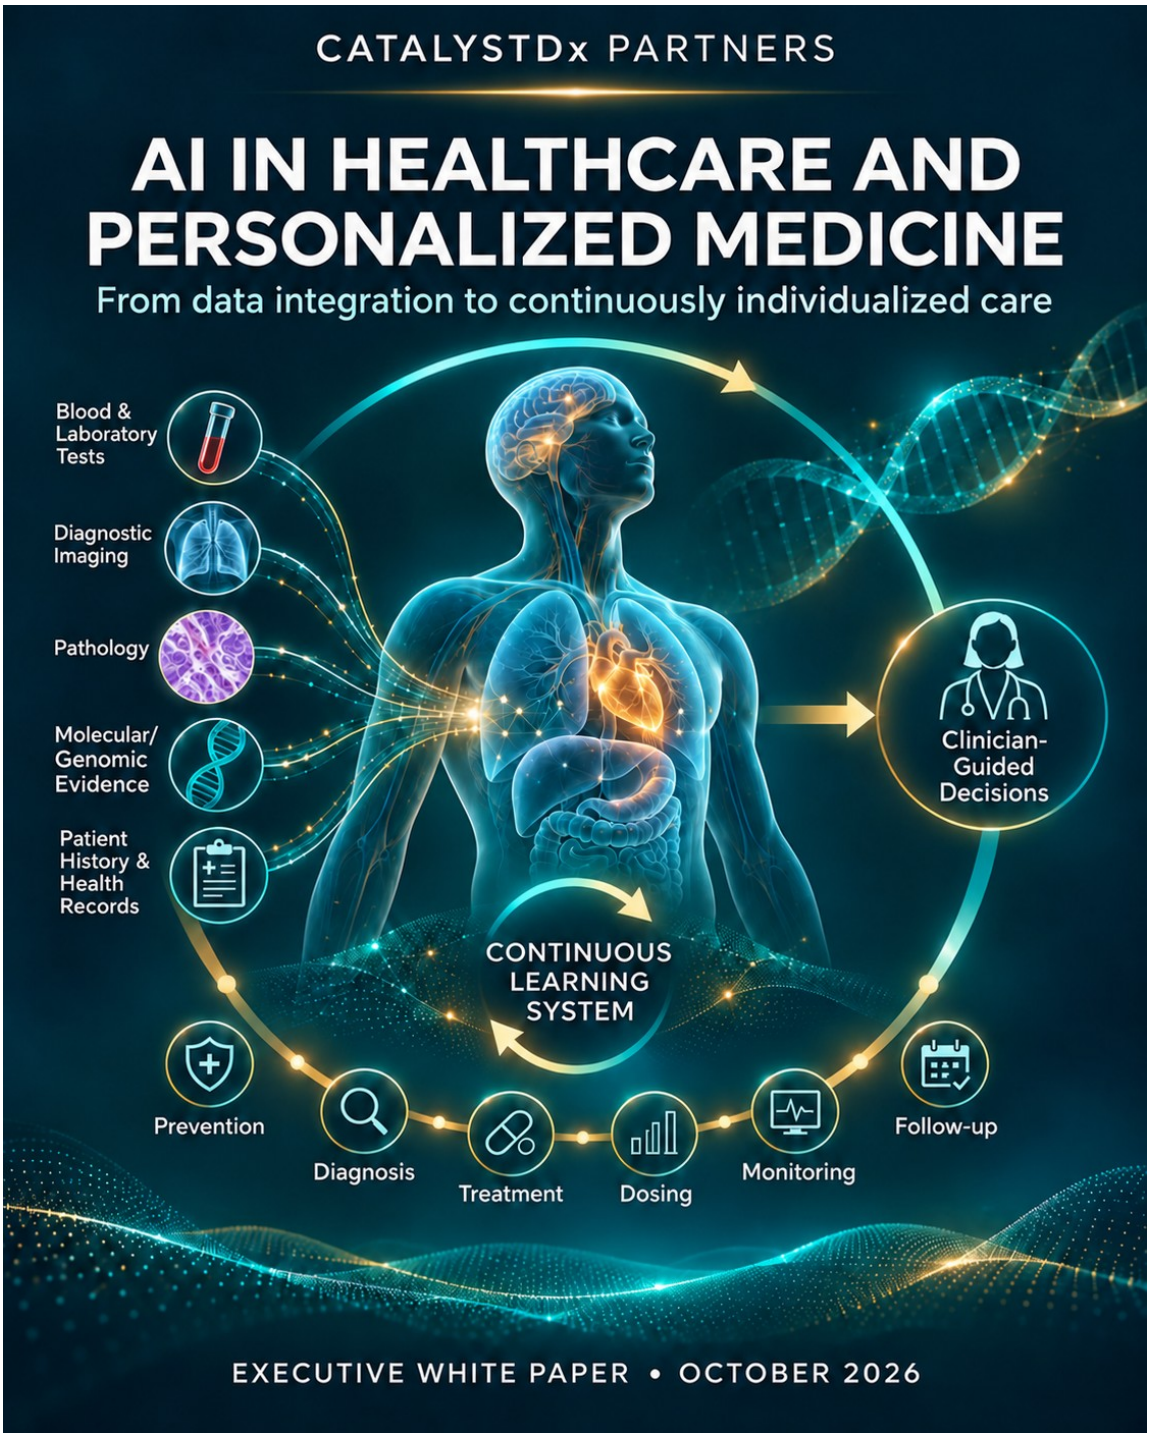


AI creates strategic value in personalized medicine when it turns multimodal patient data into governed, continuously updated clinical decisions that improve care.

Peter Serpentino | Founder, CatalystDx Partners

Educational and strategic analysis. Not medical advice.

© 2026 CatalystDx Partners. All rights reserved.

Contents

Executive Summary

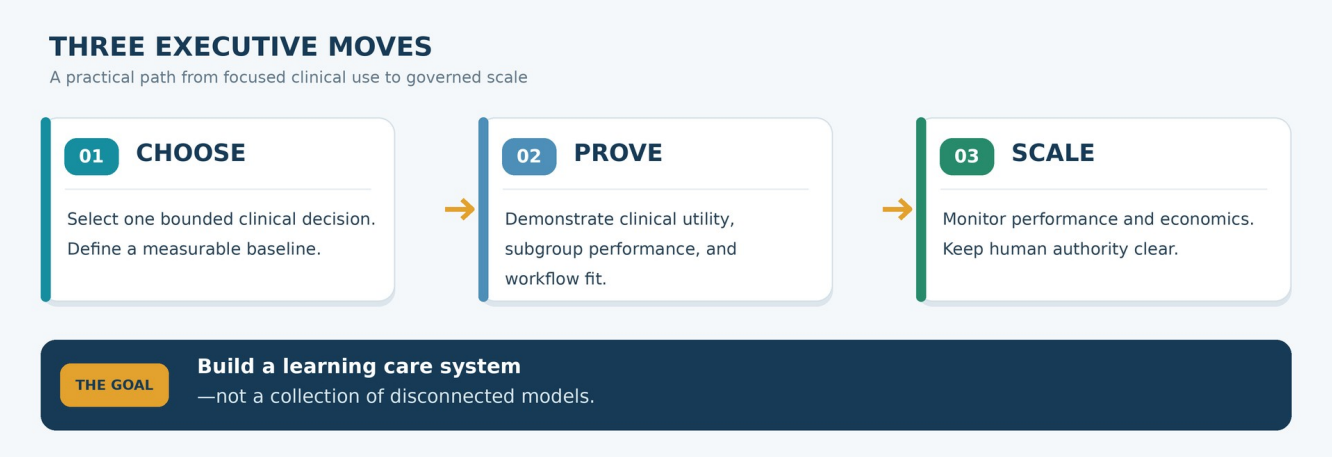
1. The strategic shift: from precision tests to personalized care systems
2. Why AI is the enabling layer
3. The multimodal data foundation
4. High-value clinical and R&D applications
5. Generative AI and agentic workflows
6. Evidence: moving from model performance to clinical utility
7. Governance, regulation and trust
8. Operating model and economics
9. A three-horizon implementation roadmap
10. Outlook: toward continuously learning medicine

Conclusion

References

Executive thesis and action plan

The strategic opportunity is to build a governed learning system that converts multimodal patient data into decisions that measurably improve care.



Executive Summary

CORE THESIS AI creates value in personalized medicine when it integrates heterogeneous patient signals into an actionable clinical decision, then updates that decision as new evidence arrives.

Healthcare is moving from episodic, population-average decision making toward a model in which prevention, diagnosis and treatment are increasingly tailored to an individual's biology, disease state, prior response, behavior and environment. Genomics made the ambition of precision medicine tangible; artificial intelligence is making the operating model scalable. The combination is creating a new clinical architecture: multimodal data are converted into dynamic patient representations, those representations inform risk and treatment choices, and longitudinal outcomes feed back into the system.

The transition is no longer theoretical. By September 2026, the U.S. Food and Drug Administration (FDA) reported more than 1,600 AI-enabled medical devices authorized for marketing in the United States, spanning imaging, cardiovascular care, pathology, obstetrics, neurology and other specialties. At the same time, the NIH All of Us Research Program has demonstrated the scale and diversity now possible for precision-medicine research, with clinical-grade whole-genome sequences linked to longitudinal health data and hundreds of millions of previously unreported genetic variants. These developments are expanding both the supply of clinically relevant data and the number of regulated AI systems capable of acting on it. [1–3]

The next phase will be defined less by stand-alone algorithms and more by integrated clinical systems. Those systems will combine genomics and other omics, EHR data, imaging, digital pathology, laboratory results, medication history, wearables and patient-generated data. Multimodal models can then support earlier detection, disease subtyping, pharmacogenomic decision support, therapy selection, outcome prediction, adaptive monitoring and trial design. In research and product

development, AI is also accelerating target discovery, biomarker development, patient stratification and evidence generation. [4–7]

However, greater personalization does not automatically produce better care. Models that use many variables may reproduce biases in their training data, learn patterns specific to small cohorts, or lose accuracy as patient populations, clinical workflows, and standards of care change. They may identify associations without showing that one factor causes another, and clinicians may over-rely on recommendations even when they are wrong. Generative models raise further concerns: their outputs may vary, include fabricated information, expose sensitive data, or lack traceable sources. Evaluation must therefore extend beyond retrospective accuracy. Before an AI system moves from providing information to making recommendations or taking actions with greater clinical influence, evidence should show that it works as designed (technical validity), performs reliably for the intended patients (clinical validity), improves decisions or outcomes (clinical utility), works across relevant populations and care settings (generalizability), and can be monitored and managed throughout its use (lifecycle assurance). [8–12]

For healthcare providers, diagnostics companies and life-science organizations, the strategic implication is clear: competitive advantage will increasingly depend on the ability to build a governed learning system around patient data. The winning architecture will not be a single proprietary model. It will be an operating model that controls data rights, model validation, workflow integration, human authority, regulatory change, monitoring and continuous learning.

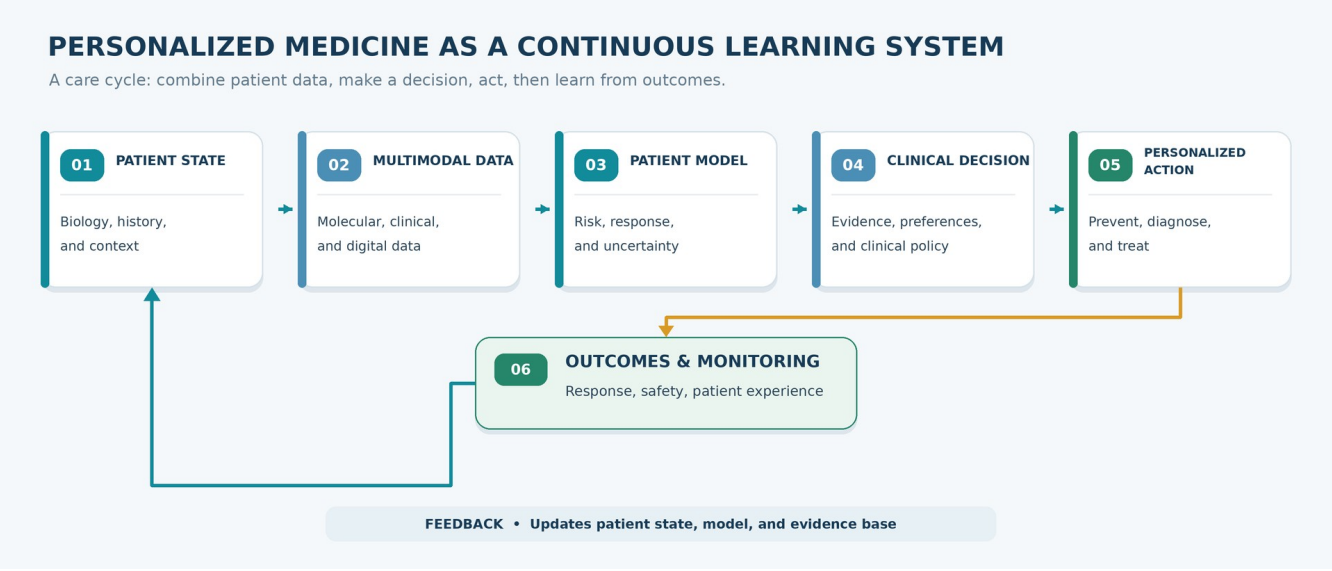


Exhibit 1 | AI turns personalized medicine into a continuous learning loop

Executive implications

- Start with a bounded clinical decision, not with "AI" as a technology program. The use case should specify who acts, on what data, at what point in the workflow, and what measurable outcome should improve.
- Build the data architecture for multimodal, longitudinal patient evidence. Connect relevant laboratory results, including blood tests; pathology; medical imaging; genomic and other molecular data; patient history and records; and patient-reported or digital measures. Preserve

each source's timing, quality, and provenance so evidence can be interpreted together as the patient's condition changes.

- Treat evidence generation as a product capability. Design validation, subgroup analysis, workflow monitoring, change control, and real-world evidence into the system from the start.
- Separate recommendation from authority. AI can summarize, rank, predict or recommend before it is allowed to execute a clinical or operational action. Human authority should be explicit and risk-adjusted.
- Protect strategic data and workflow position. As external AI vendors become embedded in clinical and R&D workflows, organizations need clear rules for data use, model portability, intellectual property, auditability and exit readiness.
- Measure value across four dimensions: clinical outcomes, patient experience, workforce productivity, and economics. A strong retrospective performance score alone does not demonstrate clinical utility or enterprise value if workflow adoption or patient outcomes remain weak.

1. The strategic shift: from precision tests to personalized care systems

Personalized medicine is often described as "the right treatment for the right patient at the right time." CatalystDx defines AI-augmented personalized medicine more broadly: the ongoing selection and adjustment of the right preventive action, diagnostic pathway, therapy, dose, monitoring plan, and follow-up cadence for an individual patient whose biology, history, preferences, and clinical condition change over time. This approach brings together relevant evidence from patient history and records, blood and other laboratory tests, imaging, pathology, molecular data, and patient-reported or digital measures. AI can help integrate these sources, surface patterns and uncertainty, and support timely, clinician-led decisions. The aim is an adaptive care system rather than a one-time match between a test and a treatment. In clinically appropriate cases, when the necessary evidence is available, CatalystDx's audacious goal is to compress the path from presentation to an evidence-informed diagnosis and appropriate treatment to one day.

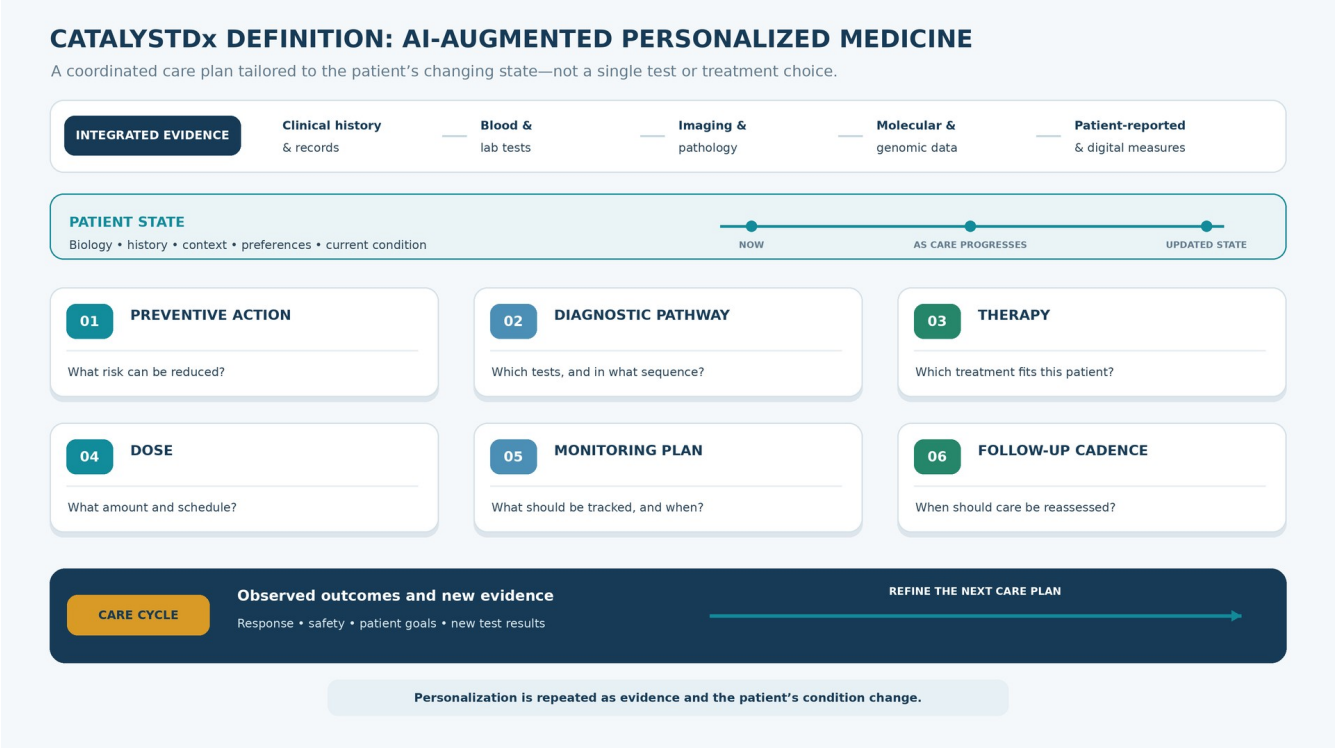


Exhibit 2 | The CatalystDx Definition of AI-Augmented Personalized Medicine

Traditional evidence-based medicine is necessarily population oriented. Randomized trials estimate average treatment effects, clinical guidelines convert those averages into standardized recommendations, and clinicians then adjust for the individual. Personalized medicine inverts part of that process. It begins with the individual patient and asks which combination of molecular, phenotypic and contextual variables makes this patient similar to—or importantly different from—the populations on which existing evidence was built.

Genomics provided the first scalable demonstration of this principle. Biomarker-defined oncology, HLA testing, CYP pharmacogenetics and molecularly targeted therapies showed that treatment benefit and risk can differ materially across biologically defined subgroups. The FDA continues to maintain extensive tables of pharmacogenomic biomarkers in drug labeling and pharmacogenetic associations, reflecting the growing role of genotype in treatment selection, dose optimization and adverse-event risk management. [13,14]

The limitation of early precision medicine was that many clinical phenotypes are not governed by one variant or one biomarker. They emerge from interacting genomic, epigenetic, transcriptomic, proteomic, metabolic, imaging, behavioral and environmental signals. The analytical problem therefore becomes high dimensional, sparse, nonlinear and longitudinal—the type of problem for which modern machine-learning methods are well suited.

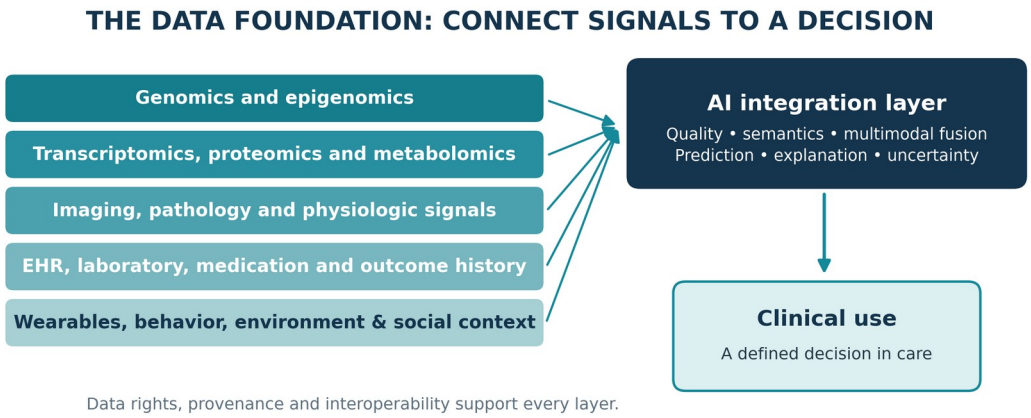


Exhibit 3 | The multimodal data stack for personalized medicine

From static stratification to dynamic patient models

The most important conceptual shift is from static stratification to a dynamic patient model. A static biomarker may classify a tumor as HER2-positive or identify a CYP2C19 metabolizer phenotype. A dynamic model can continuously update risk as treatment changes, laboratory values move, new imaging arrives, a wearable detects altered physiology or a liquid biopsy detects residual disease. The decision becomes a function of time, not only of baseline status.

This matters because disease trajectories are path dependent. Two patients with the same diagnosis may have different molecular drivers, treatment histories, comorbidities, adherence patterns and physiological reserve. A personalized system therefore needs both biological depth and temporal context.

2. Why AI is the enabling layer

AI contributes to personalized medicine through five capabilities. First, it can compress complex data into useful representations—for example, extracting quantitative features from pathology slides or learning embeddings from genomic sequences. Second, it can integrate data types that were historically analyzed in separate silos. Third, it can detect nonlinear interactions that are difficult to specify manually. Fourth, it can generate probabilistic forecasts or rankings that support decisions. Fifth, generative and agentic systems can increasingly translate those insights into workflow steps, documentation and coordinated actions.

AI capability	Personalized-medicine role	Representative output
Pattern recognition	Detect subtle disease or phenotype signals in imaging, pathology, ECGs, waveforms or molecular data	Probability, segmentation, classification, quantitative feature
Multimodal integration	Combine genomics, labs, EHR, imaging, wearables and context	Patient representation or risk state
Prediction	Estimate disease onset, response, toxicity, recurrence or progression	Risk score, expected response, time-to-event
Recommendation	Rank therapies, tests or next-best actions under defined constraints	Treatment/test recommendation with rationale

AI capability	Personalized-medicine role	Representative output
Generative synthesis	Summarize longitudinal records and scientific evidence for clinicians	Problem-oriented summary, draft note, evidence brief
Agentic orchestration	Coordinate bounded multi-step workflows with tools and policies	Scheduled test, drafted order, routed review, monitored exception

Healthcare AI capabilities are at different stages of maturity. That maturity should be judged by more than technical performance: it also depends on the strength of clinical evidence, real-world adoption, demonstrated benefit, and the level of decision-making authority the system receives. Imaging AI is among the more established categories. The FDA lists many authorized AI-enabled imaging devices, often designed for specific tasks such as detecting, measuring, or prioritizing findings. FDA authorization applies to a device's defined intended use; by itself, it does not show how widely the device is used or whether it improves outcomes in every care setting.

Ambient documentation AI listens to a clinical encounter and drafts a note for the clinician to review. Multicenter health-system evaluations have reported reductions in documentation burden and improvements in clinician experience. These tools support documentation, however; their results do not establish that AI can independently diagnose or manage a patient.

Digital twins are patient-linked computational models updated with patient data to represent a person's health or disease state and explore how it may change or respond to an intervention. A 2026 systematic review found that published healthcare applications remain dominated by simulations and that integration into routine clinical practice is rare. Autonomous care orchestration refers to AI agents that plan or coordinate multiple steps in a care workflow, sometimes using EHRs and other clinical tools. Early agent deployments remain concentrated in administrative workflows, while many clinical applications are still being evaluated in controlled settings. Both areas need stronger evidence of reliability, safety, interoperability, and clinical value before routine clinical authority is appropriate. [17,18]



Exhibit 4 | Illustrative 2026 maturity versus personalization value potential

3. The multimodal data foundation

The performance ceiling of personalized AI is often determined more by data architecture than by model architecture. Organizations can buy advanced models, but they cannot easily buy longitudinal data rights, clean phenotype definitions, normalized lab histories, high-quality imaging labels, specimen provenance, treatment-response data or the clinical context needed to interpret them.

3.1 Genomics and other omics

Population-scale genomic resources are rapidly improving the statistical foundation for precision medicine. The All of Us genomic release described 245,388 clinical-grade whole-genome sequences and more than one billion variants, including over 275 million variants previously unreported. The program's deliberate inclusion of participants historically underrepresented in biomedical research is strategically important because model validity and genetic-risk estimates can degrade when discovery cohorts do not represent intended-use populations. [2,3]

Yet genomic information is only one layer. Transcriptomics captures expression state; proteomics reflects functional molecules; metabolomics can reveal downstream pathway activity; epigenomics adds regulatory context; single-cell and spatial assays characterize heterogeneity. Reviews published in 2025–2026 increasingly frame AI as the integrative engine across these modalities, especially for pharmacogenomics, disease subtyping and precision oncology. [4–7]

3.2 Imaging and digital pathology

Medical images are already among the most mature clinical AI domains because they are information-rich, digitally available and tied to defined tasks. The FDA's authorized-device list continues to be dominated by radiology-related products, but AI is expanding into pathology, cardiovascular diagnostics, obstetrics, neurology and surgical planning. Imaging AI also contributes to personalized medicine when quantitative imaging features—radiomics—are combined with molecular or clinical variables to predict prognosis or therapeutic response. [1,19]

3.3 EHR, laboratory and longitudinal outcomes

Precision medicine becomes clinically useful only when molecular findings are connected to phenotype, treatment history and outcomes. EHR data provide that longitudinal context but introduce problems of missingness, coding inconsistency, temporal bias and site-specific workflow effects. The opportunity is substantial: a genomic signal can be interpreted differently depending on renal function, prior medication exposure, coexisting disease, age, imaging findings and current symptoms.

3.4 Wearables and patient-generated data

Wearables extend personalized medicine beyond episodic encounters. Continuous heart rate, sleep, activity, glucose and other physiological signals can reveal changes that are invisible in periodic clinic visits. Their principal value is not the sensor itself but the ability to detect deviation from an individual baseline and integrate that deviation with other clinical information. The challenge is distinguishing actionable signal from noise without increasing false alarms or clinician burden.

3.5 Interoperability and semantic integrity

A personalized-care architecture requires machine-readable, semantically stable data exchange. HL7 FHIR genomics standards are explicitly intended to support exchange of structured genomic findings, variants, haplotypes, genotypes and annotations with clinical systems. This is foundational: AI cannot reliably personalize decisions if the meaning of a laboratory result, genomic variant or treatment event changes across systems. [20]

4. High-value clinical and R&D applications

4.1 Earlier and more precise diagnosis

AI can improve diagnosis in two distinct ways. The first is perceptual—finding subtle patterns in images, waveforms or pathology. The second is integrative—combining weak signals across modalities into a stronger estimate of disease probability. The second capability is particularly relevant to personalized medicine because many early disease states do not produce a single definitive marker.

The practical design question is not simply whether an algorithm can outperform a clinician on a benchmark. It is whether the human-AI team can reduce missed disease, shorten time to diagnosis, avoid unnecessary testing and preserve performance across sites and subgroups. FDA transparency principles emphasize intended use, data characteristics, limitations, workflow context and ongoing performance precisely because these factors influence the safety of real-world deployment. [21]

4.2 Precision oncology

Oncology remains the clearest proving ground for AI-enabled personalized medicine because molecular heterogeneity is central to disease classification and therapy choice. AI can integrate somatic variants, gene expression, histopathology, imaging, prior treatments and clinical variables to support tumor subtyping, prognosis, therapy ranking and trial matching. Multi-omics reviews increasingly report that the main barrier is not model sophistication but reproducibility, external validation and translation into a defined clinical decision. [6,7]

Liquid biopsy adds a longitudinal layer. Cell-free DNA can support detection, characterization and monitoring across the cancer-care continuum. Combined with AI, serial molecular measurements could enable response assessment and recurrence detection earlier than conventional episodic approaches, although clinical utility depends on indication-specific evidence and appropriate thresholds for action. [22]

4.3 Pharmacogenomics and medication optimization

Pharmacogenomics is one of the most concrete links between molecular data and therapeutic choice. Genetic variants can alter drug metabolism, exposure, toxicity or response. FDA resources document associations used in labeling and prescribing decisions, while large biobank analyses show how population-scale genomic data can identify actionable pharmacogenetic phenotypes. AI becomes valuable as the problem expands from one gene–one drug rules to interactions among multiple genes, comorbidities, co-medications, organ function and longitudinal outcomes. [13,14,23]

4.4 Cardiometabolic and chronic disease management

Chronic disease is an attractive target because risk evolves continuously and treatment can be adjusted iteratively. AI can combine ECGs, imaging, laboratory trajectories, medication history and wearable data to estimate decompensation risk, identify subphenotypes or tailor follow-up intensity. The personalized goal is not merely prediction but risk-adjusted intervention: who requires escalation, which intervention is likely to help, and when the benefit of intervention exceeds the burden or risk.

4.5 Rare disease and diagnostic odysseys

Rare disease illustrates why multimodal synthesis matters. Individually, phenotype terms, family history, imaging and genomic variants may be nonspecific. AI-assisted variant prioritization and phenotype matching can narrow the search space, while generative tools can summarize years of longitudinal records. The system must still preserve expert genetic interpretation and avoid overclaiming uncertain variants.

4.6 Drug discovery, biomarker development and clinical trials

AI is also moving upstream into therapeutic development. A 2025 Nature Medicine review described applications across target identification, molecular design, biomarker discovery and clinical development, and a 2025 phase 2a study of an AI-discovered drug/target combination was highlighted as an important clinical milestone. For personalized medicine, the deeper opportunity is to link drug mechanism to the molecularly defined patient subgroup most likely to benefit. [24,25]

Trial design can become more adaptive as AI improves phenotyping, site selection, patient matching and endpoint measurement. Digital health technologies can also support decentralized data collection. However, enrichment strategies can inadvertently narrow generalizability if they select a population that is too optimized for the model or biomarker. Personalized trials therefore need an explicit plan for external validity and access. [26]

The definition of AI-augmented personalized medicine becomes more concrete when applied across clinical and research domains. The table maps each domain to a near-term use—such as selecting therapy, refining a diagnostic pathway, adjusting a dose, or improving trial enrollment—and identifies its potential strategic value and principal evidence risk. Together, the examples show that personalized medicine is not limited to genomics or to a single algorithm: it can combine molecular, imaging, pathology, clinical, and longitudinal information to support decisions across prevention, diagnosis, treatment, monitoring, patient engagement, and drug development.

Domain	Near-term personalized use	Potential strategic value	Primary evidence risk
Oncology	Stratify treatment using molecular, pathology, and imaging data	High	Dataset shift; small subgroup samples
Pharmacogenomics	Select therapy or dose using genotype and clinical context	High	Evidence varies by gene–drug pair

Domain	Near-term personalized use	Potential strategic value	Primary evidence risk
Radiology/ pathology	Detect findings and quantify phenotypes	High	Workflow fit; spectrum bias
Cardiometabolic	Update risk and guide monitoring	Medium–high	Alert fatigue; calibration drift
Rare disease	Prioritize diagnoses using phenotype and genotype	Medium–high	Variant uncertainty; overstated confidence
Drug development	Identify targets and biomarkers; enrich trials	High	Translational validity
Patient engagement	Provide personalized education and care navigation	Medium	Incorrect output; comprehension; unequal access

The table is intended as a strategic map, not a ranking of proven clinical outcomes or deployment readiness. “High” and “Medium–high” describe the potential value of a use case if it is validated and integrated successfully; they do not mean that the AI is clinically mature or ready to act autonomously. The evidence risks point to what must be addressed for each application: representative data and adequate subgroup sizes in oncology; gene–drug evidence in pharmacogenomics; workflow fit and generalizability in imaging; calibration and alert burden in cardiometabolic care; variant uncertainty in rare disease; and proof that biomarkers or trial-enrichment methods translate into better development decisions. Patient-engagement tools add risks around inaccurate generated guidance, patient comprehension, and equitable access.

The practical implication is that organizations should evaluate each use case against its intended decision, recipient, and workflow—not treat “AI in personalized medicine” as a single deployment category. The table helps identify where to focus validation, monitoring, human review, and investment before expanding the system’s clinical authority.

5. Generative AI and agentic workflows

Generative AI changes the interface to complex clinical information. Instead of requiring a user to navigate multiple systems, a model can synthesize the patient record, retrieve relevant evidence, draft documentation, explain options in patient-appropriate language and propose next steps. The 2025 Nature Medicine review of generative AI in medicine emphasizes both this widening task range and the need for rigorous validation. [27]

Ambient AI documentation is a useful early case because the task is valuable but bounded. Large health-system studies have associated AI scribe adoption with reductions in EHR/documentation time and modest productivity gains. This does not prove that generative AI is safe for autonomous diagnosis or therapy selection; it illustrates a broader adoption principle: begin where errors are visible, reversible and reviewable, then expand authority only as evidence matures. [15,16]

From copilots to clinical agents

The distinction between a copilot and an agent is operational. A copilot generates information for a human. An agent can pursue a goal across multiple steps using tools—querying a chart, checking criteria, drafting an order, routing it for approval, scheduling a follow-up and monitoring whether the action occurred. In personalized medicine, agents could eventually coordinate complex pathways such as hereditary-cancer evaluation, molecular tumor boards or pharmacogenomic medication review.

This architecture can create substantial value, but it also multiplies risk because errors can propagate across steps. The control system therefore becomes as important as the model: permissions, approved tools, identity, policy constraints, structured outputs, logging, human approval points and exception handling should be designed before an agent is allowed to execute clinical actions.

Workforce readiness and human-agent collaboration

Clinical agents will create value only when the people working with them know how to supervise and respond to their actions. Physicians, nurses, pharmacists, laboratory and imaging technicians, and care coordinators need role-specific preparation to understand what an agent may do, distinguish a recommendation or draft from a completed action, verify important information against its source, recognize uncertainty, and intervene when an action is incorrect, incomplete, or outside the agent's authority. They also need clear procedures for approval, escalation, override, and handoff.

This is not simply a matter of teaching staff to write better prompts. Training should use realistic clinical workflows and include missing or conflicting data, out-of-scope requests, agent errors, time-sensitive decisions, and changes to models or tools. Teams should rehearse how to maintain human authority and close the loop when an agent initiates a task. Competency should be assessed and refreshed as workflows change. Training, agent interface design, permissions, staffing, and accountability must be developed together; nominal human review is not an effective safeguard if staff lack the time or information to perform it.

6. Evidence: moving from model performance to clinical utility

A persistent weakness in healthcare AI is the tendency to treat retrospective model performance as sufficient evidence. It is not. One commonly reported measure is AUROC (area under the receiver operating characteristic curve), which measures how well a model ranks patients who experience an outcome above those who do not across possible decision thresholds. A value of 0.5 represents chance-level ranking; 1.0 represents perfect separation in the tested data. A high AUROC does not show that the model's predicted probabilities are accurate, that a particular threshold is clinically appropriate, or that using the model improves care. A model can score highly while performing less well for a patient subgroup, generating too many false positives, or disrupting workflow. Even accurate predictions may have little clinical value if they do not change decisions or improve outcomes. Evaluation must therefore include calibration, threshold-specific performance, subgroup results, workflow effects, and demonstrated clinical utility.

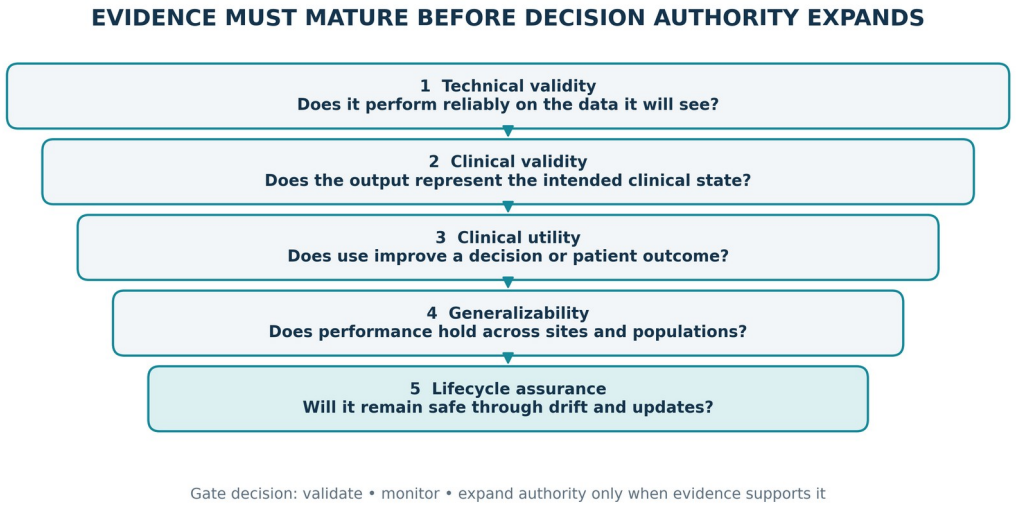


Exhibit 5 | The evidence ladder for AI-enabled personalized medicine

6.1 Technical validity

Technical validity asks whether the system is implemented correctly and performs as specified under defined conditions. It includes data quality, preprocessing, model reproducibility, software verification and analytic performance. For generative models, this also includes prompt/version controls and repeatability considerations.

6.2 Clinical validity

Clinical validity asks whether the output accurately represents the intended clinical state or predicts the intended outcome. Evaluation must match the target population and care setting, with subgroup analyses that are clinically meaningful rather than cosmetic.

6.3 Clinical utility

Clinical utility asks whether using the system improves a decision, process or patient outcome compared with usual care. This may require prospective observational studies, pragmatic trials or randomized evaluation depending on risk. The endpoint should correspond to the real value proposition—time to diagnosis, avoided adverse events, treatment response, clinician time, cost or patient experience—not merely model accuracy.

6.4 Generalizability and equity

Models can fail when moved across scanners, laboratories, EHR configurations, patient populations or clinical workflows. Diverse data are therefore not only an equity objective; they are a performance requirement. The All of Us program demonstrates the scientific value of intentionally broadening ancestry and demographic representation in precision-medicine datasets. [2,3]

6.5 Lifecycle assurance

Healthcare AI is not static. Input distributions change, clinical practice changes, vendors update models and local populations shift. The FDA's lifecycle guidance and final guidance on predetermined change control plans reflect a regulatory move toward managed iteration: planned modifications, validation methods, impact assessment and ongoing performance monitoring should be specified prospectively where appropriate. [10,11]

7. Governance, regulation and trust

The central governance problem is that AI can increase both the speed of insight and the speed of error. Personalized medicine magnifies this issue because the recommendation may depend on sensitive genomic data, uncertain biomarkers and complex probabilistic reasoning. Governance must therefore operate at the use-case level, not only at the enterprise-policy level.

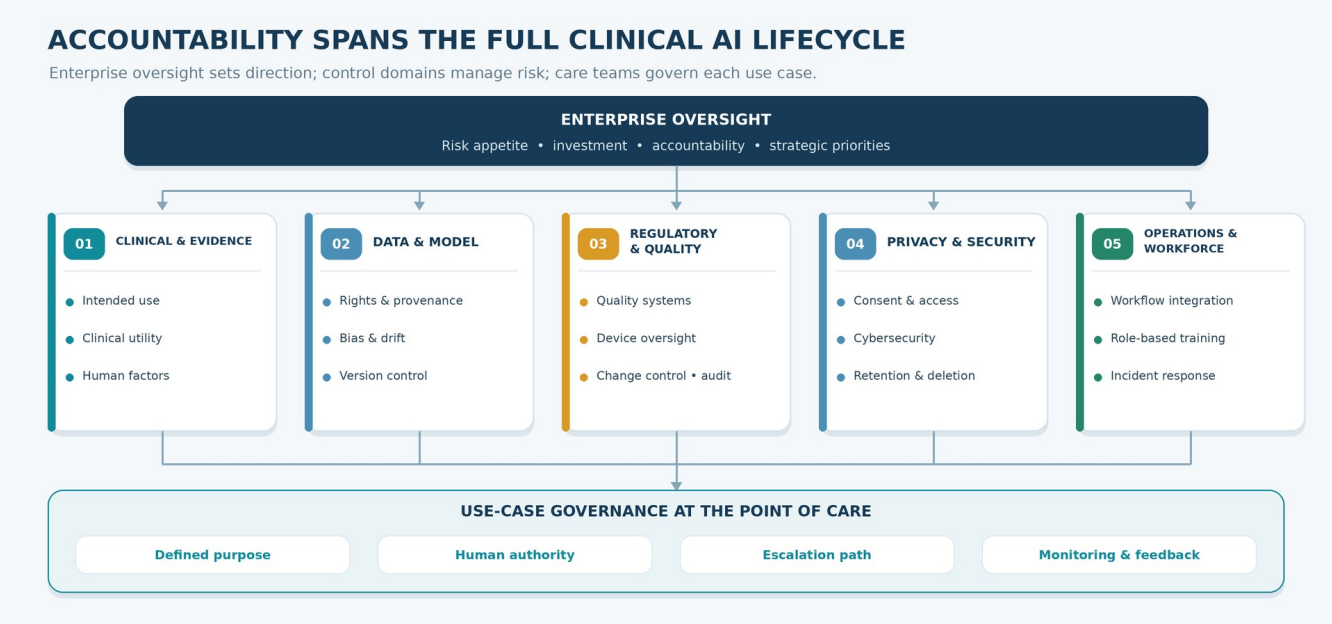


Exhibit 6 | Governance architecture for clinical AI and personalized medicine

7.1 United States: lifecycle regulation and transparency

FDA policy has evolved from an emphasis on premarket review toward a total-product-lifecycle perspective. The agency's 2025 draft guidance for AI-enabled device software functions addresses development, documentation, monitoring, bias and transparency, while its August 2025 final guidance on predetermined change control plans provides a path for specified, validated modifications without a new marketing submission for every approved change. [10,11]

For health IT, the ONC HTI-1 final rule established algorithm-transparency requirements for predictive decision support in certified health IT, giving clinical users baseline information to assess fairness, appropriateness, validity, effectiveness and safety. This is strategically important because model governance increasingly extends beyond the device manufacturer to the healthcare organization deploying the model. [9]

7.2 Europe: the AI Act and medical-product regulation

The European AI Act is now in force with phased obligations. Transparency rules for certain AI systems became applicable in August 2026, while high-risk requirements have later transition dates, including for AI embedded in regulated products such as medical devices. The European Commission has also published guidance and draft classification materials, while the European Medicines Agency has adopted a reflection paper on AI across the medicinal-product lifecycle. Global developers should expect greater convergence around risk management, data quality, human oversight, traceability and lifecycle controls even when legal mechanisms differ. [28,29]

7.3 Ethical principles

WHO guidance on AI for health and on large multimodal models emphasizes human autonomy, safety, transparency, accountability, inclusiveness and sustainability. In personalized medicine, these principles have practical consequences: patients should understand when AI materially influences care; clinicians should understand the intended use and limitations; sensitive data should be governed explicitly; and organizations should evaluate whether benefits and errors are distributed equitably. [8,30]

7.4 A minimum governance checklist

1. Intended use is narrowly stated, including target population, user, workflow and excluded uses.
2. Data provenance, consent, permissible use and retention are documented.
3. Training and validation populations are characterized, with subgroup performance reported.
4. Clinical thresholds and confidence/uncertainty are meaningful to the intended user.
5. Human authority is explicit: advise, approve, review, supervise, exception-handle or audit.
6. Model/version identity is traceable at the time of each output.
7. Local acceptance testing is completed before broad deployment.
8. Performance, drift, incidents and overrides are monitored after deployment.
9. Updates follow controlled validation and change-management procedures.
10. Vendor agreements address data use, IP, portability, security, audit rights and exit.

8. Operating model and economics

The economic case for AI-enabled personalized medicine should be built at the level of a clinical pathway. Enterprise AI programs often fail to create measurable value because the model sits outside the operational bottleneck. A useful economic model begins with the decision being improved and traces how that decision changes resource use and outcomes.

Value dimension	Illustrative metric	Why it matters
Clinical	Sensitivity at actionable threshold; adverse events avoided; response rate; recurrence detected earlier	Demonstrates utility beyond technical performance
Workflow	Minutes per encounter; turnaround time; number of handoffs; clinician override rate	Determines adoption and capacity impact
Patient	Time to diagnosis; travel burden; treatment fit; comprehension; adherence	Captures experience and access
Economic	Cost per correctly managed patient; avoided utilization; throughput; trial cycle time	Connects adoption to enterprise economics
Strategic	Proprietary data accumulation; learning rate; vendor dependence; model portability	Measures durable competitive position

The most defensible business cases typically involve one or more of four value mechanisms: avoiding a costly error, reducing time in a constrained workflow, improving conversion to an effective

treatment, or creating a learning asset that compounds over time. Personalized medicine can generate all four, but only when data capture and outcome measurement are part of deployment.

Build, buy or partner?

Most organizations should use a hybrid strategy. Foundation models, infrastructure and commodity tooling are usually better sourced externally. Patient-specific knowledge assets, proprietary assay data, workflow logic, decision policies and clinically validated integrations may be strategically important to own or tightly control. The decision should be based on differentiation, regulatory responsibility, data sensitivity, switching cost and the organization's ability to validate and maintain the system.

The build/buy/partner decision has different owners across the personalized-medicine ecosystem. A health system decides which capabilities to procure, integrate, and govern for clinical use. An AI or personalized-medicine company decides which product and infrastructure layers to build and which to source. Shared work does not blur accountability: providers define clinical policies and govern local use, while product companies remain accountable for the design and performance of the products they supply. Make the choices layer by layer, with explicit ownership for data, clinical decisions, evidence, and ongoing operations.

Build, buy, and partner choices by system layer

System layer	Build, buy, or partner posture and lead	Critical success requirements
Models and infrastructure	Buy or partner for foundation models, cloud, and commodity tools. AI companies select components; providers review them for intended use and security.	Assess fit for the defined use. Control versions, security, portability, and exit options.
Data and patient-state layer	Partner for external sources and platforms; co-develop integrations. The provider governs permitted access and use.	Connect authorized records, labs, blood tests, imaging, pathology, molecular, and patient-reported data. Preserve identity, provenance, timing, quality, and gaps.
Models and decision support	Buy or partner for standard capabilities; build selectively when distinctive data, methods, or workflow value justify ongoing maintenance.	Define the use and user. Validate calibration, subgroup performance, uncertainty, and clinical utility.
Agents and workflow	The AI company builds or configures agents; the provider co-designs workflows and sets permitted actions.	Bound tasks, tools, and permissions. Make status visible; define approval, escalation, audit, and fallback paths.
Clinical policies and authority	Provider-led. Clinical leaders and governance groups define policy; vendors implement technical controls.	Name the decision owner and output recipient. Distinguish advice from action; define review, override, exceptions, and patient involvement.
Workforce readiness	The provider leads role-based training; the vendor supports product-specific practice.	Prepare clinicians, technicians, pharmacists, and coordinators to verify outputs, recognize uncertainty, supervise tasks, escalate, close handoffs, and report incidents. Assess and refresh competency.
Validation and lifecycle governance	Shared. Suppliers document product evidence; providers validate local use and monitor performance; quality and regulatory teams manage changes.	Assess technical and clinical validity, utility, subgroup performance, and drift. Maintain change control, audit trails, incident response, and fallback procedures.
Patient engagement	Co-design with patients and care teams; the provider governs clinical content and escalation.	Use understandable, accessible, language-appropriate content. Protect privacy, prevent unsupported guidance, and provide a clear route to the care team.

System layer	Build, buy, or partner posture and lead	Critical success requirements
Value realization and scale	The provider defines the pathway and baseline; partners measure jointly; clinical and executive leaders decide whether to scale.	Track outcomes, safety, equity, experience, staff burden, access, time to diagnosis and treatment, and economics. Treat a one-day target as aspirational and limited to appropriate cases with required evidence available.

The practical rule is to buy commodity capability, build or retain control over differentiated clinical workflows and decision policies, and partner for data, expertise, and access that no single organization can efficiently provide. Buying a model does not remove the need for local validation; building an agent does not by itself create clinical value. Scale only when the technology, data, workflow, workforce, evidence, and accountability model work together.

9. A three-horizon implementation roadmap

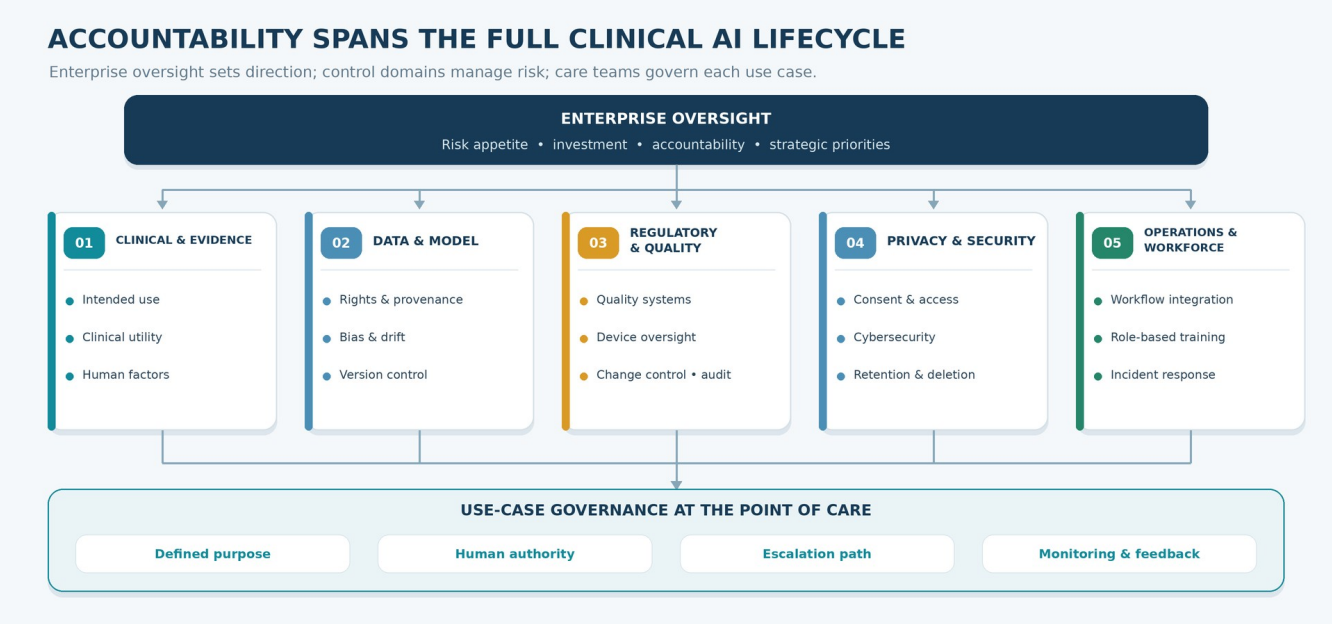


Exhibit 7 | Three-horizon roadmap from bounded pilots to a learning personalized-medicine platform

Horizon 1: establish the foundation (0–6 months)

- Select two or three use cases with a defined decision, measurable baseline and reviewable AI output.
- Map data sources, data rights, interfaces, vendors, security controls and regulatory classification.
- Create an AI/precision-medicine governance council with clinical, quality, regulatory, data and operational ownership.
- Define evidence requirements before development or procurement. Include subgroup analysis and local validation.
- Establish baseline clinical, workflow and economic metrics so benefit can be measured.

Horizon 2: prove utility and integrate (6–18 months)

- Run prospective evaluations in the intended workflow, with explicit human review and escalation.
- Integrate approved outputs into EHR, LIS, PACS, molecular-reporting or R&D systems rather than creating a separate destination.

- Build monitoring dashboards for adoption, overrides, drift, subgroup performance and safety events.
- Standardize model cards, validation reports, change-control records and vendor-risk reviews.
- Scale only those use cases that show measurable clinical or operational value.

Horizon 3: build the learning system (18–36 months)

- Link multimodal data longitudinally to create reusable patient and disease representations.
- Create governed pathways for continuous evidence generation and model recalibration.
- Introduce agent-assisted orchestration where tasks are bounded, auditable and reversible.
- Use portfolio governance to retire low-value models and concentrate investment in differentiated pathways.
- Treat the longitudinal data–evidence loop as a strategic asset, with explicit controls on vendor access and reuse.

10. Outlook: toward continuously learning medicine

Over the next several years, three trends are likely to converge. First, multimodal models will make it easier to combine molecular, imaging and longitudinal clinical information. Second, generative interfaces will reduce the friction of using these models by synthesizing records and evidence. Third, agentic systems will begin to coordinate defined workflows rather than merely produce recommendations.

The long-term destination is sometimes described as the “digital twin”: a computational representation of an individual that can simulate future states or treatment responses. The term is useful as an aspiration but should not imply more certainty than the evidence supports. In most near-term clinical settings, the more realistic architecture is a probabilistic patient model that is continuously updated, exposes uncertainty, and supports bounded decisions under clinician oversight.

This distinction matters. Personalized medicine will advance fastest when organizations resist the temptation to market simulated precision as clinical certainty. The strongest systems will be calibrated, transparent about limitations, validated in the intended context and embedded in workflows that preserve human judgment.

What will differentiate leaders

- A proprietary, longitudinal data advantage rather than a collection of disconnected AI pilots.
- Evidence-generation capability integrated into product and clinical operations.
- Interoperability that makes molecular and AI outputs usable at the point of decision.
- A governance model that permits rapid innovation without uncontrolled autonomy.
- An explicit strategy for external model/vendor dependence, data rights and portability.
- A clinical operating model that measures human-AI team performance, not algorithm performance alone.

Conclusion

AI and personalized medicine are converging because each solves a limitation of the other. Personalized medicine supplies the biological and clinical rationale for individualization; AI supplies

the computational machinery needed to integrate complex data at scale. Together they can shift care from broad classification toward dynamic, patient-specific decisions.

The opportunity is substantial, but the central challenge is no longer whether AI can generate a prediction. It is whether healthcare and life-science organizations can turn that prediction into a safe, evidence-based, economically useful and continuously monitored clinical action. The organizations that solve that operating problem will define the next era of precision healthcare.

References

1. U.S. Food and Drug Administration. Artificial Intelligence-Enabled Medical Devices. Updated 2026. FDA reports more than 1,600 AI-enabled medical devices authorized in the U.S. as of September 2026.
2. All of Us Research Program Genomics Investigators. Genomic data in the All of Us Research Program. Nature. 2024;627:340–346. PMID 38374255.
3. National Institutes of Health. 275 million new genetic variants identified in NIH precision medicine data. February 19, 2024.
4. Carrillo-Larco RM, Poulakakis-Daktylidis A. Artificial intelligence in precision medicine: transforming disease subtyping, medical imaging, and pharmacogenomics. 2025. PMID 40833051.
5. Zack M, Stupichev DN, Moore AJ, et al. Artificial Intelligence and Multi-Omics in Pharmacogenomics: A New Era of Precision Medicine. Mayo Clin Proc Digit Health. 2025;3(3):100246. doi:10.1016/j.mcpdig.2025.100246.
6. Hsu C-Y, et al. AI-driven multi-omics integration in precision oncology: bridging the data deluge to clinical decisions. Clin Exp Med. 2025. PMID 41266662.
7. Integration of artificial intelligence and multi-omics for precision medicine. Review. 2026. PMID 42669717.
8. World Health Organization. Ethics and governance of artificial intelligence for health: guidance on large multi-modal models. WHO; 2025.
9. Office of the National Coordinator for Health Information Technology. HTI-1 Final Rule: Health Data, Technology, and Interoperability—Certification Program Updates, Algorithm Transparency, and Information Sharing.
10. U.S. Food and Drug Administration. Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations. Draft guidance. January 2025.
11. U.S. Food and Drug Administration. Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence-Enabled Device Software Functions. Final guidance. August 2025.
12. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF) and Generative Artificial Intelligence Profile, NIST AI 600-1.
13. U.S. Food and Drug Administration. Table of Pharmacogenomic Biomarkers in Drug Labeling. Updated 2026.
14. U.S. Food and Drug Administration. Table of Pharmacogenetic Associations. Updated 2026.
15. Evaluation of an Ambient Artificial Intelligence Documentation Platform for Clinicians. JAMA Network Open. Published May 2, 2025. doi:10.1001/jamanetworkopen.2025.8614.
16. Changes in Clinician Time Expenditure and Visit Quantity With Adoption of Artificial Intelligence–Powered Scribes: A Multisite Study. JAMA. 2026.
17. Calcaterra V, Guardamagna L, Gatti A, et al. Digital twins in healthcare: A systematic review of current applications, frameworks, and future directions. Digit Health. 2026;12:20552076261415934. doi:10.1177/20552076261415934. PMID: 42370013.
18. Xu J, Ko JM, Kvedar JC. AI agents in clinical practice: an evidence map. npj Digit Med. 2026;9:854. doi:10.1038/s41746-026-02960-4. PMID: 42443480.
19. Wei X, et al. Advancing precision medicine: the transformative role of artificial intelligence in immunogenomics, radiomics, and pathomics for biomarker discovery and immunotherapy optimization. 2025.
20. HL7 International. Genomics in FHIR, FHIR Release 5.0.0. Clinical Genomics Work Group.
21. U.S. Food and Drug Administration, Health Canada, MHRA. Transparency for Machine Learning-Enabled Medical Devices: Guiding Principles. 2024; FDA page updated 2026.
22. Landon BV, Annapragada AV, Niknafs N, et al. Liquid biopsies across the cancer care continuum. Nature Medicine. 2025;31:4006–4021.
23. Pharmacogenetics: Opportunities for the All of Us Research Program and Other Large Data Sets to Advance the Field. 2025. PMCID PMC12262141.

24. Zhang K, Yang X, Wang Y, et al. Artificial intelligence in drug development. Nature Medicine. 2025;31:45–59.
25. Zitnik M. AI-enabled drug discovery reaches clinical milestone. Nature Medicine. 2025;31:2490–2491.
26. U.S. Food and Drug Administration. Decentralized Clinical Trials for Drugs, Biological Products, and Devices: guidance and related digital-health recommendations.
27. Teo ZL, Thirunavukarasu AJ, Elangovan K, et al. Generative artificial intelligence in medicine. Nature Medicine. 2025;31:3270–3282.
28. European Commission. AI Act: regulatory framework for artificial intelligence; implementation timeline and high-risk systems guidance. Updated 2026.
29. European Medicines Agency. Reflection paper on the use of Artificial Intelligence (AI) in the medicinal product lifecycle. EMA/CHMP/CVMP/83833/2023. Adopted 2024.
30. World Health Organization. Ethics and governance of artificial intelligence for health. WHO guidance. 2021.

About CatalystDx Partners

CatalystDx Partners provides executive advisory support at the intersection of diagnostics, R&D transformation, artificial intelligence, clinical workflow and precision medicine. This paper is intended for healthcare, diagnostics, life-science and investment leaders evaluating the strategic and operating implications of AI-enabled personalized medicine.