VERSION 0.9 · NOVEMBER 2026

Healthcare Threat Monitor in plain terms

A free, open source security watchtower for hospitals. It reads the security logs a hospital already produces, spots cyberattacks and misuse of patient records early, and keeps a person in charge of every response.

Free
No licence fee, ever. Published on GitHub under the Apache 2.0 licence.

Open
Every line of code can be read, checked and changed by anyone.

Hospital first
Built around patient records, medical devices and clinical care.

Free and open source under the Apache License 2.0. Section 11 describes how every change and release is tested.

1 · IN ONE PARAGRAPH

What we are building

Healthcare Threat Monitor is software a hospital runs on its own computers. It collects the security records that the hospital's systems already write, such as who opened which patient record, which programs ran on which computer, and which devices talked to which addresses. It looks through them continuously for signs of an attack, groups related warning signs into one incident, explains why it is worried, and lets a member of staff decide what to do. When the evidence is strong, that person can cut an infected computer off the network with one approval.

All of it is published on GitHub, free of charge, so that any hospital, clinic or health organisation in any country can download it, run it and adapt it, whatever its budget.

2 · THE PROBLEM

Why hospitals need this

Hospitals have become a favourite target for cybercriminals, and they are harder to protect than most organisations. Four things make them different.

They cannot simply switch off

An office can shut down every computer when it is attacked. A hospital cannot, because operating theatres, wards and emergency departments depend on those systems. Attackers know this and expect hospitals to pay.

Many people open patient records

Hundreds of staff legitimately read records every day. A nurse looking up a celebrity or a neighbour looks almost the same as a nurse doing her job, so simple rules cannot tell the difference.

Medical devices cannot protect themselves

Infusion pumps, patient monitors and scanners often cannot run security software. The only way to watch them is to watch the network traffic around them.

Few staff, small budgets

Many hospitals, especially smaller and public ones, have no security team at all. Commercial monitoring tools are expensive and are usually built for offices, not wards.

A false alarm also costs more in a hospital than elsewhere. Cutting a clinical workstation off the network by mistake interrupts care. So any tool that responds to attacks must be careful, and it must leave the final call to a person.

3 · THREATS

What it watches for

The platform watches for four families of threat. Each one is spotted by a mix of fixed rules, which catch known attack behaviour, and learning models, which notice activity that is unusual for that person, computer or device.

Ransomware

Criminal software that locks a hospital's files until a ransom is paid. It usually follows a recognisable path, and the platform watches each step so it can raise the alarm before the files are locked.

Warning sign	Why it matters
A document or email program starts a script	This is how a malicious attachment usually begins to run.
A program copies passwords out of memory	Attackers steal passwords so they can move to other computers.
One computer logs on to many others	The shape of an attacker spreading through the network.
Backups or backup services are deleted or stopped	Done so the hospital cannot restore its files.
Many files renamed in a minute, or a ransom note written	The encryption itself has started.

Misuse of patient records

Staff opening records they have no reason to see, sometimes called snooping. The platform learns each person's normal pattern and compares it with their colleagues in the same role.

Warning sign	Why it matters
Many records opened outside the person's own department	Normal work is mostly within one's own ward or clinic.
A flagged record opened with no care relationship	Hospitals can flag records of public figures or staff, which are common targets for curiosity.
Unusual hours or volume compared with peers	The learning model notices patterns no single rule would.

Compromised medical devices

A device taken over by an attacker, used to spy, to spread, or to flood the network.

Warning sign	Why it matters
A device talks to a computer it never talks to	Devices normally speak only to their own gateway.
A device reaches for the internet	Clinical equipment should not need the open internet.
A device sends a flood of traffic	A sign it is being used to disrupt the network.

Account misuse and other intrusions

Warning sign	Why it matters
Many failed logons for one account	Someone may be guessing a password.
Remote access from a new country	The account may have been stolen.
Administrator rights given to an ordinary account, or a former employee's account used	Classic signs of an insider or an intruder.
A phishing email reaches an inbox, or a risky attachment is delivered	The usual first step of an attack.
Lookups of known malicious websites, internal scanning, or a very large upload to the internet	Signs of control by an attacker or of data being stolen.

In total the platform ships with 21 rules and 3 learning models. Every rule can be edited or switched off, and new ones can be added, without changing any code.

4 · HOW IT WORKS

From a log line to a decision

Think of it as a smoke detector connected to a security guard. The detectors notice the smoke, the guard looks at the evidence, and only the guard decides whether to call the fire brigade. Every record goes through the same seven steps.

1	2	3	4	5	6	7
Collect	**Detect**	**Group**	**Score**	**Review**	**Respond**	**Learn**
Logs arrive from ten kinds of hospital system.	Rules and learning models look for warning signs.	Signs about the same computer or person become one incident.	The incident gets a confidence: low, medium or high.	A person reads the evidence and decides.	On strong evidence, the person can isolate the computer.	Each decision teaches the detectors.

1. Collect

The platform accepts records from ten kinds of source: patient record systems (in the FHIR AuditEvent standard), endpoint security tools, Windows and Linux computers, firewalls, VPNs, DNS, email security, medical device network traffic and cloud services. Each record is translated into one common format, so a warning sign means the same thing whatever system it came from. Every sender must sign its records with a secret key, so a stranger cannot feed in fake data.

2. Detect

Rules catch behaviour that is known to be dangerous, such as deleting backups. Each rule is linked to MITRE ATT&CK, the standard catalogue of attacker techniques. **Learning models** study recent normal activity and flag what is unusual: one watches how each member of staff uses patient records, one watches what happens on each computer minute by minute, and one watches each medical device's traffic. They need no examples of real attacks to work.

3 and 4. Group and score

Separate warning signs about the same computer or person, close together in time, are joined into one incident, so staff see one story, not fifty alarms. The incident gets a confidence score built from how many independent detectors agree.

Low	Medium	High
Kept for tuning. Nobody is disturbed.	Sent to the review queue and to the hospital's other security tools.	As medium, and isolating the computer becomes possible.

One safeguard sits on top: an incident can only reach high if at least two different kinds of log agree. A single faulty or tampered sensor can therefore never lead to a computer being cut off.

5. Review

A member of staff opens the incident and sees, in plain language, what happened, which machines and people are involved, which attacker techniques it matches, and for the learning models, which

numbers looked wrong and by how much. They then confirm it as a real attack or dismiss it, with a reason.

6. Respond

On a high confidence incident the reviewer can approve isolating the computer from the network. The platform never does this on its own. The request goes to the hospital's endpoint security product through a connector, and the computer is only marked as isolated once that product confirms it acted. Releasing it works the same way. Incidents can also be passed on automatically to a central security system (a SIEM) in standard formats.

7. Learn

When a reviewer dismisses a false alarm, the detectors that raised it learn, for that person or computer only, so the same false alarm does not keep coming back. Confirmed attacks are kept out of the learning, so the system cannot be taught that an attack is normal. Every lesson is listed in the console and can be undone.

Recorded for good

Every alert, decision and action is written to an audit log that is chained together with cryptographic fingerprints, like links in a chain. If anyone changes or deletes a past entry, the chain breaks and the platform shows exactly where. The log can be exported and checked independently.

5 · THE PARTS

What a hospital actually installs

The console

A website staff open in their browser. It shows the overview, incidents, computers and devices, rules, models, the audit log and settings.

The background service

A program that receives logs, runs the rules and models, groups incidents, and passes isolation requests to the endpoint product.

The database

Stores the logs, incidents, decisions and audit log. Standard PostgreSQL.

All three run on the hospital's own computers or its own cloud account. No hospital data is sent to us or to anyone else. The platform also includes a built in hospital simulator, so it can be tried safely with made up patients, staff and attacks before it is connected to anything real.

6 · USING IT

A day with Healthcare Threat Monitor

The people who use it day to day are the hospital's IT or security staff. In a small hospital that may be one person who checks the console a few times a day. Clinicians never need to touch it.

Morning. The IT officer opens the overview. Every log source shows as reporting, and the review queue holds two medium incidents from overnight.

- **Review.** The first is a burst of failed logons: a nurse who forgot a new password. She dismisses it as expected activity. The platform learns this pattern for that account.
- **A real signal.** The second shows a ward clerk opening sixty records outside her department, including a flagged record. The officer confirms it and passes it to the privacy lead, with the evidence and audit trail ready.
- **An attack.** In the afternoon a malicious attachment runs on a pharmacy computer. Within about two minutes an incident appears. Within about seven minutes, as more signs join it, it reaches high confidence, well before the attack reaches the stage of locking files.
- **Response.** The officer reads the evidence, approves isolation, and the endpoint product cuts the computer off the network. The attack stops on that machine. Everything is recorded in the audit log and forwarded to the regional security team.

The timings above come from the platform's tests on its simulated hospital, described in section 9.

7 · GETTING IT

How a hospital starts

1. **Try it.** Anyone with Docker installed can run the whole platform on one computer with a single command, `docker compose up`, then open it in a browser and press "Try the demo". It runs the simulated hospital, so they can launch a ransomware attack, watch it being caught and isolate the machine. Nothing real is touched.
2. **Read and check it.** Because the code is public, the hospital's IT staff, an auditor or a partner can read exactly what it does.
3. **Install it for real, with clicks.** A hospital downloads one file from the project's GitHub page: a ready made virtual machine to import into VirtualBox, VMware or Hyper-V, or a USB installer that puts the platform on a spare computer. When it starts, its screen shows the address to open and a setup code. Someone opens that address in a browser on another computer and enters the code to create the first administrator, protected by a password and an authenticator app. Nothing is typed in a terminal. It has no simulator, no test attacks and no public demo login, it is served over HTTPS, and every page needs sign in.
4. **Set it up in the browser.** A setup wizard in the console takes the hospital's departments, staff and computers from ordinary spreadsheets. The patient record system connects with a form, routers and firewalls by pointing their logging at the platform, and Windows computers with a one file installer. No programming is needed. Section 8 explains each way of plugging in data.
5. **Let it learn.** For its first week with each kind of log, the platform learns what is normal in that hospital. Clear signs of attack are reported straight away. Activity that is only unusual is reported once learning ends, so a new hospital is not buried in false alarms.
6. **Tune it.** Thresholds, rules, wording and even the landing page are edited from the console's settings, without any programming.

8 · DATA

Plugging in a hospital's data

The platform does not ask a hospital to create any new records. It reads the logs that computers, network equipment and clinical systems already write in the background. What it can watch therefore depends on what the hospital already has. Even a hospital that keeps its patient records on paper usually has computers and an internet connection, and those already keep logs.

What a hospital has, and what it can watch

If the hospital has	It can watch for
Paper patient records, a few Windows computers and an internet router	Account misuse, unusual network activity and many signs of ransomware, from the security log every Windows computer keeps and the logs most routers can send. Checks on patient record use do not apply, because there are no electronic records to misuse.
An electronic patient record system	All of the above, plus misuse of patient records. Most record systems can produce a report of who opened which record, which is all the platform needs.
Endpoint security, a firewall or email security	The fullest picture of ransomware and phishing, and the ability to isolate an infected computer through the endpoint product.
Medical devices on the network	Devices behaving unusually, from the network logs around them, without installing anything on the devices.

The platform also needs to know the hospital's departments, its staff and their roles, and its computers and devices. This is how it can tell, for example, that a nurse opened records outside her own ward.

Available today, in the browser, with no programming.

- **A setup wizard.** Staff enter the hospital's name and time zone, then upload their departments, staff list and list of computers and devices as ordinary spreadsheets, or type them in. Templates are provided, common column names are understood, and a faulty file is refused as a whole with every problem listed by row number.
- **Connect the patient record database.** For PostgreSQL, MySQL and MariaDB. The hospital asks its vendor for a read only account (the console provides the request to send), types in the address and password, picks the access log table, and checks the column matching against a few example rows. The platform then reads new entries every 30 seconds. It opens every session as read only, so it cannot change a record even if the account allows it, and it stores the password encrypted.
- **Connect a FHIR server.** Many record systems offer FHIR, the health data standard. The platform reads its AuditEvent records, signing in the way the vendor requires, most often with a key pair the platform makes and the vendor registers.
- **Upload a records report.** Without database or FHIR access, staff upload the access report their record system exports. The column matching is remembered for the next file, and the same report uploaded twice is not stored twice.
- **Routers, firewalls and servers.** They already have a setting to send their logs to an address. Staff type in the platform's address, the device appears in the console, and they accept it. Guides cover FortiGate, pfSense, OPNsense, MikroTik, Cisco, OpenWrt, Ubiquiti, Sophos and Linux servers.
- **Windows computers.** One downloaded file, run once as administrator on each computer, installs a small collector that sends the computer's security events: sign ins, new accounts and administrators, programs started, stopped services and a cleared security log. IT teams can push the same file to many computers at once.
- **A Sources page** shows whether each connection, device and computer is working, when it last read, any problem in a plain sentence, and how the learning period is going.

Other systems can still send their records through the platform's signed connection, using the reference sender in the project. That route needs someone comfortable writing short scripts.

Not supported yet. Microsoft SQL Server and Oracle databases. Hospitals using them can connect through FHIR or upload reports.

Getting started with little or no digital data

1. **List what you have.** Count the computers, note the router or firewall, and check whether patient records are electronic. This decides what the platform can watch.
2. **Try the demo first.** Run the simulated hospital to see what the platform does before connecting anything real.

3. **Start small.** Begin with the computers and the router, where most attacks first show up, then add the patient record system and other tools as they are connected.
4. **Follow the installation guide.** The project's hospital installation guide walks one person through installing, the setup wizard, connecting records, routers and Windows computers, the learning period, and backups.

9 · EVIDENCE SO FAR

How well it works on a simulated hospital

We built a simulator of a mid sized hospital: 11 departments, 197 staff on day and night shifts, 118 workstations, 10 servers, 40 medical devices and 600 made up patients, producing about 115,000 log records a day with the ordinary noise of real work. Three kinds of scripted attack were mixed in, and the platform was scored against the known answers. The test was repeated three times with different random draws.

Measure	Result
Scripted attacks detected	36 of 36 before staff feedback, and 36 of 36 after
False incidents in the review queue each day	16 to 18 before feedback, 7 to 8 after
False incidents that reached high confidence	None
Ransomware reaching the review queue	About 100 seconds after its first malicious step
Ransomware reaching high confidence	About 6 to 7 minutes, ahead of the point where files are locked
Record snooping reaching the review queue	About 8 to 13 minutes

What these numbers show. The method works and can be measured, and anyone can rerun these tests with one command and get the same numbers. In each hospital the platform first learns that hospital's normal activity, so its own numbers, especially false alarms, reflect its own systems.

10 · TRUST

Privacy, security and compliance

- **The hospital keeps its data.** It runs on the hospital's own systems. Nothing is sent to us.
- **People stay in charge.** Nothing is ever isolated automatically, and only high confidence incidents backed by two independent sources can be isolated at all.

- **Strong sign in.** Every staff account needs a password and a code from an authenticator app. Accounts lock after repeated failures.
- **Tamper evident record.** The chained audit log shows any change to past entries.
- **Open to inspection.** The code, the tests and a written threat model describing how the platform itself could be attacked are all public.
- **Supports the standards hospitals answer to.** Its functions are mapped to the HIPAA Security Rule, the NIST Cybersecurity Framework 2.0 and the HHS 405(d) health industry practices, with live evidence for each. It supports these safeguards. It does not make a hospital compliant by itself.

11 · HOW IT IS TESTED

Every change and every release, checked automatically

Healthcare Threat Monitor is tested the way a hospital would use it, from a fresh install to catching an attack and restoring a backup. These checks run automatically on every change to the code, and again on every release before its downloads are published.

- **Detection.** A measured evaluation on a simulated hospital caught all 36 scripted attacks across three independent runs, and no false alarm ever became eligible for isolation. A shorter run on every change fails if a single attack is missed.
- **Full installs.** A hospital installation is started from nothing, checked to be locked down, backed up, removed, and restored onto a new one with every account, key and certificate intact.
- **The downloads.** Each release's virtual machine and USB installer are started, switched off part way through their first start, started again, set up in a browser and given a fixed network address, and the USB stick installs onto an empty disk with Secure Boot on.
- **Windows.** The collector is installed, run and removed on a Windows computer, and what it sends is checked.
- **Attack scanning.** CodeQL reads the code, and ZAP attacks a running installation as a visitor and as a signed in administrator.
- **Connections.** Records databases, a FHIR server, uploaded reports and router and firewall logs are each read end to end with sample data written to the published standards and vendors' documented formats.

Like any security tool, it is best tried first on a hospital's own network, with a person in charge of every response, as the platform is designed. Hospitals that run it are welcome to share what they find, openly on GitHub or privately.

What comes next

1. Measurements on real, properly governed log data from hospitals and clinics that run it, to report real false alarm rates.

2. Independent reviews by hospital security teams and researchers, welcomed through the project's security policy.
3. Ready made connectors for the endpoint, firewall and patient record products hospitals use most.
4. Contributions from hospital IT teams and security researchers, who can propose improvements openly.

12 · WORDS USED IN THIS PAPER

Short glossary

Log	A record a computer system writes about something that happened, such as a logon or a file opened.
Incident	A group of related warning signs that together suggest something is wrong.
Isolation	Cutting one computer off the network so an attack cannot spread from it, while the rest of the hospital keeps working.
Endpoint product	Security software installed on computers that can, among other things, isolate them.
SIEM	A central system where large organisations collect security alerts from many tools.
Open source	Software whose code is published for anyone to read, use and improve.
Apache 2.0	A widely used open source licence that allows free use, including in commercial and public organisations.
MITRE ATT&CK	A public catalogue of the techniques attackers use, shared by security teams worldwide.