WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

Cyber Security in Healthcare: Threats, Vulnerabilities, Regulatory Landscape and a Layered Defense Framework

Shri G. M. Kulkarni

Assistant Professor
gkjesatn@gmail.com
J. E. Society's K. A. Lokapur Arts, Commerce and Science College, Athani
Affiliated to Rani Channamma University, Belagavi
Karnataka, India

ABSTRACT—

The healthcare sector has undergone rapid digital transformation through electronic health records, telemedicine, cloud platforms and the Internet of Medical Things (IoMT). While these technologies improve efficiency and patient outcomes, they have expanded the attack surface and made healthcare institutions a prime target for cybercriminals. This paper presents a narrative review of cyber security in healthcare, examining the principal threats, including ransomware, data breaches, phishing, medical-device exploitation and supply-chain compromise, together with the structural reasons for the sector's vulnerability. Three major incidents (WannaCry and the NHS, the AIIMS Delhi ransomware attack and the Change Healthcare breach) are analysed to extract recurring lessons. The regulatory and standards landscape, including HIPAA, GDPR, India's DPDP Act 2023, CERT-In directions, the NIST Cybersecurity Framework and ISO/IEC 27001, is summarised, and a layered, risk-based defense framework integrating governance, people, identity, network, endpoint, application, data and response capabilities is proposed. Emerging approaches such as zero trust, AI-based detection, federated learning and blockchain are evaluated, and practical recommendations are offered for hospitals, with attention to resource-constrained Indian institutions. The paper concludes that cyber security in healthcare must be treated as a patient-safety issue rather than solely an IT concern.

***Index Terms*—Cyber security, healthcare, ransomware, Internet of Medical Things, electronic health records, data protection, zero trust, risk management.**

I. INTRODUCTION

Healthcare is one of the most information-intensive sectors of the modern economy. Hospitals, diagnostic laboratories, insurers, pharmacies and public health agencies collectively generate, store and exchange enormous volumes of sensitive data, including demographic details, clinical histories, imaging, genomic information and payment records. Over the last two decades this information has migrated from paper files to Electronic Health Records (EHR), Hospital Information Systems (HIS), Picture Archiving and Communication Systems (PACS), cloud platforms, mobile applications and a rapidly growing ecosystem of networked clinical equipment commonly called the Internet of Medical Things (IoMT).

This digital transformation has delivered clear benefits: faster diagnosis, remote consultation, continuity of care and better analytics for population health. However, it has also expanded the attack surface available to adversaries. Unlike a stolen credit-card number, which can be cancelled and reissued, a medical record is permanent, rich in personal detail and useful for identity theft, insurance fraud, blackmail and targeted social engineering. Consequently, health records are consistently valued by criminal markets above most other categories of personal data, and healthcare has repeatedly been reported as the industry with the highest average cost per data breach [3].

What distinguishes healthcare from other critical sectors is that a cyber incident can directly endanger human life. When ransomware encrypts an EHR system, clinicians may lose access to allergy lists, medication histories and diagnostic images; surgeries may be postponed, ambulances diverted and emergency rooms forced to revert to paper. A compromised infusion pump or patient monitor represents a potential physical hazard rather than merely a data-protection problem. Cyber security in healthcare is therefore best understood as an issue of patient safety and service continuity, in addition to confidentiality.

The objectives of this paper are: (i) to characterise the contemporary threat landscape facing healthcare organisations; (ii) to analyse why the sector is structurally vulnerable; (iii) to extract lessons from major real-world incidents; (iv) to summarise the regulatory and standards environment, with particular reference to India; (v) to propose a layered, risk-based defense framework; and (vi) to discuss emerging technologies and practical recommendations, especially for resource-constrained institutions such as district and semi-urban hospitals.

The remainder of the paper is organised as follows. Section II describes the research methodology. Section III reviews related work. Section IV presents the threat landscape, Section V explains systemic vulnerabilities, and Section VI discusses case studies. Section VII covers regulation and standards. Section VIII proposes the defense framework, Section IX examines emerging technologies, Section X offers recommendations, Section XI outlines future directions and Section XII concludes.

II. RESEARCH METHODOLOGY

This study adopts a qualitative, narrative-review methodology. Peer-reviewed literature, government and regulator publications, standards documents

and industry incident reports were examined to build an evidence base. Sources were selected for relevance to healthcare cyber security, credibility of the publishing body and recency. Priority was given to systematic reviews and bibliometric studies, official investigation reports of major incidents, and primary legal and standards texts.

The analysis proceeded in four stages. First, threats were identified and grouped into a taxonomy based on attack vector and impact on the confidentiality–integrity–availability (CIA) triad. Second, vulnerabilities were mapped to organisational, technical and human causes. Third, three high-profile incidents were studied as case studies to identify recurring failure patterns. Fourth, findings were synthesised into a layered defense model aligned with the NIST Cybersecurity Framework [4]. As a narrative review, the paper does not claim exhaustive coverage, and quantitative figures cited are drawn from the referenced sources and should be interpreted as indicative.

III. BACKGROUND AND RELATED WORK

A. Security Objectives in Healthcare

The classical CIA triad takes on distinctive meaning in clinical settings. *Confidentiality* protects patient privacy and underpins trust in the clinician–patient relationship. *Integrity* ensures that diagnoses, prescriptions, dosage values and laboratory results are accurate and unaltered; tampering can cause direct clinical harm. *Availability* ensures that systems and data are accessible at the point of care, often within seconds. Many authors argue that, in contrast with corporate environments where confidentiality dominates, availability and integrity are the primary concerns in hospitals [6], [7].

B. Literature Review

Kruse *et al.* conducted a systematic review of cyber security in healthcare and found that malware, phishing and ransomware were the most commonly reported threats, while mitigation strategies emphasised staff education, access control and encryption [5]. Coventry and Branley provided a narrative review that highlighted the effect of legacy technology, interconnected devices and the human factor, and argued that security must be balanced against clinical workflow [7]. Argaw *et al.* examined the risk profile of hospitals and noted that common mitigation measures include technical controls, governance and awareness, yet adoption remains uneven, particularly in small facilities [6].

Jalali *et al.* performed a bibliometric analysis showing that the literature has grown steadily but remains concentrated on privacy and regulatory compliance, with comparatively less attention to operational resilience and medical device security [17]. Fu and Blum explained how software on medical devices introduces risks that traditional safety engineering did not consider, and called for security to be built into the device lifecycle [8]. The European Union Agency for Cybersecurity (ENISA) has similarly identified ransomware and data-related threats as leading concerns for the health sector in Europe [16].

Research on enabling technologies is also expanding. Rieke *et al.* describe federated learning as a way to train machine-learning models across institutions without centralising patient data [9]. Azaria *et al.* proposed MedRec, a blockchain-based approach to managing permissions and access to medical records [11]. The zero-trust paradigm, formalised by NIST, replaces implicit trust in network location with continuous verification of users, devices and sessions [12]. While these contributions are promising, relatively few studies evaluate them under the budget, staffing and connectivity constraints typical of hospitals in developing regions, which motivates the practical emphasis of the present paper.

IV. HEALTHCARE THREAT LANDSCAPE

Threats to healthcare organisations originate from financially motivated criminal groups, state-sponsored actors, hacktivists, malicious or negligent insiders and, increasingly, opportunistic automated attacks that scan the internet for exposed services. Table I summarises the principal threat categories, their typical vectors and the primary security properties affected.

TABLE I
Principal Threats to Healthcare Systems

Threat	Typical Vector	Primary Impact
Ransomware	Phishing, exposed RDP/VPN, unpatched servers	Availability; patient safety; data extortion
Data breach / exfiltration	Stolen credentials, misconfigured cloud storage, SQL injection	Confidentiality; regulatory penalties
Phishing / social engineering	Email, SMS, voice and fake portals	Credential theft; initial access
Insider threat	Snooping, misuse of privileges, negligence	Confidentiality; integrity
IoMT / device exploitation	Default passwords, legacy OS, unauthenticated protocols	Integrity; physical safety
Supply-chain compromise	Vendors, billing firms, software updates	All CIA properties at scale
DDoS	Botnet traffic against portals and gateways	Availability of online services
Cloud and API abuse	Weak IAM, open buckets, insecure APIs	Confidentiality; integrity

A. Ransomware and Extortion

Ransomware is widely regarded as the most disruptive threat to hospitals. Attackers gain a foothold, move laterally, disable backups and then encrypt critical servers. Modern operators practise “double extortion”: data are exfiltrated before encryption and the victim is threatened with publication if payment is refused. Because clinical operations cannot easily pause, healthcare victims are perceived to be more likely to pay, which makes them attractive targets. The consequences extend beyond ransom payments to diverted patients, cancelled procedures, forensic costs and long recovery periods.

B. Data Breaches and Insider Misuse

Breaches occur through hacking, but also through lost devices, misdirected communications, improper disposal and unauthorised access by staff. Insider “snooping” into the records of acquaintances or celebrities is a well-documented privacy violation. Misconfigured cloud storage and exposed databases have led to large-scale leaks without any sophisticated intrusion. Annual cost-of-breach studies consistently rank healthcare at or near the top of all industries by average breach cost, reflecting the sensitivity of the data and the regulatory consequences [3].

C. Phishing and Social Engineering

Phishing remains the dominant initial access vector. Clinical staff work under time pressure, handle many external communications and often share workstations, which makes them susceptible to deceptive messages. Attackers also exploit telehealth and appointment systems by impersonating hospitals,

insurers or government health schemes. Business-email-compromise attacks targeting finance departments are common in the sector, redirecting payments to fraudulent accounts.

D. Medical Devices and the Internet of Medical Things

IoMT devices such as infusion pumps, ventilators, imaging scanners, patient monitors, pacemakers and wearable sensors are increasingly networked. Many were designed for safety and reliability rather than security, run outdated operating systems that cannot be patched without vendor approval, use hard-coded credentials and communicate over protocols lacking authentication or encryption [8]. Device lifecycles of ten to fifteen years far exceed the support cycles of embedded software components. A compromised device may serve as an entry point into the hospital network, may be manipulated to produce incorrect readings or may simply be rendered unavailable.

E. Third-Party and Supply-Chain Risk

Hospitals depend on external vendors for EHR hosting, billing, transcription, diagnostics, device maintenance and cloud services. A breach at a single supplier can expose the data of many institutions simultaneously. Remote-maintenance connections granted to vendors are frequently poorly monitored. Contractual security obligations, vendor risk assessments and network segregation for third-party access are therefore essential but often under-implemented.

F. Cloud, Telemedicine and Mobile Health

The rapid expansion of telemedicine and mobile health applications accelerated during the COVID-19 pandemic. These services extend clinical data beyond the hospital perimeter to patient devices, home networks and third-party platforms that the organisation does not control. Risks include weak authentication, insecure application programming interfaces (APIs), unencrypted data in transit or at rest, unmanaged personal devices and ambiguity over shared responsibility between cloud provider and healthcare customer.

V. WHY HEALTHCARE IS STRUCTURALLY VULNERABLE

A. Structural Factors

Several characteristics combine to make healthcare an unusually difficult sector to secure.

Legacy and heterogeneous technology. Hospitals operate a mixture of modern cloud applications and decades-old clinical systems that cannot be upgraded without costly revalidation. Unsupported operating systems remain in service because replacement could disrupt clinical workflows or invalidate device certification.

Clinical priority of availability. Security controls that introduce delay, such as forced logouts, complex passwords or multi-step authentication, can frustrate clinicians in emergencies. Staff therefore adopt workarounds such as shared logins, written passwords and propped-open sessions, which undermine security. Effective controls must be designed around clinical workflow, for example through proximity-card single sign-on and "break-glass" emergency access with audit logging.

Limited budgets and skills. Security is competing with equipment, staffing and infrastructure for scarce funding. Many hospitals, particularly small and mid-sized ones, lack a dedicated information security officer or a security operations capability. The shortage of trained cyber security professionals is a global problem and is acute in smaller towns.

High value of data. Medical records combine identity, financial and health information and cannot be revoked once leaked, which sustains strong criminal demand.

Interoperability and openness. Care delivery requires sharing information across clinics, laboratories, insurers and patients. Each interface is a potential point of compromise, and standards such as HL7 and DICOM were not originally designed with strong security.

Human factors and culture. Security awareness among clinical and administrative staff is uneven. In a culture where patient care is paramount, security procedures may be seen as obstacles. Leadership attention to cyber risk is often reactive, emerging only after an incident.

C. The Human and Organisational Dimension

Technical weaknesses alone do not explain the frequency of healthcare incidents; organisational behaviour plays a decisive role. Clinicians are trained to prioritise the patient above all else, and security measures that appear to delay care are quickly bypassed. Studies of hospital environments repeatedly describe shared credentials on busy wards, unlocked workstations in corridors and the use of personal messaging applications to transmit images and reports because official channels are slow [7]. These behaviours are rational responses to poorly designed systems rather than simple negligence, and they cannot be corrected by prohibition alone.

A human-centred approach therefore treats usability as a security requirement. Security teams should observe clinical workflows, involve doctors and nurses in selecting authentication methods and provide approved, convenient alternatives to risky workarounds, such as a secure messaging platform integrated with the EHR. Security metrics should include measures of friction, for example the average time taken to log in at a shared workstation, alongside conventional indicators such as patch coverage. Leadership should communicate that reporting a mistake, such as clicking a malicious link, is encouraged and will not be penalised, because rapid reporting is often the difference between a contained event and a major incident.

VI. CASE STUDIES

A. WannaCry and the National Health Service (2017)

In May 2017 the WannaCry ransomware worm spread globally by exploiting a vulnerability in the Windows SMB protocol for which a patch had been available for some weeks. The National Audit Office reported that a substantial portion of NHS trusts in England were affected or locked out of systems, with thousands of appointments and operations cancelled and some emergency patients diverted [2]. The investigation found that the NHS had not been specifically targeted; rather, vulnerable systems were infected indiscriminately. Contributing factors included unpatched and unsupported operating systems, insufficient network segmentation and weaknesses in national-level incident coordination. The key lesson is that basic cyber hygiene (timely patching, segmentation, backups and tested response plans) would have prevented or limited the impact.

B. AIIMS Delhi Ransomware Attack (2022)

In November 2022, the All India Institute of Medical Sciences (AIIMS) in New Delhi suffered a ransomware attack that disrupted its servers and forced many services, including outpatient registration, billing and laboratory reporting, to operate manually for an extended period. The incident, widely reported in national media, drew attention to the vulnerability of India's large public hospitals, which hold millions of patient records and operate with heavily interconnected legacy systems. The episode highlighted the importance of network segmentation, offline and immutable backups, continuous monitoring and rehearsed manual fallback procedures, and contributed to a stronger policy focus on health-sector cyber resilience in the country.

C. Change Healthcare (2024)

In February 2024, a ransomware attack on Change Healthcare, a major US health-payments and claims-processing subsidiary of UnitedHealth Group, disrupted prescription processing, claims and payments across the American healthcare system for weeks. The company later disclosed that the personal and health information of a very large number of individuals, reported in the order of one hundred million or more, was affected. Public reporting indicated that the intrusion exploited a remote-access portal lacking multi-factor authentication. This case illustrates two themes: concentration risk, where a single intermediary is a single point of failure for thousands of providers, and the continued importance of foundational controls such as MFA on all external access points.

D. Common Lessons

Across these incidents recurring patterns emerge: (1) exploitation of known weaknesses rather than novel techniques; (2) flat networks that allow rapid lateral movement; (3) backups that were inaccessible, untested or also compromised; (4) inadequate third-party and external-access governance; and (5) lack of rehearsed downtime procedures. These patterns indicate that substantial risk reduction is achievable through disciplined implementation of fundamentals.

VII. REGULATORY AND STANDARDS LANDSCAPE

Legal and standards frameworks shape security obligations for healthcare organisations. Table II summarises the most relevant instruments.

TABLE II
Selected Regulations and Standards

Instrument	Scope	Focus Relevant to Healthcare
HIPAA Security Rule [1]	USA	Administrative, physical and technical safeguards for electronic protected health information; risk analysis; breach notification
GDPR [15]	EU	Health data as special-category data; lawful basis, data minimisation, security of processing, breach notification, heavy fines
DPDP Act, 2023 [13]	India	Consent-based processing of digital personal data; duties of data fiduciaries; safeguards and breach intimation; penalties
CERT-In Directions, 2022 [14]	India	Mandatory reporting of specified cyber incidents within six hours; log retention; point-of-contact designation
NIST CSF 2.0 [4]	Voluntary, global	Govern, Identify, Protect, Detect, Respond, Recover functions for risk-based programmes
ISO/IEC 27001	Voluntary, global	Information security management system; certifiable
IEC 80001-1	Global	Risk management for medical-device IT networks

In the United States, the HIPAA Security Rule requires covered entities and business associates to conduct risk analysis, implement access control, audit controls, integrity controls and transmission security, and to notify affected parties of breaches [1]. In Europe, the General Data Protection Regulation classifies health data as a special category requiring heightened protection and allows significant financial penalties for failures of security [15].

In India, the Digital Personal Data Protection (DPDP) Act, 2023 establishes obligations for entities determining the purpose and means of processing digital personal data, including reasonable security safeguards, notification of breaches and limits on retention, with consent as the central basis for processing [13]. In parallel, directions issued by the Indian Computer Emergency Response Team (CERT-In) under the Information Technology Act require covered entities to report specified categories of cyber incidents within six hours of noticing them and to retain relevant logs [14]. The Ayushman Bharat Digital Mission (ABDM) is building a national digital health ecosystem with consent-managed health-record exchange, which further increases the need for strong security and privacy architecture across participating facilities. Hospitals should confirm current rules, thresholds and timelines from official sources, since implementing rules and enforcement practices continue to evolve.

Voluntary frameworks complement binding law. The NIST Cybersecurity Framework organises security outcomes into functions that guide a programme from governance and asset identification through protection, detection, response and recovery [4]. ISO/IEC 27001 provides a certifiable management system, and IEC 80001-1 addresses risk management when medical devices are connected to IT networks. Regulators of medical devices increasingly expect manufacturers to demonstrate secure-by-design practices, vulnerability management and software bills of materials throughout the product lifecycle.

VIII. PROPOSED LAYERED DEFENSE FRAMEWORK

Drawing on the preceding analysis, this paper proposes a layered, risk-based defense model (Fig. 1) that follows defense-in-depth principles and aligns with the NIST CSF. No single control is sufficient; resilience arises from complementary safeguards such that the failure of one layer does not result in total compromise.

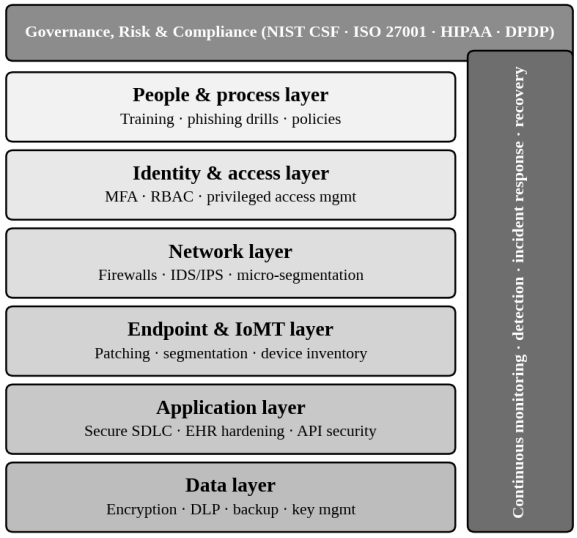


Fig. 1. Proposed layered defense framework for healthcare cyber security.

A. Governance, Risk and Compliance

Security must begin with leadership accountability. Boards and hospital administrators should designate a responsible executive, such as a Chief Information Security Officer (CISO) or equivalent, adopt a written security policy and perform periodic risk assessments that include clinical systems and medical devices. Risk registers should be prioritised by potential patient-safety impact, not only by financial loss. Incident-response, business-continuity and disaster-recovery plans must be documented, rehearsed through tabletop exercises and updated after each exercise or incident.

B. People and Process

Because human error remains a leading cause of incidents, continuous awareness programmes are essential. Training should be role-specific, short, repeated and reinforced through simulated phishing exercises with constructive rather than punitive feedback. A culture of reporting near-misses without blame improves detection. Processes should cover joiner–mover–leaver account management, secure disposal of media, vendor onboarding and change management.

C. Identity and Access Management

Compromised credentials feature in a large share of breaches. Multi-factor authentication should be mandatory for all remote access, administrator accounts, email and cloud services. Role-based access control should enforce least privilege, while privileged accounts should be vaulted and monitored. To reconcile security with clinical speed, hospitals can adopt single sign-on with proximity cards or biometrics on shared workstations and a documented break-glass mechanism that grants emergency access but triggers audit review.

D. Network Security and Segmentation

Flat networks enable rapid ransomware propagation. Clinical devices, administrative systems, guest Wi-Fi and vendor connections should be placed in separate segments with controlled traffic between them. Next-generation firewalls, intrusion detection and prevention systems, secure remote-access gateways and DNS filtering reduce exposure. Micro-segmentation and allow-list policies are particularly useful for IoMT devices whose behaviour is predictable.

E. Endpoint and Medical-Device Security

A complete and current asset inventory, including every networked medical device, is the foundation of any security programme. Endpoint detection and response tools, application allow-listing and timely patching should be applied to workstations and servers. Where a device cannot be patched, compensating controls such as isolation, virtual patching and restricted communication paths should be used. Procurement policy should require vendors to supply security documentation, support timelines, vulnerability-disclosure processes and software bills of materials.

F. Application and Data Security

EHR and other applications should be hardened, regularly tested through vulnerability scanning and penetration testing, and developed using secure software-development practices. Data should be encrypted in transit and at rest with sound key management; data-loss-prevention tools and database activity monitoring help detect abnormal extraction. Data minimisation, pseudonymisation and retention limits reduce the damage of any breach. Backups should follow the "3-2-1" principle, with multiple copies on different media, at least one offline or immutable, and restoration tested periodically.

G. Monitoring, Detection and Response

Early detection reduces dwell time and damage. Centralised log collection in a Security Information and Event Management (SIEM) platform, supported by a Security Operations Centre (SOC) or a managed detection-and-response service, enables correlation of alerts across layers. Hospitals that cannot afford a dedicated SOC may share services at a regional or state level. Response plans should define roles, communication channels, legal reporting duties (including CERT-In timelines in India), forensic preservation and clinical downtime procedures so that care can continue safely while systems are restored.

H. Risk Assessment and Prioritisation

A defense programme is only as effective as the risk analysis that guides it. A practical healthcare risk assessment begins by identifying critical assets, such as the EHR, PACS, laboratory systems, pharmacy systems, life-support devices and identity services, and then estimating for each the likelihood of compromise and the severity of impact. Unlike conventional enterprise assessments, severity should explicitly include clinical consequences: delayed diagnosis, medication error, diverted emergencies and loss of access to critical-care equipment. Table III illustrates a qualitative scoring approach that hospitals of any size can adopt using workshops with clinical, administrative and IT staff.

Scores are obtained by multiplying likelihood and impact ratings, typically on a three- or five-point scale, and the resulting ranking determines the

order in which controls are funded. Assessments should be repeated at least annually and after any major change, such as the introduction of a new clinical system, a merger or a new telemedicine service. Residual risks that cannot be reduced within budget should be formally accepted by senior management and recorded in the risk register, so that accountability is explicit.

TABLE III
Illustrative Qualitative Risk Scoring

Asset / Scenario	Likelihood	Impact	Suggested Priority Response
EHR server hit by ransomware	High	Severe	Segmentation; immutable backups; tested restoration; downtime plan
Stolen clinician credentials	High	High	MFA; privileged-access control; anomaly alerts
Compromised infusion pump	Low	Severe	Isolation; vendor patching; device monitoring
Vendor remote-access abuse	Medium	High	Just-in-time access; session recording; contracts
Misconfigured cloud storage	Medium	High	Configuration audits; encryption; least privilege
Insider snooping on records	Medium	Medium	Role-based access; audit logs; sanctions policy

TABLE IV
Prioritised Control Roadmap for Hospitals

Phase	Key Actions	Expected Benefit
1. Foundation (0–6 months)	Asset inventory; MFA on remote access and email; patching; offline backups; staff awareness; incident contacts	Removes the most commonly exploited weaknesses
2. Strengthen (6–18 months)	Network segmentation; endpoint detection; privileged-access control; vendor risk process; tested recovery	Limits lateral movement; faster recovery
3. Mature (18+ months)	SIEM/SOC or managed service; zero-trust access; device-security programme; continuous assurance	Faster detection; sustained resilience

IX. EMERGING TECHNOLOGIES AND APPROACHES

A. Zero-Trust Architecture

Zero trust assumes that no user, device or network location is inherently trustworthy and requires continuous verification and least-privilege access for every request [12]. In healthcare, zero-trust principles can be applied incrementally, beginning with strong identity and device posture checks for remote and administrative access, and extending to micro-segmentation of clinical networks. The main challenge is ensuring that verification does not impede urgent clinical tasks, so policies should be context-aware and carefully tested with clinicians.

B. Artificial Intelligence and Machine Learning

Machine-learning techniques support anomaly detection in network traffic, user behaviour analytics for insider threats, automated phishing detection and prioritisation of alerts in the SOC. For IoMT environments, behavioural baselines can identify devices communicating unexpectedly. Nevertheless, AI is double-edged: adversaries can use generative models to craft convincing phishing messages and automate reconnaissance, while the models themselves are vulnerable to data poisoning and adversarial inputs. Organisations should treat AI as an aid to, rather than a replacement for, human analysts and should secure the AI pipeline itself.

C. Federated Learning and Privacy-Enhancing Technologies

Federated learning allows multiple hospitals to collaboratively train models while keeping patient data locally, sharing only model updates [9]. When combined with differential privacy, secure aggregation or homomorphic encryption [10], it can reduce privacy risk in multi-institution research. These techniques are computationally demanding, and federated systems remain susceptible to inference and poisoning attacks, so they complement but do not replace conventional safeguards.

D. Blockchain for Access Control and Audit

Blockchain-based systems have been proposed to provide tamper-evident audit trails, patient-controlled consent and decentralised permission management for health records [11]. The immutability of distributed ledgers can strengthen integrity assurance, but concerns remain regarding scalability, energy consumption, the difficulty of erasing data under privacy law and the practical challenge of integrating with existing hospital systems. A pragmatic approach stores only references, hashes and consent records on-chain while keeping clinical data in secure off-chain repositories.

E. Post-Quantum Cryptography

Medical records have a very long confidentiality lifetime. Adversaries may harvest encrypted data today in the expectation of decrypting it once large-scale quantum computers become available. Healthcare organisations should therefore begin cryptographic inventories and plan a phased migration to quantum-resistant algorithms as standards mature.

X. CHALLENGES AND RECOMMENDATIONS

A. Challenges for Resource-Constrained Institutions

Smaller hospitals and clinics, which constitute a large share of India's healthcare delivery, face particular difficulties: limited funds, absence of dedicated IT security staff, dependence on a single vendor, use of unlicensed or outdated software, and poor awareness of legal obligations. Rural and semi-urban facilities may also have unreliable connectivity, which complicates updates and cloud-based protections. These institutions are nonetheless

increasingly connected to national platforms and insurance schemes, and thus may become weak links in the wider ecosystem.

B. Recommendations

Based on the analysis, the following recommendations are offered. *First,* treat cyber security as a patient-safety and clinical-governance matter, with board-level ownership and a budget line, rather than an IT overhead. *Second,* implement foundational controls before acquiring advanced tools: asset inventory, MFA, patching, segmentation, offline backups and staff training deliver the greatest risk reduction per unit cost, as the case studies indicate. *Third,* manage third-party risk through contractual security clauses, minimum-privilege and monitored vendor access, and exit and continuity plans for critical suppliers.

Fourth, prepare for failure. Hospitals should maintain and rehearse downtime procedures, such as paper-based registration and printed critical-patient summaries, so that care continues during an outage. *Fifth,* build a security-conscious culture through regular, practical training and non-punitive reporting. *Sixth,* adopt shared-service models: state or regional health authorities can provide centralised monitoring, threat intelligence, template policies and incident-response support to smaller facilities. *Seventh,* align with legal requirements, including data-protection and incident-reporting duties, and maintain evidence of due diligence. *Finally,* policymakers and regulators should consider incentives, minimum baseline standards for connected health facilities, support for workforce development and mandatory security requirements in the procurement of medical devices and health IT.

C. The Indian Context and the Digital Health Ecosystem

India's healthcare system is a mixture of large public institutions, a vast private sector and many small clinics and diagnostic centres. Public initiatives for digital health, including national health identifiers, consent-based record exchange and tele-consultation platforms, are rapidly increasing the volume of digital health data and the number of connected facilities. This creates a significant opportunity for continuity of care and for reaching remote populations, but it also means that the security of the overall ecosystem depends on its weakest participants. A small clinic connected to a national platform through a poorly secured computer can become an entry point that threatens many others.

Three practical priorities emerge for the Indian context. First, baseline security guidance proportionate to facility size is needed, for instance a short checklist for clinics covering licensed and updated software, strong passwords with MFA, regular offline backups, restricted access to patient records and a named person responsible for security. Second, capacity building is essential: universities, colleges and professional bodies can contribute by integrating healthcare-specific cyber security modules into computing, health-informatics and management curricula, and by offering outreach training to local hospitals. Third, awareness among patients should be strengthened, since informed individuals can recognise fraudulent calls, fake appointment links and unsolicited requests for medical or insurance information, and can exercise their rights over personal data under the prevailing legal framework [13].

D. Ethical and Patient-Trust Considerations

Cyber security in healthcare is also an ethical matter. Patients disclose intimate information on the understanding that it will be used only for their care, and a breach can lead to stigma, discrimination and loss of trust, particularly for sensitive conditions. Loss of trust may in turn cause individuals to withhold information or avoid seeking treatment, which harms both individual and public health. Organisations therefore have a duty not only to prevent breaches but also to respond transparently when they occur, by informing affected individuals promptly, explaining the likely consequences and offering practical support such as credit monitoring or guidance on recognising fraud.

Ethical considerations also arise in security operations themselves. Monitoring of staff activity to detect insider misuse must be proportionate, documented and consistent with privacy law and labour practices. The use of AI in clinical and security settings raises questions of transparency, accountability and bias. Responsible programmes should therefore embed privacy-by-design, data minimisation and clear governance of secondary uses of data, such as research and analytics, so that technical safeguards are supported by a coherent ethical framework.

Finally, ransomware payment decisions pose a difficult ethical and legal dilemma for healthcare leaders. Paying may restore services quickly and protect patients in the short term, yet it funds criminal groups, offers no guarantee of data recovery or deletion and may be restricted or discouraged by law or regulatory guidance in some jurisdictions. Decisions of this kind should be made in advance, as part of incident-response planning, with legal counsel, clinical leadership and law-enforcement engagement, rather than improvised under crisis conditions.

XI. FUTURE RESEARCH DIRECTIONS

Several areas merit further investigation. Empirical studies are needed on the actual security posture of Indian hospitals of different sizes, since published data are sparse and mostly drawn from developed economies. Research on lightweight, low-cost security architectures suited to resource-constrained environments could help narrow the capability gap. Methods for quantifying the clinical-safety impact of cyber incidents, for example on mortality and treatment delays, would improve risk assessment and justify investment. In the technical domain, further work is required on securing legacy medical devices through non-intrusive monitoring, on the robustness of AI-based clinical and security systems against adversarial manipulation, on the practicality of privacy-preserving analytics at scale and on migration paths to post-quantum cryptography. Finally, human-centred research should explore how to design security controls that are compatible with clinical workflow so that safe behaviour becomes the easiest behaviour.

XII. CONCLUSION

Healthcare's digital transformation has created major opportunities and equally significant risks. The sector combines highly valuable data, safety-critical operations, legacy technology and constrained resources, making it an attractive and vulnerable target. Incidents such as WannaCry, the AIIMS Delhi attack and the Change Healthcare breach demonstrate that failures of basic security hygiene, rather than exotic techniques, have often been decisive, and that the consequences reach patients directly.

This paper reviewed the threat landscape, explained structural vulnerabilities, examined regulatory frameworks and proposed a layered, risk-based defense model spanning governance, people, identity, network, endpoint, application, data and response capabilities. Emerging approaches such as zero trust, AI-assisted detection, federated learning and blockchain offer promise but must be adopted with realistic expectations and sound foundations. For institutions with limited resources, a phased roadmap that prioritises fundamental controls, together with shared services at the regional level, offers a pragmatic path to resilience. Ultimately, protecting health information systems is inseparable from protecting patients; sustained leadership commitment, investment and a security-aware culture are the keys to trustworthy digital healthcare.

REFERENCES

[1]U.S. Department of Health and Human Services, "The HIPAA Security Rule," 45 C.F.R. Part 164, Subpart C.
[2]National Audit Office, "Investigation: WannaCry cyber attack and the NHS," HC 414, Session 2017–2019, London, U.K., Oct. 2017.
[3]IBM Security, "Cost of a Data Breach Report 2024," IBM Corp., Armonk, NY, USA, 2024.
[4]National Institute of Standards and Technology, "The NIST Cybersecurity Framework (CSF) 2.0," NIST CSWP 29, Feb. 2024.
[5]J. Kruse, W. Frederick, K. Jacobson, and D. Monticone, "Cybersecurity in healthcare: A systematic review of modern threats and trends," *Technol.*

Health Care, vol. 25, no. 1, pp. 1–10, 2017.

[6]Y. Argaw *et al.*, "Cybersecurity of hospitals: Discussing the challenges and working towards mitigating the risks," *BMC Med. Inform. Decis. Mak.*, vol. 20, art. no. 146, 2020.

[7]L. Coventry and D. Branley, "Cybersecurity in healthcare: A narrative review of trends, threats and ways forward," *Maturitas*, vol. 113, pp. 48–52, 2018.

[8]K. Fu and J. Blum, "Controlling for cybersecurity risks of medical device software," *Commun. ACM*, vol. 56, no. 10, pp. 35–37, 2013.

[9]N. Rieke *et al.*, "The future of digital health with federated learning," *npj Digit. Med.*, vol. 3, art. no. 119, 2020.

[10]C. Gentry, "Fully homomorphic encryption using ideal lattices," in *Proc. 41st ACM Symp. Theory Comput. (STOC)*, 2009, pp. 169–178.

[11]A. Azaria, A. Ekblaw, T. Vieira, and A. Lippman, "MedRec: Using blockchain for medical data access and permission management," in *Proc. 2nd Int. Conf. Open Big Data (OBD)*, 2016, pp. 25–30.

[12]S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero Trust Architecture," NIST Special Publication 800-207, Aug. 2020.

[13]Government of India, "The Digital Personal Data Protection Act, 2023," Act No. 22 of 2023, Ministry of Law and Justice, New Delhi, India, 2023.

[14]Indian Computer Emergency Response Team (CERT-In), "Directions under sub-section (6) of section 70B of the Information Technology Act, 2000 relating to information security practices, procedure, prevention, response and reporting of cyber incidents," Ministry of Electronics and Information Technology, Apr. 2022.

[15]European Parliament and Council, "Regulation (EU) 2016/679 (General Data Protection Regulation)," *Off. J. Eur. Union*, L 119, 2016.

[16]European Union Agency for Cybersecurity (ENISA), "ENISA Threat Landscape: Health Sector," Heraklion, Greece, 2023.

[17]M. Jalali, A. Razak, W. Gordon, E. Perakslis, and S. Madnick, "Health care and cybersecurity: Bibliometric analysis of the literature," *J. Med. Internet Res.*, vol. 21, no. 2, e12644, 2019.