Notification Number: 604/2026
Notification Date: 20-08-2026

Name of the Scholar:	Mohammad Aijaz
Name of the Supervisor:	Prof. Mohammed Nazir
Name of the Department/Centre:	Department of Computer Science
Topic of Research:	Ranking and Prioritization of Security Attacks in Healthcare System
Keywords:	Healthcare security, Electronic Health Record, Threat modeling, Threat assessment, Security attack taxonomy, Cyber harms, Attack-Defense Tree, CAPEC, Social engineering, Attack path, Risk ranking, Prioritization, CVSS, EPSS

Findings

Healthcare delivery today rests on a digital foundation. Electronic Health Record (EHR) systems, cloud-hosted clinical applications, wireless body area networks, and connected medical devices have replaced paper and isolated terminals, and in doing so they have improved clinical efficiency in ways that would have been difficult to imagine a decade ago. This same connectivity, however, has quietly enlarged the attack surface of the hospital. A healthcare infrastructure is not an ordinary enterprise network; it is a safety-critical environment in which digital records and clinical procedures are tightly interwoven. A single technical exploit in such a system does not merely compromise data confidentiality, it may delay a diagnosis, corrupt a prescription record, or interrupt the continuity of care. Security in this domain is therefore inseparable from patient safety.

Despite this, the methods that are presently used to reason about healthcare security remain largely borrowed from generic information technology settings. Established frameworks such as NIST and ISO/IEC 27005 provide broad governance guidance but do not account for the interdependencies among clinical assets, technical weaknesses, and attacker behaviour. Scoring systems such as the Common Vulnerability Scoring System (CVSS) and the Exploit Prediction Scoring System (EPSS) supply useful but context-free numbers that say nothing about how a vulnerability sits within a clinical workflow. Structural approaches such as STRIDE, ATT&CK, CAPEC and TARA describe how attacks are composed, yet none of them carries an inbuilt mechanism to quantitatively rank the resulting attacks for a hospital. In practice, adversaries rarely rely on a single flaw; they chain several weaknesses together across multiple hops until a critical asset is reached, and very often the first hop is a human being rather than a machine. The prioritization of attacks in such settings has consequently remained subjective, inconsistent, and heavily dependent on the intuition of individual experts.

Taking these concerns as the point of departure, this research began with a systematic literature review of Threat Modeling and Assessment (TMA) methods, in which twenty-two methods were identified and examined against a defined set of modeling and assessment features and against their applicability and relevance to healthcare. The review established that most of the available methods are abstract and demand a trained cybersecurity expert before they can be applied at all. It further showed that threat identification and its validation have not been rigorously studied with respect to the

healthcare domain, and that existing methods provide no clear picture of how effectively system, vulnerability and attacker characterizations actually perform. These observations gave the study its direction: security knowledge for healthcare had first to be organized, then made operational, and only then could it be meaningfully ranked.

The first step towards that organization was a taxonomy of security attacks for healthcare systems, classified according to their effect on health-related data and services. Attacks were grouped into four categories, namely interception, modification, interruption and fabrication, and the harms arising from them were separately classified as patient-centric and organization-centric cyber harms. Studying the two together made visible something that isolated vulnerability analysis conceals, that is, the cascading nature of harm in a clinical environment, where a technical compromise propagates outward into treatment delays, loss of trust, financial penalty and reputational damage. The taxonomy was validated through a case study of a wireless body area network, which demonstrated how a single ransomware event unfolds into a chain of clinical and organizational consequences.

Structured knowledge is of limited value unless it can be connected to the system being protected. A semi-automated mechanism was therefore developed to map system-level information, particularly functional requirements, onto standardized CAPEC attack patterns, from which healthcare-specific Attack-Defense Trees are generated to represent multi-step attack scenarios in EHR systems. The persistent term-mismatch problem between requirement vocabulary and security vocabulary was addressed using word embeddings pre-trained on cybersecurity text, and retrieval quality was measured through mean reciprocal rank and mean average precision. The resulting trees give a system analyst a workable view of the relevant attacks and their defenses without requiring that analyst to first become a security specialist.

Because the human element remains the most frequently exploited entry point in a hospital, social engineering threats were modeled probabilistically rather than described qualitatively. Communication modalities and the principles of persuasion were represented in an attack tree to compute the Attack Occurrence Probability, while a Markov chain formulation was used to estimate the Attack Success Probability. The analysis indicated that authority-based email is the dominant path, combining high exposure with credible impersonation of administrative or technical staff, whereas face-to-face approaches occur rarely but succeed at a far higher per-interaction rate. Paths invoking two persuasion principles were found to occur less often yet succeed more reliably once initiated. Sensitivity analysis confirmed that authority is the single most influential parameter in the model, which suggests that awareness efforts in clinical settings are better directed at authority verification than at generic phishing advice.

The final stage of the work brought these strands together into a graph-based framework for ranking and prioritizing attack paths in EHR infrastructures. The framework parameterizes the system with

real vulnerability intelligence, discovers multi-hop attack paths through a constrained depth-first search, and quantifies each path using a sigmoid-normalized systemic risk score that combines path-level CVSS severity, EPSS exploitability and healthcare-specific data sensitivity weights. Applied to an enterprise-scale EHR case study, the framework identified one hundred and twenty-eight distinct attack paths and resolved them into four risk tiers, of which only seven proved critical while fifty-two were high, forty-two moderate and twenty-seven low. Systemic risk was thus found to be concentrated rather than uniformly spread, and the normalized formulation was essential in separating these few genuinely dangerous vectors from the surrounding noise that a linear additive model would have over-emphasized. The IoT gateway emerged as the entry point with the highest median risk and the greatest volatility, the authentication server emerged as the principal systemic bottleneck, and risk was seen to accumulate non-linearly with path length as an adversary traverses middleware. The tier distribution remained stable under ten per cent variation of the sensitivity weights, and the entire computation completed in well under two seconds, indicating that the approach is robust as well as computationally practical.

Taken together, these findings establish that the gap between qualitative threat modeling and quantitative risk scoring in healthcare can be bridged, and that doing so changes what security teams actually do. Instead of patching indiscriminately across a large and undifferentiated vulnerability backlog, a resource-constrained hospital may concentrate its limited defensive effort on the specific paths most likely to reach clinical data. The progression developed in this thesis, from a structured taxonomy of attacks and harms, through automated attack modeling and probabilistic treatment of human-centric threats, to path-based systemic risk prioritization, offers a coherent basis for informed security decision-making in clinical environments. Its wider implication is that security in healthcare should be treated as an element of patient safety and embedded within clinical process design, rather than delegated to information technology as a separate concern. The validation reported here was carried out in a controlled and simulated setting and should be read as a proof of concept; large-scale evaluation in live clinical infrastructures, integration of real-time threat feeds, and adaptive modeling of attacker behaviour remain the natural next steps.