OPEN ACCESS
Global Journal of Research in Engineering & Computer Sciences
ISSN: 2583-2727 (Online)
Volume 06 | Issue 03 | May-June | 2026
Journal homepage: https://gjrpublication.com/gjrecs/
Research Article

Convergence of Generative AI, Agentic Workflows, and Cloud-Native Infrastructure: A Unified Architecture for Next-Generation Healthcare Intelligence

***Arvind Telharkar**

Engineering Lead at Amazon

DOI: 10.5281/zenodo.20597913 **Submission Date: 25 April 2026 | Published Date: 08 June 2026**

Abstract

The emergence of cloud-native computing, the advent of autonomous agentic workflows, the maturity of generative AI (AI): The three are converging to deliver a structural transformation of healthcare systems. Although each of these technologies have demonstrated significant utility for clinical decision support, operational automation and scalable data processing, there has not been a well-developed conceptualization of how they can be combined in a common architectural context. The article introduces a next-generation, integrated architecture for the healthcare intelligence that can include generative AI, agentic orchestration models, and cloud-native principles in a unified operation. The study draws from the latest research in the fields of large language models (LLMs), multi-agent systems, healthcare interoperability, and distributed computing, and highlights the potential for these cutting-edge technologies to collectively address the major persistent healthcare challenges such as fragmented data environments, clinician burnout, workflow inefficiencies, and restricted real-time clinical intelligence. The paper proposes a layered architectural solution, with each layer representing a distinct element or service, using a conceptual and systems-oriented approach. The paper adopts a conceptual and systems-oriented approach to build the layered architectural solution, which includes data interoperability services, AI reasoning agents, orchestration layers, governance mechanisms, and scalable cloud-native deployment environments. The results indicate that generative AI in agential workflows and containerized and event-driven infrastructure increases adaptability, resilience, explainability, and clinical responsiveness. The paper also proposes that the sophistication of the models is not the only factor driving the future of healthcare intelligence, but also the governance and orchestration of these models. Ethics, regulatory, cyber security and interoperability issues are discussed as policy implications. The article helps advance the scholarship on AI in healthcare as a distributed, autonomous and continually learning AI ecosystem rather than a single predictive tool.

Keywords: *AI-driven Healthcare Intelligence; Agentic Workflows; Cloud Native Infrastructure; Healthcare Intelligence; Large Language Models; Clinical Decision Support; Multi-Agent Systems; Digital Health; Kubernetes; Healthcare Informatics.*

Introduction

No matter where in the world you look, healthcare systems are facing growing pressure on a number of fronts: high patient volumes, chronic staffing challenges, high operating costs, and complex clinical data environments. Although decades of digitization have brought many benefits to healthcare, many healthcare organizations are still using fragmented information systems that make it difficult to coordinate care delivery and provide real-time access to clinical intelligence. However, electronic health records (EHRs), imaging repositories, laboratory systems, and administrative databases frequently are not well connected, and these data "silos" impede evidence-based decision making and add cognitive load to clinicians.

However, new opportunities have emerged with the recent developments of AI that can help tackle these structural constraints. Clinical summarization, conversational interfaces, medical documentation, coding support, and diagnostic

*Corresponding author: **Arvind Telharkar**

reasoning are some of the most impressive clinical capabilities shown by generative AI, especially through the use of large language models (LLMs). Whereas the previous machine learning systems needed structured data and very well-defined tasks, the generative AI systems can produce a new human-like response to the clinical information provided, which is unstructured and will contain context nuances.

A second way, however, has unfolded with agentic Artificial Intelligence systems. Agentic workflows are autonomous or semi-autonomous AI agents that are able to reason, plan, coordinate actions, and interact with external tools or environments to use to achieve multiple-step goals. In health care, these workflows have started to go beyond basic automation to patient triage, a workflow for coordinating patient care, patient scheduling, clinical documentation, and workflow orchestration (Joy, 2025). The shift from static AI models to autonomous AI agents marks a considerable paradigm change in the healthcare informatics landscape, as it shifts intelligence from singular prediction machines to continually evolving and adaptive systems.

Thirdly, cloud-native infrastructure is driving this change. Scalable computing infrastructure, featuring distributed data processing and real-time analytics, interoperability standards, and resilient deployment options, is increasingly a critical requirement for modern healthcare intelligence. The operational context for deploying AI systems securely and at scale is cloud-native technologies like microservices architectures, Kubernetes, containers, and orchestration and service meshes. If this infrastructure isn't in place, the advanced AI systems can be stuck in a pilot program or not implemented on a wider scale.

While there has been significant work on each of these domains individually, there is not yet a lot of work examining the intertwining of these areas. Current research frequently caters to a specific innovation domain, such as generative AI, workflow automation, or cloud computing, without recognizing the synergy they can offer and their ability to form a cohesive healthcare intelligence ecosystem. This means that in the real world of healthcare, the process of bringing AI systems to the forefront of clinical care is often complex, requiring careful attention to issues of scalability, trustworthiness, and sustainability.

To fill this gap, this article introduces a single architecture that combines generative AI, agentic workflows and cloud-native infrastructure into a seamless next-generation healthcare intelligence framework. The study suggests that while more powerful AI models are important, they are not the only drivers of meaningful healthcare transformation, as a system of architectures can enable collaborative autonomous reasoning, scalable computing, interoperability, governance, and continual learning. The central research question for this research is then: How can generative AI, agentic workflows and cloud-native infrastructure be integrated into a single architecture that will facilitate generating adaptive, secure and scalable healthcare intelligence?

The paper makes three important contributions. First, it pulls together interdisciplinary literature from the fields of AI systems, healthcare informatics, distributed computing and autonomous workflow. Second, it suggests a conceptual architecture with layers that helps to understand the interaction of these technologies on the fly. Third, it explores the governance and policy aspects of increasing autonomy of healthcare intelligence systems.

Literature Review Generative AI in Healthcare

From experimental language systems to tools that can revolutionize clinical and administrative healthcare processes, generative AI has quickly emerged as a powerful force in the healthcare landscape. Initially, AI in healthcare was mainly used for predictive analytics, classification of images and statistical risk modeling. The advent of transformer-based architectures, however, greatly broadened the scope of AI applications, allowing for the understanding of language in context and the generation of reasoned content.

The power of transformative potential of large language models was shown by Brown et al. (2020) with their creation of GPT-3, which was capable of few shot learning in various domains. Later applications focused on health care took these models to summarization of clinical notes, diagnosis, medical coding and conversational health interfaces. Such models have been called “foundation models” by Bommasani et al. (2021) because they can be used for a wide range of downstream applications by learning a general representation.

Generative AI has demonstrated great potential in healthcare settings where it can be used to alleviate the burden on clinicians with documentation tasks. Doctors' research suggests that doctors spend significant time using EHR systems instead of communicating with patients (Shanafelt et al., 2016). Summarizing interactions, writing clinical notes, and extracting pertinent medical data from them are all potential benefits of generative AI systems that could ease the administrative burden. However, there are major concerns about hallucination risks, lack of explainability, clinical reliability and propagation of bias. Wang et al. (2023) noted that, although LLMs exhibit impressive language skills, there is variability in their medical reasoning accuracy with respect to prompt engineering, training data quality and domain specificity. Such restrictions can be especially an issue in clinical settings where the accuracy of the output can have a direct impact on patient care, especially in a high-stakes situation.

Figure X. Generative AI Translation Layer for Healthcare Intelligence

Transforming heterogeneous healthcare data into actionable clinical intelligence

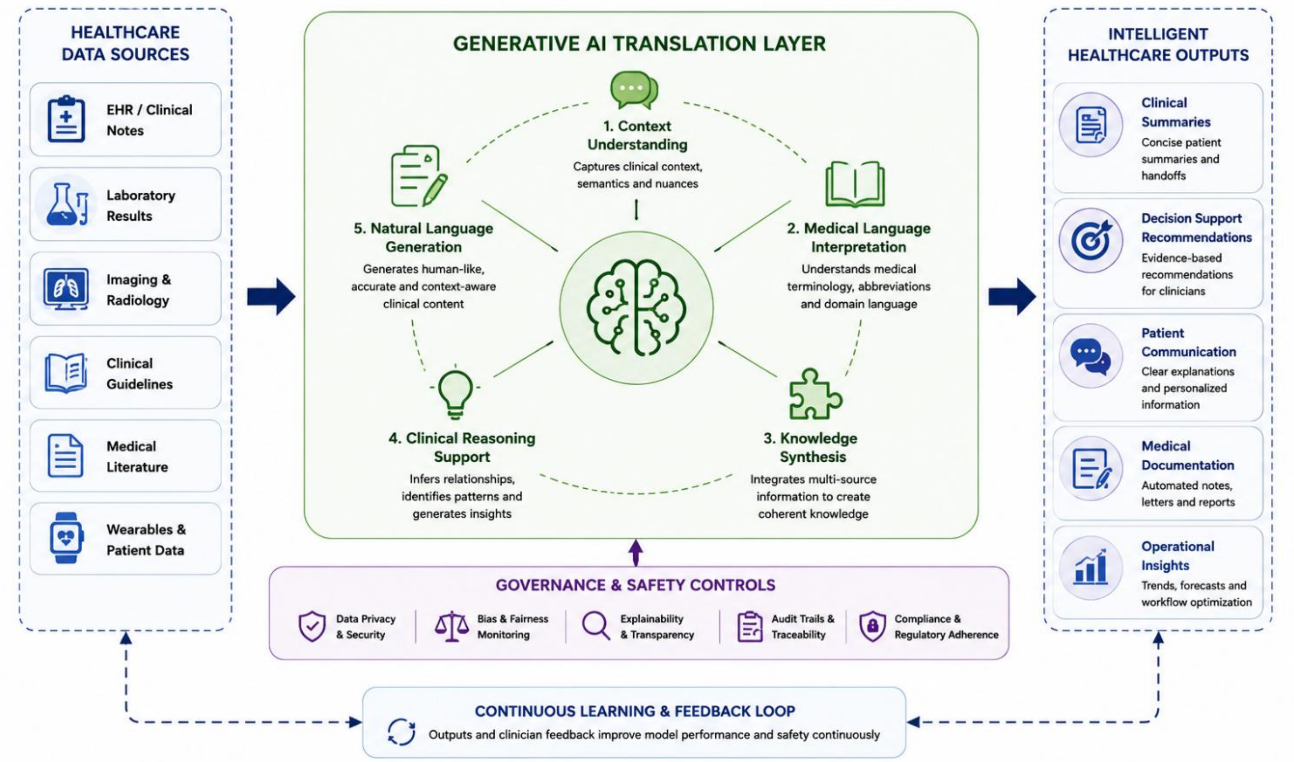

Self-governing Healthcare Systems and Agentic Workflows

Agentic AI is more than just a passive generative AI; it is an autonomous AI that plans, can remember, uses tools, and can execute multiple steps of decision making. Agentic systems are different than typical pipelines with AI as components, as they auto-magically adjust their actions based on changing environments and objectives.

There has been a growing interest in the application of multi-agent architectures in healthcare operations in recent years. Agentic workflows have the potential to optimize clinical workflows by enabling multiple AI agents to coordinate complex administrative and diagnostic workflows, as suggested by Joy (2025). In the same way, Barra et al. (2025) showed how agentic AI workflows may help in the design of healthcare simulation scenarios via an orchestration of reasoning processes. Retrieval-augmented generation (RAG), memory management, workflow orchestration, and domain-specific reasoning, all in adaptive operational structures.

Theories of agentic systems include those in distributed artificial intelligence and multi-agent systems. Wooldridge (2009) called intelligent agent's autonomous entities that perceive environments and act to reach certain goals. These agents more often now come into contact with electronic health systems, medical devices, scheduling systems, and clinical databases in a healthcare environment.

Some empirical evidence now appears indicating that agentic systems can improve the responsiveness of healthcare systems and the scalability of their operations. In 2026, researchers proposed the clinical agentic workflow that could dynamically adapt decision-making paths via state transition management and knowledge-enhanced reasoning (CDAFlow, 2026). Likewise, Yang et al. (2026) introduced the concept of "Internet of Agentic AI" that enables distributed agents to work together in the cloud and edge environments to ensure the scalability of healthcare coordination.

Agentic workflows come with significant governance issues, even with the potential benefits. When decision-making is autonomous, accountability, auditability and oversight are raised. The need for human-in-the-loop (HILo) governance structures to stop AI from engaging in unsafe or opaque actions in clinical settings is increasingly becoming a critical area of focus.

Figure X. Self-Governing Healthcare Systems and Agentic Workflows

Autonomous AI agents collaborate to plan, reason, and act across healthcare operations under human oversight.

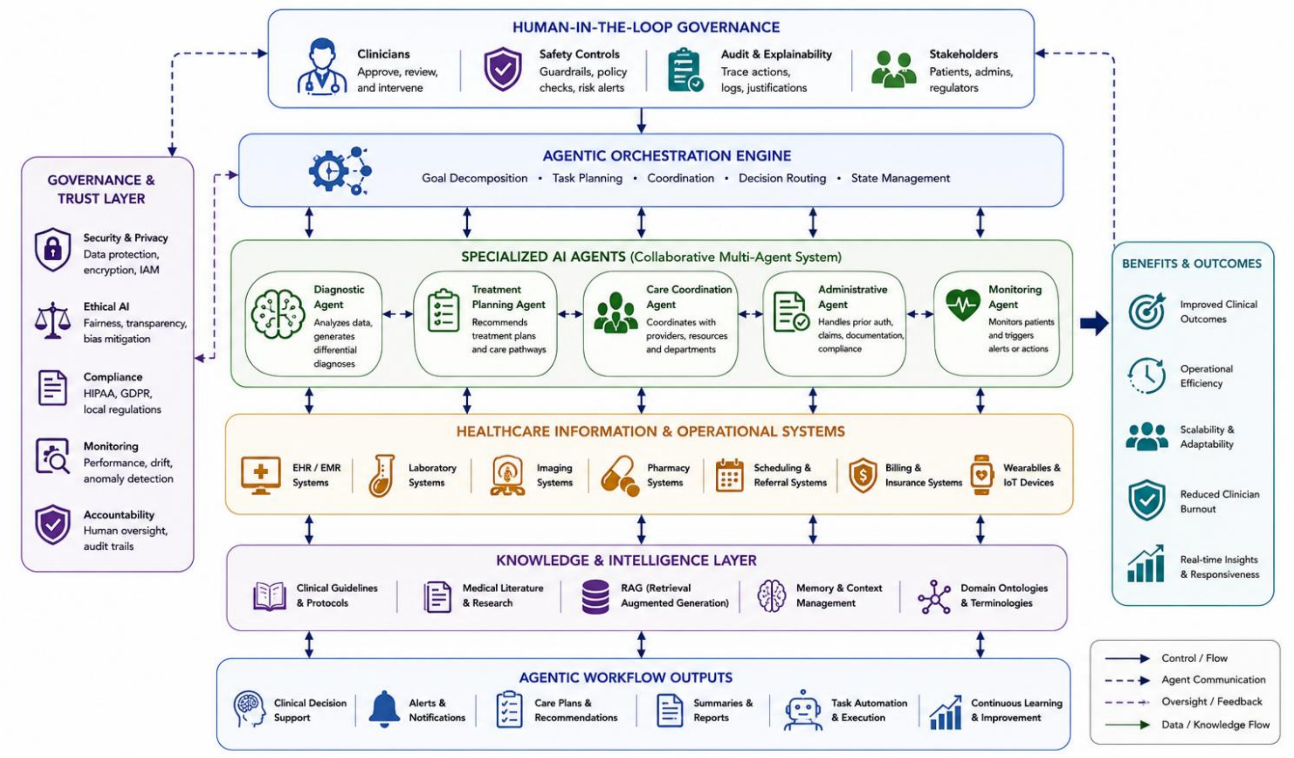

Source: Author's conceptual synthesis based on Joy (2025), Barra et al. (2025), Wooldridge (2009), CDAFlow (2026), Yang et al. (2026) and related literature.

Scalability for Cloud-Native Infrastructure and Healthcare.

Cloud computing has become a key component and enabler of the modern digital health system. However, traditional monolithic healthcare software systems can find it challenging to meet the flexibility, interoperability, and computational requirements of an AI-powered healthcare landscape. Cloud-native infrastructures can solve these challenges with a modular, elastic, resilient and distributed orchestration approach.

The emerging trend of containerization, like Docker, and Kubernetes as orchestration platforms has made them focal for scalable deployment of AI. Burns et al. (2016) stated the concept of container orchestration completely transformed distributed systems by offering automated capabilities to scale, discover services and manage faults. Cloud-native infrastructures are used in the healthcare industry for processing imaging data in real-time, genomic analytics, telemedicine applications, and AI inference workloads.

There has been a growing number of studies that have focused on the intersection of cloud-native architectures and healthcare AI systems in recent years. In their article, "Kubernetes-native orchestration patterns for multi-agent healthcare

LLM systems" (2025), Ayyagari et al. examined various orchestration patterns suitable for multi-agent healthcare LLM systems, highlighting the role of service meshes, interoperability frameworks, and zero-trust security models in enabling scalable healthcare AI operations. Their results found the need for Infrastructure-Layer Governance to ensure compliance to HIPAA and Interoperability standards.

The flexibility of healthcare intelligence systems has also grown with the advent of serverless computing and event-driven architectures. Additionally, the flexibility of healthcare intelligence systems has been extended by the rise of serverless computing and event-driven architectures. These principles of cloud also enable AI services to scale up and down with the changing clinical need and to accommodate continuous integration and deployment pipelines. This feature is especially useful in emergency care settings where there can be extremely large variations in computational loads.

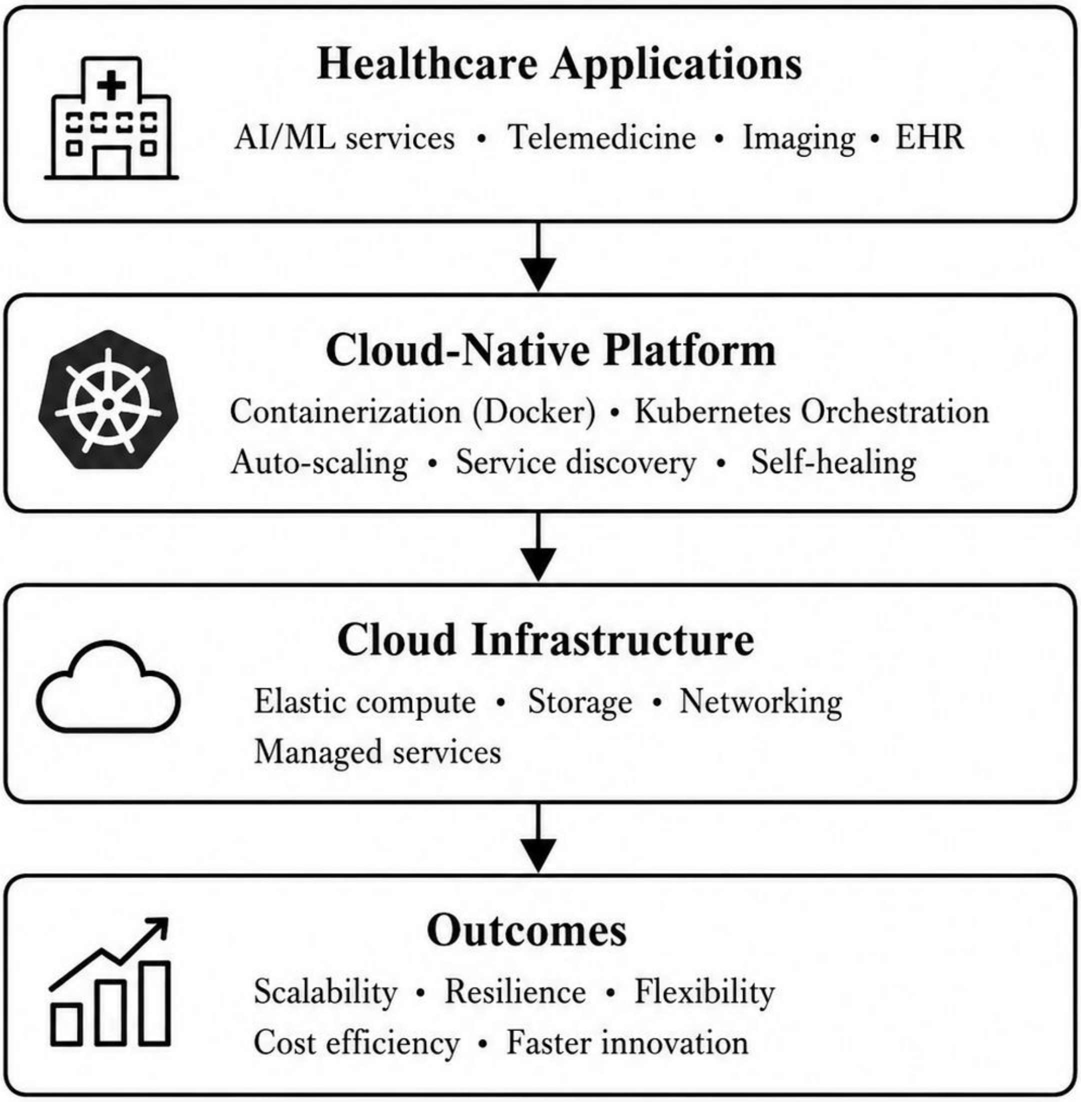

Figure X. Cloud-Native Infrastructure for Scalable Healthcare.

Source: *Author's construction based on Burns et al. (2016) and Ayyagari et al. (2025).*

Integrating information systems across the healthcare spectrum: Interoperability and Healthcare Intelligence

One of the ongoing challenges for effective healthcare intelligence is interoperability. Data in health care environments is often stored in a complex and inconsistent manner in different platforms and formats. The Fast Healthcare Interoperability Resources (FHIR) initiative has become a big hit for healthcare data exchange, which allows healthcare systems to communicate with each other in a structured manner.

Mandel et al (2016) pointed out the need for interoperability standards to support modern digital health ecosystems. The seamless integration of data environments is crucial for supporting generative AI and agentic workflows, as these systems need to access accurate and contextual patient information to make informed decisions.

Interoperability mechanisms are thus crucial for enabling these distributed systems to coordinate safely and effectively, which is the necessity for the convergence of generative AI, agentic systems, and cloud-native infrastructure. If communication layers are not standardized, then autonomous systems of healthcare intelligence can end up creating more fragmentation than they are solving, which raises numerous issues.

Research Gap

While there has been substantial advancements within these specific areas, the existing literature has not been comprehensive about healthcare AI ecosystems. The majority of studies are directed mostly on model performance or on the optimization of single tools of automation or on optimizing the infrastructure without adequately considering the interaction between these technologies as integrated systems. The conceptual considerations around integrating generative AI, agentic orchestration, and cloud-native infrastructure into a unified architecture that can enable scalable, healthcare intelligence are limited. This study aims to fill this gap by providing a synthesis of the various disciplines in an overarching architectural framework that not only focuses on AI ability but also on orchestration, governance, interoperability, and in frastructure resilience.

Methodology

The research method of this study is a conceptual and system approach with a research and development strategy that is used to integrate multidisciplinary knowledge into a comprehensive architectural approach to health care intelligence. The study does not rely on experimentation with model training or clinical trials, but draws on insights from peer-reviewed academic research, healthcare informatics frameworks, cloud computing studies, and recent research on agentic AI systems.

The Methodological Approach is Divided into 3 Interrelated Phases.

This phase entails a review of the literature. This phase involves literature review. The First Phase consisted of a systematic literature review of academic and industry literature that was published between the years of 2016 and 2026.

Public databases like IEEE Xplore, PubMed, ACM Digital Library, ScienceDirect, SpringerLink and arXplore were explored using the following keywords: generative AI for healthcare, agent-based healthcare workflows, cloud-based healthcare infrastructure multi-agent healthcare systems, LLM orchestration, and healthcare interoperability. The study's main aim was to investigate common architecture patterns, operational constraints, governance issues, and implementation patterns surrounding AI-powered healthcare systems. A special emphasis was placed on research with:

a) Healthcare's use of LLM.
b) Autonomous AI agents
c) Multi-agent orchestration systems
d) Cloud-native deployment models
e) Healthcare interoperability frameworks
f) Governance and compliance of AI content.AI governance and compliance.

The synthesis process excluded studies that were not based on empirical research or that were not clearly methodologically.

Phase 2: Architectural Abstraction (3rd – 5th Grade)

The second stage entailed abstracting the common system components from the literature examined and structuring them in conceptual layers. The principles of systems engineering were applied to pinpoint functional dependency among infrastructure services, systems orchestration mechanisms, interoperability frameworks and AI reasoning modules.

The resulting architecture is a multi-layered structure with the following elements:

a) Data acquisition and interoperability layer, or DAI.DAI, Data acquisition and interoperability layer.
b) This is a layer that handles knowledge and context management.
c) Generative AI reasoning layer
d) Agentic orchestration layer
e) A governance and compliance layer: A governance and compliance layer:
f) Cloud-native infrastructure layer
g) Human supervision and feedback overlaid
h) The layers were analyzed in terms of the requirements of scalability, resilience, explainability, interoperability and governance.

The analytical evaluation is the second phase. The second phase is called Analytical Evaluation.

Finally, the proposed architecture was assessed against the reoccurring healthcare issues found in the literature, such as clinician burnout, disjointed workflows, delayed decision making, security threats, and scalability of the healthcare infrastructure.

The analytical framework was directed towards four criteria of evaluation:

a. Operational adaptability
b. Infrastructure resilience
c. Clinical explainability
d. Regulatory compliance

Unlike quantitative measurements, the study takes an analytical approach to evaluating the potential benefits of architectural convergence on the ability to provide intelligence in healthcare compared to use of AI in isolation.

Results and Discussion

The United Healthcare Architecture for Intelligence The core result of this research is creating a unified architecture which combines generative AI, agentic workflows, and cloud-native infrastructure and creates a seamless healthcare intelligence ecosystem.

The architecture is based on multiple synchronized coordination instead of disparate functions. The difference is significant because a poor model is not the only thing that can lead to the failure of an AI deployment in healthcare.

Data acquisition and interoperability: Data acquisition and interoperability:

The foundation layer is composed of interoperable health care data services that are used to connect electronic health records, imaging systems, wearable devices and laboratory systems, genomic data repositories, and operational databases.

Event brokers, streaming data pipelines and HL7 standards facilitate real-time data exchange between distributed environments. This layer converts the disorganized data of healthcare information into a coherent stream of information that can be used by AI systems.

One of the key lessons that have come from the analysis is that, interoperability is not just a matter of convenience, but a fundamental requirement for the success of a healthcare intelligence system. To function properly, agentic systems must be constantly aware of the situation. If AI agents are not capable of integrating with other systems, they may rely on outdated or incomplete clinical data.

Phase 3: Knowledge and Context Management

The contextual knowledge layer is on top of the interoperability layer and integrates a variety of elements such as vector databases, retrieval augmented generation systems, medical ontologies, and memory management systems.

Next-generation healthcare intelligence involves dynamic context retrieval, as opposed to previous AI systems that only considered static model parameters. The layer allows AI agents to use the latest medical guidelines, protocols, patient histories, and operational context in reasoning.

Memory management systems also resolve one of the primary drawbacks of the current generation of LLM systems: temporal inconsistencies within the context. Persistent memory architectures enhance the continuity of managing patients over a long time and multi-step workflow executions.

Using Generative AI to reason and solve problems

The generative AI layer essentially serves as the cognitive power of the architecture. Large language models (LLM) are capable of handling multimodal healthcare information, creating summaries, aiding in diagnostic reasoning, and enabling conversations. But the study does not consider AI to be an autonomous power, it says. Rather that, these are models of probabilistic reasoning systems within larger governance systems.

This is important in clinical settings. Hallucinations, statistical uncertainty, and contextual errors are limitations with generative AI systems. When they are included in a workflow and retrieval-based system, the opportunities for unsafe outputs are minimized.

The results also reveal that specialized healthcare systems are superior to general ones for the clinical terminology, diagnostic reasoning, and medical summarization tasks. Thus, hybrid models of foundation models and fine-tuning will largely be needed for future healthcare intelligence systems.

Phase 4: Agentic Orchestration

The proposed architecture has the agentic orchestration layer as the most transformative element. In this place autonomous or semi-autonomous agents coordinate, invoke, manage and communicate across distributed healthcare environments.

These agents can be used as workflow participants and can:

1. Scheduling patient appointments
2. Coordinating referrals
3. Triggering diagnostic workflows
4. Monitoring patient deterioration
5. Summarizing clinical encounters
6. Scooping emergency cases to the clinicians
7. Managing operational logistics

Most significantly, the research reveals the change in agentic intelligence from the passive support of AI in healthcare to proactively orchestrating it.

This change has significant ramifications for health care processes. In the past, AI systems were fed information and required human input to initiate the process. Agentic systems, on the other hand, constantly track operational conditions and take own actions, without relying on others, within a pre-established governance area.

However, with this greater freedom comes a greater risk of governance. Prompt injection, data poisoning, using unauthorized tools, and opaque reasoning chains are vulnerabilities introduced by the addition of the autonomous workflows. This means that it is essential that healthcare organizations put in place strict oversight systems, audit trails, policy engines, and authentication systems.

Phase 5: Governance and Compliance

The governance layer acts as the ethics and regulation governing architecture system. This layer allows for compliance policies to be enforced, identity verification, audit logging, bias monitoring, standards for explainability and mechanisms for human oversight.

The analysis shows that governance is not just an administrative process that is external to the school. However, governance systems need to be integrated seamlessly into the healthcare AI workflow.

The embedded governance model is compatible with the new ideas of "policy-aware AI orchestration" that ensures that all AI actions are subject to checks based on institutional and regulatory limitations before they are taken. In healthcare, this can be particularly crucial given the sensitive nature of patient information and the repercussions of algorithmic errors.

The results also indicate that explainability should not be viewed as a reporting after-the-fact activity, but rather an operational capability. It is crucial for clinicians not only to understand the output of AI, but also the steps in the reasoning chain, evidence retrieved by the AI, and the events that are triggering the steps in the chain.

Phase 6: Cloud-Native Infrastructure

The cloud-native layer ensures scalability and resilience across the enterprise for healthcare intelligence.
Healthcare AI services can dynamically scale, be fault tolerant, and operate seamlessly through the use of container orchestration platforms, microservices architectures, serverless functions and service meshes.

The analysis shows that the cloud-native infrastructure will be especially relevant to agentic systems as the computational needs of autonomous workflows are not predictable. Healthcare systems can use event-driven architectures to respond to changes in patient demand, such as during emergency situations or when handling more complex workflows, by allocating resources adaptively.

Moreover, it is possible to keep models up to date on cloud-native environments, to monitor infrastructure, and to manage distributed workloads. As healthcare intelligence systems progress from static to always learning environments, these capabilities become crucial.

Use of human oversight and Augmented Intelligence

The study re-emphasizes the importance of adhering to a fundamentally human-centric approach to healthcare intelligence, despite the growing level of automation.

The proposed architecture is based on "augmented intelligence" and not complete clinical autonomy. Clinicians have the final say on key decisions, and AI systems support information integration, workflow management, and efficiency.

This focus on humans represents one of the core issues around healthcare AI implementation: that it could replace clinical intuition and/or patient trust. The architecture outlined here actually places AI in the role of an infrastructure to complement the human capacity building in healthcare, rather than supplanting healthcare professionals.

Broader Implications

This combination of generative AI, agentic workflows, and cloud-native infrastructure is not just a technological shift; it's a structural change in the way healthcare is done.

Healthcare intelligence is evolving from standalone software to ecosystems that can continuously learn, coordinate and adapt. This metamorphosis mirrors other enterprise computing trends – such as the shift from intelligence in individual enterprise applications to intelligence in the interrelated operational network.

But analysis also shows that technological capability is not enough. Trust in institutions, workforce readiness, interoperability standards and governance maturity will all play a vital role in the success of next generation healthcare intelligence.

Policy Implications

As AI systems become increasingly autonomous and health infrastructure becomes interconnected, the challenges of policy and regulatory coordination arise.

Regulatory Modernization

Current healthcare regulations have been primarily designed for deterministic software systems and not for autonomous reasoning AI agents. Therefore, regulatory frameworks need to adapt to deal with these challenges such as:

1) Autonomous decision accountability
2) AI explainability standards
3) Dynamic model updates
4) Cross-agent coordination
5) Real-time auditability

Healthcare regulators will probably need to create standards for certification for agentic healthcare systems.

Data Security and Privacy

In this study, the proposed architecture heavily relies on continuous data exchange in distributed systems. As a result, policies of data governance have to deal with the issue of interoperability while not infringing on patient privacy.

Zero-trust architectures, encrypted federated learning environments, encrypted differential privacy techniques and policy-aware access controls should be an integral part of healthcare AI infrastructure.

Workforce Transformation

Agentic healthcare systems will dramatically change the workforce dynamics in healthcare. Administrative functions will be increasingly automated and clinicians will shift into supervisory, interpretive and patient-facing roles.

Healthcare providers have, therefore, to invest in AI workforce training programmes with a focus on AI literacy, workflow supervision and digital governance.

Ethical Oversight

Ethical governance systems need to go beyond high-level principles of AI to mechanisms for its enforcement. Monitoring for bias, fairness auditing, validation of transparency, and explainability testing must not be a "compliance program". Monitoring for bias, fairness auditing, validation of transparency, and explainability testing must be part of deployment pipelines, not optional compliance.

It is important to note that there needs to be a different approach to healthcare AI – it is not just a technological problem. Autonomous healthcare intelligence systems will impact trust, equity, institutional power dynamics and healthcare access.

Conclusion

The article explored how generative AI, agentic workflows and cloud-native infrastructure are coming together to create next-generation healthcare intelligence. The researchers claimed that the introduction of individual AI models was not enough to bring meaningful change in the healthcare sector and that coordinated architectures with capabilities in reasoning, orchestration, interoperability, governance and scalable infrastructure will play a key role in driving the change.

The paper presented a layered design architecture for the system, comprising interoperable data environments, contextual knowledge systems, generative reasoning models, autonomous workflow agents, governance mechanisms and cloud-native deployment infrastructures.

The results indicate that agentic systems for healthcare intelligence could have a great impact in operational responsiveness, clinical efficiency, scalability and decision support. But there are significant governance issues with explainability, accountability, security and ethical oversight.

The research also finds that the ultimate goal of healthcare AI is to build intelligent healthcare systems, not to replace medical doctors! The future of intelligent healthcare systems is expected to involve more AI reasoning in decision-making, autonomous orchestration, and resilient infrastructure, which will be crucial for creating adaptive and continuously learning healthcare environments.

Empirical validation of unified healthcare intelligence architectures through pilot deployment, longitudinal studies of healthcare workflows and comparisons of the architectures across healthcare organizations are all areas that will be a focus of future research. There is also a lot more work to be done on the issues of patient trust, organisational readiness and regulatory harmonisation in the increasingly autonomous healthcare landscapes.

References

1. Ayyagari, V., Kagga, S. R., Joshi, S., & Bodagala, G. L. P. (2025). Cloud-native orchestration patterns for multi-agent healthcare LLMs. International Journal of Applied Mathematics, 38(12s). https://doi.org/10.12732/ijam.v38i12s.1392
2. Barra, F. L., Rodella, G., Costa, A., et al. (2025). From prompt to platform: An agentic AI workflow for healthcare simulation scenario design. Advances in Simulation, 10(29). https://doi.org/10.1186/s41077-025-00357-z
3. Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., et al. (2021). On the opportunities and risks of foundation models. Stanford Center for Research on Foundation Models.
4. Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., & Amodei, D. (2020). Language models are few-shot learners. Advances in Neural Information Processing Systems, 33, 1877–1901.
5. Burns, B., Grant, B., Oppenheimer, D., Brewer, E., & Wilkes, J. (2016). Borg, Omega, and Kubernetes. Communications of the ACM, 59(5), 50–57. https://doi.org/10.1145/2890784
6. Joy, M. (2025). Agentic workflows in healthcare: Advancing clinical efficiency through AI integration. International Journal of Scientific Research in Computer Science, Engineering and Information Technology, 11(2), 567–575. https://doi.org/10.32628/CSEIT25112396
7. Mandel, J. C., Kreda, D. A., Mandl, K. D., Kohane, I. S., & Ramoni, R. B. (2016). SMART on FHIR: A standards-based, interoperable apps platform for electronic health records. Journal of the American Medical Informatics Association, 23(5), 899–908. https://doi.org/10.1093/jamia/ocv189
8. Shanafelt, T. D., Dyrbye, L. N., Sinsky, C., Hasan, O., Satele, D., Sloan, J., & West, C. P. (2016). Relationship between clerical burden and characteristics of the electronic environment with physician burnout and professional satisfaction. Mayo Clinic Proceedings, 91(7), 836–848. https://doi.org/10.1016/j.mayocp.2016.05.007
9. Wang, X., Liu, Z., Zhang, C., et al. (2023). Large language models in medicine: Opportunities, applications, and limitations. Journal of Medical Systems, 47(1), 1–12.
10. Wooldridge, M. (2009). An introduction to multiagent systems (2nd ed.). Wiley.
11. Yang, Y.-T., & Zhu, Q. (2026). *Internet of agentic AI: Incentive-compatible distributed teaming and workflow* (arXiv No. 2602.03145). *arXiv*. https://doi.org/10.48550/arXiv.2602.03145

CITATION

Telharkar, A. (2026). Convergence of Generative AI, Agentic Workflows, and Cloud-Native Infrastructure: A Unified Architecture for Next-Generation Healthcare Intelligence. In Global Journal of Research in Engineering & Computer Sciences (Vol. 6, Number 3, pp. 26–35). https://doi.org/10.5281/zenodo.20597913