FROM

TECHNICAL DEBT

TO

CLINICAL TRUST

A Healthcare CIO's Operational Playbook for Modernization



Zafar Chaudry,
MD, MS, MIS, MBA, CHCIO, CDH-E, FCHIME

Sponsored by
rackspace



Zafar Chaudry

MD, MS, MIS, MBA, CHCIO, CDH-E, FCHIME

Chief Digital Health Officer, Rackspace Technology

linkedin.com/in/drzafarchaudry

Executive Biography

Dr. Zafar Chaudry is a physician-executive who has spent nearly four decades at the intersection of clinical medicine, corporate leadership, and digital innovation. He has served as CEO, CIO, Chief Digital Officer, and Chief AI Officer, and has led major transformations at multi-billion-dollar health systems in the United States, the United Kingdom, and internationally.

Most recently, he was Senior Vice President and Chief Digital, AI, and Information Officer at Seattle Children's (2017–2026). He oversaw a budget of more than $250 million and a team of 485 people. During his tenure he cut operating costs by more than $90 million and brought the institution to the industry's top benchmarks, including HIMSS Stage 7 and Most Wired Level 9. Working with Google, he led the enterprise AI and data strategy. Data insights helped Seattle Children's become the first hospital in the world to deliver opioid-free surgery. He also launched the world's first Minecraft environment in a pediatric hospital to improve patient engagement, and he supported the spin out of healthcare startups AdaptX and Brainchild Bio.

His career includes a series of firsts:

- **Cambridge University Hospitals (UK):** As CIO and CISO of one of the world's leading academic medical centers, he led the $300 million Epic EMR program, the first Epic implementation in the UK. Under his leadership, Cambridge became the first UK hospital to reach HIMSS Stage 7 and the first hospital outside the US to be named Most Wired.
- **Gartner:** As a Research Director, he advised C-suite executives and national governments across the US, EMEA, and Asia-Pacific (including Australia and New Zealand), supporting more than $3.8 billion in transformation programs.
- **Premier Group:** As CEO of this $100 million diversified conglomerate, he held full P&L responsibility and ran international operations across several sectors.

Dr. Chaudry is currently on the Board of the College of Healthcare Information Management Executives (CHIME) and a CHIME Fellow. Becker's Hospital Review has repeatedly named him a "Chief Digital Officer to Know," a "CIO to Know," and one of its "Great Leaders in Healthcare," and he was the 2022 Seattle CIO of the Year.

He holds an MD, an MBA, and master's degrees in Information Systems and Healthcare Management, along with several executive certifications in digital health leadership.

Contents

Introduction:
Healthcare Modernization Needs a Different Playbook 2

Section 1:
The Healthcare IT Paradox . . . 8

Section 2:
Clinical Trust as Infrastructure 14

Section 3:
The Triage Phase. 21

Section 4:
Designing the Hybrid Health System 31

Section 5:
The Migration Matrix. 42

Section 6:
The New Operating Model . . . 52

Section 7:
The Action Plan 65

Section 8:
The Conclusion. 77

Section 9:
Appendix 81



Introduction

Healthcare Modernization Needs a Different Playbook



Healthcare is at a tipping point. Every board wants AI, seamless digital access and bulletproof cyber resilience. Yet most CIOs are being asked to build that future on fragile legacy foundations held together by institutional memory, temporary fixes and the daily effort of people who know how to keep aging systems alive.

Modernization isn't a cloud migration, an IT refresh or a race to buy shiny new AI tools. It's the deliberate work of building a technology foundation that clinicians can trust, executives can govern, CFOs can defend and security teams can protect. The goal isn't new technology for its own sake. It's to strengthen the foundation supporting clinical care today while giving the organization the resilience and flexibility to support what comes next.

Earlier in my career as a healthcare CIO, I lived the daily grind of carrying technical debt. I inherited aging infrastructure at a time when leadership often viewed IT as a cost center and a source of friction. Clinicians complained about slow systems. Finance challenged every major funding request. My team spent far too much time reacting to incidents, keeping fragile systems operating and explaining foundational problems in language people outside IT could understand.

This is what technical debt feels like from inside the CIO's office. You're not only managing old hardware and applications that should have been retired years ago. You're also carrying the constant pressure of knowing that the systems your clinicians depend on may not be strong enough to support the

critical work they perform. Too often, the environment is kept alive by heroes working overtime and relying on knowledge that exists in the heads of a few people.

I quickly realized that if I wanted to succeed, I had to stop reacting like a firefighter and start leading as an architect. I stopped talking to the executive team about servers and switches and began framing the technology foundation as a critical clinical asset. That shift changed how I understood modernization. It was no longer an infrastructure project but the work of building a foundation clinicians could depend on when caring for patients.

I came face-to-face with what that required during a major transformation I led at a pediatric health system in the Pacific Northwest. Working with Rackspace, we modernized a mission-critical clinical platform without asking clinicians to absorb the complexity of the change. That experience changed how I thought about infrastructure. The real measure of success wasn't whether we completed a technology project but how we strengthened the foundation supporting care while allowing physicians, nurses and other clinical teams to keep doing their jobs with confidence.

That transformation shaped the operating philosophy behind this playbook. I believe the complexity of modernization belongs with the technology, operational and partner teams and not with the clinicians at the bedside.

Trying to deliver the future of healthcare on infrastructure built for the past

Healthcare CIOs are being asked to protect today's clinical environment and deliver tomorrow's capabilities at the same time. Boards want AI, digital transformation and stronger cyber resilience. Clinicians want systems that respond every time. Finance wants cost control. And security leaders want fewer unknowns. Meanwhile, the technology teams are paying for and protecting years of accumulated systems that consume budget, capacity and attention.

For modernization to succeed, the CIO has to shape a new strategy. The CTO has to turn it into architecture, migration and operations. The CMIO needs to ensure the change won't create new friction for clinicians. And the CFO needs a cost model that's predictable and defensible. All of these perspectives are legitimate, but they collide when the foundation is too fragile, complex or poorly understood to support the future the organization wants to achieve.

That's the healthcare IT paradox: the organization wants a new future while still paying for and protecting the past. AI raises the stakes for modernization because a successful pilot doesn't prove that the organization can operate the capability safely when it goes into full production. Moving everything to the cloud doesn't solve the paradox if it simply relocates unrationalized technical debt to a new environment. And deferring modernization isn't a long-term option either when aging systems continue to increase clinical, financial, security and recovery risk.

Before the organization can move forward with confidence, the CIO has to make these consequences visible so modernization becomes a shared leadership decision.

Modernization must become executable

Healthcare organizations don't need another static strategy deck that stops at recommendations. A modernization plan only has value if the organization can perform the discovery, stabilize immediate risks, design the target state, manage migration and operate the resulting services with clear accountability.

This is where the operational co-pilot model can be useful. The right partner does more than advise the CIO and leadership team. An operational co-pilot brings engineering execution, migration discipline, shared accountability and ongoing management that reduce the burden on an already overstretched internal team. You retain ownership of your clinical priorities, risk tolerance and final decisions, but you take advantage of outside expertise so you don't absorb every operational complexity by yourself.

One more important point to make about modernization. The CIO's role continues to evolve after the migration is complete. Modernization requires the CIO to become the steward of the new environment by keeping the current state visible, governing change, proving resilience, managing partners and protecting clinical trust as new demands emerge.

The six-part modernization journey

This playbook follows a six-part modernization journey. Through this playbook you'll learn to:

1. Make the healthcare IT paradox visible by connecting future demand to the technical debt holding the organization back.
2. Define clinical trust as the standard for reliability, resilience and care continuity.
3. Triage and rationalize the current state so your organization knows what it has, what should remain and what must be stabilized.
4. Design the hybrid future around workload reality rather than platform preference.
5. Move the big rocks safely through disciplined migration waves that protect clinical care and operational workflows.
6. Operate the new environment with discipline by continuously governing cost, reliability, resilience, partner performance, value and AI readiness.

We've also added Section 7, which helps you execute your six-step roadmap with a ten-dimension readiness model and a practical guide to the 18 tools in the companion Healthcare Modernization Toolkit.

The playbook and toolkit at a glance

The Healthcare Modernization Toolkit turns the six-part journey and six-step roadmap into practical working tools. The crosswalk below shows where each tool supports the modernization roadmap, the readiness dimensions it addresses and the leadership roles most likely to participate. Each tool produces the assessment, decision, plan or governance output reflected in its title and detailed in the toolkit.

Step 1: Make the Paradox Visible

Tool	Primary readiness dimensions	Primary participants
1. Modernization Readiness Assessment	All readiness dimensions	CIO, CTO, clinical/ informatics, security, finance, operations
2. Technical Debt Reality Check	Technical debt visibility; Clinical trust and reliability	CIO, CTO, clinical/ operational owner, finance, security
3. Modernization Priority Action Plan	Selected dimensions based on identified gaps	CIO, CTO, executive sponsor, accountable owners

Step 2: Define Clinical Trust as the Standard

Tool	Primary readiness dimensions	Primary participants
4. Clinical Impact Assessment	Clinical trust and reliability; Cyber resilience and recoverability	Clinical/operational owner, application owner, technology, security/ recovery

Step 3: Triage and Rationalize the Current State

Tool	Primary readiness dimensions	Primary participants
5. Current-State Inventory Checklist	Technical debt visibility; Application rationalization; Partner execution readiness	Technology/application owners, architecture, security, operations, finance
6. Clinical Criticality Matrix	Clinical trust and reliability; Cyber resilience and recoverability	Clinical/operational owner, application owner, technology, security/ recovery
7. Legacy Retirement, Archival and Interoperability Worksheet	Application rationalization; Hybrid architecture readiness; Governance and operating model	Application/business owner, records/data, integration, legal/ procurement, finance

Step 4: Design the Hybrid Future

Tool	Primary readiness dimensions	Primary participants
8. Workload Placement Matrix	Hybrid architecture readiness; Clinical trust and reliability; Cost predictability	CIO/CTO delegate, architecture, clinical/ application owner, security, data, finance
9. IoMT and Zero Trust Security Matrix	Cyber resilience and recoverability; Hybrid architecture readiness	Security, architecture, network, biomedical/ clinical engineering
10. Service Accountability and Recovery Commitments Worksheet	Governance and operating model; Cyber resilience and recoverability; Partner execution readiness	Service/clinical owner, technology, security/ recovery, legal/ procurement, co-pilot
11. Vendor Dependency and Risk Governance Matrix	Partner execution readiness; Governance and operating model; Cyber resilience and recoverability	Vendor management, service owner, technology, security, legal/procurement
12. TCO and Cost Predictability Calculator	Cost predictability; Hybrid architecture readiness	Finance, CIO/CTO delegate, architecture, workload owner, procurement

Step 5: Move the Big Rocks Safely

Tool	Primary readiness dimensions	Primary participants
13. Migration Wave Planner	Migration readiness; Clinical trust and reliability; Partner execution readiness	Migration/application leads, clinical owner, architecture, security/ recovery, co-pilot
14. Safe-Exit and Rollback Checklist	Migration readiness; Cyber resilience and recoverability; Clinical trust and reliability	Migration lead, application owner, platform, security/ recovery, clinical owner
15. Clinical Change and Adoption Readiness Tool	Migration readiness; Clinical trust and reliability; Governance and operating model	Clinical/operational owner, informatics/ nursing, training/change, support, migration

Step 6: Operate with Discipline

Tool	Primary readiness dimensions	Primary participants
16. Quarterly Governance Review Template	Governance and operating model; Clinical trust and reliability; Cost predictability; Partner execution readiness	CIO/CTO, service owners, clinical/ operations, security, finance, vendor management
17. Value Realization Scorecard	Cost predictability; Governance and operating model; Clinical trust and reliability	CIO/CTO, finance, workload owners, clinical/operations, executive sponsor
18. AI Infrastructure Readiness Checklist	AI readiness; Governance and operating model; Cyber resilience and recoverability; Cost predictability	CIO/CTO, AI/data, clinical owner, security/ privacy, architecture, finance

Use the toolkit selectively: begin with the tools associated with the most consequential readiness gaps, then carry verified evidence and decisions forward as the modernization journey progresses.

How to use this playbook

This playbook is designed for healthcare CIOs, CTOs and other leaders who need to move from modernization intention to execution. You can read it from beginning to end as a leadership journey or use individual sections to support discussions with the clinical, financial, security, legal, operational and technical stakeholders on your team.

Section 1 helps you explain what is currently holding the organization back from modernization. Section 2 helps you define reliability through the lens of clinical trust. Section 3 establishes what must be known, rationalized and stabilized in the current environment in order to move forward. Section 4 helps you determine where each workload belongs in the new environment. Section 5 explains how to move workloads safely. Section 6 establishes how the new environment should be governed and operated. Finally, Section 7 converts the journey into an actionable roadmap.

The first step in our journey is to confront the healthcare IT paradox directly. Healthcare CIOs are being asked to innovate faster, strengthen resilience, control cost and prepare for AI while supporting clinical care on infrastructure that may be too fragile, complex or expensive to carry the organization forward. That's where we begin.



Section 1: The Healthcare IT Paradox

THE CONFLICT: Legacy Debt vs. the AI Mandate



From the CIO seat, the conflict is brutal. The board demands innovation, clinicians expect systems to respond every time, and the CFO wants cost reduction. Meanwhile, the technology team is supporting years of accumulated systems that consume budget, capacity and attention. Without a disciplined plan, those demands collide.

Healthcare leaders know they can't postpone AI, digital access, cyber resilience or new models of care. Those priorities are already part of the modernization agenda. But the current foundation may still be filled with unsupported systems, fragile integrations, rising costs, incomplete documentation and applications that keep engineering teams focused on maintenance instead of progress. The organization wants a new future while still paying for and protecting the past.

Before modernization can move forward, the CIO has to make this paradox visible to the rest of leadership. That means translating technical debt into clinical, financial, operational and security consequences that leaders outside IT can see and act on. Until that happens, modernization will continue to look like another infrastructure request instead of an organizational priority.

Healthcare is being asked to deliver two futures at once

Healthcare CIOs and CTOs are being asked to protect today's clinical environment while delivering tomorrow's capabilities. The CIO has to shape the strategy without putting the EHR, clinical systems or operational platforms at risk. The CTO has to turn that strategy into architecture, sequencing, migration and day-two operations. The organization expects progress, but it won't accept disruption to care.

I remember how this paradox felt before I had a clear way to explain it. I was always in the hot seat. The board wanted greater speed and innovation. Clinicians questioned whether the systems would be available when they needed them. Finance wanted to know why infrastructure required so much investment. At the same time, my technology team was spending most of its energy keeping aging systems on life support. It was a triple bind: move faster, spend carefully and never disrupt care.

The conversation changed when I stopped describing the problem as old technology. I began translating technical debt into consequences every leader could understand: clinician time lost to slow systems, innovation capacity consumed by maintenance, security exposure created by unsupported components and recovery risk hidden inside undocumented dependencies. Once leaders could see what the debt was costing the organization, modernization stopped being an IT issue and became a shared leadership responsibility.

AI exposes weaknesses in the foundation

AI has changed the questions CIOs are expected to answer. Leadership no longer asks whether the organization will use AI. It wants to know which use cases will move first, how quickly they can scale and when the benefits will appear. The CIO has to address these demands without allowing executive enthusiasm to outrun the organization's ability to operate the new technology safely.

That requires a clear production threshold. Before an AI capability moves beyond a controlled pilot, the organization should know who owns the use case, which data sources it relies on, how it fits into the clinical or operational workflow and who's accountable when performance changes. Leadership must also understand the capacity, controls, monitoring and support required when a small experiment becomes an enterprise service.

It's risky to move AI into production before those responsibilities are settled. A promising pilot can earn executive attention quickly, but an unreliable production experience can damage confidence just as fast. The CIO's role is to establish the conditions that allow the organization to expand AI without creating new operational risk or asking clinicians to compensate for weaknesses that should have been resolved before deployment.

Technical debt is organizational risk

The accumulated technical debt rarely announces itself as a single crisis. It builds through reasonable decisions made under pressure: a refresh delayed for another priority, an application retained because no one owns the retirement decision or a workaround accepted because changing the workflow feels too disruptive. The individual choices were defensible at the time, but the accumulated result is an environment that costs more, changes more slowly and carries risks that few leaders outside IT can see clearly.

The CIO has to translate the risk of that hidden burden to every part of the leadership team:

- Finance leaders need to see the spending tied to duplicate systems, recurring fixes, aging contracts and support that no longer creates new value.
- Clinical and operational leaders need to see where slow, unstable or difficult-to-recover systems interrupt workflows and force staff to compensate.
- Security leaders need visibility into unsupported components, unclear dependencies and systems that can't be protected or restored with confidence.
- Executive leaders need to understand how legacy complexity delays strategic priorities and reduces the organization's capacity to respond to change.
- Technology leaders need a shared view of which risks require immediate stabilization and which can be addressed through the longer modernization sequence.

Rather than asking leadership to fund the replacement of an aging component in isolation, they need to see the service it supports, the consequence of failure, the people affected and the choices the organization loses by leaving it in place. The same technical issue matters for different reasons to the CFO, CISO, CMIO and the board.

Once the consequences are visible, technical debt can be prioritized rather than debated in abstract discussions. Together, leadership can decide which exposures are acceptable for a defined period, which require immediate action and which investments will return capacity to the organization. That's how the CIO turns a collection of infrastructure problems into a shared modernization agenda.

Paying interest on the past

It's called technical debt because it carries compounding interest. The organization pays this interest through emergency outages, ticket spikes, recurring fixes, staff fatigue, support exceptions, aging contracts and delayed innovation. It also pays through duplicate licenses, overlapping systems, unpredictable recovery exposure and the skilled labor consumed by keeping fragile environments alive.

Every year modernization is deferred, the work becomes harder and more expensive to achieve. Documentation disappears, institutional knowledge leaves and dependencies become more difficult to unwind. When the organization finally acts, it pays for the new environment while also untangling years of accumulated complexity.

The true cost of technical debt isn't what you see in the budget; it's what your organization can't do because too much time, funding and talent remain locked into preserving the past.

Cloud-first thinking doesn't solve the healthcare paradox

Let's set the record straight: cloud-first is a lazy strategy when it means moving unrationalized technical debt into someone else's data center. Lift-and-shift without rationalization simply gives you more expensive legacy problems. Moving an overly complex environment to a new platform doesn't automatically make it simpler, safer, less expensive or more resilient.

Public cloud can be an important part of modernization. It provides elasticity, rapid provisioning and access to advanced services. But those advantages matter only when they solve a real workload requirement. They don't compensate for applications that should have been retired, poorly understood dependencies or systems that were never designed to operate efficiently in a variable-cost environment.

Before choosing a platform, decide where each workload belongs so it supports care, performance, security, recoverability and financial discipline. Section 3 establishes that current-state discipline. Once the organization knows what should remain, Section 4 addresses where each retained workload belongs.



Doing nothing feels safer than it is

Deferring modernization can feel like the right choice. Clinical systems can't go down, budgets are constrained and internal teams are overloaded. But inaction has its consequences. Aging systems become harder to maintain and recover, knowledge becomes concentrated in fewer people and IT spends more time keeping the old environment alive instead of helping the organization move forward.

The risk of doing nothing often remains hidden until a system slows during a critical workflow, a cyber event exposes recovery gaps or a cost review reveals how much capacity is consumed by legacy support. By then, the organization may no longer be modernizing from a position of control. It's responding to a problem that chose its own time to surface. The point is to act while the organization still has choices. Doing nothing can feel safe until it's not.

The first move shouldn't be vendor selection

Once leaders accept the need to modernize, there's a natural impulse to move quickly toward a platform or partner. The right operational co-pilot can accelerate progress, provide expertise the internal team lacks and reduce the burden of transformation. But partner selection should follow a clear understanding of the environment and not be a substitute for it.

Before choosing a partner, determine what systems are operating, what they support, who owns them, what dependencies exist and what must be stabilized before anything moves. Only then can you evaluate whether a prospective operational co-pilot provides the discovery, engineering, migration and ongoing management the program actually requires. Section 3 covers the inventory, classification and stabilization work that creates that foundation.

Once the paradox is visible and leadership understands the cost of carrying the technical debt, the organization is ready to define the standard that should govern every decision that follows: what will ensure clinical trust.



CIO Summary — Section 1

The Healthcare IT Paradox: Legacy Debt vs. the AI Mandate

What this means for the CIO

You must make the paradox visible before the organization can act. Translate technical debt into clinical, financial, operational, security and board-level consequences, especially as AI and digital transformation increase pressure on a foundation that may already be fragile.

Questions to ask now

- Where is technical debt limiting our ability to support AI, digital transformation, cyber resilience and clinical operations?
- Which risks are visible only inside IT today but could become executive, financial or clinical risks later?
- How much time, funding and skilled capacity are tied up in preserving the past?
- Are we making infrastructure decisions from current-state evidence or reacting to pressure for cloud, AI or cost reduction without a clear foundation?

Helpful tools from the Healthcare Modernization Toolkit:

Use the **Modernization Readiness Assessment** to establish a fact-based view of where the organization is ready to proceed and where evidence, ownership or dependencies remain weak.

Use the **Technical Debt Reality Check** to translate the most consequential technical debt into clinical, operational, financial, security and workforce consequences that leaders outside IT can understand and act on.

Then use the **Modernization Priority Action Plan** to convert the most important readiness gaps into three to five accountable priorities with owners, next actions and evidence of completion.

If an AI use case is already being considered or piloted, use the **AI Infrastructure Readiness Checklist** early to identify production-readiness gaps before executive enthusiasm outruns the organization's ability to support the service safely.

Section 2: Clinical Trust as Infrastructure

THE PHILOSOPHY: Reliability Is a Clinical Metric

Healthcare technology is often discussed in terms of platforms, applications, networks, storage and recovery. But that isn't how clinicians experience the organization's infrastructure. Clinicians experience technology through the small moments that shape care: a chart opens without delay, an order moves through the workflow, an image appears when it's needed and the application supporting a decision behaves as expected.

The best healthcare IT infrastructure is completely invisible. When it works well, clinicians barely notice it. That's exactly how it should be. When it slows down, becomes unstable or disappears, the interruption may begin in the data center or network, but the impact is felt in the exam room, emergency department, operating room, inpatient unit or pharmacy workflow.

That's why infrastructure reliability has to be treated as part of clinical trust. The issue isn't about whether a platform appears healthy on a dashboard. It's about whether clinicians can depend on the complete service while care is being delivered.

Infrastructure is invisible until it fails

A physician signs in and the patient record is available. A nurse enters medication information and the order moves through the workflow. A radiologist opens an image without delay. Step by step, patient after patient, the infrastructure preserves the rhythm of care without demanding attention.

That rhythm changes quickly when a system slows down or stops responding. What appears to IT as a performance issue can feel very different to a clinician making a decision under pressure. A delayed chart creates uncertainty, a lagging application interrupts concentration and an unavailable system forces the care team to decide what information it trusts and how to continue care without the normal tools.

I've walked emergency department floors and watched clinicians write patient vitals on paper napkins because the EHR was lagging. IT could have logged that as a minor latency issue but clinical operations experienced it as a failure of system support. The system pushed part of the burden back onto the care team at the exact moment when they needed it to work.

Clinicians are remarkably adaptive. They'll improvise whenever it's necessary to keep patient care moving. But don't mistake clinician adaptability for operational health. Every workaround is an early warning indicator that the infrastructure is no longer supporting the workflow as intended.

Reliability should be a clinical metric

Healthcare organizations track availability, response times, incident counts and recovery objectives. Those measures are clearly important, but they're incomplete unless they show what happens at the point of care. A server can be available while a clinician waits for an image. An application can pass a technical check while an order stalls in the workflow. And a recovery target can be met while the complete clinical service remains unusable.

Clinical Trust Metrics connect technical performance to uninterrupted patient care. They help leaders measure whether clinicians can obtain the information they need, complete the workflow, continue operating during disruption and trust the restored environment when the service returns.

To ensure clinical care is reflected in the way performance is measured, here are the six Clinical Trust Metrics:



Bedside Application Latency:

How long clinicians wait for critical applications, records, images or orders during actual care delivery.



Clinical Workflow Recovery Time:

The time required to restore the complete workflow, including identity, interfaces, data and downstream services, not simply the primary application.



Time-to-Clinical-Context:

How long it takes a clinician to obtain the information needed to understand the patient and make the next decision.



Workaround Frequency:

The number and type of informal processes clinicians create because they don't trust the normal system to perform reliably.



Zero-Interruption Care Hours:

The amount of time critical workflows operate without system-related interruption, forced workaround or degraded performance.



Recovery Confidence:

Whether testing has demonstrated that critical services, dependencies and data are restored safely and used by clinicians.

Later sections apply these metrics selectively according to the decision at hand. These measures aren't intended to replace technical metrics but to translate them into usable information. Uptime matters because care continues whether the system is available or not. Latency matters because delays interrupt attention and extend the time required to complete care. Recoverability matters because the organization has to know how care will continue and when the complete service will return safely.

Earlier in my career, I learned that talking to executives about servers, storage or warranty status rarely changed the conversation. What gained their attention was explaining the consequences in clinician time lost, workflows delayed, patient-safety exposure and capacity that could no longer be used to improve care. Clinical Trust Metrics give clinical, operational, financial and technology leaders a common language for those discussions.



Downtime changes clinician behavior

We've all experienced the immediate effects of downtime. Systems become unavailable, calls flood the help desk and departments activate continuity procedures. The service is eventually restored and most of the organization resumes normal work. But the longer-term effect on clinicians is harder to see.

Once clinicians believe a system won't be there when they need it, they plan around that uncertainty. They print information in advance, maintain paper alternatives, duplicate documentation or create informal processes to protect the next patient from the next interruption. Those responses introduce inconsistent information, documentation gaps, extra work and processes outside normal controls.

Infrastructure may be restored in an hour, but clinician trust takes much longer to recover. Clinicians remember the moments when the system forced them to improvise during care. Workarounds should therefore be measured against Clinical Trust Metrics and investigated. Only then will they reveal where availability, performance or recovery plans aren't reflecting the reality of the clinical environment.

Performance at the point of care matters

Downtime receives the most attention, but performance degradation is just as damaging to clinical confidence. A slow login, delayed chart, lagging image or repeated pause may not meet the threshold for a major incident, yet those delays accumulate across hundreds or thousands of interactions.

Clinicians don't experience storage, networks, databases and application layers separately. They experience the entire chain as one system. A hospital-wide response-time average may look acceptable while a surgeon waits for images during peak operating hours or a nurse repeatedly reauthenticates during medication administration.

This is where Bedside Application Latency and Time-to-Clinical-Context are useful. They direct attention to the moments most sensitive to delay and help clinical leaders describe where technology is interrupting care. The goal is to measure performance where care is delivered, not only where infrastructure is operated.

Technology friction also adds to clinician burden. It's not the only cause of burnout, but repeated delays, duplicate documentation and unreliable workflows increase cognitive load and consume time that should be spent with patients. Infrastructure decisions either reduce that burden or add to it.

Cyber resilience is part of clinical trust

A cyber event can make normal recovery assumptions unreliable. Production systems may be isolated or under investigation, and a traditional secondary environment may not be considered safe to use. The organization may need clean, isolated recovery capabilities that allow essential clinical services to operate while the broader environment is assessed.

Disaster recovery asks whether the service can be restored. Cyber recovery also asks whether the restored systems and data can be trusted. Clinical leaders don't need every technical detail, but they do need confidence that critical systems have been identified, backup data is usable, recovery priorities reflect care impact and clinicians can operate in a degraded state while restoration is underway.

That means a recovery plan that's on paper isn't enough. The organization has to prove that the complete clinical workflow can return safely. How long can care continue in a degraded state? Which dependencies must return with the primary application? Can clinicians trust the restored data? Who confirms that the workflow is ready to resume? Recovery Confidence and Clinical Workflow Recovery Time metrics make those questions visible.

Section 4 shows you how recoverability should be designed into the architecture, and Section 6 addresses how it must be tested and governed continuously. The essential point is that resilience is part of the organization's promise to continue care when the environment is under stress.

Service levels must be defined in clinical terms

Not every healthcare application carries the same level of risk. A system supporting emergency care, medication administration, imaging, surgery or inpatient documentation requires a different availability and recovery commitment than a lower-priority administrative application.

A traditional SLA may promise 99.9 percent availability or a four-hour recovery time. However, that isn't enough information for a clinical leader. The commitment must explain which workflow may be unavailable, how clinicians will continue working, what information may be delayed or lost, what must be reconciled and when the complete workflow will be safe to resume.

The same principle applies to infrastructure and vendor decisions. Storage commitments should be tied to image-retrieval performance. Network resilience should be tied to continuous telemetry and device connectivity. Disaster recovery objectives should reflect the maximum interruption the clinical workflow can tolerate. And vendor contracts should make clear who operates the service, who escalates, who communicates with clinical leaders and who has authority to invoke recovery.

I've seen governance situations become ineffective when too many people were involved, yet no one had authority to make decisions. If you want clinician confidence, then accountability matters as much as the metric itself. Sections 4 and 6 apply that responsibility model to architecture, contracts and day-two clinical operations.



What clinical trust looks like during modernization

I saw this clinical trust standard tested during a major modernization project at a pediatric health system. Working with Rackspace, we moved a mission-critical clinical platform while protecting the workflows clinicians depended on.

The most meaningful feedback came after the transition. Clinicians asked me whether any downtime had actually occurred. The new environment was responsive, patient information remained accessible and the operational complexity of the change stayed largely invisible to the people delivering care.

That experience was reinforced by a reported +48 end-user IT Net Promoter Score, providing a broader indicator of confidence in the technology organization. While no single measure can define clinical trust, the combination of dependable performance, minimal disruption and positive user sentiment showed what becomes possible when the technology environment supports people instead of asking them to compensate for it.

That's what clinical trust looks like in practice. It wasn't just that the infrastructure changed successfully but that physicians and nurses didn't have to compensate for that change. The technology, operational and partner teams carried the disruption risk instead of transferring it to the bedside. In practical terms, the transition protected Zero-Interruption Care Hours and reinforced confidence in the clinical environment.

That should be the standard for any healthcare modernization effort: clinicians experience a more dependable service, not the complexity required to create it.

Clinical trust should be the standard for modernization

Healthcare organizations modernize to replace aging infrastructure, improve security, reduce data-center burden, create more predictable operations and prepare for digital and AI capabilities. Those are legitimate goals. But the higher standard is whether the systems that support the care environment become more dependable.

Clinicians should trust that the systems they use will be available and responsive when care requires them. Clinical Trust Metrics give the organization evidence for those judgments.

That's when infrastructure stops being viewed only as an operating expense and becomes part of the care environment. Are clinicians experiencing less uncertainty? Are workarounds declining? Can care continue during disruption? Are recovery priorities aligned with what the workflow actually requires? Can the organization prove that the complete service will return safely?

Once clinical trust becomes the new standard, it must be protected. To accomplish that the organization must understand what systems exist, which workflows depend on them, where disruption would create the greatest risk for clinical care and what needs immediate stabilization. That's the work of triage in Section 3.



CIO Summary — Section 2

Clinical Trust as Infrastructure

What this means for the CIO

Infrastructure reliability is a clinical concern. Availability, latency, recoverability and resilience shape workflow stability and the organization's ability to continue care during disruption. The CIO must connect technical evidence to the point-of-care experience and use it to guide investment, architecture and operating decisions.

Questions to ask now

- Which clinical workflows are most vulnerable to latency, downtime or failed recovery?
- Are we measuring the complete workflow or only platform availability?
- Where have clinicians created workarounds because they don't trust the normal system?
- Do recovery commitments reflect the interruption care can tolerate?
- Can we prove critical workflows and their dependencies can be restored safely?

Helpful tools from the Healthcare Modernization Toolkit:

Use the **Clinical Impact Assessment** to connect availability, latency, downtime and recovery to the clinical workflows they affect, identify the Clinical Trust Metrics that matter most and establish the clinical evidence that should guide criticality, recovery and future service decisions.

Section 3: The Triage Phase

THE ACTION: Stabilize Before You Transform



By the time you reach the triage phase, the case for modernization should be clear. Technical debt is limiting your organization's ability to support care, manage risk and prepare for what comes next. Clinical trust is now the standard for deciding what must be understood, changed and protected.

Triage is where the principles established in Sections 1 and 2 meet the facts. It should produce visibility into the current environment, rationalization decisions about what should remain and stabilization of the risks that can't wait. Done correctly, triage creates more than an inventory. It produces a current-state decision package that helps you protect the systems that matter, remove workloads that no longer belong and design the future state from verified evidence rather than unsupported assumption.

The journey to modernization can feel overwhelming at first because the environment has grown complex over decades. The good news is you don't need perfect documentation before the work begins. Few health systems have it. However, you do need enough reliable information to make responsible decisions, expose the unknowns that could change the plan and avoid committing to a platform, partner or migration sequence before you understand what you're moving.

Triage must come before transformation

Modernization can appear to be moving forward as soon as leaders compare platforms, discuss cloud destinations or invite partners to propose solutions. Those activities become necessary later on. But, if done too soon, they create the illusion of progress while ownership, dependencies, recovery requirements and clinical impact remain unclear.



Before I'm comfortable committing to a platform or an execution partner, I want answers to five questions: What systems does the organization operate? What does each system support? Who's accountable for it? What does it depend on? What happens if it fails? If you can't explain what a workload supports, who owns it, what it depends on and how care continues when it fails, you aren't ready to move it.

Technical debt has a way of hiding surprises. A workload that looks straightforward may support an undocumented clinical process. An application that appears ready to retire may still feed another system through an old interface. A four-hour recovery target may look acceptable until a clinical leader explains what four hours without the service would mean. Triage exposes these realities before they become architecture flaws or migration emergencies.

A useful triage process should answer three leadership questions:

1. What's true about the current environment?
2. Where's the greatest clinical, operational, financial or security exposure?
3. What must be stabilized, retired, clarified or investigated before target-state decisions can be made?

Once those questions are answered, the CIO can present a credible risk narrative, the CTO can shape the architecture and the CFO can understand the investment logic. Until then, the organization may have urgency, but it doesn't yet have enough control over the environment to modernize responsibly.

Start with a current-state inventory you trust

Most healthcare technology environments weren't designed as a single system. They grew through infrastructure refreshes, departmental purchases, acquisitions, clinical projects, regulatory requirements and temporary fixes that became permanent. Some applications were retired cleanly. Others remained because ownership was unclear, dependencies were undocumented, historical data was still needed or changing them felt too disruptive.

The result is usually an environment that no one fully understands, even though many people understand parts of it. A current-state inventory makes that environment visible in a form leaders can use. At minimum, it should identify:

- Applications, infrastructure platforms and data stores
- Business, technical and clinical owners
- Interfaces, integrations, APIs and shared services
- Identity, access and authentication dependencies
- Third-party vendors, hosting providers and support dependencies
- Contracts, licensing, data ownership and exit constraints
- Lifecycle, warranty and vendor-support status
- Utilization, performance and incident history
- Backup, recovery and continuity requirements
- Data-retention, archival and interoperability obligations
- Medical-device and IoMT connectivity where relevant
- Known costs and major cost drivers
- Users, locations and workflows supported
- Architecture exceptions and reliance on individual institutional knowledge

The current-state inventory is useful only when it distinguishes verified facts from working assumptions and unresolved questions. That discipline prevents an assumption from becoming a design decision. An application believed to be inactive may still support a time-sensitive report. A system may appear to have an owner until an incident reveals that no one is accountable for recovery, funding or vendor escalation. And an application may look ready to retire until clinicians explain that they still need access to years of historical data.

I've seen an initial inventory create more questions than answers. That's actually good news. It's evidence that the organization is replacing a comfortable story with a more honest picture. The inventory doesn't have to be complete on day one, but it must become more trustworthy as the team validates what it finds.

Don't turn the inventory into a technical spreadsheet that only IT can interpret. Connect every system to the function it supports. Does it contribute directly to patient care? Does it support revenue cycle, scheduling, research, supply chain, patient access or administration? Who uses it, where is it used and what happens when it is slow or unavailable? Those answers turn an inventory into a decision package leaders can use.

Classify systems by clinical and operational criticality

Once the environment is visible, you need a disciplined way to decide what matters most. Healthcare organizations often describe too many systems as critical. Every department depends on technology, and no leader wants to suggest that an important workflow is expendable. But if everything is critical, the label stops guiding investment, protection, recovery or sequencing.

Criticality should be based on consequence, not on which department has the loudest voice. Instead of asking whether a system is important, ask what stops, slows down or becomes unsafe when the capability is unavailable. Consider whether latency interrupts care, whether downtime forces a workaround, how long the workflow can operate safely in a degraded state and which dependencies must return before the service is fully restored.

As you classify systems, use four broad categories:



Mission-critical clinical systems.

Care can't continue safely or effectively for long without them. These may include systems supporting emergency care, medication administration, imaging, surgery, inpatient documentation or other time-sensitive clinical workflows.



Operationally critical systems.

They may not support bedside care directly, but they're essential to patient access, revenue cycle, scheduling, staffing, communication, supply chain or other functions that keep the health system operating.



Important but recoverable systems.

An interruption creates inconvenience or reduced productivity, but the organization can continue safely for a longer period using defined alternatives.



Low-value, redundant or obsolete systems.

They may be underused, duplicated, unsupported or no longer aligned with organizational needs.

Each classification influences service levels, protection, recovery, support, cost and migration priority. A mission-critical clinical platform may require greater resilience and a shorter recovery window than a lower-priority administrative service. An operational platform may require a different protection model because of its financial or patient-access implications.

Don't make these decisions in isolation. Technology teams understand the architecture and dependencies. But clinical leaders understand care impact; operational leaders understand workflow consequences; and security teams understand vulnerability and recoverability. The classification becomes credible when those perspectives are reconciled into one shared definition that's carried into architecture and migration.



Find the hidden fragility

The most visible application isn't always where the greatest risk lives. A major clinical platform appears stable while an aging database, unsupported operating system, undocumented interface or out-of-warranty component sits beneath it. Triage looks below the front-end experience and identifies the conditions that make the complete service strong or fragile.

Common signs of hidden fragility include:

- Undocumented interfaces, APIs and dependencies
- Single points of failure
- Unsupported software or aging operating systems
- Infrastructure beyond warranty or lifecycle
- Manual workarounds that compensate for system limitations
- Applications with unclear ownership
- Environments dependent on one or two individuals
- Recurring incidents and emergency fixes
- Backups or recovery procedures that have never been successfully tested
- Vendor dependencies missing from the system record
- Medical devices connected through aging, shared or poorly segmented networks
- Temporary architecture exceptions that became permanent operating practices

These risks are easy to underestimate because the environment continues operating despite them. A system works every day until a component fails, a vendor misses an obligation or a key employee leaves. The event doesn't create the fragility. It exposes it.

Use incident and support data to show where the environment is already struggling. Repeated tickets, performance complaints, capacity problems, support escalations and emergency interventions often point to hidden weakness. Clinical and operational workarounds matter just as much. A department maintaining its own manual process is telling you it no longer trusts the normal service.

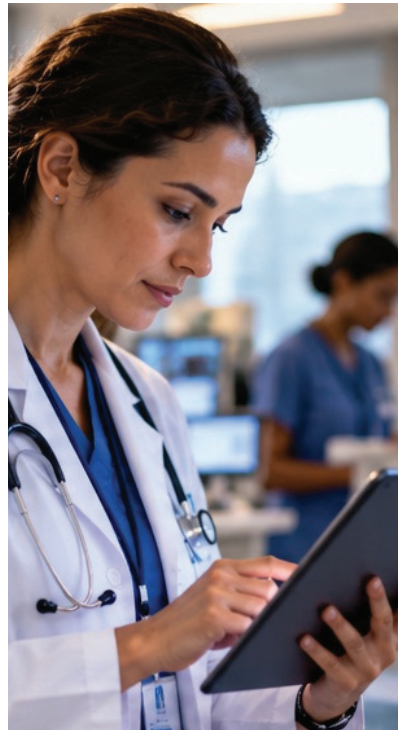
The purpose of triage is to find the weaknesses before the weaknesses find you.

Rationalize before you design the future

Don't assume that every system in the current environment belongs in the future one. Moving everything to the new environment preserves the complexity you're trying to remove. A redundant, rarely used or poorly owned application doesn't become more valuable because it runs on a newer platform. It simply carries the same cost, risk and support burden into a new home.

Rationalization helps you decide whether to:

- Retire it because it no longer provides enough value to justify its cost and risk.
- Replace it because the function remains necessary but the current application is no longer sustainable.
- Consolidate it because overlapping systems provide similar capabilities across departments.
- Retain it temporarily because clinical, contractual, operational or timing considerations make immediate change too risky.
- Replatform or modernize it because the capability remains valuable but its current technical form limits the future.
- Migrate it because it remains necessary and is suitable for the target environment.



Legacy systems rarely disappear on their own. A department chair resists retiring a niche application. A replacement exists while the old workflow never fully moved. Historical data may still be needed or an interface may be poorly documented. The organization keeps paying because retirement feels more disruptive than another year of maintenance. Without an accountable decision-maker, unnecessary applications remain indefinitely.

Treat "retain temporarily" as a controlled exception, not a permanent category. Every temporary retention decision should have an accountable owner, a clear reason, a review date, an exit condition and a plan for the data, interfaces, users and support that remain attached to the system.

Plan for archival and interoperability before retirement

Retiring an application doesn't automatically eliminate its cost or risk. The organization still has to decide what historical data must remain accessible, how long it must be retained, who needs access and how it will be retrieved. The data may need to be extracted, archived, converted or migrated. Read-only access may be sufficient in some cases; in others, clinical, legal or operational needs may require more.

The same discipline applies to interfaces and interoperability. Determine which data exchanges must continue, whether an API or archive can replace the legacy connection and when the original platform can be shut down. An application that remains online solely to provide historical access isn't retired. It's an expensive read-only footprint that still has to be funded, secured, recovered and governed.

The rationalization decision needs evidence about usage, ownership, data retention, workflow impact, interoperability, contractual obligations, supportability and total cost. Done correctly, rationalization is one of the first places modernization creates measurable value. Every system retired or consolidated reduces what the organization must fund, secure and govern.

I saw the financial impact of that discipline during the pediatric health-system transformation. As part of the broader modernization effort, we used aggressive asset triage to challenge which legacy applications still justified their cost and operating burden. That rationalization work contributed to a multimillion-dollar reduction in legacy application maintenance costs. The lesson wasn't simply that retirement saves money. Every unnecessary application removed also reduced the number of systems the organization had to support, secure, recover and govern.

Make vendor dependency visible

Healthcare services rarely depend on one internal team or one platform. Vendors host, operate, integrate, secure or support critical parts of the service. Triage identifies those dependencies before they become migration or recovery surprises.

For each critical service, determine which vendors are involved, what each party is responsible for and what happens if one of them fails to perform. Confirm whether the health system has access to documentation, credentials, data and recovery procedures. Understand the escalation path, contractual remedies and any dependency that creates risk across several applications.

Section 6 turns this information into an ongoing vendor-governance discipline while Section 3 makes the exposure visible.

Stabilize the highest-risk areas first

Triage sometimes uncovers risks that can't wait for migration.
A mission-critical application has recurring availability problems. An unsupported platform creates an immediate security exposure. A medical-device network lacks adequate segmentation. And a recovery process exists on paper but fails under testing.

Stabilization means reducing immediate exposure enough for the organization to proceed deliberately. It includes replacing failing infrastructure, strengthening backup coverage, documenting a critical dependency, improving monitoring, segmenting a vulnerable network, resolving an ownership gap or testing recovery.

Keep the distinctions clear: stabilization reduces immediate risk, rationalization determines what should remain and modernization creates the future state. A temporary stabilization fix mustn't become the next permanent exception.

These actions may not look transformational, but they create breathing room. I've seen teams so consumed by incidents that they had little capacity left for modernization. Every planning request competed with the immediate need to keep an unstable environment running. Stabilization reduces that pressure and restores enough capacity for the team to support the larger program.

Finally, treat stabilization as a decision gate. You're ready to proceed when the most immediate clinical, availability, security and recovery exposures are reduced to an acceptable level; ownership is clear; unresolved risks are documented; and the team has enough operating capacity to participate in design and execution. Stabilization doesn't eliminate every weakness but it prevents the most serious ones from undermining the next stage.

Bring the right stakeholders into validation

Triage begins with technical discovery, but the people who live with the consequences must validate what the team finds. This doesn't require a large standing committee but it does require the right expertise at the point of decision.

- Clinical leaders should confirm which workflows depend on a system and what happens when it becomes slow or unavailable.
- Operational leaders should identify consequences for patient access, scheduling, staffing, revenue cycle and supply chain.
- Finance should validate cost, investment constraints and the business implications of keeping or changing a service.
- Security should assess exposure, protection and recovery.
- Legal and procurement should identify contractual, licensing and accountability constraints.
- Records, data or compliance leaders should validate retention, archival and access obligations.
- Biomedical or clinical-engineering leaders should participate when medical devices and IoMT dependencies are involved.
- Vendor-management representatives should validate third-party responsibilities and escalation paths.

Any of these perspectives can change the plan. A technical team classifies an application as low priority until a clinical leader identifies a hidden dependency. A recovery target looks acceptable until operations explains the work that will accumulate during the interruption. A system appears ready to retire until records leaders identify a retention obligation. A technically attractive option becomes impractical because of licensing, data ownership or vendor constraints.

Early validation of these issues strengthens executive support because the modernization story is no longer owned by IT alone. The organization shares one understanding of what's critical, what's fragile and what action is necessary.

Turn triage into a decision package

Before the health system begins target-state design, triage should produce a concise current-state decision package containing:

- A verified current-state inventory
- Clinical and operational criticality classifications
- Rationalization decisions for each major workload
- Data-retention, archival and interoperability requirements
- Vendor and third-party dependencies
- Immediate stabilization priorities and completed actions
- Confirmed owners, dependencies and recovery requirements
- Unresolved questions requiring additional investigation
- The workload and service requirements that must shape the target architecture

This package gives the CIO a fact-based business and risk narrative and the CTO the evidence required to design the future state. It helps the CFO understand the investment logic and shows clinical, operational and security leaders how the organization intends to protect the care environment.

The decision package shouldn't become a static artifact because systems change, ownership shifts, interfaces are added, vendors change and temporary exceptions accumulate. The inventory and dependency map must be maintained as operating assets so the organization doesn't rebuild the same uncertainty it worked to remove. Section 6 establishes the governance rhythm for keeping that information current.

Triage is complete when the environment is visible enough to support defensible design decisions, rationalization choices are owned and the highest-risk exposures have passed the stabilization gate. The next step is to decide where each retained workload should live, how it should be protected and what operating commitment it requires. That evidence becomes the starting point for designing the hybrid health system in Section 4.



CIO Summary — Section 3

The Triage Phase: Stabilize Before You Transform

What this means for the CIO

Before approving a platform decision, migration program or operational co-pilot, you need a trustworthy baseline of what you have and what you intend to keep. Triage creates visibility, forces rationalization decisions and establishes the stabilization gate that must be passed before architecture and migration proceed.

Questions to ask now

- Do we have a current-state inventory we trust?
- Which systems are mission critical, operationally critical, recoverable, redundant or obsolete?
- Where are the undocumented dependencies, vendor exposures, ownership gaps and single points of failure?
- Which workloads should be retired, replaced, consolidated, retained, modernized or migrated?
- Have we defined what happens to historical data and interfaces when a system is retired?
- What must be stabilized before target-state design and migration can proceed safely?

Helpful tools from the Healthcare Modernization Toolkit:

Use the **Current-State Inventory Checklist** to establish a dependable baseline of workloads, owners, dependencies, support status, costs and recovery requirements.

Use the **Clinical Criticality Matrix** to classify services according to the consequences of disruption, and the **Legacy Retirement, Archival and Interoperability Worksheet** to make defensible decisions about what should be retired, replaced, consolidated, retained, modernized or migrated—and what must happen to data and interfaces before legacy systems can be shut down.

Section 4: Designing the Hybrid Health System

THE BLUEPRINT: Architecture for Reality, Not Hype



Triage establishes what's operating, what must be stabilized and what shouldn't follow the organization into the future. Now the work shifts from diagnosis to design. For every retained workload, you need a defensible answer to three questions: Where does it belong? How will the complete service be protected and recovered? Who will be accountable for operating it after the move?

This is where modernization becomes a workload-placement discipline. The goal isn't to favor public cloud, private cloud or the data center. The goal is to match each workload to the environment that best supports its clinical importance, performance requirements, data characteristics, security obligations, recovery needs and cost model.

Start with the constraints that actually matter. Is the workload latency-sensitive at the point of care? Does it sit close to large volumes of EHR, imaging or research data that would be expensive or slow to move? Does it need to remain available when wide-area connectivity fails? Does demand fluctuate enough to benefit from elastic capacity? Does it require specialized cloud services? Can its operating cost be governed predictably? What recovery commitment does the clinical workflow require?

The answers to these questions will lead different workloads to different destinations. A core clinical platform requires stable latency, dedicated capacity and continuous operational support. Imaging environments are heavily influenced by data gravity and retrieval performance. Research and AI workloads benefit from elastic or specialized compute. Administrative capabilities may fit well in SaaS. Some clinical locations require local survivability or edge capability when connectivity can't be assumed.

That's why the future state should be hybrid by design. Hybrid isn't a compromise or a temporary stop on the way to a single destination. It's a deliberate service model in which placement, protection and operating accountability are matched to the realities of each workload.

Hybrid isn't a compromise

Healthcare has moved through cycles where one platform is presented as the inevitable answer. Public cloud remains extremely valuable, but no single environment satisfies every clinical, operational and financial requirement. The architecture has to reflect the workload, not the platform of the moment.

Core clinical systems need stable latency, dedicated capacity, continuous availability and clear accountability around the clock. PACS and enterprise imaging environments carry large datasets and are highly sensitive to data movement and response time. Research teams need scalable compute for a defined period and then release it. Administrative capabilities may be well suited for SaaS. Clinical locations require local survivability or edge capability when the consequence of disconnection is too high.



A disciplined hybrid model matches each retained workload to the environment that supports it best. Some belong in private or managed environments. Some benefit from public-cloud elasticity and specialized services. Some are best delivered through SaaS. Some require local capability. Others should never reach the placement discussion because triage determined they should be retired, replaced or consolidated.

The objective isn't to take sides in the cloud debate. It's to answer a more useful question: what does this workload require to support care reliably, recover quickly and remain financially defensible?

Start with workload placement, not platform preference

The triage decision package should already identify the criticality, ownership, dependencies, recovery requirements and rationalization status of each workload. Use that evidence to answer the placement questions that follow:

- What clinical or operational function does the workload support?
- Who uses it, where is it used and what level of response does the workflow require?
- What data does the workload create, consume or exchange?
- How sensitive is the data and where is it permitted to reside?
- How much latency can the workflow tolerate at the point of care?
- Must the workload continue if wide-area connectivity, a cloud region or a central service is unavailable?
- How quickly must the complete clinical workflow be restored, and what data-loss tolerance is acceptable?
- Which interfaces, identity services, databases, imaging systems and shared platforms must remain close to it?
- Does the workload have high data gravity that makes repeated data movement slow, expensive or risky?
- Is demand stable, variable, seasonal or temporary?
- What are the full economics after licensing, support, storage growth, network traffic, resilience and data movement are included?
- Who will operate the workload around the clock and be accountable for its performance?
- What happens if the selected vendor, service or connection becomes unavailable?

Clinical Trust Metrics should be designed in at this point in the process, not reporting measures added later. Bedside Application Latency determines whether an application or dataset can tolerate remote placement. Time-to-Clinical-Context reveals whether EHR, imaging, identity and integration services need to remain close together. Clinical Workflow Recovery Time shapes redundancy and dependency-aware restoration. Recovery Confidence influences the protection model. These measures translate architecture choices into what clinicians will experience at the point of care.

However, don't treat this initial placement as a permanent label. A legacy application may remain where it is until a replacement and archival strategy are ready. A research workload may begin in public cloud and later move when demand becomes predictable. A service may become a better SaaS candidate after its data and integration requirements are clearer. A good decision gives you a defensible answer today without trapping the organization there forever.

Private and managed environments remain important

Public cloud receives much of the attention in modernization discussions, but private and managed environments remain essential in healthcare. For some workloads, predictable performance, dedicated capacity and clear accountability matter more than elasticity. The most modern choice is the one that best fits the clinical reality of the service.

Core clinical platforms are often strong candidates for private or managed environments. They may serve thousands of users and support workflows that can't tolerate inconsistent performance. Demand may be stable, but the consequence of interruption is high. In that situation, dedicated capacity, healthcare-specific controls, defined recovery commitments and a named party responsible for the service can be more important than the ability to scale on demand.

Cost predictability also matters. A fixed or clearly defined service cost may be easier to govern than an operating model that changes with compute, storage growth, data transfer, resilience and support activity. That doesn't make private cloud automatically superior. It just means the organization can choose it when the workload requirements and operating economics justify the decision.

We chose a private, managed approach for the mission-critical clinical platform at the pediatric health system where I served as CIO. The decision wasn't based on which cloud model sounded more modern. It started with the workload.

The platform required an extremely high availability commitment, predictable performance, healthcare-specific security and recovery controls, around-the-clock access to specialized Epic expertise and a cost model finance could plan around. Rackspace designed a HIPAA-compliant, HITRUST-certified hosted private-cloud environment with a 99.999 percent uptime commitment, 24x7x365 support from certified Epic specialists and fixed-fee pricing.

Just as important, the platform was designed as a complete clinical service rather than a collection of servers. Network configuration, security, performance, redundancy, recovery, support and commercial accountability were considered together. That approach should be the lesson for architecture: the right destination is the one that fits the workload's clinical, technical, operational and financial requirements.

Public cloud creates value when the workload fits

Public cloud is an essential part of a modern healthcare architecture when a workload can use its strengths. It provides rapid access to elastic compute, advanced analytics, development platforms and specialized services that would be difficult or expensive to build internally. It supports research collaboration, temporary processing needs, patient-facing applications, AI experimentation and services with variable demand.

Use public cloud because the workload benefits from it, not because the strategy began with a platform mandate. Elasticity creates value when demand changes. Advanced services create value when the organization has the skills and governance to use them. Rapid provisioning creates value when ownership, security, data and cost controls are already clear.

But be aware that the real economics often emerge after the workload is in production. Compute, storage, network traffic and managed services often grow faster than expected. Large clinical and imaging datasets can create egress charges and performance challenges. A tightly coupled application may require far more surrounding infrastructure than the original estimate assumed. And a team that lacks the operating discipline to monitor consumption can turn agility into unbudgeted cloud sprawl.

AI and analytics make the distinction between experimentation and production especially important.

Public cloud may be the right place to test a model or access scalable compute. But a successful demonstration doesn't prove that the production architecture will work. Before a pilot becomes an everyday service, determine where authoritative data resides, what latency the workflow requires, how the service will be secured and supported, and whether the cost remains defensible at scale.

Used selectively and governed well, public cloud helps the organization move quickly. Used as an automatic destination, it can create the next generation of technical debt in a more variable and expensive operating model.

Consider SaaS and local survivability where appropriate

Private and public cloud don't complete the hybrid model. SaaS removes infrastructure responsibility for capabilities that are standardized, well integrated and supported by a service model the organization can govern. It's a strong fit for administrative, collaboration or specialized clinical functions when data ownership, integration, identity, recovery, performance and exit provisions are understood.

However, the health system must remain accountable for the service even when a vendor operates it. You still need to know where the data resides, how access is controlled, how the service integrates with clinical workflows, what happens during an outage and how the organization retrieves its data if the relationship ends.

Local survivability addresses a different healthcare reality. Clinical care can't assume that wide-area connectivity or a central service will always be available. Some workflows require local access to essential information, device coordination or limited operational capability during a broader disruption. The requirement should be determined by the clinical consequence of disconnection, not by a blanket edge strategy.

Local or edge capability is necessary when telemetry and medical-device workflows require very low latency, when essential information must remain available during network isolation or when a care site needs to operate in a degraded mode while central services are restored. The point isn't to place every clinical workload at the edge but to identify where losing connectivity would create an unacceptable interruption to care.

While a hybrid architecture may sound like a collection of platforms, it's actually a set of deliberate service placements, each chosen because it supports a particular combination of care, performance, resilience, cost and accountability.

Data gravity and application behavior shape the design

Placement categories provide direction, but application behavior determines whether the design will work. Architecture diagrams make workloads appear more independent than they are. In the operating environment, applications are tied to databases, storage, identity services, integration engines, reporting platforms, security controls and clinical workflows. A workload rarely moves alone.

Data gravity is especially important in healthcare. Historical EHR databases, PACS and enterprise imaging archives, genomics data, research datasets and long-term clinical records can become so large that moving them repeatedly introduces latency, transfer time, egress charges, duplicated storage, synchronization risk and recovery complexity. The data effectively anchors the compute and applications around it.

An application appears portable on a diagram but becomes impractical to move once the surrounding data and dependencies are understood. A workload considered suitable for public cloud may need to remain close to a large clinical archive. An application that looks independent may have to move with identity, integration and database services. And a proposed migration may need to occur in phases because the surrounding environment is tightly coupled.

Application behavior matters just as much. Some legacy systems depend on a specific operating system, database or hardware configuration. Some integrations rely on response times that were never documented. Some applications perform well in testing but behave differently under real clinical volume. Others carry licensing terms that change when the workload moves.

Architecture must reflect how the application behaves under real clinical demand, not just how it appears on a diagram. Difficult dependencies must be resolved, moved with the workload or deliberately accommodated before the placement decision becomes final.

Design IoMT and zero-trust controls into the architecture

Healthcare architecture must account for the medical devices and connected systems that sit outside the traditional application portfolio. Many devices operate on aging or shared networks, can't support modern authentication, are difficult to patch and depend on local services that may not be visible in the application inventory.

The target state should identify and segment medical-device traffic, control machine identities and service accounts, apply least-privilege access and reduce opportunities for lateral movement. It should also protect the data flows connecting local devices, clinical platforms and cloud services. A device that can't support a modern control still needs compensating protection around it.

Zero trust starts with a simple operating assumption: never trust, always verify. No user, device or network connection should be treated as safe by default simply because it sits inside the health system's traditional perimeter. The **IoMT and Zero Trust Security Matrix** in the companion toolkit helps your team evaluate those requirements in more detail.



Design security and recoverability into the target state

Whatever environment a workload enters, the architecture must explain how the organization will protect and restore the complete service. Security, backup, disaster recovery and cyber recovery are not controls added after placement. They're part of the placement decision itself.

Backup creates protected copies. Disaster recovery restores services after an infrastructure or site failure. Cyber recovery assumes that production and conventional recovery environments may both be compromised. These are different processes and require different designs.

A cyber-resilient architecture requires immutable or otherwise protected copies, isolated or air-gapped recovery capability, clean-room restoration, identity recovery and dependency-aware sequencing. The organization must validate that systems and data are trustworthy before reconnecting them to clinical operations.

Restoring the primary application isn't enough if identity, integrations, imaging, communications or downstream workflows remain unavailable. The protection model should define where copies reside, how frequently data is captured, which dependencies must return with the application and how essential clinical workflows continue before full restoration.

Ask the following operational questions during the design process: Which services return first? How long can the workflow operate in degraded mode? Who decides that restored data can be trusted? Where will recovery occur if the primary and secondary environments are under investigation? Recovery Confidence and Clinical Workflow Recovery Time should be designed into the architecture, not discovered during an incident.

Section 6 addresses how the organization tests and manages these capabilities continuously after migration. Section 4 makes sure the future state supports them in the first place.

Contracts, service levels and governance are part of the architecture

An architecture diagram shows platforms, connections and security boundaries. However, it doesn't tell you who answers the phone at 2 a.m., who has authority during an incident, what a vendor is obligated to do or how the organization exits the relationship if the model no longer works. Those are architecture questions because they determine whether the service can be operated as designed.

The target state should define the service commitment required for each workload category. A mission-critical clinical service requires continuous support, rapid escalation, short recovery objectives and greater redundancy. A lower-priority administrative service requires a different protection and support model. The commitment should reflect the clinical consequence of failure rather than a standard service package.

Recovery time and recovery point objectives should be translated into clinical terms. A four-hour recovery target means certain workflows are unavailable for up to four hours, alternate procedures are required and some information needs to be recreated. A service-level percentage isn't enough if it doesn't explain what clinicians expect at the point of care.

The design should specify what responsibilities remain with the health system, what's performed by a managed-services or cloud partner, who monitors the service, who communicates with clinical leaders and who invokes recovery. It should address data ownership, portability, exit rights, third-party cyber risk, contractual remediation triggers and authority during multi-vendor incidents.

I've seen technically sound designs stall for reasons that had little to do with the technology. Finance challenged an untested assumption. Legal found an unresolved liability. Procurement uncovered a licensing constraint. Clinical leaders didn't trust the service commitment. Settling those issues during design prevents them from becoming operational disputes later.

> Design becomes an executable sequence. Once the target architecture is clear, you must decide which workloads move first, what must move together, how clinical workflows will be protected and what the organization will do when a migration doesn't perform as expected. That's the work of the migration matrix in Section 5.

The operational co-pilot model

Healthcare organizations don't need another strategy deck that leaves their overstretched internal team to absorb the operational burden. What they need is an operational co-pilot model that continues through design, engineering, migration, accountability and day-two management.

An operational co-pilot helps validate discovery, design the architecture, build the environment, establish connectivity and security controls, prepare the migration, test the complete service and support the transition into ongoing operations. The partner brings engineering capacity and operating discipline while the health system retains responsibility for clinical priorities, risk tolerance, architecture standards, data governance and final decisions.

The pediatric transformation we did with Rackspace reinforced another important lesson: architecture and operating responsibility have to be designed together. A strong target state isn't enough if no one is clearly accountable for operating, protecting and recovering the complete service after the move.

A modern architecture therefore needs a modern responsibility model. It must define not only where workloads run, but who's accountable for keeping them secure, available, recoverable and financially sustainable—and who has authority when those commitments are at risk.



The right architecture creates optionality

No architecture can predict every future requirement. Clinical workflows will change, threats will evolve, AI capabilities will mature and financial models will shift. The goal isn't to predict every change but to preserve enough flexibility, knowledge and control that the organization doesn't have to start over each time.

In a disciplined hybrid model, critical services operate in environments designed for reliable performance and predictable support. Research, analytics and AI workloads use scalable services where those capabilities create value. Local functions preserve essential operations during connectivity loss. And workloads move or change as their technical, clinical and financial characteristics evolve.

Optionality protects the organization from turning today's platform decision into tomorrow's technical debt. Be cautious about designs that make data difficult to retrieve, place all operational knowledge with one party or depend on proprietary services without a credible exit path. Preserve data portability, internal knowledge and the ability to revisit placement when economics or workload behavior changes. Optionality doesn't mean avoiding commitment. It means understanding the commitment well enough that meaningful choices remain.

Once the target architecture is clear, design has to become an executable sequence. You must decide which workloads move first, what must move together, how clinical workflows will be protected and what the organization will do when a migration doesn't perform as expected. That's the work of the migration matrix in Section 5.



CIO Summary — Section 4

Designing the Hybrid Health System: Architecture for Reality, Not Hype

What this means for the CIO

The future-state architecture should be based on what each workload requires, not on one preferred platform. Clinical criticality, latency, data gravity, local survivability, security, recoverability, cost and operational accountability should determine placement. The same decision must also define how the complete service will be protected, operated and governed.

Questions to ask now

- Which workloads require predictable performance, dedicated capacity and managed reliability?
- Which workloads genuinely benefit from elasticity, advanced services or rapid provisioning?
- Have we accounted for the data gravity of EHR, imaging, research and historical archives?
- Which clinical workflows require local survivability or edge capability?
- Does the architecture distinguish backup, disaster recovery and cyber recovery?
- Are vendor responsibilities, exit rights and multi-vendor incident authority clear?
- Does the design preserve future options without creating the next generation of technical debt?

Helpful tools from the Healthcare Modernization Toolkit:

Use the **Workload Placement Matrix** to determine where each retained workload belongs based on clinical criticality, latency, data gravity, dependencies, security, recovery, scalability, operating accountability and cost.

Use the **IoMT and Zero Trust Security Matrix** to identify required security controls and compensating protections, the **Service Accountability and Recovery Commitments Worksheet** to define service and recovery responsibilities, and the **Vendor Dependency and Risk Governance Matrix** to clarify third-party obligations, escalation, remediation and exit readiness.

Use the **TCO and Cost Predictability Calculator** to test whether the proposed architecture is financially understandable and sustainable. If an AI capability is being designed or expanded during this phase, use the **AI Infrastructure Readiness Checklist** to test whether the proposed architecture, data, governance, support and cost model can support production use before the capability advances.

Section 5: The Migration Matrix

THE EXECUTION: Moving the Big Rocks Without Disrupting Care

A strong architecture tells you where you're going and what you intend to move. However, it doesn't guarantee that you'll get there safely. The real test begins when the team separates a workload from its current environment, moves its data and integrations, validates the destination and returns the complete service to clinical and operational use.

Architecture determines where a workload belongs but migration determines whether the organization can move it without losing control. An undocumented interface surfaces hours before cutover. A department uses the application during a period everyone believed was inactive. A vendor misses a connectivity milestone. A licensing condition changes the sequence. A recovery process that looked sound on paper fails during the rehearsal. The migration matrix exists to expose those conditions before they reach the bedside.

You can't eliminate every surprise but you can keep a surprise from becoming a clinical disruption. That requires a migration model built around readiness, dependencies, complete-service testing, clinical adoption, decision authority, communication and care continuity. A technically successful cutover isn't enough if the workflow doesn't perform as expected when clinicians return to it.

Confirm that the workload still deserves to move

Rationalization began during triage, but execution requires a final checkpoint before a workload enters a migration wave. New information may have changed the original decision and the organization shouldn't spend money moving an application that no accountable leader is prepared to defend.

Before each workload proceeds, confirm that:

- An accountable business, clinical and technical owner still supports it.
- The clinical or operational function remains necessary.
- The disposition decision is current.
- Historical-data access, retention and archival requirements are understood.
- Contracts, licensing and vendor responsibilities permit the move.
- The destination remains appropriate for the workload and its dependencies.
- The value of migration still justifies the cost, effort and risk.

This checkpoint prevents a migration plan from becoming a conveyor belt that carries every legacy decision into the future state. Rationalization now becomes an execution gate rather than a repeated discovery exercise.

The migration matrix connects priority and readiness

An inventory tells you what exists and a target architecture tells you where retained workloads belong. The migration matrix helps you determine what can move safely, what must move together and what isn't ready yet.

The matrix should consider:



Clinical criticality.

How directly does the workload affect patient care and what happens if it's slow or unavailable during the transition?



Operational importance.

How much does it affect patient access, revenue cycle, scheduling, staffing, supply chain or communication?



Technical complexity.

How many applications, interfaces, databases, identity services, devices and data stores are connected to it?



Dependency grouping.

Which services must move before it, with it or immediately after it?



Destination readiness.

Has the target environment been built, secured, sized and tested under production-like demand?



Recovery readiness.

Are backup, restoration, cyber recovery, rollback and continuity procedures proven?



Clinical change readiness.

Are clinicians and operational teams prepared for the workflow change, training, support and degraded-mode procedures?



Vendor readiness.

Are application, network, cloud, interface, device and support vendors prepared, with clear responsibilities and escalation paths?



Commercial readiness.

Are cost, contract, licensing, data ownership and partner obligations settled?

One of the easiest mistakes to make is confusing urgency with readiness. A fragile clinical system may be the organization's highest priority, but that doesn't mean it belongs in the first wave. Its dependencies may still be unclear, its rollback path untested, its destination incomplete or its clinical timing unacceptable.

The CIO and CTO should be able to explain not only why a workload matters, but why it's ready now. They should also be able to state what evidence is still missing, who owns the gap and what conditions would cause the wave to wait. Urgency determines priority but only evidence should determine readiness.



Design migration waves around risk and learning

The first migration waves establish the organization's execution discipline. They should prove that the team plans carefully, communicates clearly and makes decisions under pressure without creating unnecessary clinical exposure.

Early waves should be meaningful enough to test the model but controlled enough that the consequences remain manageable. If the workloads are too simple, the organization learns very little about how the process performs under real conditions. If they're too consequential, the program takes on risk before the operating model is ready.

Every wave should have explicit entry criteria. Before it proceeds, the organization should confirm that dependencies are understood, destination capacity is available, security controls are active, vendors and the operational co-pilot are ready, contracts are settled, clinical timing is acceptable, recovery has been tested, users are prepared and decision-makers are available.

Every wave should also have exit criteria. The workload isn't complete because the technical cutover is finished. Performance must be stable, integrations must work, protection and recovery must be active, support ownership must transfer, Clinical Trust Metrics must remain within agreed thresholds, clinical and operational workflows must be accepted and open issues must have accountable plans.

Timing should also reflect the reality of the health system. A technically available weekend may overlap with a regulatory deadline, revenue-cycle close, major clinical upgrade, seasonal surge or staffing constraint. Build the migration calendar around the needs of the organization, not only around when the technical team can perform the work.

Each completed wave should improve the next one. Runbooks should become clearer, testing more realistic, communication more precise, vendor coordination stronger and decisions faster. A migration wave is more than a move. It's a controlled test of the organization's ability to change without damaging clinical trust.

A workload may be technically ready while the clinical organization isn't

Technical readiness is only one part of migration readiness. A destination can be fully built, a runbook complete and every interface pass testing while physicians, nurses and operational teams remain unprepared for the workflow change.

Before go-live, assess whether the following questions have been answered:

- Which roles and departments will experience a change in workflow?
- Has training been completed and retained?
- Do users understand what will be different when the service returns?
- Are local clinical champions and at-the-elbow support available?
- Do teams understand degraded-mode and downtime procedures?
- Are help-desk and escalation routes clear?
- Does clinical sentiment reveal unresolved concerns that could affect adoption or safety?

Don't mistake attendance at a training session for readiness. The organization needs evidence that users understand the new workflow, know where to get help and can continue care if the migration takes longer than planned.

Readiness is demonstrated through behavior, support coverage and confidence in degraded-mode procedures—not a list of attendance. **The Clinical Change and Adoption Readiness Tool** in the companion toolkit is designed to make those conditions visible before the cutover window opens.

Protect clinical and operational workflows

As a wave moves from planning into execution, clinical trust becomes an operating requirement. A workload may appear technically ready while the workflow it supports remains vulnerable.

Before the move, identify who uses the system, where it's used, which steps depend on it and when the service is most sensitive to interruption. Understand what information users need, what connected services must remain available and how work will continue if the event runs long.

Be sure to find out:

- What will stop during migration?
- What will slow down?
- What could become unsafe or create uncertainty?
- Which degraded-mode process will be used?
- Who has been trained on that alternate process?
- How long can the workflow safely remain degraded?
- How will information created during the interruption be reconciled?
- Who will confirm that the complete workflow has been restored?

Clinical and operational leaders should validate the answers because they understand the work in ways the migration team may not. They should identify unsafe timing, unrealistic contingency assumptions and workflow dependencies that don't appear in technical documentation.

Clinical acceptance must be part of the completion criteria. Infrastructure dashboards may show that the platform is available while images load slowly, orders are delayed, authentication interrupts care or users encounter errors that never appear in the infrastructure view. The workload isn't successfully migrated until the complete service performs safely at the point of care and the relevant Clinical Trust Metrics confirm that result.

Test the complete service, not isolated components

A database is reachable, an application opens and an interface passes a transaction while the full clinical workflow still fails. Testing must prove that the complete service works under conditions that resemble real use.

Test identity and authentication, databases, interfaces, imaging, orders, clinical documentation, downstream systems, backup and restoration. Simulate peak volume, degraded connectivity, failed interfaces, delayed data, partial cutover and production-like demand. Confirm that records arrived, counts reconcile, timestamps are correct, permissions remain intact and downstream services received what they expected.

Use the Clinical Trust Metrics introduced in Section 2 as validation criteria. Bedside Application Latency and Time-to-Clinical-Context should remain within acceptable ranges. Clinical Workflow Recovery Time should reflect the full service, not only the primary application. Workaround Frequency during stabilization should be visible, and Recovery Confidence should be based on demonstrated testing rather than assumption.

The organization should also define who accepts each layer of the result. Infrastructure teams confirm capacity and availability. Application owners validate configuration and data. Security teams verify controls and monitoring. Vendors confirm their components and obligations. And clinical and operational representatives confirm that the workflow behaves correctly.

For a complex or clinically important workload, rehearse the event. A rehearsal should test sequence, timing, staffing, vendor participation, communication, clinical validation, decision authority, escalation and rollback. The purpose isn't to prove that nothing can go wrong but to expose hidden assumptions while there's still time to correct them.

Use an execution-control framework

Safe migration depends on three connected controls: a detailed runbook, predefined decision and rollback authority, and disciplined stakeholder communication. The runbook defines the work; decision authority preserves control; and communication keeps the organization aligned when conditions change.

1. The runbook defines how the work happens

A migration plan describes what the organization intends to accomplish. The runbook describes the exact steps before, during and after the event.

Before migration, the runbook should confirm that dependencies are validated, backups are complete, recovery is tested, destination capacity is verified, monitoring is active, security controls are in place, vendors are ready, and clinical and operational stakeholders are prepared. During migration, it should provide the precise sequence: what changes first, when data moves, when integrations are tested, when users return and where pause points occur. It should identify the evidence required to continue from one stage to the next.

After migration, the runbook should guide technical, application, security, vendor and workflow validation. It should also capture what the team learned. For example, which steps took longer, what dependency was missed, whether communication reached the right people, whether adoption support was sufficient and whether decision thresholds worked under real conditions.

For major events, a temporary command structure provides a shared view of status. The people who diagnose, decide, communicate and protect the workflow—including accountable internal leaders, critical vendors and the operational co-pilot—must be present when they're needed, not reachable only after the event has already changed course.

2. Decision and rollback authority preserve control

No migration plan should assume that every step will perform exactly as expected. An interface may fail after passing a test. Performance may degrade under production demand. Data may not reconcile.
A vendor may miss an action. Or a clinical workflow may be affected in a way the technical team didn't anticipate.

The organization needs a safe way to pause or reverse before those problems become larger. Rollback criteria should be defined before the event begins. The team should know what conditions justify a pause, what conditions require reversal and how much time is available to decide.

Rollback must be specific to the workload. Can the application return to the prior environment? Has data changed since the move began? Will records or transactions require reconciliation? Must integrations be reversed in a particular order? Can clinicians continue working during the reversal?

Decision authority is equally important. Be explicit about who approves the start, pauses the event, invokes rollback and accepts the workload as stable. Those decisions can't be left to an undefined group once the migration is underway.

Rollback isn't evidence of failure. It's evidence that the organization values care continuity more than the appearance of staying on schedule.

3. Communication keeps stakeholders oriented

A migration can succeed technically and still damage trust if clinicians, operations, support teams or executives are surprised. Communication must be designed into the event.

Different audiences need different information:

- Executives need the current risk, status, major issues and whether a decision is required.
- Clinical leaders need to know which workflows may be affected and what contingency procedures apply.
- Frontline users need practical guidance on timing, expected behavior and how to report problems.
- Help-desk and support teams need known-issue information, escalation routes and ownership.
- Partners and vendors need clear responsibilities, checkpoints and update expectations.

Communication should occur before, during and after the event. Before the move, let stakeholders know what's changing, why, when and with what expected impact. During the event, provide updates at agreed intervals and communicate quickly when conditions change. Afterward, explain what has been validated, what remains open and when normal operations have resumed.

Honesty matters. If the organization doesn't know the cause of an issue, say so while explaining what's being done and when the next update will occur. Confidence in a move requires visible preparation, consistent information and clear accountability.



What disciplined execution looked like in practice

During the pediatric health-system transformation in the Pacific Northwest, we applied that discipline to the migration of a mission-critical clinical platform. The cutover was treated as an operating event, not simply a technical task. With Rackspace, we coordinated the technical sequence, vendor responsibilities, validation, communications and decision authority before the migration window opened. Rollback remained available if the evidence showed that continuing the migration would put care continuity at risk.

The transition was completed four weeks ahead of schedule and no emergency rollback was required. Those results demonstrated what happens when risk is surfaced early enough to manage it. That meant dependencies were understood, the destination was ready, responsibilities were clear and the team knew exactly what conditions would cause it to pause or reverse course.

That's the standard I'd apply to any major clinical migration. Success isn't simply reaching the new environment. It's reaching it under control, with the ability to stop when conditions change and with enough operational discipline that the next migration wave begins from a stronger position than the one before it.

Stabilize the workload before closing the wave

Cutover shouldn't be considered the finish line. Production volume may increase gradually, overnight jobs may not have run and some workflows may not be used until the next shift or operating cycle. Before declaring a victory, give the workload a defined stabilization period before deciding the wave is complete.

During stabilization, monitor:

- Performance and bedside latency
- Incidents and support volume
- Data reconciliation and interface behavior
- Security alerts and access anomalies
- Backup and recovery status
- Cost and consumption signals
- Clinical experience and workaround frequency
- Training and adoption issues
- Vendor response and unresolved obligations
- Support ownership and open defects

Support teams should know which issues are expected and which conditions require escalation. Temporary migration resources should remain available until ownership has transferred and the operating team supports the service without depending on informal knowledge from the project team.

Open defects and exceptions must remain visible. Distinguish between an issue that can be accepted temporarily, an issue that must be corrected before the wave closes and an issue serious enough to trigger rollback. Clinical and operational leaders should know what remains open and how it's being managed.

The wave is complete when the service has demonstrated stable performance through the required operating cycles, workflows have been accepted, protection and recovery are active, ownership is clear and unresolved items have accountable plans. This prevents a technically completed cutover from passing hidden risk into day-two operations.

Measure success by the condition of the service

Counting workloads and meeting move dates shows activity, but it doesn't show whether the organization has modernized successfully. Migration success should be measured by the condition of the service after the move and by whether each completed wave increases confidence in the next one.

Each wave should be evaluated against a concise scorecard:

- **Risk reduction:** Did the move reduce technical debt or merely relocate it?
- **Clinical continuity:** Did care and operational workflows continue safely and return fully?
- **Performance and reliability:** Does the workload perform more consistently in the new environment?
- **Recoverability:** Can the complete service and its dependencies be restored as designed?
- **Clinical adoption:** Are users working effectively in the new workflow, with fewer workarounds and manageable support needs?
- **Visibility and accountability:** Are ownership, monitoring, escalation and vendor responsibilities clearer?
- **Financial discipline:** Is the operating cost more understandable and defensible?
- **Confidence:** Did the wave increase stakeholder confidence in the next one?

The program should also measure whether its execution capability is improving. Are runbooks stronger? Are dependencies documented earlier? Are clinical stakeholders engaged sooner? Are vendors meeting their commitments? Are decisions faster and better informed? Is risk being identified before the event rather than during it? Those answers show whether migration is creating a durable organizational capability.

Once workloads are operating in the new environment, responsibility shifts again. The health system must sustain performance, cost, security, recovery, adoption and partner accountability after the migration team stands down.

That brings us to the sixth and final part of the modernization journey: operating the new environment with discipline.

> Rollback isn't evidence of failure. It's evidence that the organization values care continuity more than the appearance of staying on schedule.

CIO Summary — Section 5

The Migration Matrix: Moving the Big Rocks Without Disrupting Care

What this means for the CIO

Migration is more than a technical cutover. It's a clinical, operational, organizational and partner-coordination event that must be controlled through evidence, readiness, complete-service testing, decision authority, communication and stabilization. A workload should move because it's ready, not merely because it's urgent.

Questions to ask now

- Which workloads are ready to move, not merely urgent?
- Are clinical and operational teams prepared for the workflow impact?
- Are all dependencies, partners and vendors ready?
- Has the complete service been tested under realistic conditions?
- Who can approve, pause, reverse and accept the event?
- What evidence will show that the workload is stable, clinically usable and ready for day-two operations?

Helpful tools from the Healthcare Modernization Toolkit:

Use the **Migration Wave Planner** to sequence workloads according to readiness, dependencies, clinical timing and destination readiness rather than urgency alone.

Use the **Safe-Exit and Rollback Checklist** to establish pause and rollback triggers, decision authority and a verified return path before each migration begins.

Use the **Clinical Change and Adoption Readiness Tool** to confirm that clinicians, operational teams, support functions and degraded-mode procedures are ready for go-live and stabilization.

Section 6: The New Operating Model

THE FUTURE: From Project to Product

A successful cutover can feel like an ending. The workloads are running, the milestones are complete and the program team begins to stand down. But go-live is where the project ends and operational accountability begins. For the people who must keep the service available, secure, recoverable and financially defensible every day, the real work is just beginning.

The environment doesn't stop changing because the migration is complete. Demand keeps growing, storage keeps expanding and clinical workflows continue to develop. Plus, new devices, interfaces and data sources enter the environment. If these changes aren't governed deliberately, you begin accumulating a new form of technical debt almost as soon as the old one has been reduced.

Modernization isn't a project with a defined budget, scope and end date. The new technology foundation has to be operated as a set of continuously managed clinical and business services with clear ownership, service commitments, financial visibility, recovery requirements and a roadmap for change. The discipline that protected the migration must become the way the organization operates every day.

Modernization doesn't end at go-live

Go-live is the day your assumptions meet real users, production volume and operating pressure. A service may have passed every test and still behave differently when thousands of clinicians, staff members, devices, interfaces and downstream systems depend on it at the same time.

That's why the first 90 to 180 days matter so much. This is when the organization learns whether expected performance holds under normal clinical demand, support responsibilities work as designed, vendors respond when needed and the cost model behaves the way the business case predicted. It's also when temporary migration decisions can quietly become permanent operating practices if no one is accountable for closing them.

Watch for signs that the operating model hasn't fully taken over, including:

- Higher-than-expected incident volume.
- New clinical workarounds.
- Unclear support ownership.
- Unresolved migration exceptions.
- Cloud or storage consumption growing faster than workload value.
- Recurring vendor issues.
- Recovery procedures changing without validation.
- Project-team members continuing to resolve problems through informal knowledge.
- Day-two teams depending on personal relationships rather than documented escalation paths.

These aren't minor administrative details; they're evidence that operational accountability is incomplete.

I've seen organizations treat a successful cutover as proof that the work is finished. The environment may look stable while the people who delivered the migration continue solving problems through institutional knowledge and personal intervention. When those people move to other priorities, the organization discovers that the service still depends on the project team.

A successful move is complete when ownership has transferred and the day-two team operates the service without relying on individual memory.

Complete the operational handoff

The stabilization period introduced in Section 5 should end with a formal operational handoff. This is the point where the health system proves that the service has moved from a temporary project structure into a durable operating model.

At minimum, the handoff should include:

- Named service, application, clinical and technical owners.
- Current architecture, interface and dependency documentation.
- Open defects, accepted risks, temporary exceptions and deferred decisions.
- Performance, capacity and Clinical Trust Metric baselines.
- Support, escalation and communication responsibilities.
- Backup, disaster-recovery and cyber-recovery procedures.
- Cost assumptions, budget thresholds and ownership for variance.
- Vendor obligations, service levels and remediation paths.
- Known operational constraints and conditions that require architecture or placement to be reconsidered.
- A clear process for approving future changes.

Don't accept vague ownership. A service can have several contributors, but it still needs one accountable owner who explains its performance, risk, cost and future direction.

The CIO and CTO should be able to answer several basic questions:

- Who owns the service after the migration team leaves?
- Who monitors technical performance and clinical experience?
- Who investigates cost variance?
- Who decides whether an accepted risk remains acceptable?
- Who confirms that recovery procedures still work?
- Who holds the partner accountable?
- Who has authority when several vendors are involved and the service is still unavailable?
- Who decides when the architecture or workload placement should be revisited?

If those answers remain unclear, the organization has completed a move but hasn't completed the modernization process.

Governance becomes an operating discipline

During migration, governance is visible because every decision feels urgent. Dependencies are tracked, risks are escalated and leaders know that a delayed decision can stop a wave. After go-live, the urgency fades. That's when the organization is most vulnerable to rebuilding the complexity it worked so hard to remove.

Technical debt usually returns one small decision at a time. A new application duplicates an existing capability. A temporary exception never expires. A cloud service expands without a cost review. An interface changes and no one updates the dependency map. A vendor retains remote access after the work is complete. A recovery procedure is changed but never retested. Each choice appears manageable in isolation until the organization is once again carrying an environment no one fully understands.

The new operating model should govern six responsibilities:

1. **Architecture governance.** Confirm that new and changing workloads still fit the placement principles established in Section 4.
2. **Clinical and operational governance.** Connect technology decisions to care delivery, workflow impact, adoption and clinician experience.
3. **Security and resilience governance.** Evaluate identity, protection, recovery, compliance, IoMT exposure and third-party risk.
4. **Financial governance.** Make consumption, cost ownership, forecast variance and value visible.
5. **Service governance.** Review performance, incidents, capacity, Clinical Trust Metrics, partner commitments and planned improvements.
6. **AI governance.** Review data, models, clinical validity, privacy, security, ethics, cost, drift and safe use in production.

Good governance doesn't mean creating a large committee for every decision but it does involve the right people with the authority to decide. A small operating group can handle most reviews, bringing in clinical, security, finance, legal, procurement, privacy, data, biomedical or vendor-management expertise when the decision requires it.

Governance also requires evidence. Leaders should review technical performance and clinical experience together. They should examine incident patterns, workaround frequency, cost trends, recovery results, security findings, vendor performance, architecture exceptions, unresolved decisions and progress against the modernization business case.

Governance should become the discipline of seeing risk early enough to act while the organization still has choices.

Keep the environment visible

Section 3 established the inventory, ownership, criticality and dependency baseline required to modernize responsibly. That baseline isn't a static project document because applications change, interfaces are added, data flows evolve and temporary exceptions accumulate. If the inventory isn't maintained, the new environment will gradually become as poorly understood as the old one.

The operating model requires updates whenever a service changes. At minimum, the organization should maintain:

- Service and workload ownership.
- Clinical and operational criticality.
- Architecture and placement decisions.
- Interfaces, APIs and shared services.
- Identity and access dependencies.
- Vendor relationships.
- Recovery requirements.
- Archival and retention obligations.
- Architecture exceptions.
- Local-survivability requirements.
- IoMT and device dependencies.
- Cost ownership.
- Temporary decisions and their exit conditions.

This is especially important in a hybrid environment where one clinical workflow may depend on managed infrastructure, public cloud, SaaS platforms, local devices, identity services and several vendors at the same time.

Treat the current-state decision package as an ongoing operating asset. It tells leaders what the organization is running today, not what was true when the migration program began. Remember, if you stop maintaining visibility, you begin rebuilding uncertainty.

Govern cost and value together

Modernization makes technology spending more visible, but cost predictability has to be managed. A hybrid environment combines managed-service fees, public-cloud consumption, software licensing, support tiers, storage growth, data-transfer charges, backup and recovery capacity, security services, network services, internal staffing and specialist vendor support. Each behaves differently over time.

The business case approved before migration was a forecast. Once workloads are live, the invoices reveal what the environment actually costs. Compare forecasts with actual performance before storage growth, data movement, unused capacity, premium support or new cloud services become budget surprises.

Healthcare FinOps should extend beyond public cloud. Finance, technology and workload owners need a shared view of what is being consumed, what clinical or business value it supports and who's responsible when cost or usage changes. Cost ownership should be attached to services and workloads rather than hidden inside a general infrastructure budget.

A disciplined cost-governance rhythm includes:

- Forecast-to-actual review by service and workload.
- Budget thresholds and alerts for consumption, storage, data movement and premium services.
- Named ownership for investigating and correcting material variance.
- Visibility into underused capacity, overlapping licenses and duplicate services.
- Review of whether reliability, recovery and security commitments remain appropriately funded.
- Reassessment of workload placement when economics or application behavior materially change.
- Clear distinction between necessary resilience costs and avoidable waste.

Don't let optimization become indiscriminate cost cutting. Some costs exist because clinical reliability, cyber recovery and around-the-clock support require them. The objective isn't to spend less but to understand what causes spending to change, what value that spending supports and what action will be taken when the model no longer holds.

Cost governance tells you what the service is consuming but value realization tells you whether that consumption is producing the result leadership approved.

A modernization program can finish on time and still fail to deliver its promised outcomes. To stay on track, the operating model needs a **Value Realization Scorecard** that compares actual results with the original business case at defined post-implementation milestones.

The scorecard should track:

- Actual operating cost compared with the forecast.
- Legacy applications retired and support expense removed.
- Incident volume, severity and recurring root causes.
- Clinical performance, workflow improvement and reduction in workarounds.
- Recovery performance and confidence in cyber resilience.
- Vendor performance and service-level compliance.
- Internal capacity redirected from maintenance to clinical, data, security and innovation priorities.
- Time required to deliver new capabilities after modernization.

The purpose of the scorecard is to keep the operating model accountable to the outcomes leadership approved. When expected value isn't appearing, the organization should understand why and decide whether to correct the service, revisit workload placement, change the vendor relationship, address an adoption or workflow problem, modify the operating model or reassess the original assumption.

Modernization remains credible when leadership understands both what the environment costs and what the organization is receiving in return.

Prove reliability, resilience and security continuously

Don't assume that reliability and recoverability proven at go-live will remain intact. We know that applications change, data grows and interfaces are modified. Any of these changes can affect performance, protection and restoration.

Therefore, reliability should be reviewed through both technical and clinical evidence. A platform may meet its uptime commitment while clinicians experience slow logins, delayed images, repeated authentication, inconsistent application behavior, workflow interruptions or workarounds that remain after an incident is marked resolved. The operating model must connect service data with what users experience at the point of care.

Clinical Trust Metrics should be part of regular service review:

- Bedside Application Latency.
- Time-to-Clinical-Context.
- Zero-Interruption Care Hours.
- Clinical Workflow Recovery Time.
- Workaround Frequency.
- Recovery Confidence.

These measures help leaders judge whether the complete service is supporting care, not merely whether an infrastructure component is available. Operating evidence should also change service commitments when care requires something different. If experience shows that a workflow can't tolerate the agreed recovery time or latency threshold, revise the commitment rather than defend a number that no longer reflects clinical reality.

Recovery capability requires the same discipline. Backups should be monitored, restored and validated. Critical applications should be tested with identity, interfaces, data, imaging and other dependencies. Cyber-recovery environments should remain isolated, controlled and ready to use. Clean-room exercises should prove that systems and data are trustworthy before they are reconnected and that essential clinical operations can function while the broader environment is still being assessed.

Security must also evolve with the operating environment. New identities, devices, integrations and data flows create new exposure. Security, architecture and operations teams should review these changes together rather than treating security as a final approval. The platform may be resilient but the organization must keep proving that it is.



Operate IoMT and zero trust as living disciplines

Medical-device and zero-trust controls are part of that continuous proof. A device inventory becomes outdated quickly; vendor remote access remains active after a service event; unsupported devices are moved to a different network without updating segmentation; and machine identities and service accounts persist long after their original purpose disappears.

To manage these changes, the operating model should regularly review:

- Connected-device inventory.
- Network segmentation.
- Unsupported or unpatched devices.
- Machine and service identities.
- Access exceptions.
- Anomalous traffic.
- Vendor remote access.
- Compensating controls.
- Changes in device location or clinical use.
- Dependencies on local services and connectivity.

The **IoMT and Zero Trust Security Matrix** should be treated as a recurring governance tool. Zero trust isn't complete when the architecture is approved. It must remain visible in the way identities, devices, workloads and vendors are governed every day.

Evolve the internal team and the operational co-pilot relationship

One of the most important outcomes of modernization is a change in how the internal technology team spends its time. When skilled people are consumed by aging hardware, repeated incidents and undocumented exceptions, the organization loses more than efficiency. It loses capacity for clinical enablement, resilience, data improvement and innovation.

A modern platform or managed-service relationship creates an opportunity to redirect that capacity, but the shift won't happen by itself. Roles, expectations and development plans have to change with the environment.

Infrastructure specialists may move toward architecture, performance management, capacity planning, automation and partner governance. Application teams may work more closely with clinical and operational leaders to improve workflows. Security teams may devote more attention to resilience, threat preparation, identity, recovery and third-party risk. And data teams may focus on making information more useful and trustworthy for digital and AI initiatives.

This gives the organization much more capability to innovate but doesn't mean outsourcing the responsibilities that matter most. The health system still needs internal leaders who understand the clinical mission, know how services fit together and challenge vendors with evidence. While an operational co-pilot can operate infrastructure, monitor services, execute recovery and provide specialized expertise, it can't own:

- Clinical priorities.
- Risk tolerance.
- Architecture standards.
- Financial authority.
- Data governance.
- Final responsibility for care continuity.

The health system must also decide what knowledge can be transferred and what knowledge must remain internal. Before operational work moves to a partner, document service architecture, dependencies, clinical context, decision rights, escalation paths, known constraints, recovery procedures, accepted risks, vendor relationships and historical decisions that still shape the service.

One of the benefits of modernization is reducing the dependence on individual heroes without losing the knowledge that kept the old environment running. Earlier in my career, firefighting felt like proof of commitment because the team repeatedly rescued the organization from instability.

Modernization changes that measure of value. The team should be judged less by how many fires it puts out and more by whether services remain dependable, costs are understood, clinicians are supported, recovery is proven, partners are accountable and the organization is ready for what comes next.



Govern the operational co-pilot after go-live

A signed contract tells you what a partner promised but day-two operations show you what the relationship is worth. In a hybrid environment, partner management becomes a continuous operating discipline covering performance, incidents, cost, recovery, security, capacity, change, improvement and exit readiness. Section 4 established why the operational co-pilot model matters. Section 6 is where the health system proves whether the model is working.

The pediatric transformation I led reinforced why the operating model matters after the project team leaves. Moving the clinical platform was only one milestone. The longer-term value came from establishing a model in which specialized operational expertise, monitoring, escalation and infrastructure management continued around the clock while the health system retained accountability for clinical priorities and the service itself.

That changed where the internal team focused its attention. Instead of carrying every layer of infrastructure operation directly, technology leaders concentrated on clinical priorities, architecture, governance, resilience and the next set of improvements.

That's the real test of an operational co-pilot. The relationship shouldn't merely get you through a migration. It should leave the health system with a more sustainable way to operate—one in which responsibilities are explicit, performance is continuously governed and internal expertise is directed toward the work the organization most needs its people to own.

The RACI, service levels, recovery objectives and escalation paths established during architecture and migration should remain active management tools. Regular service reviews should examine:

- Performance and Clinical Trust Metrics.
- Incident trends and recurring root causes.
- Recovery testing and security findings.
- Capacity and cost.
- Planned changes and open obligations.
- Improvement commitments.

Partners should be expected to identify emerging risk and contribute to improvement, not simply report after a service has failed.

Vendor dependency must also remain visible. **The Vendor Dependency and Risk Governance Matrix** tracks service-level performance, third-party cyber risk, contractual remediation triggers, concentration risk, data portability, exit readiness and accountability during multi-vendor incidents. In a complex ecosystem, the health system needs to know who leads when several vendors are involved and the clinical service is down.

Even with the right partner relationship, the health system retains accountability for the outcome. Your team must understand the service, challenge performance, preserve enough internal knowledge to make informed decisions and ensure that partner priorities remain aligned with clinical and organizational needs. Outsourcing an operational function doesn't mean outsourcing clinical trust.

AI readiness becomes continuous readiness

After modernization, pressure to move faster with AI usually increases. Leadership may assume that the new foundation has removed the constraints. Don't let that pressure recreate the same gap between ambition and readiness that modernization was meant to close.

There's a hard boundary between a controlled AI pilot and an enterprise-grade clinical AI service. A pilot may use limited data, limited users, temporary compute, a controlled workflow, manual oversight and a short operating period. But production AI is different. It may operate continuously, serve a much larger population, interact with real clinical or operational workflows, depend on multiple data sources and generate outputs people begin to rely on. That changes the operating responsibility completely.

Production AI requires five disciplines that must be managed together:

1. Infrastructure and data pipelines

A production service needs reliable compute, storage, high-throughput ingestion, data normalization, network capacity and observability. Where clinical response times require it, the service may also need local or edge capability. A model that performs well in a sandbox can fail under production volume, real-time workflow demands or unstable data movement. The infrastructure must support the service continuously. The organization should understand where authoritative data resides, how data moves, what latency the workflow can tolerate, how demand changes, what capacity is required, how the service behaves when a dependency is unavailable and whether the cost remains defensible at scale.

2. Data governance and quality

The organization needs authoritative data sources, named owners, quality standards, ingestion controls, cleansing, normalization and retention rules. It should have processes for identifying missing, incomplete or conflicting data and examine whether the data adequately represents the populations affected by the use case. Poorly governed data doesn't become trustworthy because a sophisticated model processes it.

3. Lineage and traceability

The operating model should preserve evidence about source data, model versions, inputs, outputs, user actions, material changes and performance over time. Not every model can explain every output perfectly. The organization still needs enough traceability to investigate a concern, support clinical and legal review and understand what changed when performance deteriorates. Lineage needs to be part of operational accountability.

4. Clinical, ethical and algorithmic oversight

Oversight should match the risk of the use case. A multidisciplinary group may include clinicians, informatics leaders, data and AI leaders, security, privacy, legal, risk and ethics expertise. The group should review intended use, clinical validity, representativeness, bias, model drift, safety, escalation and conditions for continued use. It should also have authority to limit, suspend or retire a service when the evidence no longer supports its use. Oversight without decision authority isn't true governance.

5. Privacy and security

The operating model should address zero-trust access, de-identification, external model and API use, patient-data exposure, vendor handling of prompts and outputs, auditability, access to underlying data, retention of model inputs and outputs, and security across the full data path. Privacy and security controls must follow the data through the service rather than stop at the boundary of the health system.

Support, safe degradation and cost

Production AI also requires support and safe degradation. Users need to know how to report a concern, who investigates it, how the capability is limited or suspended, what alternate workflow applies, who communicates when an output shouldn't be trusted and how the organization returns to normal operation after the issue is resolved.

The health system must distinguish between an experiment that can be stopped without consequence and a service clinicians or staff have begun to depend on. Cost must be governed as part of readiness. Inference volume, storage, networking, model monitoring, specialist support, external API consumption and data-pipeline operations can change the business case quickly. A pilot that appears inexpensive may become costly at production scale.

AI-enabled services should be reviewed through the same performance, risk, cost and value disciplines applied to every other production service. The objective is to preserve confidence that AI can be operated as responsibly as the infrastructure beneath it.

The goal is a virtuous cycle

A strong operating model allows modernization to build momentum throughout many levels of the organization. Rationalization reduces what the organization must fund, secure and recover. Better workload placement improves performance and cost. Clear accountability makes decisions and incidents easier to manage. Proven recovery reduces uncertainty. Strong operational co-pilot governance improves service quality. Accurate architecture and dependency information reduces surprises. Internal teams gain capacity to improve workflows, resilience, data and innovation. Production AI and other emerging services can be evaluated from a stronger foundation.

From the CIO seat, modernization doesn't eliminate incidents or difficult decisions. Healthcare technology will always carry pressure because the organization depends on it so heavily. But a modernized environment should create fewer surprises, less panic during change, clearer accountability, better evidence, stronger recovery and greater confidence that the organization responds when something goes wrong. It's about building a virtuous cycle that continues to improve over time.



From architect to steward

At the beginning of this journey, I described moving from being a firefighter to becoming an architect. The new operating model adds another responsibility: stewardship.

The firefighter keeps fragile systems alive. The architect creates a deliberate path forward. But the steward makes sure the environment continues to deliver after the architecture is built and the migration team leaves. To me, stewardship means:

- Keeping the environment visible.
- Governing change.
- Proving resilience.
- Managing cost.
- Holding partners accountable.
- Maintaining internal knowledge.
- Reassessing decisions when evidence changes.
- Protecting clinical trust as new services and risks emerge.

Modernization is durable only when those responsibilities become part of the way the organization operates. Becoming a steward ensures that will happen.

While the six-part modernization journey is complete, the Action Plan in Section 7 helps you turn it into a reality. It provides a roadmap, a readiness model and a practical guide to the tools that help CIOs determine what's ready to move, what's not and what must happen before you make the next major decision.

CIO Summary — Section 6

The New Operating Model: From Project to Product

What this means for the CIO

Modernization isn't complete when the workloads move. You need an operating model that continuously governs service performance, clinical experience, cost, security, recovery, architecture, partner accountability, value realization and AI readiness. The test is whether day-two teams operate the environment without relying on the migration team, whether leaders see how services are performing and what they cost, and whether the organization adapts without rebuilding the same uncertainty and technical debt it worked to remove.

Questions to ask now

- Who owns each service after the migration program ends?
- Can day-two teams operate the service without relying on project-team memory?
- Are clinical experience and technical performance reviewed together?
- Is the current-state inventory still accurate as services, dependencies and ownership change?
- Can we explain what drives spending and what value each major service delivers?
- Are complete-workflow recovery and cyber-recovery capabilities tested continuously?
- Are IoMT, identity and vendor-access controls governed as living disciplines?
- Are operational co-pilot and vendor obligations actively managed, including remediation and exit readiness?
- Has internal capacity actually shifted from firefighting toward clinical, data, resilience and strategic work?
- Are AI-enabled services governed as production capabilities with clear ownership, lineage, oversight, privacy, support and safe-degradation procedures?

Helpful tools from the Healthcare Modernization Toolkit:

Use the **Quarterly Governance Review Template** to continuously govern clinical and service performance, resilience, cost, architecture, partner performance and operating ownership after go-live.

Use the **Value Realization Scorecard** to compare actual outcomes with the modernization business case and identify corrective action when expected value isn't fully realized.

Use the **AI Infrastructure Readiness Checklist** to determine whether AI capabilities have the production infrastructure, data, governance, oversight, support and cost controls required to move safely beyond controlled pilots.

Section 7: The Action Plan

APPLYING THE SIX-PART JOURNEY

The first six sections of this playbook explain the operating logic of the healthcare modernization journey. Section 7 converts the logic of those sections into decisions, evidence, owners and readiness gates. This Action Plan has four parts: a six-step roadmap, a ten-dimension readiness assessment, a directory of the companion Healthcare Modernization Toolkit and a practical way to begin.

Each part serves a different purpose. The roadmap defines the sequence of decisions into six straightforward steps. The readiness assessment shows whether the organization has enough verified evidence and operating maturity to make those decisions safely. The toolkit directory lists the 18 working documents needed to close the gaps. And the final section turns the findings into a focused 90-day agenda.

Part I. The Healthcare Modernization Roadmap

The six sections of the playbook become the abbreviated six steps of the roadmap below. Use this to visualize your path and get started on your journey.

1 Make the paradox visible

- **Decision:** Is technical debt materially limiting the organization's ability to support care, financial discipline, resilience, AI and future demand?
- **Required evidence:** A concise risk narrative connecting infrastructure debt to clinical, operational, financial, security and innovation consequences; visibility into the time, funding and skilled capacity consumed by legacy support; and a small number of priority exposures.
- **Accountable participants:** CIO, CTO, clinical leadership, finance, security and executive sponsors.
- **Ready to advance when:** Leaders share a fact-based understanding of what the organization is carrying, what continued delay will cost and why major platform or AI decisions shouldn't proceed from assumption. An accountable sponsor has authorized current-state discovery.

2 Define clinical trust as the standard

- **Decision:** What must the technology foundation do reliably so clinicians can deliver care without avoidable friction or uncertainty?
- **Required evidence:** A view of the workflows most sensitive to latency, downtime and failed recovery; conventional infrastructure measures connected to relevant Clinical Trust Metrics; evidence of workarounds; maximum acceptable care interruption; and shared protection and restoration priorities.
- **Accountable participants:** CIO, CTO, CMIO or equivalent clinical informatics leadership, nursing and operational leaders, security and service owners.
- **Ready to advance when:** Leaders can discuss technical performance and recovery in terms of workflow impact. They agree which capabilities require the highest protection, what degraded operation is acceptable and what evidence will prove that the complete clinical workflow has been restored.

3 Triage and rationalize the current state

- **Decision:** What is operating today, what should remain and what must be stabilized before larger change?
- **Required evidence:** A trusted current-state inventory; ownership and dependency information; clinical and operational classifications; vendor and third-party dependencies; data-retention, archival and interoperability obligations; IoMT dependencies where relevant; rationalization decisions; immediate stabilization actions; and unresolved questions that could alter architecture or migration.
- **Accountable participants:** Technology and application owners, clinical and operational representatives, finance, security, legal, procurement, records or data leaders, biomedical or clinical engineering and vendors where appropriate.
- **Ready to advance when:** The organization can distinguish what should be retired, replaced, consolidated, retained, modernized or migrated. Historical data and interfaces have a defined disposition, the highest-risk weaknesses have been reduced and the remaining unknowns are visible.

4 Design the hybrid future

- **Decision:** Where should each retained workload live, how should it be protected and who will operate it?
- **Required evidence:** A workload-placement model; target architecture; data-gravity analysis; latency and local-survivability requirements; private, managed, public-cloud, SaaS and edge decisions; IoMT and zero-trust controls; disaster- and cyber-recovery design; service commitments connected to clinical impact; partner responsibilities; cost assumptions; and portability or exit considerations.
- **Accountable participants:** CIO, CTO, architecture, security, clinical and application owners, finance, legal, procurement, data leaders and prospective operational co-pilots.
- **Ready to advance when:** The organization can explain why every major workload belongs in its selected environment, how it will be protected and restored, who will operate it, what it will cost and what conditions would cause the decision to be revisited.

5 Move the big rocks safely

- **Decision:** Which workloads should move in each wave and what controls will protect care when conditions change?
- **Required evidence:** A risk-based migration sequence; entry and exit criteria; dependency groups; complete-service testing; clinical change and adoption readiness; vendor readiness; tested runbooks; rollback triggers; decision authority; communication plans; clinical acceptance and stabilization requirements.
- **Accountable participants:** Migration and application teams, clinical and operational owners, security, support, communications, finance, legal or procurement where required, vendors and the operational co-pilot.
- **Ready to advance when:** The workload is ready. The complete service has been tested, clinicians and support teams are prepared, required vendors are ready, decision-makers are available and the organization can pause or reverse the event without losing control.

6 Operate with discipline

- **Decision:** How will the health system sustain clinical trust, cost control, resilience, partner performance, value realization and future readiness after the program ends?
- **Required evidence:** Named service owners; a formal operational handoff; a governance rhythm; Clinical Trust Metrics; forecast-to-actual cost review; value-realization measures; recovery-test evidence; IoMT and zero-trust operating controls; partner and vendor reviews; maintained architecture and dependency information; internal-team role changes; and an operating process for production AI and other emerging services.
- **Accountable participants:** CIO, CTO, service and workload owners, finance, security, clinical and operational leaders, architecture, data and AI leadership, vendor management and the operational co-pilot.
- **Ready to advance when:** The operating organization no longer depends informally on the migration team. Performance, clinical experience, cost, risk, recovery, vendors, value and AI-enabled services are reviewed regularly, and modernization has become an ongoing stewardship discipline.



Part II. Healthcare Infrastructure Modernization Readiness Assessment

The readiness assessment evaluates whether your organization has enough verified evidence and operating maturity to proceed. The CIO and CTO should begin by scoring the dimensions independently. A small cross-functional group should then compare the scores, discuss material differences and agree on a working assessment.

Focus on the lowest dimensions, any score based more on assumption than evidence and the dimensions that function as gates for major change. Clinical trust and reliability, cyber resilience and recoverability, migration readiness, decision authority and operating ownership require acceptable evidence regardless of the total score.

Readiness dimensions

1. **Technical debt visibility**

   Does the organization know where technical debt exists, what clinical, operational, financial, security and workforce risk it creates, how much capacity it consumes and what continued delay will cost?

2. **Application rationalization**

   Does the organization know which applications should be retired, replaced, consolidated, retained, modernized or migrated, and are archival, retention, interoperability and temporary-retention decisions defined?

3. **Clinical trust and reliability**

   Does the organization connect availability, latency, performance and recovery to complete clinical workflows, Clinical Trust Metrics, workarounds and clinician burden?

4. **Cyber resilience and recoverability**

   Can the organization protect and restore critical services during a cyber event through protected copies, isolated recovery, identity restoration, dependency-aware sequencing and validation that systems and data are trustworthy?

5. **Hybrid architecture readiness**

   Can workloads be placed according to clinical criticality, latency, data gravity, local survivability, security, IoMT, cost, portability and operating accountability?

6. **Migration readiness**

   Does the organization have the sequencing, complete-service testing, clinical change readiness, vendor readiness, runbooks, rollback, communication, decision authority, acceptance and stabilization discipline needed to move safely?

7. **Governance and operating model**

   Does the organization maintain ownership, inventories, dependencies, architecture exceptions, Clinical Trust Metrics, vendor-risk governance, value realization and continuous risk review after modernization?

8. **Cost predictability**

   Can leaders explain and govern forecast-to-actual spending across infrastructure, cloud, storage, data movement, backup, licensing, managed services and workload ownership without weakening reliability or recovery?

9 AI readiness

Does the organization have the production infrastructure, data pipelines, quality, lineage, oversight, privacy, security, support, safe-degradation procedures, observability and cost controls required to move AI beyond pilots?

10 Partner execution readiness

Does the organization know where outside expertise or operating capacity is needed, whether partners can execute rather than only advise, how multi-vendor accountability will work and how performance, remediation and exit readiness will be governed?

Maturity levels

Score each of the dimensions listed above from 1 to 5 based on the maturity levels listed below.

Level 1: Reactive

The organization responds to problems as they arise. Visibility is limited, documentation is incomplete and modernization is usually driven by crisis, expiring assets or a major incident.

Level 2: Aware

The organization understands that technical debt and modernization risk exist, but ownership, dependencies, cost, clinical impact, vendor exposure or recovery evidence remain incomplete.

Level 3: Preparing

The organization is actively assessing the environment, classifying systems, engaging stakeholders, stabilizing risk, defining the future state and closing the most consequential readiness gaps.

Level 4: Execution Ready

The organization has a credible roadmap, rationalized workload plan, governance structure, operating and partner model, migration sequence, cost assumptions, clinical readiness and risk controls.

Level 5: Continuously Optimizing

Modernization operates as an ongoing discipline. Clinical trust, cost, reliability, security, recovery, placement, value, vendor performance and AI operations are continuously reviewed and improved while new technical debt is actively prevented.

Scoring and interpretation

Total the ten scores for a maximum of 50 points. The total score is useful for direction and comparison over time, but it isn't a certification and doesn't authorize a specific migration. A health system may score well overall and still be unready to move a mission-critical service if clinical continuity, cyber recovery, decision authority or operating ownership remain weak.

Scores based on assumption should be treated as unverified. Any differences among stakeholders are useful because they reveal unclear ownership or inconsistent understanding. Repeat the assessment as the program progresses so the score reflects current evidence rather than an outdated snapshot.

Score of 10–19: Reactive / High Risk

Begin with discovery, stabilization and stakeholder alignment before major platform or migration decisions.

Score of 20–29: Aware / Early Readiness

Strengthen inventory, ownership, criticality, rationalization, archival planning, recovery evidence and governance.

Score of 30–39: Preparing / Moderate Readiness

Close the most consequential gaps in execution planning, clinical adoption, cost predictability, partner accountability, resilience or migration readiness.

Score of 40–45: Execution Ready

Move into disciplined sequencing and execution while continuing to validate evidence, workload-level readiness, clinical readiness and operating ownership.

Score of 46–50: Continuously Optimizing

Focus on clinical trust, service improvement, value realization, cost management, resilience testing, AI operations, vendor performance and prevention of new technical debt.

Turn the score into an action plan

A readiness score should produce an owner, an action and visibility into the next decision. Use the results to separate three kinds of work:



Immediate stabilization priorities: Present conditions that create risk to care continuity, security or recoverability, such as failing infrastructure, unproven recovery, high-risk IoMT exposure, ownership gaps or unreliable clinical workflows.



Readiness gaps that must be closed before execution: Conditions that may not create immediate exposure but make a future decision unsafe, such as incomplete dependencies, unresolved archival obligations, weak clinical adoption planning, unclear vendor accountability, untested rollback or immature AI governance.



Optimization opportunities: Improvements that increase efficiency, value or maturity after the environment is stable enough to proceed, such as cost refinement, workload re-placement, automation, improved Clinical Trust Metrics and redirection of internal capacity to higher-value work.

Limit the first action plan to the three to five items that most improve care continuity, security, recoverability, decision quality and execution readiness. When you are ready, move to Part IV.

Part III. Healthcare Modernization Toolkit Directory

The companion Healthcare Modernization Toolkit contains the working documents used to put the playbook into action. The tools are organized in the same order as the six-step modernization roadmap. Each tool connects to a roadmap step, closes a readiness gap and enables a specific decision.

Modernization Readiness Assessment

Use at the beginning of the journey and at defined reassessment points. It scores the organization across the ten readiness dimensions and helps a cross-functional leadership group identify the most consequential gaps. It enables agreement on where attention is required before the next major decision.

Technical Debt Reality Check

Use to translate aging infrastructure, unsupported systems, recurring incidents and deferred decisions into clinical, operational, financial, security and workforce risk. It enables an executive-level case for discovery and stabilization.

Modernization Priority Action Plan

Use after the Modernization Readiness Assessment and Technical Debt Reality Check to convert selected gaps into named priorities, owners, target scores, first actions, decisions enabled and evidence of completion. The Readiness Assessment identifies the gaps; the Priority Action Plan turns the selected gaps into accountable work.

Clinical Impact Assessment

Use to connect availability, latency, downtime and recovery to workflows, users, Clinical Trust Metrics and care consequences. It enables clinically grounded criticality, service-level and stabilization decisions.

Current-State Inventory Checklist

Use to create a dependable view of applications, infrastructure, owners, dependencies, vendors, contracts, support status, cost, retention and recovery requirements. It enables evidence-based triage and target-state planning.

Clinical Criticality Matrix

Use to classify systems by care impact, operational importance, acceptable interruption, degraded-mode capability and recovery priority. It enables protection, recovery and sequencing decisions.

Legacy Retirement, Archival and Interoperability Worksheet

Use before retiring or replacing a system. It defines what historical data must remain accessible, how it will be retained or migrated, how users will retrieve it and when the original platform can be shut down. It enables a defensible retirement decision without preserving a costly read-only footprint.

Workload Placement Matrix

Use to compare private or managed environments, public cloud, SaaS, local survivability or edge, and continued placement in the current environment against workload requirements. It enables a defensible placement decision.

IoMT and Zero Trust Security Matrix

Use during design and ongoing governance to assess medical-device connectivity, network segmentation, identities, access, unsupported devices, vendor remote access and zero-trust controls. It enables architecture and risk-treatment decisions across the hybrid environment.

Service Accountability and Recovery Commitments Worksheet

Use to connect accountability, service levels, recovery commitments, escalation, Clinical Trust Metrics and internal-versus-partner responsibilities. It enables clear operating and recovery commitments.

Vendor Dependency and Risk Governance Matrix

Use during design, migration and day-two operations to map vendor dependencies, performance commitments, third-party cyber exposure, remediation triggers, multi-vendor incident authority, data portability and exit readiness. It enables accountable partner governance.

TCO and Cost Predictability Calculator

Use during design and after implementation to compare modernization options with the full cost, operating burden and risk of the status quo, then support forecast-to-actual review. It enables financially defensible placement and operating decisions.

Migration Wave Planner

Use to turn workload decisions into a risk-based sequence with dependencies, readiness, entry and exit criteria, timing, testing, clinical acceptance and stabilization requirements. It enables a defensible wave decision.

Safe-Exit and Rollback Checklist

Use before each migration event to define rollback triggers, authority, decision timing, reversal steps, data reconciliation, clinical contingency and return-to-service validation. It enables the organization to pause or reverse without losing control.

Clinical Change and Adoption Readiness Tool

Use before go-live to assess workflow disruption, training capacity, local champions, support coverage, degraded-mode readiness and clinical sentiment. It enables a clinical-readiness decision rather than relying on technical readiness alone.

Quarterly Governance Review Template

Use after go-live as the recurring agenda and dashboard for clinical experience, service performance, cost, resilience, partner accountability, architecture exceptions and improvement decisions. It enables continuous stewardship.

Value Realization Scorecard

Use at defined post-implementation milestones to track actual cost savings, operational efficiency, clinical throughput, reliability, application retirement, workforce capacity and other outcomes against the original business case. It enables corrective action when expected value isn't appearing.

AI Infrastructure Readiness Checklist

Use before moving an AI capability beyond a controlled pilot. It determines whether the use case has the data, compute, pipelines, lineage, security, oversight, support, observability, cost and suspension criteria required for production. It enables a responsible production decision.



Part IV. How to Begin

The Action Plan is most useful when it leads to a focused working session and a limited set of owner-driven actions. You don't need perfect information before you begin. Missing information, conflicting scores and unresolved ownership are themselves readiness findings.



Before the session

The CIO and CTO complete the Modernization Readiness Assessment independently. They gather the best available current-state evidence, identify which dimensions score lowest, mark any score based on assumption and flag weaknesses involving clinical trust, cyber recovery, migration readiness, decision authority or operating ownership.



During the session

A small cross-functional group compares scores, discusses material differences and agrees on a working assessment. Participants should include the CIO and CTO plus clinical or informatics leadership, security, finance, operations and application or architecture leaders.

The session should produce:

- An agreed working score for each readiness dimension.
- The three most consequential gaps.
- Any immediate stabilization priorities.
- The first toolkit activities.
- Executive and working owners.
- Decisions that must wait for further evidence.



After the session

Create a 90-day stabilization and readiness agenda. This isn't the full modernization plan. Its purpose is to reduce immediate risk, validate unknowns, clarify ownership, improve decision quality and prepare for the next commitment.

For every priority, identify:

- The related readiness dimension.
- Current and target readiness level.
- Owner.
- Why the action matters.
- The first practical step.
- The decision it will enable.
- Evidence of completion.
- The reassessment date.

Don't try to fix every dimension at once. Select the three to five actions that most improve care continuity, security, recoverability, decision quality and execution readiness.

CIO Summary — Section 7

Action Plan

What this means for the CIO

The Action Plan gives you a disciplined way to separate the urgency of modernization from your actual readiness. The roadmap defines the sequence of decisions, the readiness model exposes weak evidence and risk, and the toolkit converts those findings into named actions with owners, decisions and completion criteria.

Questions to ask now

- Which readiness gaps create immediate risk to care continuity, security or recoverability?
- Which scores are supported by verified evidence and which are based on assumption?
- Is a specific weakness being concealed by a strong overall readiness score?
- What three to five actions would most improve our ability to proceed responsibly?
- Who owns each action, what decision will it enable and what evidence will prove it's complete?
- When will we reassess readiness?

Helpful tools from the Healthcare Modernization Toolkit:

Use the Modernization Readiness Assessment to identify the organization's most consequential readiness gaps across the full modernization journey.

Then use the **Modernization Priority Action Plan** to convert the three to five most important gaps into a focused 90-day agenda with accountable owners, next actions, decisions enabled and evidence of completion.

From there, use the appropriate toolkit tools to strengthen the evidence required for each modernization decision.

Section 8: The Conclusion

Modernization is How Care Continues



For me, the most important result of modernization wasn't the new environment itself, although completing the work was deeply satisfying. The more meaningful change was the way the organization operated around it. We could see the environment more clearly. Decisions were easier to explain. Ownership was more explicit. The technology team was no longer forced to begin every day from a position of uncertainty, wondering which fragile dependency or undocumented exception might become the next clinical disruption.

That's what a successful modernization journey should create. It doesn't remove complexity, eliminate incidents or guarantee that the next disruption will be easy. But it does give you a stronger foundation to face those realities without losing control. Leaders have more choices because they understand the environment. Operations act with better evidence and less improvisation. Clinicians focus more fully on care because they're not being asked to absorb the complexity caused by the technology.

The real outcome isn't a finished technology environment but the ability to keep changing without weakening clinical trust.

The foundation beneath the future

As we stated at the beginning of this journey, healthcare leaders are being asked to prepare for AI, advanced analytics, digital care and new models of clinical engagement while still carrying infrastructure, applications and operating practices built for a different era. We defined that tension as the healthcare IT paradox. The pressure to move faster is real, but so is the risk of building the future on a foundation that's poorly understood, difficult to recover and increasingly expensive to support.

The modernization journey isn't about replacing old technology with new. It's about rebuilding confidence that the technology supporting care will perform when clinicians and patients need it most. That confidence comes from operating evidence built on reliable access, predictable performance, protected data, tested recovery, fewer workarounds and clear accountability when something goes wrong.

AI makes the modernization journey more important than ever. But an organization doesn't become ready for clinical AI by purchasing a platform or completing a pilot. It's ready when the operating environment supports the service with trusted data, resilient infrastructure, clear ownership, appropriate oversight and the ability to limit or stop the service safely when the evidence requires it.

See reality clearly

One of the biggest differences I experienced on this journey was the ability to see the environment clearly and honestly. Before the work began, too many decisions were shaped by assumptions, incomplete documentation and knowledge held by a small number of people.

Once we had a clear view of the applications, infrastructure, dependencies, vendors, costs and clinical workflows, the conversations changed. We stopped debating what we thought was true and began making decisions based on what the evidence showed us.

That clarity creates a different kind of confidence in the CIO seat. You explain risk to the executive team without hiding behind deeply technical language. Finance understands why certain investments matter. Clinical leaders see how infrastructure and architecture decisions affect care. Security and operations distinguish immediate exposure from longer-term improvement. The choices may still be difficult, but they're no longer concealed by uncertainty.

Clinical trust remains the measure

The clearest evidence of success is the way clinicians experience the environment after the change. When technology becomes more dependable, it fades into the background. Clinicians don't begin their shift expecting the system to slow down, disappear or force them into another workaround. They focus on patients.

That's why clinical trust remains the standard. Conventional infrastructure metrics matter, but they're incomplete unless they're connected to the clinician's workflow. Availability means little if a clinician can't retrieve the patient context quickly enough to make a decision. Recovery isn't complete if the application is online but identity, imaging, interfaces or downstream processes remain unavailable. A technically successful service can still fail the clinical environment.

At the pediatric health system where I served as CIO, the strongest sign of success was that the transition was invisible to clinicians. The organization changed significantly, but the care continued. The complexity remained where it belonged—with the technology, operational and partner teams—rather than being transferred to the people delivering care. That's the standard healthcare modernization should pursue.

From firefighter to architect to steward

My own journey from firefighter to architect, and ultimately to steward, changed how I understood technology leadership.

The firefighter keeps fragile systems alive. The work is urgent, visible and often heroic. But heroics aren't the same as resilience. An organization that depends on people repeatedly rescuing unstable services hasn't solved the underlying problem.

The architect creates a deliberate path forward. The CIO connects clinical priorities, technology, economics, risk and execution. The organization stops making isolated infrastructure decisions and begins designing a foundation around workload reality, care continuity and accountable operations.



The steward makes sure that foundation continues to deliver after the project ends. Stewardship requires continuous visibility, tested recovery, cost discipline, vendor accountability, maintained architecture and dependency information. It also requires the willingness to revisit decisions when the evidence dictates it. It's the responsibility to prevent the new environment from becoming the next legacy environment.

Modernization is durable when the organization moves beyond completing projects and develops the discipline to govern the environment over time.

Share responsibility without diluting accountability

The path to modernization can't be the responsibility of IT alone. Clinical leadership, operations, finance, security, data and AI leaders, procurement, legal, biomedical teams, vendors and executive sponsors all share this responsibility. The CIO must orchestrate the work because the CIO is one of the few leaders positioned to see how infrastructure, applications, data, workflows, economics and future strategy connect.

But this shared participation makes accountable decisions possible. Clinical leaders define what care can tolerate. Finance understands the economics and consequences of delay. Security proves that critical services can be protected and recovered. Technology leaders make dependencies, tradeoffs and operating requirements visible. Executive sponsors decide which risks the organization will accept and which it will remove. When this shared participation comes together, it makes a stronger organization.

The same principle of shared participation applies to an operational co-pilot. Outside expertise accelerates discovery, architecture, migration, recovery and day-two operations. A capable partner provides skills, scale and execution capacity that the health system can't build alone. But a partner can't own the health system's clinical priorities, risk tolerance or final accountability.

Operational responsibility should be shared, but clinical trust can't be outsourced.

The next responsible action

As the CIO, you have a substantial task ahead of you. Modernization isn't for the faint of heart, especially when every decision must protect care while moving the organization forward. The work requires calculated choices, sustained leadership and the discipline to keep going when progress feels slower than the demands placed on you.

But you don't have to solve everything at once. Begin by assessing your organization's readiness honestly. Identify the three to five actions that most improve care continuity, security, recoverability, decision quality and execution readiness. Assign an owner to each action. Define the evidence that will prove it's complete. Be explicit about the decision it enables. Then reassess as the organization learns.

While delaying may seem like the safest choice, it has its consequences. Technical debt continues to consume funding and skilled capacity. Operational knowledge erodes. Unsupported systems become harder to recover. Clinician frustration accumulates. And AI ambitions move further ahead of the foundation required to support them. The answer isn't to rush into another transformation program but to begin with honest evidence and disciplined priorities.

You may not be able to modernize the entire environment this year. But you can make the risk visible, stabilize what's unsafe and rebuild clinical trust one service at a time.

Section 9: Appendix

From Infrastructure Risk to Clinical Confidence

How a Pediatric Health System Modernized Its Clinical Technology Foundation

This case study illustrates how the principles in the playbook were applied to a mission-critical clinical workload for a pediatric health system in the Pacific Northwest. It consolidates the pediatric proof points referenced throughout the playbook and shows how workload fit, migration discipline, operational support and day-two accountability worked together in practice.

This case study isn't a recommendation that every health system follow the same platform choice. Instead, the value lies in the decision process used to define the clinical requirements, select an environment that fits the workload, control the transition and establish an operating model capable of sustaining the result.

The situation: A mission-critical platform under growing pressure

Epic was the pediatric health system's core electronic medical record and a critical part of patient-care operations. Physicians, nurses and operational teams depended on the platform continuously to access information and complete clinical workflows. Even a brief interruption could affect care delivery, hospital operations and confidence in the technology supporting the bedside.

At the same time, the organization was expanding its services and patient population. Demand on the Epic environment was growing, and the internal technology team was carrying the operational burden of infrastructure management, specialized support and around-the-clock availability. Leaders wanted to redirect more of that capacity toward clinicians, patient-care priorities, research and future innovation.

The challenge wasn't simply to refresh hardware; the health system needed a more dependable foundation and a more sustainable way to operate it without weakening performance, security, recoverability or clinical continuity.

Leadership evaluated potential Epic hosting partners against four practical criteria: availability, security, expertise and cost. Those criteria kept the decision grounded in clinical reality. The organization needed an environment that could deliver predictable performance, protect sensitive information, provide continuous access to specialized Epic expertise and support financial planning over time.

The decision: Choose for clinical reality, not platform preference

The health system selected Rackspace Technology to migrate, host and manage its Epic deployment in a hosted private-cloud environment. The choice reflected the needs of the workload. Epic required stable performance, strict data control, high availability, clear recovery commitments and specialized operational expertise. Those requirements mattered more than a cloud-first mandate.

The target environment was designed as a HIPAA-compliant, HITRUST-certified hosted private cloud with a 99.999 percent uptime commitment, 24x7x365 support from certified Epic specialists and security controls intended to protect sensitive patient data. A fixed-fee pricing model also provided a more predictable commercial structure than a variable consumption-based approach.

This was architecture and operations designed together. The health system and Rackspace worked across network configuration, security controls, performance, redundancy, recovery, support and cost. The operating model treated Epic as a complete clinical service rather than a collection of servers and individual components.

That systems view aligned the technology design with the way clinicians and hospital operations actually depended on the platform. The environment was selected because it fit the workload's clinical, technical, financial and operational requirements.

The execution: Control the change without transferring disruption

Moving a mission-critical clinical platform required more than a technical cutover. The transition had to protect the complete service which included access, performance, integrations, support, recovery and the clinical workflows that depended on them.

Preparation for the migration focused on readiness and control.
The health system and Rackspace coordinated the technical sequence, vendor responsibilities, communications, validation and decision authority required for the event. Rollback remained available as a safety mechanism, with the ability to pause or reverse the migration if conditions threatened care continuity.

The transition concluded four weeks ahead of schedule. More important, clinical staff later asked whether downtime had actually occurred. That meant the migration didn't transfer the complexity of the change to clinicians because the technology, operational and partner teams absorbed it behind the scenes.

The health system also reported zero emergency cutover rollbacks. That result didn't make rollback planning unnecessary but it showed that the event was prepared and controlled while a safe-exit option remained available if the evidence required it.

The broader modernization program also applied radical asset triage. Leadership challenged which applications still needed to be retained, which could be retired and where unnecessary operating cost could be removed. That work contributed to a multimillion-dollar reduction in legacy application maintenance costs and reduced the risk of reproducing the old environment on newer infrastructure.



The operating model: Extend accountability beyond cutover

The engagement didn't end when the migration finished. Rackspace assumed responsibility for ongoing Epic infrastructure management and provided continuous access to certified specialists. That allowed the internal technology team to focus more heavily on clinicians, patient-care priorities, research and strategic improvement rather than carrying the full operational burden of the underlying platform around the clock.

The fixed-fee model also improved cost predictability. Healthcare leaders often discover that a technically sound service becomes difficult to govern when consumption, storage, network and support costs are unclear. A predictable commercial model gave the organization a more defensible basis for planning while preserving defined service commitments.

The relationship reflected the operational co-pilot model described in this playbook. Rackspace provided embedded expertise, execution and day-two management. The health system retained ownership of clinical priorities, risk decisions, architecture standards and final responsibility for care continuity.

Outside capacity absorbed operational complexity, but it didn't supersede the health system's mission. That distinction is essential because while operational responsibility can be shared, clinical trust can't be outsourced.

The results: Clinical trust supported by operating evidence

The transformation produced evidence across five dimensions, including what was delivered, what users experienced and what organizational capacity changed.



Clinical continuity. The migration was effectively invisible to clinicians. The platform remained available at the point of care and the organization completed the transition without an emergency rollback. A reported +48 end-user IT Net Promoter Score provided a broader indicator of user confidence in the technology organization.



Reliability and performance. The new environment was described as responsive and reliable, with no lag. The 99.999 percent uptime commitment supported continuous access to critical patient information. The organization also reported shorter EHR image-retrieval times, a useful clinical performance indicator.



Operational capacity. Moving Epic infrastructure management to a specialized operational co-pilot reduced the burden on the internal team. Technology leaders could redirect expertise toward clinicians, patient care, research and strategic improvement rather than sustaining the underlying environment around the clock.



Financial discipline. The fixed-fee operating model improved predictability. Separately, radical asset triage and application rationalization achieved a multimillion-dollar reduction in legacy application maintenance costs.



Future readiness. The modernized foundation expanded the organization's ability to consider capabilities beyond the immediate Epic environment, including high-performance computing for research, from a more stable, secure and operationally confident base.

Six lessons for healthcare CIOs

While each modernization project is unique, here are a few lessons that may apply to your environment and current situation.

1. Begin with clinical requirements.

Availability, security, expertise, cost and point-of-care access provided a more useful decision framework than platform ideology. The selected environment fit the workload because it reflected how the clinical service needed to perform.

2. Evaluate the complete service.

Epic couldn't be separated from networking, security, identity, redundancy, recovery, specialized support, cost and day-two operations. Architecture decisions became stronger when the organization evaluated the clinical system as a whole.

3. Rationalize before you migrate.

Radical asset triage reduced the risk of carrying unnecessary applications, dependencies and maintenance costs into the future state. Moving everything would have preserved some of the legacy burden rather than removing it.

4. Treat migration as a clinical event.

Detailed planning, testing, communication, decision authority, rollback readiness and clinical validation protected the care environment. The most meaningful evidence was that clinicians didn't experience the operational complexity behind the change.

5. Select for day-two capability.

A migration partner can help complete a project but an operational co-pilot also supports the environment after the project team leaves, provides specialized expertise, meets service commitments and helps the organization continue improving.

6. Measure success through clinical trust.

The infrastructure changed, but the real outcome was dependable access, responsive performance, fewer operational burdens and greater confidence in the technology supporting care. That's the standard every modernization program should be designed to achieve.



Want the toolkit? Go to this page to get it.

The playbook shows you what to decide. The Healthcare Modernization Toolkit gives you the working tools to decide it: 18 tools across the six steps of the journey, from the first readiness check to the AI production gate.

Every tool is built to end the same way: a decision, a named owner, a next action, evidence of completion and a review date. Use them in order, or start with the gap that worries you most.

Get the toolkit free:

https://go.rackspace.com/lp/cp/healthcare-cloud-epic-hosting





About Rackspace

Rackspace® is the operator of the full enterprise AI stack from governed private cloud to AI inference and agents in production. With an Outcomes-as-a-Service model built on secure infrastructure, data foundations and forward-deployed engineering, we deliver business results for regulated and mission-critical industries where governance, sovereignty and uptime are non-negotiable. Learn more at www.rackspace.com.

© 2026 Rackspace US, Inc. :: Rackspace®, Fanatical Support®, Fanatical Experience® and other Rackspace marks are either service marks or registered service marks of Rackspace US, Inc. in the United States and other countries. All other trademarks, service marks, images, products and brands remain the sole property of their respective holders and do not imply endorsement or sponsorship.

Modernization is How Care Continues

Modernization isn't a cloud migration, an IT refresh or a race to buy shiny new AI tools. It's the deliberate work of building a technology foundation that clinicians can trust, executives can govern, CFOs can defend and security teams can protect. The goal isn't new technology for its own sake. It's to strengthen the foundation supporting clinical care today while giving the organization the resilience and flexibility to support what comes next.

The six-part modernization journey

- Make the paradox visible
- Define clinical trust as the standard
- Triage and rationalize the current state
- Design the hybrid future
- Move the big rocks safely
- Operate with discipline

About the author



Zafar Chaudry

MD, MS, MIS, MBA, CHCIO, CDH-E, FCHIME
Chief Digital Health Officer, Rackspace Technology

Dr. Zafar Chaudry is a globally distinguished C-suite executive with a career spanning nearly four decades at the intersection of clinical practice, corporate leadership, and digital innovation.