

Creation and Control of a Digital Workforce

Why healthcare organizations should define, create, and run their Digital Workforce on compute they control

Fabio A. Thiers, MD, PhD
ConverGen AI, New York, NY, USA

ABSTRACT

AI systems are shifting from providing capability to performing work, and from products an organization buys to Digital Workers it must define and manage. A Digital Worker needs a clear, organization-specific definition, governing knowledge, access, authorization, supervision, and eventual retirement. Those are not product features. They are management acts belonging to the organization whose work is performed. They also depend on local context and accumulated organizational memory, material sensitive enough that it should stay inside the organization. This article argues that healthcare organizations should build and manage their Digital Workforce themselves, on infrastructure they control, while AI models and commodity tooling remain interchangeable. This approach does not require writing software, but it does require a Digital Operations System (DOPs) that the organization runs, enabling it to create each Digital Worker, govern its knowledge and data, and operate within a boundary that can be inspected, modified, or revoked. The article also outlines seven capabilities organizations should retain, a sourcing continuum in place of a build-or-buy binary, and how a role-based management system lets models change without redefining the work.

Keywords: healthcare AI; Digital Operations; Digital Operations System (DOPs); Digital Workers; Digital Labor; capability ownership; sourcing boundary; model routing; execution locality; organizational memory; prior authorization; AI governance.

Key points

- **A Digital Worker is a role, not a product.** It has to be defined, given governing knowledge, granted access, authorized, supervised, and retired. Those are management acts belonging to the organization whose work is performed.
- **Specification control and operating control are different.** A supplier can offer genuine control over how a role is defined while the role logic, context, and evidence stay inside a boundary the organization does not govern.
- **Organization-controlled does not mean writing software.** The organization defines each role, governs the knowledge and data behind it, and runs the work on a Digital Operations System (DOPs) installed on its own premises or in a cloud capacity held under its own tenancy and terms.
- **Digital Operations can accumulate role by role.** One Digital Worker produces reusable identity, validation, evidence, and escalation patterns that later roles extend rather than recreate.
- **AI can lower the burden of managing AI.** An AI advisor embedded inside the management system can continuously handle discovery, documentation, and analysis, while accountable people retain every consequential decision.

A Digital Worker is a role, not a product

What AI systems are asked to do has changed. For most of the past decade, they supplied capability, and a person decided what to do with it. Increasingly, they perform the work themselves, holding credentials, moving across systems, and producing outcomes that the organization answers for.

Consider what a prior authorization Digital Worker does. It reads chart documentation, applies one payer's medical policy as modified by one contract, navigates a portal whose behavior differs by plan, recognizes which cases exceed its scope, and routes those to a named human reviewer. Very little of that is model behavior. Almost all of it describes how one organization works.

That difference matters because a tool and a worker are governed differently. A tool is selected, installed, and configured. A worker is defined, given the knowledge it needs, granted access, authorized to act within limits, supervised, corrected, and eventually retired. Those are management acts rather than procurement decisions, and they belong to the organization whose work is being performed.

Healthcare implementation evidence has been consistent on this point for years. Trustworthy healthcare AI lifecycle guidance treats traceability, usability, robustness, and local evaluation as properties of the deployment, not the model.¹ Performance in one setting doesn't automatically transfer to another, which is why ongoing local validation is recommended instead of a one-time external validation.² Hospital implementation reviews identify process, people, information, and technology factors, rather than the model, as decisive.³ The reason is structural, as healthcare work happens inside a sociotechnical system where people, tasks, tools, setting, and organizational conditions jointly produce outcomes.⁴⁻⁵

The operative questions therefore have little to do with what an organization buys. Who defines what this worker does and may never do. Which knowledge governs it, and who keeps it current. Which systems and data it may reach. Who can change it, narrow it, or stop it. Where its work runs, and what record it leaves. Most organizations already report using AI somewhere, yet agent deployment that actually performs work remains in the single digits across most business functions.⁶ The gap sits precisely where those questions go unanswered.

Creation is the part that cannot be delegated

Creating a Digital Worker means translating the organization's own work. In this context, the word creation carries a misleading association with model development. However, in practice, bringing a Digital Worker into operation follows a sequence in which model work is a small component.

Most of that work is specification of local practice. Which payers, which contract terms, which documents count as evidence, which exceptions recur, who reviews what, what a good outcome looks like, and what must never happen. A supplier can provide templates, tooling, automation, and considerable expertise. They are not well positioned to provide the organizational knowledge that one receiving office still requires a fax, that a plan's published policy is qualified by a contract amendment, or that one service line's documentation habits produce a recurring failure mode.

The strategic shift, then, is not that every healthcare organization should become an AI company. It is that creating a Digital Worker increasingly consists of translating the organization's own work into a governed operational role. Consensus best-practice work for prior authorization AI points the same way, treating scope, human verification, provenance, and escalation handling as properties to specify and test locally rather than inherit from a product category.⁷

This direction is becoming visible in the products themselves. Epic previewed a platform in March 2026 through which health systems configure agents using a visual builder, equip them with local policies and knowledge bases, and deploy them on their own timeline.⁸ By mid-2026, a small early-adopter group had built first agents.⁹ The reported figures come from an important supplier and selected customers, so they indicate direction.

The direction is what matters. If the creation step is being redesigned as configuration of institutional work, then the durable asset is not any individual Digital Worker. It is the organization's ability to perform that translation repeatedly, and to keep the result under management afterward.

What the organization should keep doing

Seven capabilities follow from that position. Each is a management ability rather than a right or product feature, because it is exercised through ordinary operating decisions.

1) Role definition is the ability to determine what work exists, its purpose and boundaries, its prohibited uses, its accountable supervisor, and the conditions under which it expands, narrows, or retires. Context governance involves choosing which policies, contracts, and procedures are authoritative, how they rank when they conflict, when they expire, and which exceptions apply.

2) Operational identity is the ability to grant, constrain, audit, and revoke a Digital Worker's credentials and permissions. Systems operating within environments that hold protected health information must be identifiable, least-privileged, monitored, and revocable.

3) Access is the agentic operation that manages tool access, privilege escalation, and autonomous action as distinct security surfaces.^{10–12}

4) Evidence is the ability to trace consequential work from sources through actions, approvals, overrides, versions, and outcomes. This capability carries legal weight. Liability analysis for healthcare AI turns substantially on organizational choices such as selection, integration, validation, supervision, maintenance, and how staff relied on the system, rather than on the model alone.^{13–14}

5) Reconstructability plays a key role after an incident, as the question is rarely whether the model was generally good, but whether this organization gave this system this authority in this workflow, under these controls, with this evidence. An organization can retain that exposure while lacking the record needed to show it acted reasonably. All of this exists to prevent responsibility without reconstruction.

6) Organizational learning is the ability to keep validated corrections and operating lessons available to future work rather than dispersing them into a supplier's product.

7) Delegated authority is the ability to decide what a Digital Worker may do independently, on what evidence, and when that authority narrows or stops. That determination is tied to consequences and local performance, and it aligns with regulatory expectations that professionals can independently review the basis of a recommendation rather than simply rely on it. Strategic optionality is the ability to change models, tools, or suppliers without rebuilding the operating logic of the work.¹⁵

Taken together, these seven describe a management layer through which Digital Labor is defined, authorized, supervised, evidenced, and improved, which is what a Digital Operations System (DOPs) provides. The organization holds the capabilities. The technology beneath them is a separate question (Fig. 1).

Control the role, data, and where it runs

Organization-controlled means the organization defines the roles, governs the knowledge and data behind them, and runs the work on compute it controls, on its own premises or in cloud capacity held under its own tenancy and terms. It does not mean the organization writes the software. Very few will, and none needs to. It means the creation of each Digital Worker, and the environment in which that worker operates, sit inside a boundary the organization can inspect, change, and revoke.

In practice, the control that matters resolves into three concrete things: what the Digital Worker is, what data it touches, and where its work executes.

The first is the role itself. An organization that specifies its own Digital Workers is not deploying product-category agents. A supplier's prior authorization agent necessarily encodes assumptions about which payers matter, which documents count, how much latitude the worker has, and when a case should stop. Those assumptions were made for a market, not for one operation. An organization that writes the role decides what the worker may read, which sources outrank which, what it may never do, which cases must reach a human, and what evidence each action leaves. Off-the-shelf capability supplies a component, not a role that matches the work.

The second is the data flow. Because a Digital Worker holds credentials and moves across systems, someone must decide which systems it can reach, which fields may leave the governance boundary, what is retained and for how long, what is logged, and what is deleted. An organization that owns those decisions can answer where its clinical and operational data went and why. One that inherits them from a product configuration answers from a supplier's documentation, a materially weaker position when a patient, a regulator, or a plaintiff asks.




FIGURE 1 - Digital Operations as an organization-controlled layer. Supporting technology can be sourced or self-hosted, while the organization defines the roles, controls the compute, and accumulates management assets as roles are added.

The third is where the work runs. A management system built around roles rather than around any particular model can operate models the organization hosts itself and models it reaches through an API. It can also move work fluidly between multiple AI models. That flexibility is becoming more valuable, not less. Independent tracking now places several suppliers, open and closed, in the same top tier of general capability, with the remaining margin at the frontier measured in low single digits.¹⁶⁻¹⁷ Capability varies by task rather than moving as a single quantity, releases arrive continuously, and the range of AI models an organization can run on its own infrastructure keeps widening.

This is where building around roles rather than models proves beneficial. If capability varies by task, and costs, latency, and data handling differ, the best approach isn't to assign a role to a single model but to route each task to the most suitable model. Routing choices based on costs, while considering consequences, is a managerial decision. This requires information that a Digital Operations System already maintains: the nature of the work, governing factors, whether it involves protected data, and what must remain recoverable. Organizations can then assign most critical steps to models they control, delegate less critical ones elsewhere, and adjust these assignments as capabilities and prices change. The role remains unchanged, but the execution environment becomes a governed decision linked to that role.

Execution locality also bears on compliance, though not in the way it is often described. Any supplier that receives or maintains protected health information on the organization's behalf is a business associate requiring an agreement, and each adds a relationship that must be contracted, assessed, monitored, and reconstructed after an incident. Running a model on infrastructure the organization controls removes a party from that chain and narrows the disclosure surface. It does not lower the obligation, because safeguards attach to whoever operates the infrastructure and what gets examined after an incident is the whole chain of access, logging, and contracts.¹⁰

The precedent is quality, not software

There is a familiar parallel. A hospital can buy quality software, hire consultants, and undergo accreditation, and it still cannot outsource quality. Quality lives inside the organization's own management system, or it does not exist. Security follows the same logic. Managed services are routine, while the organization remains accountable for its security posture, access decisions, and incident response.

Digital Operations belongs to that family rather than to software procurement. The discipline can be supported extensively from outside, but the authority cannot be transferred, which is why the sourcing question is about what to keep rather than what to buy.

Management layers are becoming standard

These requirements are being recognized well beyond any single framework. Healthcare accreditation has moved to the organizational level, with a certification launched in June 2026 that evaluates governance, data management, risk and bias reduction, monitoring and validation, and transparency and training at the healthcare organization level, explicitly declining to certify individual AI products.¹⁸ Standards work has turned to agent identity, through a national initiative launched in February 2026 and an accompanying concept paper on applying established identity and authorization standards to software and AI agents in enterprise settings.¹⁹⁻²⁰

Major platform suppliers have converged on comparable architecture, describing control planes built around agent registries, identity and access control, lifecycle management, audit logging, and observability. The production question has shifted from whether an agent can be built to whether a fleet can be run reliably at scale.²¹⁻²³ These are commercial architectures rather than neutral evidence. Their significance is narrower and still useful. Registry, identity, lifecycle, evidence, and control are becoming expected infrastructure.²⁴⁻²⁵

The capability can begin with one role

A specialty group, community hospital, or regional plan may reasonably ask whether this is accessible at their size. The architecture does not require enterprise scale from the beginning, provided there is accountable ownership and fundamental technical controls. The capability can be built incrementally, role by role, since effectively managing a Digital Worker in one role creates resources that subsequent roles can leverage.

A first role involves creating a written definition that includes an owner, supervisor, scope, and exclusions, resulting in a role template and record pattern. The information set is similar to what you gather when planning, selecting, and onboarding a new employee. The process also requires definition of authoritative sources and freshness rules, which generate a context inventory and a source hierarchy. Additionally, a machine identity with least-privilege permissions is necessary, leading to a reusable credentialing pattern.

Validation cases and classified errors yield an evaluation set and an error taxonomy. Human approvals and escalations yield a supervision interface and an evidence schema. Performance and reviewer-burden measures yield a scorecard and a monitoring cadence. Incident and change handling yield rollback and revalidation triggers.

The claim is not that this is inexpensive. It is that the marginal management burden should decline. A second and third Digital Worker should not require rediscovering identity patterns, evidence schemas, validation methods, or escalation design. Digital Operations develops into infrastructure because these components are reused, allowing the capability to expand alongside the workforce rather than before it. This type of capability builds through intentional, repeated investments in routines, not just experience.²⁶

Lowering the cost of translation

If translation is the scarce activity, the useful question is what can make it cheaper. Two mechanisms are currently visible, and they work on different parts of the problem.

The first is forward deployment. Suppliers have increasingly organized around embedded engineering teams that work inside a customer's operation, learn the workflow, and build until the system runs in production. It is an effective answer, and its spread is evidence that translation, rather than model capability, is the binding constraint.

Its limits are limits of supply, as such people are scarce and are typically allocated to large episodic contracts, so the knowledge produced sits with the engineer and the supplier unless something on the organization's side captures it. The more practical constraint is time. A human engagement has a fixed number of hours, so translation gets triaged. The highest-value workflow receives careful attention, and the rest are left to whatever documentation already exists. No organization would staff a consultant to sit with a scheduling supervisor for a full day extracting the exception rules she carries in her head, even though those rules determine whether a Digital Worker in that queue will behave correctly.

The second mechanism follows from that constraint. The forward-deployed function can be instantiated as an embedded AI advisor inside the Digital Operations System (DOpS) rather than as a visiting team. It can inventory existing AI-enabled work, interview operational owners, draft role definitions and record fields from

standard operating procedures, identify candidate governing sources, and flag duplicates, stale versions, conflicts, and ownership gaps. It can generate representative test cases and failure taxonomies, summarize evidence, package escalations, detect drift and recurring errors, and propose candidate lessons.

What changes is availability, not insight. An embedded AI advisor has the time and the patience the work actually requires. It can spend as long as it takes on a queue nobody would fund a consultant to study, return to the same supervisor a week later when a new exception appears, and be as present for the fourth Digital Worker as for the first. Scale comes from that persistence, not from any single interaction being better than what a skilled engineer would manage.

Human team members remain essential, and their role is unchanged. Any recommendation that would alter role scope, context authority, access, delegated authority, or incident status requires accountable approval, and advisory outputs should carry their provenance. What shifts is who carries the labor of discovery, documentation, and analysis, not who carries the decision.

Some of that burden can also be removed before an organization starts. Much of what makes a first Digital Worker expensive is the surrounding apparatus of identity patterns, evidence schemas, and control baselines rather than the role itself, and it can arrive already implemented. The arrangement is best described as conformant with local regulations and relevant ISO standards (42001, 27001, 27701) by construction and by design, with controls in place by default and audit records generated as work occurs. Compliance itself remains the organization's, since obligations follow whoever operates the system. What changes is the distance to it.

Two limits deserve statement. An AI advisor may lack the cross-customer pattern knowledge a supplier's engineers accumulate, which remains a real advantage of the external model, and advisory automation propagates poor assumptions faster when the inputs are weak. This is a proposed management architecture, not a validated healthcare outcome.

Choosing the sourcing boundary

None of this makes internal implementation automatically preferable. A well-designed managed service under clear terms can be more responsible than a poorly staffed internal build, and an internal build fails wherever governance, context, validation, or supervision are weak. The test is not where the payroll sits. It is whether authority, evidence, context, accountability, and exit are clear enough for the risk the work carries.

Several operating patterns are legitimate, and they differ mainly in what the organization must still hold. What moves an organization along that range is the character of the work, not a preference about suppliers. How consequential an error would be, how much dependable performance rests on local contracts and exceptions, whether the worker proposes or transacts, whether it holds production credentials and touches protected information, whether corrections accumulate into reusable knowledge, and what changing supplier would cost.

As work moves toward the demanding end of those drivers, the case for an organization-controlled management layer strengthens. At the same end, so does the case for executing the most sensitive steps on infrastructure the organization governs (Fig. 2).




FIGURE 2 - Choosing the sourcing boundary. What differs between patterns is not whether the organization uses external capability but which management responsibilities it retains, and where each task executes.

Optionality and the cost of hollowing

Organizations that retain these capabilities gain something more specific and durable than independence. They gain the ability to evaluate a new model against their own cases, compare it with the one in use, and maintain a rollback path, allowing them to treat technical progress as a replaceable component rather than a change to their work operating logic. Governance frameworks for enterprise healthcare AI point in the same direction, treating lifecycle management and risk mitigation as organizational processes rather than as properties of a purchased system.²⁷⁻²⁸

The failure mode is worth naming because it is quiet. Organizational hollowing occurs when throughput rises while the ability to explain how work is performed declines, because role logic, context, and operating lessons accumulate in systems the organization does not govern. This is not typically highly visible, but the executive team can easily assess this risk by asking what would survive a change of supplier.

The executive test

The test focuses on continuity rather than procurement. Six questions, answerable from the system of record rather than from the implementation team's recollection, indicate whether operating a digital workforce is leaving the organization more capable than it found it (Table 1).

If the organization can change technical components while retaining role definition, context, evidence, learning, and authority, then external technology is functioning as a replaceable capability. If changing supplier means losing the operating logic of the work, the supplier has become part of the organization's memory.




TABLE 1 - The executive continuity test. Six questions answerable from the Digital Operations System of record. If the last answer includes the role itself, the capability has not been retained.

Close: build and own the capability

Control also has a cultural dimension. Staff trust Digital Labor more readily when it is visibly part of the organization, when they know which Digital Workers exist, who supervises them, and how to correct them. Clinicians need confidence that professional boundaries hold, risk leaders that events can be reconstructed, and compliance leaders that policy is enforced rather than merely stated.

External capability and organizational capability complement rather than rival each other. For commodity, low-consequence, easily reversible work, external management is often the reasonable choice.

As work becomes more sensitive, more locally differentiated, and more authority-bearing, an organization increasingly needs a management layer it governs in its own name. Not because vendors are dangerous, but because accountability does not transfer with the technology.

The durable advantage, then, is not owning every AI component. Models will change, and the ones in use a year from now may not be the ones in use today. What persists, if the organization builds it, is the ability to define the work, govern the knowledge that shapes it, decide what may be done independently, reconstruct what happened, keep what was learned, and replace the technology underneath without forgetting how its work is performed. That capability is not purchasable in a single transaction. It is accumulated one governed role at a time, and it remains the organization's after the technology has moved on.

Glossary of key concepts

- **Capability ownership.** The retained ability to define, contextualize, authorize, supervise, evidence, improve, and replace the Digital Workers. Distinct from technology ownership, which concerns who builds, hosts, licenses, or operates the underlying components.
- **Organization-controlled.** The organization defines the roles, governs the knowledge and data behind them, and runs the work on compute it controls, either on its own premises or in cloud capacity held under its own tenancy and terms. It does not require the organization to write the software.
- **Sourcing boundary.** The deliberate division between externally supplied capability and organization-retained management responsibility, set by consequence, local differentiation, delegated authority, privileged access, accumulated learning, and switching cost.
- **Execution locality.** Where a given task is computed and on whose infrastructure. Models may be hosted by the organization or accessed through an interface, and work can move between them without changing roles.
- **Consequence-based model routing.** Directing individual tasks to models selected by consequence, data class, cost, and fitness rather than committing a whole role to one model, with the allocation governed and recorded against the role.
- **Conformant by construction.** A management baseline in which controls are implemented by default and evidence is generated as work occurs rather than reconstructed afterward. It shortens

the distance to compliance, while the obligation remains with the organization operating the system.

- **Embedded AI advisory.** The forward-deployed engineering function is instantiated as AI inside the Digital Operations System rather than supplied solely as a visiting team. It reduces discovery, translation, documentation, validation-design, and evidence-review burden, and its value comes from continuous availability rather than outperforming a skilled engineer. Accountable people approve every consequential decision.
- **Exportability.** The demonstrated ability to carry role definitions, context libraries, evidence, and learning history across models and suppliers in usable, documented form.
- **Organizational hollowing.** The gradual loss of an organization's ability to explain how its work is performed, even as throughput rises, because operating logic and accumulated experience reside in systems it does not govern.
- **Organizational sovereignty.** The condition when these capabilities are retained: effective authority over the operating system of the organization's Digital Labor, while external capability remains freely usable beneath it.

Live Article information

This PDF is a dated snapshot of a continuously maintained Live Article. Material updates undergo editorial review, and substantive public revisions are released when warranted. The current maintained article and evidence base are available through the ConverGen Digital Operations Advisor: www.convergen.ai/advisor/

Peer-review status: This public edition has not undergone external peer review.

Competing interests: The author is the founder and CEO of ConverGen AI, which develops and commercializes a Digital Operations System (DOPs) for healthcare organizations.

AI-assisted tools: AI-assisted tools supported research synthesis, drafting, editing, and figure production. The author reviewed and takes responsibility for all content.

Author ORCID: [0009-0008-1814-1983](https://orcid.org/0009-0008-1814-1983)

References

1. Lekadir, K. et al. FUTURE-AI: international consensus guideline for trustworthy and deployable artificial intelligence in healthcare. *BMJ* 388, e081554 (2025). <https://doi.org/10.1136/bmj-2024-081554>
2. Youssef, A. et al. External validation of AI models in health should be replaced with recurring local validation. *Nat. Med.* 29, 2686–2687 (2023). <https://doi.org/10.1038/s41591-023-02540-z>
3. Kamel Rahimi, A. et al. Implementing AI in hospitals to achieve a learning health system: systematic review of current enablers and barriers. *J. Med. Internet Res.* 26, e49655 (2024). <https://doi.org/10.2196/49655>
4. Carayon, P. et al. Work system design for patient safety: the SEIPS model. *Qual. Saf. Health Care* 15, i50–i58 (2006). <https://doi.org/10.1136/qshc.2005.015842>
5. Sittig, D. F. & Singh, H. A new sociotechnical model for studying health information technology in complex adaptive healthcare systems. *Qual. Saf. Health Care* 19, i68–i74 (2010). <https://doi.org/10.1136/qshc.2010.042085>
6. Stanford Institute for Human-Centered Artificial Intelligence. The 2026 AI Index Report. Stanford University (2026). <https://hai.stanford.edu/ai-index/2026-ai-index-report>. Accessed 16 Aug 2026.
7. Coalition for Health AI. Best Practice Guide and Testing & Evaluation Framework: Prior Authorization (2026). <https://www.chai.org>. Accessed 16 Aug 2026.
8. Epic Systems. Agent Factory. Previewed at the HIMSS 2026 Global Health Conference & Exhibition, March 2026. Accessed 16 Aug 2026.
9. Epic Systems. Epic Almanac, volume 4 (2026). Reported early-adopter agent development at Loma Linda University Health and ECU Health. Accessed 16 Aug 2026.
10. U.S. Department of Health and Human Services. The HIPAA Security Rule. <https://www.hhs.gov/hipaa/for-professionals/security/index.html>. Accessed 16 Aug 2026.
11. OWASP Gen AI Security Project. Agentic AI: Threats and Mitigations (2026). <https://genai.owasp.org/resource/agentic-ai-threats-and-mitigations/>. Accessed 16 Aug 2026.
12. National Institute of Standards and Technology. The NIST Cybersecurity Framework (CSF) 2.0. NIST CSWP 29 (2024). <https://doi.org/10.6028/NIST.CSWP.29>
13. Mello, M. M. & Guha, N. Understanding liability risk from using health care artificial intelligence tools. *N. Engl. J. Med.* 390, 271–278 (2024). <https://doi.org/10.1056/NEJMe2308901>
14. Office of the National Coordinator for Health Information Technology. Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing (HTI-1) Final Rule (2024). <https://www.healthit.gov/topic/laws-regulation-and-policy/health-data-technology-and-interoperability-certification-program>. Accessed 16 Aug 2026.
15. U.S. Food and Drug Administration. Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff. <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software>. Accessed 16 Aug 2026.
16. Stanford Institute for Human-Centered Artificial Intelligence. Technical Performance. In: The 2026 AI Index Report. Stanford University (2026). <https://hai.stanford.edu/ai-index/2026-ai-index-report/technical-performance>. Accessed 16 Aug 2026.
17. Bengio, Y. et al. International AI Safety Report 2026. Preprint at <https://doi.org/10.48550/arXiv.2602.21012> (2026).
18. Joint Commission. Responsible Use of AI in Healthcare (RUAH) Certification (2026). <https://www.jointcommission.org/en-us/certification/responsible-use-of-ai-in-healthcare>. Accessed 16 Aug 2026.
19. National Institute of Standards and Technology. Announcing the AI Agent Standards Initiative for Interoperable and Secure Innovation (17 February 2026). <https://www.nist.gov/artificial-intelligence/ai-agent-standards-initiative>. Accessed 16 Aug 2026.
20. National Institute of Standards and Technology, National Cybersecurity Center of Excellence. Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization. Concept paper (5 February 2026). <https://csrc.nist.gov/pubs/other/2026/02/05/accelerating-the-adoption-of-software-and-ai-agent/ipd>. Accessed 16 Aug 2026.
21. Microsoft. Agent 365. Generally available 1 May 2026. <https://www.microsoft.com/en-us/microsoft-agent-365>. Accessed 16 Aug 2026.
22. Google Cloud. Gemini Enterprise Agent Platform: govern your agents. <https://docs.cloud.google.com/gemini-enterprise-agent-platform/govern>. Accessed 16 Aug 2026.
23. Amazon Web Services. Well-Architected Agentic AI Lens. <https://docs.aws.amazon.com/wellarchitected/latest/agentic-ai-lens/agentic-ai-lens.html>. Accessed 16 Aug 2026.
24. International Organization for Standardization. ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system (2023). <https://www.iso.org/standard/81230.html>. Accessed 16 Aug 2026.
25. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1 (2023). <https://doi.org/10.6028/NIST.AI.100-1>
26. Zollo, M. & Winter, S. G. Deliberate learning and the evolution of dynamic capabilities. *Organ. Sci.* 13, 339–351 (2002). <https://doi.org/10.1287/orsc.13.3.339.2780>
27. Bodnari, A. & Travis, J. Scaling enterprise AI in healthcare: the role of governance in risk mitigation frameworks. *npj Digit. Med.* 8, 272 (2025). <https://doi.org/10.1038/s41746-025-01700-4>
28. Teece, D. J., Pisano, G. & Shuen, A. Dynamic capabilities and strategic management. *Strateg. Manag. J.* 18, 509–533 (1997). [https://doi.org/10.1002/\(SICI\)1097-0266\(199708\)18:7<509::AID-SMJ882>3.0.CO;2-Z](https://doi.org/10.1002/(SICI)1097-0266(199708)18:7<509::AID-SMJ882>3.0.CO;2-Z)