
Certified Specialist Programme in Medical Device Cybersecurity

Compliance and Auditing in Medical Device Cybersecurity

Risk assessment is the systematic process of identifying, estimating, and evaluating the potential for harm that may arise from the use of a medical device, especially when the device is connected to a network or communicates wirelessly. In the context of cybersecurity, risk assessment focuses on the likelihood that a threat will exploit a vulnerability and the impact that exploitation could have on patient safety, data integrity, and device functionality. For example, a risk assessment for an insulin pump might examine how a malicious actor could alter dosage settings, leading to hypoglycemia or hyperglycemia. The output of a risk assessment is typically a risk matrix or a risk register that quantifies each identified risk and informs mitigation strategies.

Vulnerability refers to a weakness in the device hardware, software, or process that could be exploited by an adversary. Vulnerabilities can be introduced at any stage of the product lifecycle, from design through manufacturing and post-market support. Common examples include hard-coded passwords, unencrypted communication channels, outdated operating system components, and insecure firmware update mechanisms. Identifying vulnerabilities requires techniques such as static code analysis, penetration testing, and threat modeling.

Threat is any potential circumstance or actor that could cause an unwanted incident. In medical device cybersecurity, threats may be external hackers, insider employees, nation-state actors, or even accidental misuse by clinicians. Each threat is characterized by its capability, intent, and opportunity to act. For instance, a ransomware group targeting hospital networks represents a high-capability, high-intent threat that could compromise device availability.

Threat model is a structured representation of the possible threats, the assets they target, and the pathways they might use. A threat model helps designers visualize how an adversary could move from a point of entry to a critical function. One widely used approach is the STRIDE model (Spoofing, Tampering, Repudiation, Information disclosure, Denial of service, Elevation of privilege). Applying STRIDE to a wireless pulse oximeter might reveal that spoofing of the device's identity could allow an attacker to feed false oxygen saturation values to a clinician's monitor.

Risk mitigation involves selecting and implementing controls to reduce the likelihood or impact of identified risks to an acceptable level. Controls may be technical (e.G., Encryption, authentication), procedural (e.G., Secure development practices), or organizational (e.G., Training). The effectiveness of mitigation is evaluated during verification and validation activities. For a cardiac rhythm management system, risk mitigation could include implementing mutual TLS for remote firmware updates, thereby preventing unauthorized code

injection.

Secure Development Lifecycle (SDL) is a set of practices integrated into each phase of product development to embed security from the outset. The SDL includes requirements definition, secure architecture, threat modeling, secure coding, code review, testing, and incident response planning. Embedding an SDL helps organizations comply with standards such as IEC 62304, which mandates software lifecycle processes, and IEC 62443, which provides guidance on industrial automation and control system security.

Defense in depth is a layered security strategy that employs multiple, overlapping controls so that failure of one control does not lead to a total breach. In a medical device context, defense in depth might combine network segmentation, device authentication, encrypted data storage, and runtime integrity checks. This approach aligns with the principle of "least privilege," ensuring that each component has only the access necessary to perform its function.

Encryption is the process of converting data into a coded form that can only be deciphered with the appropriate key. Encryption protects data at rest (e.G., Stored patient records) and data in transit (e.G., Communication between a bedside monitor and a central server). The use of strong, approved algorithms such as AES-256 and TLS 1.2 Or higher is required by many regulatory frameworks, including the EU Medical Device Regulation (MDR) and the US FDA's guidance on post-market management of cybersecurity.

Authentication verifies the identity of a user, device, or system before granting access to resources. Authentication mechanisms may be based on something the user knows (password), something the user has (smart card), or something the user is (biometric). Multi-factor authentication (MFA) is increasingly recommended for remote access to medical devices, as it reduces the risk of credential-theft attacks.

Authorization determines what actions an authenticated entity is permitted to perform. In medical devices, role-based access control (RBAC) is a common authorization model. For example, a nurse may be allowed to view vital signs, while only a biomedical engineer may be authorized to modify device firmware settings.

Audit trail is a chronological record of system activities, capturing who performed what action, when, and where. Audit trails are essential for forensic analysis after a security incident and for demonstrating compliance during regulatory inspections. In a ventilator, the audit trail might log each parameter change, user login, and firmware update, providing traceability for both safety and security assessments.

Audit scope defines the boundaries of an audit, specifying which processes, systems, and documentation will be examined. A well-defined scope ensures that auditors focus on critical areas such as network interfaces, data handling, and risk management files. For a multi-modal imaging platform, the audit scope could include the PACS interface, the device's internal software, and the hospital's network segmentation policies.

Audit plan outlines the objectives, methodology, resources, schedule, and criteria for evaluating compliance. The plan typically includes a checklist derived from applicable standards (e.G., ISO 14971, IEC 62443) and a

risk-based prioritization of audit activities. A detailed audit plan helps auditors allocate time efficiently, especially when assessing complex devices with numerous software modules.

Nonconformance is a deviation from a specified requirement, whether it be a regulatory clause, a standard, or an internal procedure. Nonconformances are identified during audits, inspections, or internal reviews and must be documented, investigated, and corrected. For instance, discovering that a device's firmware update process does not verify digital signatures would be a nonconformance to the security control requirements of IEC 62443-4-2.

Corrective Action is the step taken to eliminate the cause of a detected nonconformance or other undesirable situation. Corrective actions are distinguished from simple fixes; they address root causes to prevent recurrence. In the firmware update example, a corrective action might involve redesigning the update mechanism to include signature verification and updating the development procedures to mandate code signing for all future releases.

CAPA stands for Corrective and Preventive Action. CAPA is a systematic approach used by manufacturers to investigate nonconformances, implement corrective actions, and establish preventive measures that reduce the likelihood of similar issues arising again. CAPA processes are required by 21 CFR 820.100 and are scrutinized during FDA audits. Effective CAPA documentation includes problem description, root-cause analysis (often using tools such as fishbone diagrams or the 5 Why method), action plans, verification activities, and closure records.

Root-cause analysis is a methodical investigation to determine the underlying reasons for a problem. In cybersecurity, root-cause analysis might reveal that a vulnerability arose because a third-party library was not kept up-to-date, or because secure coding guidelines were not enforced during code reviews. Conducting thorough root-cause analysis is essential for developing robust preventive actions.

Vulnerability management is the ongoing process of identifying, evaluating, prioritizing, and remediating vulnerabilities throughout the device lifecycle. This includes regular scanning, patch management, and coordination with suppliers. A mature vulnerability management program often aligns with the NIST Cybersecurity Framework's "Identify" and "Respond" functions, ensuring that new threats are promptly addressed.

Patch management refers to the systematic deployment of software updates that fix known vulnerabilities. In medical devices, patch management must balance the need for rapid remediation with the regulatory requirement to validate any change that could affect safety or performance. The FDA's guidance on post-market management of cybersecurity emphasizes establishing a documented patch-management process, including risk assessment of each patch and verification that the update does not introduce new hazards.

Incident response is the organized approach to detecting, analyzing, containing, eradicating, and recovering from a cybersecurity event. An incident response plan (IRP) should define roles and responsibilities,

communication channels, escalation procedures, and post-incident activities such as lessons-learned analysis. For a pacemaker that experiences an unexpected reset due to a malicious command, the IRP would guide the manufacturer's technical team in isolating the device, gathering forensic evidence, notifying regulators, and providing a firmware fix.

Forensic analysis involves collecting and examining digital evidence to understand the nature and impact of a security breach. Forensic techniques may include memory imaging, log analysis, and reverse engineering of malicious code. In the medical device realm, preserving the integrity of patient data while conducting forensic analysis is critical, as any tampering could compromise clinical decisions.

Regulatory compliance denotes adherence to laws, regulations, and standards that govern the design, manufacturing, and post-market activities of medical devices. In cybersecurity, compliance is not optional; it is a legal obligation that ensures devices meet minimum safety and security thresholds. Key regulatory references include the US FDA's 21 CFR 820 (Quality System Regulation), the EU MDR, and the International Organization for Standardization (ISO) standards such as ISO 14971 (risk management) and ISO 27001 (information security management).

Medical Device Regulation (MDR) is the European Union's regulatory framework that replaced the Medical Device Directive (MDD) in May 2021. MDR introduces stricter requirements for cybersecurity, mandating that manufacturers perform a risk-based assessment of cyber threats and demonstrate that security measures are proportionate to the identified risks. MDR also requires a post-market surveillance plan that includes monitoring of cybersecurity incidents.

21 CFR 820 is the United States Food and Drug Administration (FDA) Quality System Regulation that establishes current good manufacturing practice (cGMP) requirements. Subpart G addresses "Design Controls," which include risk management and verification activities that directly impact cybersecurity. Subpart D covers "Device Master Record" and "Device History Record," where audit trails and change controls must be documented.

IEC 62304 is the International Electrotechnical Commission standard for medical device software lifecycle processes. It defines three safety classes (A, B, C) based on the severity of possible harm. For devices in class C, more rigorous verification, validation, and documentation are required, which includes cybersecurity testing and risk analysis. IEC 62304 also requires that software changes, including security patches, be tracked and reviewed.

IEC 62443 is a family of standards addressing security for industrial automation and control systems, widely adopted for medical device cybersecurity. The series is divided into general concepts (IEC 62443-1-1), system requirements (IEC 62443-2-4), and component requirements (IEC 62443-4-2). IEC 62443-4-2 specifically outlines secure development lifecycle practices, such as secure coding, vulnerability handling, and secure configuration management.

ISO 14971 is the risk management standard for medical devices. Although it predates modern cybersecurity

concerns, ISO 14971's methodology can be extended to include cyber-related hazards. The standard requires a risk management file that documents the identification of hazards, estimation of risk, evaluation of acceptability, and implementation of risk controls. Updating the risk management file to incorporate new cyber threats is a key compliance activity.

ISO 27001 is the information security management system (ISMS) standard. While not specific to medical devices, ISO 27001 provides a framework for establishing, implementing, operating, and continually improving an organization's information security posture. Many manufacturers adopt ISO 27001 to demonstrate systematic handling of confidential patient data and to support broader cybersecurity governance.

NIST Cybersecurity Framework (CSF) is a voluntary set of industry-wide best practices, standards, and guidelines for managing cybersecurity risk. The framework consists of five core functions: Identify, Protect, Detect, Respond, and Recover. Mapping medical device cybersecurity activities to the NIST CSF helps organizations align internal processes with a recognized model, facilitating communication with regulators and partners.

Identify function of the NIST CSF focuses on developing an organizational understanding of assets, data, and capabilities. For a medical device manufacturer, this includes maintaining an inventory of devices, software components, and third-party libraries, as well as documenting the data flows between the device and external systems.

Protect function involves implementing safeguards to limit or contain the impact of a potential cybersecurity event. Controls may include encryption, access control, training, and secure configuration baselines. In practice, protecting a bedside infusion pump might involve disabling unused network ports and enforcing strong password policies.

Detect function emphasizes the ability to identify cybersecurity events promptly. Detection mechanisms can be intrusion detection systems (IDS), log monitoring, or anomaly-detection algorithms. For a connected imaging system, continuous monitoring of network traffic can reveal unexpected data exfiltration attempts.

Respond function outlines the steps to contain and mitigate an incident once detected. This includes incident response planning, communication protocols, and forensic investigation. A well-defined response for a compromised telemetry device ensures that patient care is not disrupted while the breach is contained.

Recover function focuses on restoring normal operations after an incident. Recovery activities may involve deploying a clean firmware image, re-validating the device's safety, and documenting lessons learned. Post-incident recovery also includes updating the risk management file to reflect new insights.

Supply chain security addresses the risks associated with third-party components, software libraries, and services used in device development and manufacturing. Threats in the supply chain can manifest as

malicious code injection, counterfeit parts, or compromised manufacturing processes. Implementing a supplier security program, including due-diligence questionnaires and contractual security clauses, mitigates these risks.

Software Bill of Materials (SBOM) is a formal, machine-readable inventory of all software components, including open-source libraries, used in a device. An SBOM enables rapid identification of vulnerable components when new threats emerge. Regulatory bodies such as the FDA are increasingly encouraging the use of SBOMs to improve transparency and accelerate vulnerability remediation.

Secure configuration refers to the process of establishing device settings that minimize exposure to threats. This includes disabling default accounts, enforcing strong encryption, and applying least-privilege principles to services. Secure configuration baselines should be documented and validated as part of the device's verification activities.

Hardening is the practice of strengthening a system by reducing its attack surface. Hardening techniques may involve removing unnecessary services, applying security patches, and configuring firewalls. For a portable ultrasound device, hardening might include disabling Bluetooth connectivity when not in use and restricting USB ports to authorized maintenance tools.

Penetration testing is a controlled, simulated attack performed to evaluate the security posture of a device or system. Penetration testing can be internal (performed by the organization's own security team) or external (conducted by independent consultants). Results are documented in a report that highlights findings, severity ratings, and remediation recommendations.

Static code analysis involves examining source code without executing it to detect security weaknesses such as buffer overflows, injection flaws, and insecure cryptographic usage. Automated tools can scan large codebases and generate findings that developers can address early in the development cycle.

Dynamic code analysis evaluates software behavior while it is running, often using techniques such as fuzz testing or runtime monitoring. Dynamic analysis can uncover vulnerabilities that are not apparent in static analysis, such as logic errors that only manifest under specific input conditions.

Fuzz testing (or fuzzing) is a technique that provides random, malformed, or unexpected inputs to a program to trigger crashes, memory leaks, or other anomalous behavior. Fuzz testing is especially valuable for communication protocols and parsers in medical devices, where malformed packets could lead to denial-of-service conditions.

Secure boot ensures that a device only runs firmware that has been cryptographically signed by a trusted authority. Secure boot prevents the execution of malicious or tampered code during the device's power-on sequence. Implementing secure boot is a key requirement of IEC 62443-4-2 for class C devices.

Runtime integrity checking continuously monitors the device's software environment for unauthorized

changes, such as unexpected modifications to memory or code sections. Techniques include checksum verification, hash comparisons, and trusted execution environments (TEE). Runtime integrity checks help detect attacks that attempt to inject malicious code after boot.

Trusted execution environment (TEE) is a secure area of the main processor that isolates sensitive code and data from the rest of the system. TEEs can protect cryptographic keys, patient data, and critical control algorithms from tampering. Leveraging a TEE can simplify compliance with standards that demand strong protection of confidential information.

Data integrity ensures that information is accurate, complete, and unaltered during storage, transmission, or processing. In medical devices, data integrity is vital for clinical decision-making. Techniques such as message authentication codes (MAC) and digital signatures verify that data has not been modified in transit.

Data confidentiality protects information from unauthorized disclosure. For devices that store or transmit patient health information (PHI), confidentiality is mandated by privacy regulations such as HIPAA in the United States and GDPR in the European Union. Encryption, access controls, and proper key management are the primary mechanisms to achieve confidentiality.

Key management encompasses the generation, distribution, storage, rotation, and revocation of cryptographic keys. Poor key management can undermine even the strongest encryption algorithms. Organizations should follow best practices such as using hardware security modules (HSM) for key storage and implementing automated key rotation schedules.

Privacy impact assessment (PIA) evaluates how a device's data handling practices affect the privacy of individuals. A PIA is required under GDPR and is increasingly expected by other regulators. The assessment identifies data flows, assesses legal bases for processing, and proposes mitigation measures such as anonymization or consent mechanisms.

Anonymization removes personally identifiable information (PII) from data sets, reducing privacy risk. In the context of medical device telemetry, anonymization might involve stripping patient identifiers before transmitting usage statistics to the manufacturer for performance monitoring.

Consent management handles the process of obtaining, recording, and managing user consent for data collection and processing. For devices that collect location or biometric data, explicit consent may be required before enabling certain features. Consent records must be stored securely and be auditable.

Post-market surveillance (PMS) is the systematic process of monitoring a device's performance and safety after it has been released to the market. PMS includes the collection of field safety corrective actions (FSCA), adverse event reporting, and cybersecurity incident tracking. Effective PMS allows manufacturers to detect emerging threats and initiate timely mitigations.

Field safety corrective action (FSCA) is a remedial measure taken to reduce the risk of a device that is

already in use. FSCA may involve software updates, hardware recalls, or user notifications. For cybersecurity, an FSCA could be a mandatory firmware patch that resolves a critical vulnerability discovered after deployment.

Vigilance reporting refers to the mandatory communication of adverse events and safety concerns to regulatory authorities. In the European Union, vigilance reporting is required under the MDR, while in the United States, the FDA's Medical Device Reporting (MDR) system serves a similar purpose. Cybersecurity incidents that affect device safety must be reported through these mechanisms.

Medical Device Reporting (MDR) (not to be confused with the European MDR) is the US FDA's system for manufacturers and other stakeholders to submit reports of device-related adverse events. When a cybersecurity breach leads to patient harm or a significant safety risk, the incident must be reported under MDR regulations.

Risk acceptability is the determination that a residual risk, after all mitigations, is within the organization's tolerance levels and complies with regulatory expectations. Acceptability decisions must be documented in the risk management file and reviewed by senior management. In cybersecurity, a residual risk might be deemed acceptable if a vulnerability is low-impact, has a low likelihood of exploitation, and is mitigated by compensating controls.

Compensating control is an alternative security measure employed when a primary control cannot be fully implemented. For example, if a device cannot support full-disk encryption due to hardware constraints, a compensating control could be frequent data backups and strict physical access restrictions.

Risk ranking assigns a priority level to identified risks based on their severity and likelihood. Ranking helps allocate resources to the most critical risks first. Common ranking methods include a simple high/medium/low scale or a numeric matrix that multiplies severity by probability.

Residual risk is the remaining level of risk after all planned risk controls have been applied. Residual risk must be evaluated to ensure it is acceptable. In cybersecurity, residual risk may be influenced by the effectiveness of detection mechanisms and the speed of incident response.

Threat intelligence is the collection and analysis of information about current and emerging cyber threats. Threat intelligence feeds can be internal (e.g., Logs from deployed devices) or external (e.g., Industry sharing platforms, vulnerability databases). Leveraging threat intelligence enables proactive updates to risk assessments and mitigation strategies.

Vulnerability disclosure is the process by which researchers or users report discovered security flaws to the responsible party, typically the manufacturer. A coordinated disclosure program encourages responsible reporting and may include timelines for public disclosure after a fix is available. Many manufacturers now operate "bug bounty" programs to incentivize external researchers.

Bug bounty programs offer monetary rewards to security researchers who identify and responsibly disclose vulnerabilities. Implementing a bug bounty can expand the pool of expertise assessing a device's security, supplementing internal testing efforts.

Security by design is the principle of integrating security considerations from the earliest stages of product development. This contrasts with "bolt-on" security, where protections are added after the fact. Security by design aligns with the concept of "privacy by design" and is a cornerstone of modern regulatory expectations.

Security by default ensures that a device ships with the most secure configuration possible, requiring users to explicitly enable less-secure features if needed. For a wireless glucose monitor, this means shipping with encryption enabled and default passwords disabled.

Security testing encompasses a broad set of activities designed to validate that security controls function as intended. Types of testing include unit testing of cryptographic modules, integration testing of communication protocols, and system-level testing of threat scenarios. Security testing should be performed throughout the development lifecycle and repeated after any significant change.

Verification and validation (V&V) are core components of the software development process. Verification confirms that the product was built correctly according to specifications, while validation confirms that the built product meets user needs and intended uses. In the cybersecurity context, V&V includes confirming that encryption algorithms are correctly implemented and that authentication mechanisms resist known attacks.

Change control is the formal process for managing alterations to a device's design, software, or documentation. Change control ensures that any modification, including security patches, undergoes impact analysis, review, approval, and documentation. Failure to control changes can lead to non-compliance with 21 CFR 820 and IEC 62304.

Document control ensures that the correct versions of policies, procedures, and technical files are available to authorized personnel. Document control is essential for audit readiness, as regulators will review records such as risk management files, design history files, and CAPA reports.

Design history file (DHF) is a compilation of records that demonstrate the design controls applied to a medical device. The DHF includes design inputs, design outputs, verification and validation results, and design changes. For cybersecurity, the DHF must contain evidence of threat modeling, security testing, and risk assessments.

Device master record (DMR) contains the specifications, manufacturing processes, and quality procedures needed to produce a device. The DMR also includes software build instructions and configuration settings. Maintaining an accurate DMR is critical for traceability and for demonstrating compliance during audits.

Device history record (DHR) documents the production history of each individual device, including serial numbers, manufacturing dates, and any post-manufacturing modifications. The DHR can be used to trace which devices received a particular security patch, supporting both recall management and regulatory reporting.

Regulatory audit is an inspection conducted by a governmental authority, such as the FDA or a Notified Body, to verify that a manufacturer's quality system and products meet applicable regulations. Auditors may focus on cybersecurity aspects by reviewing risk management documentation, security testing results, and post-market surveillance data.

Internal audit is a systematic, independent examination performed by the organization itself to assess compliance with internal policies and external standards. Internal audits help identify gaps before they are discovered by regulators, allowing corrective actions to be taken proactively.

Audit checklist is a structured list of items that auditors use to verify compliance. Checklists are derived from the requirements of standards such as IEC 62443-4-2, ISO 14971, and 21 CFR 820. A well-crafted checklist ensures that auditors cover critical security controls, including authentication, encryption, and secure update mechanisms.

Audit evidence consists of records, observations, and artifacts that demonstrate compliance. Evidence may include test reports, configuration files, training logs, and interview transcripts. Collecting and preserving audit evidence in an organized manner speeds up the audit process and reduces the risk of findings.

Audit finding is a documented observation that a requirement has not been met. Findings are categorized by severity (e.g., Major, minor) and must be addressed through corrective actions. For example, an audit finding might note that the device's firmware does not verify digital signatures, constituting a major nonconformance.

Audit closure occurs when all identified findings have been resolved, verified, and documented. Closure requires evidence that corrective actions have been implemented and that the device continues to meet safety and security requirements. Auditors sign off on the closure, indicating that the audit is complete.

Regulatory submission is the package of documentation submitted to a regulatory authority to obtain market clearance or approval. Submissions for cybersecurity often include a risk management file, a security development plan, and evidence of testing. In the United States, a 510(k) or De Novo submission may contain a dedicated cybersecurity section.

510(k) submission is the US FDA pathway for demonstrating that a new device is substantially equivalent to a legally marketed predicate device. When the predicate includes cybersecurity controls, the new device must show comparable or improved security measures. The submission must address any new cyber threats that have emerged since the predicate's clearance.

De Novo submission is used for novel devices that have no appropriate predicate. De Novo submissions require a comprehensive risk management file, including a detailed cybersecurity risk analysis and mitigation plan. Because De Novo devices often introduce new technology, regulators scrutinize the security architecture closely.

Premarket approval (PMA) is the most stringent FDA pathway, requiring extensive clinical data and a thorough review of safety and efficacy. For high-risk devices, the PMA dossier must contain a robust cybersecurity risk assessment, evidence of secure development practices, and a post-market surveillance plan.

Clinical evaluation assesses the device's performance and safety in a clinical setting. While traditionally focused on physiological outcomes, modern clinical evaluations also consider cybersecurity aspects, especially when device functionality directly impacts patient care. For example, a clinical trial of a remote monitoring system must verify that data integrity is maintained throughout the study.

Human factors engineering (HFE) studies how users interact with a device, aiming to reduce use errors. In cybersecurity, HFE ensures that security features are usable and do not lead to work-arounds. Designing intuitive authentication workflows and clear warning messages helps maintain security without compromising usability.

Usability testing validates that users can operate the device safely and effectively. Including security-related tasks in usability testing uncovers potential gaps, such as confusing password policies that cause users to write passwords on sticky notes—a practice that undermines security.

Incident classification categorizes security events based on severity, impact, and scope. Common classifications include "low," "moderate," "high," and "critical." Accurate classification guides response priorities and reporting obligations. A critical incident, such as a ransomware attack that disables life-support functions, triggers immediate regulatory notification.

Risk communication involves informing stakeholders—patients, clinicians, regulators, and partners—about identified risks and the measures taken to mitigate them. Transparent risk communication builds trust and satisfies regulatory expectations for disclosure. For a device with a known vulnerability, manufacturers may issue a safety notice outlining the risk and recommended actions.

Security policy is a formal document that defines an organization's security objectives, responsibilities, and procedures. Policies cover areas such as access control, incident response, and data protection. A security policy must be approved by senior management and reviewed regularly to remain effective.

Security procedure provides detailed, step-by-step instructions for implementing the security policy. Procedures might describe how to apply a firmware patch, conduct a vulnerability scan, or perform a user access review. Procedures are the operational backbone that ensures consistent execution of security controls.

Security awareness training educates employees about cybersecurity threats, safe practices, and their role in protecting devices. Training should be tailored to different roles—engineers, clinicians, and support staff—and include phishing simulations, secure coding workshops, and incident-reporting drills.

Phishing simulation is a controlled exercise where employees receive mock phishing emails to gauge their susceptibility. Results inform targeted training and help reduce the risk of credential compromise that could be leveraged against medical devices.

Secure coding standards are guidelines that developers follow to avoid common programming errors that lead to vulnerabilities. Examples include CERT C Secure Coding Standard and OWASP Secure Coding Practices. Adherence to secure coding standards is often required by IEC 62443-4-2.

Code signing digitally signs software binaries to verify their authenticity and integrity. Code signing is a prerequisite for secure boot and for ensuring that only authorized firmware is installed on a device. The signing process must use a trusted certificate authority and protect private keys in a secure environment.

Certificate authority (CA) is an entity that issues digital certificates used for authentication and encryption. In medical device ecosystems, a private CA may be established to manage certificates for device-to-server communication, reducing reliance on public CAs that may not meet regulatory requirements.

Public key infrastructure (PKI) is the framework that manages public-key certificates, keys, and related services. PKI supports secure communications, device authentication, and code signing. Implementing a PKI requires policies for certificate issuance, renewal, revocation, and lifecycle management.

Certificate revocation list (CRL) is a list of certificates that have been revoked before their expiration date. Devices must check the CRL or use the Online Certificate Status Protocol (OCSP) to ensure that they do not trust compromised certificates. Failure to manage revocation can expose devices to man-in-the-middle attacks.

Online Certificate Status Protocol (OCSP) provides real-time verification of a certificate's revocation status. OCSP is more efficient than downloading a full CRL and is recommended for devices with limited bandwidth.

Secure firmware update is the process of delivering new software to a device in a way that ensures authenticity, integrity, and confidentiality. Secure update mechanisms typically involve encrypted transport, digital signatures, and rollback protection. Manufacturers must document the update process and provide a means for emergency patches.

Rollback protection prevents a device from being downgraded to an older, potentially vulnerable firmware version. Rollback protection can be implemented by storing a version number in a protected area of non-volatile memory and refusing to install firmware with a lower version identifier.

Emergency use authorization (EUA) is a regulatory pathway that allows the use of unapproved medical

devices during public health emergencies. Even under EUA, manufacturers must implement basic cybersecurity controls to protect patients and data. The rapid deployment of devices during emergencies highlights the need for pre-established security foundations.

Cybersecurity maturity model assesses an organization's capability to manage cybersecurity risks. Models such as the Capability Maturity Model Integration (CMMI) or the NIST CSF maturity levels provide a roadmap for improvement. Achieving higher maturity levels demonstrates to regulators and customers that the organization is proactive in managing cyber risk.

Continuous monitoring involves ongoing observation of devices, networks, and processes to detect security events in real time. Continuous monitoring may employ automated tools that collect logs, analyze traffic patterns, and generate alerts. The goal is to reduce the time between detection and response, minimizing potential harm.

Log retention specifies how long audit logs and security events must be stored. Retention periods are often dictated by regulatory requirements and may range from six months to several years. Secure log storage must protect integrity and confidentiality, preventing tampering or unauthorized access.

Secure lifecycle management integrates security activities throughout the entire product lifecycle—from concept to disposal. This includes planning, design, development, verification, production, post-market support, and end-of-life decommissioning. Secure lifecycle management aligns with both IEC 62304 and IEC 62443, ensuring that security is not an afterthought.

End-of-life (EOL) disposal addresses the safe retirement of devices, ensuring that sensitive data is erased and that hardware cannot be repurposed maliciously. Secure disposal may involve physical destruction of storage media or certified data wiping procedures. Proper EOL handling prevents data leakage and mitigates the risk of counterfeit parts.

Supply chain risk assessment evaluates the security posture of vendors, subcontractors, and third-party software providers. The assessment includes reviewing the supplier's security policies, certifications (e.g., ISO 27001), and incident history. High-risk suppliers may require additional controls such as contractual security clauses and independent audits.

Supplier security agreement is a contractual document that defines security expectations, responsibilities, and reporting requirements for a supplier. Agreements often include clauses for vulnerability disclosure, audit rights, and compliance with specific standards (e.g., IEC 62443).

Third-party component validation verifies that external libraries and modules meet the organization's security criteria before integration. Validation may involve reviewing the component's source code, checking for known vulnerabilities, and confirming that the component is maintained by a reputable vendor.

Security risk register is a living document that tracks identified cybersecurity risks, their status, mitigation

actions, and residual risk. The register is updated as new threats emerge or as mitigations are implemented. Maintaining a risk register supports ongoing risk management and audit readiness.

Security governance refers to the structures, policies, and processes that provide oversight of an organization's security program. Governance includes defining roles such as Chief Information Security Officer (CISO), establishing security committees, and setting strategic objectives aligned with business goals.

Chief Information Security Officer (CISO) is the senior executive responsible for establishing and maintaining the organization's security strategy, ensuring compliance, and coordinating response to incidents. In a medical device company, the CISO works closely with regulatory affairs, quality, and engineering to embed security throughout product development.

Security metrics are quantitative measures used to assess the effectiveness of security controls. Common metrics include mean time to detect (MTTD), mean time to remediate (MTTR), number of vulnerabilities found per code line, and percentage of devices with up-to-date patches. Tracking metrics helps demonstrate continuous improvement to regulators.

Mean time to detect (MTTD) measures the average time taken to discover a security incident after it occurs. Reducing MTTD is a priority for organizations seeking to limit the impact of attacks. Implementing real-time monitoring and automated alerting can improve this metric.

Mean time to remediate (MTTR) quantifies the average time required to resolve a security incident from detection to closure. Efficient MTTR depends on well-defined incident response procedures, clear communication channels, and available resources for rapid patch development.

Security incident log records details of each security event, including timestamps, affected assets, actions taken, and outcomes. The log supports root-cause analysis, regulatory reporting, and trend analysis. Maintaining a comprehensive incident log is essential for audit evidence.