GRANICUS

Public Records as a Public Service

From fulfillment to sensemaking: designing for transparency, risk, cost and trust

Alex Pettit, Senior Solutions Consultant | Granicus | September 2026

Cover Letter by Cameron Chadsey, Strategic Consultant & Enterprise Accounts Lead (AZ, NV & NM)

Working thesis: A public records request is often a proxy for a question: "What am I trying to understand about government?" It is not a single administrative transaction. It is an enterprise service journey that begins when information is created and ends when responsive records are delivered, a lawful denial is understood or a dispute is resolved. A mature model should also learn from completed requests so government can answer recurring questions more quickly, connect related information and help the next person make sense of what is already public.

Risk posture: Transparency is not a risk to eliminate, it is a fundamental function of our government. The risks are delay, inconsistency, privacy failure, security exposure, uncontrolled cost, avoidable litigation and loss of trust. Those risks can be transferred only in part, avoided only in part and never reduced to zero. The practical strategy is mitigation: better records stewardship, enforceable retention and archival practices, clearer policy, reusable releases, measurable workflow, neutral dispute resolution and carefully governed automation.

Granicus thought leadership. This document is not legal advice and does not recommend a specific technology product.

Cover Letter:

As Nevada's Public Records Task Force works toward its recommendations to the Legislature, we wanted to share several resources that we hope may provide useful context as you consider the issues currently before the Task Force.

Granicus supports one of the largest public records management communities in the country, with more than 900 public records customers across state and local government: including 80+ state agencies, more than 110 of the 300 largest U.S. cities, and approximately 45% of the nation's largest counties. Through that community, we have the opportunity to learn from public records professionals working through many of the same questions Nevada is considering today - across jurisdictions of very different sizes, structures, laws, and operating environments.

One of the clearest lessons we've learned through this experience is that request volume alone does not tell the full story. Our Public Records Complexity Benchmark Report was developed to better understand the changing demands of public records work using anonymized operational data from agencies across the country. Since 2018, the data shows a 161% increase in request volume and a 72% increase in overall complexity, alongside a 438% increase in responsive file sizes, 173% growth in video files, and 31% increase in clarification emails. Those trends point to a broader challenge. The burden associated with a request increasingly depends not simply on how many requests an agency receives, but on the types and volume of records involved, redaction requirements, legal review, coordination across departments, communication with requesters, and the growing variety of digital records government creates and maintains.

Across the nation, our work has afforded us the opportunity to observe different approaches to many of the questions Nevada is considering today: how to meaningfully measure public records performance, how to understand administrative burden, how to improve consistency without diminishing the authority of individual public bodies, and how to make information more accessible to the public. That experience has also raised a broader question: Can public records programs evolve beyond processing individual requests toward making public information easier to find and understand in the first place? The accompanying discussion paper, *Public Records as a Public Service: From Fulfillment to Sensemaking*, explores that question and several others relevant to Nevada's current work. It considers approaches to statewide measurement, administrative and redaction burden, fee methodology, dispute resolution, proactive disclosure, reuse of previously released information, shared terminology and governance, and the responsible use of emerging technologies - while preserving the legal authority of the agencies and records professionals responsible for disclosure decisions.

We recognize that the Task Force is considering policy, governance, transparency, and operational questions - not conducting a technology-selection process. We are sharing these materials in that same spirit. We do not presume to have Nevada's answers. Rather, we hope that what we have learned from public records professionals and governments confronting many of these same challenges can provide useful perspective as Nevada develops its recommendations. If these materials raise questions or there are areas where additional context would be useful, we will welcome the opportunity to share more of what we are seeing, provide additional research or examples from other jurisdictions, or simply participate in the conversation.

Thank you for your work on behalf of Nevada and for the opportunity to contribute to this effort.

Executive Summary

Public records laws begin with a simple democratic premise: government exists to serve the needs of the people and the people insist on remaining informed so that they may maintain control over the instruments that they have created. The operating reality is no longer as simple as providing a paper file. A modern request can touch email, collaboration tools, text messages, social media, body-worn camera video, cloud storage, line-of-business systems, generative AI prompts and outputs, scanned files and sometimes even records held by contractors. Depending on governing law and context, newer forms of machine-generated content can create the same questions of custody, retention and disclosure as older forms of communication. A request may require multiple custodians, legal review, redaction, fee decisions and repeated communication throughout the release process.

The proliferation of technology has grown the workload in scale and complexity. Granicus' 2025 Public Records Complexity Benchmark reports request volume up 161% and its overall complexity index up 72% since 2018 among the agencies in its dataset. File sizes increased 438%, video files 173%, and clarification emails 31%. The same benchmark reports total time spent down 10%, suggesting that process and automation can offset some burden even as the work becomes harder, although the report does not isolate the cause of that decline.

Across the country the policy conversation is shifting beyond a binary choice between transparency and administrative burden. Public records administration is also an exercise in risk management. Government can avoid some preventable demand through proactive disclosure, transfer some processing through vendors or shared services and accept that a residual amount of difficult work will always remain. The dominant strategy, however, should be mitigation: reduce the probability and impact of failure without reducing the right of access.

A second opportunity is to stop treating every request as a snowflake. The federal FOIA framework already recognizes that different words can point to substantially the same records. Its proactive-disclosure rule applies when released records have become, or are likely to become, the subject of later requests for substantially the same records, and U.S. Department of Justice guidance tells federal agencies to look for requests involving the same or similar records. The design implication is broader than exact duplicate detection. A modern intake layer can use semantic similarity, sometimes described informally as fuzzy matching, to recognize that differently worded questions may seek the same record set or overlap substantially. The unit of reuse should be the information sought, not the wording of the request. The system can then offer prior releases, related requests or a narrower path without extinguishing the requester's right to proceed.

This paper proposes a public records semantic commons: a federated discovery and relationship layer over prior releases, open data, request logs, records indexes and other already-public information. The records remain with their lawful custodians. The commons provides a shared minimum vocabulary, metadata and relationship model so information described differently by separate agencies or source systems can still be discovered together. It does not redefine statutory terms such as "public record" or "exemption." Rather, it connects subjects, programs, records series, entities, dates, source-system labels, prior requests and releases while preserving provenance and the original authoritative record. A resident, journalist, policy advocate, business or agency employee could ask a question in ordinary language, discover what government has already released and create a new request only when existing information is insufficient.

Artificial intelligence can strengthen that model, but its role should be bounded. AI can help find prior releases, route a request, identify similar records requests or disputes, summarize public material, suggest clarification and help a Public Records Advocate or records officer research a question. An optional context layer could also surface current legislation, public meeting material, official government announcements and cited external reporting that may help a requester or records officer understand why a topic is appearing now. Official sources should be preferred where available, and external context should be visibly separated from the official records environment. AI should never use outside context to infer a requester's motive, alter scope, change fees or make a disclosure decision. It should not make final exemption decisions, determine a fee waiver, label a requester abusive or replace the human official accountable for a legal decision.

Nevada is a timely case study. Its Legislature considered an independent Public Records Ombudsman in 2025, removed those provisions by amendment and instead created a Public Records Task Force to study exemptions, broad requests, redaction burden, requester costs, dispute resolution, fees, data collection and related issues. The Task Force report is due October 31, 2026. Nevada therefore has an unusual opportunity to address not only how disputes are resolved, but how the entire public records service should learn and improve over time.

Architectural intent: Risk mitigation is how the operating model protects transparency; sensemaking is why the model exists. A mature public records model does not ask only, "Did we close the request?" It also asks, "Did the requester get what they were trying to understand, what did this request teach us and can the next person find the answer without starting over?"

1. Public Records as an Enterprise Service and a Risk-Management Problem

Public records administration is usually described in legal and operational terms. It is also a risk-management discipline. A request creates exposure across legal compliance, privacy, security, cost, staff capacity, litigation and public trust. The central management question is not whether those risks can be eliminated. They cannot. The question is how government chooses to respond to them.

Risk response	What it means in public records	Practical implication
Avoid	Reduce preventable demand or prevent the failure before it exists. This does not mean avoiding access.	Publish high-demand records, govern records upstream, prevent use of unmanaged channels and make existing public information easy to find.
Transfer	Move selected processing, hosting, redaction or specialist work to another capable party.	Shared services, vendors and insurance can transfer work or some financial exposure, but statutory accountability generally remains with government.
Mitigate	Reduce the likelihood or impact of delay, error, inconsistent treatment, privacy failure and litigation.	Standard workflow, enforced records retention and archival requirements, common fee rules, shared minimum metadata, automation, training, reuse, neutral dispute resolution, auditability and human oversight.
Accept	Recognize that some cost, difficult cases, adversarial scrutiny and litigation are inherent in transparent government operations.	Budget for residual risk, define escalation paths and do not distort ordinary service design around the rarest or most adversarial request.

Recommended posture: Mitigation should be the default. Avoid preventable work through lawful proactive disclosure and sound records stewardship, transfer commoditized tasks when sensible and accept residual risk, but design the operating model primarily to reduce friction, inconsistency, exposure and cost without reducing access.

Sensemaking is the architectural "why." Risk management describes how government makes the public records function sustainable, but the purpose is not efficient document production for its own sake. The purpose is to help a person understand what government knows or did, while preserving lawful custody, privacy, confidentiality, security and accountable human judgment.

2. The Current-State Journey: One Question, Many Front Doors and Systems

A requester usually experiences one question: "I am trying to understand something government knows or did." In a federated state, there may be no single intake point at all. The requester may find it necessary to submit substantially the same question to several public bodies because no one can know in advance which agency holds the authoritative record or whether different custodians hold different parts of the responsive material. Each inquiry can create a separate tracking number, often in different systems or spreadsheets, with different scope interpretations, fee estimates, searches, reviews and responses. Separate public bodies and request systems often have no visibility into related requests submitted elsewhere, so substantially similar inquiries can trigger redundant work without either the requester or the agencies recognizing the overlap. The burden of assembling a comprehensive and cohesive answer is therefore pushed back onto the requester.

Figure 1. Current-State Public Records Journey: One Question, Many Front Doors

REQUESTER: "I am trying to understand something government knows or did."

↓ **same question submitted through separate channels** ↓

Governor's Office	Executive Agency	Board / Commission	Independent Office
Separate intake	Separate intake	Separate intake	Separate intake
Scope + fee estimate	Scope + fee estimate	Scope + fee estimate	Scope + fee estimate
Search + legal review	Search + legal review	Search + legal review	Search + legal review
Production / partial	No records / referral	Partial / denial	Production / denial

↓ **requester must reconcile overlapping and inconsistent outcomes** ↓

OVERLAPS • GAPS • DIFFERENT FEES • DIFFERENT TIMELINES • DIFFERENT DISCLOSURE DECISIONS

CURRENT-STATE FAILURE MODE: the public asks one government question, while the operating model creates multiple independent cases. Duplicate work is performed, prior answers remain isolated, related requests are invisible to one another and the requester becomes the integration layer.

NO SHARED REQUEST VISIBILITY • NO COMMON METADATA • NO CROSS-AGENCY SEMANTIC MATCHING

The risk is amplified when government presents many independent front doors. The same question can be sent to a governor's office, an executive agency, a board or commission and an independently elected office, each of which may interpret scope differently, calculate a different estimate, search different systems and reach a different disclosure decision. Some variation is legitimate because custody, records series and exemptions differ. But when the requester must know the organization chart in advance and then reconcile overlapping responses afterward, the architecture itself becomes a public-trust risk. Inconsistent answers can look like concealment and materially different fee estimates can seem punitive, even when both arise from fragmented process rather than bad intent.

What the journey reveals

- The public records process begins before the request. Poor records classification, inconsistent retention and archival practices, mixed personal and government accounts, inaccessible formats and unclear custody become fulfillment problems later.
- The requester sees one government. Fulfillment may require navigating many organizational units, technologies and vocabularies inside that government.
- Routing should be an internal government responsibility. A requester should not need to understand an organization chart to ask a lawful question.
- The most expensive work often occurs after intake: search, collection, legal review, redaction and coordination across systems.
- Even producing the estimate consumes staff time. Staff may have to identify custodians, sample repositories, estimate search and redaction effort and consult legal or technical specialists before the requester decides whether to proceed. Without a common methodology, similar requests can produce materially different estimates across public bodies.
- A single front door does not require centralized custody. Intake, routing, minimum metadata, status communication and fee methodology can be standardized while agencies, boards and independently elected offices retain lawful authority over their own records, classification, retention and archival schedules.
- A closed request should not automatically become a dead transaction. When lawful, a responsive release, its sanitized request metadata and its relationships to prior or subsequent requests should improve the next request and the agency's future response.

3. From Pull to Push: The Self-Service Lineage

Public records requests are a pull mechanism. An individual decides what matters, formulates a question and compels government to respond within the law. Over the last three decades, governments have also built push mechanisms that publish information before a request is made. Those approaches are related, but they solve different problems.

Era / approach	What government publishes	Primary value	Important limitation
Electronic reading rooms and frequently requested records	Previously released records and other information likely to be requested repeatedly	Avoid duplicate processing and give immediate access	Publication is only useful when the public can find and understand the material.
Performance transparency: CitiStat / StateStat	Executive priorities, agency performance measures and progress	Accountability, management focus and public visibility	The executive chooses what is measured and displayed. Push transparency can be selective.
Open data / open checkbook	APIs, machine-readable datasets and transaction-level expenditure information	Self-service analysis and broad reuse	Internal schemas and codes may not match the questions people naturally ask.
Modern request portals	Prior releases, FAQs, trending topics and request status	Deflection, reuse, consistency and lower processing cost	Most reuse remains within one jurisdiction or one request-management corpus.

Era / approach	What government publishes	Primary value	Important limitation
Public records semantic commons / sensemaking layer	Federated prior releases, open data, request logs, records indexes and related public information	Answer questions before a new request is made, connect context across agencies and create institutional memory	Requires records classification and indexing, shared metadata, operational semantics, provenance, privacy controls and disciplined human governance.

The federal FOIA experience provides a particularly direct precedent. The 1996 Electronic FOIA Amendments added a requirement to make frequently requested records electronically available. The 2016 FOIA Improvement Act codified a practical "rule of three": records that have been requested three or more times must be made available electronically. The policy logic is explicit, process the record once and reduce the need for repeated individualized requests.

Maryland StateStat represented a different kind of push. Launched in 2007 and modeled on Baltimore CitiStat, it collected performance information from agencies, used that information in regular accountability meetings and later made substantial data public. Open-checkbook initiatives followed a similar logic in financial transparency. Socrata Open Checkbook, for example, was implemented in states such as Iowa, Delaware, Nevada and South Dakota to expose transaction-level expenditures so the public could search spending without requesting the underlying ledger each time.

Push does not replace pull: A dashboard can be accurate and still be selective. A published dataset can be complete and still be difficult to interpret. Proactive disclosure should reduce unnecessary requests, but the legal right to ask a question government did not choose to publish remains essential.

The tire problem: availability is not the same as answerability

Consider a simple municipal question: "How much did the city spend on tires?" A resident may regard tires as one logical category. The accounting system may not. Tires for a garbage truck, police vehicle and electric-utility bucket truck may be represented by different stock codes, purchasing categories, vendors or departments. A transaction-level open-checkbook portal can expose every relevant payment while still making the answer difficult or impossible to assemble with confidence.

This is a semantic problem. Government data is structured around internal administration. People ask questions using ordinary concepts. Open data reduced the barrier to obtaining the rows. The next problem is helping people map those rows to the question they are actually trying to answer.

The employee-vendor problem: correct data can still produce a misleading answer

A second example is more subtle. In many government financial systems, an employee who is reimbursed for travel or another expense must be created in the accounts payable module as a vendor or payee so the government can issue payment. A query of the vendor master record can therefore return employees alongside outside suppliers. Cross-referencing the vendor list against the employee list can produce an alarming but technically accurate observation that large numbers of employees are also "vendors" of the government. The process did not necessarily reveal self-dealing or an ethical violation. It revealed an accounting construct required to issue payment.

That distinction matters for a semantic commons. The tire example is largely a vocabulary and classification problem. The employee-vendor example is a process-driven semantic collision: the source system uses "vendor" correctly for payment processing but differently from how a resident, journalist or policymaker would normally understand the word. The public discovery layer therefore cannot simply rename fields or impose a universal tag. It must preserve the source-system meaning and provenance, add plain-language context and expose relationships so the user can distinguish administrative structure from the real-world concept being investigated.

4. Reuse Before Rework: Previous Requests as Public Infrastructure

A previous public-records release should be treated, when lawful, as an information asset rather than a dead transaction. If the same records have already been reviewed, redacted and released, government should not automatically repeat the entire process for the next requester. The federal frequently-requested-records rule reflects this principle, and modern request platforms already implement practical versions of it.

Granicus Records Request Management, for example, can flag similar requests, link or merge redundant requests, reuse prior responsive packets and surface public content that may satisfy a requester before submission. Those capabilities are practical precursors to a broader model. The larger opportunity is to make reuse possible across agencies and across types of public information, not only within one request-management database.

Design principle: Process once, reuse when lawful. Preserve the requester's right to proceed when a prior release is incomplete, stale, differently scoped or legally distinguishable.

A public records semantic commons

The proposed semantic commons is not a central warehouse that takes custody away from agencies. It is a federated discovery and relationship layer. Each agency remains responsible for its records, exemptions, retention, archival schedules and legal decisions. The commons makes already-public information easier to discover through shared minimum metadata, semantic indexing and explicit relationships among questions, requests, records series, releases, datasets, source-system concepts and custodians. It should be able to recognize that two differently worded requests occupy the same information neighborhood, that a later release supersedes an earlier one, that several agencies hold different parts of the same subject and that an internal term such as "vendor" may have a specific operational meaning that differs from ordinary language.

Metadata domain	Illustrative elements
Subject and plain-language intent	topic, program, entity, place, date range, question type and natural-language concept
Custody and provenance	source agency, records series, source system, authoritative URL and release date
Legal status	public release, redactions applied, exemption categories, appeal or dispute status
Content characteristics	format, file count, size, video/audio, OCR status and machine readability
Request and information relationships	related requests, overlapping scope, superseded releases, responsive datasets, dashboards, records series and adjacent custodians
Operational semantics	source-system label, plain-language meaning, administrative purpose, known caveats and mapping to related concepts
Discovery aids	plain-language summary, keywords, taxonomy terms and semantic index or embedding
Lifecycle	version, retention, review date and whether the material remains current

This creates institutional memory and, equally important, shared visibility. A question about a subject can be compared with prior releases, open-data sets, related requests and other agencies' published material. The records officer can see that similar work may already have been done elsewhere, and the requester can see relationships that help sharpen the question or locate an existing answer. The goal is not a universal government taxonomy. It is enough shared semantic structure to let different taxonomies and source-system meanings be related without erasing their differences.

Should every request become public?

Not necessarily in raw form. A mature design should distinguish three things: the protected request case file, a public request ledger and a public release library. The case file may contain requester contact information,

first-party requests, sensitive correspondence or facts that should not be proactively exposed. The public ledger can publish a sanitized description, custodian, dates, disposition, fee information, exemption category and appeal status. The release library can publish records already determined appropriate for public reuse.

Federal FOIA logs provide a useful precedent. Many publish the subject, tracking number and disposition so another person can see that records on a topic have already been processed, while excluding first-party requests and requester identity. A denied request can still teach the system something without publishing the very information the denial was meant to protect. The design objective should be maximum reusable transparency consistent with privacy, security and the governing records law, not automatic publication of every raw request.

5. From Search to Sensemaking: The Digital Records Advocate

The next generation of public-records self-service should not merely ask, "Do you want to submit a request?" It should first ask, "What are you trying to understand?" That is the service-design expression of the paper's architectural "why." The request is a means to an information outcome, not the outcome itself. A good system should help a person discover what is already public, see how related information fits together, clarify what remains unknown and create a formal request only when the existing public record is insufficient.

This is already beginning at the federal level. FOIA.gov now uses logic and machine learning to match a natural-language query to agencies, published FOIA logs and frequently requested documents. The tool can suggest existing public records before a requester creates a new request. That is an early version of the semantic-matching concept proposed here.

There is also a bounded role for external context. A Digital Records Advocate could optionally search current legislation, public meeting agendas, official announcements and other government sources first, then supplement them with clearly cited external reporting when useful. It might say, in effect, "This topic is receiving attention now; these events may help explain the records you are seeing." That can support sensemaking, but the context layer should be visibly separate from the official records layer, identify its retrieval date and cite its sources. It should never infer a requester's political, commercial or personal motive as a basis for routing, fees or disclosure.

Sensemaking is a design assumption, not an eligibility test. Many requesters are genuinely trying to understand something, but a journalist may deliberately test inconsistent government answers, an advocate may be building a case, a commercial requester may seek reusable data and an automated agent may submit requests at a scale no human could. Some activity may also be intentionally duplicative or disruptive. The architecture should be resilient to those behaviors through reuse, shared visibility, auditability, capacity planning and only those limits authorized by law. It should not require government to decide whether a requester's purpose is sufficiently sincere before honoring a lawful request.

Figure 2. Future-State Sensemaking Loop

1. Ask in ordinary language	"What am I trying to understand?"
2. Search the public records environment	Prior releases + open data + request logs + dashboards + records indexes
3. Match, relate and clarify	Show likely matches, map relationships to prior requests, releases, datasets, source-system concepts and other custodians, explain overlaps or gaps and ask a narrowing question if needed
4. Route a new request only if necessary	Send to the correct custodian with useful context and a clear audit trail
5. Fulfill and publish when lawful	Release, redact, explain and add reusable metadata or public material back to the corpus
6. Close the sensemaking loop	"Was this what you were looking for?" "This is related. Would it help?" Feed the result back into metadata, reuse and proactive disclosure.

YOUR QUESTION	Prior release	Open dataset	Related request	Other agency

1. Ask in ordinary language	"What am I trying to understand?"		
same subject / scope	same concept / data	overlap / prior treatment	adjacent custody

Illustrative relationship map: the goal is not to tell the requester what to think. It is to expose the information neighborhood around the question, including prior treatment, related data, source-system meanings and adjacent custody, so the requester can sharpen the inquiry or answer it without a new search.

AI should advise the Public Records Advocate, not become the advocate

Artificial intelligence is well suited to search, comparison and synthesis. A Public Records Advocate or ombudsman role also requires neutrality, judgment, procedural fairness and accountable legal interpretation. Those responsibilities should remain human. A useful design principle is simple: AI can be the advocate's researcher, clerk and navigator. It should not be the human advocate who makes final interpretations, decisions or determinations.

AI can assist with	Human accountability should remain with
Find similar requests, releases and prior decisions	Final interpretation of exemptions and disclosure obligations
Summarize facts, correspondence and public records	Release authorization and final denial
Surface statutes, policy and relevant precedent with citations	Balancing tests and legally discretionary judgments
Suggest clarification language or mediation options	Fee waivers, fee disputes and final settlement decisions
Detect inconsistency across agencies or similar requests	Designation of a requester as abusive, acting in bad faith or otherwise subject to a legally recognized restriction
Draft a neutral explanation and create an audit trail	Final appeal, enforcement or adjudicatory decisions
Surface cited current events, legislation and public trends as a separate context layer	Determine legal relevance, infer requester motive or change scope, fees or disclosure based on external context

6. What the Data Says: Complexity, Capacity and Concentration

The Granicus 2025 benchmark is useful because it attempts to quantify what records officers experience: two agencies receiving the same number of requests can face radically different workloads. A request for a routinely published report is not operationally equivalent to a request spanning years of email, video and multiple custodians. The benchmark groups complexity into scale, diversity and connectivity and then measures characteristics of the records and the human activity required to fulfill them.

Benchmark signal	Reported change since 2018	Why it matters
Request volume	+161%	More demand entering the system
Overall complexity index	+72%	Requests are harder even after accounting for volume
File size	+438%	Multimedia and large electronic productions change review and delivery
Video files	+173%	Public-safety records add specialized redaction and storage demands
Clarification emails	+31%	Scope ambiguity creates iterative labor and delay
Total time spent	-10%	Process and automation may be offsetting some burden, but the benchmark does not establish causation

Source: Granicus, 2025 Public Records Complexity Benchmark Report, supplied with this discussion draft.

Connecticut: workload is concentrated by mission

The Connecticut workbook supplied for this draft contains 71,024 requests across participating state-agency sites during 2023 and 2024. It is most useful for understanding concentration rather than statewide growth because participation expands materially across the reporting period.

Agency	Requests in supplied dataset	Share
Department of Emergency Services and Public Protection	30,270	42.6%
Department of Children and Families	11,091	15.6%
Workers' Compensation Commission	7,847	11.0%
Department of Energy and Environmental Protection	7,316	10.3%
Department of Correction	5,653	8.0%
Department of Public Health	4,337	6.1%

Those six agencies account for approximately 93.7% of the request volume represented in the workbook. A statewide policy based on an "average request" would therefore be misleading. Public safety, child welfare, environmental regulation, corrections and public health produce different records, different redaction burdens and different risk profiles.

7. How States Govern Public Records Disputes

A scan of all 50 states shows a strong dominant design, but it is not a single dispute model. The dominant design is federated fulfillment with variable escalation. In every state reviewed, the public body or custodian that holds the record remains the starting point for the request. Centralized state bodies, where they exist, generally provide advice, mediation, legal review, administrative appeal or enforcement after the agency-level process has begun.

Dominant design: Federated fulfillment, variable escalation. The state may centralize policy, assistance or review, but it generally does not centralize custody and fulfillment of every state and local public-records request in one office.

The states can be understood along a continuum. The categories overlap, and the appendix assigns a primary working model only to make comparison easier.

Escalation model	What happens when the request goes wrong	Representative examples
Court-first / limited statewide review	The requester generally relies on judicial enforcement, sometimes after internal agency action.	Alabama, California, Colorado, Louisiana, Missouri, Montana, South Carolina and West Virginia
Attorney General / prosecutor / agency-head review	A state legal office or senior agency official reviews withholding or can investigate and enforce.	Delaware, Kentucky, Nebraska, North Dakota, Oregon, Rhode Island, Texas and Vermont
Advisory / ombuds / mediation	A neutral or advisory body helps interpret the law, mediate or improve practice without replacing the courts.	Florida, Idaho, Indiana, Maine, Tennessee, Virginia, Washington and Wyoming
Administrative / quasi-judicial appeal	A commission, board, special master or administrative office can hear disputes and issue decisions or orders.	Connecticut, Hawaii, Illinois, Iowa, Maryland, Massachusetts, Mississippi, New Jersey, Ohio, Pennsylvania, South Dakota and Utah

The policy implication is important. A statewide semantic commons or common service standard would not require a state to strip agencies of lawful custody. It could sit above a federated fulfillment model as a shared discovery, metadata, measurement and guidance layer.

Federated custody also does not require fragmented intake. Iowa's Department of Management Division of Information Technology makes an Open Records Request Portal available to state agencies and describes its purpose as providing a consistent process across agencies, including the ability for constituents to review responses to prior questions and requests. Montana went further for agencies reporting to the governor: its Office of Public Information Requests was created to establish common policies and a centralized intake and response process while agencies continue to own the underlying records. At the federal level, FOIA.gov provides a consolidated submission front door while agencies retain their own response-processing authority and systems.

Figure 3. Target Operating Architecture: One Front Door, Federated Custody

ONE FRONT DOOR / STATEWIDE INTAKE		
Semantic match & prior-release search • cross-agency request visibility • common request ID / metadata • relationship map • routing • common status and fee methodology		
Agency / Custodian A	**Agency / Custodian B**	**Board / Independent Custodian**
Search • collect • review • redact Lawful custodian retains authority	Search • collect • review • redact Lawful custodian retains authority	Search • collect • review • redact Lawful custodian retains authority
SHARED REQUEST LEDGER + RELEASE LIBRARY + SEMANTIC COMMONS		
Coordinated result: responsive records / partial release / denial + explanation + fee basis + related public materials		
Feedback: "Was this what you were trying to understand?" → improve metadata, reuse and proactive disclosure		

The front door is unified; custody and legal authority remain federated.

8. Nevada Is Already at the Decision Point

Nevada is not beginning from a blank sheet. Assembly Bill 128 was introduced in January 2025 as a proposal to create an independent Office of Public Records Ombudsman. The introduced bill contemplated mediation, complaint investigation, training, public education and final determinations that could address access or fees. Those sections were deleted by amendment.

The enacted law created the Public Records Task Force instead. Its mandate includes exemptions, broad requests, redaction burden, requester financial burden, alternative dispute resolution and methods for collecting and sharing request data. It must also recommend approaches to custody, sensitive information, access and denial, penalties, a standardized fee schedule and protections against retaliatory litigation. Its report is due October 31, 2026.

Nevada did not miss the ombudsman idea: It considered one and chose, at least for the current cycle, to study the issue through a task force. The question now is not whether an ombudsman model exists elsewhere. It is what combination of dispute resolution, service standards, fee rules, data and accountability Nevada wants to institutionalize.

9. A Fee Schedule Is a Policy Choice, Not Just an Accounting Exercise

A standardized fee schedule sounds simple until the state defines what a requester should pay for. Paper and media costs are relatively easy. Search labor, legal review, redaction, specialized video processing, contractor expense and system extraction are harder. The policy question is which costs are ordinary overhead, which are incremental to the request and which should be subsidized as part of government's transparency obligation.

A useful statewide approach does not necessarily require one flat price. It does require common definitions and explainability. A requester should understand why an estimate exists, what work it includes, what options exist to narrow the request and how to challenge a disputed fee without immediately filing suit.

Design principle	Why it matters
Define cost consistently	Separate ordinary agency overhead from incremental fulfillment cost.
Estimate before work begins	Let the requester narrow scope or proceed knowingly.
Reward narrowing	Make clarification a service feature rather than an adversarial exchange.
Differentiate specialty records	Video and unusual formats may justify different treatment, but the method should be published.
Create a low-cost challenge path	Fee disputes should have neutral review before litigation when feasible.
Publish the methodology	Transparency about the fee model makes agency practice easier to compare and defend.

10. Upstream Records Stewardship Is Part of the Architecture

Public-records modernization is often treated as a fulfillment project, but some of the most consequential failures are created months or years before a request arrives. Personal accounts, ephemeral messaging, ungoverned collaboration tools, AI-generated content, inconsistent retention, unclear custody, inaccessible formats and poorly classified records can all become disclosure problems later. Identity, account design, retention, archival practice, classification, custody and system architecture are therefore transparency controls. A request-management system can improve fulfillment, but it cannot fully repair poor information stewardship after the fact.

11. From Fulfillment to Sensemaking: A Three-Horizon Maturity Model

The future state does not require a leap directly from email inboxes to autonomous AI. It can be built in horizons, with each horizon creating value on its own and reducing the risk of the next.

Horizon	Primary objective	Core capabilities	Success looks like
1. Transaction control	Make each request visible, lawful and manageable.	Intake, routing, deadlines, fees, redaction, secure release, audit trail and status communication.	Fewer missed deadlines, better consistency and defensible processing.
2. Reuse and proactive disclosure	Stop re-processing information government has already made public.	Prior-release reuse, duplicate linking, release-to-all, common metadata, request logs, proactive publishing and neutral assistance.	More self-service, fewer redundant requests and lower unit cost.
3. Semantic sensemaking	Help people find meaning across public information before creating a new request.	Cross-agency semantic commons, natural-language search, relationship mapping, Digital Records Advocate, related-content suggestions, feedback loops and AI-assisted research.	People understand answers across organizational boundaries while formal request rights and human legal decisions remain intact.

Self-service and automation create new risks too

Making public-records access easier is not risk-free. Technology reduces some forms of friction and can create others. The current federal debate is already seeing this tension. In a 2026 survey reported by the FOIA Advisory Committee, 46% of respondents believed they had received AI-generated requests, 50% were unsure and only 4% said no. The Committee cautioned that the deeper issue may be the volume and frequency enabled by technology generally, not AI alone.

Self-service / AI risk	Why it matters	Control
Automated high-volume requests	Scripts or AI agents can make it cheap to create large numbers of tailored requests, including requests designed to assemble bulk corpora for commercial analysis or model training. A workflow designed for human-scale demand can be overwhelmed even when each request is facially lawful.	Publish reusable datasets or APIs when appropriate, distinguish bulk-data access from case-based request fulfillment, measure request patterns and use only those rate or abuse controls the governing law permits. Do not make requester purpose a condition of access where the law does not allow it.
Deliberately duplicative or disruptive activity	Not every request is primarily a search for understanding. Observable behavior may include repeated near-duplicate submissions, conflicting multi-agency requests or automated activity intended to test or consume capacity. Similar behavior can also arise from legitimate journalism, oversight or commercial reuse, so motive is an unreliable control point.	Design around observable workload and lawful rights, not inferred motive: semantic duplicate detection, shared request visibility, prior-release reuse, bulk-data channels, auditability, capacity planning and only those restrictions or sanctions authorized by law.
Deflection becomes denial	A system may suggest a "similar" release that is not actually responsive.	Always preserve a clear path to submit the request and explain why an item was suggested.
Stale or incomplete publication	Old releases can appear authoritative after facts or law change.	Version, date, review and link every result to the authoritative source.
Privacy through aggregation	Separate lawful releases may become sensitive when combined.	Privacy review, metadata minimization and careful publication rules.
Bad semantics	Poor taxonomy can confidently map an ordinary question to the wrong government category.	Expose provenance, allow correction and use human review for uncertain matches.
Political curation or ranking bias	Push systems can favor what government wants people to see.	Separate search relevance from policy promotion and preserve uncensored request rights.
AI hallucination or legal overreach	A fluent answer can misstate law, exemptions or the contents of a record.	Ground answers in source material, cite provenance and require human legal decisions.
Digital divide and accessibility	A digital-first path can disadvantage people who need another channel.	Maintain accessible, offline and assisted alternatives.
Metadata leakage	Request descriptions can reveal identity, investigations or sensitive context.	Define what metadata may be public and separate public indexes from protected case data.
Rebound effect	Better access can increase legitimate demand faster than savings arrive.	Plan capacity, measure deflection honestly and treat higher use as a service signal rather than automatic failure.

12. Measure Outcomes, Not Just Request Count

Nevada's Task Force specifically calls for methods to collect and share public-records data. That is an opportunity to avoid a common mistake: using request volume as the primary measure of performance. Volume describes demand. It does not describe difficulty, timeliness, quality, fairness or whether a requester actually found what they needed.

Measurement domain	Examples
Demand	Requests received, repeat requests, subject categories and channels
Timeliness	Time to acknowledge, first substantive response, median closure and 90th-percentile closure
Complexity	Repositories searched, custodians involved, file count/size, multimedia and redaction effort
Interaction	Clarification cycles, partial productions, status contacts and requester inactivity
Cost	Estimated fees, final fees, waived/reduced fees, staff hours and external costs
Outcome	Full release, partial release, no records, withdrawal, denial and proactive-source referral
Reuse	Prior release matched, related request discovered, duplicate linked, successful self-service resolution, content reused and proactive publication triggered
Sensemaking	Requester confirms answer, related material accessed, cross-agency relationships discovered, clarification reduced and new request avoided
Dispute	Mediation, appeal, reversal, litigation, penalties and attorney fees
Learning	Recurring exemptions, recurring fee disputes, training needs and new metadata or records-management changes

13. Questions Worth Asking Before Choosing an Operating Architecture

- What is the desired public-records service standard, and which parts should be common across every public body?
- Which risks is the state trying to avoid, transfer, mitigate or accept, and which risks cannot legally be transferred away from the public body?
- Where should the first neutral review of a dispute occur: inside the agency, at the Attorney General, with an ombudsman, before an administrative body or only in court?
- Which fulfillment costs should be treated as the ordinary cost of transparent government and which may reasonably be charged to the requester?
- Should statewide fee policy standardize exact prices or standardize definitions, calculation methods and disclosure of the estimate?
- What previously released records should become reusable public assets, and how should the state decide when a release is too stale, specific or sensitive to reuse?
- What common metadata, operational semantics and relationship rules would let agencies compare requests and releases without centralizing custody of the underlying records?
- Could a resident ask one natural-language question and discover relevant material, related requests and source-system context across state agencies, counties and cities before deciding whether a formal request is needed?
- How should AI assist a records officer or ombudsman while preserving human accountability for legal judgment?
- How should the state govern official business conducted through personal email, text messaging, collaboration tools, AI assistants and vendor platforms?
- What minimum request data should every public body collect so policymakers can evaluate workload, delay, cost, reuse and disputes?
- How should the architecture respond to automated, duplicative or strategically burdensome use without turning requester motive into an access test?
- What should constitute success: faster responses, lower cost, fewer disputes, greater access, better sensemaking, higher successful self-service resolution before formal submission or a balanced combination?

14. Closing: Mitigate the Risk, Improve the Service

The traditional public-records model is transactional. A requester asks, government searches, lawyers review and a record is released or denied. That model will remain necessary because the right to ask an unanticipated question is fundamental to meaningful transparency.

But the transaction does not have to remain isolated. A mature operating model can learn from every completed request. It can publish what is frequently sought, reuse what has already been lawfully released, connect related information, standardize metadata, identify where agencies differ and provide neutral assistance before a disagreement becomes litigation.

That is the risk-management opportunity, but risk mitigation is the operating discipline rather than the ultimate purpose. Government cannot transfer away the duty to be transparent. It should not attempt to avoid scrutiny, and it should accept that some difficult work is unavoidable. What it can do is mitigate the avoidable cost and risk created by fragmented systems, inconsistent policy, repeated work and poor information discovery. The architectural "why" remains sensemaking: helping people understand government more reliably without weakening lawful protections or accountable human judgment.

Closing thought: Responsible stewardship of public records should be treated as a public service, not a treasure hunt. The goal is not simply to lawfully deliver or withhold requested records, but to help the public make sense of government while protecting lawful privacy, confidentiality, security, capacity and public trust.

Appendix A. Fifty-State Public Records Dispute-Resolution Scan

This is a working policy classification, not a legal opinion. It identifies the primary statewide overlay that a requester is likely to encounter after or alongside the public body that holds the record. Models overlap, local ordinances may add processes and subject-specific laws can create exceptions. The comparative baseline is the Reporters Committee for Freedom of the Press Open Government Guide, supplemented by current official state sources for specialized offices and recent structural changes.

State	Initial responsibility	Primary statewide escalation / assistance	Working model
Alabama	Public body / custodian	No general administrative appeal identified; judicial enforcement is the primary route. Attorney General opinions may address open-records questions.	Court-first
Alaska	Public body / custodian	State-agency denials may be appealed to the agency head; the general Alaska Ombudsman can assist with administrative complaints.	Internal appeal + ombuds assistance
Arizona	Public body / custodian	No general records-law administrative appeal; the general Ombudsman-Citizens' Aide may investigate state administrative acts.	Court-first + general ombuds
Arkansas	Public body / custodian	No general FOIA ombudsman; court is the principal enforcement route. Limited Attorney General review exists for certain personnel/evaluation matters.	Court-first
California	Public body / custodian	No statewide administrative appeal for general public-records denials; judicial enforcement. Some local governments add sunshine procedures.	Court-first
Colorado	Public body / custodian	No statewide public-records ombudsman or administrative appeal body for general CORA disputes; judicial enforcement.	Court-first
Connecticut	Public body / custodian	Freedom of Information Commission hears complaints, uses settlement/ombuds processes and can order relief.	Quasi-judicial commission
Delaware	Public body / custodian	Requester may petition the Attorney General for a determination of whether a FOIA violation occurred; court remains available.	AG review
Florida	Public body / custodian	Judicial enforcement plus a voluntary Open Government Mediation Program in the Attorney General's Office.	Court-first + AG mediation
Georgia	Public body / custodian	Court enforcement plus Attorney General open-government assistance, informal mediation and enforcement authority in appropriate matters.	Court-first + AG mediation/enforcement
Hawaii	Public body / custodian	Office of Information Practices reviews denials, issues opinions, investigates and provides guidance under UIPA and Sunshine Law.	Administrative review
Idaho	Public body / custodian	Judicial enforcement plus an executive-branch Public Records Ombudsman for informal mediation and guidance.	Court-first + ombuds
Illinois	Public body / custodian	Attorney General Public Access Counselor accepts requests for review, mediates and may issue binding or advisory opinions.	AG administrative review

State	Initial responsibility	Primary statewide escalation / assistance	Working model
Indiana	Public body / custodian	Public Access Counselor provides education, investigates complaints and issues advisory opinions; court provides enforcement.	Advisory + court
Iowa	Public body / custodian	Iowa Public Information Board offers informal resolution and mediation, investigates complaints and can issue orders through contested cases.	Quasi-judicial board
Kansas	Public body / custodian	No general requester merits appeal body; Attorney General, county attorney or district attorney may investigate and enforce.	Prosecutorial enforcement + court
Kentucky	Public body / custodian	Attorney General reviews open-records appeals; a decision has legal force unless timely appealed to circuit court.	AG administrative review
Louisiana	Public body / custodian	No general administrative appeal; judicial enforcement. Attorney General provides education but not a general requester appeal process.	Court-first
Maine	Public body / custodian	No statutory general administrative appeal; Attorney General FOAA Ombudsman assists with complaints and interpretation; court enforces.	Court-first + ombuds assistance
Maryland	Public body / custodian	Public Access Ombudsman mediates; covered disputes may then go to the Public Information Act Compliance Board, with court also available.	Mediation + administrative review
Massachusetts	Public body / custodian	Supervisor of Public Records provides an administrative appeal process and may refer noncompliance for enforcement.	Administrative review
Michigan	Public body / custodian	Requester may pursue an optional appeal to the head of the public body or proceed to circuit court.	Internal appeal or court
Minnesota	Responsible authority / public body	Commissioner of Administration issues advisory data-practices opinions; Office of Administrative Hearings can hear certain data-practices complaints.	Advisory + administrative review
Mississippi	Public body / custodian	Ethics Commission accepts complaints, mediates, holds hearings and can order production and impose statutory remedies.	Quasi-judicial commission
Missouri	Public body / custodian	No general administrative requester appeal identified; judicial enforcement is the principal route.	Court-first
Montana	Public body / custodian	No general public-records ombudsman or administrative appeal under the act; judicial enforcement.	Court-first
Nebraska	Public body / custodian	Attorney General can review a written denial and pursue enforcement when a public body refuses to disclose.	AG review/enforcement
Nevada	Public body / custodian	NRS Chapter 239 provides district-court relief; no current general statewide ombudsman or administrative appeal body. A temporary task force is studying alternatives.	Court-first; reform under study
New Hampshire	Public body / custodian	Right-to-Know Ombudsman is a statutory alternative to superior court, though the office was reorganized in 2025 and the position has been in transition.	Ombuds review or court
New Jersey	Public body / custodian	Government Records Council offers mediation and adjudication; requester may instead pursue Superior Court.	Quasi-judicial council

State	Initial responsibility	Primary statewide escalation / assistance	Working model
New Mexico	Public body / custodian	No formal general administrative appeal; Attorney General or district attorney may enforce and requester may go to court.	Court-first + public enforcement
New York	Public body / custodian	Internal appeal to the agency head or designated appeals officer; Committee on Open Government provides advisory opinions and training; court follows.	Internal appeal + advisory
North Carolina	Public body / custodian	Judicial enforcement for access disputes; State CIO/designee has a limited role mediating certain fee disputes.	Court-first + limited fee mediation
North Dakota	Public body / custodian	Attorney General may review written denials and open-government violations and issue opinions.	AG review
Ohio	Public body / custodian	Court of Claims public-records complaint process uses mediation and a special master for expedited review; mandamus remains an alternative.	Special judicial review
Oklahoma	Public body / custodian	Judicial enforcement is primary; Attorney General opinions interpret the Open Records Act and guide public bodies.	Court-first + AG guidance
Oregon	Public body / custodian	State-agency denial appeals go to the Attorney General and local-body denials generally to the district attorney; Public Records Advocate offers facilitated dispute resolution, training and advice.	AG/DA review + advocate mediation
Pennsylvania	Public body / open-records officer	Office of Open Records hears administrative appeals and issues final determinations for most covered bodies.	Administrative appeal
Rhode Island	Public body / custodian	Internal appeal to chief administrative officer, then Attorney General complaint or court.	Internal + AG review
South Carolina	Public body / custodian	No general statewide ombudsman or administrative appeal identified; judicial enforcement.	Court-first
South Dakota	Public body / custodian	Office of Hearing Examiners reviews certain public-records disputes.	Administrative review
Tennessee	Public body / custodian	Office of Open Records Counsel provides assistance, informal advisory opinions and dispute support; courts enforce access rights.	Advisory + court
Texas	Governmental body	Attorney General Open Records Division issues rulings when governmental bodies seek to withhold requested information; litigation remains available.	AG legal review
Utah	Government entity / records officer	Agency-level appeal precedes review by the Government Records Office Director; Government Records Ombudsman provides mediation; judicial review follows.	Administrative review + ombuds
Vermont	Public body / custodian	Agency-head administrative appeal, then judicial review.	Internal appeal + court
Virginia	Public body / custodian	FOIA Advisory Council provides guidance, advisory opinions and training but no binding dispute authority; court enforces.	Advisory + court
Washington	Public body / records officer	Court enforces access; Attorney General Open Government Ombuds assists and can provide nonbinding review of certain state-agency denials.	Court-first + AG ombuds

State	Initial responsibility	Primary statewide escalation / assistance	Working model
West Virginia	Public body / custodian	No general statewide public-records ombudsman or administrative appeal identified; judicial enforcement.	Court-first
Wisconsin	Public body / custodian	Judicial enforcement; Attorney General and district attorneys have enforcement/advisory roles but no general requester ombudsman.	Court-first + public enforcement
Wyoming	Public body / custodian	Public Records Ombudsman mediates disputes and has statutory authority regarding timelines and certain fee issues; district court remains available.	Ombuds with statutory authority

District of Columbia comparison: the Office of Open Government provides FOIA/open-government guidance and training, while denials may also follow the District's FOIA appeal process. It is not included in the 50-state count above.

Appendix B. Proposed Semantic Commons Design Principles

The semantic commons is intentionally a governance and information-discovery concept, not a product specification. A state could implement parts of it through standards, open-data catalogs, request-management systems, search services or shared metadata without moving the underlying records into one repository.

Principle	Design implication
Federated custody	The lawful custodian remains authoritative. The commons indexes or links to public material rather than becoming the legal system of record.
Provenance first	Every result identifies source agency, source system or published location, release date and version.
Common minimum metadata	Agencies agree on a small shared vocabulary for subject, custodian, time period, release status, format, provenance and relationships while preserving agency-specific taxonomies.
Semantic, not schema-dependent	Natural-language questions can map across different agency taxonomies, codes, operational labels and data models without rewriting the authoritative source.
Reuse is conditional	A prior release can be suggested but does not extinguish the right to request records when scope, time period or legal context differs.
Public by design, private by rule	Only metadata and records appropriate for public discovery enter the public layer; requester identity and protected case information remain separated.
Human legal authority	AI can rank, summarize and suggest. Humans remain accountable for exemptions, redactions, fees, waivers and disputes.
Feedback creates learning	Completed requests, requester feedback, recurring topics and cross-agency relationships improve metadata, proactive publication, training and future routing over time.
Three-layer publication model	Separate the protected request case file from a sanitized public request ledger and the public release library so reuse does not expose requester identity or protected context.
Operational meaning	Preserve source-system terminology and provenance while adding plain-language context so administrative labels such as “vendor” are not mistaken for a broader real-world meaning.
Behavior, not inferred motive	Design resilience around observable duplication, volume and workload. Do not make lawful access depend on government deciding whether a requester is sincere, adversarial, commercial or politically motivated.

Illustrative metadata record

The example illustrates why metadata matters. The resident’s concept “tires” can span multiple internal purchasing categories. The public layer should help bridge that vocabulary difference while preserving links to the authoritative records.

Appendix C. Selected Sources and Research Notes

1. Granicus, 2025 Public Records Complexity Benchmark Report. User-supplied benchmark report. Key reported signals used here: request volume +161%, complexity index +72%, file size +438%, video +173%, clarification emails +31% and total time spent -10% since 2018.

2. Connecticut ARV Report. User-supplied workbook. 71,024 requests represented across participating agency sites in 2023-2024. Used for concentration analysis, not as a clean statewide trend series.

3. Oregon Senate Bill 106 (2017). User-supplied enrolled bill establishing the Office of Public Records Advocate and Public Records Advisory Council, including facilitated dispute resolution, training and policy review.

4. Reporters Committee for Freedom of the Press, Open Government Guide. Comparative baseline for state public-records laws and appeal/enforcement models. Source

5. U.S. Department of Justice, Proactive Disclosure of Non-Exempt Agency Information. Explains the 1996 e-FOIA frequently-requested-records requirement and the efficiency rationale behind proactive disclosure. Source

6. U.S. Department of Justice, FOIA Improvement Act of 2016 summary. Documents codification of the three-request threshold for electronic proactive disclosure. Source

7. FOIA.gov, How the FOIA Search Tool Works. Describes logic and machine-learning matching against agencies, published FOIA logs and frequently requested documents. Source

8. National Archives, 2024-2026 FOIA Advisory Committee. Current federal work on volume, frequency and AI-generated requests. Source

9. National Archives, Volume and Frequency Subcommittee Final Report. Reports perceptions of AI-generated requests and cautions that request volume/frequency may be the broader technology issue. Source

10. Maryland State Archives, Governor's StateStat Office. Documents StateStat's 2007 origin, performance-management role and transparency/accountability mission. Source

11. Socrata / Tyler Data & Insights, Open Checkbook Overview. Describes transaction-level expenditure publication and self-service exploration. Source

12. Granicus Records Request Management (GovQA). Current product capabilities include redundant-request linking/merging, similar-request alerts, prior-response reuse and self-service deflection. Source

13. Nevada AB 128, introduced version. Introduced January 2025 as an independent Office of Public Records Ombudsman proposal. Source

14. Nevada 2025 Statutes, Chapter 228 (AB 128). Enacted Public Records Task Force, including its study topics and October 31, 2026 reporting deadline. Source

15. Nevada Revised Statutes Chapter 239. Current statewide public-records statute, including application to governmental and local governmental entities. Source

16. Nevada Revised Statutes Chapter 268. Current city powers framework explaining Dillon's Rule and Nevada's matters-of-local-concern modification for incorporated cities. Source

17. Oregon Office of the Public Records Advocate. Current dispute-resolution, training and consultation role. Source

18. Iowa Public Information Board. Current mediation, investigation and contested-case enforcement model. Source

19. Maryland Public Access Ombudsman and PIA Compliance Board. Current mediation and administrative-review framework. Source

20. Utah Government Records Office. Current post-2025 administrative review structure and records ombudsman function. Source

21. Ohio Attorney General, Public Records Complaint Process. Describes Court of Claims complaint, mediation and special-master process. Source

22. Texas Attorney General, Open Records Division. Current legal ruling and assistance framework under the Texas Public Information Act. Source

23. U.S. Department of Justice, Proactive Disclosure of Non-Exempt Agency Information. Explains the federal requirement and DOJ practice of identifying requests involving the same or similar records, not merely exact duplicate language, and using prior releases to reduce repetitive processing. https://www.justice.gov/oip/oip-guidance/proactive_disclosure_of_non-exempt_information

24. Iowa Department of Management, Division of Information Technology, Open Records Request Portal. Describes a state-supported portal intended to provide a consistent process across agencies and allow constituents to review responses to prior questions and requests. https://dxtraining.iowa.gov/open-records-request-portal

25. Montana Department of Administration, Office of Public Information Requests. Describes centralized intake and response, common policies and enterprise request processing for supported executive-branch agencies. https://doa.mt.gov/opir/

26. U.S. Department of Justice, Office of Information Policy, FOIA Library and request logs. Provides a useful public-ledger precedent: request subject, tracking and disposition are published while requester identity and first-party requests are excluded. https://www.justice.gov/oip/available-documents-oip

27. U.S. Department of Justice, National FOIA Portal / FOIA.gov. Federal consolidated intake allows submission from a common website while agencies retain their own response-processing authority and systems. https://www.foia.gov/