consequences for an unsupported denial than for disclosure, regardless of anyone's intent — and ordinary caution responding to ordinary incentives produces exactly the pattern in the data.**

Why this reading fits the numbers better than the alternative

IF THE PROBLEM WERE BAD APPEALS…	…BUT WHAT THE RECORD SHOWS
The denial rate would rise as volume rose	It fell from 28% to 20% in the year volume rose 29%.
Fewer appeals would end in release	The share ending in release rose every year from 2022 to 2026.
Growth would concentrate among heavy filers	The ten largest filers held 11.4% of the docket in 2024 and 11.6% in 2025.
Agencies would be defending exemptions	Their most common position, 839 times, was that they did not have the record.

WHAT WOULD CHANGE THE INCENTIVE

Not a restriction on who may appeal. Something closer to: **publish agency-level overturn rates**, so a pattern of denials that do not survive review is visible; and give OOR the ability to attach a consequence when a denial is reversed, rather than leaving the only remedy to a court case most requesters cannot afford to bring.

PART IV · 23 A QUESTION THE COMMONWEALTH CAN ANSWER WITH ITS OWN DATA

23 Who advises the agency may matter more than who files the request

The preceding section describes an incentive that applies to every agency. This one describes a difference between them — and it is the single most testable proposition in this report.

There is a second asymmetry in how agencies are advised, and it is arithmetic rather than motive. Most Commonwealth agencies are advised by salaried in-house counsel. Most municipalities are advised by a contracted solicitor billing by the hour.

SALARIED IN-HOUSE COUNSEL

A denial that is appealed and reversed costs the agency nothing additional. The salary is fixed whether the record goes out on day one or after a year of appeal.

CONTRACTED SOLICITOR

The same denial generates a response, an appeal, briefing and any subsequent court filing — each separately billable. The cost is borne by the municipality, not by the adviser recommending the position.

THE QUESTION, STATED AS A QUESTION

This is a hypothesis, not a finding, and nothing in this report establishes it. It is also testable with data the Commonwealth already holds: **do municipalities that rely on contracted outside solicitors show higher outright denial rates, higher appeal rates and higher total legal expenditure than comparable municipalities with salaried in-house counsel, controlling for request volume and record complexity?** OOR already publishes appeals by agency type, and the Legislative Budget and Finance Committee has already applied a costing methodology to this statute once. Nobody has asked the question.

If the answer is yes, some meaningful share of what is described as citizen abuse of the Right-to-Know Law is not citizen behaviour at all. It is a fee structure that makes denial the path of least resistance for whoever is advising the agency — whether or not the denial is defensible, and without anyone intending harm.

WHY THIS QUESTION BELONGS IN FRONT OF THE GENERAL ASSEMBLY

It costs nothing to answer. The appeal data are already published by agency type; municipal legal expenditure is already reported; and the Legislative Budget and Finance Committee has already built a costing model for this statute. **Every input exists, and no new obligation on any requester or any agency is required to run the test.** If the hypothesis is wrong, the record will show it and the question is closed. If it is right, it identifies a driver of cost that no restriction on filing would reach.

PART IV · 24 WHAT TO MEASURE FIRST

24 Six questions the Commonwealth cannot currently answer about its own open-records system

Every figure in Part I of this report had to be derived by reading case records one at a time, because the Office of Open Records does not publish the outcome data in a form that answers these questions. None of the six requires new authority. Each requires only that an existing disposition be recorded and reported.

1 When an agency loses, what did it actually get wrong?

OOR reports how many appeals were granted. It does not report how many of those orders rest on a finding that the agency produced *no evidence at all*. That number had to be counted by hand for this report: **713 orders in five years** in which the agency filed nothing on appeal.[15] A distinction that basic should not depend on a private citizen counting it.

2 How often are records handed over only because an appeal was filed?

In the published statistics, a withdrawal after the agency produces the record is indistinguishable from a withdrawal because the requester gave up. They are opposite events. Until **production during appeal** is recorded as its own disposition, every annual report will understate what appeals accomplish — and overstate how often denials were correct.

3 Which agencies are reversed repeatedly — and does anything change?

No overturn rate is published by agency. In the public data, a body whose denials are reversed in consecutive years is indistinguishable from one that has never been appealed. An agency-level reversal rate is the single cheapest accountability measure available, and it already exists inside OOR's docket system.

4 What was in the 12,412 appeals nobody ever decided?

Across 2009–2026, **27.7% of all appeals** closed on a procedural defect — a missed deadline, a form filled in wrong — with no ruling on whether the record was public.[14] Nobody knows how many of those involved plainly public records, because nobody looks. It is the largest unmeasured outcome in the system.

5 Does it matter who advises the agency?

Most Commonwealth agencies are advised by salaried in-house counsel; most municipalities by a contracted solicitor billing by the hour. A denial that is appealed and reversed costs the first nothing additional and generates billable work for the second. **Do municipalities using contracted solicitors deny at higher rates than comparable municipalities with in-house counsel?** OOR already publishes appeals by agency type. Nobody has crossed it against how the agency is advised.

6 What does a denial that could not be justified cost the agency that made it?

In practice, nothing — unless the requester can prove bad faith in Common Pleas, which almost none can afford to attempt.[2] If the General Assembly is weighing new obligations on people who file appeals, the symmetrical question is what obligation attaches to a denial the agency could not defend.

WHY THESE SIX AND NOT OTHERS

Each can be answered from records the Office of Open Records already creates. None asks any agency to do anything it is not already doing — only to record which of several existing outcomes occurred, and to publish the totals.

PART IV · 25 WHAT WOULD ACTUALLY HELP

25 Reforms that address the denial side, not only the filing side

Bath Borough has asked the General Assembly to act on AI-assisted filings. The record supports acting — but on both sides of the exchange. These are ordered by how little new law each would require.

1 · PUBLISH STRUCTURED OUTCOME DATA

Separate procedural dismissals, merits denials, grants, partial grants, production during appeal, withdrawal after production, and mediation — by requester type and agency type. This is a reporting change, not a statutory one, and it makes the other four measurable.

2 · REQUIRE A REAL DENIAL NOTICE

A denial should have to identify the record category, the legal basis, the search actually performed, and the segregability decision. Most of the cost both sides complain about is created at that moment, by a response written without those four things in it.

3 · ATTACH A CONSEQUENCE TO A REVERSED DENIAL

An agency currently faces no adverse outcome from denying a record it should have released. The statute already authorizes penalties for bad faith;[2] what is missing is any remedy the Office of Open Records itself can apply when it reverses a denial — leaving a court action most requesters cannot afford as the only route.

4 · VERIFICATION DUTIES BOTH WAYS

Require citation and quotation verification of every filing, from any party, regardless of what tools produced it — and give open-records officers access to the same source-linked body of final determinations and court opinions. An officer who can see what the law already requires released, *before* denying, generates fewer appeals.

ON SANCTIONS

Tie them to demonstrable false submissions or repeated rule violations — not to the number of requests, the number of appeals, or the use of a particular technology. A rule written against a tool ages badly; a rule written against false statements does not.

THE CLOSING POINT

Bath is entitled to criticize citation errors, and those criticisms were fair. But criticism of how an appeal was written cannot substitute for the outcome of the dispute, and a rising appeal count cannot prove its own cause. **The data do not establish a single cause for rising appeal volume, and this report does not claim one.** What they do show is that a substantial and increasing share of appeals ends with records being released — which makes any explanation focused only on requester conduct or technology incomplete. The question worth putting to the legislature is not "why are more people appealing?" but **"why are so many of those appeals changing the result?"**

APPENDIX A FULL OUTCOME SERIES, 2009–2026

Every year in the corpus, with both denominators

The complete year-by-year series behind Sections 4 and 5, published in full so that any reader can check the screen, recompute the rates, or challenge the classification rather than accept a headline percentage. Read the whole series: the favorable-outcome rate was also high in 2010–2012 and again in 2018. The honest reading is not that something new began in 2022, but that **a substantial share of appealed agency responses has produced access for as long as OOR has published dockets** — and that the share now sits at the top of its historical range.

YEAR	ALL ROWS	REMOVED AT SCREEN	REACHED REVIEW	ORDERED RELEASED	FORMAL RATE	PRACTICAL ACCESS	PRACTICAL RATE
2009	1,018	133	**885**	196	22.1%	199	**22.5%**
2010	1,272	117	**1,155**	366	31.7%	395	**34.2%**
2011	1,623	657	**966**	314	32.5%	315	**32.6%**
2012	2,180	964	**1,216**	385	31.7%	388	**31.9%**
2013	2,422	1,058	**1,364**	330	24.2%	352	**25.8%**
2014	2,019	627	**1,392**	387	27.8%	412	**29.6%**
2015	2,879	1,136	**1,743**	510	29.3%	553	**31.7%**
2016	2,108	683	**1,425**	284	19.9%	344	**24.1%**
2017	2,319	604	**1,715**	347	20.2%	462	**26.9%**
2018	2,302	623	**1,679**	430	25.6%	535	**31.9%**
2019	2,536	541	**1,995**	393	19.7%	523	**26.2%**
2020	2,710	413	**2,297**	464	20.2%	616	**26.8%**
2021	2,913	502	**2,411**	500	20.7%	641	**26.6%**
2022	3,034	678	**2,356**	454	19.3%	600	**25.5%**
2023	3,094	843	**2,251**	445	19.8%	658	**29.2%**
2024	3,299	929	**2,370**	516	21.8%	773	**32.6%**
2025	3,748	1,077	**2,671**	569	21.3%	889	**33.3%**
2026 YTD	3,308	827	**2,481**	603	24.3%	883	**35.6%**
All years	**44,784**	**12,412**	**32,372**	**7,493**	**23.1%**	**9,538**	**29.5%**

"Removed at screen" covers threshold and procedural dispositions: premature, untimely, jurisdictional, noncompliant, insufficient under §1101, and improvidently docketed. Rates are computed on the "reached review" column. Shaded rows mark 2022–2026; 2026 is year-to-date through August 19.

TWO OOR COUNTS, MEASURING TWO THINGS

This table counts OOR decision and docket records by resolution year, so its 2025 total (3,748) differs from OOR's official count of appeals *filed* (3,970). Both come from OOR. Volume claims use the published filing totals; outcome rates use the docket records.

OLDER YEARS ARE UNDERCOUNTED

Early docket descriptions are far less detailed. A 2011 entry may read only "dismissed as moot," never recording whether records were supplied during the appeal. Those were not counted as favorable, so the early years are understated.

THE FORMAL COLUMN IS THE COMPARABLE ONE

Formal grants and partial grants are recorded consistently across all eighteen years, and that column moves with the practical-access column in every recent year. The trend does not depend on how docket language is classified.

APPENDIX B PRIMARY SOURCES

Primary sources used in this report

Every number in this report comes from one of four places: the Right-to-Know Law itself, OOR's published annual reports, OOR's public docket system, or Bath Borough's own website. All are listed here with links.

[1] **Right-to-Know Law, 65 P.S. § 67.708(a) — burden of proof rests on the agency**
https://www.openrecords.pa.gov/Documents/RTKL/PA_Right-To-Know_Law.pdf

[2] **Right-to-Know Law, 65 P.S. §§ 67.1304–1305 — attorney fees and civil penalties, available only from a court and only on a finding of bad faith or wilful conduct**
https://www.openrecords.pa.gov/Documents/RTKL/PA_Right-To-Know_Law.pdf

[3] **Pennsylvania Office of Open Records — 2025 Annual Report**
https://www.openrecords.pa.gov/Documents/AnnualReport2025.pdf

[4] **Pennsylvania Office of Open Records — 2024 Annual Report**
https://www.openrecords.pa.gov/Documents/AnnualReport2024.pdf

[5] **Pennsylvania Office of Open Records — Annual Reports archive (2009–2023 volumes, used for the filing series)**
https://www.openrecords.pa.gov/AnnualReports.cfm

[6] **Bath Borough — "A Requester Markets AI Legal-Research Software…" (Aug. 2026)**
https://bathborough.org/rtk-corpus-mike-os-and-bath-borough-a-documentary-timeline/

[7] **Bath Borough — "Artificial Intelligence and the Right-to-Know Law: Bath Borough Experience" (Oct. 15, 2025)**
https://bathborough.org/artificial-intelligence-and-the-right-to-know-law-bath-borough-experience/

[8] **Bath Borough — AP 2026-2665 public docket page and filing list**
https://bathborough.org/docket-2026-2665-active/

[9] **Bath Borough — Docket AP 2024-0001 (settlement / production)**
https://bathborough.org/docket-2024-0001-closed/

[10] **OOR Docket AP 2022-2675 — Long v. Bath Borough**
https://www.openrecords.pa.gov/Appeals/DocketSheet.cfm?docket=20222675

[11] **OOR Docket AP 2023-1083 — Long v. Bath Borough**
https://www.openrecords.pa.gov/Appeals/DocketSheet.cfm?docket=20231083

[12] **OOR Docket AP 2023-1598 — Long v. Bath Borough (status: Affirmed by Common Pleas)**
https://www.openrecords.pa.gov/Appeals/DocketSheet.cfm?docket=20231598

[13] **OOR Docket AP 2025-1943 — Long v. Bath Borough**
https://www.openrecords.pa.gov/Appeals/DocketSheet.cfm?docket=20251943

[14] **OOR case records, complete population — 44,784 decision and docket rows, 2009 through August 19, 2026, retrieved from the OOR public docket system**
https://docs.google.com/spreadsheets/d/1e2Jqd0Iu0m4SA8gOoQnjbx4Ey5UJKiPvJGpMIVS5eR0/edit

[15] **OOR release-order review — all 2,587 orders releasing records, 2022 through August 19, 2026, classified against the text of each decision**
https://docs.google.com/spreadsheets/d/1eSNPtEE3QQjhwhidIha4Si3Tv5b4nvyOcfb4LIDYaOA/edit

[16] **PA RTK Wizard — public product page and legal/privacy disclosures**
https://partkwizard.famous.chatgpt.site/

[17] **Schnee Legal Services — Pennsylvania Right-to-Know Law practice page**
https://www.schneelegal.com/right-to-know-law

APPENDIX B · CONTINUED STANDARDS AND LITERATURE RELIED ON

Standards and published literature

Each entry is labelled with what kind of source it is. The evidence and accounting entries are **design and validation criteria** — not certifications, accreditations, or endorsements of any software described in this report.

[18] Alsbrook, Margie, "Untangling Unreliable Citations," 37 *Georgetown Journal of Legal Ethics* 415 (2024) **LEGAL SCHOLARSHIP — SECONDARY AUTHORITY**
https://www.law.georgetown.edu/legal-ethics-journal/wp-content/uploads/sites/24/2024/12/GT-GJLE240021.pdf

[19] Desbien, Emily R., "With Great Power Comes Great Responsibility: Court Sanctions for Misuse of GenAI," 38 *J. Am. Acad. Matrim. Law.* 211 (2025) **LEGAL SCHOLARSHIP — SECONDARY AUTHORITY**
https://www.aaml.org/wp-content/uploads/38_Comment-1.pdf

[20] Charlotin, Damien — *AI Hallucination Cases Database*. Continuously updated; 1,934 cases recorded as of August 19, 2026 **RESEARCH DATABASE**
https://www.damiencharlotin.com/hallucinations/

[21] Federal Rule of Evidence 901, *Authenticating or Identifying Evidence* **RULE OF EVIDENCE**
https://www.law.cornell.edu/rules/fre/rule_901

[22] Federal Rule of Evidence 902, *Evidence That Is Self-Authenticating* — including 902(13) and 902(14), added by the 2017 amendments **RULE OF EVIDENCE**
https://www.law.cornell.edu/rules/fre/rule_902

[23] National Institute of Justice, *Electronic Crime Scene Investigation: A Guide for First Responders*, 2d ed., April 2008 (NCJ 219941) **GOVERNMENT GUIDANCE**
https://nij.ojp.gov/library/publications/electronic-crime-scene-investigation-guide-first-responders-second-edition

[24] NIST Special Publication 800-86, *Guide to Integrating Forensic Techniques into Incident Response* (August 2006) **GOVERNMENT GUIDANCE**
https://csrc.nist.gov/pubs/sp/800/86/final

[25] Scientific Working Group on Digital Evidence (SWGDE) — best-practice documents for digital-evidence acquisition, hashing and verification **PROFESSIONAL BEST PRACTICE**
https://www.swgde.org/

[26] U.S. Department of Justice and Administrative Office of the U.S. Courts, Joint Electronic Technology Working Group, *Recommendations for Electronically Stored Information (ESI) Discovery Production in Federal Criminal Cases* (February 2012). Non-binding **GOVERNMENT GUIDANCE**
https://www.justice.gov/media/884436/dl?inline

[27] Pennsylvania Department of Community & Economic Development — Local Government Publications and Documents library **GOVERNMENT GUIDANCE**
https://dced.pa.gov/local-government/publications-and-documents/

[28] Governmental Accounting Standards Board — an independent standard-setter under the Financial Accounting Foundation, not a government agency **PROFESSIONAL STANDARD**
https://www.gasb.org/standards-and-guidance

[29] U.S. Government Accountability Office, *Government Auditing Standards: 2024 Revision* (the Yellow Book), GAO-24-106786. The 2018 revision, GAO-21-368G, remains in effect for engagements beginning before December 15, 2025 **GOVERNMENT STANDARD**
https://www.gao.gov/products/gao-24-106786

[30] U.S. Government Accountability Office, *Standards for Internal Control in the Federal Government* (the Green Book), GAO-25-107721, issued May 15, 2025, superseding GAO-14-704G **GOVERNMENT STANDARD**
https://www.gao.gov/products/gao-25-107721

[31] AICPA, *Statement on Standards for Forensic Services No. 1* (FS Section 100), effective for engagements accepted on or after January 1, 2020 **PROFESSIONAL STANDARD**
https://www.aicpa-cima.com/resources/article/statement-on-standards-for-forensic-services-ssfs-faq

[32] Association of Certified Fraud Examiners — *Occupational Fraud: A Report to the Nations* series. A research study, not a standard **PROFESSIONAL REFERENCE**
https://www.acfe.com/report-to-the-nations

These are cited only for the proposition that they exist and state a control expectation — never for any claim that a named body reviewed, approved or certified anything described here. Where an authority is superseded or has a future effective date, the entry says so.

APPENDIX C METHODOLOGY

How every figure in this report was produced

Two datasets, two published series, and one classification pass. Each is described here in enough detail that a reader can reproduce the counts or challenge the rules.

Methodology in plain English

Two OOR sources are used and never mixed. **Statewide volume and the requester and agency categories come from OOR's published annual reports. Outcome rates come from OOR's own case records** — 44,784 rows, 2009 to August 19, 2026, retrieved from the Office of Open Records' public docket system, because OOR publishes no outcome breakdown by year.

That second set is not a sample: it is the **complete population** for the period, every row traceable to a document on OOR's own site. What has been added is a classification layer, stated below so any reader can reproduce or contest it.

Two outcome measures are used. **Ordered released** means OOR granted the appeal in whole or in part. **Practical access** means an OOR order to release *or* a docket description expressly recording the agency producing or providing records during the appeal. Both are computed after removing appeals thrown out on procedure — premature, untimely, jurisdictional, noncompliant, insufficient under §1101, improvidently docketed — from the denominator in every year. Withdrawals, moot dispositions, transfers and ordinary merits denials were *not* removed, which makes the resulting rates conservative.

The release-order review

Section 7 rests on a separate pass: every OOR order releasing a record from January 1, 2022 to August 19, 2026 — 2,587 orders out of 16,483 appeals — read and classified for whether the decision states the agency failed to carry its burden, or made no submission. Both quote OOR's own reasoning rather than inferring it.

WHAT "PRACTICAL ACCESS" MEANS

The docket shows access-related relief after the dispute reached OOR. It does not mean every agency acted unlawfully, or that a record would have been permanently withheld without appeal.

SCOPE OF THE BATH SECTIONS

A selected case study tied to this report's question — whether appeals can be necessary because agency responses change after review. Not a win/loss scorecard.

2026 figures are year-to-date through August 19, 2026.

APPENDIX D BOROUGH STATEMENTS AND CASE SOURCES

Bath Borough statements being answered

These are the Borough's own publications. They are cited **only for what the Borough publicly claimed.** What actually happened in each matter is established from OOR determinations, court orders, invoices and productions — not from these pages.

B1 RTK Corpus, MIKE OS, and Bath Borough: A Documentary Timeline
https://bathborough.org/rtk-corpus-mike-os-and-bath-borough-a-documentary-timeline/

B2 Right-to-Know: The Costs Behind One Requestor
https://bathborough.org/right-to-know-one-man-is-costing-his-community-a-fortune/

B3 Artificial Intelligence and the Right-to-Know Law: Bath Borough Experience
https://bathborough.org/artificial-intelligence-and-the-right-to-know-law-bath-borough-experience/

B4 History of OOR Appeals with Bath Borough
https://bathborough.org/office-of-open-records-appeals-vs-bath-borough/

Supplemental Borough pages

B5 Borough docket page — AP 2023-1598
https://bathborough.org/docket-2023-1598-pending/

B6 Borough docket page — AP 2024-0001
https://bathborough.org/docket-2024-0001-closed/

B7 Borough docket page — AP 2025-1943
https://bathborough.org/docket-2025-1943-pending/

B8 Borough docket page — AP 2026-2665
https://bathborough.org/docket-2026-2665-active/

B9 Notable Statements of 2024
https://bathborough.org/notable-statements-of-2024/

B10 Notable Statements of 2025
https://bathborough.org/notable-statements-of-2025/

CAPTURE AND CURRENCY

The Borough's pages and totals are actively updated. Each page cited here should be captured as a dated PDF before transmission, with its title, URL, access date, page count and hash recorded, so that what was answered remains fixed even if the published page changes.

READING RULE

Statewide disposition statistics and the selected Bath case studies answer **different questions.** Neither is used as a substitute for the other, and **neither records AI use.** Where a case-study statement rests on a party's assertion rather than an official finding, the text says so.