Chapter 17

Security Analytics and Monitoring of Medical Devices

By Paul Koster

DOI: 10.1561/9781680838237.ch17

Published in *Cyber-Physical Threat Intelligence for Critical Infrastructures Security* by John Soldatos, Isabel Praça and Aleksandar Jovanović (eds.). 2021. ISBN 978-1-68083-822-0. E-ISBN 978-1-68083-823-7.

Suggested citation: Paul Koster. 2021. "Security Analytics and Monitoring of Medical Devices" in *Cyber-Physical Threat Intelligence for Critical Infrastructures Security*. Edited by John Soldatos, Isabel Praça and Aleksandar Jovanović. pp. 375–390. Now Publishers. DOI: 10.1561/9781680838237.ch17.

now
the essence of knowledge

Cybersecurity risks are increasing for connected medical devices, e.g., MRI, large systems for catheterization laboratories ("cath labs") with imaging equipment for diagnosis and treatment, and small devices like insulin pumps. Consequently, security requirements, norms, standards and regulations have been increasing, being a joint responsibility of medical device manufacturers and healthcare providers. Emerging aspect is to complement protection with detection and security monitoring powered by analytics.

Security monitoring can address risks that surface during the long lifecycle of a device. For example, security controls may stop functioning correctly over time, e.g., due failed patches or misconfigured firewalls. Similarly, the operational environment may pose threats, e.g., attacks from the network or (un)intentional misuse by people operating it.

The healthcare sector can benefit from security analytics and monitoring concepts in other domains e.g., IT. However, medical devices face several challenges such as strict medical validation requirements and complex lifecycle management, which requires a tailored approach.

This chapter outlines an approach for security monitoring powered by analytics to enhance the security posture of medical devices and its operational environment. Implementation experiences demonstrate feasibility. Empirical results show further that medical device security control status can be monitored with high accuracy and low false positive rate. Security monitoring of the operational environment is also promising.

The approach demonstrates potential to integrate in larger cyber threat management systems. The perspective of the medical device nicely complements other monitoring solutions such as network monitoring.

The expected impact on medical device security and its operating environment is very positive. Over time this can grow as medical device logging and log export capabilities are extended as part of their design, enabling more monitoring.

17.1 Introduction

The healthcare sector faces an increasing cybersecurity risk over the last few years. This also affects medical devices, e.g. MRI, large systems for catheterization laboratories ("cath labs") with imaging equipment for diagnosis and treatment, and small devices like insulin pumps. This can be attributed to increasing connectivity of medical devices to computer networks and convergence of technologies in the healthcare sector that has exposed vulnerable devices and software applications to security attacks. Furthermore, highly infectious and damaging malware and for-profit cybercrime are on the rise, which also affects the healthcare sector.

Abovementioned risks affect hospital assets: IT, patient data, and medical devices. The attacks that target medical devices are concerning as they have potential impact on clinical care and safety. A device infected with malware has the potential to disrupt hospital operations, expose sensitive patient information, compromise other connected devices, and harm patients. A compromised X-ray device could cause radiation overdose or uncontrolled movement of mechanical parts thus physically harming patients and clinical staff.

Ensuring medical device security is crucial, and requirements as well as recommendations from FDA, NIST, ENISA and EU MDR, etc. have increased over time. These need to be considered during device design and development. Ensuring security of medical devices is a joint responsibility of medical device manufacturers and healthcare providers [20]. The manufacturers need to apply security by design approach. They also have the potential to provide security monitoring services to help their customers maintain adequate level of security thus reducing risks. The healthcare providers need to use appropriate technical, physical, and procedural means to maintain a secure environment in which the devices will operate.

Insufficient maintenance may leave operational issues undetected and unresolved, both in terms of cybersecurity posture, but also in terms of patient care operations.

Security monitoring powered by analytics supports above joint responsibility. Monitoring the security posture of medical devices and their operational environment allows the associated security risks to be managed. For this, the field can borrow from fields where these concepts are already more established, e.g. monitoring of network infrastructure and enterprise IT. However, medical devices face specific challenges that require a tailored solution.

17.2 The Need for Security Monitoring

17.2.1 Medical Device Cyber Security

In recent years the state of the art in cybersecurity for medical devices has been catching up with other domains [1, 2]. This follows medical device reaching a tipping point with software-driven functionality and increasing connectivity of devices [3]. Before, cybersecurity for networked medical devices has been often "bolted on" at the end of the design cycle, rather than integrated as a key factor of the product development and value creation process [4]. Consequently, medical devices got challenged by basic cybersecurity hygiene that must be addressed during early engineering and design.

To get medical devices' cybersecurity state to a proper level, experts from academia and industry put together guidance, and regulatory bodies sharpened their expectations. For example, ENISA presented recommendations on security good practices for technical security measures for smart hospitals including networked medical devices [5]. Similarly, Haigh and Landwehr present a building code – organized in 10 categories – that provides a basis for reducing the risk that software used to operate medical devices is vulnerable to malicious attacks. Yet another proposal to secure medical devices proposes a platform approach and reference architecture, specifying requirements for security mechanisms and functionality [7]. On the topic of security monitoring, ENISA recommends implementing monitoring and intrusion detection mechanisms as part of state-of-the-art measures [5].

As for regulatory bodies, The FDA (Food & Drug Administration) introduced guidance to medical device manufacturers. The initial FDA guidance on Premarket Management on Cybersecurity in Medical Devices [8] identified general principles to be applied together with explicit requirements for cybersecurity functions reflecting priorities on addressing urgent issues like hard-coded password use. The new draft version takes this to a new level and expands on the general

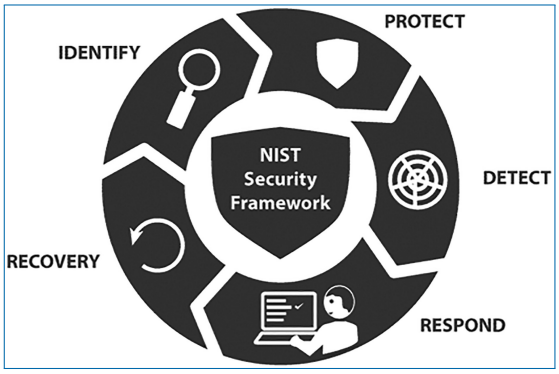

Figure 17.1. NIST CyberSecurity framework [24].

principles and risk assessment, recommends the application of NIST Cybersecurity Framework, and describes the specific design features and cybersecurity design controls it believes should be included in the design of a trustworthy device [9]. The design controls are grouped in the categories "Identify and Protect Device Assets and Functionality" and "Detect, Respond, Recover: Design Expectations". The FDA also published its guidance for Postmarket Management of Cybersecurity in Medical Devices to address and manage cybersecurity for devices after being on the market [10]. Although the state of the art of security analytics for medical devices is very much subject to research, in order to reach sufficient maturity for broad deployment, its application is stimulated and expected by the FDA: "Medical devices may not be capable of detecting threat activity and may be reliant on network monitoring. Manufacturers should consider the incorporation of design features that establish or enhance the ability of the device to detect and produce forensically sound postmarket evidence capture in the event of an attack. This information may assist the manufacturer in assessing and remediating identified risks" [10].

Similarly in Europe, regulators and industry step up their cybersecurity efforts, e.g. by providing medical device manufacturers guidance on how to fulfil requirements from the EU MDR (Medical Device Regulation) with regard to cybersecurity [20]. This includes the expectation to "consider design features that will allow the device to detect, resist, respond and recover from cybersecurity attacks" [21]. Specifically, monitoring of cybersecurity information sources for identification and detection of cybersecurity vulnerabilities and risk is identified with referencing ISO/IEC 27035 for incident response detection and (real-time) security monitoring and analysis [21].

17.2.2 Related Work

Security analytics is an emerging topic in healthcare with several aspects being explored. For example, Implantable Medical Devices (IMD) is a class of medical devices for which the concept of anomaly detection has been explored [13, 14]. IMDs are a special class because they typically have a small and well-defined functional scope, perform life-critical functions, have restricted and infrequent external communication, and are very resource constraint. One approach is to monitor the IMD externally, particularly the radio-frequency wireless communications, for anomalies [15, 16].

Anomaly and intrusion detection in case of other classes of medical devices raises similar questions. Logical candidate areas for research are malware detection beyond current pattern/signature-based approaches, host firewall enhancements with network anomaly detection capabilities, etc. Empirical studies must determine the effectiveness of such methods and how to optimally leverage them as a security control. The same applies to the translation of these concepts from network-/host-level to medical application-level. Since false positives may disturb the medical function of the device, detection is likely to be the primary function, and prevention can be recognized as a secondary derived function. When done right, this can be particularly powerful in combination with remote monitoring where the combined data may be used for security intelligence and risk management purposes. The combined availability of heterogeneous logs may enable a multi-analysis approach to study complex security events [17]. However, promising scientific results are lacking until now. Chaundry *et al.* present a middleware approach to postmarket surveillance of devices to provide the operational details of the devices to the manufacturers, and give device manufacturers the means to closely monitor the functioning of devices, upgrade devices, patch security vulnerabilities and monitor device performance thereby enhancing health care outcomes [18].

At macro scale, analysis has been done, such as the prevalence of security risks within the clinical setting, based on publicly available databases maintained by the Food and Drug Administration (FDA) to evaluate recalls and adverse events related to security and privacy risks of medical devices [19].

17.2.3 Challenges

Medical devices pose several challenges that need a tailored approach. Security monitoring and analytics solutions from other domains cannot be applied as-is for medical devices. First, medical devices have (very) long lifecycles compared to e.g. enterprise IT. Medical devices may be in use for many years, up to 15–20 years for large imaging modalities like MRI. As a consequence, all software must be

supported for a long time: clinical software, operating system, and also third party security software. To complicate things more, many medical devices continue to be used after the end-of-support date.

Second, medical devices are single purpose appliances. They use COTS (Commercially of the Shelf) software components such as Windows or Linux operating systems, but are tailored to their clinical function and cannot be administered as any other piece of IT equipment. For safe operation they need to be serviced by trained and certified service engineers. Similarly, due to their specific function, risks do not have to apply by mere fact that one of its components has a vulnerability. As a consequence, the field of medical devices and consequently also the hospital eco-system are very heterogeneous.

Third, modification of medical devices is subject to strict validation as invalidated modifications of a medical device can adversely affect performance or safety [20]. As a consequence, it is not possible to just make changes to system configuration to e.g. enable logging or install security agent software. The bigger the change and associated potential risks, the more effect required for validation.

These challenges significantly affect the solution space and timeline for security analytics and monitoring for medical devices. All three disfavor the addition of security monitoring software agents to the medical device. It also directs to a gradual introduction in the installed base, starting with devices that are (partially) capable today and expanding over time for new or upgraded products, accepting that it will take time. Architecturally, it points to approaches were the least impact is made to the device, e.g. implement the necessary logging extensions on the devices but perform analytics as well as monitoring external to the device.

To overcome some of the limitations that come with above approach, device-based monitoring can be complemented with passive network-based monitoring. It should be noted that monitoring externally observable behavior is complementary but not a replacement as it is limited in the insight it can offer. It is expected to further decrease over time as encrypted communication becomes the norm also for medical devices. Active network-based vulnerability scanning should not be used against medical devices in operation to avoid accidental affecting performance or safety.

17.2.4 Requirements

Security analytics and monitoring for medical devices aims at enabling monitoring the security relevant device internals, e.g. its security posture. Furthermore, it may enable contribute to monitoring its operational environment, e.g. the network it is part of or how operators use the medical device. Basis of the approach is that the medical device is the source of the observation, i.e. monitoring from perspective

of the medical device. Below we present selected, mostly functional, requirements. A more extensive consideration of requirements can be found in [23].

For vulnerability detection, the solution should inform relevant stakeholder, e.g. operators, of passively detected system vulnerabilities and weaknesses, even if they are not being actively exploited. For example, the security analytics solution should detect security functions on medical devices that are not functioning correctly or if the devices have unpatched vulnerable components.

For alert generation and remediation, the solution should generate timely alerts for the detected security events and send them to relevant stakeholders that can take remediation action.

For input to risk management model, the solution should provide insights and statistics about the security posture of the devices and its environment that becomes input to the security risk management model of the devices.

For post-incident analysis, the solution should facilitate forensic investigations of security incidents.

For anomaly and device misuse detection, the solution should be able to detect suspicious events from the logs, e.g. improper user behavior or malicious network traffic.

For security trend analysis, the solution should analyze logs from medical devices to detect trends that indicate potential security attacks such as malware infestation or data leakage.

As a non-functional requirement, for efficient serviceability, only few false positives are tolerated for alerts relating to medical device security posture. Of course, accuracy is important, but the time spent on a false alert goes at the expense of other service actions required to treat patients, particularly considering that some service actions must be performed on-site at medical devices and outside patient treatment moments. These two risks must be balanced per aspect that is monitored.

17.3 Architecture

Medical devices are the primary source of security relevant data for security analytics and monitoring. This is input for security models that define what is being monitored and embed the analytics logic for detection and alerting. As such, security models are the core function in the security analytics and monitoring architecture for medical devices, depicted in Figure 17.2 with other functions.

17.3.1 Medical Devices as Data Source

Medical devices are the primary source of security relevant data for security analytics and monitoring. They are also the primary target of monitoring. Although medical

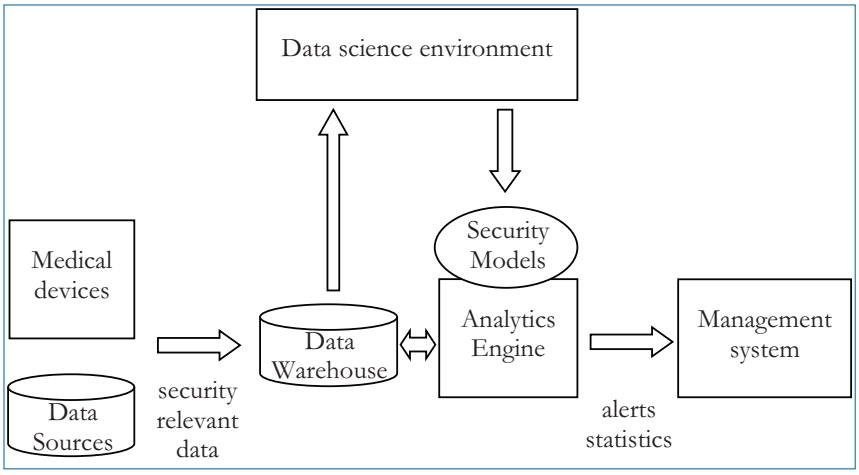

Figure 17.2. Security analytics and monitoring architecture.

devices nowadays typically have logging and auditing capabilities, their capabilities must be adapted and extended for the new purpose.

First, the right data must be identified, logged and made available for analytics. Security relevant data included in scope includes operating system and application log data, security configuration data, security controls usage data, etc.: operating system (security) event logs, (host) firewall logs and its configuration, endpoint protection logs, relevant (parts of) medical application logs, etc.

Several challenges apply to log availability. For example, existing logging primarily serves different purposes. Application logs may be focused on debug traces for support and development, not security. And the operating system may do extensive logging, but not necessarily have the optimal set enabled for security yet.

A further challenge with existing logging relates to the intended recipients of the information. Identifiable information of medical device users, such as hospital staff and patients, must remain on the device/in the hospital. This, for example, can mean that audit data collected for e.g. IHE ATNA [27] or DICOM [28] cannot be used as-is. The overall logging and auditing strategy of medical devices should support differentiation among recipients. In some cases, this means selecting certain log files for purposes. In other cases, this means inclusion or exclusion of information, e.g. de-identification of certain logs when exported from a device.

For both challenges holds that new logging facilities can take these into account as requirements from the start.

Second, logging and log export must be realized in a way that does not negatively impact the performance of the medical device. In the basis, logging and log exporting can be a resource-light operation that does not influence the clinical function of the device or patient safety. However, certain (third party) logging

functionalities may impose such an influence. For example, operating systems might be configured for excessive logging that under certain conditions, e.g. a high volume of network traffic, affect the performance. Similarly, methods to exports logs may differ, e.g. methods to export Windows Event Logs vary greatly in resource usage. This must be carefully addressed particularly when retrofitting in existing products.

Third, data quality and data robustness must be sufficient for the purpose of security analytics and monitoring. Existing logs, local to the device, may be sub-optimal or not meet the requirements for security analytics and monitoring. For example, for security event timestamps correct time is essential [27]. Furthermore, timestamps should be captured in a format that captures universal and local time to properly relate events, i.e. use UTC (Coordinated Universal Time) with local offset according to ISO 8601/RFC 3339 [11]. Similarly, logs should not depend on the system locale or language, use proper encoding (e.g. UTF-8) and strictly adhere to data exchange standards (XML, JSON, etc.). Use of standards would be useful here, but it is recognized that standards leave a gap in this field and in their absence application specific formats must be used.

17.3.2 Other Data Sources

Other data sources complement data from medical devices to maximize utility for security analytics and monitoring. A first category of this concerns product design data. This includes the software bill of material, supported and secure versions of the software, the baseline configuration as the product is designed and shipped, as well as insight in valid configurations. Furthermore, insights from the product security risk assessment may be input to analytics and alerting models.

A second category concerns threat intelligence and vulnerability data: CVE databases, vendor patch data, etc. Combined with above this enables very risk-oriented tailored monitoring.

A third category comes from customer support and complaint data. Such feedback data can reveal patterns over time, which can motivate new monitoring models or product design changes.

Of course, much more data may be as an input to the analytics process. A direction to explore is incorporation of information from the device operational environment.

17.3.3 Data Warehouse

To develop high quality security models security analytics requires sufficient historic data. This data must cover a large enough set of equipment stored in a

data warehouse. For certain types of medical equipment this is ideally equal to the complete installed base.

To add new security relevant data to the data warehouse, log data files are received, stored in a data lake, and ETL-ed in the data warehouse.

The data warehouse, regardless if it is a generic relational DBMS or dedicated security solution, is a high performance, highly scalable analytic database that receives daily log files from medical devices over mutually authenticated secure channels, and stores it in the database system.

Similarly, a data warehouse is needed for operational security monitoring, i.e. apply developed security models on incoming data. Yet, typically less historic data is needed than for abovementioned data science purposes to develop models. They may require some historic data, but typically not very far back in time. However, if forensics are needed in follow-up of an alert, then historic data may be useful.

17.3.4 Data Science Environment

In the data science environment, data scientists with security subject matter expertise or vice versa use the tooling of their liking on the historic data in the data warehouse to develop and validate security models. These models will run on the analytics engine in the production environment.

17.3.5 Analytics Engine

The security analytics engine runs the security analytics models on data from the data warehouse to generate alerts and statistics for the management system. Security models range from basic static rule-based to complex machine learning models. Models can be implemented in generic software languages (Python, JAVA), languages for statistics and data analytics (e.g. R) or security platform specific languages.

The analytics engine executes the models on pre-defined intervals, e.g. daily, a few times a day, and continuously for specific models for (near) real-time monitoring.

Generated alerts go to the applicable security management system. The Syslog protocol [25] enables a basic generic interface, but ideally dedicated interfaces are utilized to realize better (semantic) interoperability.

17.3.6 Management Systems

Security management systems enable the follow-up on security model outputs by parties responsible for the remediation. The shared responsibility between medical

device manufacturers and care delivery organizations for secure operation of medical devices means that multiple parties may be involved and responsible for their part.

The manufacturer service case system may be the target for alerts with predefined actions on non-functioning security controls. Subsequently, field or remote service engineers fix the problem, e.g. by correcting a configuration, installing a patch or updating/reinstalling the software component, etc.

The manufacturer SOC (Security Operations Center) platform may be the target for qualified security incidents and anomalies relating to the medical device. SOC operators, security experts, product security officers, etc. subsequently analyze these and define remediation action.

The manufacturer security risk assessment system may be the target for security and risk statistics. Product security officers can subsequently update the risk assessment and manage the risk for medical devices. The statistics may also go to installed base security KPI tracking dashboard. This enables for example the effectiveness of security controls at an aggregate level.

The hospital cyber threat management system (CTMS) may be the target for alerts relating to the operational environment of the medical device. Its SOC, CISO, IT, etc. staff can subsequently process and remediate the problem. For example, the network may be reconfigured to meet the medical device installation instruction again, investigate operator behavior, change a device setting of a configurable security setting, or investigate security hygiene in a department.

Of course, above systems should be appropriately integrated to ensure security alerts and information timely reaches the intended recipient for follow-up. The abstractly described systems will be implemented by a combination of SIEM, CTMS, SOAR, UTM, etc. functionalities, tools and services. These systems also enable the practical assignment, handover and escalation between parties jointly responsible for the security, e.g. the medical device manufacturer, the hospital and other security service providers. The systems also enable integration with other monitoring solutions, e.g. for asset discovery and network monitoring, operational technology and industrial control system monitoring, physical security monitoring, etc. [12].

17.4 Security Models

Analytical security models form the core of the medical device security monitoring approach. These models take security relevant data as input and as output create actionable security alerts or information [23].

17.4.1 Monitor Device Operational Environment

Medical devices operate in a hostile environment with threats originating from the networks and physical access. Devices have a view on this and can contribute their perspective on security relevant events, e.g. relating to status of environmental security prerequisites, network threats and (ab)use of the device.

In this category, a variety of experimental, proof-of-concept and prototype security models have been created in context of the SAFECARE project [12, 23].

A first model detects network traffic on SMB ports that should not be there accordingly to deployment guidance for the medical device, e.g. because a firewall is expected between the hospital network and the system. This model addresses the relatively high risk these Windows file sharing ports pose with WannaCry and NotPetaya as examples. The potential value of a security model over a basic firewall rule is that the security model can be optimized for (near) zero false positives. By validating the model against historic data e.g. anomalous but not malicious behavior can be ignored, which may happen for example during servicing. Such model can generate alerts with high confidence, a clear set of possible root causes, and actionable steps for remediation.

In a variant also alerts can be generated for suspicious traffic on these ports, which is then accompanied with a lower confidence indicator. The CTMS or SOC operator can then consider this in combination with signals from other monitoring systems.

Another model alerts on suspicious queries for patient demographic data and files from the medical device to the PACS system, e.g. excessive use of wildcard queries and anomalous query and retrieval patterns. The model can generate an alert with relevant contextual information such as the authenticated user of the device. Such anomaly detection model will be accompanied with lower confidence indicator and the alert should be considered in perspective of other signals.

A third model detects suspicious login events, e.g. anomalous patterns of failed, successful and emergency logins. This model exploits that organizations often have particular workflows around medical devices. Yet, false positives are to be expected as deviations are likely. Therefore, the same arguments hold as above and a lot of historic data and machine learning will be necessary to develop a model with utility, i.e. achieve sufficient detection capability with acceptable false positive ratio.

The takeaway for security monitoring of the operational environment of medical devices is that the concept has been demonstrated. However, models must be matured and validation at scale must be done. Variety in the operational environment make it more complex to make highly reliable models, as is expected for anomaly detection.

17.4.2 Monitor Device Security Control Status

Medical devices come with security controls designed in. However, over the lifetime of the medical device things may happen that render security controls ineffective. From the IT domain it is known that end-point protection solutions eventually degrade or fail over time when left unattended [22]. Customer configurable options and service actions may also affect security controls, e.g. adaptations to firewall configurations that expose the system more than necessary to achieve the intended integration, failed upgrades, etc.

Models in this category range from experimental to mature and validated. One such validated model monitors if antivirus is functioning and utilizing up-to-date virus definition files. The model here is optimized to avoid false positives, considering service workflows and validated configuration for the medical device. For example, after a system reinstall the first virus definition file update may take some time, virus definition updates may be held back due to known incompatibility, etc. The resulting model generates alerts with high confidence, accurate root cause issue description, and concrete service action to remedy the problem by service staff.

Another model monitors host firewall status and configuration. Analytics here distinguishes between good states and high risk non-compliant states. The model exploits that in practice there are patterns in configuration changes for particular purposes and a number of anti-patterns, contributed by subject matter expertise, which can be learned and captured in a model. Alerting on these can timely bring the device back into a complaint state with good security posture, while preserving the functionality intended to enable by the configuration change.

Yet another model monitors software/firmware releases and patch levels. Monitoring here contributes to timeline installation of updates and patches, failed installations, de-installations, installation of unvalidated patches, restoration of images lacking patches and many more exception cases that may leave devices temporarily or for good in an insecure state. The model leverages accurate information on installed software and patches of the medical device as well as baseline data on supported versions for the medical device.

For all models in this category holds that high accuracy, low false positive rates can be achieved. In combination with good actionability to remedy issues, it thereby offers an efficient method to maintain good security posture.

Several of above models have been empirically validated with real-world conditions and data. The findings confirm that monitoring of medical device security control status is feasible and enables effective and efficient remediation. In other words, detection meets accuracy and false positive requirements and produces actionable alerts. It further confirms earlier findings on degradation of end-point

protection solutions when left unattended to not only apply to the IT domain [22], but also to the medical device setting. Furthermore, with the monitoring models in place, such failures are detected and remediated to a point where the problem is practically absent. This leads to the conclusion that security monitoring provides an effective security control to manage and reduce risk related to security controls of medical devices.

17.4.3 Support Security Risk Analysis

Security risk analysis can also benefit from security monitoring. It is a relatively straightforward extension once security monitoring and analytics is in place. It can support security risk analysis in a data-driven, fact-based and qualitative manner.

Models here are derived from the models designed for monitoring the device security control status. Aggregate statistics on the (in)correct functioning of security controls can be fed into the periodic updates of product security risk analysis. Subsequently, additional risk mitigation measures can be taken where needed.

This category also offers potential to collect statistics on the occurrence of threats and attacks and also feed this into the risk management process.

17.5 Conclusion

This chapter presented an approach for security analytics and monitoring for medical devices. The concept is feasible and contributes to the security of medical devices and their operational environment. It addresses the medical device specific challenges and requirements.

Some medical devices today already allow for (some) security monitoring. To grow the monitoring and analytics the logging and log export capability of medical devices must be extended as part of their design. Over time this increases the ability to monitor more medical devices and monitor them to greater extent.

It has been empirically found that security monitoring improves the security posture of monitored systems. Particularly, pro-active monitoring of key security controls such as end-point protection and firewalls contributes to undisturbed functioning. This can be done with high accuracy and low false positives. This enables efficient detection and remediation, e.g. by service engineers certified to service the medical device.

Security monitoring of the operational environment is also promising. However, it is limited by the ability of the medical device to observe events in or from its environment. It can observe how it is used or operated and alert on suspicious behavior. Like anomaly detection this typically does not lead to direct action, but to events that can be considered by security operations staff in combination with

other events. For this purpose, the medical device security monitoring can integrate with e.g. the cyber threat management system of the hospital.

Acknowledgements

The SAFECARE project has received funding from the European Union's H2020 research and innovation programme under Grant Agreement no. 787002.

References

[1] Suzanne Schwartz, *et al.*, The Evolving State of Medical Device Cybersecurity, Biomedical Instrumentation & Technology, Vol. 52, Issue 2, pp. 103–111, AAMI, 2018.
[2] Mandeep Khera, Think Like a Hacker: Insights on the Latest Attack Vectors (and Security Controls) for Medical Device Applications, Journal of Diabetes Science and Technology, Vol. 11, Issue 2, pp. 207–212, SAGE, 2017.
[3] A. J. Burns, M. Johnson, Peter Honeyman, A brief chronology of medical device security, Communications of the ACM, Vol. 59, Issue 10, pp. 66–72, ACM, 2016.
[4] George Tanev, Peyo Tzolov, Rollins Apiafi, A Value Blueprint Approach to Cybersecurity in Networked Medical Devices, Technology Innovation Management Review, Vol. 5, Issue 6, 2015.
[5] ENISA, Smart Hospitals: Security and Resilience for Smart Health Service and Infrastructures, 2016.
[6] Tom Haigh, Carl Landwehr, Building Code for Medical Device Software Security, IEEE, 2015.
[7] Steven Harp, Todd Carpenter, John Hatcliff, A Reference Architecture for Secure Medical Devices, Biomedical Instrumentation & Technology, Vol. 52, Issue 5, pp. 357–365, AAMI, 2018.
[8] FDA, Content of Premarket Submissions for Management of Cybersecurity in Medical Devices – Final Guidance. 2014.
[9] FDA, Content of Premarket Submissions for Management of Cybersecurity in Medical Devices – Draft Guidance, 2018.
[10] FDA, Postmarket Management of Cybersecurity in Medical Devices – Final Guidance, 2016.
[11] IETF, Date and Time on the Internet: Timestamps, RFC3339, July 2020.
[12] John Soldatos, James Philpot and Gabriele Giunta (eds.), Cyber-Physical Threat Intelligence for Critical Infrastructures Security, Ch. 10, Integrated

Cyber-Physical Security Approach for Healthcare Sector, Boston–Delft: Now Publishers, 2020.
[13] G. Zheng, *et al.*, Ideas and Challenges for Securing Wireless Implantable Medical Devices: A Review, IEEE Sensors Journal, Vol. 17, Issue 3, pp. 562–576, 2017.
[14] Access Control Schemes for Implantable Medical Devices: A Survey. Wu, Longfei, *et al.* 2017. 5, s.l.: IEEE, 2017, IEEE Internet of Things Journal, Vol. 4, pp. 1272–1283.
[15] M. Zhang, A. Raghunathan, N. K. Jha, MedMon: securing medical devices through wireless monitoring and anomaly detection, IEEE transactions on biomedical circuits and systems, Vol. 7, Issue 6, 2013.
[16] Nader Sehatbakhsh, *et al.*, Syndrome: Spectral analysis for anomaly detection on medical IoT and embedded devices, IEEE International Symposium on Hardware Oriented Security and Trust, Washington, IEEE, pp. 1–8, 2018.
[17] Julio Navarro, *et al.*, HuMa: A multi-layer framework for threat analysis in a heterogeneous log environment, Foundations and Practice of Security, pp. 144–159, 2017.
[18] Junaid Chaudhry, *et al.*, POStCODE Middleware for Post-Market Surveillance of Medical Devices for Cyber Security in Medical and Healthcare Sector in Australia, 12th International Symposium on Medical Information and Communication Technology (ISMICT), IEEE, 2018.
[19] Daniel B. Kramer, *et al.*, Security and Privacy Qualities of Medical Devices: An Analysis of FDA Postmarket Surveillance, PLoS One, Vol. 7, 2012.
[20] Medical Device Coordination Group Document, Guidance on Cybersecurity for Medical Devices, 2019–16, December 2019.
[21] Medical Device Cybersecurity Working Group, Principles and Practices for Medical Device Cybersecurity, IMDRF/CYBER WG/N60FINAL:2020, 18 March 2020.
[22] Absolute Software, 2019 Endpoint Security Trends Report, 2019.
[23] Brinda Hampiholi, Paul Koster, Specification of e-Health Device Security Analytics, SAFECARE deliverable D5.7, August 2019.
[24] NIST, Cybersecurity Framework, version 1.1, 16 April 2018.
[25] IETF, The Syslog Protocol, RFC 5424, 2009.
[26] IHE, Audit Trail and Node Authentication (ATNA), IHE IT Infrastructure Technical Framework, 2019.
[27] IHE, Consistent Time (CT, IHE IT Infrastructure Technical Framework, 2019.
[28] NEMA, DICOM Security and System Management Profiles, PS3.15 2021a, Appendix A.5.3 DICOM Specific Audit Messages, 2021.