SIGNAL 1

It’s Time to Start Thinking about Your Agentic Tech Stack

Health systems are no longer managing a handful of AI applications. They are preparing to operate hundreds of AI agents developed by internal teams, embedded within enterprise software, and delivered by an expanding ecosystem of technology partners.

Unlike the recent wave of gen AI applications, these agents won't simply create content. They will retrieve information, make decisions, coordinate workflows, and take action across clinical, operational, and administrative systems.

Three shifts make this moment different from anything healthcare has navigated before:

1. **Agentic AI will reshape the economics of care.** Agents can autonomously carry out a wide range of tasks, creating the potential for enormous improvements in productivity and enabling revenue growth without proportional workforce growth.
2. **Health systems are becoming AI builders, not simply AI buyers.** Organizations are increasingly developing and customizing AI capabilities themselves while also adopting agents from an expanding ecosystem of software vendors.
3. **Enterprise AI is becoming modular rather than vendor defined.** Health systems can select the foundation models, agent platforms, cloud providers, and applications that best fit their needs instead of relying on a single software vendor's system.

Together, these shifts mean health systems **must take a deliberate approach to building their enterprise AI infrastructure.**

Cloud providers and data platforms have published reference architectures for years. What healthcare lacks is a common blueprint for how enterprise AI should be governed and operated, and, more importantly, a point of view on what has to be true for that governance to work.

This paper presents that blueprint. As health systems start building and deploying agents of their own, every AI stack will look different, but each one needs the same operational layer underneath it: an Agent Control Plane that governs, operationalizes, and continuously improves every agent, regardless of who built it or where it runs. The point of view we put forward here has been informed by conversations with over a dozen CIOs and architects from health systems at the forefront of agentic AI adoption.

A Reference Architecture for Enterprise Agentic AI

The reference architecture below illustrates the core components required to operate agentic AI at enterprise scale. It combines the foundational technologies already present in most health systems with the new operational capabilities required to govern, secure, monitor, and continuously improve AI agents. While every organization's implementation will differ, the functions represented here are broadly applicable regardless of the vendors, models, or agent frameworks selected. Wherever possible, we include examples of technologies currently available in each layer to make the architecture concrete and actionable.

Our goal in providing this is to help health systems understand what is necessary to safely operate agents at scale in a healthcare setting and to identify the critical components they should be looking to invest in.

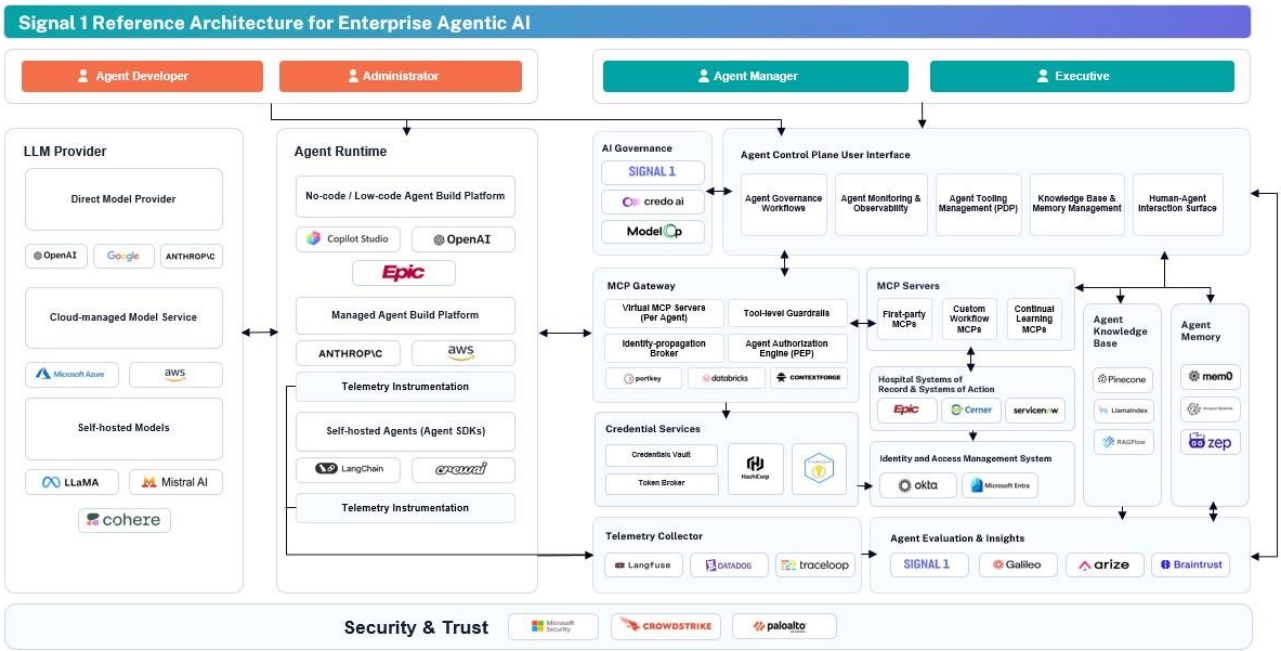

The architecture depicted above is best explained by following the journey of a single agent:

Before the first governed agent is deployed, the health system establishes the shared infrastructure that every agent will inherit. The **MCP Gateway** becomes the controlled path between agents and hospital systems, supported by **Credential Services** and integrated with the hospital's existing **Identity and Access Management System**. The **Telemetry Collector** begins capturing activity from the outset, while the **Agent Control Plane User Interface** provides the central surface through which people govern and oversee the fleet. Beneath it all, the hospital's existing **Security & Trust** capabilities continue to provide runtime, data, network, and threat protection. The control plane does not replace the security stack; it extends it to agents as a new class of actor.

An agent may then be created through any of the health system's **AI Agent Builders**, whether an enterprise platform, a no-code tool, or a self-hosted framework, and may use **Model Providers** delivered directly, through a managed service, or within the hospital's own environment. Before the agent can operate, it is onboarded to the **Agent Registry**, where its owner, purpose, risk tier, approval status, evaluation requirements, human-review policy, and access scope are established through **Agent Governance Workflows**.

Once approved, the agent is equipped with the capabilities and context it needs. The agent's governed **Virtual MCP** server exposes approved functions, which connect the agent to the appropriate **Hospital Systems of Record** and **Systems of Action,** such as the EHR. Access is limited to what the agent needs for its assigned role, while **Knowledge Base & Memory** grounds it in the hospital's current policies, procedures, workflows, escalation paths, and relevant prior context. This is where a generic agent becomes an agent configured for a specific organization and task.

When the agent begins work, each action passes through the **MCP Gateway**. The control plane identifies the agent and, where applicable, the person on whose behalf it is acting; verifies that the action is authorized; applies guardrails; and brokers the credentials required by the downstream system. The hospital's systems remain authoritative. The control plane governs the path into them and preserves accountability across it.

As the agent reasons and acts, its traces and tool calls flow through the **Telemetry Collector**, creating an auditable record of what it did, which systems it used, under whose authority it acted, and what happened next. **Agent Monitoring & Observability** allows teams to track agents' behavior, performance, reliability, and cost. The **Human-Agent Interaction** surface gives clinicians and operational staff a place to work directly with agents, review human-in-the-loop requests, approve or reject proposed actions, and intervene when human judgment is required.

The lifecycle does not end when the task is complete. The control plane evaluates whether the agent acted correctly against the health system's own policies, workflows, and situational context, then combines those findings with operational outcomes and human feedback. The resulting insights flow back into the **Agent Memory layer** and are served back to the agents via **Continual Learning MCP tools,** closing the loop from operation to improvement.

Open standards such as MCP and OpenTelemetry preserve portability across AI Agent Builders, Model Providers, and vendors as the market evolves. We designed this architecture to be practical: it names real product options for each major component and supports health systems whether they assemble the stack from multiple vendors or adopt more integrated platforms.

The Path Forward

Every health system's AI stack will evolve differently, but the operational foundation used to govern it shouldn't have to. A common Agent Control Plane is necessary to allow organizations to operate new agentic AI capabilities.

This paper is a starting point, not the full picture. For the complete technical blueprint with detailed design patterns, integration paths, and implementation guidance for each layer of the Agent Control Plane, see our technical companion document, ***Reference Architecture for Enterprise Agentic AI in Healthcare***.