


The Adoptiv Method™

Turning user-built AI into enterprise-grade, adopted systems — assembled from a governed library, shipped with controls, engineered to be used.



Philip Morales · VP, Healthcare AI & Digital Transformation

Digital Intelligence Advisory Services (DIAS™)

Version 1.4 (Public Edition) · August 2026

Public edition — cleared for external distribution.

CONTENTS

Abstract.....	3
Figure 1 — The Adoptiv Method™ ecosystem.....	4
1 The problem: value stalls twice.....	5
2 The prototype-to-production gap.....	5
3 The Adoptiv Method™: four stages, one motion.....	5
4 The reference library: assemble, don't reinvent.....	6
5 Guided assembly: meeting each user where they are.....	6
6 Regulatory controls and governance.....	6
7 Platform and technical foundations.....	9
8 The adoption engine.....	9
9 One discipline, applied across regulated workflows.....	10
10 Benefits and the metrics that prove them.....	10
11 Tooling across the lifecycle.....	10
12 Conclusion: from risk to moat.....	11
References and reference-validation statement.....	12

DOCUMENT CONTROL

Field	Detail
Document	The Adoptiv Method™ — White Paper
Version	1.4 (Public Edition) · v1.4-P
Date	31 August 2026
Supersedes	Version 1.3 (Public Edition), 27 August 2026, v1.3-P
Classification	Public — cleared for external distribution
Changes in this edition	The companion guide referenced in Sections 6 and 12 is renamed HAGAP™ (Healthcare AI Governance and Adoption in Practice), and the request subject line updated to match. The cover and document-control bylines are corrected to VP, Healthcare AI & Digital Transformation, matching the byline used across the series. Figure 1 is rebuilt: the retired GDS mark removed, the version stamp and byline aligned to the cover, and the validation banner restated as eleven sourced claims — nine verified, two as-reported — matching the reference-validation statement. The Section 10 pilot-to-production target is restated as an absolute rate. One source listed but not cited in the text is removed. No change to the method or to any sourced claim.
Author	Philip Morales · VP, Healthcare AI & Digital Transformation · Digital Intelligence Advisory Services (DIAS™)

ABSTRACT

Enterprises are investing billions in artificial intelligence, yet most of that investment stalls twice — once before production, and again before adoption. MIT NANDA's 2025 study *The GenAI Divide* found that 95% of generative-AI pilots deliver no measurable P&L impact. S&P Global Market Intelligence reports that 42% of companies abandoned most of their AI initiatives in 2025 — up from 17% a year earlier, while the average organization scraps 46% of proofs-of-concept before they reach production. Gartner puts the average prototype-to-production cycle at roughly eight months, with only 48% of AI projects making the crossing at all. The barrier is not ambition or tooling. It is the absence of a repeatable, governed path from what a domain user builds to a production-grade system that people actually adopt. The Adoptiv Method™ is that path: a four-stage discipline — Build, Industrialize, Governed Deploy, Adopt & Scale — in which users assemble from a governed library of pre-approved agents and agentic flows, engineering hardens them through a quality gate, governance ships them with regulatory controls enforced on every component, and a structured adoption engine turns deployment into sustained usage. Each cycle compounds: every hardened deployment sharpens the reference library so the next build starts further ahead.

In healthcare, that path now has an external standard to meet. The Joint Commission and the Coalition for Health AI published the first national framework for responsible health AI use in September 2025, CHAI released its governance playbooks in May 2026, and the Joint Commission opened a voluntary certification in June 2026. The gap those instruments address is stark: CHAI reports that 72% of health systems are already using or piloting AI tools while fewer than 30% have a formal AI governance framework in place. Section 6 maps the method against all three.

FIGURE 1 — THE ADOPTIV METHOD™ ECOSYSTEM



The loop, the governed library, and the controls that make AI adoptable — at a glance.

1 The problem: value stalls twice

Two failure points define the enterprise AI landscape. The first is the prototype-to-production gap — impressive demos that never survive the journey to a governed, reliable system. The second is the production-to-adoption gap — systems that ship but go unused, quietly abandoned by the very teams they were built for.

The evidence is consistent across issuers. MIT NANDA finds that 95% of GenAI pilots produce no measurable P&L impact, and that only 5% of integrated systems create significant value. S&P Global Market Intelligence, surveying more than 1,000 enterprises across North America and Europe, reports that the share of companies abandoning most of their AI initiatives more than doubled in a single year — from 17% to 42% — with 46% of proofs-of-concept scrapped before production. Gartner's data completes the picture: an eight-month average prototype-to-production cycle and a 48% production rate.

Healthcare carries a second version of the same problem. CHAI reports that 72% of health systems are already using or piloting AI tools while fewer than 30% have a formal AI governance framework in place, against more than \$4.5 billion in healthcare AI market spending in 2025. Capability has outrun control, which is why the governance instruments described in Section 6 arrived when they did.

These are not model-quality problems. They are operating-model problems — the absence of a standard path, shared guardrails, and a deliberate approach to adoption. The Adoptiv Method™ is built to close both gaps by design.

2 The prototype-to-production gap

Prototyping has been democratized. Low-code, prompt, and agent-builder tools let domain users stand up a working demo in days, capturing hard-won workflow knowledge directly from the people who own the problem. Production is a different discipline entirely: guardrails, evaluations, security, compliance, integration, reliability, monitoring, and governance. Between the two sits the gap where most AI dies — built, but never shipped, and never adopted.

Closing the gap is not a matter of better prototyping tools or a bigger model. It requires an explicit, repeatable route from the user's build to a governed system, and then a second route from that governed system to real, measured adoption. The Adoptiv Method™ supplies both, with clear ownership at every step.

3 The Adoptiv Method™: four stages, one motion

The method organizes work into four stages, each with an accountable owner. The business user prototypes; engineering productionizes; governance ships; champions drive adoption. A quality gate sits before deployment, and a reuse loop runs after it — feeding reusable patterns, evaluations, and guardrails back into the library so every subsequent build starts further ahead.

Stage	Owner	What happens
01 Build	Business + domain experts	Domain users assemble a functional prototype from pre-approved agents, prompts, and templates in low-code and agent-builder tools — capturing workflow knowledge in days, not quarters.
02 Industrialize	Platform engineering	Guardrails, evaluation, security, integrations, observability, and scale — the rigor that turns a demo into a dependable system.
03 Governed Deploy	Engineering + business	Production rollout with monitoring, compliance, tenant isolation, human-in-the-loop, and SOC 2 controls — auditable from day one.
04 Adopt & Scale	Business + champions	Drive real usage, measure value, and expand across teams — the stage where deployment becomes adoption and ROI.

Table 1 — The four stages of the Adoptiv Method™, with accountable owners.

The full process runs Build → Industrialize → Validate → Deploy → Adopt & Measure, with the quality gate between Validate and Deploy. Because each hardened deployment sharpens the reference library, the discipline compounds: the next workflow ships faster, safer, and adopts quicker than the last.

4 The reference library: assemble, don't reinvent

At the center of the method is a governed reference library. Users do not build from a blank canvas; they assemble from approved parts. The library is what makes speed and safety compatible rather than opposed.

Pre-approved agents

A catalog of vetted, versioned agents — the building blocks users assemble from. Each agent is documented, evaluated, and carries a known cost/performance profile, so a domain user can compose a solution without re-deriving safety or quality from scratch.

Governed agentic flows

Pre-approved, multi-step agentic workflows available in the library and ready to configure. Rather than wiring a novel agent chain by hand, a user adopts a proven flow — intake-and-triage, extract-and-validate, resolve-and-retrieve — and tailors it to the workflow at hand. The flow arrives with its guardrails, escalation paths, and human-in-the-loop checkpoints already defined.

Standards, patterns, and enablement

Accessibility standards (WCAG), prompt-engineering and agent-design patterns, human-in-the-loop and disclosure conventions, and naming, versioning, and documentation standards travel with every component. Playbooks, golden-path examples per user profile, evaluation and red-team suites, and governance checklists round out the library so teams inherit institutional judgment instead of rediscovering it.

5 Guided assembly: meeting each user where they are

A governed library only delivers if users can navigate it. The Adoptiv Method™ guides each user from prompt to production through golden paths, in-context help, and guardrails that lighten as capability grows. Three profiles anchor the experience:

- Novice — consumes. Little technical background; uses AI through finished interfaces and guided templates. Needs simple guided UX, strong guardrails, and in-context help.
- Educated dabbler — configures. An AI-literate domain professional who tweaks prompts and assembles prototypes in builder tools. Needs a safe sandbox, reusable components, and a clean handoff to engineers.
- Power user — builds. Technically fluent, with deep domain and tooling expertise; builds sophisticated workflows, chained agents, and evaluations. Needs governed production access, APIs, the agent library, and version control. Autonomy grows and guardrails lighten as capability increases — the same governed library serves the novice assembling a template and the power user composing a chained-agent workflow, each guided to the appropriate boundary.

6 Regulatory controls and governance

In regulated industries, controls cannot be bolted on after the fact. The Adoptiv Method™ maps domain frameworks into the library so that compliance is enforced on every agent and flow before it ships, not audited afterward. Until recently that claim had to be made against general frameworks, because healthcare had no AI-specific standard of its own. It now does — and it comes with a certification, which turns the claim from a statement of good practice into something an organization can be assessed on.

The healthcare baseline: RUAIH, the playbooks, and certification

Three things arrived in sequence. In September 2025 the Joint Commission and the Coalition for Health AI (CHAI) published Guidance on the Responsible Use of AI in Healthcare (RUAIH), the first national framework of its kind in the United States. In May 2026 CHAI released governance playbooks defining the baseline controls behind it. In June 2026 the Joint Commission opened a voluntary RUAIH certification, available to any healthcare organization whether or not it holds Joint Commission accreditation. The scope is wider than most readers assume. RUAIH defines health AI tools as clinical, administrative and operational solutions applying algorithmic methods — predictive, generative or combined — across direct and indirect patient care, care support services, and care-relevant operations. Revenue cycle management,

coding, prior authorization and care quality management are named explicitly. This is not a clinical-decision-support framework; it reaches the administrative workflows where most enterprise AI actually lands.

One boundary matters for any platform vendor. The certification is organization-level and does not validate or certify individual products or tools. A method cannot be certified. What it can do is make the organization running it certifiable — a narrower claim, and a defensible one.

The seven RUAIH elements

#	Element of Responsible Use of AI in Healthcare
1	AI Policies and Governance Structures
2	Patient Privacy and Transparency
3	Data Security and Data Use Protections
4	Ongoing Quality Monitoring
5	Voluntary, Blinded Reporting of AI Safety-Related Events
6	Risk and Bias Assessment
7	Education and Training

Table 2 — The seven elements, as published by the Joint Commission and CHAI, September 2025.

The eight CHAI playbook domains

RUAIH sets seven elements of responsible use. The CHAI playbooks operationalize them, and the architecture is worth stating precisely because it is widely misreported. CHAI publishes four core domains — AI Policy, Organizational Structures, Organizational Resources, and Organizational Processes — with five subdomains beneath Organizational Processes. The eight elements commonly cited, and used in the mapping below, are the first three domains plus those five subdomains. CHAI itself presents both framings; the playbook tools are numbered against the four-plus-five structure. The playbooks were developed with more than 150 health AI leaders across more than 100 organizations and provide the framework for achieving the certification.

- Domain 1 — AI Policy
- Domain 2 — Organizational Structures
- Domain 3 — Organizational Resources
- Subdomain 4.1 — Responsible AI Lifecycle Management and Use
- Subdomain 4.2 — Risk and Impact Assessments
- Subdomain 4.3 — Responsible Data Management and Use
- Subdomain 4.4 — Third Party Management
- Subdomain 4.5 — Education, Training, and Feedback Two of the eight have no direct RUAIH element. Organizational Resources and Third Party Management are additions rather than restatements, and they are the two most likely to catch an organization that believed it was already compliant.

How the four stages map

Adoptiv stage	RUAIH elements	CHAI domains	Where the method does the work
01 Build	6 Risk and Bias Assessment	AI Policy; Risk and Impact Assessments	Assembly from pre-approved parts moves risk and bias assessment to library admission, so it happens once per component rather than once per prototype
02 Industrialize	4 Ongoing Quality Monitoring; 3 Data	Responsible AI Lifecycle Management and Use	Guardrails, evaluations, security and observability are the hardening layer; monitoring hooks are

Adoptiv stage	RUAH elements	CHAI domains	Where the method does the work
	Security		built in rather than retrofitted
03 Governed Deploy	1 Governance; 2 Privacy and Transparency; 5 Event Reporting	Organizational Structures; Responsible Data Management and Use	The quality gate is the governance body exercising risk-based oversight, with disclosure and an event-reporting route as conditions of shipping
04 Adopt & Scale	7 Education and Training; 4 continued monitoring	Education, Training, and Feedback	The adoption engine is the education element instrumented, with usage measured rather than assumed
The reference library	1, 3 and 6 jointly	Third Party Management; AI Policy	A governed catalog is how a component is assessed once and reused under a standing control instead of re-litigated per project
Partially addressed	—	Organizational Resources	Four of the domain's five areas are engaged — staffing and workforce planning, tooling investment, sustaining governance over time, and shared resources. Budgeting is not, and the method does not claim it

Table 3 — Stage-to-element mapping. Authors' analysis; not a Joint Commission or CHAI statement.

The mapping above is the authors' analysis, not a Joint Commission or CHAI statement, and is offered as a working view rather than a compliance opinion. (Judgement.) Two features of the guidance align particularly closely with the method's structure. RUAH asks that post-deployment monitoring be risk-based and scaled to setting — the question it poses is how close a tool sits to patient care decisions and what happens if it performs poorly, which is the same graduated logic the reference library applies at admission. And it names CHAI's Applied Model Card as the instrument for capturing vendor risk, bias and limitation information, which is precisely the artifact a library entry should carry.

Organizational Resources: a partial answer, and where it stops

The Organizational Resources domain covers budgeting for AI governance, staffing and workforce planning, technology and tooling investments, sustaining governance over time, and leveraging shared resources and consortia. It exists because early digital-health standards over-indexed on organizations with large budgets, leaving regional and community providers without an actionable path. It is the domain most likely to expose an organization that has written a policy it cannot staff. A methodology cannot supply budget, and this paper does not claim otherwise. What it can do is change the shape of the requirement, and it reaches four of the domain's five areas. (Judgement.)

On staffing and workforce planning, the three user profiles in Section 5 are a workforce model: they say what capability is needed at which tier, and assembly from a governed library moves the first stage onto domain users who already exist, so scarce platform engineering is spent on hardening rather than on building each prototype from nothing. On tooling investment, Section 11 maps the categories by lifecycle phase. On sustaining governance over time — the area most organizations underestimate — the compounding library is the mechanism: because controls attach to components at admission rather than to projects at review, a governance body of fixed size covers a growing portfolio instead of re-reviewing every project from scratch. And on shared resources, a governed reference library is itself a shared-service model, which is how a multi-entity operator gives three business units one control set without building governance three times.

Budgeting is the exception, and it is not a small one. An organization with no governance function to make more efficient, and no route to fund one, is not helped by a method that makes review cheaper. That is an organizational question, and it should be answered before a methodology is selected rather than after.

The same treatment is applied to all eight elements in HAGAP™ — Healthcare AI Governance and Adoption in Practice — which sets out for each one what the framework requires, what it leaves to the organization, and what works in practice — including a cross-cutting section on change management and adoption. It is provided on request rather than published: write to philip@philipmoralesai.com with the subject line “HAGAP”.

The wider control set

RUIAH and the playbooks sit above, not instead of, the component-level regulation the library already enforces:

- Healthcare — HIPAA, HITECH, HITRUST, SOC 2, FDA SaMD/PCCP, and ONC, supported by controls tooling such as Vanta, Drata, ComplyAssistant, Censinet, Paubox, and Protenus. The Joint Commission's Responsible Use of Health Data (RUHD™) framework applies where data is used secondarily.
- Financial services — SOX, PCI-DSS, GLBA, SEC/FINRA, BSA/AML, and SR 11-7 (model risk), supported by NICE Actimize, ComplyAdvantage, Archer, LogicGate, AuditBoard, and Kosli.

Spanning both domains, AI governance tooling — Credo AI, Holistic AI, OneTrust AI Governance, and IBM watsonx.governance — is aligned to the NIST AI Risk Management Framework and the EU AI Act for model risk, bias, and auditability. A shared governance board sits at the production gate, where business value, IT risk, and engineering readiness are signed off together. Compliance agents run as part of the quality gate, checking each agent and flow against the applicable framework and producing an immutable audit trail from day one.

7 Platform and technical foundations

The method rests on a platform designed to industrialize AI safely and repeatably. Several design principles are applied by default rather than by exception.

Data aggregation and semantic layer

A dedicated aggregation and semantic layer sits between raw data and the intelligence and experience layers. Curated aggregate marts, consistent metric and measure definitions, and cross-entity roll-ups give every agent, flow, and dashboard a single, governed source of truth — so a figure means the same thing whether it surfaces in an agent's reasoning, a compliance check, or an executive view. This layer is what lets adoption and value be measured consistently across workflows.

Core design principles

- Standards-first interoperability with a canonical data model, and a clean separation of ingestion from business logic.
- Medallion data architecture with data quality, lineage, and master-data management as first-class concerns.
- Event-driven and batch duality so real-time and scheduled workloads share the same governed foundation.
- Zero-trust security — least privilege, strong identity, and encryption in transit and at rest.
- Governance by design — catalog, classification, retention, and AI/model governance, with human-in-the-loop autonomy calibrated to risk.
- Observability, scalability, and resilience throughout, with deploy-anywhere flexibility across cloud, on-premises, and hybrid, and explicit regulatory mapping.

8 The adoption engine

Deployment is not the finish line — adoption is. The Adoptiv Method™ treats adoption as an engineered outcome, aligned to the Prosci ADKAR change-management model: Awareness → Desire → Knowledge → Ability → Reinforcement moves each user from first exposure to sustained habit.

- Champions and a Center of Excellence sustain capability, consistency, and reuse across the organization through a certified champion network.
- Phased rollout starts with one workflow that works, then expands — reaching materially higher adoption than big-bang launches.
- Value realization instruments usage from day one and ties every deployment to a measured business outcome, feeding learnings back into the build.

Adoption is measured, not assumed: active-user rate, the share of target workflows adopted, time-to-value, and sustained usage at 90 days.

9 One discipline, applied across regulated workflows

The same repeatable process applies across verticals, with domain-specific guardrails, integrations, and compliance controls.

- Healthcare — prior-authorization automation. A clinical champion builds an intake-and-criteria agent; the platform adds EHR integration and an audit trail; the flow is validated and deployed under governance — yielding faster determinations, fewer manual touches, and a full audit trail. RUAIH names prior authorization explicitly within its scope, so this workflow sits inside the governance perimeter, not beside it.
- Financial services — KYC and onboarding. An operations power user builds an extraction-and-risk-flag workflow; the platform adds AML controls and logging; validated and deployed — yielding faster onboarding, consistent compliance, and a clear audit trail.
- Customer experience — tier-1 service automation. A support lead builds a resolution-and-retrieval agent; the platform adds CRM integration and human handoff; validated and deployed — yielding higher deflection, faster resolution, and agents freed for complex cases.

10 Benefits and the metrics that prove them

The method is designed to be measurable. Targets below are illustrative planning baselines benchmarked to the third-party baselines cited in this paper; they vary by workflow and are not guarantees.

Metric	Adoptiv Method target	Industry baseline
Time to production	Weeks	~8 months (Gartner)
Pilot → production rate	80%+ of piloted workflows reach production	48% reach production (Gartner)
Cost per deployment	↓ 30–50% by the 3rd build, via reuse	—
Catalog reuse	50%+ of new builds from approved parts	—
Active adoption	70%+ of target users, sustained 90 days	—
Time-to-value	< 90 days from deploy to measured ROI	—
Certification readiness (healthcare)	Evidence produced as a by-product of the gate, not assembled for an audit	RUAIH certification, five standard areas

Table 4 — Representative KPI targets versus industry baselines.

The last row is a change in kind rather than degree. Because the quality gate produces an immutable audit trail, and because controls attach to components rather than projects, the evidence a certification review asks for is generated in the course of shipping rather than reconstructed afterward. (Judgement. The mechanism follows from the method's structure; no organization has yet been certified using it.)

Beyond the numbers, the compounding structure is the durable advantage. The business user prototypes and the engineering leader productionizes, so the method serves both the champion and the economic buyer in a single motion. Every hardened deployment sharpens the library, adoption is engineered in rather than hoped for, and the whole system is instrumented on a dashboard — turning the prototype-to-production gap from an enterprise's biggest risk into its moat.

11 Tooling across the lifecycle

The method is tool-agnostic but opinionated about categories. Representative examples by phase:

- Build (for users) — Copilot Studio, Vertex AI Agent Builder, Bedrock AgentCore, Dify, prompt/agent canvases, and a test-data sandbox.
- Industrialize (for developers) — LangGraph, CrewAI, OpenAI / Claude Agent SDKs, GitHub and MLflow, CI/CD, and LangSmith / Langfuse.

- Validate & QA (AI-native testing) — agentic QA platforms such as TestMu AI, Mabl, Functionize, Applitools, testRigor, and Braintrust / Ragas. As AI writes code at machine speed, quality engineering must keep pace across web, mobile, and AI applications.

12 Conclusion: from risk to moat

The enterprises that will capture value from AI are not those with the best models, but those with a repeatable, governed path from build to adoption. The Adoptiv Method™ supplies that path: a governed library of pre-approved agents and agentic flows, guided assembly for every user, regulatory controls enforced on every component, and an adoption engine that turns shipped systems into sustained usage.

In healthcare that path now has a destination it can be measured against. HAGAP™ takes each of the eight elements in turn and sets out what the framework requires, what it leaves to the organization, and what works in practice; it is available on request at philip@philipmoralesai.com. The constraint was never ambition or tooling — it was the missing path. The Adoptiv Method™ closes the gap, and compounds a little further ahead with every cycle.

References and reference-validation statement

Reference-validation statement

Tier 1 — Externally sourced statistics and framework facts. This paper carries 11 externally sourced items: 9 verified against the issuing organization's own publication, 2 as-reported. Verified: 95% of GenAI pilots deliver no measurable P&L impact (MIT NANDA, 2025); 42% of companies abandoned most AI initiatives, up from 17%, and 46% of proofs-of-concept scrapped before production (S&P Global Market Intelligence / 451 Research, 2025); the ~8- month prototype-to-production cycle with a 48% production rate (Gartner, 2024); the seven RUAIH elements, the definition of health AI tools, and the risk-based monitoring test (Joint Commission and CHAI, 17 September 2025, retrieved in full); the eight playbook elements, their development basis and their role as the certification framework (CHAI, 27 May 2026); the four-core-domain and five-subdomain architecture (CHAI AI Governance workgroup); the published coverage of the Organizational Resources domain (CHAI, AI Governance in Community Health, July 2026); the certification launch date (Joint Commission, 1 June 2026); and CHAI's Creative Commons licence terms. As-reported: the five certification standard areas, the statement that applicants need not be Joint Commission accredited, and that the certification does not validate individual products, all taken from Fierce Healthcare of 2 June 2026 because the Joint Commission certification page could not be retrieved directly; and the figures of 72% of health systems using or piloting AI, fewer than 30% with formal governance, and \$4.5 billion in 2025 healthcare AI spending, which CHAI states without onward citation.

Tier 2 — Frameworks, standards, and tools. Named frameworks (HIPAA, HITECH, HITRUST, SOC 2, FDA SaMD/PCCP, ONC, RUHD™, SOX, PCI-DSS, GLBA, SEC/FINRA, BSA/AML, SR 11-7, NIST AI RMF, EU AI Act) and vendor tools are real and current, cited as representative of their category rather than as an exhaustive or endorsed list.

Tier 3 — Method targets and analysis. Adoptiv Method™ KPI targets and outcome ranges are illustrative planning baselines benchmarked to the Tier 1 sources. They vary by workflow and are not guarantees. The stage-to-element mapping in Section 6, the Organizational Resources position, and the certification-readiness claim in Section 10 are the authors' analysis, marked "Judgement" inline, and are not endorsed by the Joint Commission or CHAI.

Use of third-party material. CHAI publishes its frameworks, playbooks and tools under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 licence. This paper describes and attributes those materials in the authors' own words and reproduces no CHAI table, worksheet, tool or template. Readers should obtain the playbooks from CHAI directly.

Sources

- MIT Project NANDA (MIT Media Lab). The GenAI Divide: State of AI in Business 2025.
- S&P Global Market Intelligence / 451 Research. Voice of the Enterprise: AI & Machine Learning, 2025.
- Gartner. Survey: Generative AI Is the Most Frequently Deployed AI Solution in Organizations (May 2024) — 48% production rate; ~8-month prototype-to-production cycle.
- Joint Commission and Coalition for Health AI. Guidance on the Responsible Use of AI in Healthcare (RUAIH), 17 September 2025.
- Coalition for Health AI. Comprehensive Governance Playbooks release, 27 May 2026, chai.org — eight domains of baseline controls.
- Joint Commission. Responsible Use of AI in Healthcare certification, launched 1 June 2026. Five certification standard areas as reported by Fierce Healthcare, 2 June 2026.
- Coalition for Health AI. AI Governance workgroup page, chai.org — four core domains with five subdomains under Organizational Processes.
- Coalition for Health AI. AI Governance in Community Health: A Practical Roadmap from CHAI, July 2026 — per-domain coverage and getting-started actions; adoption and governance figures; CC BY-NC-ND 4.0 licence terms.
- Prosci. ADKAR Change-Management Model.
- NIST. AI Risk Management Framework (AI RMF 1.0).
- European Union. Artificial Intelligence Act (Regulation (EU) 2024/1689).

HAGAP™, The Adoptiv Method™, Vitruvian Health™, Compass IQ™, and DIAS™ are trademarks of Digital Intelligence Advisory Services. RUAIH™ and RUHD™ are trademarks of Joint Commission enterprise. © 2026 Digital Intelligence Advisory Services. All rights reserved. This public edition may be shared and quoted with attribution; it may not be altered or resold.