REVISTA CIENTÍFICA SOL21
STANDARD OPEN LITERATURE - ISSN 3086-089X

OBSERVABILITY AS A PRECONDITION FOR DETECTION IN HOSPITAL ENVIRONMENTS: A TELEMETRY AND RESPONSE MODEL UNDER CLINICAL CONSTRAINTS

OBSERVABILIDADE COMO PRÉ-CONDIÇÃO DA DETECÇÃO EM AMBIENTE HOSPITALAR: UM MODELO DE TELEMETRIA E RESPOSTA SOB RESTRIÇÃO ASSISTENCIAL

LA OBSERVABILIDAD COMO PRECONDICIÓN DE LA DETECCIÓN EN ENTORNOS HOSPITALARIOS: UN MODELO DE TELEMETRÍA Y RESPUESTA BAJO RESTRICCIÓN ASISTENCIAL

Marcelo de Melo Marcon[1]

e2138
https://doi.org/10.67553/sol21.v.2i1.38

PUBLISHED: 10/2026

ABSTRACT
Intrusion detection literature assumes that an organization can contain an incident as soon as it identifies one. In hospital settings this assumption does not hold: isolating an endpoint, revoking a credential or segmenting a network may interrupt ongoing care and harm patients. This article proposes an observability and response model for healthcare organizations in which containment is conditioned by the clinical criticality of the affected asset. Design Science Research is adopted as the strategy, developing a model-type artifact from a documentary base of technical standards, government agency publications and peer-reviewed literature, including studies linking ransomware attacks to adverse clinical outcomes. The results comprise the formulation of clinical constraint as a property of the environment; the demonstration that mean time to respond has a floor imposed by patient safety, so that mean time to detect is the only genuinely compressible variable; a seven-layer telemetry architecture; response stratified into four levels of autonomy; detection use cases prioritized by clinical impact; five operational indicators; and a mapping against ISO/IEC 27001:2022, the NIST Cybersecurity Framework 2.0 and ISO 22301. The article concludes that under clinical constraint, observability determines which responses exist, not merely when they begin.

KEYWORDS: Observability. Intrusion detection. Hospital security. Care continuity. Incident response.

RESUMO
A literatura sobre detecção de intrusões pressupõe que a organização possa conter um incidente assim que o identifica. Em ambiente hospitalar, esse pressuposto não se sustenta: isolar um endpoint, revogar uma credencial ou segmentar uma rede pode interromper assistência em curso e produzir dano ao paciente. Este artigo tem por objetivo propor um modelo de observabilidade e resposta aplicável a organizações de saúde, no qual a contenção é condicionada pela criticidade clínica do ativo afetado. Adota-se a Design Science Research como estratégia, com desenvolvimento de artefato do tipo modelo a partir de base documental composta por normas técnicas, publicações de agências governamentais e literatura revisada por pares. Os resultados compreendem a formulação da restrição assistencial como propriedade do ambiente; a demonstração de que o tempo médio de resposta possui piso imposto pela segurança do paciente, de modo que o tempo médio de detecção é a única variável efetivamente compressível; uma arquitetura de telemetria em sete camadas; a estratificação da resposta em quatro níveis de

[1] Doutor Honoris Causa em Ciências Exatas pela World Christian University (WCU), Mestre em Direção Estratégica em Tecnologias da Informação pela Universidad Europea del Atlántico, pós-graduado em MBIS – Segurança da Informação e em Engenharia de Redes e Sistemas de Telecomunicações. Possui mais de 20 anos de experiência profissional em Tecnologia da Informação e Comunicação (TIC), com atuação em ambientes críticos nos setores de saúde, governo e telecomunicações. Sua experiência abrange Segurança da Informação, Cibersegurança, Proteção de Dados, Gestão de Riscos, Datacenters, Virtualização, Linux, Redes, Governança, Continuidade de Negócios e Melhoria de Processos. É instrutor da Escola Superior de Guerra – ADESG/Campinas, membro do Comitê Avançado e Altos Estudos de Tecnologia e Inovação da Digital Law Academy, Diretor do Comitê de Segurança da APDADOS e Imortal da Academia Brasileira de Ciências, Artes, História e Literatura (ABRASCI).

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

REVISTA CIENTÍFICA SOL21
STANDARD OPEN LITERATURE - ISSN 3086-089X

autonomia; casos de uso de detecção priorizados por impacto clínico; cinco indicadores operacionais; e o mapeamento contra a ISO/IEC 27001:2022, o NIST Cybersecurity Framework 2.0 e a ISO 22301.

PALAVRAS-CHAVE: Observabilidade. Detecção de intrusões. Segurança hospitalar. Continuidade assistencial. Resposta a incidentes.

RESUMEN
La literatura sobre detección de intrusiones supone que la organización puede contener un incidente tan pronto como lo identifica. En el entorno hospitalario ese supuesto no se sostiene: aislar un endpoint, revocar una credencial o segmentar una red puede interrumpir la asistencia en curso y causar daño al paciente. Este artículo propone un modelo de observabilidad y respuesta para organizaciones de salud en el que la contención está condicionada por la criticidad clínica del activo afectado. Los resultados comprenden la formulación de la restricción asistencial como propiedad del entorno; la demostración de que el tiempo medio de respuesta tiene un piso impuesto por la seguridad del paciente, de modo que el tiempo medio de detección es la única variable comprimible; una arquitectura de telemetría en siete capas; la estratificación de la respuesta en cuatro niveles de autonomía; casos de uso priorizados por impacto clínico; cinco indicadores operativos; y el mapeo contra la ISO/IEC 27001:2022, el NIST Cybersecurity Framework 2.0 y la ISO 22301.

PALABRAS CLAVE: Observabilidad. Detección de intrusiones. Seguridad hospitalaria. Continuidad asistencial. Respuesta a incidentes.

INTRODUCTION

A contemporary hospital depends on information systems in virtually every care process: the electronic health record organizes prescribing and clinical documentation, the imaging system distributes examinations, and infusion pumps, monitors and ventilators communicate over the network. When these systems become unavailable, the result is not administrative inconvenience: it is diagnostic delay, suspended procedures and ambulance diversion.

Empirical evidence on this causal chain is no longer anecdotal. Research using United States Medicare claims data found that in-hospital mortality rises by 34 to 38 percent among patients already admitted when a ransomware attack begins, alongside a 17 to 24 percent decline in hospital volume during the first week (Neprash et al., 2026). A qualitative study at four affected organizations documented longer emergency department stays, delayed testing and additional patient transfers (van Boven et al., 2024). A cyberattack on a hospital is therefore a patient safety event before it is an information security event. Financial cost tracks clinical cost: healthcare remains, for the thirteenth consecutive year, the sector with the highest average cost per data breach, at 6.64 million dollars per incident against a global average of 4.99 million, and the mean time to identify and contain a breach rose to 247 days (IBM, 2026).

The consolidated organizational response is to expand the detection and response stack. This architecture is appropriate and necessary, but it carries an assumption the hospital environment invalidates: that once an incident is identified, the organization can contain it. In a hospital, canonical

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

containment actions carry clinical cost. Isolating an endpoint may mean isolating the workstation displaying an imaging report during a procedure; revoking credentials in bulk may block access to the health record during an emergency admission; quarantining a segment may silence telemetry from intensive care monitors. None is inadmissible in absolute terms; all are conditioned by what is happening to the patient at that moment.

From this we formulate the proposition that guides the work: in healthcare organizations, mean time to respond has a floor imposed by patient safety, whereas mean time to detect has no equivalent floor. The two variables, treated as symmetrical by the literature, are not equally compressible. As a consequence, observability ceases to be instrumentation supporting detection and comes to determine which responses are available: the earlier and the more contextualized the observation of an event, the larger the set of clinically acceptable interventions, because containment can still occur before the asset enters care use.

We define the following research question: how should observability and incident response be structured in a healthcare organization where conventional containment actions are clinically constrained? The general objective is to propose a model that conditions response to the clinical criticality of the affected asset. The specific objectives are: (i) to characterize clinical constraint as a property of the environment; (ii) to demonstrate the asymmetry between detection and response times; (iii) to propose a telemetry architecture with clinical context; (iv) to stratify response by level of governed autonomy; and (v) to define indicators and map the model against normative references. The research is justified because the detection literature rarely treats the cost of containment as a constraint of the problem; because healthcare organizations keep orchestration platforms in informational mode only, for fear of clinical impact; and because continuity of care is a legal interest distinct from confidentiality. The model was built from public literature and professional experience, without identifying the organization, for the reasons given in subsection 2.3.

1 - THEORETICAL FRAMEWORK

1.1 Monitoring, observability and telemetry

The three terms are used interchangeably in practice, yet they designate distinct capabilities. Monitoring is the collection of predefined signals to indicate state and trend, and it presupposes knowing in advance what to ask. Observability is the capacity to infer the internal state of a system from its measurable outputs, including for questions not anticipated in the instrumentation (Sridharan, 2018). Telemetry is the set of data that components emit. The distinction matters because an attack rarely triggers a foreseen alert: monitoring reports that a server is unavailable; observability makes it possible to ask, afterward, why it became unavailable and which other assets displayed the same pattern. The triad supporting this capability comprises metrics, suited to service

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

level alerts (Beyer et al., 2016); logs, the basis of forensic analysis (Kent & Souppaya, 2006); and traces, which relate cause and effect across components.

1.2 The detection stack and asset context

Intrusion detection and prevention systems inspect traffic or host activity for signatures and anomalies, and may operate passively, merely alerting, or inline, blocking (Scarfone & Mell, 2007). The distinction is decisive in a clinical network, because inline blocking introduces a point of failure in the data path of care equipment. The security information and event management platform correlates events from heterogeneous sources, and is the layer at which detection ceases to be local and becomes systemic; endpoint detection and response offers process visibility and containment on the host itself; network detection and response learns traffic behavior without depending on an installed agent, an attribute relevant to clinical devices that do not admit additional software; and orchestration chains these capabilities into predefined flows. The NIST recommendations, reorganized in 2025 to align with the Cybersecurity Framework 2.0, treat preparation, detection, analysis, containment, eradication and recovery as governed activities rather than purely technical stages (Nelson et al., 2025). It is in this dimension that the present work intervenes.

Correlating events without asset context produces alerts that are technically correct and operationally useless. Knowing that a host showed anomalous authentication is different from knowing that this host is the reporting server of the operating theater, on which three clinical services depend. This context resides in the asset inventory, which ISO/IEC 27001:2022 treats as an organizational control (International Organization for Standardization [ISO], 2022). In a hospital environment the inventory acquires an additional dimension: beyond technical criticality, each asset carries a clinical criticality expressing the care consequence of its unavailability. These are two scales, and the second is the one that conditions response.

1.3 Hospital particularities and time indicators

Three characteristics distinguish the hospital network from a corporate one. The first is the presence of connected medical devices and operational technology that do not admit a security agent and may behave unpredictably under active scanning — so their security depends on passive observation and compensating controls. The second is the asymmetry between maintenance windows and required availability: critical services operate continuously, and the window to apply a patch is narrow and frequently canceled by clinical demand, accumulating a debt that arises not from negligence but from care priority. The third is segmentation: the zones and conduits approach of the IEC 62443 series groups assets by security requirement and mediates communication through controlled channels (International Electrotechnical Commission [IEC], 2020), separating the clinical zone from the administrative one without interrupting legitimate flows. A data protection regime also

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

applies: hospital security logs contain patient identifiers, subjecting telemetry to the principles of purpose, necessity and minimization of the Brazilian General Data Protection Law (Lei n. 13.709, 2018).

Mean time to detect measures the interval between compromise and its identification; mean time to respond, between identification and effective containment. They are the most widespread indicators in security operations, and market literature tends to treat them as symmetrical targets for optimization. According to IBM (2026), organizations that deploy automation extensively contained incidents 65 days faster — a finding habitually read as a direct recommendation to automate. The results section examines why that reading requires qualification in a care setting.

2 - METHODOLOGY

2.1 Research design and construction stages

The research is qualitative in approach, applied in nature and documentary in procedure. We adopt Design Science Research as the strategy, appropriate when the purpose is not to describe a phenomenon but to produce an artifact that solves an identified organizational problem (Hevner et al., 2004). The product is a model, understood as a representation that organizes constructs and relations so as to guide decisions in a concrete situation. The path follows the activities of Peffers et al. (2007), of which problem identification, definition of solution objectives, design and development, and analytical demonstration were conducted, in five stages: survey and qualification of sources; characterization of clinical constraint; derivation of the asymmetry between the time indicators; structuring of the telemetry architecture and response stratification; and definition of indicators with the normative mapping.

2.2 Documentary sources and qualification criteria

We established a four-level qualification hierarchy. At the first level, technical standards and publications by government agencies. At the second, peer-reviewed academic literature, with priority for studies establishing an empirical relation between cyber incident and clinical outcome. At the third, telemetry reports and large-scale sector research, admitted when they describe sampling method and collection period. At the fourth, material from solution vendors, admitted for conceptual characterization and never as a quantitative basis. Numerical data without a reachable primary source were discarded.

2.3 Ethical and information security considerations

This work uses no patient data, clinical records, professional identification data or any personal information, and therefore does not fall within the situations requiring submission to a Research Ethics Committee, nor does it use operational data from any identified organization.

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

Omitting the organization of origin is not a formality of anonymity but an information security requirement: publicly describing the network topology, detection coverage or response limitations of a specific hospital is information of operational value to an adversary, and its disclosure would increase the risk to the very organization the work seeks to protect. The model is presented at the level of a generic architecture, without parameters, suppliers, versions or thresholds that would permit inference about a specific environment.

2.4 Limitations of the study

Three limitations must be made explicit. The Design Science Research cycle was not completed: the artifact was designed and demonstrated analytically but not subjected to controlled empirical evaluation, so its usefulness is argued rather than demonstrated. Anonymization prevents presenting the historical indicator series that would support a before and after comparison. And sector heterogeneity is relevant: hospitals differ in size, care profile and maturity, and the model assumes an environment with computerized critical care services and some installed security capability.

2.5 Declaration of artificial intelligence use

The authors declare the use of a generative artificial intelligence assistant in surveying and screening documentary sources, organizing the material analytically and drafting the manuscript. The definition of the theme and scope, the formulation of the central proposition, the decisions on research design, the source qualification criteria and the final validation are the exclusive responsibility of the authors. No personal, sensitive, clinical, confidential or organizationally identifying data were provided to the system. Outputs were reviewed by checking each reference against its primary source and verifying quantitative data against the originating reports. The prompts, parameters and methodologically relevant records are archived by the authors and available on request, subject to legally protected information. The system was responsible for screening sources, deriving the asymmetry between the time indicators, preparing the tables and producing the draft text; the authors were responsible, beyond review and validation, for defining the theme from the master's research in progress, excluding a thematic strand without support in that prior research, anonymizing the operational environment in full, and validating the central proposition. The authors assume full responsibility for the accuracy, integrity and final content of the work.

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

REVISTA CIENTÍFICA SOL21
STANDARD OPEN LITERATURE - ISSN 3086-089X

3 - RESULTS AND DISCUSSION

3.1 Clinical constraint as a property of the environment

The first contribution treats the clinical cost of containment as a structural property of the environment, not an exception negotiated case by case.

Table 1

Conventional containment actions, care risk and admissible alternatives

Conventional action	Risk on a clinically critical asset	Admissible alternative
Endpoint isolation	Interruption of a workstation in care use, with loss of access to a report, prescription or image during a procedure.	Selective isolation preserving routes to essential clinical services, or reinforced monitoring until the procedure ends.
Credential revocation	Blocking access to the health record during emergency admission or shift handover.	Staged revocation, beginning with sessions and privileged access, preserving a contingency credential under audit.
Segment quarantine	Silencing of patient monitoring telemetry and of communication by care equipment.	Restriction of conduits by allow-list, maintaining communication foreseen in the device profile.
Process blocking	Interruption of software driving equipment in clinical use.	Blocking restricted to assets not in care use, triggered conditionally on operational state.
Active scanning	Unpredictable behavior of an embedded medical device under stimulus not anticipated by the manufacturer.	Passive discovery through traffic analysis and correlation with a vulnerability database by model and version.
Preventive shutdown	Suspension of procedures and diversion of patients.	Controlled degradation with a previously exercised clinical contingency.

Note. Compiled by the authors based on International Organization for Standardization (2022), National Institute of Standards and Technology (2024) and Nelson et al. (2025).

Two observations follow. No alternative in the right-hand column is free: all require more information about the asset and its operational state. Isolating an endpoint requires only its address; isolating it selectively requires knowing which communications sustain ongoing care. And that information is not obtained at the moment of the incident: it must be available beforehand, in the inventory and in the telemetry. Clinical constraint thus converts into an observability requirement.

3.2 The asymmetry between detection and response times

Response time comprises the interval between identification of the incident and its containment. In a corporate environment it is bounded by organizational factors — staff availability, quality of procedures, degree of automation — all amenable to improvement through investment. In a care setting a limitation of a different nature is added: the clinical window. If containment requires an ongoing procedure to finish, response time cannot be reduced below the duration of that procedure, however mature the security operation may be. This floor is not inefficiency to be corrected; it is the correct expression of patient priority. Detection time has no analogous limitation: reducing it depends on telemetry coverage, correlation quality and asset context — variables under

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

organizational control that do not compete with care delivery. One additional passive sensor interrupts no surgery.

The practical consequence inverts the current intuition. The habitual reading of the finding that automation reduces containment time by 65 days (IBM, 2026) is that hospitals should automate response. The correct reading under clinical constraint is that they should automate only where clinical criticality permits, concentrating the remaining investment in detection. Every unit of time saved in detection converts fully into advantage, because it shifts containment to a moment when the set of admissible actions is still broad: detecting during reconnaissance permits blocking at the perimeter; detecting after encryption forces a choice between clinical harm and informational harm.

3.3 Proposed telemetry architecture

The third contribution structures the telemetry chain into seven layers organized by decision point: each answers a choice that, if poorly resolved, compromises the following ones.

Table 2

Layers of the telemetry architecture and their decision points

Layer	Function	Critical decision and clinical requirement
Generation	Emission of signals by hosts, applications, network and clinical devices.	Define what to instrument without requiring an agent on assets that do not admit one; favor native emission and network observation.
Collection	Capture through agents, syslog, flow exporters and traffic mirroring.	Passive collection as the default in the clinical zone; mirroring rather than inline insertion in the care data path.
Transport	Delivery of data to the consolidation point.	Encrypted channel independent of the care network, so telemetry survives degradation of the observed segment.
Normalization	Standardization of fields and formats.	Reliable time synchronization and consistent asset identifier; without these, correlation is illusory.
Enrichment	Addition of inventory, criticality and exposure context.	Incorporate clinical criticality as an attribute distinct from technical criticality.
Correlation	Systemic detection through rules and behavior.	Prioritize by potential clinical impact, not by technical severity alone.
Presentation	Operational dashboards and response triggering.	Expose the operational state of the asset to the analyst at the moment of the containment decision.

Note. Compiled by the authors based on Kent and Souppaya (2006), Scarfone and Mell (2007), Sridharan (2018) and International Organization for Standardization (2022).

The normalization layer deserves emphasis: correlating events across distinct sources depends on reliable temporal ordering. If clocks diverge, the platform may invert the sequence of cause and effect and fail to recognize a chain that did occur. Time synchronization, treated as an infrastructure detail, is a precondition of systemic detection.

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

3.4 Response stratification by governed autonomy

The fourth contribution resolves the impasse that leads healthcare organizations to acquire orchestration platforms and keep them inert, and which arises from treating automation as a binary decision.

Table 3

Levels of response autonomy and assignment criteria

Level	Behavior	Recommended application
N0 — Observation	Logging and alerting, with no action on the asset.	Clinical devices in care use and life support systems, in every case.
N1 — Reversible action	Actions that do not interrupt function and can be undone without residual effect, such as log elevation and evidence capture.	Any asset, including those of maximum clinical criticality.
N2 — Action subject to approval	Effective containment conditioned on human authorization with a defined deadline and automatic escalation.	Assets of high clinical criticality not in care use at the moment of decision.
N3 — Automatic action	Immediate containment without human intervention, with auditable logging and predefined reversal.	Administrative and infrastructure assets without direct care dependency.

Note. Compiled by the authors based on Nelson et al. (2025) and National Institute of Standards and Technology (2024).

Level assignment is not static: it depends on the operational state of the asset at the moment of the event. An operating theater workstation runs at N0 during a procedure and may run at N2 outside it. This requires care state to be a telemetry signal rather than information obtained by telephone during the incident — which reinforces the argument of subsection 3.1.

3.5 Detection use cases and proposed indicators

The fifth contribution reorders detection use cases by care impact rather than by frequency or ease of implementation.

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

REVISTA CIENTÍFICA SOL21
STANDARD OPEN LITERATURE - ISSN 3086-089X

Table 4

Detection use cases ordered by care impact

Use case	Correlated signals	Care consequence avoided
Ransomware precursors	Repetitive communication with an external destination, changes to recovery mechanisms and unusual elevation of write activity.	Simultaneous unavailability of health record, imaging and laboratory, with suspension of procedures.
Privileged credential compromise	Authentication outside the usual time and origin pattern, changes to privileged groups and successive second-factor failures.	Unrestricted access to clinical systems and capacity to disable controls.
Deviation in clinical device behavior	Communication with a destination or protocol outside the profile established for the device model.	Interference with care equipment and propagation from an asset without an agent.
Lateral movement within the clinical zone	Anomalous access between workstations in the same zone and internal port scanning.	Extension of the compromise to clinical decision support services.
Exfiltration of health data	Atypical outbound volume, unusual compression and an unfamiliar destination, correlated with access to a clinical repository.	Exposure of sensitive patient data, with regulatory and trust consequences.

Note. Compiled by the authors based on MITRE Corporation (2024), Nelson et al. (2025) and van Boven et al. (2024).

The ordering differs from the usual one in one respect: exfiltration, treated as the principal risk in most prioritization matrices because of its regulatory consequence, appears last — its consequence is patrimonial and reputational, whereas the first four affect the capacity to deliver care. This inversion expresses the treatment of care continuity as a security property in its own right.

Table 5

Proposed operational indicators

Indicator	Operational definition	Interpretation
Detection time segmented by criticality	Mean time to detect calculated separately for assets of high clinical criticality and for other assets.	Reveals whether telemetry coverage tracks clinical importance.
Telemetry coverage by zone	Proportion of active assets emitting telemetry effectively ingested, calculated by zone.	Measures the fraction of the surface actually observed; it is the limiting indicator for the others.
Admissible window latency	Interval between detection and the first moment at which containment becomes clinically acceptable.	Quantifies the response floor imposed by care delivery and separates clinical constraint from operational slowness.
Response distribution by autonomy level	Proportion of incidents handled at each level from N0 to N3.	Concentration at N0 suggests an orchestration platform acquired and not used.
Precision on clinical assets	Ratio of confirmed alerts to issued alerts, calculated only for assets of high clinical criticality.	Noise in this band is especially damaging, since it induces distrust where containment is most delicate.

Note. Compiled by the authors.

The third indicator is the most useful in dialogue with clinical leadership: it separates two frequently conflated causes of delay in containment — the security operation was slow, or the clinical window did not permit action. Without that distinction, every delay is attributed to the technical team, and the discussion about the care floor never takes place.

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

3.6 Normative mapping and discussion

Table 6

Mapping of model components against normative references

Component	ISO/IEC 27001:2022	NIST CSF 2.0	Continuity
Asset context and criticality	5.9 Inventory of information and other associated assets	Identify	ISO 22301 — business impact analysis
Telemetry architecture	8.15 Logging; 8.16 Monitoring activities	Detect	—
Segmentation by zones and conduits	8.20 Networks security; 8.22 Segregation of networks	Protect	IEC 62443 — zones and conduits
Stratified response	5.24 Information security incident management planning and preparation	Respond	ISO 22301 — response procedures
Clinical contingency and controlled degradation	5.29 Information security during disruption	Recover	ISO 22301 — continuity strategies
Governance of response autonomy	5.2 Information security roles and responsibilities	Govern	—

Note. Compiled by the authors based on International Organization for Standardization (2022), National Institute of Standards and Technology (2024) and International Electrotechnical Commission (2020). Control identifiers should be checked against the full text of the standards before use in an audit process.

The model engages critically with the detection literature, which concentrates effort on classification accuracy, assuming that the subsequent action is indifferent to context. The hospital environment shows that it is not: two alerts with the same statistical confidence demand distinct conduct if one falls on an administrative server and the other on the workstation displaying an image during a procedure. Accuracy is necessary and insufficient.

A second point concerns the standing of availability: the conventional triad assigns equal weight to the three properties, yet investment concentrates on the first two. Findings associating ransomware with increased in-hospital mortality (Neprash et al., 2026) suggest that in healthcare availability is the property whose violation produces irreversible harm: breached confidentiality generates a serious but reparable consequence; care interrupted in a time-dependent condition does not. A third point is the economics of defense: no layer of the model depends on advanced detection, since extending telemetry coverage, qualifying asset context and stratifying autonomy act on the set of available actions rather than on the capacity to distinguish the adversary.

We acknowledge, finally, an unresolved tension. The model requires the clinical operational state of each asset to be a telemetry signal available to the analyst at the moment of decision, and obtaining that signal implies integrating clinical and security systems, widening the interconnection between environments that segmentation recommends keeping separate. There is a conflict between the need for context and the principle of segregation. The practical resolution lies in transmitting only the state — whether the asset is in care use — without clinical content, but the secure design of that integration remains open.

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

REVISTA CIENTÍFICA SOL21
STANDARD OPEN LITERATURE - ISSN 3086-089X

4 - FINAL CONSIDERATIONS

This work proposed an observability and incident response model for healthcare organizations, built through Design Science Research, in which containment is conditioned by the clinical criticality of the affected asset.

We characterized clinical constraint as a structural property of the environment, mapping six containment actions against the corresponding clinical risk and identifying admissible alternatives, all dependent on greater prior information about the asset. We demonstrated that response time has a floor imposed by the clinical window, whereas detection time has no analogous limitation, which reorients investment priority. We proposed a seven-layer telemetry architecture, stratified response into four levels of autonomy, and defined five indicators, with the normative mapping.

The theoretical contribution consists in incorporating the clinical cost of containment as a constraint of the detection problem and in demonstrating the asymmetry between the established time indicators. The practical contribution lies in the model and the indicators, notably admissible window latency, which separates operational delay from clinical constraint. The limitations were made explicit in the methodology, notably the absence of controlled empirical evaluation.

As a future agenda, we suggest four directions: empirical evaluation of the model in organizations of different sizes, completing the Design Science Research cycle; investigation of the secure design of integration between clinical and security systems to transmit operational state without clinical content; study of admissible window latency as an indicator, including its distribution by specialty; and analysis of the regulatory implications of care unavailability in Brazil.

REFERENCES

Beyer, B., Jones, C., Petoff, J., & Murphy, N. R. (2016). *Site reliability engineering: How Google runs production systems*. O'Reilly Media.

Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly, 28*(1), 75–105. https://doi.org/10.2307/25148625

IBM. (2026). *Cost of a data breach report 2026*. IBM Security.

International Electrotechnical Commission. (2020). *IEC 62443-3-2: Security for industrial automation and control systems — Security risk assessment for system design*.

International Organization for Standardization. (2019). *ISO 22301:2019 — Security and resilience: Business continuity management systems: Requirements*.

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.

International Organization for Standardization. (2022). *ISO/IEC 27001:2022 — Information security, cybersecurity and privacy protection: Information security management systems: Requirements*.

Kent, K., & Souppaya, M. (2006). *Guide to computer security log management* (NIST SP 800-92). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-92

Lei n. 13.709, de 14 de agosto de 2018. (2018). *Lei Geral de Proteção de Dados Pessoais (LGPD)*. Presidência da República. https://www.planalto.gov.br/ccivil_03/_ato2015-2018/2018/lei/l13709.htm

MITRE Corporation. (2024). *ATT&CK for enterprise*. https://attack.mitre.org

National Institute of Standards and Technology. (2024). *The NIST Cybersecurity Framework (CSF) 2.0* (NIST CSWP 29). U.S. Department of Commerce. https://doi.org/10.6028/NIST.CSWP.29

Nelson, A., Rekhi, S., Souppaya, M., & Scarfone, K. (2025). *Incident response recommendations and considerations for cybersecurity risk management: A CSF 2.0 community profile* (NIST SP 800-61r3). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-61r3

Neprash, H., McGlave, C., & Nikpay, S. (2026). Hacked to pieces? The effects of ransomware attacks on hospitals and patients. *American Economic Journal: Economic Policy, 18*(1), 256–281. https://doi.org/10.1257/pol.20240594

Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems, 24*(3), 45–77. https://doi.org/10.2753/MIS0742-1222240302

Scarfone, K., & Mell, P. (2007). *Guide to intrusion detection and prevention systems (IDPS)* (NIST SP 800-94). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-94

Sridharan, C. (2018). *Distributed systems observability: A guide to building robust systems*. O'Reilly Media.

van Boven, L. S., Kusters, R. W. J., Tin, D., van Osch, F. H. M., De Cauwer, H., Ketelings, L., Rao, M., Dameff, C., & Barten, D. G. (2024). Hacking acute care: A qualitative study on the health care impacts of ransomware attacks against hospitals. *Annals of Emergency Medicine, 83*(1), 46–56. https://doi.org/10.1016/j.annemergmed.2023.04.025

This article is published in Open Access under the Creative Commons Attribution 4.0 International (CC-BY) license, which allows unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited.