Doi: 10.32481/djph.2026.03.15

Maria in 2035:

Delaware as a Living Laboratory for AI-Enabled Public Health

Neil G. Hockstein, MD[1] and Patrick J. Callahan, JD[2]

1. Chair, Delaware Health Care Commission

2. Founder Keel3.ai, Co-Founder Acellus Health

Delaware's Opportunity: Small State, Big Bet

Delaware is small enough to coordinate and large enough to matter. It is a state that can align health systems, payers, employers, and public agencies, can share measures of outcomes and cost, can reduce unwarranted clinical variation, and can learn faster than larger jurisdictions. In chronic disease, the biggest wins rarely come from a single breakthrough drug. They come from reliable measurement, early detection, consistent follow-through, and fewer gaps between what evidence recommends and what patients receive.

Our public health need is clear. Diabetes and related chronic conditions drive preventable complications and a large share of healthcare spending. State health reports underscore that diabetes remains widespread and consequential for morbidity, mortality, and quality of life across communities.[1]

Delaware's structural assets including a nationally recognized health information exchange and only six health care systems providing both acute and outpatient care lay the foundation for success. Coupled with a culture of collaboration and demography that mirrors the country, Delaware can serve as a model for the nation.

This commentary argues for a "small state, big bet" strategy: use Delaware's compact geography, data infrastructure, and cross sector alignment to build a true learning health landscape. A learning landscape is one that continuously improves by turning care data into knowledge and then into better practice. It is not a dashboard. It is a disciplined feedback loop.[2]

To make the stakes concrete, start with one patient. Meet Maria.

Maria in 2025: the System's Predictable Failure Mode

Maria is 56 and lives in Sussex County. She has type 2 diabetes and hypertension. Her story is not dramatic. That is why it matters. Many of the most expensive failures in healthcare are quiet and ordinary. They happen when a manageable problem is allowed to evolve into a crisis.

In today's care model, Maria's experience is episodic. It depends on appointments, phone calls, and portal messages that assume time, transportation, and digital comfort. When she develops a minor illness, her blood glucose becomes more variable. She feels tired, drinks less water, and skips a walk. These early signals are common, but they are easy to miss when data are scattered across settings.

Doi: 10.32481/djph.2026.03.15

Care becomes a sequence of disconnected encounters: urgent care, emergency department, a rushed follow up, and sometimes a delayed medication change. Each step may be reasonable in isolation, yet the overall experience becomes costly, variable, and administratively burdensome. This burden is not evenly distributed. People with fewer resources have less slack to absorb delays, missed work, and complex instructions.

The administrative layer compounds the clinical risk. Prior authorization and documentation requirements can delay care, consume clinician time, and shift work onto patients. National surveys show physicians report that prior authorization often delays necessary care and can lead to adverse events.[3]

Most clinicians do not oppose prudent utilization management. They oppose opaque, inconsistent processes that add friction without improving outcomes. When administrative delay is routine, it becomes a hidden clinical factor. In a chronic disease economy, delay is not neutral. It is a predictable driver of deterioration.

So what would a better system look like, if we designed it to notice early and respond early?

Maria in 2035: Continuous, Anticipatory Care with a Human Clinician at the Center

It is 2035. Maria wakes up congested with a low-grade fever. Ten years earlier she might have waited until symptoms forced an urgent visit. Now the system notices first.

Overnight, her continuous glucose monitor shows rising variability. Her blood pressure cuff shows a modest but meaningful increase. These are early risk signals that deserve attention, especially for a patient whose baseline data are well understood. Maria's health partner checks in with a plain language prompt and two questions: what are you feeling, and do you want to address this now while it is still small?

Maria says yes.

Within minutes she sees a short summary: what is changing, what typically happens to patients like her, and what actions reduce risk. She is presented options that include expected benefit, time cost, and estimated out of pocket cost. The system displays what is known, what is assumed, and what is uncertain. This transparency is essential for trust.

A brief virtual visit is scheduled the same day. The clinician sees a coherent timeline rather than a blank screen. The recommended plan is transparent, with sources and relevant data highlighted. The clinician adjusts medications and orders a test. Coverage is confirmed quickly because the administrative process is interoperable. Maria's condition stabilizes. The best outcome is the quiet one: no emergency visit, no avoidable complication, and less time lost for both patient and clinician.

In this future, the human clinician remains central. AI supports monitoring, triage, and administrative routing. It does not replace clinical judgment or the therapeutic relationship. If a patient is anxious, grieving, confused, or in pain, the most important intervention is still human attention.

Better systems create the promises that reduce preventable deterioration and liberate clinicians to practice at the top of their training.

Doi: 10.32481/djph.2026.03.15

What Makes this Plausible: Learning Health Systems and Better Use of Data

The concept of a learning health system is not new. The Institute of Medicine described a national pathway toward continuously learning care: data from practice become knowledge, and knowledge becomes better practice.[4]

What has changed is the feasibility of operationalizing the learning loop at scale. Interoperable data exchange, increased availability of home monitoring, and modern AI methods make earlier detection and faster coordination more realistic. However, feasibility is not the same as readiness. A system can be technologically capable and still untrustworthy if governance is weak.

Agentic AI matters because it goes beyond generating text. It can plan, sequence tasks, monitor trends, and coordinate actions across systems under defined constraints. When used carefully, it can reduce the administrative drag that currently consumes clinician capacity and slows care. When used carelessly, it can amplify bias, obscure accountability, and erode trust.

Delaware's advantage is that it can build statewide governance, with clear standards and measurable pilots. Once evidence proves efficacy, Delaware can provide a path to scale.

Delaware's Assets: a Head Start on Interoperability and Alignment

Delaware has an unusually strong foundation for statewide coordination, especially compared with states where health data remain siloed among competing systems.

First, Delaware has a mature statewide health information exchange. DHIN has operated for years and has been recognized as an early statewide model for enabling data exchange.[5]

Second, Delaware has clear public health need. State reports document the persistent burden of diabetes and other chronic disease, along with the downstream cost of complications.[6]

Third, Delaware can align stakeholders more quickly than larger states. That is the core of the "small state, big bet." A compact state can establish shared measures, launch pilots across a meaningful share of the population, and iterate quickly.

Finally, national interoperability initiatives provide additional leverage. TEFCA aims to expand trusted exchange and could support broader connectivity as Delaware scales its learning system approach and FHIR has matured from simple data exchange to allowing for actionable real-time interoperability.[7]

The Required Scaffolding: Governance, Privacy, Auditability, and Accountability

Technology is the easy part. Legitimacy is the hard part.

A critical and often overlooked dimension of this transformation is the regulatory boundary around health data. The major AI companies entering the healthcare space are

Doi: 10.32481/djph.2026.03.15

not themselves HIPAA-covered entities. HIPAA governs health plans, clearinghouses, and healthcare providers and their business associates, but once patient information leaves a covered entity and enters consumer-facing platforms, health apps, wearable ecosystems, or standalone AI tools, it falls outside HIPAA's protections. In those contexts, data can be governed instead by general consumer privacy policies and terms of service, which often permit broad secondary uses, including algorithm training and product development. This structural asymmetry: information generated within a clinical encounter is tightly regulated, but once exported to non-covered digital environments, it may be reused to refine models or train large-scale AI systems without explicit, encounter-specific patient consent. As AI companies expand deeper into health-related services, clarifying these regulatory gaps - and aligning them with patient expectations of confidentiality - will become increasingly urgent.

To create a learning health landscape supported by agentic AI, Delaware needs a public trust framework that answers four questions clearly.

First, who can access which data, for what purpose, and under what oversight? Second, how is privacy protected, including data minimization and security controls? Third, how are AI supported recommendations audited, explained, and monitored for bias, safety, and unintended consequences? Fourth, who is accountable when systems fail, including workflows that incorporate AI outputs?

A practical way to structure this is to use an established risk framework for AI and adapt it to healthcare context. The NIST AI Risk Management Framework provides a widely recognized approach to mapping, measuring, and managing AI risks across the lifecycle.[8]

In practice, this means every pilot should have a documented purpose, defined boundaries, performance metrics, monitoring plans, and clear human accountability. It should also include pre deployment testing, post deployment surveillance, and explicit criteria for pausing or stopping a deployment if harms emerge.

This framing helps avoid a common error. People argue about whether AI is good or bad. The real issue is whether an AI enabled workflow is governed, measurable, and accountable.

Reducing Administrative Waste: Prior Authorization as a Test Case

If Delaware wants immediate wins that matter to clinicians and patients, start with administrative burden. Prior authorization is a natural target because it sits at the intersection of cost, access, and clinician capacity.

Nationally, prior authorization is widely reported as a major burden. It is also an area where federal policy is pushing for greater interoperability and more standardized, electronic processes.[9]

Delaware can treat this as a measurable pilot. Choose a defined set of chronic disease services where evidence based pathways are clear. Measure time to decision, appeals, reversal rates, clinician time spent, and patient outcomes. Publish results transparently. When the process improves care and reduces burden, expand. When it fails, fix it before scaling.

Doi: 10.32481/djph.2026.03.15

This is also a strong use case for carefully constrained automation. For low risk, high volume services that are strongly evidence based, the default should be rapid approval with audit rather than slow approval with delay. A learning system earns trust by showing its work and by measuring what happens after policy changes.

Equity: the Non-Negotiable Goal

Chronic disease burden and administrative friction fall hardest on under resourced communities. A system that makes care continuous and anticipatory can narrow equity gaps, but only if it is designed to do so.

Delaware should require that pilots report outcomes stratified by geography and key demographic variables, and that mitigation plans are built into deployment rather than added after problems emerge. Equity should be a design requirement.

This is also why the Maria vignette matters. It is a story about preventing predictable harm, reducing wasted time, and making care easier to navigate.

A Call to Action for Delaware's Leaders

Delaware can lead the nation by proving that a state can reduce chronic disease complications through a learning health landscape that is both technologically capable and governance first.

That requires coordinated action. Health systems, already committed to population health, can provide data for transparent dashboards based on clinician designed clinical workflows. Payers should align incentives to reward reduced variation and remove avoidable administrative friction. Employers should support benefit designs that favor evidence based, high value care. Public health leaders should establish data governance rules that protect trust while enabling learning. Technology partners should be invited into structured, measurable pilots with clear guardrails and accountability.

A practical starting roadmap is to select two or three chronic disease pathways, define shared metrics, stand up governance and audit processes, and launch pilots that include both clinical outcomes and administrative burden measures. If pilots show reduced complications, improved patient experience, and lower total cost of care, Delaware can scale with confidence. If they do not, Delaware can iterate without having overcommitted.

Maria's 2035 story is what happens when a system chooses to notice earlier, coordinate better, and learn continuously. Delaware can choose that future.

Mr. Callahan may be contacted at Patrick.callahan@keel3.ai.

References

1. Holliday, C. S., & Gabbay, R. A. (2025, February 27). Breaking barriers: CDC and American Diabetes Association unite to combat diabetes. *Preventing Chronic Disease*, *22*, 240273. https://doi.org/10.5888/pcd22.240273 PubMed

2. McDonald, P. L., Foley, T. J., Verheij, R., Braithwaite, J., Rubin, J., Harwood, K., . . . van der Wees, P. J. (2024, January 25). Data to knowledge to improvement:

Creating the learning health system. [PubMed https://doi.org/https://doi.org/10.1136/bmj-2023-076175]. *BMJ (Clinical Research Ed.)*, *384*, e076175.

3. American Medical Association. (2025). *2024 AMA prior authorization physician survey*. https://www.ama-assn.org/system/files/prior-authorization-survey.pdf
4. Committee on the Learning Health Care System in America. Institute of Medicine. *Best Care at Lower Cost: The Path to Continuously Learning Health Care in America*. Smith, M., Saunders, R., Stuckhardt, L., McGinnis, J.M., eds. National Academies Press; 2013. Accessed February 22, 2026. https://pubmed.ncbi.nlm.nih.gov/24901184/
5. Delaware Health Information Network. (n.d.). About DHIN. Delaware Health Information Network. https://dhin.org/about/
6. Delaware Department of Health and Social Services. (2025). The impact of diabetes in Delaware, 2025. Division of Public Health; Delaware Department of Human Resources, Statewide Benefits Office. https://dhss.delaware.gov/wp-content/uploads/sites/10/2025/07/2025-Impact-of-Diabetes-in-DE_FINAL.pdf
7. Assistant Secretary for Technology Policy. (2026). Trusted exchange framework and common agreement. https://healthit.gov/policy/tefca/
8. National Institute of Standards and Technology. (2023). Artificial intelligence risk management framework (AI RMF 1.0). U.S. Department of Commerce, NIST AI 100-1. https://nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf
9. Kyle, M. A., & Song, Z. (2023, July 27). The consequences and future of prior-authorization reform. [PubMed]. *The New England Journal of Medicine*, *389*(4), 291–293. https://doi.org/10.1056/NEJMp2304447 PubMed