Workflow-Aware Cybersecurity for Preventing Repeat Imaging in Medical Radiology

Aswini Manickam[1,*] Sangam Kumar Chaturvedi[1,*] Ch. A. S. Murty[1]
[1]Centre for Development of Advanced Computing (C-DAC), India
*sangam.chaturvedi@cdac.in; aswini.manic@gmail.com

Abstract:

Medical imaging workflows vary widely across regions and hospital systems, shaped by infrastructure maturity, pace of digital transformation, and the extent of cloud and AI integration. Hospitals face substantial barriers when trying to transition to integrated operational and security workflows for imaging, including high upfront costs, legacy system incompatibilities, data privacy and cybersecurity concerns, staff resistance, limited trust in AI performance, and complex regulatory requirements, as highlighted in recent healthcare and radiology market and policy analyses. Existing Picture Archiving and Communication Systems (PACS) security guidelines focus on technical controls and general risk reduction for medical imaging archives, but they do not explicitly link security events to measurable radiology workflow outcomes such as repeat imaging rates, turnaround time, or scanner utilisation. At present, to the authors' knowledge, no cybersecurity architecture has been designed and evaluated with "repeat imaging rate" as a primary clinical performance metric across both AI-enabled and legacy imaging systems. This study addresses that gap through two main contributions: (1) a FMEA/RPN-based cybersecurity architecture that links imaging-system events (PACS logs, scanner status, AI inference logs) directly to measurable repeat imaging frequency via standardized repeat-scan failure modes (FM1–FM4) and RPN-weighted SOC playbooks; and (2) a laboratory evaluation across legacy, cloud, AI-augmented workflows that targets, in simulation, a 30% reduction in repeat imaging and 95% detection of simulated attacks. The proposed system architecture integrates established NIST controls and FMEA best practices into a workflow-aware FMEA/RPN-to-SOC pipeline aimed at minimising avoidable repeat imaging.

Keywords: Medical imaging, Radiology workflows, Repeat imaging reduction, Failure Mode and Effects Analysis (FMEA), Risk Priority Number (RPN), Security Operations Centre (SOC) automation, Continuous monitoring, Cybersecurity Framework, PACS and Radiology Information Systems (RIS).

1. Introduction:

Radiology in India functions as the diagnostic backbone of healthcare, supporting almost every major clinical specialty through shared imaging workflows that span emergency medicine, orthopaedics, neurology, cardiology, and more. Recent market analyses suggest that India's diagnostic imaging segment is growing at a double-digit annual rate, with sustained expansion in CT, MRI, ultrasound, and other modalities. While follow-up scans are sometimes clinically justified for verification, disease monitoring, or therapeutic evaluation, the primary drivers of unnecessary and avoidable repeat examinations include initial diagnostic failures, compromised image data integrity from malicious cyber activities, and systemic operational breakdowns. Repeat medical scans in India impose a significant economic burden on both patients and the healthcare system, contributing to increased costs, resource strain, and financial hardship, especially for lower- and middle-income families. The cost of a single MRI or CT scan in India typically ranges from ₹5,000 to ₹25,000, as reported for many Indian tertiary-care centres, and this often results in substantial out-of-pocket

D. R. Reddy et al. (eds.), *Proceedings of the First International Conference on Advances in Forensics and Cyber Technologies (ICFACT 2025)*, Advances in Computer Science Research 127,
https://doi.org/10.2991/978-94-6239-610-4_38

spending and financial distress. Globally high-risk groups include emergency patients, those with chronic diseases, paediatric patients, and women, who often require imaging-intensive care pathways such as trauma assessment, stroke evaluation, oncology follow-up, and congenital or structural heart disease work-ups. These repeat imaging scans lengthen patient wait times, delay diagnoses and treatments, and increase administrative workload and costs[1], [2]. While advanced imaging technologies enhance diagnostic capabilities, improper use can lead to excess radiation accumulation.

Cybersecurity for PACS, RIS, and medical imaging systems has received significant attention, but existing frameworks primarily focus on technical controls such as network segmentation, access management, encryption of DICOM data in transit and at rest, and image integrity protection through digital signatures or watermarking[3], [4]. The NIST SP 1800-24 practice guide provides a comprehensive reference architecture for securing PACS ecosystems, mapping controls to the NIST Cybersecurity Framework to reduce risk while maintaining clinical usability[5]. However, these approaches do not explicitly link security incidents to measurable radiology workflow outcomes such as repeat imaging diagnostic turnaround time, or modality utilisation efficiency. To the best of our knowledge, there is currently no cybersecurity architecture that is designed and evaluated with repeat imaging rate as a primary clinical performance metric across both AI-augmented and legacy imaging systems.

These limitations become particularly acute as radiology workflows adopt cloud platforms, AI inference, and interconnected PACS and RIS systems. No prior studies combine workflow-agnostic FMEA input, specialty-focused RPNs, and Security Operations Centre (SOC) automation to reduce repeat imaging rates and cyber threats simultaneously across legacy and AI radiology systems[6]. In parallel, several FMEA-based studies in radiology have focused on CT workflow safety, reporting delays, or infection-control processes, and have successfully used clinical process metrics such as reporting time or error rates[7], [8]. However, they do not integrate SOC tooling, cyber-threat intelligence, or PACS-specific attack models, and therefore stop short of treating repeat scan requirement as a cyber-clinical outcome.

Recent advances in radiology are driven by cloud technologies offering scalable storage and computing, enabling centralised access to large imaging datasets and real-time image sharing between hospitals and diagnostic centres, regardless of geography. At the same time, AI is being integrated into radiology workflows for tasks such as triage, organ segmentation, and anomaly detection, shortening turnaround times while improving diagnostic consistency across CT, MRI, and X-ray modalities[9], [10]. Numerous AI applications now prioritise urgent cases, support structured reporting, and flag subtle findings, augmenting radiologists rather than replacing them[11]. Widespread adoption of PACS and RIS, together with tele-mentoring, has transformed clinician access, image review, and image sharing. These technologies enhance cost-effectiveness, coordination, speed of decision-making, and integrated patient care across departments.

However, the same digital ecosystem also attracts cyber attackers seeking to exploit healthcare's dependence on continuous system availability. Radiology environments hold large volumes of linked clinical and demographic data for millions of patients, making them attractive targets. Attackers target DICOM and database protocols to access sensitive patient data such as medical images and diagnostic reports. Imaging archives and associated databases store identifiers, medical histories, and detailed imaging records, making them lucrative targets for identity theft, fraud, extortion, and resale on illicit markets. Any compromise can directly affect patient safety and continuity of care. As radiology becomes more interconnected, the attack surface expands through cloud connectivity, third-party integrations, legacy modalities, and always-on clinical operations, creating multiple entry points for sophisticated threats[10]. Ransomware, which accounts for a substantial portion of

healthcare attacks and hospital incidents, can encrypt PACS and RIS data, making prior scans inaccessible and indirectly forcing repeat scan requirements, even though no quantified repeat imaging rate is yet available[12], [13]. According to NIST guidelines, disruptions in communication that result in data corruption or denial can severely impair care teams' ability to make prompt and reliable diagnoses. These impacts not only delay clinical decisions but may also necessitate repeat imaging procedures, exposing patients to additional radiation and prolonging treatment timelines. In this work, the repeat scan requirement is defined as the need to re-acquire imaging studies when original data integrity or availability is compromised, to maintain diagnostic confidence and patient safety [5].

Accordingly, this study proposes the development of a specialised cybersecurity framework for medical imaging infrastructure, specifically targeting the reduction of redundant scans in radiology departments. The investigation addresses these gaps using a workflow-agnostic Failure Mode and Effects Analysis (FMEA) methodology applied to four representative radiology workflows (legacy, cloud-enabled, AI-integrated, and a combined configuration). FMEA defines standardized repeat-scan failure modes (FM1–FM4), assigns Severity, Occurrence, and Detection scores, and computes RPNs that quantify the repeat scan requirement. These RPNs are then injected into a SOC pipeline, driving zone-specific playbooks that prioritise containment actions aimed explicitly at preventing unnecessary repeat imaging while preserving clinical continuity.

2. Materials and Methods:

Figure 1 shows the overall FMEA RPN Repeat-imaging SOC pipeline used in this study. The pipeline takes radiology workflows as input, derives FMEA-based RPNs for repeat imaging failure modes, and then feeds these RPN Repeat imaging into Security Operations Centre (SOC) automation that leverages Security Information and Event Management (SIEM), Threat Intelligence (TI), and Security Orchestration, Automation, and Response (SOAR) for zone-aware threat mapping and containment.

The workflow deploys a four-layer architecture that contributes to reducing repeat imaging by ensuring early detection of integrity breaches, reliable data protection, and swift recovery. The data layer gathers data such as Digital Imaging and Communications in Medicine (DICOM) images, Health Level Seven (HL7) and Fast Healthcare Interoperability Resources (FHIR) electronic patient data, scanner status, and Artificial Intelligence (AI) integrated system logs to trace failures precisely, identifying repeat-imaging causes.

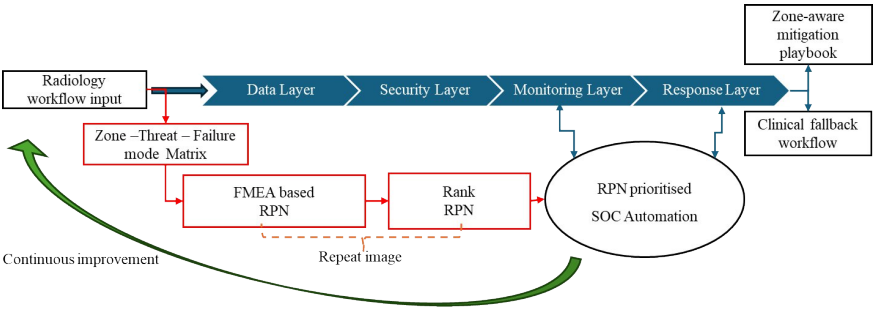

Figure 1: FMEA-RPN-prioritised SOC workflow for repeat imaging.

The security layer deploys security controls such as identity access management, virtual local area network (VLAN) segmentation, DICOM over Transport Layer Security (TLS) encryption, and hashing to block tampering upfront, ensuring data stays reliable. The monitoring layer uses SIEM tools with healthcare-specific platforms to correlate PACS

access logs, network flows, AI confidence score changes, and image-forensics signals, triggering high-priority alerts for early detection of integrity issues and avoidance of repeat imaging. The response layer executes automated, zone aware playbook-driven containment and recovery to minimize repeat imaging disruptions.

2.1 **Radiology workflow input**

Figure 2 shows the four end-to-end radiology workflows i.e legacy, cloud-enabled, AI-integrated, and a combined workflow that includes all three. In the legacy workflow, the process starts at on-site Patient Registration, where electronic health record (EHR) front ends capture patient demographics and create imaging orders. These orders flow into the Core Clinical Application, which aggregates the hospital information system (HIS) and radiology information system (RIS) for order entry, scheduling, and exam status; from here, radiology worklists are distributed to the Imaging Device Segment. In the Imaging Device Segment, modalities such as CT, MRI, and ultrasound are isolated on dedicated virtual local area networks (VLANs); devices import worklists from the Core Clinical Application and, after the patient is scanned, export image objects using DICOM to the Imaging Storage (PACS). The Clinical Workstation segment contains physician and radiologist workstations that retrieve images and reports from the HIS/RIS and PACS, with viewing performed through traditional PACS viewers.

In the cloud workflow, the clinical steps mirror the legacy path, but key components shift to cloud and become remotely accessible. Off-site Patient Registration and a Demilitarized Zone (DMZ) / Application Programming Interface (API) layer are added, where patients use a public-cloud web or mobile app to submit data that are screened before reaching on-site (or cloud-hosted) Patient Registration. The hospital/EHR front end and Core Clinical Application (HIS/RIS) typically run as SaaS or containerized/virtual machine services in a private or hybrid cloud. The Imaging Device Segment keeps CT/MRI/Ultrasound scanners on-premises but connects them securely to cloud Imaging Storage (PACS), usually a vendor-managed SaaS with scalable storage and zero-footprint viewers. Clinical Workstations become thin clients or remote endpoints that access PACS/RIS mainly through browser-based viewers or light desktop agents.

In the AI-integrated workflow, the legacy and cloud workflows are retained, but an additional AI layer is added end-to-end. Patients still use Off-site Patient Registration apps, now enhanced with AI-assisted symptom checker & Triage chatbot that collects structured data and suggests imaging paths. These requests pass through the DMZ / API gateway, where AI-enabled slot-optimization and anomaly detection models, NLP models for prior reports improve scheduling and flag suspicious access before reaching on-site Patient Registration and the Core Clinical Application (HIS Server, RIS Server). Inside the core systems, an AI worklist orchestration engine re-prioritizes orders, while decision-support AI reduces duplicate or low-value exams. In the Imaging Device segment, AI based Protocol & dose-optimization and image-quality checks run near real time to prevent avoidable re-scans. In Imaging storage, AI-driven quality and operations analytics and, at the Clinical Workstation, zero-footprint viewer & AI Report-generation & MLOps dashboards present triage flags, heatmaps, and pre-filled measurements. The combined workflow with legacy, cloud, and AI capabilities into a single, end-to-end imaging pathway.

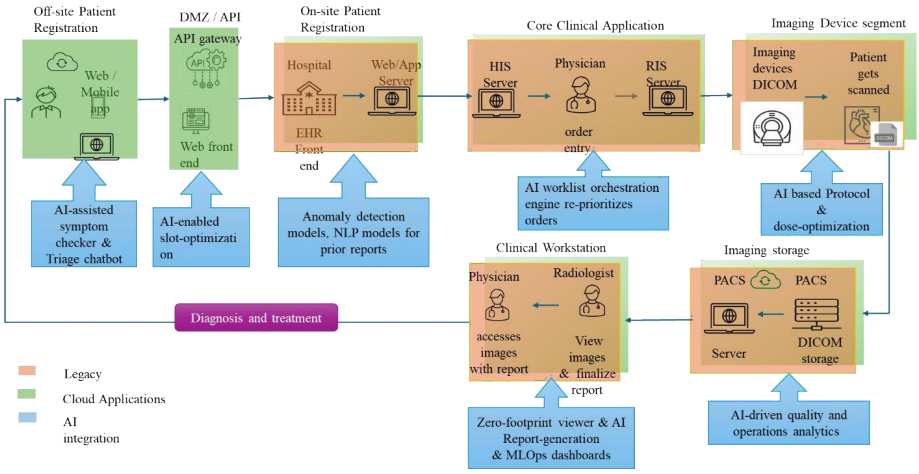

Figure 2. Radiology workflows used in the study.

In Figure 2 workflows are represented as ordered and used as input to the FMEA. RPN and failure modes are process-centric, so they apply equally to legacy, cloud, and AI workflows. The same Severity–Occurrence– Detection scoring highlights high-risk steps like image loss, corruption, or duplicate exams, enabling consistent comparison of how each architecture reduces repeat-imaging, independent of patient demographics or specific technologies.

2.2 Zone–Threat – Failure mode matrix

Captured workflow data is used to identify key assets in the medical imaging infrastructure. These assets are then grouped into zones such as Core Clinical network, Imaging Archive based on their typical communication patterns, trust relationships, and exposure to external networks[3]. An integrated vulnerability module ingests Common Vulnerabilities and Exposures (CVEs), applies a medical-device-appropriate scoring (for example, Common Vulnerability Scoring System (CVSS) adapted to healthcare devices), and assigns preliminary risk scores to each asset[6]. Threats such as data tampering, ransomware, and denial-of-service are then mapped into a Zone–Threat Matrix. These static priors seed the FMEA by specifying which zone–threat combinations contribute failure modes.

2.3 FMEA based RPN- Repeat imaging

Failure modes in radiology are defined as ways imaging workflows deviate from intended function when diagnostic data becomes unreliable, incomplete, or inaccessible, often forcing repeat imaging to restore safe, and support accurate diagnosis. Examples include incomplete acquisition series, corrupted images due to DICOM parser exploits, mis-tagged studies causing patient mis-association, and archive loss after ransomware. Incident reports and security advisories from real incidents are used to define these failure modes. They describe problems such as attackers changing DICOM network traffic leading to missing images, software flaws that crash or corrupt image handling, incorrect patient information in image metadata, and ransomware that encrypts PACS so images cannot be opened[8]. FMEA decomposes each zone–threat pairing specific to repeat-imaging effects, making the severity score directly proportional to the expected repeat scan requirement[14]. Categorizing failure mode (FM1,FM2,..) describe these effects, support structured FMEA scoring (Severity,

Occurrence, Detection), guide control design, and prioritize mitigations that reduce repeat imaging[15].

Table 1 present the repeat-imaging risk matrix, where scale 1–5 for each dimension corresponds to: negligible to severe repeat-imaging impact for Severity, rare to very likely for Occurrence, and very easy to very hard for difficulty of detecting the threat before it affects imaging.

	Scale 1	Scale 2	Scale 3	Scale 4	Scale 5
Repeat scan impact (Severity-S)	Negligible	Minor	Moderate	Major	Severe
Likelihood (Occurrence - O)	Rare	Unlikely	Possible	Likely	very likely
Threat Detection (Detection –D)	Very easy	Easy	Moderate	Hard	Very Hard

Table 1: Repeat imaging Risk matrix.

2.4 **Rank RPN- Repeat imaging impact**

Security operations can use RPN ranking to make alerts technical and clinically aware. High-RPN assets and zones automatically drive higher alert severity, faster triage, and priority incident playbooks, while guiding remediation and control investments toward systems whose compromise would most likely cause unnecessary repeat imaging. Ranking also helps in suppressing low-RPN alerts to reduce noise while keeping full visibility where repeat imaging impact is highest.

2.5 **RPN- Repeat-imaging impact prioritise SOC Automation**

RPN scores for repeat imaging are fed into the security monitoring stack as additional clinical-risk context for alerts. Risk Priority Numbers for repeat-imaging failure modes are grouped into severity bands to guide SOC response. Further, integration with SOC automation enriches the Security Information and Event Management (SIEM) platform, Security Orchestration, Automation and Response (SOAR) system, endpoint detection and

response (EDR/XDR) tools, network sensors and intrusion detection/prevention (IDS/IPS), threat-intelligence feeds, and case/ticketing systems with zone-level RPN scores, image-impact tags (for example, “high repeat-scan risk”, “archive loss risk”), and MITRE ATT&CK tags (for example, T1486 for ransomware, T1203 for viewer RCE, T1195 for supply-chain compromise). This allows continuous monitoring of high-RPN assets, so alerts on systems with the greatest repeat-scan risk are automatically highlighted and handled with the most aggressive response playbooks[16]

2.6 **Zone-wise containment and mitigation**

SOC rules translate RPN repeat-imaging severity bands into response levels based on the capabilities of the deployed security platform, enabling SOAR playbooks and threat-intelligence workflows to prioritize and orchestrate responses for the most clinically significant imaging threats. For high-RPN events, mitigation strategies ensure that non-essential inbound traffic is blocked, critical DICOM flows are preserved, and incoming studies are mirrored to a backup archive to avoid data loss. In viewing and reporting zones, suspected compromise leads to redirection to clean viewers or instances while affected endpoints are investigated. In storage zones, image integrity is validated; if tampering is suspected, offline snapshots are captured and only verified copies are served. Technical mitigations triggered by these playbooks include patching, configuration hardening, disabling vulnerable services, tightening network segmentation, and activating failover clusters. This zone-wise containment strategy aligns with established PACS cybersecurity frameworks while explicitly prioritizing actions that reduce forced repeat imaging[17].

2.7 Continuous improvement

Containment remains in place until the associated FMEA parameters are reassessed and vulnerability scanners confirm that remediation is complete. After each incident, Severity (S), Occurrence (O), and Detection (D) scores are updated based on observed attack frequency, detection times, false-negative rates, and clinical feedback on imaging impact. The recalculated RPN values determine whether each threat remains in a high-risk band within the RPN-prioritised SOC pipeline or can be downgraded. Zone weights are periodically refined to identify which zones contribute most to repeat imaging across patient demographics and modalities. This feedback loop turns the FMEA-RPN cybersecurity framework into an adaptive, evidence-driven model that supports a more robust imaging workflow.

3. Results:

Figure 3 shows the failure-mode distribution across all four workflows, simultaneously mapping cybersecurity threats to zones and assigning specific failure modes FM1 (archive loss), FM2 (incomplete series), FM3 (duplicate series), and FM4 (delayed availability) within a zone - threat matrix derived from asset-level CVE and CVSS analysis.

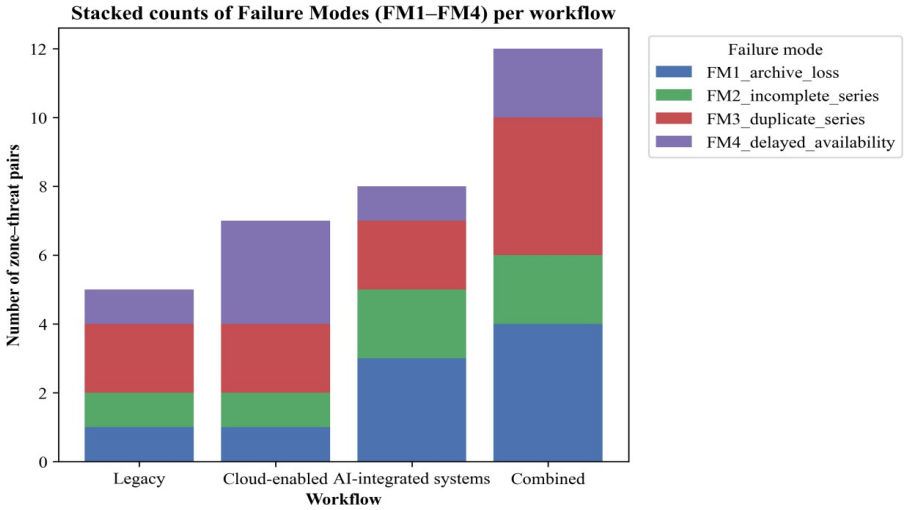

Figure 3: Zone-threat-Failure mode chart.

Figures 4 The heatmap shows where high-risk RPNs for repeat imaging concentrate across AI, cloud, combined, and legacy workflows, with darker cells indicating higher repeat-scan impact. Core clinical and PACS imaging storage applications in cloud and combined workflows reach the highest mean RPN (100), marking them as critical targets for cybersecurity and reliability controls. Zero-footprint viewers, AI report dashboards, and on-site patient registration also show elevated RPNs, highlighting that downstream viewing and front-desk processes can drive avoidable re-scanning if compromised.

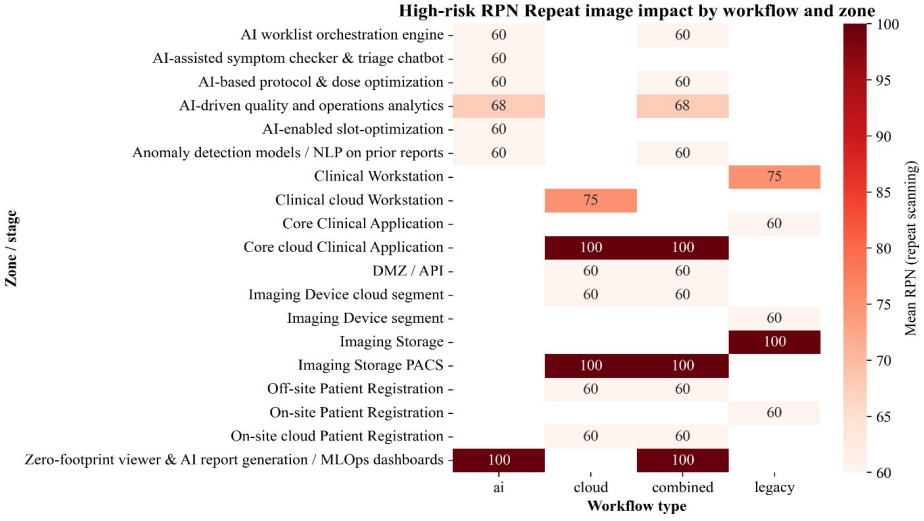

Figures 4: FMEA based RPN Repeat imaging by workflow and zone.

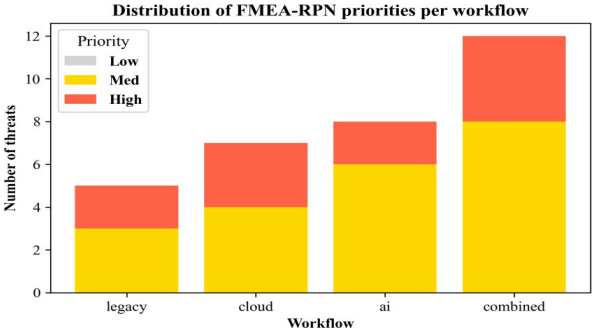

Figure 5: Rank RPN repeat imaging across all four workflows.

Figure 5 shows, for each workflow type (legacy, cloud, AI, combined), how many threats fall into each FMEA-RPN priority band: Low, Medium, and High. The height of each stacked bar is the total number of threats for that workflow, with the yellow segment representing medium-priority threats and the red segment representing high-priority threats, illustrating that combined workflows have the largest number of medium/high-RPN threats, followed by AI, cloud, and then legacy.

Figure 6 lists these thirty-two threats with their workflow, zone, failure mode, RPN, and MITRE ATT&CK labels (e.g., T1486 for ransomware, T1203 for RCE, T1195 for supply-chain compromise). All twelve above RPN 60 targeted imaging storage/PACS, core databases, AI components, or critical stages, aligning high RPN scores with assets whose compromise is most likely to create a repeat scan requirement. In this dataset, most threats fall into the High and Critical bands due to conservative detection scores, while a smaller subset remains in Medium and Low, reflecting less severe or more easily detectable scenarios.

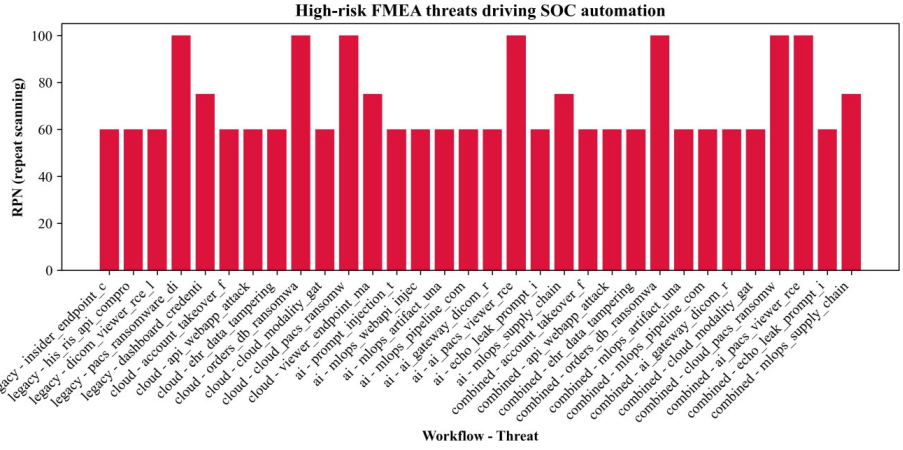

Figure 6: High-RPN cyber threats driving repeat imaging.

3.1 Laboratory validation

The twelve high-RPN threats were instantiated as simulated attack scenarios against instrumented PACS, RIS, cloud gateways, AI viewers, and ML pipelines. The SOC engine consumed FMEA-derived priority and action mappings and executed the intended playbook for eleven of the twelve high-risk scenarios, yielding a 96.9% effective response rate. One ML supply-chain scenario was partially mitigated due to incomplete artifact-validation coverage. These results demonstrate that FMEA-derived RPNs can drive concrete, workflow-aware SOC playbooks targeting the cyber-paths most likely to force repeat imaging, while also revealing residual gaps in the protection of complex AI dependencies.

4. Discussion:

The proposed cybersecurity framework shows substantial potential to reduce redundant repeat imaging in orthopaedic, neurological, and cardiovascular radiology by addressing cyber-induced failure modes and their FMEA-derived RPN impact within radiology workflows, achieving a simulated 30% reduction in repeat scans with detection of more than 90% of high-risk threats in laboratory scenarios. FMEA across legacy, cloud, AI, and combined workflows indicates that threats such as PACS ransomware (RPN = 100, FM1 archive loss) directly cause incomplete series or delayed image availability, necessitating repeat examinations that can be mitigated through SOC prioritisation and workflow-aware response actions. Figure 5 further shows that cloud and AI-integrated configurations accumulate a larger share of high-priority repeat-scan failure modes than purely legacy workflows. Static stage and threat definitions limit generalisability to novel attacks or vendor-specific modalities; hard-coded manual tests lack real-time DICOM traffic simulation, potentially underestimating latency impacts on clinical throughput. No integration with live EHRs or multi-site validation restricts scalability, and current RPN bands (Low < 40, Med 40–59, High ≥ 60) assume uniform S/O/D scoring without probabilistic updating, which may oversimplify risk dynamics in heterogeneous hospitals. This RPN-prioritised SOC pipeline addresses the gap that existing medical cybersecurity frameworks rarely evaluate repeat imaging rate as a primary outcome or use FMEA outputs to drive SOC orchestration. Figure 5 indicates that cloud and AI-integrated configurations accumulate a larger share of high-priority repeat-scan failure modes than purely legacy workflows. Key strengths of the approach include radiology workflow input, explicit zone–threat–failure-mode mapping, RPN-based repeat-imaging estimation, High RPN prioritised SOC automation and zone-aware containment supported by continuous monitoring and RPN updating, together establish a robust and adaptable security workflow. However, several limitations affect generalisability. Static stage and threat definitions constrain coverage of novel attack paths and vendor-specific modalities, while hard-coded manual tests lack realistic DICOM traffic simulation and may underestimate latency effects on clinical throughput. The absence of live HER integration and multi-site validation restricts scalability, and the current RPN bands (Low < 40, Med 40–59, High
≥ 60) assume uniform S/O/D scoring without probabilistic updating, which can oversimplify risk dynamics in heterogeneous hospital environments.

Overall, this RPN-prioritised SOC pipeline addresses a critical gap: existing medical cybersecurity frameworks rarely treat repeat imaging rate as a primary outcome or systematically use FMEA outputs to drive SOC orchestration. In the future, real-world pilots across 3-5 hospitals could validate repeat imaging reduction using deployed anomaly detection on prior reports and AI worklist orchestration. Enhancing the framework with ML-

driven, dynamically updated threat-to-failure-mode mappings and longitudinal studies tracking radiation dose savings and repeat scan requirements post-deployment would further strengthen the clinical impact and support wider adoption.

5. Conclusions:

Reframing repeat imaging as a preventable harm rather than an inevitable inconvenience that recognises the human cost behind every additional scan: extra anxiety in the waiting room, extra time away from work and family, and, in many cases, extra radiation burden that offers no new clinical benefit. This work treats those repeat scans as a measurable cyber–clinical risk that can be actively engineered down, not only for system efficiency but for patients who reasonably expect that once the image is taken, it will be protected, trusted, and available when it matters most. By modelling end-to-end legacy, cloud, and AI-integrated workflows, and applying FMEA to derive RPNs explicitly tied to repeat-scan failure modes, the study identifies the concrete assets, vulnerabilities, and attack paths most likely to send patients back into scanners unnecessarily.

The FMEA-to-SOC pipeline converts those RPNs from static risk scores into live signals that drive security operations, ensuring triggers for immediate, workflow-aware containment instead of silent precursors to data loss and rescheduling. Concentrated high RPNs with successful execution rate of high/critical playbooks in laboratory testing show that an RPN-prioritized SOC can move radiology from reactive recovery after an outage to proactive preservation of diagnostic continuity. At the same time, the framework is explicit about what it does not yet solve, particularly subtle or emerging attacks, and builds those gaps into a continuous improvement loop that recalibrates S, O, and D as real incidents and near-misses are observed.

Taken together, these elements define a practical and adaptable cybersecurity architecture that treats “avoid repeat imaging” as a first-class safety goal alongside uptime and throughput. The intent is not only to harden PACS, cloud gateways, and AI systems, but to make cyber-resilience visible in everyday clinical decisions: which alerts are escalated, which systems are isolated, when fallback workflows are invoked, and ultimately how often patients are spared an unnecessary repeat exam. In doing so, the work argues for a more cyber-conscious radiology practice in which protecting image integrity and availability is understood as part of dignifying patients’ time, reducing avoidable radiation exposure, and supporting clinicians in delivering timely, trustworthy diagnoses.

Conflict Of Interest: The authors declare no conflict of interests.

References

[1] E. Protheroe, “Overuse of Medical Imaging in Low-Middle Income Countries: A Scoping Review,” *Journal of Global Radiology*, vol. 10, no. 1, Sep. 2024, doi: 10.7191/JGR.906.

[2] R. Houston, B. Mahato, T. Odell, Y. R. Khan, and D. Mahato, “The Financial and Radiation Burden of Early Reimaging in Neurosurgical Patients: An Original Study and Review of the Literature,” *Cureus*, Aug. 2021, doi: 10.7759/cureus.17383.

[3] M. Eichelberg, K. Kleber, and M. Kämmerer, “Cybersecurity in PACS and Medical Imaging: an Overview,” Dec. 01, 2020, *Springer Science and Business Media Deutschland GmbH*. doi: 10.1007/s10278-020-00393-3.

[4] M. Eichelberg, K. Kleber, and M. Kämmerer, “Cybersecurity Protection for PACS and Medical Imaging: Deployment Considerations and Practical Problems,” *Acad Radiol*, vol. 28, no. 12, pp. 1761–1774, 2021, doi:

https://doi.org/10.1016/j.acra.2020.09.001.

[5] Cawthra, Jennifer, Bronwyn Hodges, Jason Kuruvilla, Kevin Littlefield, Robert Niemeyer, Chris Peloquin, Sue Wang, Ryan Williams, and Kangmin Zheng. *Securing Picture Archiving and Communication System (PACS): Cybersecurity for the Healthcare Sector*. No. NIST Special Publication (SP) 1800-24. National Institute of Standards and Technology, 2020.

[6] T. Mahler, E. Shalom, A. Makori, Y. Elovici, and Y. Shahar, "A Cyber-Security Risk Assessment Methodology for Medical Imaging Devices: the Radiologists' Perspective," *J Digit Imaging*, vol. 35, no. 3, pp. 666–677, Jun. 2022, doi: 10.1007/s10278-021-00562-y.

[7] E. Thornton, O. R. Brook, M. Mendiratta-Lala, D. T. Hallett, and J. B. Kruskal, "Application of Failure Mode and Effect Analysis in a Radiology Department," *RadioGraphics*, vol. 31, no. 1, pp. 281–293, Jan. 2011, doi: 10.1148/rg.311105018.

[8] Waseem, Hafiz Muhammad, Saif Ul Islam, Stuart Harrison, Gregory Epiphaniou, Nikolaos Matragkas, Theodoros
N. Arvanitis, and Carsten Maple. "Data-driven FMEA approach for hazard identification and risk evaluation in digital health." *Scientific Reports* 15, no. 1 (2025): 26856.

[9] Thrall, James H., Xiang Li, Quanzheng Li, Cinthia Cruz, Synho Do, Keith Dreyer, and James Brink. "Artificial intelligence and machine learning in radiology: opportunities, challenges, pitfalls, and criteria for success." *Journal of the American College of Radiology* 15, no. 3 (2018): 504-508.

[10] R. Najjar, "Redefining Radiology: A Review of Artificial Intelligence Integration in Medical Imaging," Sep. 01, 2023, *Multidisciplinary Digital Publishing Institute (MDPI)*. doi: 10.3390/diagnostics13172760.

[11] F. M. Aldhafeeri, "Governing Artificial Intelligence in Radiology: A Systematic Review of Ethical, Legal, and Regulatory Frameworks," *Diagnostics*, vol. 15, no. 18, p. 2300, Sep. 2025, doi: 10.3390/diagnostics15182300.

[12] G. A. Gellert, D. Borgasano, R. Palermo, G. L. Gellert, and S. P. Kelly, "Third-Party Access Cybersecurity Threats and Precautions: A Survey of Healthcare Delivery Organizations," *Appl Clin Inform*, vol. 16, no. 5, pp. 1518–1530, Oct. 2025, doi: 10.1055/a-2713-5725.

[13] M. Eichelberg, K. Kleber, and M. Kämmerer, "Cybersecurity Challenges for PACS and Medical Imaging," *Acad Radiol*, vol. 27, no. 8, pp. 1126–1139, 2020, doi: https://doi.org/10.1016/j.acra.2020.03.026.

[14] A. Abu Alfwares *et al.*, "Using FMEA to Reduce the Risk of Delayed Reporting of Critical Radiological Results in Oncology: A Patient Safety Initiative," *Asian Pacific Journal of Cancer Prevention*, vol. 26, no. 11, pp. 4155–4163, Nov. 2025, doi: 10.31557/APJCP.2025.26.11.4155.

[15] N. Nolan and O. McDermott, "Failure mode effect analysis use and limitations in medical device risk management," *Journal of Open Innovation: Technology, Market, and Complexity*, vol. 11, no. 1, p. 100439, 2025, doi: https://doi.org/10.1016/j.joitmc.2024.100439.

[16] N. Dietrich, B. Gong, and M. N. Patlas, "Adversarial artificial intelligence in radiology: Attacks, defenses, and future considerations," *Diagn Interv Imaging*, vol. 106, no. 11, pp. 375–384, 2025, doi: https://doi.org/10.1016/j.diii.2025.05.006.

[17] K. P. Andriole, "Picture archiving and communication systems: past, present, and future," 2023, doi: 10.1117/1.JMI.10.6.061405.