WHITE PAPER

An AI-Enabled Knowledge Layer for HHS Evidence

Conselara Consulting · 2026

Contents

Executive Summary

The U.S. Department of Health and Human Services (HHS) has significant data and evidence assets, but this information is not connected in a way that allows HHS to safely and accurately use AI to facilitate discovery, synthesis, dissemination, and use (HHS, 2025a). Clinical records, research findings, claims data, and public health surveillance exist in separate systems with inconsistent standards. Evidence is difficult to share, validate, and apply. The result: slower science, fragmented decision-making, and missed opportunities to improve health outcomes.

An AI-enabled knowledge infrastructure can help close this gap. Rather than replacing human expertise, AI can help the workforce find, summarize, review, and apply evidence more efficiently. Sustainable AI adoption requires governance, connected infrastructure, workforce training, and feedback loops that allow the system to improve over time. Getting this right is the opportunity. Getting it wrong carries real consequences: when AI systems produce outputs that look valid but are not, they erode public trust and undermine the very evidence base they were meant to strengthen.

This paper makes the case for an AI-enabled knowledge infrastructure as the connective layer HHS needs. The operating model is to connect approved HHS knowledge sources, use AI to retrieve and summarize relevant evidence with source traceability, and require human review, governance controls, and audit logs before outputs inform research, policy, operations, clinical guidance, or public communication. Based on our analysis, we recommend that HHS focuses on four priorities:

- **Build it right.** Establish a named evidence management layer and shared interoperability standards that allow AI to operate safely across HHS data assets.
- **Govern it responsibly.** Treat privacy and security as ongoing lifecycle processes and establish shared frameworks for validating AI-enabled tools before deployment.
- **Grow with it deliberately.** Invest in living AI literacy infrastructure that keeps the workforce current as tools and policies evolve.
- **Test high-value use cases.** Pilot agentic AI systems (AI workflows that automate monitoring and routing while preserving human review) for living systematic reviews to demonstrate measurable, scalable benefits.

Conselara Consulting specializes in the intersection of health evidence infrastructure, AI governance, interoperability, and workforce enablement for federal health agencies. We are prepared to discuss your agency’s priorities, help assess implementation risk, and identify priority use cases. We welcome the opportunity to explore practical next steps for a governed AI evidence management pilot.

SECTION 01

Introduction

The U.S. Department of Health and Human Services (HHS) is intent on pursuing gold standard science to inform decisions that will advance health outcomes for Americans. Using digital technologies, particularly artificial intelligence (AI), is a strategic focus in this pursuit (HHS, 2025a). HHS also recognizes that equipping its workforce with the knowledge and skill to use AI effectively is a necessity for high-quality evidence going forward, with role-based training pathways a key goal across all HHS staff (HHS, 2025a).

Modernizing the evidence generation-to-use pipeline to take advantage of AI effectively is equally essential, since that pipeline is the mechanism by which rigorous, reproducible science is produced, shared, and used. This paper identifies an AI-enabled knowledge infrastructure as a critical missing layer in HHS's modernization efforts, one that, alongside governance reform and workforce development, is necessary to connect the evidence generation-to-use pipeline at scale. The AI-literate HHS workforce will work interdependently with this infrastructure as depicted in Figure 1. The paper's objectives are to discuss significant challenges and how to address them, provide Conselara's relevant experience as proof points, and depict how this infrastructure will evolve. At the core of this infrastructure is a simple operating model: connect approved HHS knowledge sources, use AI to retrieve and summarize relevant evidence with source traceability, and require human review, governance controls, and audit logs before outputs inform research, policy, operations, clinical guidance, or public communication.

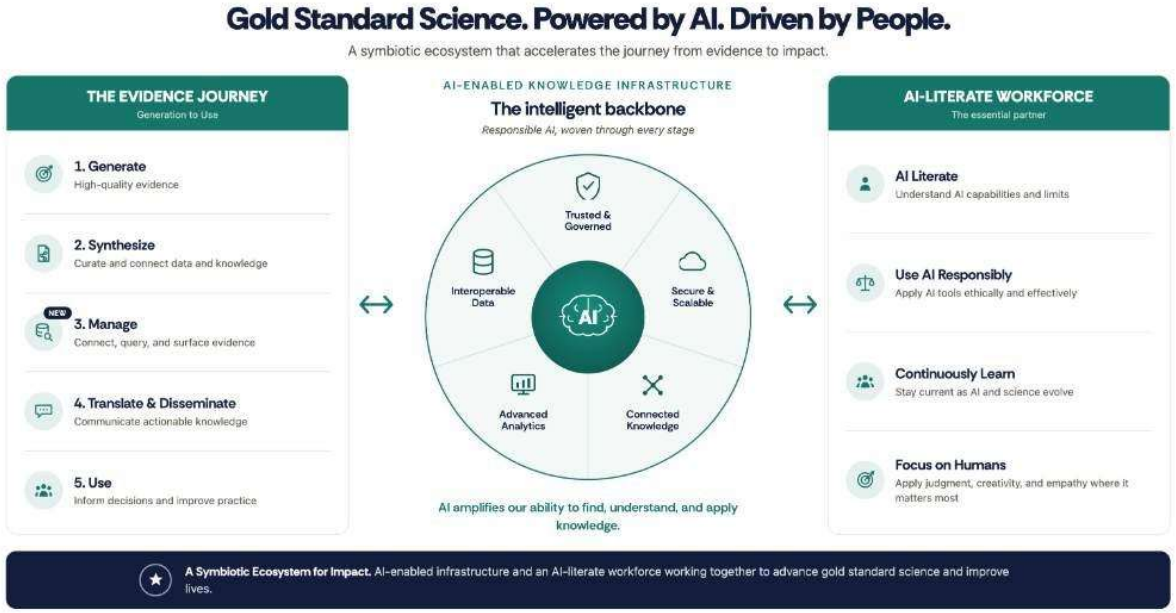

Figure 1. *The AI-Enabled Knowledge Infrastructure as the connective layer in the HHS evidence generation-to-use pipeline.*

SECTION 02

Technical Background & Strategic Context

Gold standard science is reproducible. Reproducibility must apply to any AI models used within the scientific process. AI outputs are only valuable if replicable and verifiable across relevant datasets and contexts. A knowledge infrastructure supports reproducibility by preserving source links, metadata, version history, and review records for AI-assisted outputs.

The advantages of integrating AI into the scientific process are increased scale, timeliness, and precision within the evidence generation-to-use pipeline.

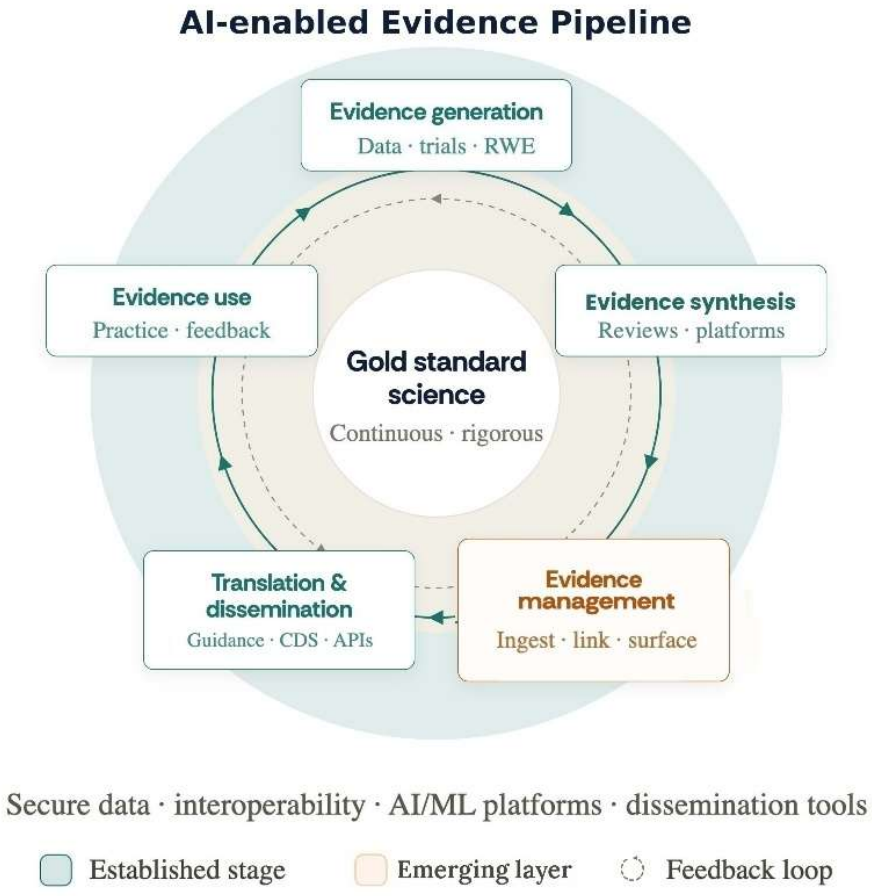

Figure 2. *The HHS AI-enabled evidence pipeline. Five stages operate within a shared AI-enabled knowledge infrastructure; evidence management is an emerging layer not yet established at HHS department scale. The dashed arc represents the feedback loop that defines a learning health system.*

Evidence Generation

AI can assist with hypothesis generation by scanning large datasets, with trial design and recruitment, and can render real-world data more usable for research by structuring it via natural language processing.

Evidence Synthesis

Evidence is synthesized to enable rigor in decision making. Systematic reviews represent the most comprehensive form of evidence synthesis, but creating one takes a long time, and new evidence may make it outdated by publication. Organizations like the Agency for Healthcare Research and Quality (AHRQ) have begun implementing living systematic reviews (LSRs) as a remedy (Chou et al., 2025). Unlike traditional reviews, which represent a snapshot of the evidence at one point in time, LSRs are continuously monitored and updated as new relevant studies are published, keeping guidance current. AI assists LSRs by accelerating initial literature screening and continuously monitoring for new evidence, freeing human reviewers to focus on the judgment-intensive work of synthesis, quality assessment, and interpretation where their expertise remains essential. AI makes the systematic review process faster and more sustainable, enabling the same rigorous human oversight to operate at greater scale and frequency.

Evidence Translation and Dissemination

Creating health guidance based on a systematic review is a key function of translation. Sharing that guidance via integration into clinical workflows, policy development, public health response, program oversight, grants management, and research portfolio monitoring, usually as clinical decision support (CDS), is the critical dissemination step. Evidence must also be disseminated before formal reviews are complete, so researchers stay current. AI enables timely, personalized dissemination targeting the appropriate recipient at the appropriate point in the decision process.

Evidence Use

Dissemination does not always result in evidence use. While not an explicit focus of this paper, evidence use is mentioned here because applying evidence in practice allows a learning health system to improve over time. Data about evidence use provides important feedback for further evidence generation.

Evidence Management

A distinct and emerging capability sits between evidence synthesis and dissemination: evidence management using an AI-enabled infrastructure that ingests, links, and surfaces knowledge across agency systems for rapid discovery and analysis. Evidence management tools are not traditional synthesis platforms (that produce new reviews) nor dissemination tools (that deliver

findings to end users); they are the connective layer that makes knowledge findable and actionable. The Food and Drug Administration's (FDA) Harmonized AI & Lifecycle Operations for Data (HALO) and Evaluation of Life Science Applications (Elsa) 4.0 are early examples at the agency level (FDA, 2026b). HALO consolidates FDA application and submission data and systems; Elsa's integration with HALO allows for AI-enabled data analysis. While HALO and Elsa are FDA-centric, their integration exemplifies the idea of ingesting, linking and surfacing data and knowledge for rapid discovery and analysis enabled by AI.

HHS needs this capability at the department level. Without it, agencies will continue to generate evidence that exists in isolation, difficult for AI systems and decision-makers to locate, assess, and reuse. The volume of health research is growing faster than manual discovery and synthesis processes can keep pace (Pezzullo & Boccia, 2024). AI-enabled evidence management is a practical mechanism by which HHS can ensure its evidence assets are linked, discoverable, and capable of analysis.

HHS does not yet appear to have a fully integrated department-wide, continuous, AI-enabled learning system unifying all steps in the evidence pipeline. HHS is providing the strategic impetus through its OneHHS approach, which aims to eliminate information silos and enable AI-ready access to HHS-funded studies and data, including electronic health records (EHR), administrative healthcare claims data, surveillance, and human services data (HHS, 2025a). Building that infrastructure is the practical step we advocate for throughout this analysis. The urgency extends beyond strategy documents. Active Requests for Information from the Centers for Medicare and Medicaid Services (CMS) on the use of AI tools in Medicare (CMS, 2026), from FDA on real-time AI-enabled clinical trials (FDA, 2026a), and from the Office of the National Coordinator for Health IT (ONC) on accelerating AI adoption in clinical care (ASTP/ONC, 2025) signal that federal investment in AI-driven modernization is no longer theoretical; it is moving toward procurement and implementation. An AI-enabled knowledge infrastructure is the foundational capability needed to support and scale those efforts.

A fully integrated, department-wide AI-enabled evidence management capability does not yet appear to exist. Building it (alongside the governance structures to oversee it and the workforce capacity to use it) is the central challenge this paper addresses.

SECTION 03

Challenges and How to Address Them

The challenges described below are not indictments of current practice; they are the known obstacles that any organization navigating AI adoption must plan for. For each challenge, the recommendations that follow are designed to be implemented proactively, before problems arise, as well as to address gaps that already exist. Addressing these challenges requires more than technology. A well-designed knowledge infrastructure, paired with governance reform and sustained investment in workforce AI literacy, is the combination that makes the difference between AI adoption that scales and AI adoption that stalls.

Cross-Cutting Challenges

These challenges are "cross-cutting" because they affect every stage of the evidence pipeline shown in Figure 2: generation, synthesis, management, translation and dissemination, and use, rather than any single stage. The challenges below are the foundational conditions, data, governance, security, and workforce readiness, that the entire operating model depends on. The Evidence Pipeline Challenges section that follows addresses obstacles specific to one part of that pipeline.

Research data can be fragmented across systems that do not easily connect, and be stored in different formats with uneven quality, making it difficult to use effectively for AI and advanced analysis. Consistent tools are also needed to track and validate how data and AI models are used across programs. Without shared standards and clear documentation, it is difficult to reproduce results, compare findings, or scale successful approaches (NIST, 2020a; OMB, 2022).

Interoperability

The challenge

One of the biggest technology hurdles across HHS is making different data systems work together meaningfully. Each agency collects data for its own mission, meaning the data is not just in different formats, it reflects different ways of thinking. A clinical record captures a patient visit, a coder translates that visit into billing codes for an administrative claim, research data follows strict study designs, and public health data often shows incomplete, real-world trends. Even when systems are connected, these differences make it hard to combine and interpret the data reliably.

The challenge is also about meaning. Core concepts like diseases, outcomes, and risk factors are defined differently across datasets, creating the risk that AI may perform well in one setting but fail in another, producing insights that are difficult to reproduce, compare, or trust.

There is reason for cautious optimism. Several major federal initiatives are collectively moving the field toward a common data language:

Initiative	Agency	What It Does
United States Core Data for Interoperability (USCDI) v7 draft	ONC	Expands standardized data elements health IT systems must support, strengthening the foundation for cross-agency exchange
Trusted Exchange Framework and Common Agreement (TEFCA) security onboarding	ONC	Tightens requirements for joining the nationwide health information exchange network, raising the baseline for secure, trusted data sharing

Initiative	Agency	What It Does
Fast Healthcare Interoperability Resources (FHIR) for regulatory submissions	Food and Drug Administration (FDA)	Accepts FHIR-formatted data in regulatory submissions, signaling federal convergence on modern interoperability standards
Patient Access Application Programming Interface (API) mandates	CMS	Requires health plans to make patient data available through standardized interfaces, enabling more consistent access to claims and clinical information

These initiatives create the conditions for AI-enabled knowledge management, but they do not by themselves give staff a usable way to discover and apply evidence across systems.

The challenge is no longer whether standards exist but managing their consistent implementation across agencies with different missions and legacy systems. That implementation gap is precisely where a well-designed knowledge infrastructure delivers value.

How to address it

HHS's OneHHS vision points in the right direction: true interoperability requires shared definitions, strong metadata, and clear data provenance, not just shared formats (HHS, 2025a; ONC, n.d.-a). Open, widely used approaches like Health Level Seven's (HL7) FHIR standards, guidance from ONC, and the Findability, Accessibility, Interoperability, and Reusability (FAIR) data principles offer practical paths forward (HL7 International, n.d.; ONC, n.d.-a; Wilkinson et al., 2016).

First, HHS can reinforce shared standards. FHIR gives agencies a common API-based way to move data, but standards alone are not enough. Agencies need consistent implementation with agreed-upon definitions, so data means the same thing everywhere (ONC, n.d.-a). Second, HHS can require better data documentation: the FAIR principles provide a clear model, with every dataset accompanied by clear descriptions, common vocabularies, and provenance records that help both people and AI systems understand and trust the data (HL7 International, n.d.). Finally, HHS should treat interoperability as an ongoing cross-agency effort, investing in governance, updating legacy systems, and supporting open, collaborative approaches.

CONSELARA IN PRACTICE: FHIR IMPLEMENTATION AND FEDERAL STANDARDS DEVELOPMENT

Conselara piloted a cancer screening planning application integrated directly into a major health system's Epic EHR using standard FHIR APIs, with no custom integration required. The pilot used Clinical Quality Language (CQL) logic to make prevention recommendations computable and produced concrete findings about the real-world limits of standards-based write-back, exactly the implementation gap described above between standards adoption and operational interoperability.

Conselara also developed and maintained a national repository of federal healthcare data standards including HL7, National Council for Prescription Drug Programs (NCPDP), X12, and Accredited Standards Committee (ASC) standards, and team members contributed to a

federal interoperability standards panel that shaped the national health IT agenda before FHIR's widespread adoption.

Workforce AI Literacy

The challenge

HHS recognizes that its workforce needs the knowledge and skill to use AI effectively, and providing appropriate role-based training across the entirety of HHS is a significant challenge. Uncertainty around using new, AI-enabled knowledge infrastructure tools, how results should be interpreted, and where human review remains essential creates real operational risk. Moreover, as AI continues to evolve, AI guidance, policies, and training materials are also likely to change quickly, making this an ongoing challenge rather than a one-time effort.

How to address it

Employees benefit most when guidance clearly explains what AI knowledge is expected for different roles. Aligning workforce expectations with existing AI strategies and federal human capital reviews helps create a shared understanding and reduces uncertainty about what skills matter most (HHS, 2025a; OMB, 2024).

Stronger coordination between policy owners, training teams, and learning management systems keeps courses current, authoritative, and aligned with published requirements. Federal guidance emphasizes that AI adoption depends on training and workforce capability, yet courses can quickly become outdated as policies evolve (OMB, 2024; Tabassi, 2023).

The most durable investment HHS can make is infrastructure that self-updates. An AI literacy infrastructure that continuously ingests new guidance from OMB, NIST, HHS policy owners, and agency-level directives, and automatically flags outdated training content for review, is qualitatively different from a static course catalog. Just as LSRs keep clinical evidence current, a living AI literacy infrastructure keeps workforce knowledge current.

CONSELARA IN PRACTICE: APPLIED AI LITERACY IN REGULATED FEDERAL OPERATIONS

Conselara has conducted Large Language Model (LLM) comparative evaluations to assess model suitability for federal research and communications tasks, vetted AI-powered platform modules for compliance with federal security and accessibility standards, and piloted ML-driven anomaly detection tools in production federal environments. These experiences, assessing tools for compliance fit and managing the human-AI handoff, directly reflect the workforce readiness challenges federal health agencies face at scale.

Conselara has piloted an AI-enabled knowledge base, applying an ingest-embed-retrieve-synthesize pattern that automatically surfaces updated content when new material is added. This same pattern is directly applicable to the living AI literacy infrastructure HHS needs.

Privacy

The challenge

HHS manages some of the federal government's most sensitive data under laws such as the Health Insurance Portability and Accountability Act (HIPAA). AI tools rely on large, detailed datasets to identify trends, creating tension with rules that limit how much information can be shared or reused (Alder, 2026; HHS, 2025a). Traditional de-identification approaches remain a critical foundation for privacy protection but must evolve alongside AI. Modern AI systems can link multiple data sources and infer identities indirectly, especially for small populations or rare conditions, meaning de-identification techniques that were sufficient for earlier technologies now need to be paired with additional safeguards such as differential privacy, access controls, and audit trails to remain effective (Alder, 2026; NIST, 2020a).

Accountability is equally challenging. Agencies must be able to trace how data were used and how AI outputs were produced, but many existing systems were not built to support detailed data tracking or routine audits. Guidance from the Centers for Disease Control and Prevention (CDC) and CMS highlights the importance of auditability and lifecycle oversight for AI systems that handle sensitive data, but aligning these requirements with legacy technology remains difficult across HHS (CDC, 2026c; CMS, 2025; Tabassi, 2023).

How to address it

HHS can better align privacy and AI by adopting shared, modern data practices that allow information to be used safely without unnecessary restrictions. Building on existing HIPAA protections (HHS, n.d.), this includes department-wide guidance based on the National Institute of Standards and Technology (NIST) AI Risk Management Framework (NIST AI RMF) and the NIST Privacy Framework (NIST, 2020a; Tabassi, 2023), standardizing how agencies handle, share, and protect data so teams are not working from different rules (GSA, n.d.-a).

Agencies should consistently track where data comes from, how it is used, and how AI systems produce results. Maintaining data inventories, documenting model inputs and outputs, and enabling routine reviews make it easier to share and reuse data responsibly (GSA, n.d.-a). Privacy-enhancing technologies, such as secure data environments, federated learning, and differential privacy, can enable analysis without exposing sensitive information, though each carries implementation trade-offs in performance, scalability, and residual risk (ONC, 2025; Tabassi, 2023). HHS also needs updated guidance that reflects how modern AI systems operate, including expectations for audit trails, oversight of multi-step AI workflows, and the use of third-party AI vendors.

> **CONSELARA IN PRACTICE: OPERATING AI ENVIRONMENTS FOR SENSITIVE FEDERAL DATA**
>
> Conselara has completed full Authority to Operate (ATO) processes for government web systems and operated HIPAA-, FISMA-, and NIST-aligned environments across multiple federal health IT programs. This experience informs our understanding of where privacy guidance gaps create real friction for AI adoption, including the audit trail gaps and access

control limitations that make federated approaches operationally attractive but technically complex in legacy federal environments.

For programs where cloud AI is restricted by data residency or authorization requirements, Conselara operates a production on-premises AI inference environment running frontier-class language models with no data transmitted to external APIs.

Security

The challenge

AI systems must operate within the Federal Information Security Modernization Act (FISMA) and Federal Risk and Authorization Management Program (FedRAMP) frameworks, aligned with the NIST SP 800-53 (NIST, 2020b), and CMS policy emphasizes that AI systems must meet federal security standards before handling sensitive data (CMS, 2025). The challenge is that AI systems are not static; they evolve through retraining, updates, and changing data inputs. Traditional one-time approval processes were not designed for this model, but they remain a valuable starting point. The answer is not to abandon them, but to extend them with continuous monitoring and periodic reassessment so that the rigor of the original authorization is maintained as systems change (CDC, 2026c; CMS, 2025; OMB, 2022).

Security is increasingly integrated with governance and responsible AI practices. HHS must ensure AI systems are auditable, transparent, and continuously evaluated. CMS guidance requires accountability for AI outputs and ongoing human oversight, while the HHS AI Strategy underscores the importance of building trust through strong governance and risk management (CMS, 2025; HHS, 2025a).

How to address it

HHS can strengthen AI security by applying existing federal requirements, such as FISMA, FedRAMP, and NIST SP 800-53, in flexible, risk-based ways that reflect how AI systems are developed, updated, and used over time (GSA, n.d.-b; NIST, 2020b). Security should be treated as an ongoing lifecycle process. Continuous monitoring, automated security checks, and zero trust approaches help systems remain protected as they evolve (OMB, 2022). Since AI systems are rapidly evolving, HHS cannot rely on reviewing every individual update. Instead, agencies should focus on establishing strong guardrails, clear limits on how AI can be used, and ongoing oversight processes that keep systems secure, transparent, and aligned with policy goals (NIST, 2020b; Tabassi, 2023).

CONSELARA IN PRACTICE: CONTINUOUS SECURITY IN PRODUCTION FEDERAL ENVIRONMENTS

With more than a decade of experience supporting federal health agencies, Conselara helps clients embed AI into existing operations using proven engineering and security practices, rather than treating it as a standalone effort. Conselara integrates AI into already authorized environments, minimizing disruption while maintaining continuity.

> For one HHS entity, Conselara maintains HIPAA, FISMA, NIST, and Federal Information Processing Standards (FIPS) compliance across its ecosystem, implementing Amazon Web Services (AWS) native tools for secure developer access, with continuous vulnerability scanning and role-based access controls throughout. At another HHS entity, Conselara supported the migration of 70+ enterprise systems to AWS, hardening systems to Center for Internet Security (CIS) Level 2 benchmarks and maintaining continuous readiness for Government Accountability Office (GAO), FISMA, and Federal Information System Controls Audit Manual (FISCAM) audits.
>
> Rather than starting from scratch, agencies can build on existing security authorizations. FedRAMP High cloud services provide ready foundations for AI workloads. Conselara also operates an internal air-gapped AI environment with no data transmitted to external APIs, positioned for sensitive federal environments where data residency and boundary compliance are essential.

Taken together, these cross-cutting challenges reflect a common underlying problem: HHS agencies are operating with systems, standards, and training infrastructure that were not designed for the AI era. Interoperability gaps make it hard to connect the data AI needs. Privacy and security frameworks built for earlier technologies are strained by AI's ability to link and infer across sources. And the workforce lacks a consistent, current foundation for knowing when and how to use AI responsibly. Addressing these challenges is not optional; it is the precondition for the infrastructure and reforms described in the sections that follow.

Evidence Pipeline Challenges

The challenges below are specific to individual stages of the evidence pipeline rather than the system as a whole.

Evidence Synthesis

The challenge

A growing number of AI-enabled tools for literature review, screening, and summarization are becoming available, but HHS lacks a consistent, Department-wide approach to evaluating and validating them. Common evaluation methods, pilot environments, and performance benchmarks are essential to confirm that tools produce accurate and consistent results. Establishing shared methods for testing and comparing tools is critical to meeting HHS goals to accelerate research and ensure decisions are based on reliable, up-to-date evidence (AHRQ, 2022; AHRQ, 2025).

How to address it

HHS can speed adoption of AI evidence synthesis tools by creating a clear, department-wide evaluation and approval process. Building on AHRQ's work, HHS can define standards for accuracy, consistency, transparency, and oversight that tools must meet before use in systematic reviews (AHRQ, 2025). Shared testing environments allow head-to-head comparisons using common datasets and evaluation methods, monitored over time. This is

especially important for LSRs, which can benefit from agentic workflows that automate literature monitoring and routing, continuously screening for new studies, flagging candidates for human review, and triggering the update process without manual re-initiation, while all evidentiary judgments remain with human reviewers. Aligning with the NIST AI RMF ensures reliability as tools evolve (HHS, 2025a; Tabassi, 2023).

Human oversight in agentic LSR workflows requires more than a general commitment to review. Designated human subject matter experts should receive structured queues of studies flagged by the AI agent, with the agent's reasoning made visible at each step so human reviewers can critically assess rather than simply accept the agent's outputs (Tabassi, 2023). All screening decisions should be logged in an auditable trail, and when confidence falls below defined thresholds, the workflow should automatically escalate to human review. These design requirements should be established before pilots launch and treated as evaluation criteria alongside accuracy and efficiency metrics.

A related emerging capability is evidence management infrastructure. FDA's Elsa 4.0, a generative AI platform for scientific and regulatory reviews within a FedRAMP High environment, and HALO, which links data from more than 40 sources for a unified submission history view, demonstrate what shared, well-governed AI-enabled knowledge infrastructure can achieve. Together they show how agencies can reduce manual work, improve consistency, and support faster discovery and decision-making (FDA, 2026b). HHS needs this capability at the department level.

CONSELARA IN PRACTICE: COMPUTABLE EVIDENCE INFRASTRUCTURE AND RESEARCH MONITORING AT SCALE

Conselara developed functional and technical requirements for computable evidence and guidance tools for a federal health agency, including a recommendation summary browser and a rapid-access knowledge portal. In parallel, Conselara implemented automated publication discovery to systematically monitor a portfolio of more than 750 active research projects and surface emerging literature, an operational precursor to the LSR monitoring infrastructure described in the preceding recommendations.

The team also coordinated evidence dissemination across 130+ practice-based research networks serving 67,000+ clinicians and 52.7 million patients, providing direct experience with the scale at which evidence management infrastructure must ultimately operate.

Conselara applies this same approach internally. Argus, an AI-enabled knowledge platform built and operated in-house, gives staff cited, natural-language access to the firm's capabilities, past performance, and expertise in place of institutional memory. A parallel demonstration extends the same pattern to the 750-plus project portfolio referenced above, spanning more than 8,000 records in total, and returns cited answers and exact analytics rather than search results alone. Both systems operate today, illustrating at working scale the infrastructure this paper recommends building department-wide.

Evidence Translation & Dissemination

The challenge

HHS agencies are generating strong evidence across healthcare, research, and public health programs, but there is increasing need to present it in formats that are clear, timely, and actionable. Agencies often rely on manual processes and static files to distribute findings, which can delay updates and create version control issues. Limited use of modern tools like APIs and automated reporting makes it harder to keep information current and widely accessible. Expanding shared dissemination tools, standard formats, and cross-agency coordination mechanisms would help HHS deliver clearer, more timely information and maintain public trust (HHS, 2025a). Without evidence management infrastructure connecting the earlier stages of the pipeline, the translation and dissemination problem is compounded: agencies cannot effectively share what they cannot first find, validate, or synthesize at scale.

How to address it

HHS can improve how research is used by modernizing how findings are presented. Shared dashboards, automated updates, and plain-language summaries that highlight actionable conclusions help ensure research informs decisions in real time (HHS, 2025a). Expanding the use of common formats, APIs, and reusable tools reduces delays and makes it easier to distribute updates as new findings emerge, consistent with 21st Century Cures Act interoperability goals (ONC, n.d.-a).

Shared platforms, common dissemination standards, and cross-agency coordination reduce duplication and support consistent communication, supporting better decision-making, stronger program performance, and increased public trust (GSA, n.d.-a; Tabassi, 2023).

Collectively, these recommendations advocate for a direction, not a blueprint, and several important questions remain open. They do not address funding mechanisms, resolve inter-agency authority questions, or empirically demonstrate return on investment at HHS department scale. These are not reasons to delay; they are precisely the questions a cross-agency working group, as recommended in Section 05, should be chartered to answer.

> **CONSELARA IN PRACTICE: EVIDENCE DISSEMINATION AT NATIONAL SCALE**
>
> For an HHS agency, Conselara manages and disseminates information from a portfolio of 750+ funded research projects, developing digital platforms, translating complex findings into accessible content for clinicians, researchers, policymakers, and the public, and facilitating national webinars that regularly draw 500 to 1,500 or more participants from across the health IT ecosystem.
>
> Conselara has also integrated AI into dissemination workflows, including automated publication discovery and LLM-assisted drafting with program team review at each step.

In aggregate, the evidence pipeline challenges point to the same gap: HHS is generating more research than its current tools can reliably synthesize, validate, and deliver to the people who need it. Closing that gap requires not just better technology, but shared standards for evaluating

AI tools, infrastructure that keeps evidence moving through the pipeline in near-real time, and a commitment to getting findings into the hands of clinicians, policymakers, program leaders, researchers, public health officials, and the public before they become outdated.

SECTION 04

From Current State to Future Vision

The evidence generation-to-use pipeline is at an inflection point. The connective layer that would allow these investments to compound, a shared AI-enabled knowledge infrastructure supported by governance and workforce development, is what we advocate for building. AI is beginning to reshape each stage, but progress is uneven and the pipeline does not yet flow seamlessly with appropriate feedback loops. Some capabilities are operational today; others require foundational infrastructure investments before they can scale. The table below summarizes where HHS agencies stand now, what is achievable in the near term with the right AI-enabled infrastructure in place, and the longer-term vision the field is moving toward. Human oversight remains essential across all stages, requiring an AI-literate workforce supported by automatically updated learning infrastructure.

Focus area	Current state	Near term	Future vision
Data	Fragmentation in data exists with pockets of interoperability. Privacy and security infrastructure is also siloed typically by individual HHS agency.	Work is progressing on interoperability, privacy, and security infrastructure across HHS agencies; use of real-world data is becoming more common in evidence generation.	Interoperability, privacy, and security issues are well-handled across agencies and real-world data are incorporated into evidence generation.
Evidence synthesis, management, and translation	Evidence synthesis is largely manual and evidence discovery via evidence management systems is just beginning.	Evidence synthesis and evidence management systems are being used and piloted in more agencies. LSRs are growing but automating translation and dissemination is still a work in progress.	Evidence management systems are integrated across HHS for efficient evidence ingestion and analysis. Evidence updates based on LSRs move through governed workflows into appropriate clinical, policy, program, or public health channels with appropriate reviews.
Evidence use	Little to no integration of evidence use and outcomes into the evidence generation-to-use pipeline.	Evidence use and outcomes are not yet fully integrated but pilots are initiated.	The evidence generation-to-use pipeline has emerged into a learning ecosystem where use and outcomes inform further generation.

Focus area	Current state	Near term	Future vision
AI literacy	Infrastructure for workforce AI literacy that includes auto-updated role-based training pathways does not exist.	Infrastructure for workforce AI literacy that includes auto-updated role-based training pathways is being piloted.	Role-based training pathways for the AI-literate workforce are automated with appropriate review as AI evolves.

SECTION 05

Call to Action

The infrastructure described in this paper is not theoretical for Conselara. We are actively developing and testing evidence management components in our own R&D environment, including retrieval-augmented generation pipelines, automated literature monitoring, and knowledge synthesis workflows operating in an air-gapped, on-premises AI system. A focused pilot with a federal health agency partner is a concrete, near-term possibility.

To advance this vision, Conselara suggests HHS consider the following actions:

- A cross-agency AI-enabled knowledge infrastructure working group under the OneHHS initiative
- Department-wide guidance on applying the NIST AI Risk Management Framework to evidence synthesis, evidence management, and dissemination tools
- Cross-agency LSR pilots that test agentic workflows for automated literature monitoring, human-reviewed candidate flagging, and routing of updates into clinical workflows
- A living AI literacy infrastructure with auto-updating, role-based training pathways tied to OMB, NIST, and HHS policy changes
- Shared APIs, evidence management standards, and dashboard formats for evidence dissemination across HHS agencies

These investments, taken together, create the connective infrastructure for a continuously learning evidence ecosystem in which data, evidence synthesis, dissemination, and feedback flow more seamlessly across HHS agencies and programs.

For agencies ready to take a focused first step, Conselara has created a starting framework for a near-term pilot. First, identify a high-value evidence workflow, such as LSR monitoring, research portfolio synthesis, or evidence-based guidance retrieval. Second, inventory approved sources and define access, privacy, security, and review rules. Third, build a limited retrieval and summarization prototype within an authorized environment. Fourth, evaluate accuracy, source traceability, time saved, and user confidence against defined benchmarks. Finally, plan for scale only after pilot evidence supports it.

SECTION 06

About Conselara

Conselara supports federal health agencies at the intersection of evidence infrastructure, interoperability, secure cloud engineering, AI governance, and dissemination. As a minority-owned small disadvantaged business, we provide services through established contract vehicles such as the General Services Administration (GSA) Multiple Award Schedule (MAS), making it easy for partners to engage with us. Our work spans the full evidence generation-to-use pipeline, from interoperability infrastructure and secure cloud engineering to evidence dissemination, digital health platforms, and AI-enabled knowledge systems. We focus on aligning mission objectives, business priorities, and technology capabilities to deliver results that are both effective and sustainable, drawing on more than a decade of hands-on experience across the federal health IT ecosystem.

CONNECT WITH US

Conselara Consulting · 3 Bethesda Metro Center, Suite 700, Bethesda, MD 20814

301-961-1695 · info@conselara.com · conselara.com

GSA Multiple Award Schedule: 47QTCA22D0051

CAGE Code: 7K3M4 · UEI: Q77DLNA8FLX7

Designations: Small Business · Minority-Owned Small Business · Small Disadvantaged Business

References

1. Agency for Healthcare Research and Quality. (2022, September). *Inclusion of nonrandomized studies of interventions in systematic reviews of intervention effectiveness: An update.* Methods Guide for Effectiveness and Comparative Effectiveness Reviews. https://effectivehealthcare.ahrq.gov/products/collections/cer-methods-guide
2. Agency for Healthcare Research and Quality. (2025, August). *Evidence-based practice center reports.* https://www.ahrq.gov/research/findings/evidence-based-reports/index.html
3. Alder, S. (2026, February 8). *HIPAA, healthcare data, and artificial intelligence*. The HIPAA Journal. https://www.hipaajournal.com/hipaa-healthcare-data-and-artificial-intelligence/
4. Centers for Disease Control and Prevention. (2026c, March 13). *CDC AI Strategy.* https://www.cdc.gov/ai/strategy/index.html
5. Centers for Medicare and Medicaid Services. (2025). *Guidance for responsible use of artificial intelligence (AI) at CMS.* https://security.cms.gov/policy-guidance/guidance-responsible-use-artificial-intelligence-ai-cms
6. Centers for Medicare and Medicaid Services. (2026). *Request for information: AI tools for Medicare experience modernization.* https://sam.gov/workspace/contract/opp/570eab65b31b4c93b556185df473266c/view
7. Chou, R., Ahmed, A. Y., Dana, T., Morasco, B. J., Bougatsos, C., Fu, R., et al. (2025). *Living systematic review on cannabis and other plant-based treatments for chronic pain: 2025 update.* AHRQ Publication No. 25-EHC036. https://effectivehealthcare.ahrq.gov/products/plant-based-chronic-pain-treatment/living-review
8. HL7 International. (n.d.). *Interoperability. The Standard:* The Official Blog of HL7. https://blog.hl7.org/topic/interoperability
9. National Institute of Standards and Technology. (2020a). *Privacy framework (Version 1.0).* https://www.nist.gov/privacy-framework
10. National Institute of Standards and Technology. (2020b). *Security and privacy controls for information systems and organizations* (NIST SP 800-53, Rev. 5). https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final
11. Office of Management and Budget. (2022, January 26). *Moving the U.S. government toward zero trust cybersecurity principles* (OMB M-22-09). https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf
12. Office of Management and Budget. (2024, March 28). *Advancing governance, innovation, and risk management for agency use of artificial intelligence* (OMB M-24-10). https://www.whitehouse.gov/wp-

content/uploads/2024/03/M-24-10-Advancing-Governance-Innovation-and-Risk-Management-for-Agency-Use-of-Artificial-Intelligence.pdf

13. Office of the National Coordinator for Health Information Technology. (n.d.-a). *Introduction: HL7 FHIR. Interoperability Standards Platform.* https://isp.healthit.gov/fhir-introduction-0
14. Office of the National Coordinator for Health Information Technology. (2025*). Privacy and security.* https://healthit.gov/privacy-security
15. Pezzullo, A. M., & Boccia, S. (2024*). Growth and challenges in biomedical scientific publishing: the phenomenon of mega-journals.* European Journal of Public Health, 34(Supplement 3), ckae144.677. https://doi.org/10.1093/eurpub/ckae144.677
16. Tabassi, E. (2023). *Artificial intelligence risk management framework (AI RMF 1.0) (NIST AI 100-1).* https://doi.org/10.6028/NIST.AI.100-1
17. U.S. Department of Health and Human Services. (2025a). *HHS artificial intelligence (AI) strategy.* https://www.hhs.gov/sites/default/files/hhs-artificial-intelligence-strategy.pdf
18. U.S. Department of Health and Human Services, ASTP/ONC. (2025, December 23). *Request for information: Accelerating the adoption and use of artificial intelligence as part of clinical care.* Federal Register. https://www.federalregister.gov/documents/2025/12/23/2025-23641/request-for-information-accelerating-the-adoption-and-use-of-artificial-intelligence-as-part-of
19. U.S. Department of Health and Human Services, Office for Civil Rights. (n.d.). *Health information privacy*. https://www.hhs.gov/hipaa/index.html
20. U.S. Food and Drug Administration. (2026a, April 28). *FDA announces major steps to implement real-time clinical trials.* https://www.fda.gov/news-events/press-announcements/fda-announces-major-steps-implement-real-time-clinical-trials
21. U.S. Food and Drug Administration. (2026b, May 6). *FDA expands AI capabilities and completes data platform consolidation*. https://www.fda.gov/news-events/press-announcements/fda-expands-ai-capabilities-and-completes-data-platform-consolidation
22. U.S. General Services Administration. (n.d.-a). *Data.gov.* https://www.data.gov
23. U.S. General Services Administration. (n.d.-b). *FedRAMP.* https://www.fedramp.gov
24. Wilkinson, M. D., Dumontier, M., Aalbersberg, I. J., Appleton, G., Axton, M., Baak, A., et al. (2016). *The FAIR guiding principles for scientific data management and stewardship*. Scientific Data, 3, Article 160018. https://doi.org/10.1038/sdata.2016.18