Subsalt

WHITE PAPER

Health Data Governance Wasn't Built for AI

How de-identification unlocks scalable health data access for enterprise AI

JULY 2026

Table Of Contents

Introduction

A SUBSALT WHITEPAPER

Across research, operations, and analytics, manual approvals cannot scale to meet the volume and velocity of AI-generated data requests — so without a different approach, the gap between data demand and access will grow with every new AI use case.

HIPAA's Privacy Rule protects identifiable patient data through purpose-based use limitations and minimum-necessary access. Health systems enforce these restrictions through project-based governance review processes. That model was built to serve hundreds of data and research projects per year; analytical AI tools have the potential to generate hundreds of requests per day. Across research, operations, and analytics, manual approvals cannot scale to meet the volume and velocity of AI-generated data requests — so without a different approach, the gap between data demand and access will grow with every new AI use case.

De-identification is the pathway HIPAA provides to resolve the governance bottleneck. De-identified data is excluded from HIPAA's definition of protected health information, so it falls outside the Privacy Rule's restrictions. Therefore, health systems can enable broad, permissive access to de-identified data for AI tools and agents without sacrificing patient privacy or risking HIPAA non-compliance.

However, HIPAA's two de-identification pathways have not been operationally viable for this purpose. The Safe Harbor method requires removal of important dates and geographic detail. Expert Determination, the other acceptable method, is often too expensive and time-consuming to deploy for routine internal data access. In both pathways, degradation of statistical signal and static, inflexible data models pose operational challenges to de-identification as an internal capability for health systems.

Advances in privacy-preserving data techniques have changed the economics of de-identification. It has become possible to quantify residual re-identification risk of a de-identification process across datasets, unlocking on-demand de-identification backed by Expert Determination. By combining synthetic data and a secure runtime for calculating metrics, health systems can now provide access to full-granularity patient data sets without sacrificing the precision required to publish or inform treatment decisions. This turns a multi-month data access queue into a query-time privacy approval.

The AI Mandate Has Outpaced Operational Readiness

Shrinking margins and threatened funding cuts have pushed health systems to adopt enterprise AI tools quickly, but almost half of health systems aren't operationally ready.

Americans spend $5.3 trillion per year on healthcare — 18 percent of GDP.[1] But median operating margins at US hospitals fell to -0.6% in early 2026.[2] The lapse of enhanced federal insurance subsidies, and enacted Medicaid reductions projected at $911 billion over ten years, will exacerbate these financial challenges.[3,4] As a result, 84% of hospital and health system executives cite financial pressure as their biggest threat, and 72% list "lowering operational costs" as their most important long-term priority.[5]

In response, health systems have turned to technology, tripling AI investments in 2025 to chase operational efficiencies. In one survey of hospital CIOs, 94% said falling behind on AI would put their organization at a competitive disadvantage.[6]

While financial survival drives the executive mandate, the clinical frontline is pushing for AI to solve its workforce crisis: With nearly two-thirds of US physicians citing bureaucratic tasks such as charting as the primary reason for burnout,[7] ambient AI scribes have spread faster than any prior clinical AI: The Permanente Medical Group's physicians used them across more than 2.5 million patient encounters in their first 15 months, saving nearly 16,000 hours of documentation time,[8] and a six-system study found clinician burnout fell from 52% to 39% within 30 days of adoption.[9] At the same time, operational leaders are deploying machine learning to forecast patient throughput, predict length of stay, and optimize bed availability.

Yet only 52% of health systems consider themselves operationally ready to integrate AI,[10] and the governance processes built to protect patient privacy and establish HIPAA compliance conflict with the data requirements of AI tools and agents.

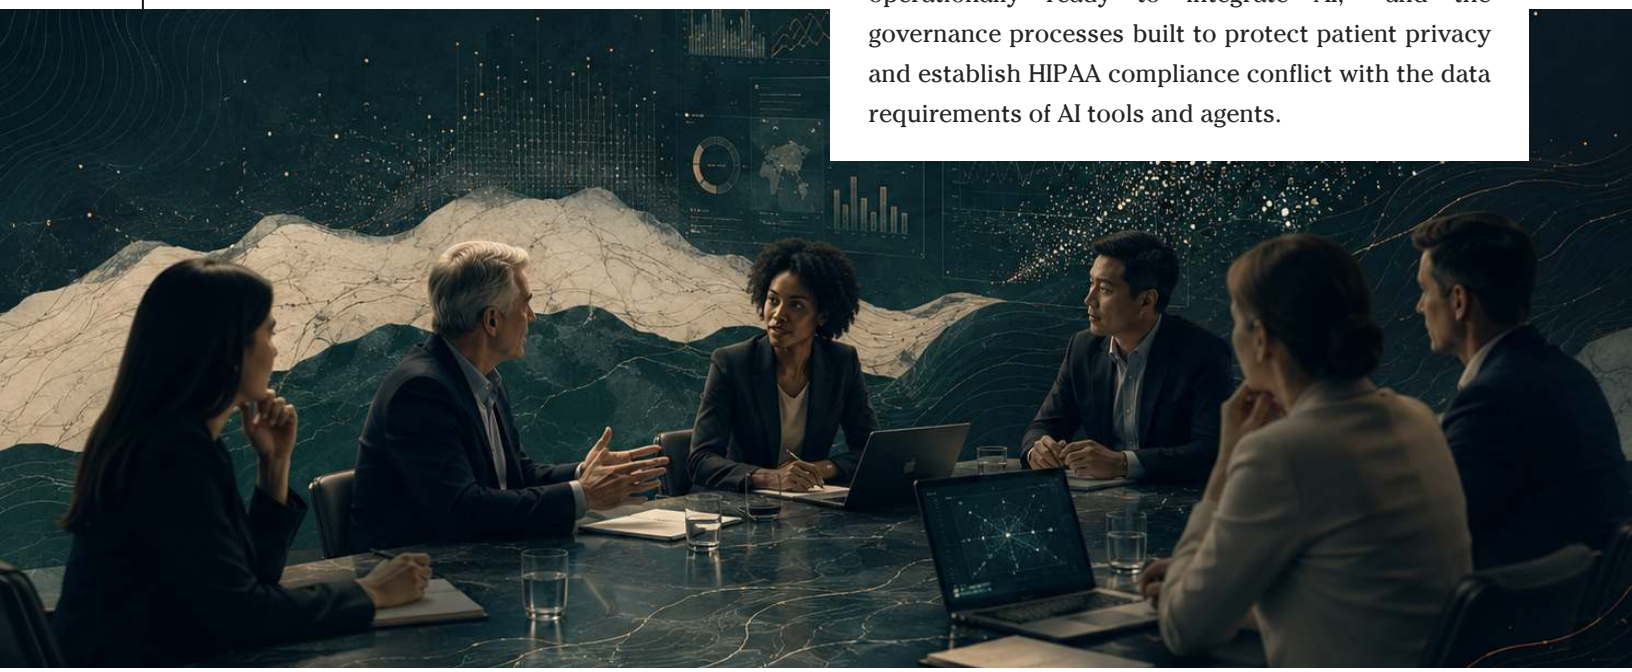

Agentic AI Magnifies AI Data Access Challenges

Enterprise AI tools and agents rapidly increase the number of health system employees that can productively analyze enterprise data, which increases data request volume.

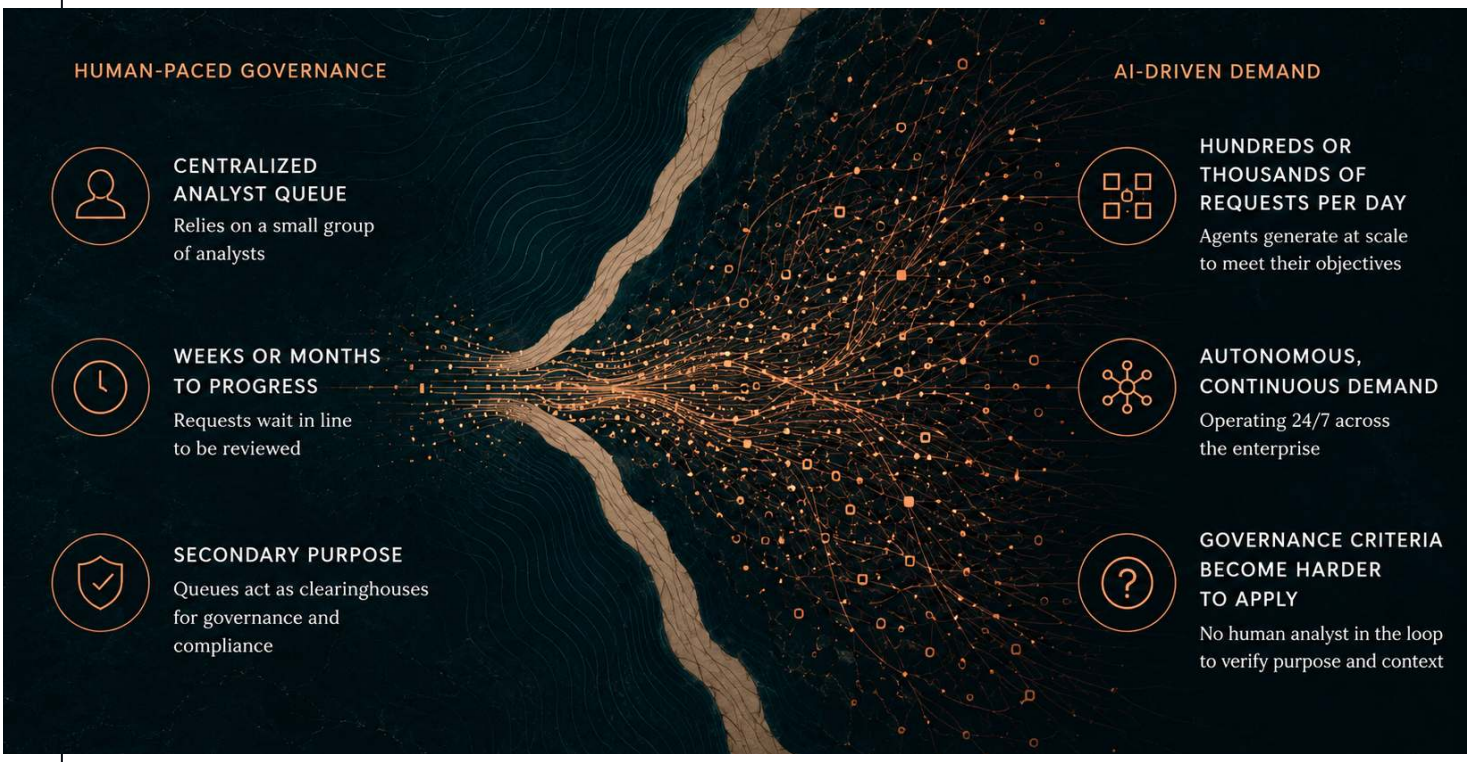

Access to patient data inside health systems has been heavily centralized: a relatively small group of analysts review and service requests to analyze patient data. As a result, subject matter experts building strategies to reduce readmissions or model the impact of a new payer contract often wait weeks or months for their requests to progress to the front of an analyst's project queue.

AI data science capabilities allow subject matter experts to analyze patient data directly without help from an analyst. This has the potential to speed up data projects, but it would also make it harder to enforce data use controls.

That's because analytics queues serve a secondary purpose as centralized clearinghouses for governance and compliance processes.

Agentic AI will have an exponentially greater impact on the challenges posed by AI data access. Deloitte expects half of enterprises currently using generative AI to pilot autonomous agents by 2027,[11] which will strengthen the need for high-throughput data governance processes. Agentic systems can generate hundreds or thousands of data requests per day to fulfill their objective. Without a human analyst in the loop, the provenance and justification for these requests will be difficult to translate into traditional data governance criteria.

To support AI and agentic workloads, health systems must either increase the throughput of the clearinghouse, or build a distributed data governance model to serve distributed data analysis tools. But HIPAA makes it near-impossible to scale centralized clearinghouses through automation.

Unlike data security controls, which can be observed and verified through machine logs, a data request cannot be deterministically mapped to an acceptable purpose.

Why Agentic AI Breaks Purpose-Based Governance Models

HIPAA's purpose-based data privacy framework makes it hard to automate data access and data governance workflows.

HIPAA protects patient data through two regimes: the Security Rule and the Privacy Rule. The Security Rule, much like the control catalogs of SOC 2 or HITRUST, requires specific technical and organizational cybersecurity controls to protect patient data. The Privacy Rule restricts the use of protected patient data ("PHI") by purpose.[12] For example, a covered organization may use or disclose patient information for treatment, payment, or operations (TPO), but may not sell patient data without explicit authorization from the individual.

Unlike data security controls, which can be observed and verified through machine logs, a data request cannot be deterministically mapped to an acceptable purpose.

The problem is that a request's purpose is a function of the requester's intent, which is not observable. Without a mechanism for automated purpose detection, there is no reliable way to scalably review HIPAA compliance for new data requests.

The Privacy Rule also limits data use and disclosure to the "minimum necessary" data required to fulfill the user's purpose.[13] To enforce this requirement, governance teams create a dataset that fulfills a requester's needs, which requires evaluation of the user's goals and the enterprise's data assets. This second contextual evaluation makes it even harder to automate data approvals for HIPAA-covered data.

To illustrate the challenge, a single analysis — for example, the 30-day readmission rate for heart-failure patients, by payer — can be permitted, restricted, or prohibited depending on the user's intent and use case. No mapping of permissions — by user, role, or data requested — can reliably enforce a purpose limitation.

Applied to enterprise AI tools, this means that most prompts from most users will require individualized review before touching PHI. Applied to AI agents, even prompt-based reviews will often be insufficient.

Given an objective like "find opportunities to reduce cost," an agent may spawn multiple analyses that span permitted and prohibited purposes. Reviewing utilization data by service line: permitted. Reviewing health system employees' medical records to inform layoffs: prohibited. So approving an agent's objective serves as a "maximum necessary" data authorization, inverting the Privacy Rule's requirements and authorizing prohibited sub-projects.

Approving an agent's objective serves as a "maximum necessary" data authorization, inverting the Privacy Rule's requirements and authorizing prohibited sub-projects.

How AI Creates A Governance Math Problem

Human review scales by staffing; AI demand scales by usage, prompts, and agents.

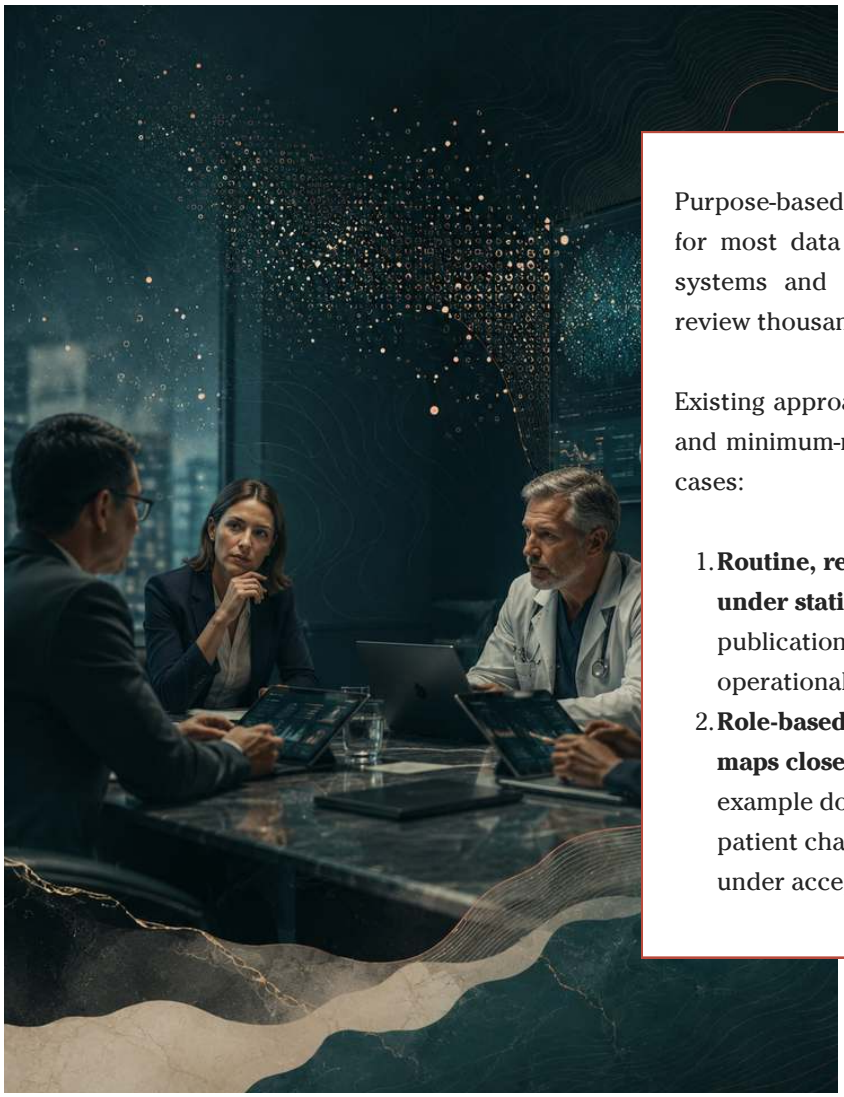

Purpose-based limitations require purpose evaluation for most data requests generated by AI, but health systems and hospitals don't have the capacity to review thousands of prompts per day.

Existing approaches to managing purpose restrictions and minimum-necessary data access work well in two cases:

1. **Routine, recurring disclosures of patient data under static, approved protocols;** for example, publication of annual financial reports or operational dashboards; and
2. **Role-based access policies for users whose role maps closely to a single, acceptable purpose;** for example doctors and nurses need to access patient charts during an appointment, which falls under acceptable use for treatment.

Ambient scribe deployment fits a role-based access pattern — health systems deploy ambient scribe products into treatment-specific workflows where they inherit an unambiguous purpose from their users, doctors and nurses. But many AI workloads don't meet either criterion, including the rollout of tools like OpenAI's ChatGPT and Anthropic's Claude.

Advocate Health, one of the largest nonprofit health systems in the country, intends to roll out Anthropic's suite of AI tools to all of its 167,000 employees. This includes 90,000 employees who are not doctors, nurses, or medical staff.[14] This kind of enterprise-wide rollout of general-purpose AI tools generates use cases that are non-recurring, dynamic, and generated

by employees whose roles cannot be easily mapped to a single acceptable purpose.

The most analogous existing process for reviewing non-routine, non-recurring data access in healthcare is research data governance. But throughput of these systems is low: Cedars-Sinai's honest-broker committee produced 487 data extracts for researchers from 2018 to 2023, with annual throughput growing from 51 projects to 121, and per-extract labor — data engineering, cohort definition, quality assurance, and review together — falling from 12.5 hours to 9 hour work year.[15]

An AI rollout in a large health system creates at least 10,000 new potential data requestors — in the largest systems, closer to 90,000. If just one percent of those users ask a single question a day that benefits from access to patient data — 100 reviewable requests, a floor, not a forecast — extending today's nine-hour review unchanged would require 158 full-time reviewers (assuming requests arrive every day of the year and a 2,080-hour work year).

Even if per-request labor could be cut to a single hour, the queue would still require 18 full-time staff dedicated to triaging AI data requests.

No one, including the health systems that run these committees, proposes putting individual AI prompts through an honest-broker review. That is precisely the point: the current model has no unit of review small enough for a prompt.

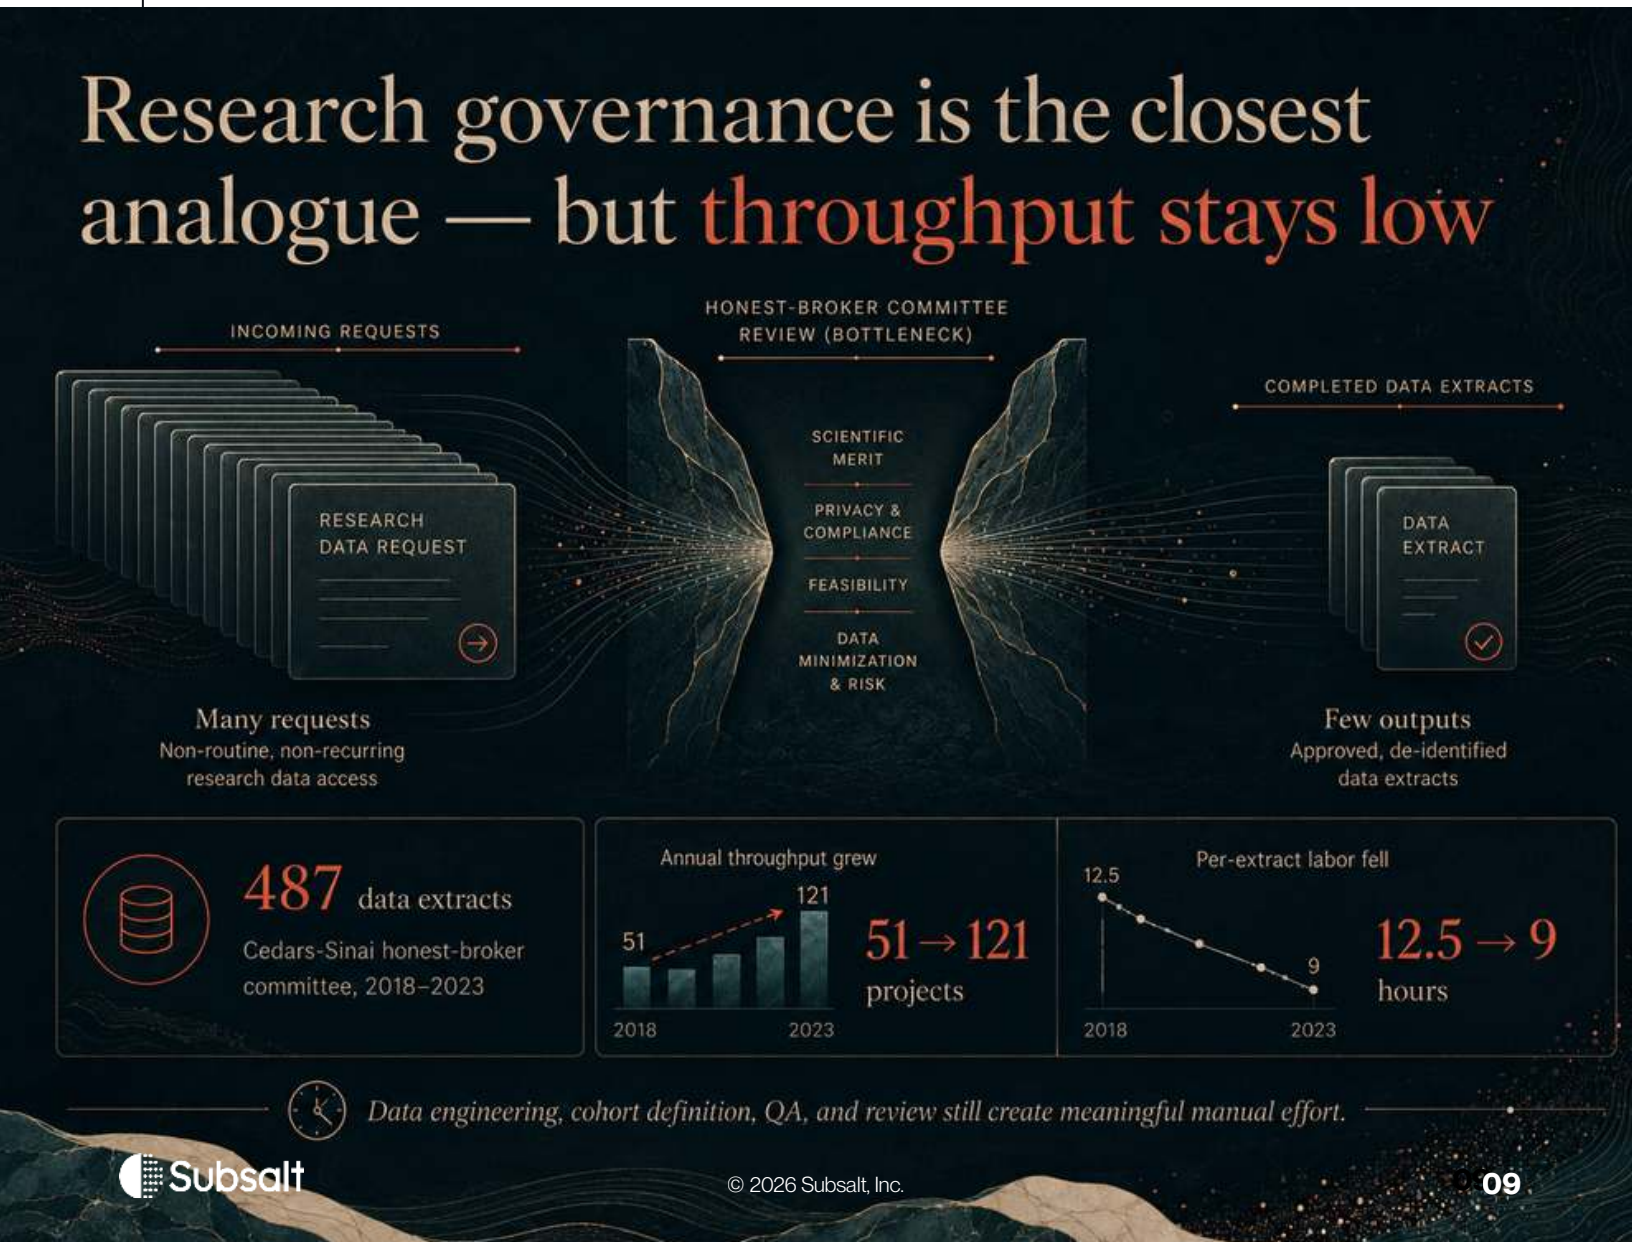

How De-Identification Decouples Data Access from Purpose Limitations

HIPAA provides a scalable data governance mechanism for new AI workloads: de-identification.

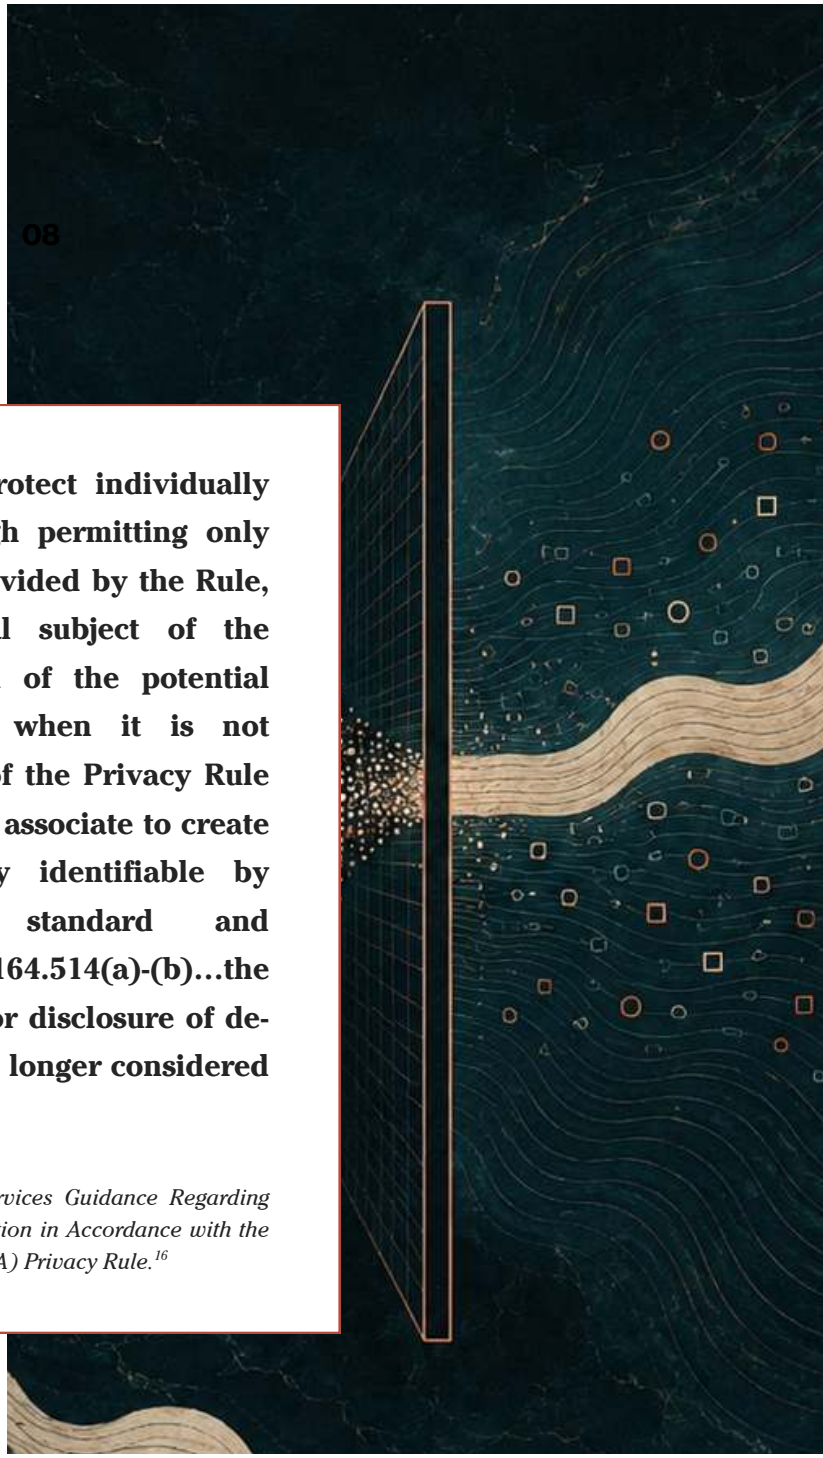

HIPAA's Privacy Rule does not restrict de-identified data:

> **The Privacy Rule was designed to protect individually identifiable health information through permitting only certain uses and disclosures of PHI provided by the Rule, or as authorized by the individual subject of the information. However, in recognition of the potential utility of health information even when it is not individually identifiable, §164.502(d) of the Privacy Rule permits a covered entity or its business associate to create information that is not individually identifiable by following the de-identification standard and implementation specifications in §164.514(a)-(b)…the Privacy Rule does not restrict the use or disclosure of de-identified health information, as it is no longer considered protected health information.**
>
> *— United States Department of Health and Human Services Guidance Regarding Methods for De-identification of Protected Health Information in Accordance with the Health Insurance Portability and Accountability Act (HIPAA) Privacy Rule.*[16]

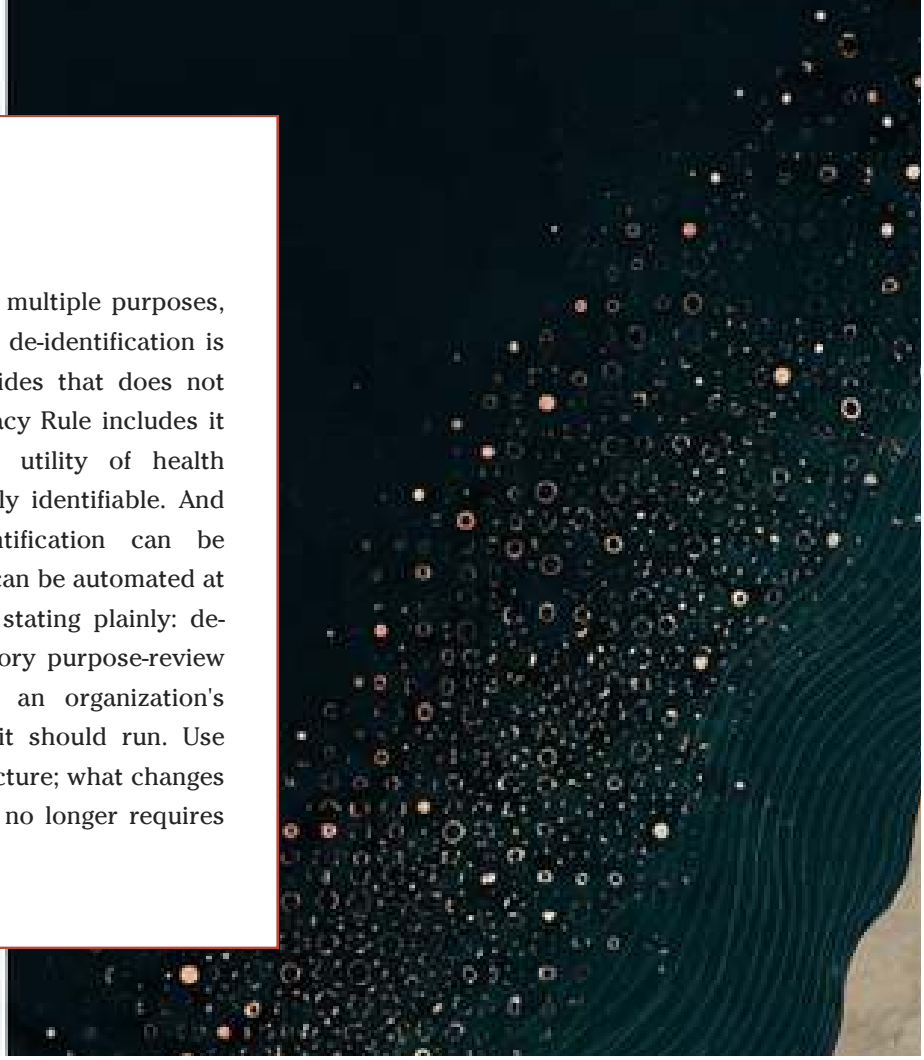

De-Identification is the Only Mechanism

For analytical work that can span multiple purposes, roles, and permission boundaries, de-identification is the only mechanism HIPAA provides that does not require purpose review. The Privacy Rule includes it precisely in recognition of the utility of health information that is not individually identifiable. And unlike purpose review, de-identification can be programmatically evaluated, so it can be automated at runtime. One boundary is worth stating plainly: de-identification resolves the regulatory purpose-review bottleneck; it does not replace an organization's judgment about which analyses it should run. Use policy still sits above the infrastructure; what changes is the foundation below it, which no longer requires human review of every request.

Why Traditional De-Identification Models Fail the AI Test

HIPAA's two approved methods of de-identification have never scaled into reliable, everyday governance infrastructure.

HIPAA provides two acceptable methods for data de-identification: the Safe Harbor (§164.514(b)(2)) and Expert Determination (§164.514(b)(1)).[17]

Safe Harbor

The Safe Harbor method defines a set of eighteen prohibited data elements that must be removed from any identifiable dataset to achieve de-identification. Because it defines specific rule-based data transformations to de-identify a dataset, Safe Harbor is easy to implement and widely used.

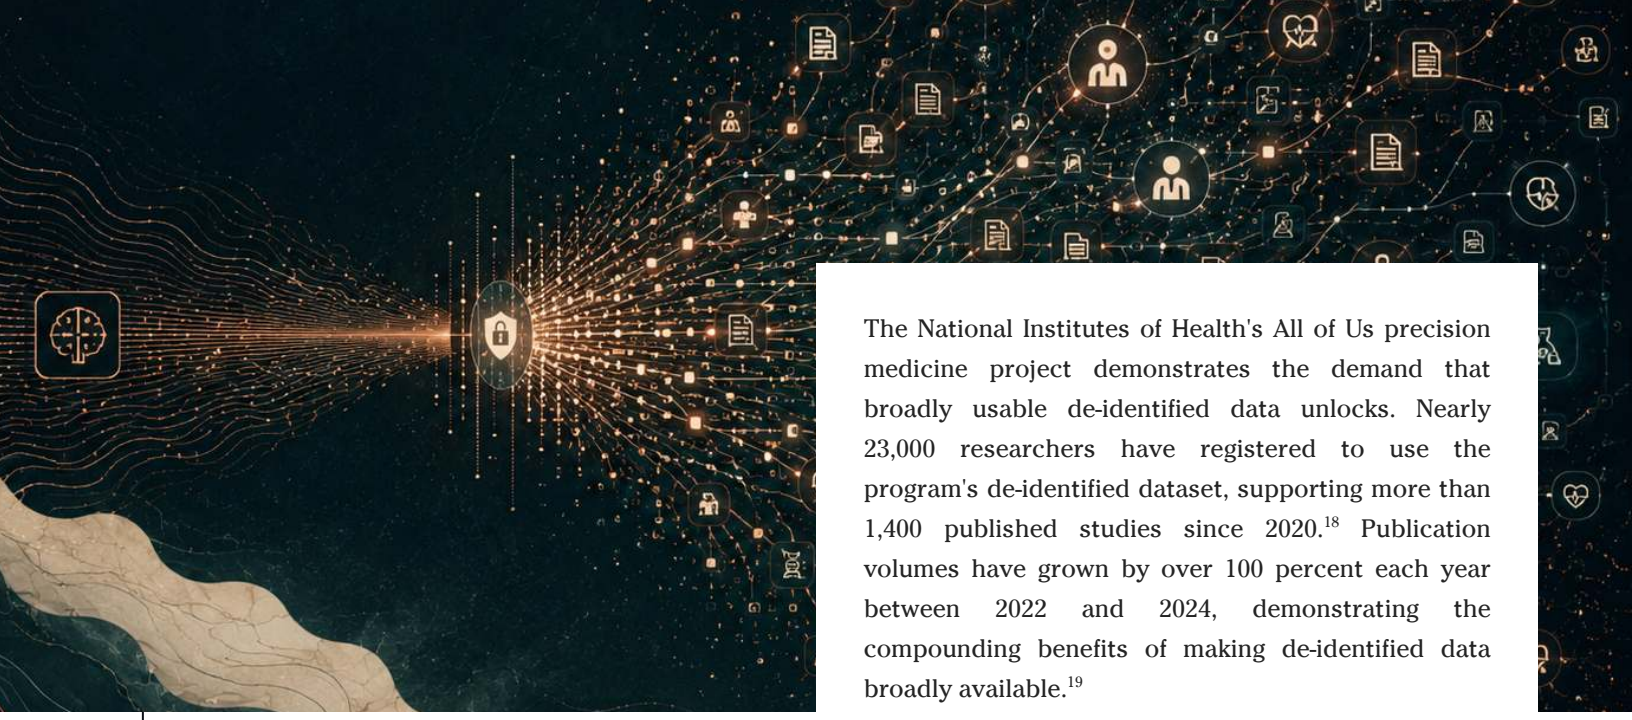

But the transformations it requires, including suppressing dates more specific than a year and geography finer than three-digit zip codes, often destroy valuable analytical insights.

Suppressing zip codes can make it difficult or impossible to understand specific underserved patient cohorts. And removing or shifting specific dates can destroy longitudinal or seasonal patterns that are essential for forecasting.

Expert Determination

The Expert Determination method is process-based rather than rules-based. Rather than prescribing specific data transformations, it requires that a person with appropriate statistical and scientific expertise examine a dataset to determine the risk that an anticipated recipient could re-identify individuals in that dataset. If the statistical expert determines that re-identification risk is "very small" for the anticipated data recipients, the dataset satisfies HIPAA's de-identification requirements. This provides flexibility to retain valuable information that the Safe Harbor method suppresses. But expert determination projects require weeks or months of expensive statistical consulting work, which increases project timelines and costs.

The National Institutes of Health's All of Us precision medicine project demonstrates the demand that broadly usable de-identified data unlocks. Nearly 23,000 researchers have registered to use the program's de-identified dataset, supporting more than 1,400 published studies since 2020.[18] Publication volumes have grown by over 100 percent each year between 2022 and 2024, demonstrating the compounding benefits of making de-identified data broadly available.[19]

All of Us also demonstrates what that access costs under the traditional model. It is a ten-year project tied to a $1.5 billion federal grant, it relies on a static data model that cannot be easily adapted to every study design, and access runs through registration tiers, data use agreements, and a controlled analysis environment.[20]

For commercial data assets or foundation model training data, this kind of consulting engagement makes sense: a standardized data asset can be commercialized and sold to many buyers or used to train a model.

But where downstream use is diverse and unpredictable, inflexible data models work poorly. For these use cases, de-identification has to be both faster and more flexible. Until recently, that was impossible.

Where downstream use is diverse and unpredictable, inflexible data models work poorly. For these use cases, de-identification has to be both faster and more flexible.

De-Identification as Governance Infrastructure

Recent advances in technology have made it possible to de-identify data at runtime, enabling flexible, distributed access to de-identified data derived from PHI.

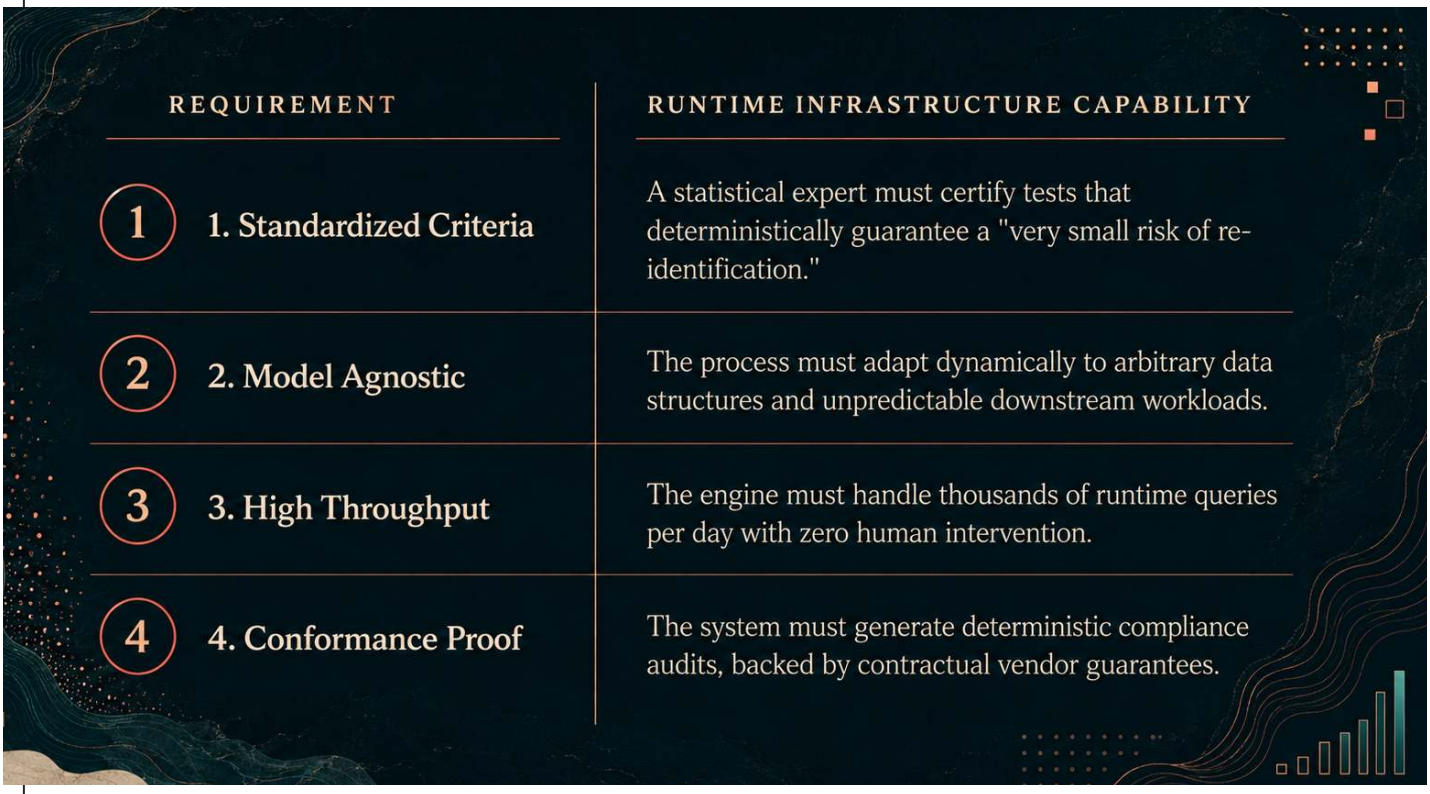

To govern data in real time at the volume, velocity, and diversity of AI-driven data requests, de-identification needs four things:

1. **A statistical expert must define and certify a set of statistical tests and criteria** that can be consistently and deterministically applied to produce datasets that carry a "very small risk of re-identification";

2. **The de-identification process must apply to arbitrary data models** to serve use cases across different data assets and analytical workloads;

3. **The de-identification process must be able to handle thousands of queries** per day and produce precise, accurate results;

4. **The de-identification process must produce deterministic proof of conformance** with the tests and criteria that establish "very small risk of re-identification;" because the HIPAA Covered Entity is liable for inadequate de-identification, where de-identification is provided by a vendor, this requirement should be supported by contractual guarantees of conformance.

Traditional data masking-based de-identification techniques fail the first criterion. Because they rely on an expert evaluation of the re-identification risk created by individual data features and combinations of features, the de-identification "process" is bespoke for every dataset. As a result, these methods cannot be evaluated by a single set of statistical tests across datasets.

Standardized de-identified data assets such as commercial claims databases or public research repositories fail the second criterion. They use a standard data model and static data elements, making them unsuitable for many downstream use cases. And rule-based systems, the most common alternative to de-identification, fail the third criterion because they require purpose review for at least a subset of queries.

No single technique satisfies all four criteria. Two, orchestrated together, do.

Synthetic Data Generation

The first technique is synthetic data generation. A synthetic dataset is a fictional lookalike population — sometimes called a “digital twin.”

Synthetic data contains no real patient records, but its fictional patients share the statistical characteristics of a real cohort or population. Because these statistically generated patients don't exist, patient re-identification from synthetic data relies on probabilistic inference attacks, which can be quantified consistently across datasets and data models. This allows the consistent application of expert-defined re-identification tests to verify de-identification.

As a result, synthetic data satisfies the first, second, and fourth criteria for scalable de-identification. But the third criterion prevents reliance on synthetic data alone: because synthetic data does not perfectly match real patient data, it can contain small but meaningful statistical deviations from the source patient data – which means analysis performed on purely synthetic data is not widely accepted for high-impact decisions, regulatory filings, or research publication.

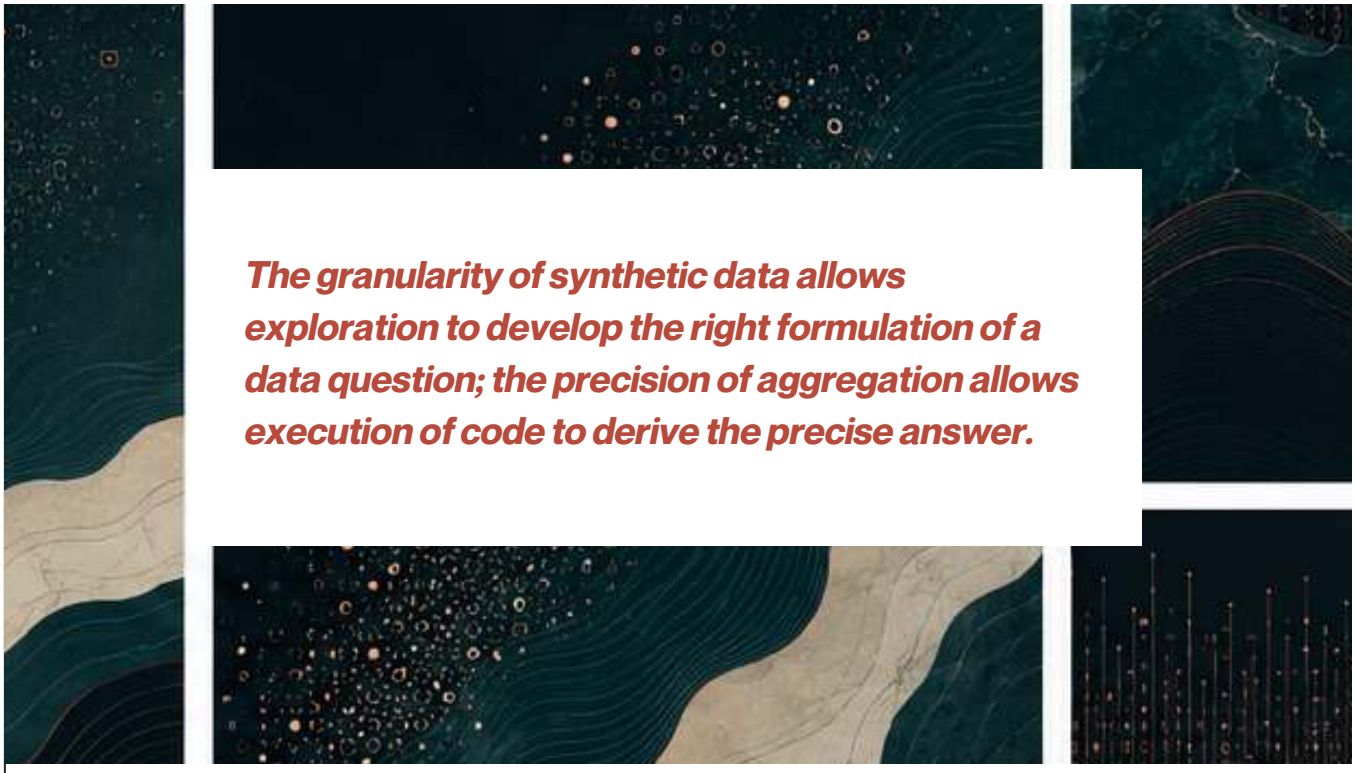

The granularity of synthetic data allows exploration to develop the right formulation of a data question; the precision of aggregation allows execution of code to derive the precise answer.

De-Identification by Data Aggregation

De-identification by data aggregation addresses the weakness of synthetic data generation. Where synthetic data generation provides granular, row-level data for analysis but potentially imprecise results, de-identification through aggregation sacrifices all granularity to provide publishable precision.

For example, a set of rows may convey too much individual patient information to be released, but releasing a row count or average of the values in those rows eliminates the privacy risk to the individual patients. Attacks to re-identify patients in aggregate statistics are possible through differencing attacks, but like synthetic data, the risk posed by the strongest attacks can be measured and mitigated programmatically. The granularity of synthetic data allows exploration to develop the right formulation of a data question; the precision of aggregation allows execution of code to derive the precise answer.

Orchestrated together, the two techniques deliver accurate, precise, de-identified data to AI agents and tools at runtime — with evidence of HIPAA compliance attached to every result. This turns Expert Determination from a consulting engagement into code: statistical tests certified once by an expert, then executed and verified at every query. It allows for distributed data access without risking data misuse or patient privacy, and it is the only approach that satisfies all four criteria for governing AI data access.

This turns Expert Determination from a consulting engagement into code: statistical tests certified once by an expert, then executed and verified at every query.

Achieving AI-Native Data Governance

HIPAA's restrictions on data use will not change to make AI adoption easier. Successful enterprise AI integration therefore requires an AI-native approach to data governance.

Health systems and hospitals operationalizing AI should assume that purpose-based restrictions will apply to AI workflows and invest in the transition from centralized to distributed data governance.

Centralized human clearinghouse governance models like honest broker reviews cannot scale to meet the volume and velocity of AI-generated data requests. A more scalable approach is necessary to achieve the projected benefits of AI in healthcare.

About Subsalt

Subsalt created runtime de-identification, deployable in a health system's existing cloud environment. Our de-identification process is governed by a privacy report prepared by Datavant: roughly 150 pages defining the statistical methods and quantitative thresholds our outputs must meet to qualify as de-identified. Those tests are applied to every dataset the system produces, and we back conformance contractually. The result: our customers can put AI to work on their most important financial and operational challenges without weakening HIPAA compliance or patient privacy.

References

1. Centers for Medicare & Medicaid Services, Office of the Actuary, "National Health Expenditures 2024 Highlights," January 2026 (2024 national health expenditures: $5.3 trillion; 18.0% of GDP). cms.gov/files/document/highlights.pdf

2. Strata Decision Technology, StrataSphere monthly healthcare financial benchmarks, January 2026 data (median year-to-date operating margin –0.6% in January 2026, a twelve-month low); see also HFMA, "Hospital margin trends in 2026 show rising cost pressure." stratadecision.com · hfma.org

3. KFF, "Allocating CBO's Estimates of Federal Medicaid Spending Reductions Across the States: Enacted Reconciliation Package," July 2025 (CBO: $911 billion net reduction in federal Medicaid spending, 2025–2034, under P.L. 119-21), citing Congressional Budget Office, "Estimated Budgetary Effects of Public Law 119-21," Publication 61570, July 2025. kff.org

4. KFF, "There is No Drop-Dead Date for an ACA Tax Credit Extension, But Coverage Losses Will Mount as the Clock Ticks," December 2025 (enhanced premium tax credits expired December 31, 2025; out-of-pocket premiums for subsidized enrollees projected to rise 114% in 2026). kff.org

5. B.E. Smith / AMN Healthcare, "Healthcare Leadership Trends for 2026 Report," February 2026 (national survey of 703 hospital and health system executives: 84% cite financial pressures as their greatest threat; 72% say lowering operating costs is their most critical long-term priority); via Michigan Health & Hospital Association, "The State of Healthcare Leadership: Risks, Reality and Readiness," March 2026. besmith.com · mha.org

6. Qventus, "Beyond the Pilot: How CIOs Are Operationalizing AI Across Health Systems in 2026," April 2026 (94% of CIOs say AI delays will put their organization at a competitive disadvantage; hospital CIOs tripled AI spending to $1.4 billion in 2025), via HIT Consultant. hitconsultant.net

7. Medscape, "Physician Burnout & Depression Report 2024" (survey of 9,226 US physicians across 29+ specialties; 62% cite too many bureaucratic tasks — charting, paperwork — as the primary contributor to burnout). medscape.com

8. The Permanente Medical Group, "Ambient Artificial Intelligence Scribes: Learnings after 1 Year and over 2.5 Million Uses," NEJM Catalyst, 2025 (7,260 clinicians; 2,576,627 patient encounters; 15,791 hours of documentation time saved, October 2023–December 2024). catalyst.nejm.org/doi/full/10.1056/CAT.25.0040

9. Olson KD, et al., "Use of Ambient AI Scribes to Reduce Administrative Burden and Professional Burnout," JAMA Network Open, 2025;8(10):e2534976 (quality-improvement study of 263 ambulatory clinicians across six US health systems; burnout fell from 51.9% to 38.8% after 30 days of use). doi.org/10.1001/jamanetworkopen.2025.34976

10. Guidehouse / HIMSS, "2026 Healthcare AI Trends," February 2026 (52% of health systems consider themselves operationally ready to implement AI; survey of 50 healthcare leaders). guidehouse.com

11. Deloitte, "Autonomous generative AI agents," TMT Predictions 2025, November 2024 (25% of enterprises using generative AI to launch agentic AI pilots or proofs of concept in 2025, growing to 50% by 2027). deloitte.com

12. HHS.gov, "Uses and Disclosures for Treatment, Payment, and Health Care Operations"; 45 CFR §164.501 (health care operations); 45 CFR §164.506 (uses and disclosures for TPO); 45 CFR §164.508(a)(3) (marketing requires individual authorization); and 45 CFR §164.508(a)(4) (sale of PHI requires individual authorization). hhs.gov · ecfr.gov

13. 45 CFR §164.502(b) and §164.514(d), HIPAA Privacy Rule, Minimum Necessary Standard (§164.514(d)(2): role-based access policies for internal uses; §164.514(d)(3): standard protocols for routine, recurring disclosures and individual review for the rest). ecfr.gov

14. Anthropic, "PwC is deploying Claude to build technology, execute deals, and reinvent enterprise functions for clients," 2026 (Advocate Health deploying Claude toward its full 167,000-person workforce). anthropic.com/news/pwc-expanded-partnership; clinical workforce composition per Advocate Health, "Who We Are" (35K+ doctors/aligned and medical staff; 42K nurses). advocatehealth.org/about

15. Keller, M.S., et al., "The Honest Enterprise Research Broker: Facilitating Ethical, Efficient, and Secure Access to Health Data for Research," Applied Clinical Informatics, 2025 (Cedars-Sinai HERB Committee: 487 extracts 2018–2023; 51→121 requests/yr; 12.5→9 hours per extract). doi.org/10.1055/a-2499-4090

16. U.S. Department of Health and Human Services, "Guidance Regarding Methods for De-identification of Protected Health Information in Accordance with the HIPAA Privacy Rule." hhs.gov/hipaa/for-professionals/special-topics/de-identification

17. 45 CFR §164.514(b), HIPAA Privacy Rule, De-identification standards (Safe Harbor: §164.514(b)(2); Expert Determination: §164.514(b)(1)). ecfr.gov

18. NIH All of Us Research Program, "The All of Us Research Program Enters the Multiomics Era with Landmark Data Release," July 2026 (nearly 23,000 registered researchers; more than 1,400 peer-reviewed publications since the Researcher Workbench launched in 2020). allofus.nih.gov

19. CDC Public Health Genomics and Precision Health Knowledge Base, "All of Us Publication Finder" (annual publication counts as retrieved July 2026: 126 in 2022; 304 in 2023; 706 in 2024). https://phgkb.cdc.gov/PHGKB/allOfUsPublicationFinder.action

20. 21st Century Cures Act, Pub. L. No. 114-255, §1001 (2016) (authorizing $1.455 billion for the Precision Medicine Initiative / All of Us, FY2017–FY2026); NIH, "All of Us Research Program Backgrounder" ("a total of $1.5 billion over ten years"); data model: All of Us Research Program, "Data Types and Organization" (Curated Data Repository standardized to the OMOP Common Data Model v5.3). congress.gov · nih.gov · support.researchallofus.org

Subsalt

Health Data, Safe to Use by Default

Give AI safe access to your data without the queue → subsalt.ai/contact

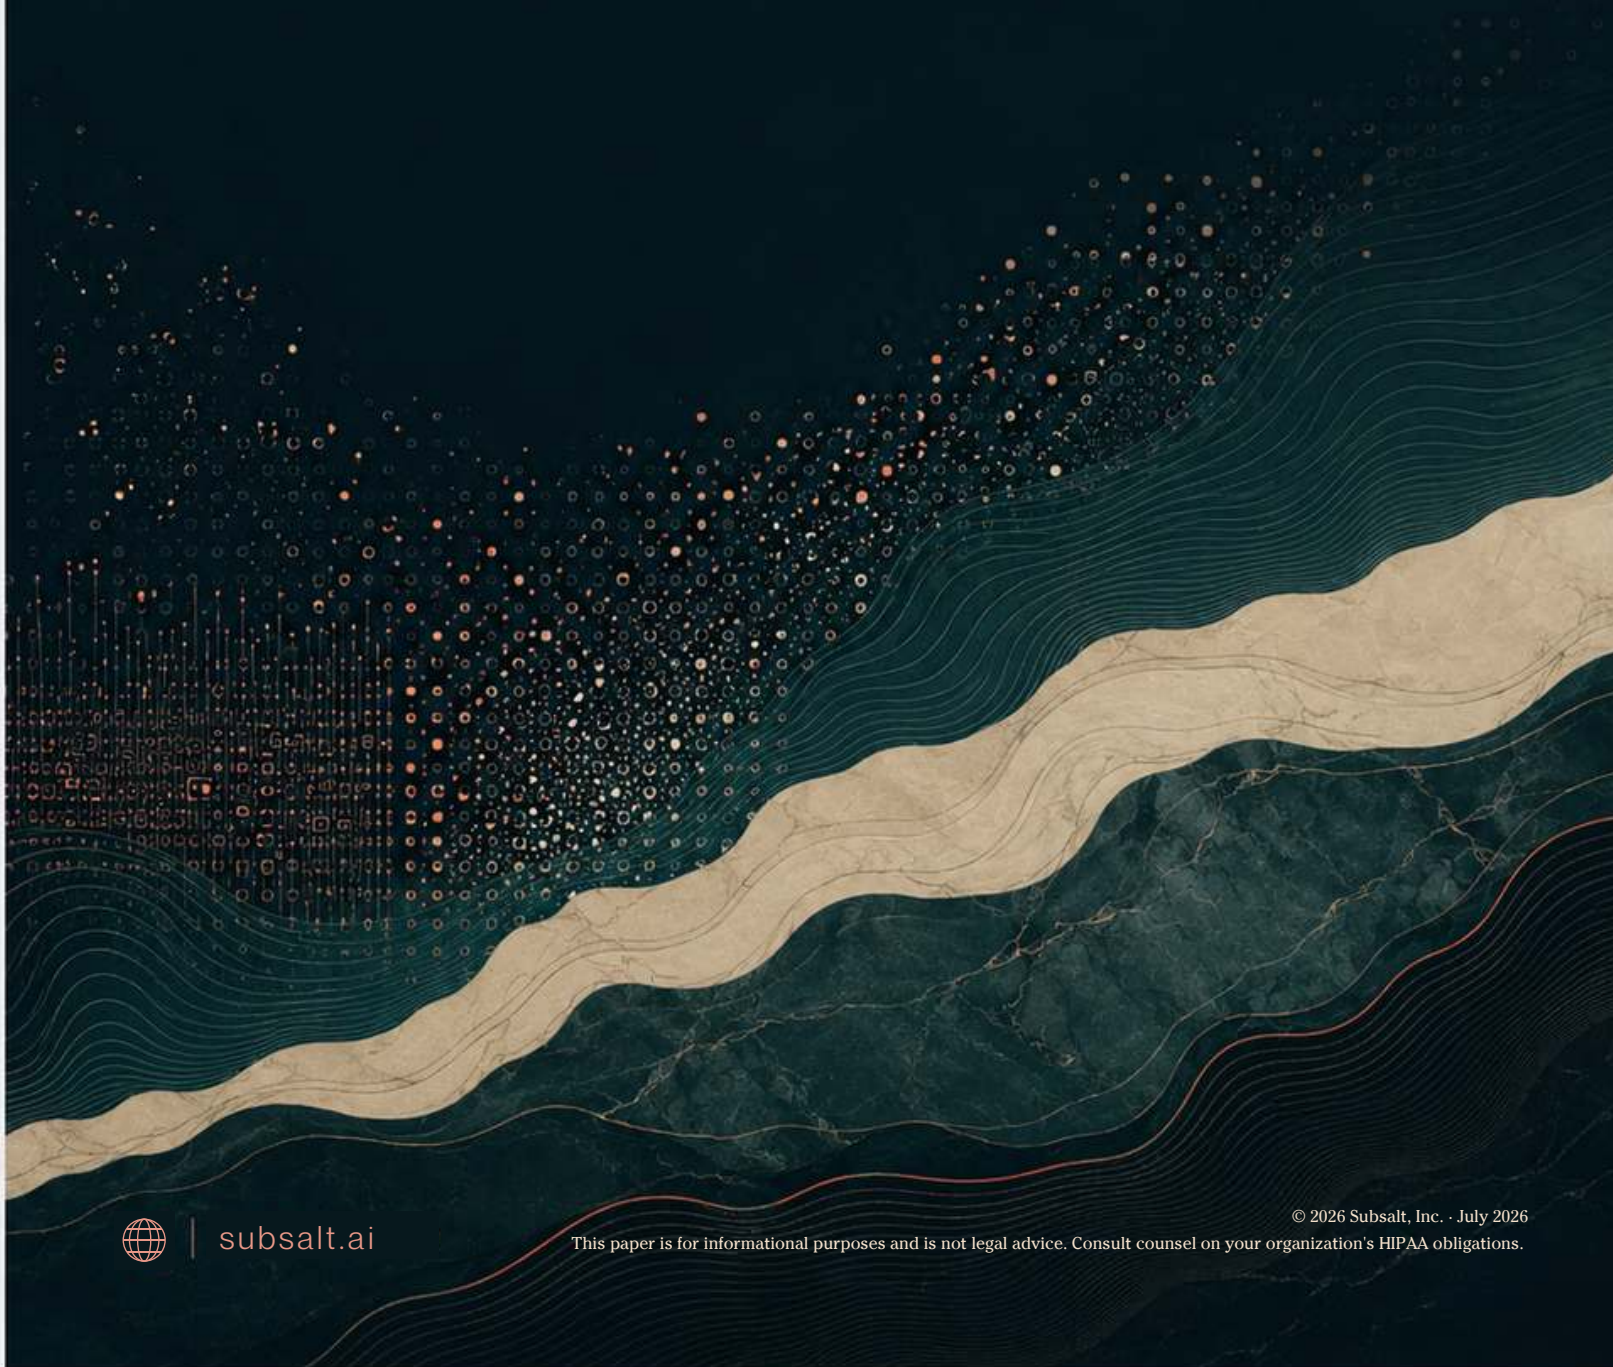

subsalt.ai