

Toward a safe Integrated Clinical Environment: A communication security perspective

Denis Foo Kune,
Yongdae Kim
{foo,kyd}@cs.umn.edu
University of Minnesota

Krishna
Venkatasubramanian,
Insup Lee
{vkris,lee}@cis.upenn.edu
University of Pennsylvania

Eugene Vasserman
eyv@ksu.edu
Kansas State University

ABSTRACT

With a vision emerging for dynamically composable and interoperable medical devices and information systems, many communication standards have been proposed, and more are in development. However, few include sufficiently comprehensive or flexible security mechanisms to meet current and future safety needs. In this work, we enumerate security requirements for the communication stack of a medical composition framework. With a focus on safety and security, we then survey existing medical and non-medical communication standards and find significant gaps between required properties and those that can be fulfilled even by non-trivial combinations of currently standardized protocols. We conclude that, at the moment, medical interoperability requires a “full-stack” communication solution with designed-in security and extensibility features. This paper is meant to inform future work on building such a comprehensive protocol stack or standardizing protocols and protocol suites that satisfy the security properties needed for safe and secure next-generation device coordination.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General—Data communications, Security and protection; K.6.5 [Computing Milieux]: Security and Protection—Authentication, Unauthorized access; J.3 [Computer Applications]: Life and Medical Sciences—Medical information systems

Keywords

Medical device, integrated clinical environment, security

1. INTRODUCTION

Recent years have brought increased attention to security vulnerabilities in standalone medical devices [12, 17]. The next step is the challenging problem of security and safety

of *interconnected and dynamically composable* medical systems. Various agencies and standards bodies, including the U.S. Food and Drug Administration, have signaled that *the future of medical technology lies in medical device interoperability*, such as a “system of systems” that can integrate information from multiple clinical sources in a context-sensitive way to guide patient care or prevent common critical mistakes [26, 33]. Such systems can reduce the cost of care and ultimately save lives by providing functions such as clinical decision support, inference and early warning, adverse interaction detection, alarm aggregation, and false alarm suppression. While there is a general agreement that security must play a part, few existing standards mention specific security considerations or mechanisms for medical systems [6, 8, 19, 24]. Even when discussed, security standards are incomplete, optional, or both, preventing strong security guarantees even when implementing standards-mandated methods. *Gaps in available standardized security mechanisms* can lead to failures in the safety of resulting systems in the presence of malicious insider or outsider adversaries. The purpose of our work is to: (1) draw attention to this increasingly important problem, (2) describe security requirements for communication in integrated clinical environments, and (3) show the gaps between requirements and features provided by currently standardized protocols.

Interoperability Architecture. In this work, we focus on the ASTM F2761 standard architecture [6] shown in Figure 1, also known as the MD PnP Integrated Clinical Environment (ICE). ICE provides a means for enabling manufacturer-agnostic interoperability between arbitrary medical devices. Logically ICE is separated into the *Supervisor, Network Controller, and devices*, although many components may be implemented by the same physical hardware. Logging and external interfacing, such as off-site patient Electronic Health Records (EHRs), are also supported by dedicated logical components.

Devices perform sensing and/or actuation automatically or on command, i.e. a device may take a blood pressure reading or infuse medication. Coordinating devices may temporarily suppress a high blood pressure alarm if all other patient vital signs are normal and the just-infused medication is known to elevate blood pressure. Currently, devices from different manufacturers cannot communicate except in very limited ways, so even this simple level of coordination is hard to achieve without a standardized interoperability protocol. ICE allows such coordination — each device communicates with the Network Controller, a sort of “medical

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MedCOMM'12, August 13, 2012, Helsinki, Finland.

Copyright 2012 ACM 978-1-4503-1478-7/12/08/12/08 ...\$15.00.

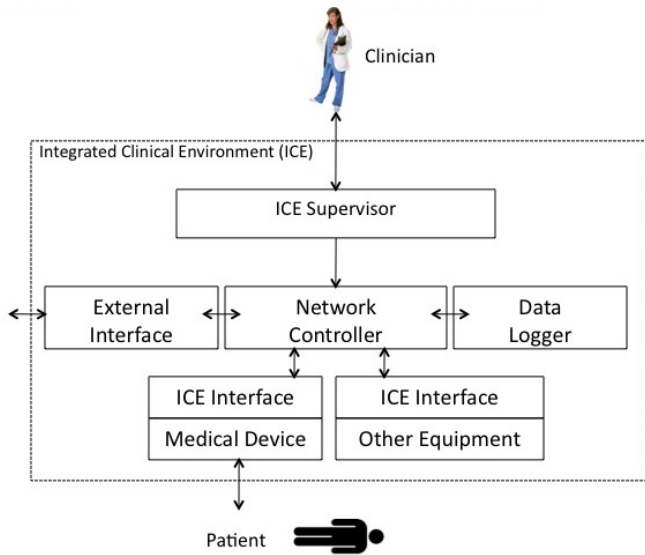


Figure 1: Interoperability architecture of MD PnP ICE

router” which does not have any medical/clinical functionality itself, but is responsible for data routing, translation, and quality of service (QoS) enforcement, facilitating communication between devices and the Supervisor. The Supervisor is responsible for executing “clinical workflows,” from common and easily scriptable tasks such as taking blood pressure at pre-defined intervals and recording the results, to more complex procedures like medication interaction monitoring and suppression of likely false alarms. Each component has different connection security, authentication and authorization, logging, and physical protection requirements, which need to be considered in the overall system architecture.

Threat Model. Because ICE makes medical systems *composable at deployment time*, we do not a priori know the target network topology, communication protocols, or transmission media (e.g. wireless or wired). Thus we assume a strong adversary in clinical care contexts. Adversaries can eavesdrop on all communication and arbitrarily delay, inject, re-order, or forge packets traveling through the network. In addition, we assume that the attacker can be present in the physical vicinity of the patient to carry out an immediate or delayed attack on the devices or the system.¹ Authorized medical staff members are trusted.

2. SECURITY REQUIREMENTS

We consider the design of secure interoperable medical systems by focusing on security requirements specific to each OSI communication layer, as well as cross-layer requirements. We evaluate available communication standards against those requirements while keeping the assumptions about deployed systems to a minimum. Table 1 summarizes these requirements and organizes them by the layer at which they must be addressed. Requirements marked with a “†” span multiple layers. We note that physical tamper resistance/evidence, while important, are not protocol issues and are thus out of scope for this paper. The requirements are:

¹Proximity allows an attacker to physically harm the patient, but we assume that more subtlety or delay is desired to e.g. prevent detection.

Table 1: Requirements statified at OSI layers 2, 4 and 7. Items marked “†” span multiple layers.

Requirement	Layer 2	Layer 4	Layer 7
1 Medium Access	✓		
2 Session Security		✓	✓
3 Data Provenance			✓
4 User Authentication		✓	
5 Data Access Control			✓
6 Logging	✓	✓	✓
7 Alerts [†]	✓	✓	✓
8 Device Management [†]	✓		✓
9 User Management [†]			✓

1. **Secure medium access control:** An attacker with access to the wired or wireless medium should not be able to generate forged layer 2 protocol messages that would be accepted by a receiving interface. Confidentiality may be added as needed.
2. **Secure sessions:** Applications hosted by devices should be able to set up end-to-end secure (confidential, authenticated, and timely) communication channels.
3. **Authenticity of application objects:** Each application should be able to authenticate and determine the trustworthiness of its remote communicating principals. Note that, even if a device is trusted on a network, its applications may not be trusted to generate/access certain data blocks.
4. **User authentication:** The system should ensure that the medical staff and patients are properly identified before granting the appropriate level of access.
5. **Access control of application data:** The system should provide granular access control to application data blocks to enable clinicians and patients to retain control as well as record access to those data blocks.
6. **Timely and secure logs:** Security events may generate logs, depending on the applicable policy. Those logs should be timestamped and transmitted to a central repository in a timely fashion to enable both reconstruction of past events and estimation of likelihood of future events. Logging is usually done at the application layer only. Lower layer “log-worthy” behaviors are forwarded up the stack to be application layer, which knows where and how to record the information. Once generated, logs should be immutable and maintain accountability for log access.
7. **Alerts for unexpected behavior** († cross-layer): The system should support the generation and delivery of alerts based on local policy. Alerts generated at each layer of the communication stack may be reported up to the next layer, or directly to the application layer for logging and/or user notification.
8. **Device provisioning and management** († cross-layer): The system should support binding of devices to a facility-local trusted keying infrastructure to track their lifecycle, ensuring revocation when required.
9. **User management** († cross-layer): The system should support time-aware role-based access controls for clinicians and patients, giving appropriate access to data blocks at the appropriate time.

3. SURVEY OF EXISTING PROTOCOLS

We focus on single-patient systems (each patient has a personal coordination infrastructure) with devices connected in a star topology with the coordinator, such as the ASTM F2761 network controller [6], in the center. At layer 2, interface pairing only provides pairwise single-hop protection, but data that is transferred over a shared hospital network remains unprotected.² We thus leverage the properties of layer 4 protocols for end-to-end protection of multi-hop communication. Finally, we use layer 7 protocols (which amalgamate layer 5 and 6 functionalities as well) to ensure data-level security granularity between components which may only be partially trusted.

In this section, we summarize existing healthcare-specific communication standards, as well as existing layer 2, 4, and 7 protocols specified in those standards, and compare them with our requirements. We begin with a short description ISO 11073 and the Continua standards, discuss the specified layer 2 and 4 protocols, and then examine layer 7 standards, including IHE and HL7.

3.1 Cross-stack/full-stack protocols

ISO/IEEE 11073. The ISO/IEEE 11073 family of standards is composed of 4 main groups, namely Device Data, Application Services, Internetworking, and Transport, numbered by group. The document *00101-2008, wireless guidelines* [18] refers to sections typically covered in parts 305xx, including parts of mobile cellular networks, wireless broadband, WLAN, and WPAN. The document looks at 2 main categories for security — data and network; physical security was considered out of scope.

The *Data security* recommendations are limited to the HIPAA requirements where system integrators and operators are ultimately responsible for the risk analysis and choice in the appropriate security mechanisms. We note that some protocols mentioned in this document do not provide adequate building blocks due to weak or broken schemes, e.g., [30]. The document also recommends the use of encryption only after the patient identifiers have been included, which might leave the data vulnerable at prior stages. A follow-up recommendation suggests avoiding security mechanisms between the sensor and the amplifier due to concerns of burdening the processor with cryptographic calculations. The document focuses on encryption techniques and provides very few details on message integrity, thus only partially addressing requirements **3**, **2**.

For *network security*, the document focuses on 3 pieces: authentication (presumably of users, which would address requirement **4**), encryption, and firewalls. It mentions 802.1x protocols for authentication (requirement **3**) and AES (in modes of operation specified in current 802.11-series wireless protocols) for encryption and integrity addressing requirement **1** if the appropriate modes are chosen. This section also discusses malware, but doesn't make recommendations short of referring to the FDA's cybersecurity efforts. It goes through denial of service (DoS) attacks and mentions intrusion detection and prevention as mitigation mechanisms, which could address **7** requirement, though it is not clear. Substitution attacks are looked at as a network security issue, and recommends the use of message au-

²The proposed model might not be a direct fit to existing clinical networks. We thus consider mechanisms that are independent of lower layers to allow a smooth migration path.

thentication/integrity codes such as AES in CCMP mode of operation, thus partially addressing requirements **1** and **2**.

30x series (transport) documents including *30200 (cabled)*, *30300 (infrared)*, *30400 (Inter-LAN)* mention little in regard to security, perhaps due to having no built-in security mechanisms specified. At the time of writing, a document for *305xx (Wireless)* has not yet been issued, but may be more promising due to security mechanisms already built into the wireless protocols considered.

*20x series (Internetworking)*³ includes a security section that appears to still be in draft, and no publicly circulated copy was available at the time of writing.

The 11073 group of standards though spanning all the layers of OSI stack, only appears to partially satisfy a subset of the requirements.

Continua. The reference architecture [9] mentions the Bluetooth Health Device Profile [7] for the wireless interface, the USB Personal Healthcare Devices [11] protocol for the wired interface, and the IEEE 11073 Personal Health Device standard [8] for the application data format. Although security is identified as a technical issue, it is not yet clearly addressed in the Continua effort and is therefore not evaluated in this work.

3.2 Layer 2 and 4 protocols

Since most medical standards either mandate or recommend certain data link and transport protocols, it is useful to summarize the available standards, in this regard, here.

Wired. ISO/IEEE 11073 has specifications for cabled serial connections, e.g., RS-232 [1], Ethernet (802.3 family [10]), USB [34] and Firewire [21]. Part 30200 [22] specifies a transport profile for cabled connections. It defines the physical layer, but inherits the upper layers from IrDA [23]. Thus, there are no security mechanisms built into the profile. At the physical or data link layer the document appears to assume that data security is provided if physical security is achieved. Similarly, USB [34] and FireWire [21] do not specify security mechanisms. They assume physical security and rely on the upper layers to provide any of those services. For example, secure USB storage relies on the upper layer for encryption and authentication, and physical countermeasures for key storage. The IEEE 802.3 family of protocols [10] as well do not specify any security mechanisms, relying on physical security instead. Thus, under our threat model, with an attacker having access to the medium, the protocols specified above were not designed to and therefore do not meet any of the requirements from section 2.

Wireless. Due to space constraints, we will focus on the most widely deployed protocols mentioned in the P11073-00101 document [18], namely cellular networks, 802.11 [20], and 802.15 [2, 35] families. Cellular networks can be grouped by generation: **2G** includes GPRS and EDGE on GSM networks, **3G** has UMTS and CDMA2000, and **4G** incorporates LTE and WiMax. All have varying levels of security, but 2G networks cannot be considered secure [30]. The schemes of UMTS [4] and LTE [3] have not yet been reported to have significant vulnerabilities. In wireless local area networks (WLAN) we consider the 802.11 [20] family, which is widely deployed. With the current use of WPA2 and 802.1x authentication, the 802.11 protocols appear to address requirements **1** with WPA2 and **3** with 802.1x. Wireless personal area networks (WPAN), Bluetooth, and

³ISO/IEEE 11073-20500 Security - Framework and overview

Table 2: Summary of Layer 7 standards addressing our security requirements. “*” denotes partial fulfillment.

Requirement	IHE	HL7	ICE	11073
2 Session Security	Yes*	–	–	Yes*
3 Data Provenance	–	–	–	Yes*
4 User Authentication	Yes*	–	–	Yes
5 Data Access Control	Yes*	Yes	–	–
6 Logging	Yes	Yes	–	–
7 Alerts	–	–	–	–
8 Device Management	–	–	Yes	Yes*
9 User Management	Yes*	–	Yes	–

802.15.4 [35] provide security mechanisms including device authentication, message encryption, and integrity, thus addressing requirements **1**, **3**. However, the key exchange and interface pairings must be controlled by the upper layers, parts of which are not explicitly mandated by 802.15.4, leaving requirement **8** incomplete.

Due to the nature of radio frequency communication, disruptions in the medium (e.g. jamming) is almost always possible. A large body of work exists on jamming detection, avoidance and resistance, e.g. [36, 37], but these mechanisms are not explicitly mandated in current standards, and so we do not consider them in our evaluation.

Transport. At the transport layer, the two main secure protocols we consider are TLS v1.2 [14] for streams and DTLS v1.2 [15] for datagram-based protocols. The architecture for key/certificate distribution and update is not addressed at this layer and is instead considered at the application layer.

TLS v1.2 provides unidirectional or mutual authentication for secure transport sessions, allowing devices to authenticate in an end-to-end session if they both have certificates signed by a trusted entity, addressing requirements **3**, **2** and **4** but leaving requirement **8** to the implementer. It supports cryptographic algorithms known to be secure at the time of writing, including support for signing and stream encryption. As far as we are aware, TLS sessions are considered secure and can provide confidential authenticated end-to-end L4 (OSI transport) channels between devices as long as it is supported by good certificate management. *DTLS v1.2* is the datagram counterpart of stream-focused TLS and would similarly address requirements **3**, **2** and **4**. As of the writing of this document, some successful attacks on implementations of DTLS v1.2 [5] have been disclosed, and have been accounted in our requirement assessment Table 3.

3.3 Layer 7 (application) protocols

At layer 7 we consider two existing standards for medical device interoperability; Integrating the Health Enterprise (IHE) [24], a healthcare industry consortium that publishes standards to improve the way computer systems in healthcare share information, and Health-Level 7 (HL7) [19]. Table 2 summarizes the extent to which these protocols satisfy the requirements specified in Section 2.

3.3.1 Integrating the Health Enterprise Standards

IHE defines a large number of profiles that describe solutions to specific interoperability issues among medical devices. The current profile list covers a wide range of issues from Anatomic Pathology to Radiology, but only two of the eleven profiles deal directly with security issues when medical devices interoperate, addressing issues at and above the transport layer.

The **Audit Trail and Node Authentication (ATNA)** profile establishes security measures for patient confidentiality, data integrity, and caregiver accountability [25]. It specifies access control, security audit logging and secure inter-device communication. The profile defines the notion of a *Secure Node* (SN) which enables secure interaction with other nodes and uses access control mechanisms in conjunction with user authentication to secure user-to-node interaction. The SN shares the most similarities with the Network Controller in ICE [6]. All aspects of the SN device are assumed to be secure including its filesystem and OS. All Secure Nodes interacting with one another are collectively called a *Secure Domain* (SD), which can be established at the hospital or departmental level, or at some other level of granularity. All machines within this Secure Domain are assumed to be “host-authenticated”, i.e. known to the operating facility. The ATNA profile has two requirements — node authentication and auditing.

The authentication aspect has two parts, the first requiring node-user interaction authentication, and the second authenticating inter-node interaction. The node-user authentication in ATNA uses an access control mechanism (though details are left to the implementer) to determine the level of access the user gets to various applications on the node based on the user credentials, thus only partially addressing requirements **4**, **9**, and **5**. Inter-node authentication is certificate-based, requiring mutual authentication, and proposes the use of TLS (assuming v1.2) for end-to-end secure channels between the nodes. The generation and maintenance of the certificates for individual devices is not specified. Finally, the profile does not mandate confidentiality/encryption and is more focused on ensuring the integrity of the channel. However, as it makes confidentiality optional, it only partially addresses requirement **2**.

The *Logging and Audit Trail* aspect of ATNA ensures that all security-related events are logged by the SNs. This is usually done in a centralized repository recording events including accesses to a patient’s personal health information (PHI), the user performing the access, and node or user authentication failures.

The ATNA profile mandates the use of the DICOM vocabulary [13] for auditing purposes, extended by RFC 3881 [32]. These standards provide the data definitions for reporting security and privacy events.⁴ The audit messages (or audits) are sent to the repository for storage using the standard Syslog protocol defined in RFCs 5424[16] and 5426 (Syslog over UDP) [31]. Note that so far confidentiality or integrity mechanisms have not been specified. To alleviate these problems, an alternative has been proposed in RFC 5425 [27], that sends Syslog messages over TLS, addressing requirement **6**. During normal operation, every user login attempt to an SN generates an audit event for both successful and failed actions. IHE also defines an ATNA Radiology-option which is an extension of the profile for radiology purposes. Its requirements mirror those of the base ATNA profile except it mandates that communication between SNs be encrypted, given the sensitivity of the radiology information.

The **Enterprise User Authentication (EUA)** profile has two main tasks — (1) Provide centralized authentication management for users thereby enabling single sign-on

⁴Note that the auditing framework assumes that all the devices and systems are time-synchronized and have the correct timestamp for every event record.

Table 3: Summary of protocols which address various requirements at different layers. “ Δ ” and “*” denote optional and partial addressing, respectively.

Requirement	Layer 2	Layer 4	Layer 7
1 Medium Access	11073*, 802.15.4, 802.11i-2004		
2 Session Security		TLS, DTLS	11073*, IHE Δ
3 Data Provenance	802.15.4, 802.11i-2004	TLS, DTLS*	11073*, IHE*
4 User Authentication		TLS, DTLS*	11073, IHE
5 Data Access Control			HL7
6 Logging			HL7 Δ , IHE Δ
7 Alerts	–	11073*	–
8 Device Management	11073*, 802.15.4 Δ	–	11073*, ICE*, IHE*
9 User Management			IHE*

over the healthcare enterprise (2) Seamlessly allow a users’ context to be transferred between applications on a single machine [25] using a users’ authentication credentials, thus enabling interoperability between applications in an authenticated manner. Authentication in EUA is done using Kerberos, the centralized key distribution center-based scheme that provides a user with a ticket — a temporary key for a user and a service to communicate in a secure session [28]. One of the services of the Kerberos system is an authenticated user access to the Context Manager (CM) in the machine to which the user is trying to log on. The CM and different client applications use the specifications of the Clinical Context Object Workgroup (CCOW) [19] to provide seamless movement of a user’s context between applications on a single machine. The EUA profile is primarily used to improve the effectiveness of addressing requirement **4**. Therefore overall, IHE profiles comes pretty close to addressing the security requirements at layer 7, though they are somewhat deficient in terms of sessions and completely in terms of alerts requirements

3.3.2 Health Level 7

HL7 provides a framework for exchange, management, and integration of electronic health information to support clinical practice and management of healthcare delivery services. Interoperability in HL7 is supported by standardizing at five levels of standards abstraction: conceptual (e.g. RIM), document (e.g. CDA), messaging (e.g. HL7 v2.x and HL7 v3), application (e.g. CCOW), and service (e.g. Arden). Most of the discussion below focuses on HL7 v3 which contains more details of the security specification than HL7 v2[19]. Security in HL7 is defined in v3 as a service standard in the form of *Privacy, Access and Security Service (PASS)*. The focus of these standards is from an information standpoint and not from that of individual medical devices.

Privacy, Access and Security Service (PASS) defines a set of loosely-coupled service components that enable confidentiality and integrity of healthcare information. The PASS-Audit service describes, at a conceptual level, the requirements that relate to the functional behavior of auditing in a healthcare environment. The service provides two capabilities that would address requirement **6**: (1) Audit submission in response to events generated by Audit Event Sources, and (2) Retrieval of audit records with respect to access of personal health information. Further, it specifies that the audit service must have the ability to validate any requests that can be submitted and it must establish a secure communication channel with the querying entity. Audits (events) can be generated by users, information systems or devices. The model used is a generalization of the one used

in DICOM which is based on RFC 3881 as referenced in the ATNA profile described above.

The PASS-Access Control service presents functionalities required for access to resources in a distributed healthcare setting. The document also specifies the lifecycle of the policies involved in access control. Both these are currently in the form of unconstrained conceptual specification and do not provide any implementation details [19] The access control system is responsible for generating audit record based on security relevant information, which would address requirements **5** and **6**. In general terms, the access control system suggested by HL7 is *Role-based Access Control*. In this regard, HL7 does not provide any specific list of roles or permissions, which are left to the implementers. HL7 only provides a framework for role engineering, using scenario-based approaches as described in [29].

Table 2 illustrates the requirements satisfied by these standards, the ISO/IEEE 11073 family, and the ICE standard.

4. DISCUSSION

It is clear that currently available standards do not cover all of our requirements. (See Table 3.) With our threat model that assumes physical access to both the wired and wireless medium of the clinical environment, we find that some wireless protocols may have an advantage with built-in security mechanisms. A VPN could address **1** by providing defenses on the virtual medium and preventing altered packets protected within the VPN, but it would add an additional layer. Additionally, all nodes will need to support the particular version of deployed VPN, and appropriate certificates/keys need to be deployed to the devices to know that they are trusted. Even with proper mechanisms to defend the medium, we still need mechanisms to secure end-to-end session to carry the application payload. In this case, the Layer 4 requirements lead us to conclude that TLS or DTLS are appropriate, if adequately supported by a certificate distribution and update mechanisms.

At the application layer, the focus on dynamic composability makes it difficult to consider security in detail. Both IHE and HL7 pay some attention to security, but the mechanisms suggested are only partial solutions. IEEE 11073 defines vertical profiles through the communication stack with different data-link components, each with different security properties. Due in part to those options for transport layers and a lack of specification of how application-level user access controls feed into security at the lower layers of the stack, IEEE 11073 does not seem to fulfill all the requirements we extracted. We also note that none of the standards surveyed account for a mechanism to do secure time synchronization that would be critical to address requirements

7 and 6. This is the reason in Table 3, we mark the IHE and HL7 satisfaction of requirement 6 as partial.

Addressing all proposed requirements would be a formidable task. However, we were surprised at the large gaps that exist in the mapping between the requirement set and the mechanisms specified in the standards surveyed. It suggests that a lot more attention should be paid to developing secure protocols as those standards evolve.

Acknowledgments

This work was supported in part by the National Science Foundation NSF award CPS-1035715 and the NIH/NIBIB Quantum program grant NIH-1U01EB012470.

References

- [1] Interface between data terminal equipment and data circuit terminating equipment employing serial binary data interchange. TIA/EIA-232-F, 1997.
- [2] Wireless medium access control (MAC) and physical layer (PHY) specifications for wireless personal area networks (WPANs). *IEEE P802.15.1/D6*, 2005.
- [3] 3GPP TS 36.201 V10.0.0 — LTE physical layer; General description (Release 10), 2010.
- [4] 3GPP TS 36.331 V10.3.0 — Evolved Universal Terrestrial Radio Access (E-UTRA); Radio Resource Control (RRC); Protocol Specification, 2011.
- [5] N. AlFardan and K. Paterson. Plaintext-recovery attacks against datagram TLS. In *NDSS*, 2012.
- [6] ASTM F-29.21. Medical devices and medical systems — essential safety requirements for equipment comprising the patient-centric integrated clinical environment (ICE), 2009.
- [7] Bluetooth, SIG. Health device profile v1.0. 2008.
- [8] M. Clarke, D. Bogia, K. Hassing, L. Steubesand, T. Chan, and D. Ayyagari. Developing a standard for personal health devices based on 11073. In *EMBS*, 2007.
- [9] Continua health alliance. <http://www.continuaalliance.org/>.
- [10] CSMA/CD access method and physical layer specifications. *IEEE 802.3-2005*, 5, 2005. (Revision of 802.3-2002).
- [11] Definition, U.S.B.D.C. USB personal healthcare device profile. *V1.0*, available at <http://www.usb.org/developers/devclassdocs/Personal-Healthcare-1.zip>, 2007.
- [12] T. Denning, Y. Matsuoka, and T. Kohno. Neurosecurity: Security and privacy for neural devices. *Neurosurgical Focus*, 27(1), 2009.
- [13] The DICOM standard. <http://medical.nema.org/standard.html>.
- [14] T. Dierks and E. Rescorla. The transport layer security (TLS) protocol version 1.2. RFC 5246, 2008.
- [15] J. Fischl, H. Tschofenig, and E. Rescorla. Framework for establishing a secure real-time transport protocol (SRTP) security context using datagram transport layer security (DTLS). RFC 5763, 2010.
- [16] R. Gerhards. The syslog protocol. RFC 5424, 2009.
- [17] S. Hanna, R. Rolles, A. Molina-Markham, P. Poosankam, K. Fu, and D. Song. Take two software updates and see me in the morning: The case for software security evaluations of medical devices. In *HealthSec*, 2011.
- [18] Health informatics — point-of-care medical device communication — technical report — guidelines for the use of RF wireless technology. *IEEE Unapproved Draft P11073-00101/D03*, 2007.
- [19] Health level seven international. <http://www.hl7.org/>.
- [20] IEEE. Wireless LAN medium access control (MAC) and physical layer (PHY) specifications, 1997.
- [21] IEEE standard for a high-performance serial bus. *IEEE 1394-2008*, 2008.
- [22] IEEE standard for health informatics — point-of-care medical device communication — part 30200: Transport profile — cable connected. *ISO/IEEE 11073-30200:2004(E)*, 2004.
- [23] IEEE standard for health informatics — point-of-care medical device communication — part 30300: Transport profile — infrared wireless. *ISO/IEEE 11073-30300:2004(E)*, 2004.
- [24] Integrating the healthcare enterprise. <http://www.ihe.net/>.
- [25] IHE Integration Profiles. Technical Report, 2011.
- [26] K. Lesh, S. Weininger, J. Goldman, B. Wilson, and G. Himes. Medical device interoperability-assessing the environment. In *HCMDSS-MDPnP*, 2007.
- [27] F. Miao, Y. Ma, and J. Salowey. Transport layer security (TLS) transport mapping for syslog. RFC 5425, 2009.
- [28] C. Neuman, S. Hartman, and K. Raeburn. The Kerberos network authentication service (V5). RFC 4120, 2009.
- [29] G. Neumann and M. Strembeck. A scenario-driven role engineering process for functional rbac roles. In *SACMAT*, 2002.
- [30] K. Nohl. Wideband GSM sniffing. <http://events.ccc.de/congress/2010/>, 2010.
- [31] A. Okmianski. Transmission of syslog messages over UDP. RFC 5426, 2009.
- [32] G. L. Simmons. Security audit and access accountability message XML data definitions for healthcare applications. RFC 3881, 2004.
- [33] D. Tillman and L. Kessler. Medical device interoperability — assessing the environment. In *HCMDSS-MDPnP*, 2007.
- [34] USB 2.0 specification, 2000.
- [35] Wireless medium access control (MAC) and physical layer (PHY) specifications for low rate wireless personal area networks (LR-WPANs). *ANSI/IEEE*, 802(4), 2003.
- [36] A. Wood and J. Stankovic. Denial of service in sensor networks. *Computer*, 35(10), 2002.
- [37] W. Xu, W. Trappe, Y. Zhang, and T. Wood. The feasibility of launching and detecting jamming attacks in wireless networks. In *MobiHoc*, 2005.