

Organizational Context as Infrastructure

Why local knowledge is the control layer for Digital Labor

Fabio A. Thiers, MD, PhD

ConverGen AI, New York, NY, USA

ABSTRACT

AI models do not know the organization by default. They may possess broad medical, administrative, and procedural knowledge, but dependable healthcare work depends on local operational context: organizational policies, contracts, workflow rules, source authority, patient and case state, supervisory expectations, and other institution-specific knowledge that cannot be inferred from general AI model training. This article argues that organizational context is not merely information available to an AI Digital Worker. It is the managed control layer that governs how Digital Labor is performed. Information becomes operational context only when its relevance, authority, freshness, scope, and permitted use have been specified for a particular role and action. Digital Operations therefore manages context as infrastructure through role-specific context specifications, source hierarchy, freshness requirements, completeness thresholds, stewardship, and evidence of context use. These controls allow Digital Workers to operate within the organization's actual operating conditions rather than relying solely on intrinsic AI capabilities. This organizational context infrastructure maintains the reliability, explainability, and governability of Digital Labor as AI models and operational conditions continue to evolve.

Keywords: healthcare AI; organizational context; context stack; source hierarchy; freshness; missing context state; context stewardship; context specification; Digital Operations; Digital Operations System (DOPs); Digital Workers; Digital Labor; data provenance; sociotechnical systems.

Key points

- **Context is more than retrievable information.** It becomes operational context when its relevance, authority, scope, freshness, and permitted use are specified for a Digital Worker.
- **Missing context is an operational state.** Absent, stale, conflicting, low-authority, or inapplicable information should trigger the Digital Worker to verify, gather, ask, abstain, defer, route, or escalate rather than confidently continue AI inference.
- **Reliable Digital Workers depend on role-specific context specifications.** Source hierarchy, freshness, completeness thresholds, and missing-context responses determine how work proceeds under local organizational conditions.
- **Shared context becomes strategic infrastructure.** Governed sources, update triggers, and evidence can be reused across a portfolio of Digital Workers while preserving role-specific access and control.
- **Context stewardship is a Digital Operations function.** Local knowledge requires explicit ownership, governance, versioning, and operational evidence so that Digital Workers consistently use the appropriate sources for each task.

Models do not know the organization by default

A capable model can know many general facts and still fail to properly perform accountable work in an organization. It may understand what prior authorization is, how an appeal is structured, or what information commonly appears in a referral. That does not mean it knows which payer policy is active for this plan, which EHR field is authoritative for this clinic, which documentation template this service line uses, what the supervisor wants escalated, or whether a portal response is stale. General capability and local competence are different properties, and the second cannot be inferred from the first.

Healthcare implementation research indicates that technology performance depends on factors such as local setting, workflow, users, organizational complexity, and the surrounding sociotechnical system¹⁻⁷. Additionally, repeated local validation is

recommended because performance can differ between institutions and over time.⁸

Generative AI intensifies this dependency because fluent outputs can mask missing or inappropriate organizational context while remaining highly persuasive.⁹ A Digital Worker may therefore produce a persuasive answer while relying on an outdated policy, an incomplete record, or a source that is authoritative for a different plan, jurisdiction, or task. Retrieval-grounded systems reduce, but do not eliminate, this risk because retrieval depends on selecting authoritative sources, assessing their applicability to the task, and recognizing when relevant context is incomplete, conflicting, or unavailable.¹⁰ Human reviewers also help, but cannot be assumed to catch every persuasive error because automation bias can undermine verification when outputs appear authoritative.^{11,12}

The organizational implication is therefore straightforward. Local context cannot be treated as an informal collection of documents or a retrieval corpus attached to an AI system. It must be managed as operational infrastructure, with explicit definitions of authority, freshness, applicability, stewardship, and evidence of use. Only then can Digital Workers operate within the institution's actual operating conditions rather than relying on generalized inference.

Context is not just data

Data is part of context, but context is broader. A lab result, claim status, payer rule, clinic SOP, patient message, locally approved review threshold, and queue note can all become relevant context. They differ in authority, freshness, applicability, scope, provenance, and actionability, and those differences determine what a Digital Worker may safely do with each of them.

Operational context is not defined by what information exists within an organization, but by what information has been deliberately prepared to support a specific action. Information becomes operational context only when its relevance, authority, applicability, freshness, and permitted use have been specified for a particular Digital Worker role and task. Context is therefore not simply retrieved; it is operationally defined.

A Digital Worker does not need all information. It needs the right information for the role. A referral worker needs network status, scheduling rules, symptoms that affect urgency, language access, transportation constraints, and documentation requirements. A denial worker needs claim history, the denial reason, contract terms, payer policy, medical documentation, prior submissions, and appeal deadlines. A documentation worker needs chart facts, service-line templates, payer-specific requirements, and clinician review thresholds. Each role draws on a different slice of the organization's knowledge, at a different tempo, under different constraints.

The worker therefore needs a context stack: a layered set of local knowledge and operating state that defines what is true, what is current and applicable, which source governs, and what should happen when information is missing or conflicting (Fig. 1).

The context stack

A practical context stack includes at least six layers. **Core records** anchor the stack: patient identity, encounters, orders, claims, benefits, notes, documents, messages, authorizations, and clinical facts relevant to the role. **Policy and rule context** includes payer policy, plan rules, benefit design, contract terms, local SOPs, documentation templates, compliance constraints, and the source hierarchy itself. **Workflow state** carries the case forward: queue status, prior attempts, current task, handoffs, deadlines, pending approvals, and next action.

System and partner state describes the environment in which the worker acts: the EHR build, portal behavior, API status, fax or document-exchange evidence, partner response reliability, and known failure modes. **Patient and access context** captures what safe and equitable service requires: language, caregiver authority, disability accommodation, transportation, portal access, health literacy, and timing constraints. **Supervisory judgment** completes the stack through documented local expectations about what accountable leaders require to be reviewed, escalated, questioned, or prohibited.

The layers have different operating forms and tempos. Policies, contracts, and approved SOPs are relatively persistent institutional context. Patient facts and claim histories are case-specific context. Queue status, portal responses, pending handoffs, and system availability are ephemeral operating states. Payer contracts may change yearly, policies monthly, system state daily, and workflow state within a single case. DOpS must therefore manage context as a living operating layer, with update triggers, validation, retention, and evidence requirements matched to each layer. Each role draws a different slice of the same stack, which is why context requirements are specified at the role level rather than the enterprise level (Fig. 2).

Most of these context layers already exist within healthcare organizations. Digital Operations does not recreate them; it identifies the subset required for accountable work, assigns operational meaning to those assets, and makes them consistently consumable across Digital Workers.




FIGURE 1 - Organizational context stack. Six layers of local knowledge support a Digital Worker role, from core records to supervisory judgment. The layers change at different tempos, which is why context must be operated as a living layer rather than archived as a static knowledge base.

Source hierarchy and freshness

A Digital Worker needs a rule for which context governs a particular action. Source hierarchy is not one universal ranking; it is a task-specific precedence rule. An executed contract may govern reimbursement while a current payer policy governs medical-necessity documentation. A signed clinician order may govern the requested service while a local SOP governs the internal workflow. A portal response may provide current operational evidence without being legally dispositive. Each priority must therefore be encoded with its scope and applicability, and conflicts that cannot be safely resolved must remain visible rather than being silently collapsed by the model.

Freshness defines whether a source is current enough for action, but recency alone is insufficient. The source must also apply to the correct patient, plan, service, jurisdiction, effective period, and task. A payer policy used to submit an authorization may need

a version and effective-date check at submission. A patient address used for communication may require verification at the moment of action. A queue note may become stale after a handoff or external response. A recently retrieved document may still be superseded, incomplete, or inapplicable. Freshness is therefore shorthand for task-specific current validity.

Data-provenance research establishes the importance of tracing where information came from, how it moved, and how it was transformed.^{13,14} Digital Operations extends that principle into action-time control. For each material action, the worker should preserve the source, timestamp, version, authority level, applicability, and any conflict or gap that shaped the decision (Fig. 3). For clinical decision-support functions within its scope, FDA guidance similarly emphasizes that healthcare professionals must be able to independently review the basis for a recommendation.¹⁵




FIGURE 2 - Role-level context slices. Three Digital Worker roles draw on the same six-layer stack, but each requires a different slice of it. Greyed cells mark layers that are secondary for that role. Context requirements are therefore specified per role, not per enterprise.

Missing context is an operational state

A reliable Digital Worker should not fill missing organizational context with confident inference. Missing context should be treated as an operational state, and it should trigger one of several designed responses: verify, gather more information, ask a human, abstain, defer, route to another role, or escalate.

The principle is simple but consequential. A human worker may know that an absent document usually means the case should wait. A generative AI system, left uninstructed, may treat the same absence as a gap to be bridged by plausible reasoning, consistent with the failure modes described in the hallucination literature.⁹ NIST's Generative AI Profile likewise identifies confabulation and information-integrity risks that require explicit management.¹⁶ The same operational risk appears when context is stale, conflicting, low-authority, inapplicable, or outside the role's scope.

DOPs should therefore encode context-completeness thresholds by task. A worker may draft but not submit when policy validity is

unknown. It may retrieve but not conclude when patient identity is ambiguous. It may summarize but not send when language preference or caregiver authorization is unclear. It may escalate when two authoritative sources conflict. Each threshold converts a silent judgment into an explicit, auditable rule and defines what work remains permissible while the gap is resolved.

Context ownership and stewardship

Organizational context needs accountable owners. Payer-policy context may be owned by revenue cycle or contracting. Clinical documentation templates may be owned by clinical operations. Access rules may be owned by security. Workflow SOPs may be owned by operations. Equity and language-access context may require input from patient experience and compliance. Material supervisory judgment belongs with accountable service-line leaders and should be converted from tacit practice into documented, reviewable operating expectations before it becomes machine-usable context.

If no one owns context, Digital Workers inherit fragmentation. They will use whatever sources are easiest to retrieve, not necessarily the sources the organization trusts, and they will learn from old cases, informal summaries, or stale documents because no one maintains the authoritative layer.

Context stewardship is therefore a Digital Operations function. Data-stewardship research supports the need for accountable organizational ownership; Digital Operations extends that responsibility into operational controls for versioning, validity review, source retirement, conflict resolution, access, update triggers, and evidence that workers used the appropriate source at the right time.¹⁷ Each layer should carry a named steward, explicit update triggers, and a review cadence (Fig. 4).

This discipline is consistent with AI management-system and risk-management frameworks that emphasize governed, documented, monitored information and risk controls^{16,18-20}, and it extends to local operating knowledge the structured documentation practices developed for models and datasets.^{21,22}

Stewardship should be proportional to operational consequence. Not every organizational document requires the same level of governance. Frequently changing, safety-critical, or operationally significant context warrants stronger stewardship than static reference information. Digital Operations therefore focuses governance effort on the context that materially influences accountable work rather than attempting to curate every organizational asset.

Context management scales through reuse

One concern is that managing organizational context for Digital Workers could become an impractical documentation exercise. That is not the objective. Healthcare organizations already manage large bodies of operational knowledge through policies, contracts, payer rules, workflow definitions, identity systems, provider directories, clinical pathways, templates, and standard operating procedures. Digital Operations does not create this knowledge; it organizes and governs the subset that materially influences accountable work.

Most operational context is shared across many Digital Workers rather than maintained independently for each role. A payer-policy library, provider directory, language-access rules, or referral playbook may support dozens of Digital Workers simultaneously. Individual roles do not duplicate these assets. Instead, each role defines a Context Specification that identifies which governed context sources it requires, the authority assigned to those sources, applicable freshness requirements, and the expected response when required context is unavailable or conflicting.

Context management should also be incremental. Organizations do not need to define enterprise-wide context before deploying their first Digital Worker. They begin with the context required for one role, govern what materially affects that work, and progressively expand reusable context assets as additional Digital Workers are introduced. As the portfolio grows, the marginal effort shifts from creating new context toward reusing and refining existing governed context.

AI itself can substantially reduce the effort required to build and maintain this infrastructure. AI-enabled Digital Operations Systems (DOPs) can assist organizations by iteratively identifying candidate context sources, extracting relevant operational rules from existing policies and procedures, detecting inconsistencies, recommending missing context, monitoring source freshness, and proposing updates for human review. Human stewards remain accountable for approving operational context, but AI can reduce the cognitive load and manual effort required to discover, organize, and maintain the organizational context infrastructure.

The objective is therefore not to curate every piece of organizational knowledge. It is to identify, govern, and continuously maintain the relatively small subset of operational context that materially influences accountable work, making it reusable across many Digital Workers while preserving organizational ownership and accountability.




FIGURE 3 - Source hierarchy and the current-validity gate. The hierarchy defines which source is authoritative when sources conflict; the current-validity gate decides whether that source is current and applicable to the patient, plan, service, jurisdiction, effective period, and task at hand. Failing either test triggers a designed response rather than confident inference.

Context as a shared operational infrastructure

The first Digital Worker may justify assembling context for one workflow. The fifth and the tenth reveal the strategic value of shared context infrastructure. Payer-policy sources, contract terms, documentation requirements, identity rules, and escalation norms can be governed once and reused across multiple roles. Repeated discovery is reduced; stewardship cost is amortized; new validation begins from known sources; and a policy or contract change can be propagated to every dependent worker. Each subsequent deployment starts further ahead without giving every worker access to everything.

This is where organizational context differs from generic data strategy. The goal is to focus on access to the local knowledge that is required for accountable work, while keeping it current, authoritative, applicable, machine-usable, and governable. Interoperability and health-data-exchange standards, including FHIR, USCDI, and the CMS prior-authorization requirements, can improve the availability and structure of some inputs.²³⁻²⁵ They do not by themselves create the role-level context specification, source precedence, validity gates, or missing-context responses that Digital Labor requires.

Every Context Layer Has a Named Steward

If no one owns context, Digital Workers inherit fragmentation — they use whatever is easiest to retrieve, not what the organization trusts

CONTEXT LAYER	ACCOUNTABLE STEWARD	UPDATE TRIGGER
LAYER 6 Supervisory judgment	Accountable service-line leaders	Leadership change, incident review, or a shift in escalation expectations
LAYER 5 Patient and access context	Patient experience, with Compliance	Language-access, accommodation, or equity policy revision
LAYER 4 System and partner state	Information technology, with Security	EHR build, portal or API change, or a partner failure mode observed
LAYER 3 Workflow state	Operations	Workflow redesign, queue restructuring, or handoff change
LAYER 2 Policy and rule context	Revenue cycle and contracting; clinical operations for templates	Contract cycle, payer policy bulletin, or SOP revision
LAYER 1 Core records	Health information management, with IT	System of record change, data migration, or field authority change

STEWARDSHIP IS A DIGITAL OPERATIONS FUNCTION

Versioning Freshness review Source retirement Conflict resolution Access control Update triggers


Evidence that the right source was used at the right time

FIGURE 4 - Context stewardship map. Each layer of the context stack carries a named accountable steward and an explicit update trigger. Stewardship — versioning, freshness review, source retirement, conflict resolution, and access control — is a Digital Operations function, not an informal responsibility.

Context failure is often mistaken for model failure

When Digital Workers fail, the visible error may look like a model mistake. The underlying cause may be context failure: the worker did not know which source was authoritative, that a policy had changed, that a local exception applied, that a patient needed language assistance, that a portal value was stale, or that a supervisor expected review.

The distinction matters for improvement (Fig. 5). An organization that treats every failure as a model-performance issue may change models without repairing the operating dependencies that produced the error. Consistent with the broader literature on hidden technical debt in machine-learning systems, leaving source hierarchy, validity checks, ownership, and escalation unchanged may allow the same underlying weakness to recur in another form.²⁶



Close: context is the control layer

Digital Labor becomes local through context. It becomes bounded through task-specific source hierarchy, current-validity rules, completeness thresholds, permitted-use constraints, and escalation. It becomes accountable when the organization can show which context was used, why it was considered authoritative and applicable, what was missing or conflicting, and who reviewed the exceptions.

This is why context should be treated as infrastructure. It is not a document library attached to an AI system; it is the control layer through which organizational knowledge and operating state become usable for authorized work. Without that layer, capable models can guess, overgeneralize, or act on incomplete information, and failures may appear as surprising model errors even when their cause lies in the surrounding work system.

With a governed context layer, Digital Workers operate inside the institution's actual conditions: its contracts, policies, systems, patients, workflows, and accountable judgment. Once models cross a threshold of capability, the strength of this local layer often determines whether that capability can be trusted in consequential work. DOPs makes the layer explicit, maintained, testable, and evidenced so that context quality can improve across the Digital Worker portfolio rather than being rebuilt informally for each deployment.

Glossary of key concepts

- **Organizational context.** Local knowledge and operating state required to perform work safely and effectively, including policies, rules, workflows, patient and case facts, source hierarchy, system state, and documented supervisory judgment.
- **Context stack.** The layered structure of core records, policy and rule context, workflow state, system and partner state, patient and access context, and supervisory judgment that supports a role.
- **Source hierarchy.** A task-specific precedence rule that determines which source governs when sources conflict, including the scope and applicability of that authority.
- **Freshness.** The requirement that context be current and applicable enough for the action being taken, with validity thresholds set by task.
- **Missing context state.** A designed operational state in which absent, stale, low-authority, inapplicable, or conflicting context triggers verify, gather, ask, abstain, defer, route, or escalate rather than inference.
- **Context stewardship.** Named ownership and maintenance of local knowledge as a governed operating asset, including versioning, freshness review, and source retirement.
- **Context specification.** The role-level definition of required sources, scope, authority, applicability, freshness, completeness, permitted use, missing-context responses, evidence, and stewardship for a Digital Worker.

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. Greenhalgh, T. et al. Beyond adoption: a new framework for theorizing and evaluating nonadoption, abandonment, and challenges to the scale-up, spread, and sustainability of health and care technologies. *J. Med. Internet Res.* 19, e367 (2017). <https://doi.org/10.2196/jmir.8775>
2. Beede, E. et al. A human-centered evaluation of a deep learning system deployed in clinics for the detection of diabetic retinopathy. *Proc. 2020 CHI Conf. Hum. Factors Comput. Syst.* 1-12 (2020). <https://doi.org/10.1145/3313831.3376718>
3. Kelly, C. J., Karthikesalingam, A., Suleyman, M., Corrado, G. & King, D. Key challenges for delivering clinical impact with artificial intelligence. *BMC Med.* 17, 195 (2019). <https://doi.org/10.1186/s12916-019-1426-2>
4. Carayon, P. et al. Work system design for patient safety: the SEIPS model. *Qual. Saf. Health Care* 15 Suppl. 1, i50-i58 (2006). <https://doi.org/10.1136/qshc.2005.015842>
5. Holden, R. J. & Carayon, P. SEIPS 101 and seven simple SEIPS tools. *BMJ Qual. Saf.* 30, 901-910 (2021). <https://doi.org/10.1136/bmjqs-2020-012538>
6. 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>
7. Salwei, M. E. & Carayon, P. A sociotechnical systems framework for the application of artificial intelligence in health care delivery. *J. Cogn. Eng. Decis. Mak.* 16, 194-206 (2022). <https://doi.org/10.1177/15553434221097357>
8. 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>
9. Ji, Z. et al. Survey of hallucination in natural language generation. *ACM Comput. Surv.* 55, 248 (2023). <https://doi.org/10.1145/3571730>
10. Lewis, P. et al. Retrieval-augmented generation for knowledge-intensive NLP tasks. *Adv. Neural Inf. Process. Syst.* 33, 9459-9474 (2020). <https://proceedings.neurips.cc/paper/2020/hash/6b493230205f780e1bc26945df7481e5-Abstract.html>
11. Goddard, K., Roudsari, A. & Wyatt, J. C. Automation bias: a systematic review of frequency, effect mediators, and mitigators. *J. Am. Med. Inform. Assoc.* 19, 121-127 (2012). <https://doi.org/10.1136/amiajnl-2011-000089>
12. Lyell, D. & Coiera, E. Automation bias and verification complexity: a systematic review. *J. Am. Med. Inform. Assoc.* 24, 423-431 (2017). <https://doi.org/10.1093/jamia/ocw105>
13. Sembay, M. J., de Macedo, D. D. J., Júnior, L. P., Braga, R. M. M. & Sarasa-Cabezuelo, A. Provenance data management in health information systems: a systematic literature review. *J. Pers. Med.* 13, 991 (2023). <https://doi.org/10.3390/jpm13060991>
14. Johns, M. et al. Data provenance in biomedical research: scoping review. *J. Med. Internet Res.* 25, e42289 (2023). <https://doi.org/10.2196/42289>
15. U.S. Food and Drug Administration. Clinical Decision Support Software: Guidance for Industry and Food and Drug Administration Staff (2026). <https://www.fda.gov/regulatory-information/search-fda-guidance-documents/clinical-decision-support-software>. Accessed 3 Aug 2026.
16. National Institute of Standards and Technology. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1 (2024). <https://doi.org/10.6028/NIST.AI.600-1>

17. Rosenbaum, S. Data governance and stewardship: designing data stewardship entities and advancing data access. *Health Serv. Res.* 45, 1442-1455 (2010). <https://doi.org/10.1111/j.1475-6773.2010.01140.x>
18. 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>
19. International Organization for Standardization. ISO/IEC 42001:2023 Artificial intelligence - Management system (2023). <https://www.iso.org/standard/81230.html>. Accessed 3 Aug 2026.
20. International Organization for Standardization. ISO/IEC 23894:2023 Artificial intelligence - Guidance on risk management (2023). <https://www.iso.org/standard/77304.html>. Accessed 3 Aug 2026.
21. Mitchell, M. et al. Model cards for model reporting. *Proc. Conf. Fairness Accountability Transparency* 220-229 (2019). <https://doi.org/10.1145/3287560.3287596>
22. Gebru, T. et al. Datasheets for datasets. *Commun. ACM* 64, 86-92 (2021). <https://doi.org/10.1145/3458723>
23. HL7 International. FHIR Release 5 specification (2023). <https://www.hl7.org/fhir/>. Accessed 3 Aug 2026.
24. Assistant Secretary for Technology Policy/Office of the National Coordinator for Health Information Technology. United States Core Data for Interoperability Version 7 (USCDI v7) (2026). <https://isp.healthit.gov/sites/default/files/2026-07/USCDI-Version-7-July-2026.pdf>. Accessed 3 Aug 2026.
25. Centers for Medicare & Medicaid Services. 2024 CMS Interoperability and Prior Authorization Final Rule (CMS-0057-F) (2024). <https://www.cms.gov/initiatives/burden-reduction/overview/interoperability/policies-regulations/cms-interoperability-prior-authorization-final-rule-cms-0057-f>. Accessed 3 Aug 2026.
26. Sculley, D. et al. Hidden technical debt in machine learning systems. *Adv. Neural Inf. Process. Syst.* 28, 2503-2511 (2015). <https://proceedings.neurips.cc/paper/2015/hash/86df7dcfd896fcdf2674f757a2463eba-Abstract.html>
27. Obermeyer, Z., Powers, B., Vogeli, C. & Mullainathan, S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 447-453 (2019). <https://doi.org/10.1126/science.aax2342>