AI agents inside the hospital

What they can read, what they can reach, and who decides

IN BRIEF

01 Hospitals have put AI to work on the patient record at scale, and in every published case a clinician reads and signs what it drafts. Agents that act alone are arriving in booking, reminders and coding. Nothing clinical runs unsupervised.

02 Hospital work is the patient record first and the hospital's own protocols second. An agent has to read both, and most tools reach neither.

03 Public catalogues of ready-made connections do not contain a hospital's record system. It is reached through the interfaces the hospital already runs, and health has a standard for them.

04 A hospital needs one place where every tool an agent can call is listed, every call is checked against policy, anything that would change the record waits for the person who signs, and the record of it all can be put in front of an auditor.

05 All of this can run on infrastructure the hospital controls, in the country its law names. A few hospitals already say theirs does.

IN THIS PAPER

01 Where hospitals are today 2
02 The record is in nobody's catalogue 3
03 Three doors, one gate 4
04 Six ways into a hospital's systems 5
05 What an agent may do 6
06 Seven questions before an agent touches a patient record 6
07 What regulators already ask 7
08 Where to start 8

PART 01

Where hospitals are today

5,000
staff at Singapore General Hospital have used its note tool

39 sites
in New Zealand's public health service with a licensed scribe

44 days to 1
for a clinic letter in a Brisbane hospital's pilot

12,000
AI agents created by Singapore's healthcare professionals

The scale is in documentation. At Singapore General Hospital, over 5,000 healthcare staff had used the hospital's note tool for more than 67,000 medical and administrative notes by December 2025. New Zealand's public health service has licensed a scribe at 39 sites, and reports out-of-hours documentation down by up to 81%. At The Prince Charles Hospital in Brisbane, a six-month pilot in paediatric outpatients took letters to referring doctors from 44 days to within one day.

Staff are building their own. Singapore's health ministry said in August 2026 that healthcare professionals had created over 12,000 AI agents on the public health system's shared platform.

In every one of these, a clinician signs. In New Zealand clinicians review, edit and confirm each note, and accountability stays with them. The Royal Children's Hospital in Melbourne says clinicians always review and approve the notes. Hong Kong's Hospital Authority, trialling generative AI for medical reports in six public hospitals, has said every report is checked and authorised by the doctor. In Abu Dhabi, the physician assistant on the emirate's health information exchange leaves clinical assessment and treatment decisions with the physician.

Some run it on their own infrastructure, and say so. Asan Medical Center in Seoul puts sepsis protocols, tests and antibiotic guidelines in front of clinicians within seconds, from servers inside the hospital, on a closed network with no connection to the internet. Hallym University Sacred Heart Hospital, also in Korea, chose an on-premises system over cloud services. Oda Hospital in Japan adopted an on-premises environment for nursing summaries, and reports the median time for a summary falling from 16 minutes to 10. In Abu Dhabi the exchange's assistant runs in a cloud region inside the UAE.

It does not always work. New Zealand's public health service withdrew its first scribe, built in-house, after technical problems. Australia's device regulator has set out when a scribe is a medical device, and says it may take compliance action against anyone who supplies a non-compliant one.

So the question for a hospital is not whether AI may help with the record. It is who signs, where that is written down, and whether it holds.

PART 02

The record is in nobody's catalogue

327
servers in Docker's public catalogue on 19 November 2026

0
of them a hospital's record system

7 in 130
systems and services whose API description anyone can download

13 of 60
kinds of record one large system lets an outside caller change

The Model Context Protocol, MCP, has become the common way to give an agent a tool. A system is wrapped once as an MCP server, and any agent can use it.

The largest public catalogue of these servers is Docker's. We read every entry. On 19 November 2026 it held 327 servers. They cover databases, document stores, ticketing, monitoring, maps and web search.

Not one is a hospital's record system. Two of the 327 are about health at all: a record product for small clinics and a genomics service in a public cloud. There is no laboratory, radiology or pharmacy system, and nothing that speaks a health standard.

The record vendors' own efforts do not fill the gap yet. The largest publishes an MCP server, and it serves the vendor's documentation, not a patient's record. One integration platform that many hospitals run now speaks MCP from inside the hospital's own installation. Every large record vendor has also launched assistants of its own, for notes, for patients and for billing, sold inside its product. A hospital that keeps its record on its own infrastructure, or that must keep health data in the country by law, has to ask each one where it runs.

There is a second fact underneath this one. **Hospital record systems do not publish their API descriptions.** Of about 130 systems and services we surveyed, seven had an API description anyone can download, and every one of the seven was an open-source server or a national service.

What health has instead is a standard. The large record systems offer FHIR, a common interface with a common data model, and the two largest leave the description of theirs open on a test server for anyone to read. Every national profile we found, in Australia, New Zealand, Malaysia, Indonesia, Saudi Arabia, Japan and Taiwan, is built on the same release of it.

Three things follow from the standard. A hospital's own server describes itself: it lists the kinds of record it holds and the searches it allows, and that list is the one that counts, because each hospital switches on a different part. A system that calls it unattended signs in with a signed token, made with a private key whose public half the hospital registers, not with a shared password. And writing is narrow: on the largest vendor's test server, 13 of the 60 kinds of record can be created or changed from outside, and the rest can only be read.

So the route to a hospital's record does not run through a catalogue. It runs through the hospital's own interfaces, described by the hospital's own systems.

PART 03

Three doors, one gate

A hospital's work has documents at the edge, the record at the centre and the hospital's own rules over both. An agent platform needs a door to each.

The first door is for documents. A referral letter, a report from another hospital, a scanned form, a recording of a dictated note. On Katonic a user gives the agent the files in the conversation, and a recording is transcribed by a speech model the hospital serves. The standing documents, the policies, protocols, formulary, coding guidelines and checklists, are indexed once as knowledge, with each reader's permissions kept, so an agent answers from the document and cites the section.

One rule belongs here. **Where the record or the protocol is silent, the agent says so.** A draft discharge summary lists what the record does not hold. An answer about a policy says when the policy is silent. An agent that fills a gap in a patient's record has written something no clinician said.

The second door is the catalogue, for the edge. Katonic's catalogue of MCP servers is drawn from Docker's and carries Katonic's own additions. A server is deployed as a container inside the hospital's own cluster. This is how an agent reaches a reporting database, a document store, or a web search for published guidance. Where the hospital's integration engine speaks MCP, the engine comes in through this door too.

One caution applies to any catalogue. A container that runs inside the hospital may still call a service outside it. A search server must never be given a patient's details, and that has to be a policy on the tool, not a line in a prompt.

The third door is the hospital's own APIs, for the centre. The hospital gives Katonic the description of an interface it already runs: an OpenAPI file, or the address of its FHIR server, which describes itself. Each operation becomes a tool an agent can call.

- The description is checked before anything is created. The hospital sees how many tools it will produce and any warnings.
- For a FHIR server, the hospital ticks the kinds of record an agent may use. A discharge summary needs eight of them, not all.
- Tools that read and search are ready to use. Tools that would create or change a record are created switched off.
- A search returns the records themselves, a page at a time, and not the raw envelope the standard wraps them in.
- A tool can be tried with real values, or as a dry run that shows the exact request and sends nothing.
- Each connection has its own sign-in, including the signed token that record systems expect from another system, a limit on calls per minute, and a time limit.
- A document can travel with a call: a scanned referral attached, or a report fetched from the record.

No agent reads an image. Where imaging is involved, the agent reads the report.

All three doors lead to one gate. A document, a tool from a catalogue server and a tool from an imported interface sit under the same policies, with the same record.

THREE DOORS, ONE GATE

Door 01	Door 02	Door 03	All three
Documents.	**The catalogue.**	**Your own APIs.**	**One gate.**

An agent builder sees tools. The hospital's governance committee sees one inventory.

PART 04

Six ways into a hospital's systems

Most of what moves inside a hospital does not move through a modern API. Results, orders and admissions travel between systems as messages, and images in a format of their own. The work of giving an agent a way in is done in products the hospital already owns, and it follows a small number of patterns.

	THE SITUATION	WHAT THE HOSPITAL DOES	WHAT THE AGENT PLATFORM DOES
01	**The hospital's integration engine, or a server it runs in front of the record, already offers MCP**	Switches it on and issues a credential	**Connects to it, and sets a policy on each tool**
02	**The record system describes itself through the FHIR standard**	Gives its own server's address, registers a public key, and names the kinds of record agents may use	**Reads the description, signs in with a signed token, and creates the tools**
03	**A system produces or publishes its own API description**	Exports the file from its own installation	**Imports the file**
04	**The hospital's integration engine or API gateway already fronts the system**	Registers the agent platform as one more consumer	**Imports the engine's file**
05	**A laboratory, radiology or pharmacy system that speaks only messages**	Puts a REST or FHIR face over it in the engine it already runs	**Imports the file that product generates**
06	**A system with only a screen, such as a payer's or an authority's portal**	Records the steps in a screen-automation product	**Connects to that product's tool, for reading only**

Two things follow. Nothing is replaced: the agent comes in through the same door as every other consumer of the hospital's interfaces. And the oldest systems in a hospital are usually behind pattern 5.

That work is the hospital's, and it is worth doing once for every consumer, not for agents alone.

PART 05

What an agent may do

Reaching the record is half of it. The other half is deciding what an agent may do there, and being able to show the decision afterwards.

Every call is checked before it runs. On Katonic the order is fixed: is this user allowed at all, do the inputs carry anything that looks like an injection attack, and what does the hospital's policy say about this tool.

A policy gives one of four answers. Allow it. Allow it and record it. Hold it until a person approves. Block it. A policy can cover one tool, every tool on a server, or every tool, and it can apply to the whole hospital or to one team. Where two policies meet, the stricter one wins.

A tool that would change something waits. Each imported operation is rated by what it does: reading is low risk, changing is medium, deleting is high. When an agent is published, any tool that changes something and has no policy of its own is set to wait for approval. That is how nothing enters the record unsigned: the note, the code and the request each wait for the person whose job it is. A hospital can relax that for a tool, letting a routine booking go through, say. It has to decide to, and the decision is on record.

No agent diagnoses, prescribes or reads an image. Software that suggests a diagnosis, a treatment or a triage priority is a medical device in several markets, and an approval step does not change that.

If the policy cannot be read, the call is refused. A fault in the control must not become permission.

The agent acts as the person using it. Katonic passes the signed-in user's identity with each call. The value comes from the platform, not from the model, so an agent cannot be talked into acting as someone else. Where the hospital's interface enforces it, a clinician's agent sees the patients that clinician may see.

Every call leaves a record: who asked, which agent, which tool, what the policy decided and who approved.

There is always a way to stop. An agent, a tool or a whole server can be switched off at once.

PART 06

Seven questions before an agent touches a patient record

These are worth asking of any agent platform, ours included.

01	**Is there one list of every agent in the hospital, of every tool each one can call, and of which tools change something?**
02	**Who signs before anything enters the patient record, and where is that written?**
03	**Whose identity does the agent act under when it calls a system?**
04	**What does the agent do when the record or the protocol is silent?**
05	**Is any agent doing a medical device's job, and who checked?**
06	**Where does the model run, and does patient data leave the country on any request?**
07	**Can every call be shown to an auditor, with the decision that allowed it?**

PART 07

What regulators already ask

Few rules are written for AI agents, and no health regulator in the eleven markets we read has written one. Hong Kong's privacy regulator published guidance on agentic AI on 25 August 2026, for organisations of every kind. It recommends giving an agent "the minimum access rights necessary to complete the tasks at hand", a person in the loop for decisions likely to have a significant impact on individuals, and mechanisms for traceability and auditability. A Singapore minister told a health conference in August 2026 that as AI moves from recommending to acting, assurance must move "from assuring outputs to assuring actions". Everywhere else, agents fall under texts a hospital already lives with: professional standards, medical-device rules, health-records law and privacy law.

Across the regulators we read, the same expectations recur. The clinician stays responsible and the facility stays accountable. A person can override the system, and nothing clinical runs unsupervised. The patient is told, and consents where it matters. A system is validated before use and checked in use. The hospital keeps a register of the AI it uses, and records of what it did. Vendors are assessed.

What one regulator already says

Singapore's Ministry of Health and Health Sciences Authority published version 2.0 of their guidelines on AI in healthcare in March 2026. They are guidance, not law, and they do not use the word agent. They read like a specification for section 5 of this paper.

- They say "human oversight is mandatory for all Clinical and Clinical-Ops use of AI".
- They name three levels of oversight: a person in the loop, who keeps full control; a person over the loop, who supervises and can take over; and no person in the loop.
- Deployers "should not deploy AI solutions that make clinical decisions independently without human making the final decision".
- A hospital should keep "a registry of all AI solutions deployed", and "be ready to suspend AI deployment in the case of serious adverse events".
- They draw the medical-device line by example. A language model that uses patient data "to provide clinical recommendations to clinicians" is a device. Software that transcribes a consultation and gives no diagnosis or treatment recommendation is not.

MARKET BY MARKET

The Gulf. The UAE's federal law on information technology in health says that health information about services provided in the country "may only be stored, processed, generated or transferred outside the State in the cases prescribed" by the health authorities, with a fine of 500,000 to 700,000 dirhams for a breach. In Dubai, the health authority's policy on AI binds licensed facilities and professionals: every AI solution goes through independent validation by a third party, and one that may significantly affect people's lives "must be able to be overridden or their decisions reversed by designated end users". Saudi Arabia's device regulator treats software intended to diagnose, monitor, treat or manage a condition as a medical device, and expects one approved elsewhere to be validated locally. We could not open Abu Dhabi's or Qatar's own health texts, and a hospital there should take them from its authority.

South-East Asia. Singapore is set out above. In Malaysia, the Medical Council's guideline says doctors "remain responsible for their clinical decisions", and that patient data "must not be shared on any AI system without the patient's consent". Indonesia's regulation on medical records obliges every facility to keep electronic records and to connect them to the ministry's national platform, and a facility that outsources their storage must use an operator with data storage facilities inside the country. For Thailand and the Philippines we opened no health-specific text on AI.

PART 07 · CONTINUED

Australia and New Zealand. Australia's device regulator has drawn the clearest line. Software is likely a medical device if it "identifies, recommends or suggests diagnoses", does the same for treatments or investigations, "prioritises, triages or monitors patients", or generates clinical conclusions. It is not one if it only records, transcribes, summarises and drafts for a practitioner to review, edit and confirm, and a general disclaimer does not change the answer. The practitioner regulator adds that approval of a tool "does not change a practitioner's responsibility to apply human oversight and judgment". In New Zealand, the Medical Council told doctors in March 2026 that a doctor's registration "cannot be transferred to an AI system", and the public health service bars its staff from using generative AI "for any clinical decisions, or for personalised advice to patients". New Zealand's health privacy code allows health information to be disclosed abroad only on listed grounds.

So running agents on the hospital's own infrastructure does not take away the duty to validate, to tell the patient, or to keep a clinician responsible. It takes away the questions about where the infrastructure is, who else shares it, and which country's law reaches the data.

And the controls in section 5 are the evidence these texts ask for: a list of agents and their tools, a named person who signs, a record of each decision, and a way to stop.

PART 08

Where to start

Pick work where reading is slow, the decision stays with a person, and no patient is waiting on the agent.

THREE PLACES THAT FIT

Policy and protocol questions

Staff ask what the hospital's own policy says and get the answer with the section cited, and "the policy is silent" where it is. No patient information is involved, so it can run on day one.

Incident report triage

The agent classifies each new report by the hospital's own categories and proposes the severity its own matrix gives. The patient safety team decides.

Discharge summary

The agent drafts the summary from the stay, each line with its place in the record. The discharging clinician corrects and signs.

THEN, IN ORDER

01 Name the documents and the systems the agent needs, and which of the six ways each system is reached.

02 Decide the policy for every tool, and who signs each kind of output, before the agent is built.

03 Build it with the clinicians and staff who do the work today.

04 Run it on invented patients, and then beside the manual process.

05 Measure it against that process: time, and how often a person corrected it.

06 Go live for one service and one team.

A fuller list of 31 agents across nine areas is published with this paper.

About Katonic AI

Katonic AI is the enterprise platform for serving models and building, running and governing AI agents. It runs on your machines, with your models, under your administrators. Katonic is listed in the Red Hat Ecosystem Catalog and runs on Red Hat OpenShift. The platform is in production at a top-tier bank in Asia, air-gapped, inside the bank's own perimeter, and it has run in production for about 1,100 staff of a life insurer in Asia, on the insurer's own cluster, inside its network. It was founded in Sydney in 2020 by Prem Naraindas, with its global hub in Dubai and teams in Sydney, Mumbai and Dubai.

KATONIC.AI

SOURCES

All read on 20 November 2026.

01 Singapore General Hospital, article on its note tool for clinical staff. sgh.com.sg

02 Ministry of Health, Singapore, speech at the HIMSS26 APAC conference, 24 August 2026. moh.gov.sg

03 OECD.AI, policy entry on the AI scribe in New Zealand's public health service.

04 Metro North Health, Queensland, "AI scribe pilot", quality of care report 2025.

05 The Standard, Hong Kong, report on the Hospital Authority's trial, 20 February 2025.

06 The Royal Children's Hospital Melbourne, page on its AI scribe. rch.org.au

07 Asan Medical Center, news release, 4 June 2026.

08 The Korea Times, report on Hallym University Sacred Heart Hospital, 14 August 2025.

09 Dove Press, paper on nursing summaries at Oda Hospital, 16 December 2025.

10 Gulf Business, report on the physician assistant of Abu Dhabi's health information exchange, 10 November 2026.

11 Creative HQ, New Zealand, article on AI scribes in public hospitals, 8 April 2026.

12 Docker MCP registry, 327 servers, read in full on 19 November 2026.

13 HL7, the FHIR standard, release 4, and its profile for backend services; the open test servers of two record systems. hl7.org

14 UAE Federal Law No. 2 of 2019 on information and communication technology in health fields, in the English translation of Emirates Health Services.

15 Dubai Health Authority, Policy "Artificial Intelligence in the Healthcare", DHA/HISHD/PP-09. dha.gov.ae

16 Saudi Food and Drug Authority, guidance MDS-G010 on medical devices based on AI and machine learning.

17 Ministry of Health and Health Sciences Authority, Singapore, Artificial Intelligence in Healthcare Guidelines, version 2.0, March 2026.

18 Malaysian Medical Council, Guideline on the Ethical Use of AI in Medical Practice, version 1/2025.

19 Minister of Health of Indonesia, Regulation No. 24 of 2022 on medical records.

20 Privacy Commissioner for Personal Data, Hong Kong, media statement of 25 August 2026.

21 Therapeutic Goods Administration, Australia, "Digital scribes", updated 16 October 2026. tga.gov.au

22 Ahpra and the National Boards, Australia, "Meeting your professional obligations when using Artificial Intelligence in healthcare".

23 Medical Council of New Zealand, guidance on using AI in patient care, March 2026.

24 Health New Zealand, "Using generative AI and large language models"; Health Information Privacy Code 2020, rule 12.