International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

Risk-Oriented Analysis and Classification of Network Attacks Affecting Hospital Information Systems: Evidence from a Patient-Monitoring Testbed

Crimary C. Mahinay[1]; John Augustus P. Diesto[2]; Jerald B. Babor[3]; Ray Marlou T. Georpe[4]; Welquim B. Panogaling[5]; Serafin C. Palmares[6]; Kristine T. Soberano[7]

[1,2,3,4,5]MIT Students, State University of Northern Negros, Sagay City, Negros Occidental, Philippines
[6,7]MIT Professor, State University of Northern Negros, Sagay City, Negros Occidental, Philippines
[1] crmrymhny@gmail.com; [2] johndiesto102@gmail.com; [3] jeraldbabor60@gmail.com; [4] raymarlou06@gmail.com; [5] welpanogaling12@gmail.com; [6] spalmares@sunn.edu.ph; [7] ksoberano@sunn.edu.ph

DOI: https://doi.org/10.47760/ijcsmc.2026.v15i08.003

Abstract: Hospital Information Systems increasingly rely on connected patient-monitoring devices that transmit biometric data through hospital networks. This study analyzed the WUSTL-EHMS-2020 testbed and assessed the potential risk of its recorded attacks to HIS. A quantitative descriptive secondary analysis retained all 16,318 verified records. Record-level summaries described Normal, Spoofing, and Data Alteration traffic; the primary comparison aggregated ordered records into 169 contiguous binary-label blocks (85 Normal and 84 attack). Nine network-flow and eight biometric variables were summarized, and Cliff's delta with 95% percentile block-bootstrap confidence intervals was estimated using 10,000 repetitions and seed 20260727. The dataset contained 14,272 Normal records and 2,046 attack records, comprising 1,124 Spoofing and 922 Data Alteration records. Network-flow variables showed strong block-level separation (absolute delta 0.7891–0.9655; all intervals excluded zero), while all biometric effects were negligible (absolute delta at most 0.0594; all intervals included zero). Researcher-defined assessment classified Spoofing (score 12) and Data Alteration (score 15) as High risk. The study provides a reproducible bridge from verified testbed evidence to a separate, transparent HIS risk assessment; its conclusions remain limited to the controlled dataset and its undocumented attack-category field.

Keywords: Hospital Information System, healthcare cybersecurity, risk assessment, patient-monitoring systems, Internet of Medical Things

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

I. INTRODUCTION

Hospital Information Systems (HIS) now underpin routine healthcare operations. They support patient records, laboratory and imaging results, pharmacy transactions, billing, and a broad range of other clinical and administrative work. Their reach also extends to connected patient-monitoring systems, in which medical sensors transmit physiological data through network gateways to remote servers. Hady et al. [1] examined these two layers together, combining network-traffic and patient-biometric data in an intrusion-detection study of a healthcare monitoring testbed.

As hospitals grow more dependent on connected technology, cybersecurity can no longer be treated as a concern belonging to the information-technology department alone. A successful attack may expose private records, alter information that clinicians rely on, or lock staff out of essential systems at the moment they are most needed. NIST Special Publication 800-30 frames information-security risk in terms of threats, vulnerabilities, likelihood, and the resulting effects on organizational operations, assets, and individuals [2]. This study examines those risks in one such environment. The WUSTL-EHMS-2020 dataset records traffic from a controlled healthcare-monitoring testbed in which the malicious activity consists of man-in-the-middle attacks involving spoofing and data injection [3].

Incidents at operating healthcare organizations show how quickly such risks reach care delivery. An attack may begin inside a computer system or network, but its consequences rarely stay there. Neprash et al. examined 374 ransomware attacks involving healthcare organizations in the United States and found that nearly half coincided with disruptions such as system downtime, cancelled services, and ambulance diversion [4]. Dameff et al. observed comparable spillover after an attack on a single organization: nearby emergency departments absorbed more patients, waiting times lengthened, and pressure on staff and resources increased [5]. Abbou et al. reported that surgical activity, emergency visits, and hospital occupancy fell sharply once a cyberattack disabled a hospital's computer systems, with clinical services resuming gradually as electronic medical records, laboratory systems, and imaging services were restored [6]. Together, these cases indicate that the effects of an attack can spread well beyond the system first compromised, disrupting workflows, reducing service capacity, and delaying time-sensitive care.

Healthcare cybersecurity research has advanced along two lines that seldom meet. Dataset-based work concentrates on detecting malicious network activity and measuring how well classification methods perform [1]. Risk-assessment guidance instead asks how likely a threat is, how serious its effects may be, and which risks warrant attention first [2]. Hospitals need both. Detection is only the opening step; administrators and information-technology personnel must still determine which systems, data, and services are exposed, whether information could be disclosed, altered, or made unavailable, and how the incident might disrupt operations or place patients at risk. This study connects the two perspectives by linking recorded attack characteristics to a structured assessment of what those characteristics could mean for HIS.

The dataset serves to identify the recorded attack categories and describe the characteristics observable directly in the data. The risk assessment then considers which HIS assets and functions could plausibly be affected, together with the possible consequences for confidentiality, integrity, availability, hospital operations, and patient safety. Each recorded category is rated for likelihood and impact using scales the researchers defined for this study, adapted from the likelihood-and-impact approach in NIST Special Publication 800-30 [2], and the resulting classification gives a documented basis for practical security recommendations. Because the records come from a controlled testbed rather than an operational hospital, the proportion of attack traffic reflects the design of the experiment and is not treated as evidence of how often such attacks occur in practice.

Three elements combine here in one reproducible sequence: verification of the source file against its official documentation, descriptive estimation that accounts for dependence among consecutive records, and a risk

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

classification developed separately from the dataset evidence. Detection performance is not treated as a measure of risk. Keeping testbed-derived observations apart from researcher-assigned likelihood and impact ratings holds the limitations of the controlled collection environment in view while showing how network observations can support a transparent, risk-oriented interpretation.

A. Objectives of the Study

This study analyzes the network attacks recorded in the dataset and classifies their assessed risk to Hospital Information Systems using a risk-oriented approach. Specifically, it seeks to:

1. Identify the network attack categories recorded in the healthcare-monitoring dataset.
2. Determine the hospital assets that could plausibly be affected by the identified attacks.
3. Assess the potential effects of the attacks on confidentiality, integrity, and availability, and consider their plausible implications for hospital operations and patient safety.
4. Classify the identified attacks according to their likelihood, potential impact, and overall risk level.
5. Recommend security measures appropriate to the assigned risk level.

II. METHODOLOGY

A. Research Design

This study used a quantitative descriptive design applied to an existing network-traffic dataset. It is non-experimental: no variables were manipulated, no intervention was introduced, and no hospital or clinical process was altered. The study examined recorded network activity to identify attack patterns in a patient-monitoring environment, then assessed what risks those attacks could pose to Hospital Information Systems (HIS). It does not compare security tools, test a solution, or attempt to establish cause and effect.

Using a publicly available intrusion-detection dataset was both practical and ethically appropriate. Collecting traffic from an operating hospital is difficult, since patient-monitoring devices carry sensitive clinical and biometric information, and privacy requirements constrain how such data may be collected, shared, and reused. A dataset built in a controlled healthcare-monitoring environment offers a safer route: it carries labelled normal and attack traffic without exposing identifiable patient information, and because it is public, other researchers can examine the same records and reproduce the analysis.

Testbed data carries its own limitations. A controlled environment cannot reproduce the scale, complexity, device diversity, or daily operating conditions of a working hospital network. The findings should therefore be read within the scope of the dataset, and should not be assumed to hold across hospital environments without further validation.

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

B. Research Framework

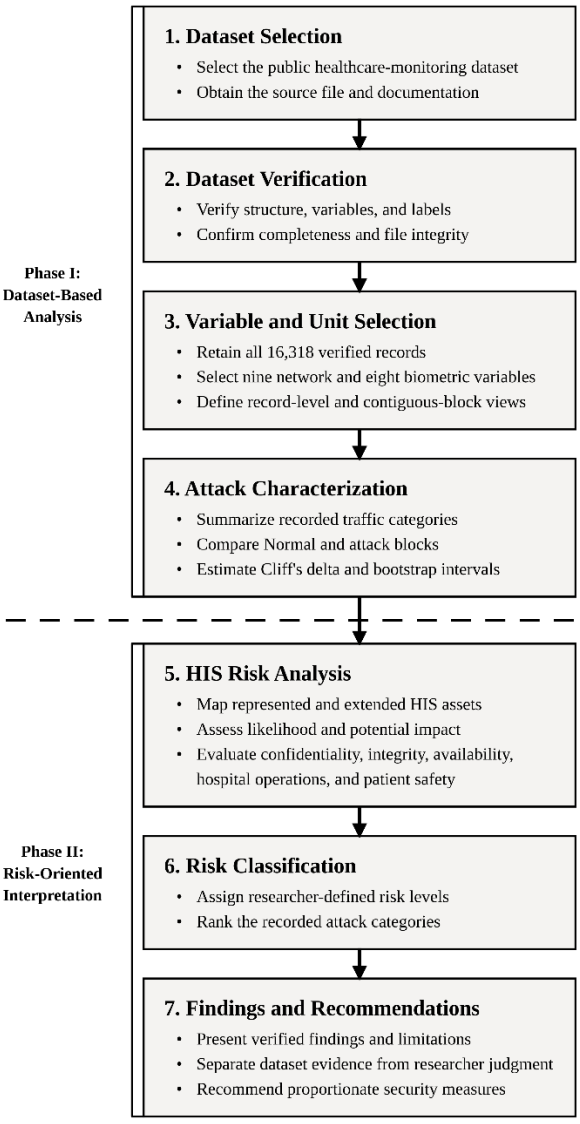

Fig. 1 Research framework of the study

Note. The dashed line separates Phase I (Dataset-Based Analysis) from Phase II (Risk-Oriented Interpretation). Activities in Phase I are based directly on the verified dataset, while activities in Phase II represent researcher-developed interpretation and were not measured directly from the dataset.

Fig. 1 sets out the research framework in seven stages across two phases. Phase I covers what was drawn from the dataset; Phase II covers what the research team inferred from those findings. The boundary between

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

them is the framework's organising principle and the reason the two are reported separately throughout this manuscript.

Phase I ran in four stages: dataset selection, dataset verification, variable and unit selection, and attack characterization. Selection obtained the public healthcare-monitoring dataset together with its source file and documentation. Verification confirmed structure, variables, labels, completeness, and file integrity. Variable and unit selection retained all 16,318 verified records, identified nine network-flow and eight biometric variables, and defined the record-level and contiguous-block views used in the analysis. Attack characterization summarized the recorded traffic categories, compared Normal and attack blocks, and estimated Cliff's delta with bootstrap intervals. Sections C through G describe each stage in full.

Phase II ran in three stages. HIS risk analysis mapped the assets represented in the testbed and the extended HIS components, then assessed likelihood and potential impact against confidentiality, integrity, availability, hospital operations, and patient safety. Risk classification assigned researcher-defined risk levels and ranked the recorded attack categories. The final stage presents the verified findings and their limitations, keeps dataset evidence separate from researcher judgment, and sets out proportionate security measures.

C. Data Source and Description

This study used the WUSTL-EHMS-2020 dataset as its sole source of data. Developed at Washington University in St. Louis, it was released publicly to support cybersecurity research in healthcare monitoring and Internet of Medical Things environments [1], [3].

The data came from the EHMS testbed, a controlled environment representing how patient information travels from medical sensors to a remote monitoring server [1]. It records network traffic during both normal system activity and controlled cyberattack scenarios. Because the experiments ran outside an operating hospital, malicious activity could be studied without involving real patients, interrupting medical services, or putting a working hospital network at risk.

The official documentation describes 44 features: 35 network-flow features, eight patient-biometric features, and one binary label marking each record normal or malicious [3]. The network-flow features cover connection duration, packet transfer, traffic volume, and communication patterns; the biometric features hold the patient measurements used in monitoring. The CSV file used here also carries an attack-category column labelling each record Normal, Spoofing, or Data Alteration, bringing the working file to 45 columns.

Pairing network and biometric data is what makes this dataset suitable here, since it places cyberattack activity inside a healthcare-monitoring setting. The recorded traffic is therefore used to identify and describe attack behaviour, while its possible effects on hospital operations, information security, and patient care are assessed separately in the risk analysis.

D. Dataset Verification

The downloaded dataset underwent a documented verification process before any analysis began. All checks ran in Python so that the reported audit values can be reproduced from the same file. Verification worked on a separate copy while the original download was preserved untouched; no cleaning, recoding, imputation, or other modification occurred at this stage.

The dataset came directly from the official page maintained by the institution that developed it [3]. Basic file properties were examined first — filename, byte size, character encoding, delimiter, and line-ending format — and a SHA-256 checksum was generated as a reproducible identifier for the CSV used here. The file measured 3,946,888 bytes, matching the size reported on the official page [3]. Its checksum was:

8359da96154fa60247df9d75e52d232077acb9886c09c36de2a0aaaee6cf2c25

Matching file size does not establish byte-level identity, since the dataset authors publish no official checksum for comparison. The checksum is included so that other researchers can confirm they hold the same CSV.

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

Structure was examined next. The file contained 16,318 records beneath a single header row across 45 columns. Column names were compared against both the official page [3] and the feature table in the original paper [1]. The official page describes 44 features: 35 network-flow variables, eight patient-biometric variables, and one binary label. The paper's table lists 34, omitting the nine addressing and record-identification fields excluded from the authors' models along with the binary label — a difference in presentation that reconciles with the downloaded file. Data types and distinct-value counts were recorded for every column, including those holding a constant value across all records.

The class-related fields were reviewed separately. The binary label held the values 0 and 1, while an additional text field classified each record as Normal, Spoofing, or Data Alteration. The two agreed across all records with no conflicting assignments; category counts appear in the Results section.

Completeness and consistency checks followed. None of the 45 columns contained missing values or empty text entries, though several biometric variables held zeros; these were counted but not acted on at this stage, and their handling is set out in Section II-F. No exact duplicate rows were found. For every record, total byte count equalled the sum of source and destination byte counts, and no record had a minimum packet size exceeding its maximum. One inconsistency did emerge: seven records carried a packet-loss count and percentage differing by more than rounding could explain. Several smaller irregularities were also noted — extra spaces in some text fields, three non-numeric values in the source-port column, and repeated values and sequence gaps in the packet-counter field. All were documented rather than corrected, so the dataset remained in its original condition. The packet-counter field warranted particular caution: being neither consistently unique nor sequentially complete, it was not treated as a record identifier.

One difference between the documented structure and the downloaded file remained. The official documentation describes 44 features; the file held 45 columns. The additional column was an undocumented attack-category field containing Normal, Spoofing, and Data Alteration. The official page describes the malicious traffic as man-in-the-middle attacks involving spoofing and data injection [3], so Spoofing corresponds directly to the official description, while the CSV uses Data Alteration rather than data injection for the second category. Neither the official page nor the original paper documents this field or explains how its values were assigned [1], [3]. The category names were therefore retained exactly as they appeared, with no assumption that Data Alteration and data injection are equivalent. Setting the additional column aside, the remaining 44 columns matched the documented structure.

E. Selection of Relevant Data for Analysis

Verification settled what the file contained; the next step was deciding what the study would use. No records were removed at this stage — selection operated only on variables.

All 16,318 verified records were retained, normal and attack traffic alike, since both were captured in the same healthcare-monitoring environment. The attack records supply the basis for describing the recorded cyberattack categories, and the normal records serve as the reference against which those categories are compared. Excluding either would remove the comparison itself. No record was dropped for its traffic class or attack category, and no attack type was weighted above another.

Selection therefore worked at the variable level, across three groups: class-related fields, network-flow variables, and patient-biometric variables.

The class-related fields distinguish normal traffic from malicious activity and identify the recorded categories. The binary label marks each record normal or malicious; the attack-category field assigns it to Normal, Spoofing, or Data Alteration. The field was retained because the study requires the recorded attack types to be examined separately, subject to the documentation gap described in Section II-D.

Nine network-flow variables were selected for statistical comparison: SrcLoad, DstLoad, SIntPkt, DIntPkt, SrcJitter, DstJitter, Dur, Load, and Rate. These cover source and destination loads, source and destination inter-

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

packet times, source and destination jitter, flow duration, total load, and transmission rate — together, the basis for comparing Normal, Spoofing, and Data Alteration traffic within the monitoring environment.

The patient-biometric variables hold the physiological measurements moving through the monitoring system. Because the network traffic carries patient-related information, retaining them allows the analysis to ask whether the recorded attacks are associated with differences in network behaviour, in biometric measurements, or in both.

Three groups of variables were excluded. Constant fields came out first: with no variation across the 16,318 records, they cannot separate normal traffic from attack traffic, or one attack category from another. Addressing and identification fields came out next — the source-port field holds values that are nearly unique across the dataset, and the packet-counter field proved neither consistently unique nor sequentially complete, so neither measures network behaviour in any meaningful sense.

The source-MAC field was excluded for a different and more important reason. The official dataset page reports that the source MAC address was used to construct the binary label, with records carrying the attacker laptop's MAC address assigned a label of 1 [3]. Including it would have let the analysis rediscover the labelling rule and present it as a characteristic of the traffic.

None of these variables were deleted. They remain in the working file so that the verified dataset stays intact and every selection decision stays reversible. The raw zero values and the seven packet-loss inconsistencies also remain unchanged; the raw biometric columns were preserved, with separate derived columns created under the zero-handling rule set out in the following subsection. The packet-loss variables do not enter the principal comparisons.

F. Data Preparation and Treatment of Data Issues

A separate analysis copy was created once the relevant variables were identified, preserving the verified source file and all raw numerical columns. No missing values were imputed, no outliers or records removed, and no source observation manually corrected. Derived fields were added only to preserve row order, identify contiguous label blocks, and implement the biometric zero-handling rule described below.

Class-related fields were reviewed so that labels and attack categories applied consistently throughout. Normal records served as the reference group; malicious records were examined under the recorded categories of Spoofing and Data Alteration, retained exactly as they appear in the dataset. Because the official documentation does not explain how the categories were assigned, no attempt was made to rename or independently verify them.

Before grouping, leading and trailing spaces were stripped from the Dir and Flgs text fields so that identical stored values grouped consistently. This changed no label, attack category, or numerical value. The selected network-flow and biometric variables were already numeric and needed no type conversion, and all numerical variables were analysed in their original units — not normalized, standardized, or transformed — because the study describes recorded values and compares them directly across traffic classes. Keeping the original units also keeps the findings traceable to the source data.

Raw zero values were retained in the source columns. For the primary biometric comparison, zeros in SpO2, Pulse_Rate, SYS, DIA, Heart_rate, and Resp_Rate were recoded as missing in separate derived columns only, on the basis that a zero is not an attainable reading from an attached sensor. No value was imputed and no complete record deleted. ST was treated differently: because it is signed and a zero may be genuine, zeros and negative values were retained in the primary ST variable, with a second ST variable treating zeros as missing for sensitivity analysis. Temp was retained unchanged and interpreted as a temperature sensor channel, since the available documentation does not establish its measurement basis.

The distribution of these zeros across traffic classes raised a further concern. For oxygen saturation, pulse rate, systolic pressure, diastolic pressure, and heart rate, zeros appeared only in normal records. Including them

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

could depress the statistics reported for the normal group and distort comparisons with Spoofing and Data Alteration.

Zero handling was therefore evaluated two ways. Record-level descriptive statistics were calculated with raw values and again after excluding zeros from the variable under examination. In the block-level analysis, derived zero-excluded columns were used for the six variables above, while ST was compared under both retained-zero and zero-excluded definitions. The raw columns and all 16,318 records remained available for traceability.

The seven records with packet-loss inconsistencies stayed unchanged in the verified working file. Because the packet-loss variables fall outside the final analytical set, these inconsistencies never entered the statistical comparisons, and retaining the original records avoided unsupported correction of either field.

Variables excluded during selection likewise remained in the verified file but were used in neither the statistical comparisons nor the later risk analysis. A separate analytical view containing only the selected variables prevented constant, identifying, or label-related fields from entering unintentionally. The final analysis thus drew on two linked views of the same 16,318 records: a record-level view for class composition and descriptive summaries, and a block-level view for dependence-aware comparisons.

G. Data Analysis and Risk Classification Procedure

The analysis ran in two stages. Phase I produced the dataset-derived findings through record-level summaries and block-level comparisons. Phase II asked what the recorded attacks could mean for Hospital Information Systems. The division keeps measured evidence apart from the likelihood, impact, and risk judgments the researchers assigned.

Phase I began by counting Normal and attack records and calculating their percentages, with the binary Label field as the primary outcome. The Attack Category field served as a secondary, exploratory outcome for describing Spoofing and Data Alteration, for the documentation reason given in Section II-D; findings resting on it are testbed-specific.

At the record level, the nine network-flow and eight biometric variables were summarized for Normal, Spoofing, and Data Alteration using count, mean, median, standard deviation, interquartile range, minimum, and maximum. These summaries describe recorded values only — the 16,318 rows are not independent inferential samples. Class membership occurs in contiguous runs, and consecutive biometric patterns repeat, so a second analysis aggregated each variable to its median within every contiguous run of the binary Label. That yielded 169 comparison blocks, 85 Normal and 84 attack. Block boundaries are an operational rule for handling dependence, derived from file order; they are not independently sampled hospital incidents.

Cliff's delta was calculated for each principal variable in the block-level normal-versus-attack comparison. Ninety-five percent percentile confidence intervals, for both the median difference and Cliff's delta, came from resampling whole blocks with replacement independently within each binary-label group, preserving group sizes. The procedure used 10,000 bootstrap repetitions and random seed 20260727. Where a block had no computable median for a derived variable, pairwise omission applied — this arose once, for Resp_Rate in a single attack block. No null-hypothesis significance test was performed.

The analysis stayed descriptive and estimation-focused throughout: no significance testing, no machine-learning classification, no prediction, no causal inference. Cliff's delta and the bootstrap intervals describe the direction, consistency, and uncertainty of group differences, and effect-size rank ordering was read alongside the observed median differences rather than on its own. No block was described as an independently observed patient, device, hospital incident, or attack occurrence.

The frozen analysis ran in Python 3, using pandas 3.0.2 for CSV ingestion, data handling, grouping, tabulation, and export, and NumPy 2.4.4 for numerical calculation and bootstrap resampling. A final reproducibility run used Python 3.12.13, pandas 2.2.3, and NumPy 2.3.5. Processed dataset, block medians, primary tables, category-decomposition tables, and console log were byte-identical across both runs.

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

The proportion of attack records describes the composition of the recorded experiment and nothing more. It was not used to estimate how often such attacks occur in operational hospitals, and it played no part in the likelihood ratings.

Phase II turned from describing the data to weighing its implications. Each recorded attack category was examined within the monitoring environment the testbed represents. The researchers identified the information assets and system functions that could plausibly be affected, then considered the possible consequences for confidentiality, integrity, availability, hospital operations, and patient safety — potential consequences, not effects observed in an actual hospital.

The assessment drew on the confidentiality, integrity, and availability objectives and the likelihood-and-impact approach in NIST Special Publication 800-30 [2]. The five-point rating scales, the scoring method, and the risk-level boundaries, however, were developed by the researchers for this study. They are adapted from that approach, not specified by it.

Ratings ran on a five-point likelihood scale from Rare to Almost Certain. Three considerations drove each judgment: how technically feasible the attack is, what level of access it demands, and what network conditions it needs. Published exposure data for comparable healthcare communication systems were unavailable, so no frequency evidence entered the assessment. Neither did the number or proportion of attack records in the dataset — those figures reflect the design and duration of a controlled experiment.

Potential impact used a separate five-point scale, Insignificant to Severe, weighing possible exposure of confidential information, alteration of transmitted data, interruption of availability, disruption of hospital operations, and consequences for patient safety. Each category's impact rating reflects the most serious credible consequence identified across those areas.

Multiplying likelihood by impact gives an overall risk score between 1 and 25. Table I presents the researcher-defined risk-classification matrix used to interpret those scores.

TABLE I
RESEARCHER-DEFINED RISK CLASSIFICATION MATRIX

Likelihood ↓ / Impact →	**1 Insignificant**	**2 Minor**	**3 Moderate**	**4 Major**	**5 Severe**
5 — Almost Certain	5 Moderate	10 High	15 High	20 Critical	25 Critical
4 — Likely	4 Low	8 Moderate	12 High	16 High	20 Critical
3 — Possible	3 Low	6 Moderate	9 Moderate	12 High	15 High
2 — Unlikely	2 Low	4 Low	6 Moderate	8 Moderate	10 High
1 — Rare	1 Low	2 Low	3 Low	4 Low	5 Moderate

Risk score = likelihood × impact. Bands: 1–4 Low, 5–9 Moderate, 10–16 High, 17–25 Critical. The scales, scoring rule, and band boundaries were defined by the researchers for this study. They are adapted from the likelihood-and-impact approach described in NIST Special Publication 800-30 [2] but are not specified by that publication.

The research team assigned all likelihood, impact, and risk ratings through discussion and agreement, reviewing each category against the dataset authors' descriptions [1], [3], the descriptive findings, the related literature, and the healthcare setting defined in the study. Where proposed ratings differed, the team returned to the evidence and worked through the reasoning until it reached a single classification.

These ratings are researcher judgments, not values measured from the dataset. No external panel of hospital cybersecurity practitioners, clinical personnel, or risk-management specialists reviewed them, and that constraint bears on how the resulting classifications should be read.

Once finalized, the categories were ranked by overall risk level. Security measures followed from how each attack works, which information assets and system functions it could affect, and what consequences could reasonably follow, with priority given to High and Critical ratings. Findings and recommendations remain

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

limited to Spoofing and Data Alteration within the controlled patient-monitoring environment the dataset represents.

H. Ethical Considerations

This study did not undergo institutional ethics review. It involved no recruitment of or interaction with human participants and relied solely on a publicly released dataset in which no individual is identifiable. No ethics approval or formal exemption was sought or obtained for the secondary analysis.

Because the study relied on secondary data, the researchers could not independently verify the arrangements under which the original biometric measurements were collected. Neither the dataset paper nor the official documentation reports the number of human subjects involved or states whether ethics approval was obtained for the original collection [1], [3]. This gap is acknowledged as a limitation of using the publicly released dataset.

The dataset authors and the official distribution source are acknowledged through citation [1], [3]. The official distribution page states no explicit licensing or redistribution conditions, so the dataset was used only for non-commercial academic research with attribution to its original authors. The researchers claim no ownership of the data and no involvement in its original collection, and neither the dataset nor any portion of it is reproduced or redistributed with this manuscript.

III. Results and Discussion

A. Dataset and Attack Distribution

The dataset contained 16,318 records. Of this total, 14,272 were classified as Normal, representing 87.46% of all records. The remaining 2,046 records, or 12.54%, were classified as attacks.

The binary label and attack-category fields were checked against each other to confirm that the classifications were consistent. They matched across the entire dataset. Every record labelled as an attack was assigned to one of the two recorded attack categories, while every normal record appeared under the Normal category. No conflicting class assignments were found.

The attack records were divided into two categories: Spoofing and Data Alteration. Spoofing accounted for 1,124 records, while Data Alteration accounted for 922. These represented 6.89% and 5.65% of the full dataset, respectively. When only the attack records were considered, Spoofing made up 54.94% of the total, compared with 45.06% for Data Alteration. The attack-category field contained only three values — Normal, Spoofing, and Data Alteration — and no other category appeared in the file.

The wording used in the dataset documentation required careful interpretation. The official source describes the malicious traffic as man-in-the-middle attacks involving spoofing and data injection [3]. However, the analyzed file did not include separate record-level labels for man-in-the-middle or data injection. Instead, the second attack category was listed as Data Alteration. Since the documentation does not explain how this field was assigned or confirm whether Data Alteration directly corresponds to data injection, the category names were reported exactly as they appeared in the file. Table II summarizes the distribution of records by traffic class and attack category.

TABLE II
DISTRIBUTION OF RECORDS BY TRAFFIC CLASS AND ATTACK CATEGORY

Classification	Category	Frequency	Percentage of full dataset (n = 16,318)	Percentage of attack records (n = 2,046)
Traffic class	Normal	14,272	87.46	—
Traffic class	Attack	2,046	12.54	100.00
Attack category	Spoofing	1,124	6.89	54.94

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

Classification	Category	Frequency	Percentage of full dataset (n = 16,318)	Percentage of attack records (n = 2,046)
Attack category	Data Alteration	922	5.65	45.06
Total dataset	**All records**	**16,318**	**100.00**	—

Note: Spoofing and Data Alteration are subdivisions of the 2,046 attack records and therefore do not represent additional records beyond those included in the Attack row. Percentages were rounded to two decimal places.

Overall, Normal traffic made up most of the dataset, while attacks accounted for slightly more than one-tenth of all records. Spoofing was the more frequently recorded attack category, exceeding Data Alteration by 202 records. The total counts for Normal and Attack traffic matched those published on the official dataset page [3]. The separate counts for Spoofing and Data Alteration, however, were established directly from the analyzed data file.

B. Network-Flow Characteristics by Recorded Traffic Category

Nine network-flow variables were first summarized at the record level across the Normal, Spoofing, and Data Alteration groups. These three-category summaries are exploratory because the attack-category field is undocumented, and the row counts are descriptive rather than independent sample sizes. Median and interquartile range are shown with mean and standard deviation because several variables were strongly skewed and contained extreme values. Table III provides this record-level descriptive context.

TABLE III
RECORD-LEVEL DESCRIPTIVE NETWORK-FLOW CHARACTERISTICS BY RECORDED TRAFFIC CATEGORY

Variable	Category	Median (IQR)	Mean (SD)
SrcLoad	Normal	240,874 (54,464)	223,852 (67,054)
	Spoofing	224,342 (42,668)	218,036 (43,637)
	Data Alteration	20,059 (4,782)	18,360 (5,123)
DstLoad	Normal	80,291 (18,163)	74,617 (22,346)
	Spoofing	74,781 (14,223)	72,679 (14,546)
	Data Alteration	6,686 (1,581)	13,389 (157,622)
SIntPkt	Normal	4.1183 (0.9763)	6.8538 (34.3738)
	Spoofing	4.4218 (0.8578)	5.2059 (4.9486)
	Data Alteration	49.4142 (10.9080)	81.3020 (397.4415)
DIntPkt	Normal	2.4945 (0.7795)	3.9596 (24.3915)
	Spoofing	2.8580 (0.7711)	3.5684 (5.0057)
	Data Alteration	48.4648 (9.3619)	85.0681 (182.6813)
SrcJitter	Normal	3.0039 (0.8346)	35.3820 (1,354.0095)
	Spoofing	3.2202 (0.7135)	3.8520 (5.0736)
	Data Alteration	5.5310 (8.3034)	14.4204 (53.7732)
DstJitter	Normal	2.4405 (0.7746)	3.9470 (23.7550)
	Spoofing	2.7975 (0.7654)	3.5167 (5.0047)

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

Variable	Category	Median (IQR)	Mean (SD)
	Data Alteration	48.3950 (9.3693)	85.8814 (169.0669)
Dur	Normal	0.0124 (0.0029)	0.0206 (0.1134)
	Spoofing	0.0133 (0.0026)	0.0156 (0.0148)
	Data Alteration	0.1484 (0.0399)	0.3558 (0.8947)
Load	Normal	321,166 (72,653)	298,469 (89,399)
	Spoofing	299,122 (56,891)	290,714 (58,183)
	Data Alteration	26,746 (6,321)	31,749 (156,912)
Rate	Normal	485.633 (109.857)	451.311 (135.189)
	Spoofing	452.301 (86.024)	439.588 (87.978)
	Data Alteration	40.442 (9.557)	65.785 (624.916)

Normal: 14,272 records; Spoofing: 1,124 records; Data Alteration: 922 records. These are record counts, not independent comparison units. IQR denotes interquartile range, and SD denotes standard deviation. Values were rounded for presentation and remained in the units recorded in the dataset. Category-specific results are exploratory.

In the exploratory record-level summaries, Spoofing remained close to Normal traffic in measurement size. Median source load, destination load, total load, and rate were about 6.9% lower; median source inter-packet time, source jitter, and duration were about 7% higher; and destination inter-packet time and destination jitter were about 15% higher. These differences describe the testbed records and do not establish operational detectability.

Data Alteration showed a clearer exploratory departure from Normal traffic. Its median source load, destination load, total load, and rate were about 8% of the corresponding Normal values. Median source inter-packet time and duration were about twelve times higher, while destination inter-packet time and destination jitter were about nineteen to twenty times higher. Median source jitter was about 1.8 times the Normal value.

In the pooled block-level primary comparison, absolute Cliff's delta across the nine network variables ranged from 0.7891 for SrcJitter to 0.9655 for DstJitter, and every 95% confidence interval excluded zero. The pooled attack-block medians for SrcLoad, DstLoad, Load, and Rate were 12.92% lower than the Normal-block medians. Timing, jitter, and duration measures were 11.31% to 30.22% higher. The large rank-based deltas indicate consistent ordering across blocks and should be interpreted together with these measurement differences rather than as evidence of equally large practical gaps.

The category-specific block decomposition supported the same directional pattern: 46 Spoofing blocks differed from 85 Normal blocks by 7.35% to 14.95% across the nine network measures, whereas 38 Data Alteration blocks showed substantially larger differences. However, the modal Flgs value also corresponded exactly with the Data Alteration blocks, so the separation may reflect a protocol or capture condition specific to the testbed. It should not be presented as evidence that Data Alteration would be easily detected in an operational hospital.

Recent classifier-centered studies using WUSTL-EHMS-2020 provide a useful methodological comparison. Naeem et al. [7] combined feature selection, data augmentation, and deep ensemble learning and reported optimal classification scores of 100% on this dataset. Wu et al. [8] used SMOTE, kernel principal component analysis, and stacked deep-learning models; their feature analysis identified DIntPkt, DstJitter, and SrcJitter as influential. The present study likewise found strong block-level separation for DIntPkt and DstJitter, but its attack blocks had a higher median SrcJitter, whereas Wu et al. associated lower SrcJitter with a greater attack probability. The difference may reflect their transformations and predictive modeling versus the present study's untransformed, block-aggregated descriptive analysis.

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

The reported classifier scores in [7] and [8] cannot be used to validate the risk scores in this study because classification performance and organizational risk measure different constructs. This study did not rebalance the classes, transform the analytical variables, train a model, or estimate detection accuracy. Instead, it quantified observed group differences and then kept the researcher-developed likelihood and impact assessment separate.

This distinction also responds to broader concerns in recent IoMT literature. Naghib et al. [9] found that WUSTL-EHMS-2020 is frequently used for IoMT intrusion-detection research, while identifying unresolved challenges involving AI models, dataset quality, evaluation, and generalizability. The present findings reinforce the need for caution: large testbed separation does not establish real-world detection performance or operational risk, and validation across additional datasets and healthcare environments remains necessary.

C. Biometric Characteristics and Zero-Value Sensitivity

The eight biometric variables were examined at the record level for Normal, Spoofing, and Data Alteration, and then in the primary block-level normal-versus-attack comparison. Table IV reports the raw record-level distributions. Table V reports the zero values by category. The derived zero-handling rule was applied without deleting complete records, and its effect was checked against the raw summaries and the ST sensitivity variant.

TABLE IV
RECORD-LEVEL DESCRIPTIVE PATIENT-BIOMETRIC CHARACTERISTICS BY RECORDED TRAFFIC CATEGORY

Variable	Normal — Median (IQR)	Spoofing — Median (IQR)	Data Alteration — Median (IQR)
Temp	27.0 (1.00)	27.0 (0.90)	26.8 (1.20)
SpO2	98.0 (0.00)	98.0 (0.00)	98.0 (1.00)
Pulse_Rate	73.0 (6.00)	75.0 (8.00)	77.0 (11.00)
SYS	144.0 (6.00)	144.0 (10.00)	142.0 (10.00)
DIA	83.0 (8.00)	83.0 (6.00)	82.0 (9.00)
Heart_rate	73.0 (6.00)	73.0 (8.00)	73.0 (8.00)
Resp_Rate	19.0 (6.00)	19.0 (5.00)	19.0 (7.75)
ST	0.3 (0.10)	0.3 (0.12)	0.3 (0.10)

Normal: 14,272 records; Spoofing: 1,124 records; Data Alteration: 922 records. These are record counts, not independent comparison units. IQR denotes interquartile range. Table IV reports raw values in the units recorded in the dataset. Temp ranged from 23.6 to 29.2, but the available documentation did not establish its measurement basis.

As Table IV shows, the central biometric values were generally similar across the three categories. Median oxygen saturation, heart rate, respiratory rate, and ST were identical for Normal, Spoofing, and Data Alteration records. Temperature varied by no more than 0.2 units. Median diastolic pressure differed by only one unit, while systolic pressure varied by two. The largest difference appeared in pulse rate, with the Data Alteration median four units higher than the Normal median. The interquartile ranges were also broadly comparable, indicating that the middle half of the values had a similar spread across the three groups.

TABLE V
ZERO VALUES IN PATIENT-BIOMETRIC VARIABLES BY RECORDED TRAFFIC CATEGORY

Variable	Normal, n (%)	Spoofing, n (%)	Data Alteration, n (%)	Total, n (%)
Temp	0 (0.000)	0 (0.000)	0 (0.000)	0 (0.000)
SpO2	2 (0.014)	0 (0.000)	0 (0.000)	2 (0.012)

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

Variable	Normal, n (%)	Spoofing, n (%)	Data Alteration, n (%)	Total, n (%)
Pulse_Rate	2 (0.014)	0 (0.000)	0 (0.000)	2 (0.012)
SYS	36 (0.252)	0 (0.000)	0 (0.000)	36 (0.221)
DIA	36 (0.252)	0 (0.000)	0 (0.000)	36 (0.221)
Heart_rate	3 (0.021)	0 (0.000)	0 (0.000)	3 (0.018)
Resp_Rate	779 (5.458)	65 (5.783)	25 (2.711)	869 (5.325)
ST	150 (1.051)	7 (0.623)	13 (1.410)	170 (1.042)

Percentages for Normal, Spoofing, and Data Alteration were calculated within their respective categories. Total percentages were calculated using all 16,318 records.

Across the eight biometric variables, the dataset contained 1,118 zero entries. Most were found in respiratory rate, which accounted for 869, followed by ST with 170. Zero values for oxygen saturation, pulse rate, systolic pressure, diastolic pressure, and heart rate appeared only among Normal records. Even within that group, however, each variable affected fewer than 0.3% of the observations. Overall, 1,044 records, or 6.40% of the dataset, contained at least one zero in a biometric field. Temperature was the only variable with no recorded zeros.

At the record level, excluding zeros from the variable being summarized did not change any of the three category medians, and no interquartile range shifted by more than 1.75 units. This statement applies only to the raw-versus-zero-excluded category summaries; it does not describe the median of block medians used in the primary comparison.

At the block level, all eight cleaned biometric Cliff's deltas were negligible (absolute delta no greater than 0.0594), and every 95% bootstrap confidence interval included zero. One attack block had no valid Resp_Rate value after zero recoding, so that variable used 83 rather than 84 attack blocks. For ST, treating zeros as missing changed the Normal median of block medians from 0.25 to 0.26 and changed delta from 0.0594 to 0.0678; both estimates remained negligible and both intervals included zero. The conclusion of no meaningful block-level biometric separation was therefore stable.

The dependence check found only 8,877 distinct biometric value combinations across 16,318 records; the most frequent combination appeared 622 times, and class membership occurred in 169 contiguous runs. Consecutive rows were therefore not treated as independent comparison units. The contiguous binary-label block was used as a dependence-aware operational unit for aggregation and resampling, while its boundaries were acknowledged as products of the recorded label sequence rather than independently sampled hospital events.

D. HIS Asset Mapping and Security-Objective Impact Assessment

The findings in Subsections A through C were drawn from traffic recorded in a controlled healthcare-monitoring testbed. To place those findings within the wider context of Hospital Information Systems, the recorded attack categories were mapped to the components represented in the testbed and to selected HIS components that could receive, display, or store the transmitted information. This was an analytical mapping only. It does not suggest that the same effects were observed in an actual hospital. The assessment was guided by the concepts of threats, vulnerabilities, and potential adverse impacts described in NIST Special Publication 800-30 [2]. Table VI identifies the components included in the analytical asset mapping.

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

TABLE VI
HOSPITAL INFORMATION SYSTEM COMPONENTS INCLUDED IN THE ASSET MAPPING

HIS component	Represented in the dataset	Basis for inclusion
Medical sensors	Yes	Origin of the biometric measurements transmitted through the testbed [1]
Monitoring gateway	Yes	Intermediate component in the recorded communication path [1]
Monitoring server	Yes	Destination of the network traffic recorded in the dataset [1]
Sensor-to-server communication path	Yes	Represented by the network-flow variables analyzed in Table III
Clinical monitoring workstation or dashboard	No	Analytical extension representing a possible interface through which healthcare personnel view monitoring information
Electronic medical record system	No	Analytical extension representing a possible downstream recipient of monitoring information in an integrated HIS
Patient-data repository	No	Analytical extension representing a possible storage location for transmitted biometric measurements

Only the first four components were directly represented in the testbed. The remaining components were included solely to place the monitored communication path within a broader HIS environment. No effect on these downstream components was directly observed or measured.

The medical sensors, monitoring gateway, sensor-to-server communication path, and monitoring server were directly represented in the testbed. Clinical monitoring workstations, electronic medical record systems, and patient-data repositories were not. They were included only as possible downstream points where monitoring information could be viewed, stored, or integrated into a broader HIS environment. Table VII summarizes the researcher-assessed effects on confidentiality, integrity, and availability.

TABLE VII
RESEARCHER-ASSESSED IMPACT ON SECURITY OBJECTIVES BY RECORDED ATTACK CATEGORY

Security objective	Spoofing	Data Alteration
Confidentiality	Potential unauthorized observation of biometric information transmitted through the intercepted monitoring path	Potential unauthorized observation of biometric information during interception and modification
Integrity	Potential misrepresentation of the source or authenticity of monitoring communications	Potential modification of information transmitted between the sensor and monitoring server
Availability	No direct availability effect was established from the recorded data	No direct availability effect was established from the recorded data

The assessments were developed by the researchers from the attack descriptions provided by the dataset authors [1], [3]. They represent potential security impacts and were not directly measured from the network-flow or biometric values.

Two exploratory category findings were relevant to the mapping. First, Spoofing differed only modestly from Normal traffic in median measurement size across the nine selected network variables, although its block rankings were consistently shifted. This does not establish detectability because the study did not develop a detection model or evaluate classification performance.

Second, Data Alteration showed a much more distinct testbed pattern, with lower loads and rates and higher inter-packet times, destination jitter, and duration. Because the category field is undocumented and Data

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

Alteration blocks aligned exactly with the modal Flgs value M, this separation may reflect a testbed-specific protocol or capture condition. It should not be generalized to an operational hospital network.

The biometric findings required more careful interpretation. Although the central values were generally similar across the three recorded categories, this does not show whether individual biometric measurements were changed during an attack. The dataset did not identify which fields, if any, were altered in each record, nor did it provide paired original and modified values. The biometric findings were therefore used only to describe the distributions recorded within each category. They were not used to determine the clinical meaning of any individual measurement.

Overall, the mapping identified confidentiality and integrity as the main security objectives that could potentially be affected by the recorded attack scenarios. The dataset did not provide direct evidence of an effect on availability. It also did not show how these possible security impacts might affect hospital operations or patient safety. Those consequences are examined separately in the succeeding risk analysis and classification.

Published studies of healthcare organizations have associated cyberattacks with system downtime, cancelled services, ambulance diversion, and reduced clinical capacity [4], [5], [6]. Those studies examined ransomware incidents, which principally affect availability, and are cited here only for the general finding that attacks on healthcare organizations have been associated with disruption to care delivery. They do not establish operational consequences specific to Spoofing or Data Alteration in patient-monitoring systems. The sources reviewed for this study did not provide direct empirical evidence on the operational effects of integrity attacks against monitoring communications.

E. Risk-Oriented Classification of the Recorded Attack Categories

Spoofing and Data Alteration were assessed using the likelihood and impact scales defined in the methodology. For each category, the likelihood rating was multiplied by the impact rating to determine the overall risk score. The resulting score was then matched to the appropriate risk level. The assessment considered the conditions that could make each attack possible, the HIS components that might be affected, the security objectives involved, and the potential consequences for hospital operations and patient safety. Table VIII presents the resulting likelihood, impact, score, and risk level for each attack category.

TABLE VIII
RISK-ORIENTED CLASSIFICATION OF THE RECORDED ATTACK CATEGORIES

Attack category	Likelihood	Impact	Risk score	Risk level
Spoofing	3 – Possible	4 – Major	12	High
Data Alteration	3 – Possible	5 – Severe	15	High

Risk score = likelihood × impact. Scores from 10 to 16 were classified as High according to the risk-classification matrix established in the methodology.

Spoofing was assigned a likelihood rating of Possible. This type of attack could occur if an attacker gained access to the monitoring communication path or successfully impersonated a trusted device or network endpoint. Weak authentication, poor network segmentation, limited device verification, and unsecured communication channels could all create opportunities for the attack.

Its impact was rated Major because Spoofing could affect both the confidentiality and integrity of patient-monitoring communications. An attacker positioned along the communication path could view transmitted biometric information or make the data appear to come from a trusted source. This could cast doubt on the authenticity of the information received by the monitoring server and reduce its reliability when presented through connected clinical systems. With a likelihood rating of 3 and an impact rating of 4, Spoofing received a risk score of 12 and was classified as High risk.

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

Data Alteration was also rated Possible. To carry out the attack, an attacker would need to intercept the communication and modify the transmitted information before it reached the monitoring server. This would require access to the monitoring path, together with weaknesses that allowed altered data to pass without being rejected or detected.

The impact of Data Alteration was rated Severe because the attack directly threatens the integrity of patient-monitoring information. Altered biometric measurements could be received, displayed, or stored as though they were genuine. Within a Hospital Information System, inaccurate monitoring data could influence patient observation, clinical escalation, prioritization, or decision-making. A likelihood rating of 3 and an impact rating of 5 produced a risk score of 15, placing Data Alteration in the High risk category.

The exploratory network-flow findings distinguished the recorded categories within this testbed, but they do not demonstrate operational detection performance. Spoofing showed modest median differences from Normal traffic, while Data Alteration showed larger differences and a testbed-specific association with Flgs. These observations were used as descriptive context only and did not determine the researcher-assigned likelihood or impact ratings.

The impact ratings assigned above therefore rest on researcher judgment rather than on measured operational or clinical consequences arising from integrity attacks on patient-monitoring communications. The assessment was also not reviewed by clinical or hospital cybersecurity practitioners.

Both attack categories were classified as High risk because each could compromise important security objectives within the patient-monitoring path. Data Alteration received the higher score because modifying biometric information presents a more direct threat to the reliability of data used for patient monitoring and clinical decision-making. It was therefore placed ahead of Spoofing in the final order of risk priority.

IV. CONCLUSION

This study examined the network attacks recorded in a publicly available healthcare-monitoring dataset and used a risk-oriented approach to assess the risks they could pose to Hospital Information Systems.

The dataset contained 16,318 records. Of these, 14,272, or 87.46%, were recorded as Normal traffic, while 2,046, or 12.54%, were identified as attack traffic. The attacks were grouped into two categories: Spoofing, which accounted for 54.94% of the attack records, and Data Alteration, which accounted for the remaining 45.06%. The binary label and attack-category fields were consistent across every record.

The analysis used the 16,318 records for composition and descriptive summaries, but it did not treat them as independent comparison units. The primary normal-versus-attack comparison used 169 contiguous binary-label blocks: 85 Normal and 84 attack. Across the nine network variables, absolute Cliff's delta ranged from 0.7891 to 0.9655 and every bootstrap interval excluded zero, indicating consistent block ordering. The exploratory category decomposition showed modest measurement differences for Spoofing and substantially larger, potentially testbed-specific differences for Data Alteration.

The biometric findings were less pronounced. All eight block-level biometric effect sizes were negligible and every bootstrap confidence interval included zero. Raw category medians were unchanged when zeros were excluded from the variable being summarized. The ST block-level sensitivity estimate changed slightly under zero exclusion, but its magnitude remained negligible and its interval continued to include zero.

Four parts of the monitoring environment were directly represented in the dataset: the medical sensors, monitoring gateway, sensor-to-server communication path, and monitoring server. Clinical workstations, electronic medical record systems, and patient-data repositories were not part of the recorded environment. They were considered only as analytical extensions of the patient-monitoring system.

Confidentiality and integrity were identified as the security objectives that could be affected by both attack categories. The recorded data did not establish an effect on availability. Plausible implications for hospital operations and patient safety were considered during the researcher-developed risk assessment, but they were not measured or established directly from the dataset.

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

Both attack categories were classified as High risk. Spoofing received a risk score of 12, while Data Alteration received a score of 15. Although both fell within the same risk level, Data Alteration ranked ahead of Spoofing in the resulting order of priority.

The study provides a reproducible path from verified record-level data to dependence-aware descriptive estimates and then to a separate researcher-developed risk assessment. The distinction is essential: the record and block results were calculated from the dataset, whereas the likelihood, impact, and risk ratings were judgments of the research team.

These conclusions must be viewed within the limits of the study. The records came from a controlled testbed rather than an operational hospital. The proportion of attack records therefore reflects the design of the experiment and should not be interpreted as the actual frequency of such attacks in practice. The available sources also do not explain how the attack-category field was created or assigned.

The likelihood and impact ratings were based on research-team judgment and were not validated by external clinical or cybersecurity practitioners. The attack-category field was undocumented, the contiguous block boundaries followed the testbed label sequence, and the Data Alteration separation may reflect a protocol or capture artifact. The findings therefore apply only to the recorded testbed and should not be generalized to operational Hospital Information Systems without further validation.

V. RECOMMENDATIONS

The recommendations in this section are based on the risk levels and possible effects identified in the study. They are grouped according to the two attack categories recorded in the dataset. Because these measures were not tested or applied in the study environment, hospitals should assess them carefully before using them in actual operations.

A. Measures Addressing Spoofing

Spoofing received a risk score of 12 and was classified as High risk. The main concern is that an unauthorized source could present itself as a trusted device or system component. This could place both the confidentiality and integrity of monitoring communications at risk.

Hospitals that use connected patient-monitoring systems should consider the following measures:

1. Use strong authentication among medical sensors, monitoring gateways, and monitoring servers. Where the technology supports it, both endpoints should verify each other's identity before monitoring data are accepted.
2. Protect biometric and monitoring data during transmission through encrypted communication channels. This can reduce the risk of unauthorized access as information moves between devices, gateways, and servers.
3. Restrict access to the patient-monitoring communication path through network segmentation and appropriate access-control rules. Whenever possible, monitoring devices and servers should be separated from general hospital network traffic.
4. Maintain an approved inventory of connected monitoring devices and their authorized identities. Any unfamiliar device or unexpected change in identity should be investigated before its transmitted data are trusted.

Together, these measures address the central risk associated with Spoofing: an unauthorized source being mistaken for a legitimate part of the monitoring system.

B. Measures Addressing Data Alteration

Data Alteration received a risk score of 15 and was also classified as High risk. It was given higher priority because changing monitoring information poses a more direct threat to the reliability of data used in patient observation.

Hospitals should consider the following measures:

1. Protect monitoring data during transmission so that unauthorized changes can be detected before the information is stored, processed, or displayed.

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

2. Reject or isolate transmissions that fail authentication or integrity checks. Questionable data should not automatically replace previously accepted information without further review.
3. Maintain detailed transmission and system logs showing the originating device, receiving system, transmission time, and relevant verification results. These records can support investigations when a measurement is questioned or suspected of being altered.
4. Establish clear procedures for responding to suspected alteration of monitoring data. These procedures should include checking the affected device and communication path, preserving relevant logs, and notifying the appropriate technical and clinical personnel.

The purpose of these measures is to prevent altered monitoring information from being accepted without detection and to support a timely investigation whenever the integrity of the data is in doubt.

C. Measures Applicable to Both Attack Categories

Spoofing and Data Alteration were both classified as High risk within the healthcare-monitoring environment represented in the dataset. Their immediate effects differ, but both involve the security of communications among medical sensors, the monitoring gateway, the communication path, and the monitoring server.

Connected patient-monitoring systems should therefore be included in hospital cybersecurity risk assessments and asset inventories. Access to monitoring devices, gateways, communication channels, and servers should be limited to authorized personnel and systems.

Security settings should also be reviewed whenever devices, network arrangements, software, or the clinical use of monitoring information change. These reviews are important because changes in the monitoring environment may introduce new weaknesses or make existing safeguards less effective.

Hospitals may also maintain a risk register for connected monitoring systems. This can provide a clear record of identified threats, affected components, existing safeguards, responsible personnel, and planned actions. It also makes it easier to review risks and update security measures over time.

D. Recommendations for Further Research

Future studies should consider the following areas:

1. Validate the likelihood and impact ratings used in this study by consulting hospital cybersecurity personnel, clinical staff, biomedical engineers, and risk-management specialists.
2. Examine how altered or falsified monitoring data could affect clinical workflows, staff decision-making, and patient safety under different operating conditions.
3. Apply the risk-oriented approach to data collected from operational healthcare environments, subject to appropriate ethical, privacy, and institutional requirements. This would help determine whether the patterns observed in the controlled testbed also appear in practice.
4. Develop and test methods for detecting Spoofing and Data Alteration using network-flow characteristics. The present study identified descriptive differences, particularly for Data Alteration, but did not develop or evaluate a detection model.
5. Investigate whether combining network-flow information with device, system, and communication logs improves the detection of attacks affecting connected patient-monitoring systems.
6. Clarify how the attack-category field in the dataset was created and assigned, since this information was not explained in the available documentation.
7. Repeat the analysis using other healthcare cybersecurity datasets to determine whether the observed attack patterns and risk priorities remain consistent across different systems and test environments.

International Journal of Computer Science and Mobile Computing
A Monthly Journal of Computer Science and Information Technology

Crimary C. Mahinay *et al*, Volume 15, Issue 8, August 2026, pg. 39-58
(An Open Accessible, ISI Indexed, Fully Refereed and Peer Reviewed Journal)
ISSN: 2320-088X
Impact Factor: 7.056

REFERENCES

[1]. A. A. Hady, A. Ghubaish, T. Salman, D. Unal, and R. Jain, "Intrusion detection system for healthcare systems using medical and network data: A comparison study," *IEEE Access*, vol. 8, pp. 106576–106584, 2020, doi: 10.1109/ACCESS.2020.3000421.

[2]. Joint Task Force Transformation Initiative, *Guide for Conducting Risk Assessments*, NIST Special Publication 800-30 Revision 1. Gaithersburg, MD, USA: National Institute of Standards and Technology, Sep. 2012, doi: 10.6028/NIST.SP.800-30r1.

[3]. Washington University in St. Louis, "WUSTL EHMS 2020 Dataset for Internet of Medical Things Cybersecurity Research," Department of Computer Science and Engineering. [Online]. Available: https://www.cse.wustl.edu/~jain/ehms/index.html. [Accessed: Jul. 27, 2026].

[4]. H. T. Neprash, C. C. McGlave, D. A. Cross, et al., "Trends in ransomware attacks on US hospitals, clinics, and other health care delivery organizations, 2016–2021," *JAMA Health Forum*, vol. 3, no. 12, Art. no. e224873, 2022, doi: 10.1001/jamahealthforum.2022.4873.

[5]. C. Dameff, J. Tully, T. C. Chan, et al., "Ransomware attack associated with disruptions at adjacent emergency departments in the US," *JAMA Network Open*, vol. 6, no. 5, Art. no. e2312270, 2023, doi: 10.1001/jamanetworkopen.2023.12270.

[6]. B. Abbou, B. Kessel, M. Ben Natan, et al., "When all computers shut down: The clinical impact of a major cyber-attack on a general hospital," *Frontiers in Digital Health*, vol. 6, Art. no. 1321485, 2024, doi: 10.3389/fdgth.2024.1321485.

[7]. H. Naeem, A. Alsirhani, F. M. Alserhani, F. Ullah, and O. Krejcar, "Augmenting Internet of Medical Things security: Deep ensemble integration and methodological fusion," *Comput. Model. Eng. Sci.*, vol. 141, no. 3, pp. 2185–2223, 2024, doi: 10.32604/cmes.2024.056308.

[8]. W. Wu, F. Harrou, S.-M. Senouci, and Y. Sun, "Improving cyber-attack detection in Internet of Medical Things using ensemble deep learning methods," *Cluster Comput.*, vol. 28, Art. no. 891, 2025, doi: 10.1007/s10586-025-05660-y.

[9]. A. Naghib, F. S. Gharehchopogh, and A. Zamanifar, "A comprehensive and systematic literature review on intrusion detection systems in the Internet of Medical Things: Current status, challenges, and opportunities," *Artif. Intell. Rev.*, vol. 58, Art. no. 114, 2025, doi: 10.1007/s10462-024-11101-w.

