A Security Architecture for Health Information Networks

Rajashekar Kailar, PhD, Business Networks International Inc,
Vinod Muralidhar, CSC Consulting

Abstract

Health information network security needs to balance exacting security controls with practicality, and ease of implementation in today's healthcare enterprise. Recent work on 'nationwide health information network' architectures has sought to share highly confidential data over insecure networks such as the Internet. Using basic patterns of health network data flow and trust models to support secure communication between network nodes, we abstract network security requirements to a core set to enable secure inter-network data sharing. We propose a minimum set of security controls that can be implemented without needing major new technologies, but yet realize network security and privacy goals of confidentiality, integrity and availability. This framework combines a set of technology mechanisms with environmental controls, and is shown to be sufficient to counter commonly encountered network security threats adequately.

Introduction

Widespread adoption of interoperable electronic health record (EHR) systems are expected to improve the quality of patient care by facilitating clinician access to accurate health information, reduced health care costs through better utilization of resources, and reduction of medical errors by automated detection of mistakes in human-entered data. Networks of healthcare organizations, enabled by interoperable EHR systems, or *health information networks*, are the means to realize these benefits on a large scale. Ensuring data security and privacy is a significant challenge to information exchange between disparate enterprises, especially when connecting over public networks such as the Internet.

Health information network architectures need to find the right balance between stringent security controls and ease of implementation. While it is crucial that data communications be highly secure and mindful of patient privacy, a pragmatic approach to information security would dictate that the technology requirements not be so complex as to inhibit growth of the network. A practical goal would be to extend the patient privacy protection of the Health Insurance Portability and Accountability Act of 1996 (HIPAA) and other regulations to the network, without burdening the administrative and technical staff of healthcare enterprises with complex implementation and operational requirements.

Based on a recent project to develop a nationwide health information network (NHIN) architecture prototype [1], this paper sets out the goals of health information network security architectures in the context of the salient services that a health information network supports. We develop a generalized data exchange model, discuss the trust model options to support secure data flows, and show the advantages of a federated trust model for large networks. The network security problem is thereby decomposed to ensuring pair-wise security between nodes, and applying policies to enforce transitive trust. Many of the components of the network security architecture are implemented using infrastructure commonly used by healthcare enterprises to connect to the Internet.

Security Goals

The following security goals are proposed for health information networks, as derived from the NHIN project [1], HIPAA security and privacy rules [12], and the Connecting for Health Common Framework [1,5,6]:

1. Protect patient data privacy by empowering individuals to control access to their own healthcare information
2. Allow only fully authenticated and specifically authorized individuals access to data
3. Preserve integrity of data sent over the network
4. Hold users and organizations accountable for all actions on the network
5. Hold each node (organization) in a network accountable for the security of the data in its custody
6. Enable the formation of larger scale networks by securely linking together health information networks

Network Data Flow Models

In a general healthcare data exchange scenario a data provider shares its data with a data consumer either directly or through an intermediary. A consumer may initiate a request for data or have data pushed to it by

the provider. The generalization can be represented graphically as shown in Figure 1:

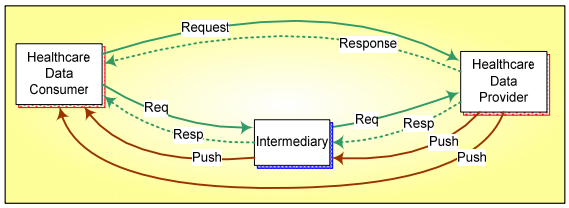

Figure 1: Data Flow between Network Nodes

Data communication between consumer and provider falls into one of the following message exchange patterns:

- *Request/response*: Data consumer sends a request message to data provider, which responds with a message containing the information that is being sought
- *One-way*: Data provider sends data to the consumer when a consumer notification event occurs. The consumer may have previously expressed interest in the data by *subscribing* to a *publish* service. Data may also be *broadcast* to multiple consumers following this pattern.

For the purpose of this analysis, both of these messaging operations have similar security requirements and constraints.

Inter-Domain Data Flows

Secure data exchange over network nodes is relatively simple if the provider, consumer and intermediary are in the same *trust domain*, where a trust domain is defined as a set of network nodes where the identity and privileges of users and devices at one node are shared with all the others. However, managing shared identities over a network tends to be impractical at large scales. Particularly, in a network of networks, the data provider is often in one trust domain, and the consumer is in another. When the data provider and data consumer do not share a direct trust relationship, they may rely on trusted intermediaries to act as their brokers as shown in Figure 2.

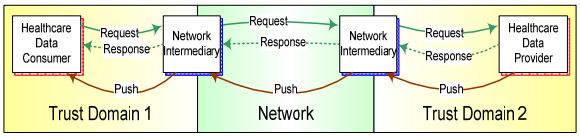

Figure 2: Inter-domain data flow

This use of trusted network intermediaries as secure messaging brokers across networks is an application of *transitive trust*, which, we propose, forms the basis of security architecture for a network of networks.

Trust Models

Network trust can be centralized, distributed or federated. A centralized model offers the advantage that the number of trust relationships to manage is of the order of n (where n is the number of nodes in the network), but creates a single point of failure and performance bottleneck at the central trust anchor. A purely distributed model requires peer-to-peer trust relationships. The number of relationships is of the order of n^2, which is challenging to establish and maintain as the network grows. The federated model supports peer-to-peer trust relationships between domains, each of which may have centralized or peer-to-peer trust. The topology of such a model is shown in Figure 3 where three networks, each a separate trust domain, connect to one another.

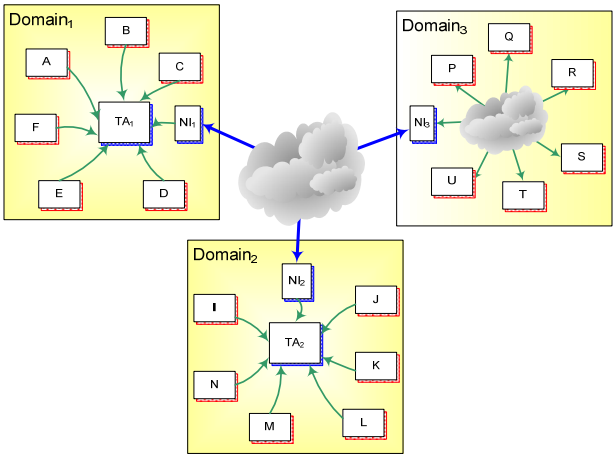

Figure 3: Federated Trust Model supports network of networks

$Domain_1$ nodes A-F use the services of a trust anchor node: TA_1, $domain_2$ nodes I-N use trust anchor TA_2, and $domain_3$ nodes P-U have a peer-to-peer trust relationship for local communication, but use NI_3 as a trust anchor for inter-domain communications. There is a peer-to-peer trust relationship between the three intermediary nodes (NI_1, NI_2 and NI_3). Although the number of peer-to-peer relationships needed for inter-domain communication security is $O[n^2]$, n is now the number of domains, and not the total number of nodes in all domains.

Considerable work has been done on identity federation and propagation of trust in federated networks [11]. However, these typically call for a common network-level identity provider, which is not a practical option on a national scale at this point. Secure communication between federated nodes across domains in our model is based on transitive trust between domain nodes and network intermediaries. Transitive trust has the advantage that it is simpler to scale to very large networks. A major disadvantage of the transitive trust model is that a breach of security at a node can propagate to other nodes and compromise security over a large segment of the network. Exposure to such breaches can be mitigated by:

- Applying least privilege principles, by limiting the trust-chain length [8, 9].
- Applying trusted path principles to ensure that there is path evidence for all transactions (e.g., path between A and S is A⇔TA_1⇔NI_1⇔ NI_3⇔S). Each node retains a copy of the transaction; the parties involved can be held accountable.
- Enforcing trust in real world by establishing policies and agreements to which all partners in a transitive trust network are legally bound.
- Applying least privilege principles to limit access to data and resources requested by remote users/processes using role based access control, with standardized role definitions across nodes and, where possible, across trust domains.

The federated model is recursive, offers higher extensibility and scalability, and provides the rationale for a nationwide network to be architected as a network of networks.

Network Security Requirements

We aggregated the security requirements for a set of representative health information network use cases [2, 3, 4], to develop a core set of security requirements, as listed in Table 1.

ID	Security Requirement
R1	Only authenticated and authorized systems shall be targets of network queries
R2	Only authenticated and authorized users and systems shall request for data over the network
R3	Data integrity shall be preserved across all communication processes within nodes and over the network
R4	Data confidentiality shall be protected over the network
R5	All access to healthcare data shall be traceable to an individual and organization
R6	Where applicable, patients shall be able to specify who can access their data, and such rules shall be enforced at all nodes
R7	Requests originating in a different trust domain shall be authenticated and authorized
R8	Data and system integrity shall be preserved at each node in the network

Table 1: Security Requirements

Security Controls

The application and network security mechanisms (also referred to as technical security controls in NIST 800-53 [10]) that are needed to satisfy the identified security requirements are listed in Table 2.

ID	Security Mechanism	Mapping to Requirements
M1	User Identity Management	R2
M2	User Authentication	R2, R4
M3	User Authorization	R2,R4,R6
M4	Auditing	R5
M5	Anonymization	R4
M6	Secure Messaging	R1,R2,R3,R4
M7	Consent Management	R6
M8	Inter-domain Security	R7
M9	System Availability and Integrity Protection	R8

Table 2: Security Mechanisms

Mechanisms addressing authentication and authorization (**M1**, **M2** and **M3**) are typically implemented in applications at the network edge. For example, each organization maintains its own user identity directory, authenticates users, and enforces security policies based on user roles and their need to access specific data (e.g., patient registration, physician's notes, etc.). The health information network leverages these edge application mechanisms rather than replicating them. The other mechanisms are necessitated more directly by network data sharing requirements, and are implemented by the health information network. The implementations may be centralized within a domain or distributed to network intermediaries.

M4 Audit: All messages that flow through network nodes shall be logged for audit and analysis. Logs also provide path evidence for transitive trust enforcement and support non-repudiation. In health information networks these logs contain sensitive patient data and must therefore be carefully partitioned to separate message level metadata (that may be queried for audits) and clinical data. Clinical data logs are invariably encrypted to protect from access by unauthorized individuals including network and database administrators.

M5 Anonymization: Anonymization services de-identify patient data for aggregation and reporting to secondary use systems such as public health or research, while maintaining the ability to re-identify the patient if required subsequently, e.g. for authorized public health investigations.

M6 Secure Messaging: Network nodes use standardized secure messaging protocols to send and receive data between each other over the Internet. Messaging transactions are conducted from computer to computer. A computer authenticates itself either 1) using user credentials (e.g., by presenting a username

and password or a certificate issued to the user), or 2) using organizational credentials (e.g., by presenting a certificate that has been issued to an organization). Certificates are the basis for secure communication over public networks and offer a practical way to implement secure messaging for health information networks using the Internet. Certificates-based two-way SSL provides confidentiality, integrity and mutual end point authentication. Higher protection levels can be achieved with digital signatures and message level encryption. However, these technologies have not yet reached the level of standardization or widespread use that SSL/TLS has. Widely interoperable secure messaging mechanisms are critical to a nationwide network operating securely as a 'network of networks', and setting the technology bar unrealistically high can be a significant barrier to progress.

M7 Consent Management: The ability of patients to specify who may access their healthcare data requires mechanisms to communicate consent information over the network. Consent may be specified as access control policies, which are transmitted along with the data on a network-wide basis. All data consumers must implement mechanisms to verifiably enforce these policies. Standard notations exist to specify these policies, of which the eXtended Access Control Markup Language (XACML) is currently the industry leading standard [7]. It should be recognized that current levels of healthcare data standardization will likely support only large-grained access control policy specification.

M8 Inter-Domain Access Control: Requests for data sent from one network to another need to include additional attributes that identify the originating domain with credentials, and the user identity, role and affiliation, so that the message recipient (or its proxy) can make access control decisions. In practice, the user identity is not used for authentication or access control decisions, but logged for audit purposes. This mechanism enforces the inter-domain authentication and authorization policies of each organization or domain.

M9 System Availability and Integrity Protection: All production environments need to have mechanisms in place to ensure that the system integrity is protected and that denial of service attacks are countered. These include mechanisms like anti-virus software, server hardening configurations, and safeguards like network firewalls, application firewalls and intrusion detection systems. Additionally, monitoring and management of security resources like certificates and key stores are essential to the reliable operation of a large network.

Environmental Assumptions

Certain threats that may undermine the security of the system can be assumed to be countered based on environment specific assumptions (also referred to as operational security controls in NIST 800-53 [10]). Table 3 lists a candidate set of assumptions that may be made about the health information network.

ID	Assumption	Justification
A1	Intermediary organizations that route messages are trusted not to disclose sensitive data	Legally binding agreements
A2	Care delivery personnel are trusted not to disclose patient data	Legally binding agreements; Doctor-patient relationship
A3	Organizations managing data repositories are trusted not to abuse data that is stored	Legally binding agreements

Table 3: Environmental Assumptions

As shown, these assumptions are backed up with legally binding agreements, such as the HIPAA business associate agreement. The agreements also specify remedial and punitive measures to be taken in event of a breach of data security or disclosure of confidential information. Connecting for Health Common Framework provides sample data sharing agreements [5] and recommended policies to cover breaches of confidentiality [6].

Threat Analysis

We test the sufficiency of the set of security mechanisms identified by assessing whether it provides counter-measures against a generic list of threats as shown in Table 4.

ID	Threat	Countermeasure / Mapping to Mechanism or Assumption
T1	Un-authorized user/system produces data	Identification, and authentication of data producer and intermediary / M1, M2
T2	Un-authorized user/system consumes data	Identification, authentication, and access control checks of consumer and intermediary / M1, M2, M5, M6, M7, M8
T3	Data integrity compromised at producer, consumer or intermediary level	Network, operating system, application, and database level integrity protections and access controls at each node / M1, M2, M9
T4	Data integrity compromised over network	Integrity protection (e.g., MD5 hash, checksums) / M6, A1

ID	Threat	Countermeasure / Mapping to Mechanism or Assumption
T5	Data confidentiality compromised over network	Encryption over network (e.g. SSL) / M6, A1, M7
T6	Information compromised at data provider, consumer, or intermediary by valid user	Audit, organizational binding/responsibility / M4, A1, A2, A3
T7	Virus, Spyware	Anti-virus, Intrusion detection systems, firewalls. / M6, M9
T8	Denial of service	Intrusion detection systems, firewalls; application level counter measures should be used after transport level ones (e.g. two-way SSL) / M6
T9	Identity spoofing	Client certificate based authentication (two-way SSL) / M1, M6

Table 4: Threats and Countermeasures

Similar threat analyses may be applied to specific environments with specific sets of threats mapped to mechanisms and assumptions. A detailed threat analysis was conducted in [1], which helped validate the set of minimum security mechanisms for the generic health information network.

Implementation

Health information networks may meet essential security and privacy requirements by implementing the nine mechanisms listed above. These are implemented in application services that support the messaging infrastructure (e.g. the network intermediary service). While interoperable message-level security standards (e.g. WS-Security), have few mature implementations today, there are widely used transport level security protocols (SSL/TLS) that adequately address secure messaging requirements, supported by appropriate environmental controls. Such a strategy was used in the health information network prototype implementation that was the basis for this paper. Use of open Internet standards, such as certificates-based two-way SSL, allows future extension of the architecture to message-level security implementations as they mature and are more widely available.

Summary

This paper proposes a security model for health information networks and a minimum set of security mechanisms needed to address security requirements of the model. These are shown to counter the security threats in a public network. The minimum set of security controls provides a simple and convenient framework for health information networks to design and implement their security architecture.

Acknowledgments

This paper is based on projects funded by the Markle Foundation Connecting for Health, and the Office of the National Coordinator for Healthcare Information Technology, U.S. Department of Health and Human Services. The authors acknowledge the sponsors of these projects, and also wish to thank Jared Adair, Lonnie Blevins, John Calladine, Carol Diamond, Don Grodecki, John Lightfoot and Clay Shirky, for reviewing and contributing valuable inputs to precursor documents from which this paper is drawn.

References

1. CSC Connecting for Health. CSC-CfH NHIN Security Architecture and Design. 2007 Jan 15 (available on request from authors)
2. Office of the National Coordinator (HHS). Harmonized Use Case for Bio-surveillance (Visit, Utilization and Lab Result Data). 2006 March 19.
3. Office of the National Coordinator (HHS). Harmonized Use Case for Electronic Health Records (Laboratory Result Reporting). 2006 March 19.
4. Office of the National Coordinator (HHS). Harmonized Use Case for Consumer Empowerment (Registration and Medication History). 2006 March 19.
5. Markle Foundation, The Connecting for Health. Common Framework M2: Model Contract for Health Information Exchange. http://www.connectingforhealth.org/commonframework/index.html, 2006 April.
6. Markle Foundation, The Connecting for Health. Common Framework P8: Breaches of Confidential Health Information. http://www.connectingforhealth.org/common framework/index.html, 2006 April.
7. OASIS, eXtensible Access Control Markup. Language, www.oasis-open.org/committees/xacml/, 2007 July.
8. V. Ramakrishna, K.Eustice, M.Schnaider. Approaches for Ensuring Security and Privacy on Unplanned Ubiquitous Computing Interactions. Workshop on Research and Challenges in Security and Privacy for Mobile and Wireless Networks, 2006 March.
9. D. Lekkas, Computer Communications. Establishing and managing trust within the Public Key Infrastructure, Vol. 26, No. 16 (2003).
10. NIST Special Publications 800-53, "Recommended Security Controls for Federal Information Systems", 2005 February.
11. SAML 2.0 and Related Work in XACML and WS-Security, http://lists.oasis-open.org/archives/security-services/200506/zip00000.zip
12. Centers for Medicare and Medicaid Services. Health Insurance Portability and Accountability Act of 1996 http://www.cms.hhs.gov/HIPAAGenInfo/