OUTCOMES ROCKET | CENSINET®

THE FUTURE STATE OF HEALTHCARE IN 2026

The Vision Of 30 Leaders Across AI, Clinical Practice, Informatics & Innovation

Table of Contents

1 Foreword

Toward a More Intelligent Healthcare System

SAUL MARQUEZ,
CEO and Founder,
Outcomes Rocket

ED GAUDET,
CEO and Founder,
Censinet

If there is one thing the last few years have made undeniable, it is that healthcare is both more **resilient** and more **fragile** than most people realize. Resilient because clinicians, care teams, and operators keep showing up, even when the system asks too much.

Fragile because the way we move information, make decisions, coordinate care, and pay for outcomes still too often depends on workarounds and hope.

In 2026, the conversation is moving from experimentation to accountability. Artificial intelligence (AI), modern data platforms, and new models of digital care are becoming part of the operating environment. The standard is whether intelligence can be integrated into healthcare in a way that is safe, measurable, secure, and worthy of trust.

This white paper, ***The Future State of Healthcare in 2026: The Vision of 30 Leaders Across AI, Clinical Practice, Informatics & Innovation,*** draws directly from semi-structured interviews conducted with 30 experts and leaders at **AIMed 2025.**

These conversations reflect the perspective of people working close to the realities of care delivery, informatics, innovation, and applied AI. They bring optimism where it is earned, caution where it is necessary, and a consistent focus on what will actually improve outcomes for patients and reduce burden for the people doing the work.

This foreword is a commitment to disciplined ambition.

The leaders featured here are not claiming that AI will solve healthcare. Their views point to something more practical and more achievable: healthcare can learn faster, coordinate better, and support human expertise instead of draining it. The system can become more proactive and more connected.

Reaching that future state requires leadership that:

- Invests in data governance, interoperability, identity, auditability, and change management - the unglamorous infrastructure that makes intelligence scalable and safe.
- Designs technology around people, not dashboards, prioritizing usability and human judgment at the point of care.
- Sets measurable standards for responsible AI, so the industry can move with urgency while maintaining accountability.

What does the future of AI look like for your organization? We welcome your perspective and invite you to join the conversation. Please share your thoughts or reach out to:

Saul Marquez: **saul@outcomesrocket.com**
Ed Gaudet: **egaudet@censinet.com**

2 Executive Summary

Healthcare in 2026 is shifting from AI as experimentation to AI as infrastructure.

Healthcare enters 2026 with a more practical sense of urgency. The focus has shifted from what AI can achieve in theory to what performs consistently in real-world clinical environments. Health systems are looking for solutions that remain reliable under real clinical volume, workforce constraints, and regulatory oversight– while meeting established expectations for privacy, security, and responsible data use.

This white paper examines six shifts shaping healthcare in 2026, informed by industry voices from across AI, clinical practice, informatics, and innovation. A full list of these industry voices is included on page 32.

1 Agentic AI enters operational workflows:

Agentic systems are moving beyond single prompts and isolated assistants into tools that can execute multi-step processes. The emphasis is on reducing handoffs and rework while making routine work more consistent and trackable.

2 Workforce productivity becomes the primary ROI metric:

The ROI conversation is becoming more concrete and defensible, with leaders prioritizing measures tied to real capacity gains. In many organizations, productivity is increasingly linked to retention and burnout reduction.

3 Multimodal AI transforms clinical decision support:

The most trusted multimodal AI deployments are designed to reduce noise, surface relevant context, and support situational awareness at the point of care.

4 Healthcare platforms consolidate the AI ecosystem:

Fragmentation is giving way to consolidation as major EHRs, cloud vendors, and healthcare platforms become the primary delivery layer for AI. In parallel, health systems are rationalizing point solutions to reduce vendor sprawl and simplify security and compliance.

5 Edge AI expands access and diagnostic capability:

Edge deployments are bringing intelligence closer to the point of care. Across remote monitoring, portable diagnostics, and frontline triage support, these systems are expanding access in underserved settings and

6 Clinician-led innovation accelerates adoption:

In leading programs, clinicians define use cases, establish performance standards, validate results in real workflow conditions, and set guardrails for safe use. This approach builds credibility with care teams and narrows the gap between pilot success and durable deployment.

3 The 2026 Healthcare Trends

The healthcare landscape in 2026 rewards execution.

Organizations are expected to improve quality and expand access while operating with tighter constraints on cost and staffing.

AI is part of that story, but only when it integrates into real workflows and performs reliably under governance, security, and clinical oversight.

The trends below reflect the changes leaders described as most defining for the year ahead.

3.1 AGENTIC AI ENTERS OPERATIONAL WORKFLOWS

The 2026 Healthcare Trends

Trends and Case Studies

Agentic AI is increasingly defined not by conversational sophistication, but by operational autonomy. The most credible implementations are **task-executing systems** that execute multi-step workflows, manage handoffs, and close loops that historically stall in inboxes, phone queues, and fragmented administrative processes.

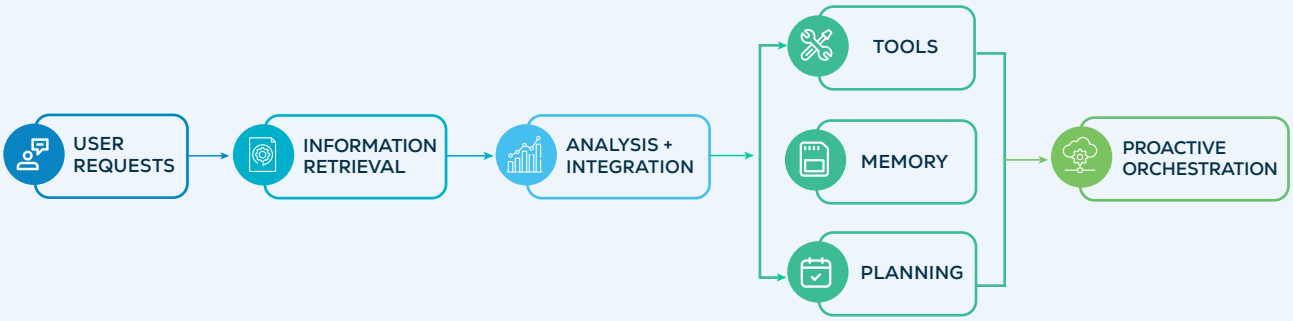

The shift is practical: **copilots compress thinking time, agents compress waiting time.** In healthcare, waiting time is often the more expensive constraint.

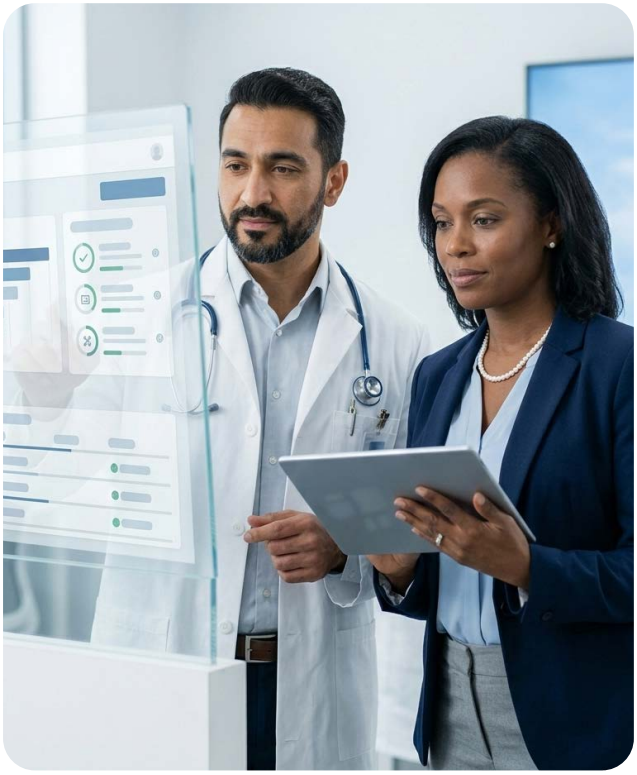

Always-on outreach that converts to scheduled care. A large health system is using AI agents to run preventive-care outreach in areas where manual calling and periodic campaigns often fall short on follow-through.

The relevant signal is operational performance, not engagement. In one deployment, agents completed the equivalent of **22 hours of outreach in 30 minutes, with 10% of contacted patients scheduled immediately. This is operationally significant because it** demonstrates sustained follow-through at scale—something healthcare has historically struggled to maintain.

The deeper implication for 2026 is that gap closure becomes a workflow capability. Once outreach is continuous, the limiting factor shifts downstream to readiness (scheduling capacity, escalation pathways, follow-up protocols, and care navigation).

61.4%

Adults 45–75 up to date on CRC screening.

CDC (PCD), 2025

28M

Adults 45–75 never screened for CRC.

CDC (PCD), 2025

60%

Medicare community beneficiaries with an AWV (2022).

CMS MCBS, 2022

Sources: *CDC Preventing Chronic Disease (2025) baseline CRC screening estimates; CMS MCBS preventive care use (2022).*

Prior authorization: removing dead time without removing accountability.

Prior authorization is a high-friction workflow defined by repeatable steps and long periods of non-value-added time. An emerging agentic pattern uses voice agents that call payers, wait on hold, complete structured intake steps, and then trigger a warm transfer when human intervention is required. The value is that it **removes the "hold-time" and reduces cycle time variability.**

The same approach can be applied to medical policy requirements. In one example, the time to compile the relevant policy requirements was described as typically 40–45 minutes, with complex cases expanding into hours, while an agent-driven approach can reduce the assembly step to seconds by preparing the required content for human submission. **In 2026, expect this category to expand rapidly** because it targets work that is structured, high-volume, and highly measurable.

Operational governance: agents that enforce guardrails continuously.

As autonomy increases, governance must move at operational speed.

A second agentic application targets compliance and risk management. It translates broad organizational requirements into enforceable micro-policies and monitors drift, hallucinations, and data lineage.

In mature implementations, governance is engineered as a control plane: logging actions, detecting exceptions, triggering intervention, and maintaining auditability while automation runs within defined boundaries.

Key Barriers

Experts consistently point to a **common set of barriers** that determine whether agentic AI becomes production infrastructure or remains a pilot.

Definition ambiguity creates misalignment.

"Agentic" is used inconsistently across the market. Practically, organizations need an internal definition anchored on autonomy level and permitted actions, supported by a simple internal taxonomy:

- **Level 1 (Assistive) -** Suggests or summarizes; no system changes.
- **Level 2 (Task executing) -** Performs scoped tasks in approved systems.
- **Level 3 (Autonomous + escalation) -** Runs multi-step workflows, handles exceptions, and escalates edge cases, with explicit approval required for high-risk actions.

Most agentic AI pilots don't fail on capability. They fail on governance, measurement, and ownership — problems that no model can solve.

Without that, teams cannot align on risk tolerance, accountability, or success metrics. **The result is scope creep and stalled decision-making.**

Data governance and PHI exposure shape architecture from day one.

Enterprise healthcare environments are highly risk-sensitive. Adoption frequently hinges on precise answers to foundational questions: What data is accessed? Where does it flows? What is stored? What is restricted? In some implementations, this has forced designs that operate outside the EHR and avoid handling PHI directly.

Agentic workflows that do not address data minimization early often face procurement friction and delayed deployment.

Measurement gaps keep organizations stuck in pilot mode.

Many "pilot failures" are measurement failures. When success metrics are not defined up front, organizations cannot prove value decisively or justify the additional governance overhead that autonomy requires.

Agentic AI raises operational and compliance stakes, so the bar for quantified impact is higher.

Risk ownership is unclear when autonomy increases.

Risk aversion can slow deployment even when the current process is clearly unsustainable. Without a named owner across operations, compliance, and security, agentic systems remain politically difficult to scale.

Recommendations and Solutions

To operationalize agentic AI, organizations should focus on a limited set of disciplined design and governance choices.

a) Start with bounded autonomy and engineered handoffs.

- Agents perform routine steps and preparation work.
- Humans approve or complete sensitive actions.
- Escalation triggers are explicit (warm transfer, review queue, exception routing) because it targets work that is structured, high-volume, and highly measurable.

This approach scales faster than attempting full autonomy and reduces the probability of high-impact errors.

b) Make PHI minimization the default deployment strategy.

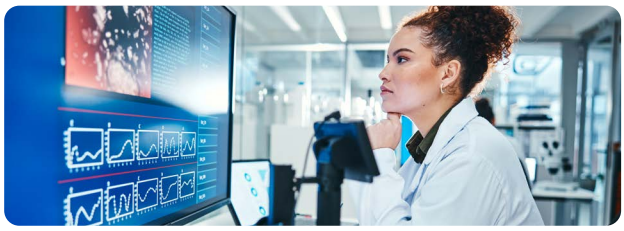

If data exposure is the core blocker, solve it structurally - limit what agents access, avoid persistence of sensitive data when possible, and implement designs that keep final PHI completion and submission within secure clinical systems.

Architectures that embed data minimization from the start encounter fewer procurement delays and governance objections.

c) Define success metrics in workflow terms before launch.

Agentic AI should be evaluated on operational KPIs leadership already tracks: cycle time, completion rates, hold time eliminated, staff minutes returned, escalation rate, error rate, and downstream impact such as denial reductions or improved scheduling conversion.

Clear pre-defined metrics prevent “pilot purgatory” and establish a credible path to scale.

d) Anchor the business case to measure administrative burden.

Prior authorization is a strong entry point because the pain is quantifiable. **The AMA’s 2024 prior authorization survey** reports that practices complete 39 prior authorizations per physician per week, with physicians and staff spending approximately 13 hours weekly on the process; 89% report it contributes to physician burnout. These metrics help stakeholders evaluate agentic initiatives with the same rigor applied to other operational investments.

3.2 WORKFORCE PRODUCTIVITY BECOMES THE PRIMARY ROI METRIC

The 2026 Healthcare Trends

Trends and Case Studies

Workforce shortages, rising clinical complexity, and the administrative weight of modern care have turned "productivity" into the most practical ROI lens.

One clear trend is that documentation relief is becoming the baseline expectation. Ambient listening is increasingly viewed as **"table stakes"** because it targets the most persistent productivity drains: note creation and after-hours cleanup.

For instance, in interviews with industry experts, ambient tools were described as **returning up to two to three hours of after-hours documentation time per clinician per day,** while improving the patient experience by allowing clinicians to focus on the interaction instead of the keyboard.

A second productivity pattern is emerging upstream of the visit. Rather than using AI to rewrite notes, some organizations are using it to collect and structure patient history before the clinician enters the room.

In one model, AI generates a detailed history and timeline, and prepares a structured summary of key issues and next steps. The downstream effect is **improved clinical efficiency and reduced burnout,** with an operational ripple effect too, such as fewer appointments running long and less schedule spillover at the end of the day.

Finally, there is a noticeable reframing of AI's role: the goal is **an AI and human partnership that filters low-value work.** One practical illustration is the idea of an **"AI peer"** that screens out low-quality submissions so human experts can focus on what requires real thought.

The same logic applies to patient preparation, where AI helps patients formulate questions before and after visits, improving the quality of the clinical encounter.

Recommendations and Solutions

If productivity is the ROI metric, then AI strategy must be built around measurable workforce outcomes.

a) Define productivity ROI as a workforce scorecard.
Use a balanced set of indicators that reflect real operational constraints:

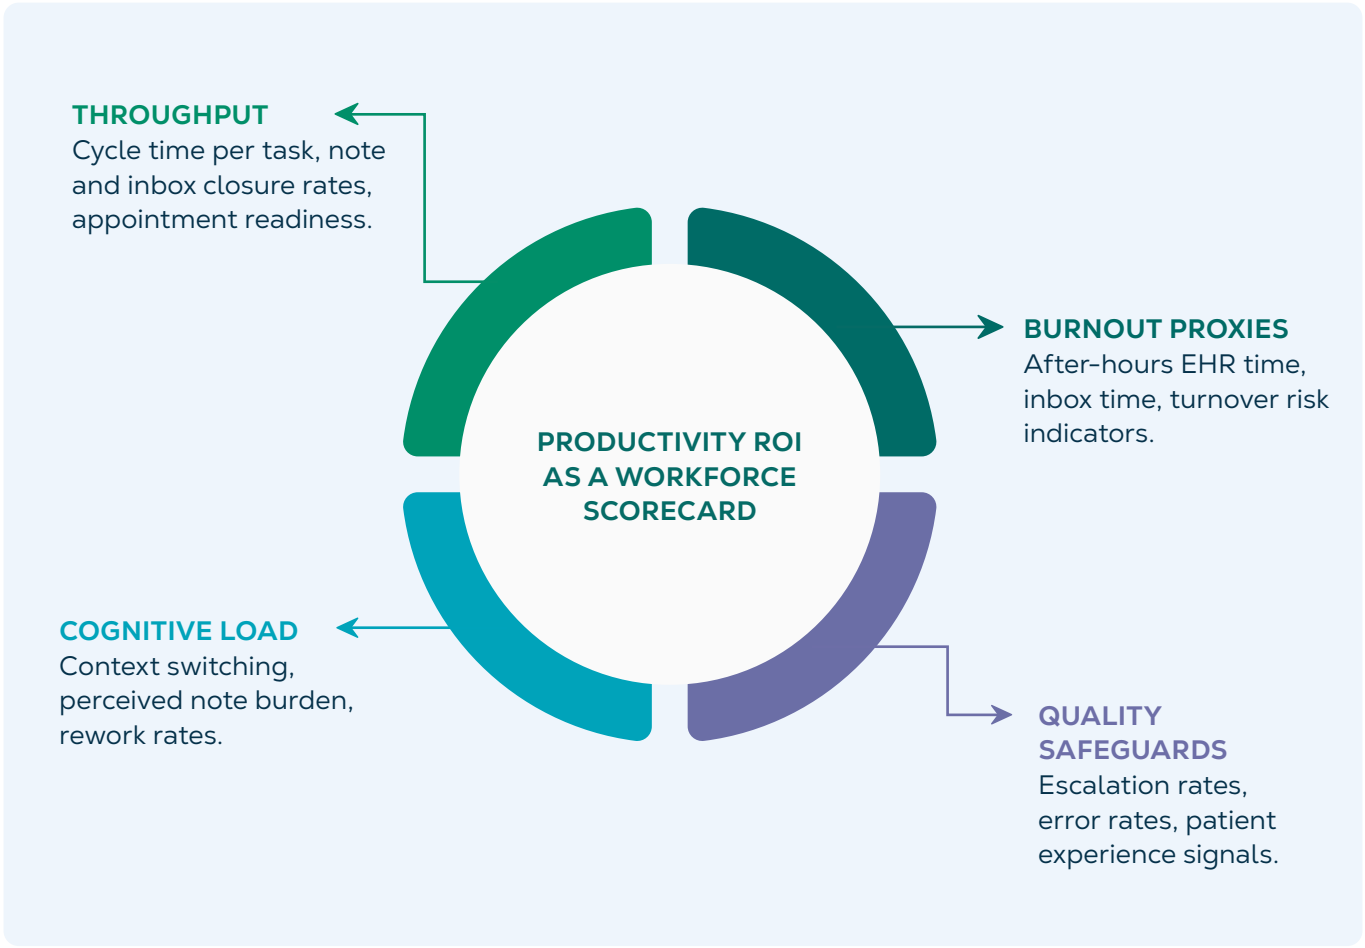

b) Move from "time savings" to "attention restoration" as the core design goal.

Treat documentation automation and pre-visit history collection as complementary. Ambient tools reduce the burden during and after visits; pre-visit structuring reduces cognitive load before the visit begins. Together, they return clinician attention to work that requires judgment, empathy, and reasoning.

c) Invest in resilience and human capital development as part of the AI program.

Normalize skill-building and peer learning so clinicians and staff feel agency. The strongest framing is "AI removes low-value drag so humans can do higher-value work."

3.3 MULTIMODAL AI TRANSFORMS CLINICAL DECISION SUPPORT

The 2026 Healthcare Trends

Trends and Case Studies

Multimodal clinical decision support is shifting from single-signal alerts to **context-aware insight engines** that fuse text, imaging, genomics, and physiologic signals into recommendations clinicians can trust and act on.

In imaging-heavy environments, organizations often deploy multiple narrow AI tools that each detect one signal well. Clinicians are still left to reconcile competing outputs.

A clear pattern emerging is **aggregation and orchestration.** Multiple models feed into a consolidated view, so clinicians get a single, **workflow-friendly output** instead of fragmented alerts. The shift to true decision support happens when the system adds an interpretation layer. It prioritizes what matters now, communicates uncertainty, and links outputs to next-step actions.

A second pattern is the **integration of omics-scale signals into decision-making,** even when patient examples are limited.

Many clinical questions are "large P, small N" problems– Thousands of features exist, but there are not many labeled cases. Multimodal value shows up when genomics is fused with clinical context, such as diagnoses, medications, labs, and comorbidities. This supports decisions about who is likely to respond, who is likely to fail, and where wasted time and avoidable costs.

Additionally, decision support **breaks most visibly at the edges:** Rare diseases, unusual presentations, and "looks-right-but-is-wrong" recommendations expose model limits. A strong technical theme is building **negative datasets and counterfactual examples** so models learn what not to do.

Key Barriers

Experts from **MedAI interviews** surfaced fundamental challenges that explain why **multimodal decision support is harder than building a high-performing model.**

- **Rare-case scarcity and "unknown unknowns"**

  High-value clinical decisions often live in low-frequency events. Without deliberate creation of counterfactual and negative examples, models can look strong in routine cases but fail when the clinician most needs support.

- **Data poisoning and provenance risks**

  As organizations draw training data and external knowledge from many sources, the risk surface expands. There is increasing attention to whether datasets are seeded with errors, whether provenance is clear, and how to detect manipulation early.

- **Reliability expectations are higher than typical AI deployments**

  Clinical decision support is held to a "near-zero surprise" standard. Drift, hallucinations, and inconsistent mappings are not minor defects. They directly erode trust.

  This is why production-ready approaches increasingly rely on hybrid systems that constrain model behavior and validate outputs.

Recommendations and Solutions

The biggest unlock is building the orchestration layer.

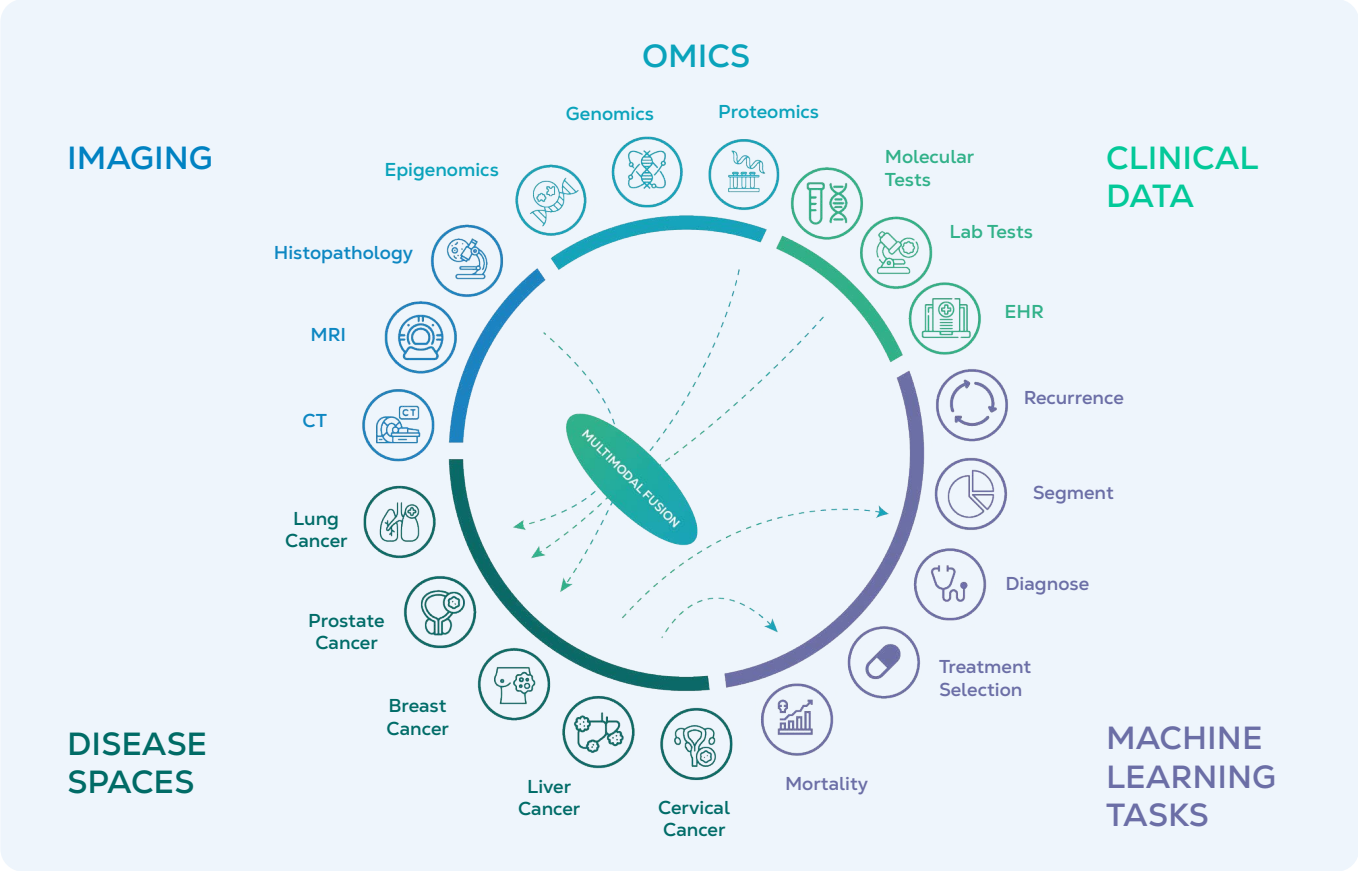

a) Treat multimodal decision support as an orchestration problem

Build modular inputs, a fusion layer, and one clinician-facing output that includes recommended actions and uncertainty. This matches the shift away from siloed point solutions toward consolidated decision support.

b) Institutionalize negative data and counterfactual evaluation

Make "**negative datasets**" a required artifact, especially for rare conditions and safety-critical pathways, so models learn what a condition is not as clearly as what it is. In stroke triage, for example, a well-built negative set helps the model tell true strokes apart from common look-alike mimics by learning the differences across imaging, key labs, and a small set of clinical cues.

Pair this with adversarial testing and provenance checks so the model learns how errors look and how they emerge.

c) Prioritize concept-level interoperability

Multimodal CDS improves when data is normalized at the level of clinical concepts rather than documents. Invest in terminology mapping and concept graphs to support coherent fusion.

d) Use hybrid pipelines to control behavior and reduce drift

Default to a pattern where language models handle narrative tasks and deterministic logic handles mappings, constraints, and validation. This improves operational stability and lowers cost, which matters for scaling.

3.4 HEALTHCARE PLATFORMS CONSOLIDATE THE AI ECOSYSTEM

The 2026 Healthcare Trends

Trends and Case Studies

Health systems are increasingly looking for a small number of **integrated AI layers** that can carry many use cases, instead of managing dozens of disconnected point solutions.

One theme is full stack standardization, as AI touches connected devices, data exchange, and clinical workflow. This requires a trust approach that spans identity, privacy, safety, and security – not as separate checkboxes but as one enforceable set of expectations across infrastructure and interoperability.

This aligns with work like **IEEE/UL 2933,** which sets a framework for clinical IoT data and device interoperability with TIPPSS principles, bridging devices and healthcare systems such as EHRs.

A second pattern is **EHR workflow anchoring**. One clinical leader described building a platform that takes voice or video, converts it to structured clinical text, and inserts it into the EHR in the right format.

He also referenced evidence that **dictation can be materially more efficient** than keyboard and mouse documentation.

In the broader market, the same direction shows up in the way **ambient documentation and generative AI are being embedded into EHR workflows** through major vendor partnerships, including Epic and Microsoft/ Nuance.

POINT-SOLUTION SPRAWL	PLATFORM OPERATING MODEL
Multiple disconnected tools for transcription, documentation, and data capture	Single integrated layer embedded into clinical workflows
High administrative burden for clinicians	Streamlined, voice- and video-first documentation
Difficult to scale and govern across the organization	Scalable via major vendor partnerships (e.g., Epic, Microsoft/Nuance)
Fragmented ROI and inconsistent data quality	Consolidated ROI with structured, EHR-ready clinical text

A third signal is the push for a **marketplace-level registry** so AI can be evaluated and monitored like any other clinical asset. As one expert argued, an AI "registry" should track how it performs over time and how it behaves across the marketplace, not only inside a single organization.

She emphasized this as an accountability mechanism, including the importance of lineage and an approach that is dynamic.

Taken together, these examples point to a **2026 environment** where "platform" is more about three capabilities that keep repeating across buyers' demands.

- Workflow-native deployment (especially EHR-embedded)
- Enterprise-grade trust controls across the stack
- Cataloging and monitoring that behaves like clinical infrastructure, not a pilot program

Key Barriers

There are some fundamental challenges that keep platforms from being “plug-and-play” in practice.

a) Static governance cannot keep up with model change

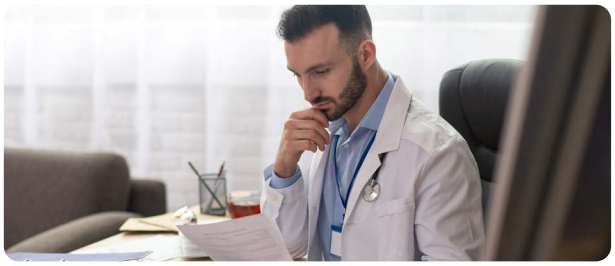

A recurring governance problem is **speed mismatch.** AI changes quickly. Documentation, model cards, and internal approvals move slowly. Jodyn pointed out that “nobody’s going to fill out a model card,” and even if they do, it risks becoming outdated as models and deployments evolve.

b) Procurement opacity and weak comparability

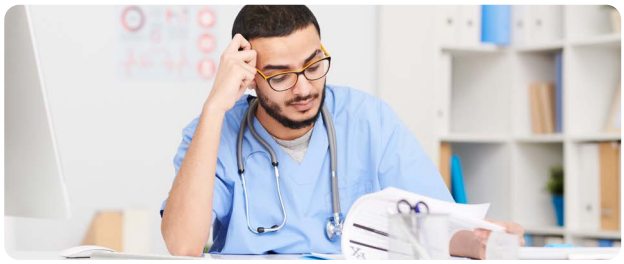

Without **shared disclosure norms,** buyers struggle to compare tools consistently. This is one reason **registries** matter. It is also why regulatory momentum is pushing **transparency** for predictive tools embedded in certified health IT.

Recommendations and Solutions

A platform strategy works best when it is treated as an operating model. The strongest implementations start by defining the “rails” that every AI capability must ride on.

a) Establish a platform reference architecture

WORKFLOW	DATA	MODEL SERVICES	GOVERNANCE
• Intake + triage	• Sources (EHR, claims, imaging)	• Model selection + routing	• Risk tiers + policies
• Human-in-the-loop	• Pipelines + access controls	• Agent runtime + tools	• Guardrails
• Exception handling	• Quality + labeling	• Orchestration	• Audit trails
• Continuous improvement	• Negative sets	• Evaluation + telemetry	• Compliance + reviews

Compliance + Reviews

Workflow Layer

EHR-native integration patterns, standardized note structures, order support handoffs

Data Layer

Governed access, provenance, auditability, and approved data products

Model Services Layer

Deployment, versioning, evaluation, rollback

Governance Layer

Risk classification, monitoring, incident response, change control

Define the minimum platform layers and their owners.

b) Build an AI registry that is operational

Treat the registry as a live system of record. Aim for a minimum record that can be updated automatically.

- Intended use and clinical domain
- Model version and deployment context
- Data lineage and dependency mapping
- Monitoring signals and event history
- Human override and escalation paths

3.5 EDGE AI EXPANDS ACCESS AND DIAGNOSTIC CAPABILITY

The 2026 Healthcare Trends

Trends and Case Studies

Edge AI is moving clinical intelligence closer to where care actually happens. Not inside a data center or an analytics team, but at the **point of care**, in community settings, and on consumer devices. Capture, inference, and action are starting to occur in the same workflow step.

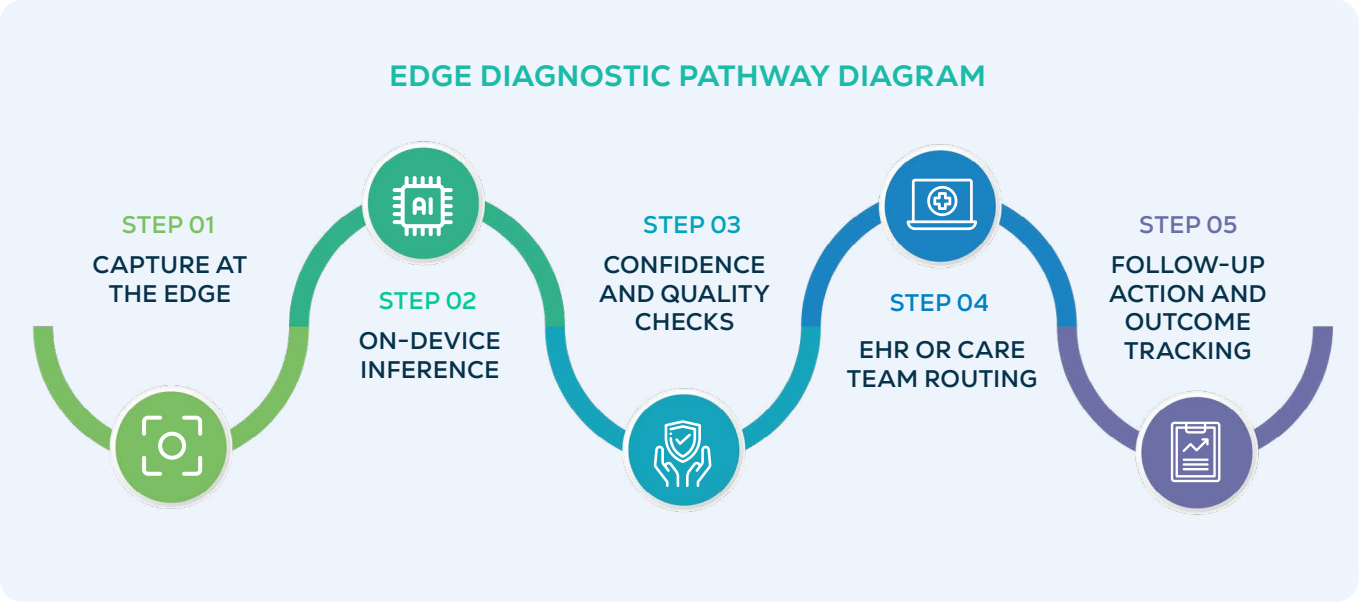

A practical illustration is how edge-enabled imaging is being designed to remove **operator variability.** In one case, diagnostic capture is becoming **"almost totally automated,"** with autofocus and auto-alignment built into the device so the system can reliably collect multiple diagnostic images and produce **AI-supported interpretation** in a standardized way.

The strategic significance is the operational design choice. When capture quality becomes consistent, the system can be deployed in more settings with **less dependence on scarce specialists.**

Edge AI also expands access by turning **passive sensing into an earlier clinical signal.** The wearable discussion makes a subtle but important point. Continuous monitoring changes when care can happen. When risk signals are detected earlier and more consistently, the clinician's role shifts from detective to decision-maker and coach.

The productivity and access gains come from spending scarce visit time **where judgment is needed,** not where data collection is missing.

The other access lever is **comprehension at the point of decision.** Edge AI is about whether a patient can understand and act on guidance, especially when language barriers, hearing challenges, or low health literacy stand between the result and follow-through.

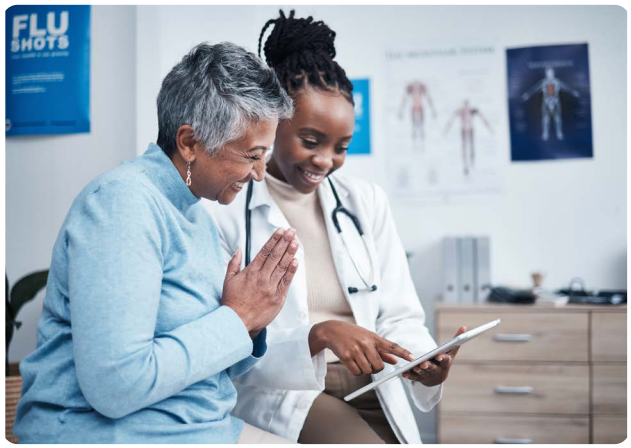

Key Barriers

Most barriers emerge after the demo. They stem from integration, governance, and real-world variability.

- **Regulatory and validation burden becomes device-bound**

  In device-based AI, outcomes are often tied to the specific hardware and workflow configuration. That raises the cost of proving performance and slows iteration, because validation requires large, representative datasets and reproducibility across real-world conditions.

- **Data representativeness at the edge is harder than it looks**

  Edge deployments amplify “messy reality” problems such as poor lighting, inconsistent positioning, varying devices, incomplete patient history, and unequal performance across populations. The edge is where bias shows up first because the environment is less controlled and how to detect manipulation early.

- **Trust and comprehension are part of diagnostic performance**

  If patients do not understand what an AI-enabled screening result means, or clinicians do not trust how it was derived, the pathway stalls. The translation and health literacy discussion is a reminder that edge AI must be designed for real patient communication.

- **Security and governance constraints tighten**

  Edge devices expand the attack surface and increase the number of locations where sensitive data can be captured. Governance also becomes more complex because you are managing a distributed network of endpoints, updates, and configurations.

Recommendations and Solutions

The edge opportunity is real, but it rewards disciplined design.

a) Start with "high-friction, high-volume" access problems.

Prioritize workflows where delays and shortages are already visible, like screening and triage.

b) Build a validation strategy that matches the device lifecycle

Treat the device plus model as a combined clinical product. Plan for performance monitoring, drift checks, and periodic re-validation.

c) Make escalation rules explicit and auditable.

Define what happens after an edge result is generated. Who is notified, what is documented, what thresholds trigger referral, and how exceptions are handled.

d) Design for capture quality before chasing model sophistication

Standardize what the device captures, how it captures it. The shift toward automated capture such as autofocus and auto-alignment is a practical blueprint.

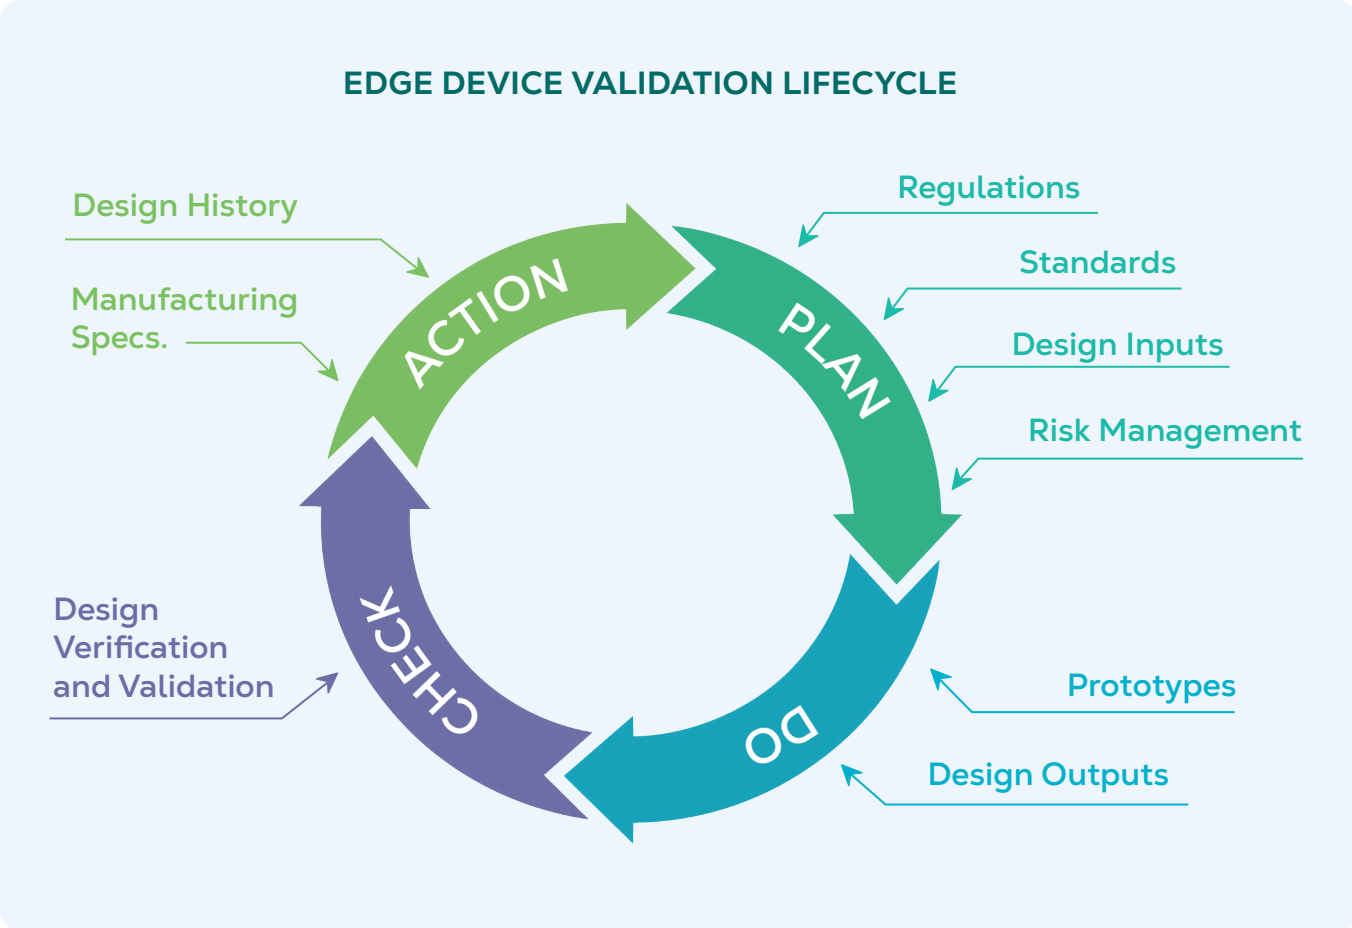

3.6 PATIENT TRUST AND CONSENT FRAMEWORKS MODERNIZE

The 2026 Healthcare Trends

Patient Trust and Consent Frameworks Modernize

Trends and Case Studies

Patient trust has become a systems problem that spans governance, security, identity, and the way consent is recorded and enforced across AI workflows. Patients are already using AI tools to interpret health information. At the same time, regulators are increasingly treating transparency and traceability as requirements. In the U.S., **ONC's HTI-1 Final Rule** elevates algorithmic transparency expectations for predictive decision support in certified health IT.

One-time, broad consent forms were built for a simpler world, where data flowed along predictable rails and clinical decisions were easier to explain. **In 2026, data is reused for new models, models evolve, and AI can act across workflows.**

Consent is moving toward granular, machine-readable directives that can be enforced automatically, updated over time, and tied to purpose and context.

Besides, two forces are driving explainability to shift from **"nice to have"** to **"required to operate."**

First, patient and public expectations. Second, the compliance environment that is increasingly demanding transparency around predictive tools used in care and operations. Think data provenance, intended use, known limitations, and governance controls that are understandable to clinicians, patients, and auditors.

Key Barriers and Solutions

The core constraints are about making it safe, governable, and acceptable in real care environments.

a) Consent is static, but AI use is continuous and context-specific

Traditional **consent assumes** a stable relationship between data and purpose. **AI breaks that assumption in several ways.** Models evolve, vendors change, and the same data can be repurposed for new workflows that were not part of the original patient conversation.

Even when a health system has strong intent, the operational question becomes hard quickly. What exactly did the patient agree to when a new model is introduced, a dataset is re-linked, or an agent begins acting across multiple tools?

This is also where trust quietly fails. Patients rarely object to innovation in the abstract.

> ***"Consent is no longer a form – it's an enforceable part of the data pipeline, tied to purpose, time, and context."***

Move toward granular, machine-enforceable consent that is tied to purpose, time, and context. Start with the highest-risk domains, like sensitive data categories or workflows where AI meaningfully influences care. A practical operating step is to make consent enforcement part of the data pipeline

b) Transparency expectations rise, but most transparency is not usable

There is a growing mismatch between what organizations publish and what people need. Patients want clarity. Clinicians want confidence. Boards and regulators want traceability. If transparency is delivered as a single artifact, it tends to satisfy no one.

This also shows up in sentiment. **Pew research** found **60% of Americans would be uncomfortable with a provider relying on AI** in their own care, and **75% believe providers will move too fast using AI before fully understanding risks.** Those numbers do not mean patients are anti-AI. They mean trust is conditional on visibility, oversight, and clear accountability.

Organizations should adopt layered transparency so each audience gets what they need, at the level they can use.

- **Patient layer** that clearly discloses AI involvement, human oversight, and what the output means in plain language.
- **Clinician layer** that communicates intended use, limitations, and the practical basis for confidence.
- **Audit layer** that documents provenance, evaluation results, monitoring signals, and governance decisions.

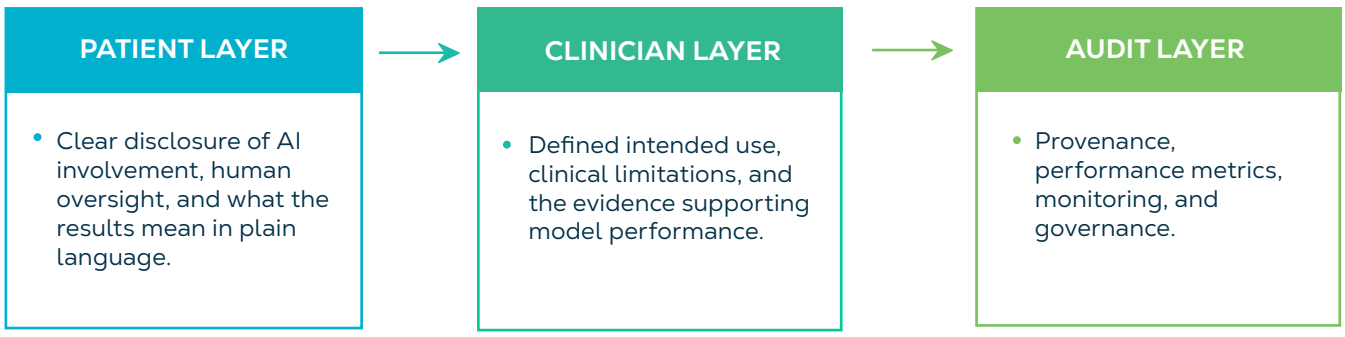

c) Agentic and LLM systems expand the attack surface and the error surface

Agentic workflows introduce a new trust problem. The system may be reading, deciding, and acting across workflows. That makes it powerful, but also susceptible to manipulation, especially through untrusted inputs.

Once untrusted data is near an agent, outputs should be treated as untrusted until screened.

It also introduced the concept of trust boundaries as the design point where defenses and validation logic should be inserted.

Constrain autonomy and add verification where it matters:

- Apply least-privilege tool access for agents and require human approval for high-impact actions.
- Implement output screening and allowlists so the system cannot drift into unexpected behavior in critical workflows.
- Treat AI like insider risk in the operating model, with continuous evaluation rather than “one and done” testing.
- Monitor for supply chain exposure, since model and plugin dependencies are part of the trust story now.

“Agentic AI systems don’t just assist – they read, decide, and act. Constrain autonomy, screen every output, and treat AI like insider risk: continuous evaluation, not a one-time test.”

4 Implementation Roadmap

How to Prepare for 2026

A PRACTICAL BLUEPRINT FOR HEALTH SYSTEMS, PAYERS, AND INNOVATORS

Most organizations will enter 2026 with the same tension. T**here is real pressure to adopt AI, and there is real fear of doing it wrong.** The gap between the two is where value is either created or lost. The healthiest path forward is to build a repeatable operating system that can scale AI safely, show impact in everyday workflows, and hold up under scrutiny when something goes sideways.

The roadmap below is designed to turn the trends into an execution plan that is measurable, auditable, and realistic in clinical environments.

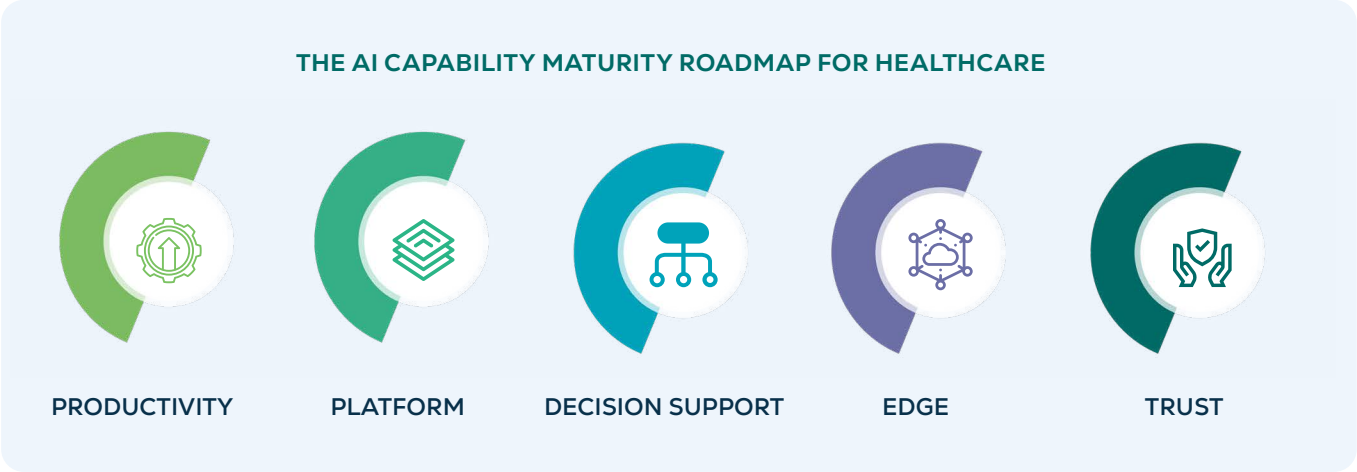

1. Treat productivity as the ROI that matters, then manage it like an operating metric

In 2026, the **strongest business case for AI is capacity.** Health systems need time back for clinicians and staff. The mistake is measuring ROI only through cost reduction, then wondering why adoption stalls. Productivity is easier to prove, easier to feel, and easier to connect to patient access and quality.

The practical move is to agree on a small set of workforce signals that the organization is willing to manage against, then make AI earn its place by moving those signals in the right direction.

When a tool saves time, decide where that time goes. If the time simply gets absorbed by more administrative work, the workforce does not experience relief, and trust evaporates fast.

2. Consolidate the platform layer before you scale the use cases

Most AI programs fail the same way. They accumulate tools faster than they accumulate reliability. The result is tool sprawl and inconsistent integration patterns that become impossible to enforce. Platform thinking flips the order. Instead of scaling use cases first, scale the rails first.

The organizations that get ahead in 2026 will **standardize how AI plugs into workflow,** h**ow data is accessed and traced, how models are deployed and updated, and how monitoring is done.** Once the rails are in place, new capabilities can be added consistently and at scale. For clinicians, maturity means outputs appear directly in workflow. For leadership, it means governance and security standards apply uniformly across all tools.

3. Treat multimodal decision support like an engine

Multimodal AI is compelling because it mirrors how clinicians actually think. They use history, labs, imaging, context, and pattern recognition together. The trap is trying to "buy multimodal" as if it were a single model. In practice, decision support works when there is an **orchestration layer** that **fuses signals** into one **clinician-ready output,** carries **uncertainty** honestly, and knows when to escalate rather than pretend. This is where smaller, **right-sized models and hybrid pipelines tend to win.** They are cheaper to run and more predictable in workflow.

The best implementations also avoid turning clinicians into auditors. The system should show its work in a way that is usable with clear reasoning, evidence links where appropriate, and limits that are explicit.

4. Make rare-case readiness a requirement

Healthcare is defined by **edge cases,** and that is where **trust is earned or lost.** If models are trained only on common patterns, they will look impressive until the first rare presentation, the first unusual combination of conditions, or the first look-alike case that requires nuance.

That is why **negative datasets** and **counterfactual testing** belong in the roadmap. The organizations that do this well treat it as a **continuous loop.** They stress-test before deployment, monitor for new failure modes after deployment, and continuously update the "what can go wrong" set as reality teaches them. This approach prevents the kind of **public failure** that slows everything.

5. Expand access with edge AI, but design the pathway

Edge AI will open doors in 2026, especially for screening, triage, and care that depends on timely capture. But edge deployments can also create a new kind of failure, where results are generated without a **reliable path to follow-through.**

The practical play is to start with a **measurable access gap,** standardize capture quality so results are consistent, and then hardwire routing and escalation so the next step happens reliably. Patient trust is part of this. If results are delivered closer to the patient, explanations have to work closer to the patient too. **Clarity, language support,** and expectation setting are part of **clinical safety** and adherence.

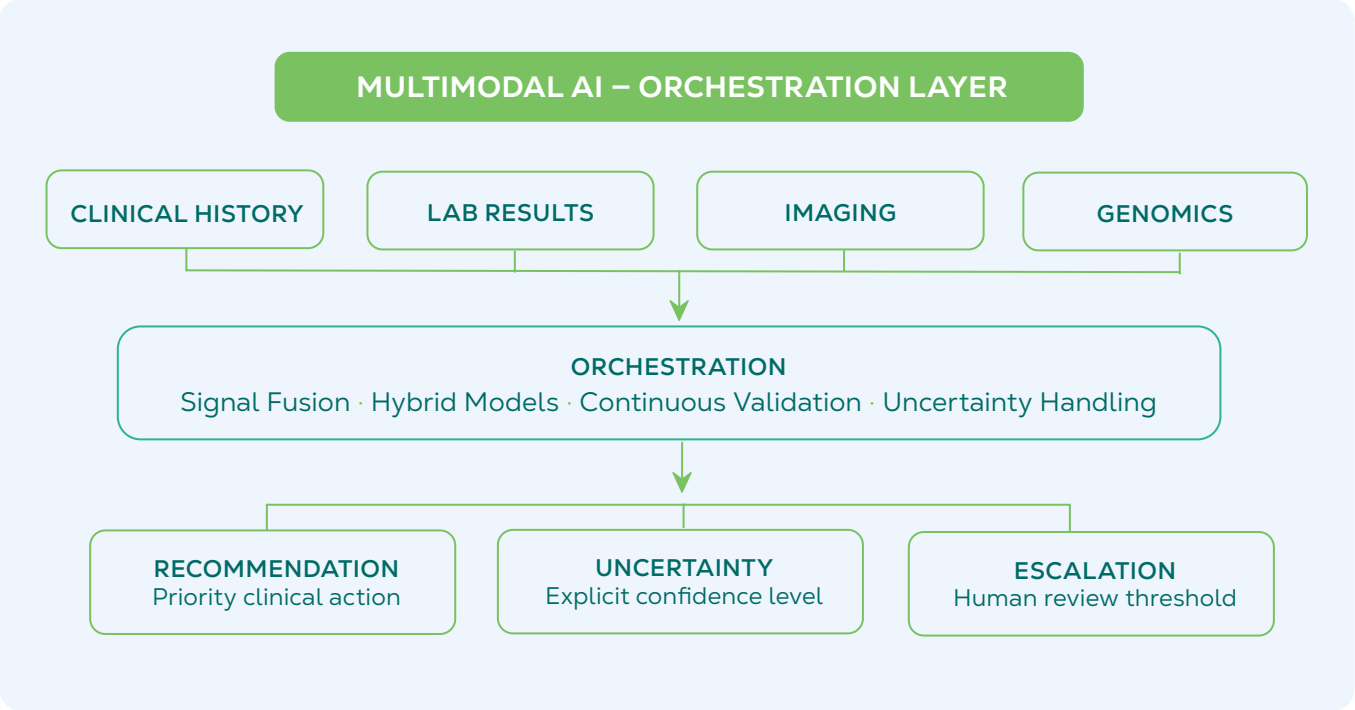

5 Conclusion

AI is moving into the operational core of healthcare. If leaders align AI investments with real workflow outcomes and build platforms that reduce fragmentation, the impact in 2026 can be meaningful and durable.

The trends in this white paper point to a clear direction.

- AI shifts from assisting to executing through agentic workflows.
- Productivity becomes the primary ROI as health systems fight capacity constraints and rising complexity.
- Clinical decision support becomes multimodal, combining text, imaging, genomics, and sensor signals, with rare-case readiness treated as safety work.
- Platform consolidation accelerates as organizations prioritize EHR-native deployment, unified governance, and registries that track what is deployed and how it performs.
- Edge AI extends diagnostic capability and access, but only when it is designed as an end-to-end care pathway.
- Trust and consent frameworks evolve because static models cannot keep up with continuously changing AI systems and data flows.

These shifts share one requirement. Health systems, payers, and innovators need to operate from shared standards for interoperability, governance, and accountability. **Organizations that build common rails and stay transparent with patients will scale faster** and with fewer setbacks.

The path to 2026 is disciplined execution across capacity, platforms, safety, and trust. Done well, AI becomes a practical force for better access, safer decisions, and a more sustainable healthcare workforce.

6 Industry Voices

Listed alphabetically

HAMED ABBASZADEGAN
Member Board of Directors
The Arizona AI Syndicate

DR. KATHRYN ALDERMAN
CEO & Founder
Intelligent Care Alliance

LIOR ESHEL ANSHER
Founder and CEO
TestDynamics

HASSAN BENCHEQROUN
CEO
Medical AI Academy/
Podcast Author AI Ready

NILESH BHANDARI
Founder and CEO
Advanced Health

CHRIS BRIGHAM
Brigham and Associates
Smart Medical History

HARVEY CASTRO
Chief AI Officer
Phantom Space

SHERRI DOUVILLE
CEO
Medigram

ROBIN FARMANFARMAIAN
Co-Founder
i-GENTIC AI

DOUG FRIDSMA
Chief Science Officer
Health Universe

APURV GUPTA
Vice President, Advisory
Services Premier, Inc.

ROSA HART
Owner
Rosa Hart Media
Consulting LLC

JASON HILL
Innovation Officer
Ochsner Health

FLORENCE HUDSON
Executive Director
Columbia University

KALI IHDE
Director, NM Ventures & Innovation
Northwestern Medicine

JUNAID KALIA
CEO / Founder
SaveLifeAI

ANTHONY LEE
Executive Communication Coach and Head
Heroic Voice Academy

DAVID LEINGANG
Director – Innovation – AI, *Data Science, and Analytics*
Ochsner Health

NICHOLAS LIN
Co-Founder & CEO
Guava

JENNIE LUNA
CEO / Founder
Nuvee

YVES LUSSIER
Chair, Biomedical Informatics & Professor of Medicine
University of Arizona

DR. JAMES MAISEL
CEO
ZyDoc

JOSE MACION
CEO
GenOp: Block-Chained Healthcare Internet

ORVIL MARTINEZ
CEO / Owner
Advance Medical Research Center, US Neurology Consulting

ARLEN MEYERS
CEO
Society of Physician Entrepreneurs

JODYN PLATT
Associate Professor
University of Michigan Medical School

RAGHUNATH SAPURAM
CEO
T-Lex USA

KIANOR SHAH
Co-Founder
Healthcare Intelligence Board

ANAHI SANTIAGO
CISO
ChristianaCare

JOSHUA TAMAYO-SARVER
VP of Innovation
Inflect Health and Vituity

STEVE WILSON
Chief AI and Product Officer
Author, Holding Time for
Steve Wilson

About OUTCOMES ROCKET

Outcomes Rocket is a healthcare-focused growth and marketing agency that turns industry complexity into measurable business results. Serving med-tech and health-tech companies, they operate through a proven three-phase framework, Discover, Define, Deliver, combining market research, go-to-market strategy, digital marketing, podcast production, and brand development to drive real ROI. With over 20 years of healthcare expertise and an average 8x return for clients, Outcomes Rocket bridges the gap between insight and execution for hundreds of healthcare brands worldwide.
Learn more at outcomesrocket.com.

About CENSINET®

Censinet is a purpose-built, AI-powered cybersecurity and risk management platform designed exclusively for healthcare. Censinet RiskOps, enables healthcare delivery organizations to manage third-party vendor risk, medical devices, supply chain, HIPAA compliance, and more through a collaborative network of over 50,000 vendors and products, helping teams stop chasing questionnaires and start proactively managing risk. Trusted by leading health systems like Intermountain Health, Cedars-Sinai, and Northwell Health, Censinet delivers speed, scale, and security where it matters most.
Learn more at censinet.com