

FEDERAL RECORDS & GOVERNANCE PRACTICE · PLAYBOOK

# The Records Mandate Playbook

How the federal electronic records mandate is reshaping agency architectures, what "AI output as a record" means, and a de-risking sequence for the agencies still inside the mandate phase.

<b>02</b>	<b>Executive summary</b>
<b>03</b>	<b>The mandate landscape, sequenced</b>
<b>05</b>	<b>Where federal agencies actually are</b>
<b>07</b>	<b>What the mandate changes structurally</b>
<b>09</b>	<b>When AI output becomes a record</b>
<b>11</b>	<b>The cost of postponement</b>
<b>12</b>	<b>A de-risking sequence</b>
<b>14</b>	<b>What this rules in and out</b>
<b>15</b>	<b>Endnotes &amp; sources</b>

## ● Executive summary

The federal electronic records mandate is not new. OMB/NARA Memorandum M-19-21 set the framework in June 2019 — all permanent and temporary records managed electronically by the end of calendar year 2022, with NARA ceasing to accept analog transfers after that date.<sup>1</sup> M-23-07 extended specific elements and clarified the ongoing obligations.<sup>2</sup> The agencies that read the memos in 2019 had every signal they needed to start the work. Many did not start the work.



What is forcing the convergence now is not the mandate itself — it is the arrival of agentic AI workloads inside federal agencies, generating records faster than records officers can classify them, on architectures that were never designed to absorb them. The mandate work and the AI work are now the same work. This playbook sets out how to do that combined work as a deliberate program rather than a forced response to the next audit. Its argument is in five parts:

- **The mandate is a layered obligation, not a single deadline.** M-19-21, M-23-07, 36 CFR Part 1236, Capstone, and the annual self-assessment cycle together define a posture most agencies have not yet reached.
- **The December 2022 deadline created a distribution, not uniform compliance.** A small group reached mandate posture operationally; a larger group hit it on paper

only; a meaningful share has not materially moved since 2019.

- **Mandate posture is an architecture, not a document.** Fully meeting the mandate changes the shape of the records architecture — ECM becomes the system of record, classification moves to creation time, disposition automates.
- **AI output is a federal record.** Once an agent produces official content, that content meets the statutory definition of a record and has to be classified, scheduled, and dispositioned — at machine tempo.
- **The cost of upfront compliance is bounded; the cost of delay is not.** Every year of postponement compounds, and AI workloads convert the deferred work from paperwork into architecture.

*"The mandate work and the AI work are now the same work. Agencies that postponed the records architecture are about to do mandate compliance and AI records architecture simultaneously — under deadline pressure they could have avoided by starting on time."*

**A note on figures.** Stage shares and cost-of-delay shapes in this playbook are drawn from FCI Advisory's engagement base across federal records-modernization programs, FY23-Q1 through FY26-Q1.<sup>6</sup> They are illustrative of the directional pattern, not point-precise extracts.

## ● 01 The mandate landscape, sequenced

The federal electronic records mandate is not a single document. It is a layered set of memos, regulations, and NARA-specific bulletins built up over fifteen years, with the load-bearing requirements concentrated in the 2019–2026 window. Reading it as one deadline is the first mistake; the obligations are cumulative.

### THE LOAD-BEARING INSTRUMENTS OF THE FEDERAL RECORDS MANDATE

Instrument	Year	What it establishes
M-19-21	2019	The transition to electronic records: all permanent and temporary records managed electronically by Dec 31, 2022; NARA stops accepting analog transfers. <sup>1</sup>
M-23-07	2022	Update to M-19-21; extends specific elements and clarifies ongoing digital records-management requirements. <sup>2</sup>
36 CFR Part 1236	standing	The regulatory floor for federal electronic records management practice. <sup>3</sup>
UERM / FRMSA	annual	NARA Universal Electronic Records Management requirements and the Federal Records Management Self-Assessment annual report. <sup>4</sup>
Capstone	2013+	The NARA approach to email-records preservation by role, per Bulletin 2013-02 and subsequent guidance. <sup>5</sup>

FCI Advisory analysis of OMB/NARA memoranda and 36 CFR Part 1236.

### Three phases of the mandate

## FIGURE 01 · THE MANDATE TIMELINE

### A fifteen-year build-up, with the load-bearing window in 2019–2026.

Most agencies are still inside the third phase — and the fourth will stack the next NARA cycle on top of unresolved AI-record decisions.



FCI Advisory analysis of OMB/NARA memoranda and 36 CFR Part 1236. The fourth phase is the expensive path for any agency that defers records modernization to a future budget.

The timeline reads in three phases. The **2010–2018** phase was preparation — Presidential directives on records management, early NARA digitization guidance, the foundation of 36 CFR 1236. The **2019–2022** phase was mandate definition — M-19-21 set the operational deadline, agency records officers stood up modernization programs, the December 2022 boundary became visible. The **2023–2026** phase is the one most agencies are still inside — post-deadline compliance posture, the M-23-07 clarifications, ongoing FRMSA reporting, and now the introduction of AI workloads that were not contemplated when the mandate was written.

The phase after this one — 2027 and beyond — is where agencies that have not yet modernized will run into the next NARA memo cycle on top of the unresolved current cycle, and where AI-generated records will compound the legacy backlog. The agencies that are ahead now will be much further ahead by then; the agencies still inside the 2019 posture at that point face the next cycle on top of an unresolved current one.

**What changes after this phase.** The 2027+ window will combine an updated NARA memo cycle with the AI-record disposition decisions agencies are improvising now. Doing both at once, under deadline, is the expensive path — and it is the default for any agency that treats records modernization as a future-budget problem.

## Where federal agencies actually are


### ● 02 Where federal agencies actually are

The mandate's December 2022 deadline did not create uniform compliance. It created a distribution. Across the federal estate, agencies sit in four readiness postures, and the distribution is uneven: a small group hit the mandate operationally, a larger group hit the paperwork, and a meaningful share has not materially moved.

FIGURE 02 · FOUR RECORDS-READINESS POSTURES

#### The deadline created a distribution, not uniform compliance.

*A small group hit the mandate operationally; the largest group is still mid-rollout — the most exposed to AI-workload arrival.*



Distribution of federal agencies across four electronic-records readiness postures as of FY26-Q1. Shares reflect FCI's engagement observation; the directional shape is the point.<sup>6</sup>

#### FOUR FEDERAL RECORDS-READINESS POSTURES (ILLUSTRATIVE DISTRIBUTION)

Posture	Share <sup>6</sup>	What it means operationally
Mandate posture · operational	~12%	Met the Dec 2022 deadline operationally and kept investing afterward. Can absorb AI-generated records on the existing stack.

Posture	Share <sup>6</sup>	What it means operationally
Mandate posture · documentation only	~31%	Hit the deadline at the documentation level. ERM software stood up; classification, Capstone coverage, and legacy backlogs remain inconsistent.
Active modernization rollout	~38%	Programs initiated late, running now. Expecting mandate posture in 18–24 months. Most exposed to AI-workload arrival mid-rollout.
Genuinely stuck	~19%	Records architecture unchanged since 2019. Recurring self-assessment gaps; modernization unfunded or stalled.

Distribution of federal agencies across four electronic-records readiness postures as of FY26-Q1. Shares reflect FCI's engagement observation; the directional shape is the point.<sup>6</sup>

The variance is not random. It correlates with three factors: whether the agency CIO and the agency records officer have a working relationship, whether the agency runs a Documentum-class enterprise content management environment that can be modernized in place (versus a fragmented landscape that has to be consolidated first), and whether records management was treated as a compliance line item or as an operational architecture decision. The agencies that made the operational architecture choice — in 2019 or 2020 — are the ones in the upper postures, and they have not regretted it.

*Readiness posture correlates with whether records management was treated as a compliance line item or as an operational architecture decision. The agencies in the upper postures made the operational choice early and built on it.*

## What the mandate changes structurally

### 03 What the mandate changes structurally

The mandate is not a procedural requirement layered on top of existing architecture. Fully meeting it changes the architecture itself. Mandate posture is a different operating architecture, not the same architecture with new documentation on top — which is why agencies that produced compliance reports without changing operational architecture remain out of posture in practice.

FIGURE 03 · BEFORE VS. AFTER MANDATE POSTURE

#### Mandate posture is an architecture, not a document.

*Completion changes the shape of the records architecture across five layers — the documentation is the artifact, not the goal.*

##### BEFORE MANDATE POSTURE

- **Storage:** paper offsite plus electronic scattered across shared drives and point systems.
- **ECM:** handles some structured records; many live outside it.
- **Classification:** partial; Capstone applied manually to a subset.
- **Disposition:** periodic NARA transfer with manual schedule review.
- **Governance:** records officer and IT operate in separate silos.



##### AFTER MANDATE POSTURE

- **Storage:** ECM is the system of record; paper digitized or scheduled for destruction.
- **ECM:** single environment with explicit classification and retention workflows.
- **Classification:** Capstone at scale, embedded in records creation.
- **Disposition:** automated against NARA schedules; AI records classified at creation.
- **Governance:** records officer and CIO operate against a shared architecture.

Five-layer comparison between pre-mandate and post-mandate federal records postures. Both are ideal-typical; real agency architectures sit on a continuum.<sup>6</sup>

Layer	Before mandate posture	After mandate posture
<b>Storage</b>	Paper offsite plus electronic scattered across shared drives, email archives, and point systems.	ECM is the system of record. Paper backlogs digitized or scheduled for destruction.
<b>ECM</b>	Handles some structured records, not all. Many records live outside ECM.	Single environment with explicit classification, retention, and disposition workflows.
<b>Classification</b>	Partial at best. Capstone applied to a subset of senior officials, manually.	Capstone applied at scale. Classification embedded in records creation.
<b>Disposition</b>	Periodic NARA transfer with manual schedule review.	Automated against NARA schedules; AI-generated records classified at creation.
<b>Governance</b>	Records officer and IT operate in separate silos with episodic coordination.	Records officer and CIO operate against a shared architecture.

Five-layer comparison between pre-mandate and post-mandate federal records postures. Both are ideal-typical; real agency architectures sit on a continuum.<sup>6</sup>

Before mandate compliance, the typical federal records architecture is a constellation: records exist, but their governance is partial. After mandate compliance, the architecture changes shape — the ECM environment becomes the system of record, Capstone is applied at scale, legacy backlogs are digitized or dispositioned, and the records officer and CIO operate against a shared architecture rather than negotiating across silos. The architecture is not optional once the agency has actually completed the work; it is what completion looks like.

**AI changes the architecture again**

The introduction of agentic AI workloads changes the architecture a second time. AI-generated artifacts have to be classified at creation time, not retroactively. Audit trails of AI decisions are themselves federal records. The records officer's role expands to include disposition decisions for AI-produced content that did not exist when the

schedules were written. Agencies still in pre-mandate architectures cannot absorb agentic AI without making both transitions simultaneously — which is the heart of why the mandate work and the AI work have become one program.

## ● 04 When AI output becomes a record

Every AI agent operating in a federal workflow is generating records. The model produces transcripts; the workflow produces decision logs; the reasoning chain produces intermediate outputs; the agent's interactions with federal systems produce a complete audit trail of every prompt, every retrieval, every action. Most of this content meets the federal definition of a record at 44 U.S.C. § 3301 — "all recorded information, regardless of form or characteristics, made or received by a Federal agency under Federal law or in connection with the transaction of public business."<sup>5</sup> The agent did not invent a new category of content; it accelerated the production of an old one.

The federal records definition has not changed. The volume arriving at it has. NARA's framework was built against an operational tempo where records were generated by federal employees taking deliberate actions over hours or days. That bound is gone: an agentic workflow can produce hundreds of decision logs an hour and thousands of reasoning traces in a day. The governance machinery sized for the slower tempo is now receiving records-eligible content at orders of magnitude beyond what it was built for.

### What a federal AI workflow actually produces

The output of a federal AI workflow is not one artifact. It is a cascade — most of which agencies have not yet classified against the records definition.

#### CATEGORIES OF AI-GENERATED ARTIFACT AND THEIR TYPICAL RECORDS TREATMENT

Artifact category	Typical treatment	Records consideration
Final AI responses	Clearly a record	Delivered outputs and formal communications; retain against existing schedules.
Agent decision logs	Gray zone	Document how the agent reached a final action; often records, rarely scheduled.

Artifact category	Typical treatment	Records consideration
Prompt–response pairs	Gray zone	Document operational decisions even if the user never saw them.
Reasoning traces	Often transitory	Intermediate model output; may be transitory, but the determination must be deliberate.
Retrieval queries	Often transitory	Queries the agent later refined or abandoned; usually administrative.
Operational metadata	Administrative	Timestamps, identity assertions, system events; short retention under existing schedules.

FCI Advisory framework, derived from federal AI deployment observation. Classification of gray-zone artifacts depends on the use case, the agency's retention schedule, and whether the artifact documents a federal decision or transaction.

### Three load-bearing schedule–design decisions

- **Records–eligibility classification.** Which categories of AI output are records, which are transitory, and which sit in the gray zone where agency-specific determination is required. Most agencies have addressed final outputs and almost nothing else.
- **The retention destination.** AI artifacts live in vendor platforms, model APIs, and observability tooling by default. Federal records have to live in records-managed environments – typically the agency's Documentum or equivalent ECM, with schedules attached. The pipeline that moves artifacts from operational platforms into records-managed environments needs to exist *before* the artifacts accumulate, not after.
- **The schedule horizons.** Some AI content qualifies for short retention (operational logs); some for longer retention as documentation of federal decisions (decision logs, final outputs); some may qualify for permanent retention under Capstone if it documents senior officials' decision-making.<sup>5</sup> The horizons must be determined deliberately, agency by agency, with NARA-aligned guidance.

**Why this connects to the mandate.** The retention destination for AI-generated records is the same ECM that mandate posture makes the system of record. An agency that has reached mandate posture already has a home for AI records; an agency that has not must build the records architecture and the AI-records pipeline at the same time. The two programs converge on the same infrastructure.

## ● 05 The cost of postponement

Records modernization is the kind of program that is easy to defer in any single budget cycle. It is invisible until it fails. It has no executive sponsor outside the records officer's organization. It returns no visible operational metric. And the deadline that nominally enforces it passed three years ago without immediate consequence. The pattern of deferral is rational inside any single year. It is not rational across the timeline the agency is actually operating against.

### UPFRONT COMPLIANCE VS. DEFERRED COMPLIANCE

Dimension	Upfront compliance	Deferred compliance
<b>Cost shape</b>	Bounded and predictable.	Unbounded; compounds each year.
<b>Duration</b>	Typically 24–36 months at agency scale.	Longer, because backlog grows while the work waits.
<b>Line items</b>	ECM modernization, data-quality remediation, integration, change management.	The same items plus accumulated audit findings and deferred maintenance.
<b>AI interaction</b>	Architecture ready to absorb AI-generated records.	AI workloads convert deferred paperwork into an architecture problem.
<b>Outcome</b>	An auditable records architecture that meets mandate posture.	Two transitions at once, under deadline, at higher total cost.

FCI Advisory framework, derived from federal records program cost observation. The curves diverge most sharply when AI workloads arrive.<sup>6</sup>

#### FIGURE 04 · TWO COST CURVES

### Upfront compliance is bounded. Deferred compliance compounds.

*The curves diverge most sharply when AI workloads arrive — converting deferred paperwork into an architecture problem.*

#### UPFRONT COMPLIANCE

### Bounded & predictable

24–36 months at agency scale. Known line items: ECM modernization, data-quality remediation, integration, change management. Outcome: an auditable architecture at mandate posture.

#### DEFERRED COMPLIANCE

### Unbounded & compounding

Longer, because backlog grows while the work waits. Same line items plus accumulated audit findings and deferred maintenance. AI workloads force two transitions at once, under deadline.

FCI Advisory framework, derived from federal records program cost observation. The total-cost gap favors the upfront path by a meaningful margin across the FCI engagement base.<sup>6</sup>

The cost of upfront compliance is a known quantity — predictable durations, predictable line items, predictable outcomes. It is real, and it is bounded. The cost of delay is unbounded: each year of postponement adds to a legacy backlog that has to be processed eventually, against schedules that continue to apply; each year an agency operates without mandate posture is a year of accumulated audit findings, deferred maintenance on legacy ECM, and missed integration opportunities with the rest of the federal modernization stack. The arrival of agentic AI converts the deferred work from a paperwork problem into an architecture problem. The agencies that delayed in 2020 are paying more in 2026; the agencies that delay through 2026 will pay more in 2028.

## ● 06 A de-risking sequence

The sequence below de-risks the combined mandate-and-AI program for an agency that is in active rollout or documentation-only posture — the two largest postures in the distribution. It treats records modernization and AI-records architecture as one program, because they now are one.

#### STEP 01 · WEEKS 0–8

##### Posture assessment and gap baseline

Establish where the agency actually sits against mandate posture — not against its 2022 documentation. Inventory ECM coverage, Capstone application, legacy backlogs, and the AI workloads already generating records. Output: a single gap baseline both the CIO and records officer sign.

#### STEP 02 · MONTHS 2–4

##### Content consolidation (if fragmented)

Agencies on fragmented content landscapes consolidate before mandate-architecture work begins. Underestimating this phase is the most common reason modernization programs stall. Documentum-class environments shorten this step materially.

#### STEP 03 · MONTHS 3–9

##### ECM as system of record

Make the ECM environment the system of record with explicit classification, retention, and disposition workflows. Apply Capstone at scale. This is the architectural core of mandate posture.

#### STEP 04 · MONTHS 6–12

##### AI-records pipeline and eligibility framework

Build the records-eligibility classification scheme for AI outputs and the pipeline that moves them from operational AI platforms into the records-managed ECM — with NARA-aligned schedules, not vendor defaults.

#### STEP 05 · ONGOING

##### Shared governance, sustained

Stand up the standing function where the records officer and CIO operate against a shared architecture. Records governance, like data quality, decays without it. This is what keeps mandate posture from eroding as new records — human and AI — arrive.

## ● What this rules in and out

Four strategic conditions reshape what federal program leadership should be funding through the next budget cycle:

- **Records modernization and AI program design are now one program, not two.** Agencies that scope these separately end up with AI architectures that generate

uncatalogued records and records architectures that cannot ingest AI output. Federal CIOs and records officers should co-sponsor the same program plan.

- **The mandate posture is the architecture, not the documentation.** The 2027 audit cycle will look at architecture, not at the documentation generated in 2022. Programs scoped to documentation will need to be rescoped to architecture.
- **Documentum and equivalent ECM environments are load-bearing for the transition.** Agencies running Documentum-class ECM have a clearer modernization path. Fragmented agencies should plan for a content-consolidation phase before mandate-architecture work begins.
- **The cost economics favor starting now.** The same work costs more each year it is postponed, and AI workloads compound the curve. Programs initiated this budget cycle land at lower total cost than the same programs initiated in two years.

*The decision is no longer whether to modernize the records architecture — the mandate decides that, and the AI workloads enforce it. The decision is whether to do the work as a deliberate program with predictable cost, or as a forced response to the next audit. The cost is similar either way. The timing changes whether the program delivers.<sup>6</sup>*

### Put this thinking to work.

FCI Advisory helps federal agencies turn records-mandate compliance and AI-records architecture into one delivered program. McLean, Virginia.

## ● Endnotes & sources

1. OMB/NARA Memorandum M-19-21, "Transition to Electronic Records," June 28, 2019. Required agencies to manage all permanent and temporary records in electronic format by December 31, 2022, and directed NARA to stop accepting analog records after that date.
2. OMB Memorandum M-23-07, "Update to Transition to Electronic Records," December 23, 2022. Extended specific elements of M-19-21 and clarified ongoing requirements for digital records management.
3. 36 CFR Part 1236 — Electronic Records Management. The regulatory floor for federal electronic records management practice.

4. NARA Universal Electronic Records Management (UERM) requirements and the Federal Records Management Self-Assessment (FRMSA) annual report.
5. Capstone approach to email-records preservation, as specified in NARA Bulletin 2013-02 and subsequent guidance. The federal definition of "records" is set out in 44 U.S.C. § 3301: "all recorded information, regardless of form or characteristics, made or received by a Federal agency under Federal law or in connection with the transaction of public business." AI-output records considerations also draw on OMB Memorandum M-24-10, "Advancing Governance, Innovation, and Risk Management for Agency Use of Artificial Intelligence," March 28, 2024.
6. Aggregate observations cited in this playbook are drawn from FCI Advisory's engagement base across federal records-modernization programs, FY23-Q1 through FY26-Q1. Stage shares and cost-of-delay figures are illustrative of the directional pattern, not point-precise extracts.