NISTIR 7497

Draft Security Architecture Design Process for Health Information Exchanges (HIEs)

Matthew Scholl
Kevin Stine
Kenneth Lin
Daniel Steinberg

National Institute of Standards and Technology
U.S. Department of Commerce

NISTIR 7497

Draft Security Architecture Design Process for Health Information Exchanges (HIEs)

Matthew Scholl
Kevin Stine
Computer Security Division
Information Technology Laboratory
National Institute of Standards and Technology
Gaithersburg, MD 20899

Kenneth Lin
Daniel Steinberg
Booz Allen Hamilton

January 2009

U.S. Department of Commerce
Carlos M. Gutierrez, Secretary

National Institute of Standards and Technology
Patrick D. Gallagher, Deputy Director

Acknowledgements

The authors, Matthew Scholl and Kevin Stine from NIST and Kenneth Lin and Daniel Steinberg from Booz Allen, wish to thank their colleagues and reviewers who contributed greatly to the document's development. A special note of thanks goes to Christina Salameh from Booz Allen for her keen and insightful assistance throughout the development of this document.

TABLE OF CONTENTS

TABLES AND FIGURES

1.0 Executive Summary

Protecting electronic patient health information is crucial to the deployment of a Health Information Exchange (HIE). As noted in the Summary of the National Health Information Network Report from the Office of the National Coordinator, "An important core competency of the HIE is to maintain a trusting and supportive relationship with the organizations that provide data to, and retrieve data from, one another through the HIE. The trust requirement is met through a combination of legal agreements, advocacy and technology for ensuring meaningful information interchange in a way that has appropriate protections."[1]

The purpose of this publication is to provide a systematic approach to designing a technical security architecture for the exchange of health information that leverages common government and commercial practices and that applies them specifically to the HIE domain. This publication assists organizations in ensuring that data protection is adequately addressed throughout the system development life cycle, and that these data protection mechanisms are applied when the organization develops technologies that enable the exchange of health information.

This operating model will help organizations that are implementing HIEs to:

- Understand major regulations and business drivers;
- Identify cross-organizational enabling services;
- Define supporting business processes (for each service);
- Develop notional architectures (as a blueprint to support services, processes, and the selection of technical solutions); and
- Select technical solutions.

2.0 Introduction

The secure exchange of electronic health information is important to the development of electronic health records (EHRs) and to the improvement of the U.S. healthcare system. While the U.S. healthcare system is widely recognized as one of the most clinically advanced in the world, costs continue to rise, and often preventable medical errors occur. Health information technology (HIT), and specifically the development of electronic health records for use in both inpatient and ambulatory care settings, have the potential for providing reliable access to health information and thereby improving the healthcare system. However, the prospect of storing, moving, and sharing health information in electronic formats raises new challenges on how to ensure the data is adequately protected.

Currently, protected health information (PHI)[2] is scattered among various parties including providers and payers, with patients maintaining limited control over the collection, access, use, and disclosure of their health information. The challenge of protecting this sensitive information is exacerbated when an electronic version of health information can be shared much more easily

[1] Summary of the NHIN Prototype Architecture Contracts, A Report for the Office of the National Coordinator for Health IT 31 May 2007

[2] This document uses the term "protected health information" (PHI) as it is suitably broad and well-understood by the health information technology community. While the term was coined and is defined by the HIPAA Security and Privacy Rules, note that the material in this document may also be instructive to healthcare entities that are not HIPAA-covered entities, or to the development of systems that will contain personal health information otherwise excluded from the definition of PHI.

than health records are exchanged today. The protection of a patients' health information is an important factor in the adoption of the EHR.

Integrating security across different business and technical layers is necessary in order to address complex data protection challenges in today's HIEs. This publication presents a five-layered architecture design process as a systematic approach to identify and implement HIE security and privacy. The five-layers, which are required elements for ensuring data protection, include: 1) policies; 2) core services; 3) business processes; 4) notional architecture; and 5) technical solutions. The security architecture design process provides a scalable and standardized methodology to guide HIE system development in the integration of data protection mechanisms across each layer, and results in a technology selection and design that satisfies high level requirements and mitigates identified risks to organizational risk tolerances.

2.1 Purpose and Scope

The purpose of this publication is to provide a systematic approach to designing a technical security architecture for the exchange of health information by leveraging common government and commercial practices and applying them specifically to the HIE domain. The publication defines the five-layers of this design process, their purposes and their relationships, and how they work together systematically to facilitate the secure exchange of protected health information.

This publication focuses specifically on the health information exchange process; it does not discuss the development of the entire information technology architecture of an HIE. Many organizations must comply with data protection laws at the local, state, federal, or international levels that will require them to conduct certain activities under specific operational parameters. While this publication does not directly address nontechnical issues such as those related to laws, regulations, and policies, specific roles and responsibilities, training, human resources issues, or nontechnical privacy issues, it does describe an architecture design process that allows for their integration into the information technology architecture of an HIE.

While the main focus of this document is security architecture, it is understood that privacy protections are essential to the collection, access, use, and disclosure of protected health information. For the purposes of this document, technical assurance of privacy is viewed as a subset of confidentiality. Implementation of security technologies that support confidentiality objectives may in turn support the technical implementation of privacy policy.

2.2 Audience

HIE executives, HIE security policy developers, HIE security architects, and technical solution providers are the principal audience for this publication. The objective is to provide a development approach through various stages of an HIE lifecycle to produce a security architecture for the identification of appropriate technologies or to serve as an evaluation model for existing HIEs.

2.3 Document Organization

The remaining sections of this document discuss the following:

- Section 3.0, HIE Contexts, describes the scope and characteristics of the four main HIE contexts discussed in this document.
- Section 4.0, HIT Security Architecture Design Process Overview, introduces the five-layer operating model to address those barriers.

- Section 5.0, Capstone Policies – Layer 1, identifies and describes many major U.S. laws, regulations, and guidelines that influence and, in many cases, drive the development of an organization's policies for ensuring secure exchange of health information.
- Section 6.0, Enabling Services – Layer 2, identifies and discusses 12 services, derived from common industry-wide practices, necessary to implement the Capstone Policies.
- Section 7.0, Enabling Processes – Layer 3, describes processes that expand the Enabling Services into detailed, HIE-specific business requirements.
- Section 8.0, Notional Architecture – Layer 4, identifies architecture design principles and constructs that will serve as inputs, along with Capstone Policies and Enabling Services and Processes, to create a Notional Architecture, the blueprint to drive technical solution decisions.
- Section 9.0, Technology Solutions and Standards – Layer 5, illustrates the steps to select the technical solutions and data standards that will satisfy the requirements specified in the Notional Architecture.
- Section 10.0, Building a Nationwide HIE using Regional HIEs, discusses using a federation of Regional HIEs to construct a Nationwide HIE with federated security services.
- Appendix A, Applying the Security Architecture Design Process, employs the five-layer design process to a specific American Health Information Community (AHIC) Use Case to illustrate the analyses and considerations that need to be made when applying this model to the exchange of health information.

3.0 HIE Contexts

There are many contexts in which health information can be exchanged. Therefore, it is important to identify under which contexts the security architecture design process presented in this publication is most applicable. Four main HIE contexts are identified in the following figure: ad hoc, regional, multi-regional and national.

Figure 1. HIE Contexts

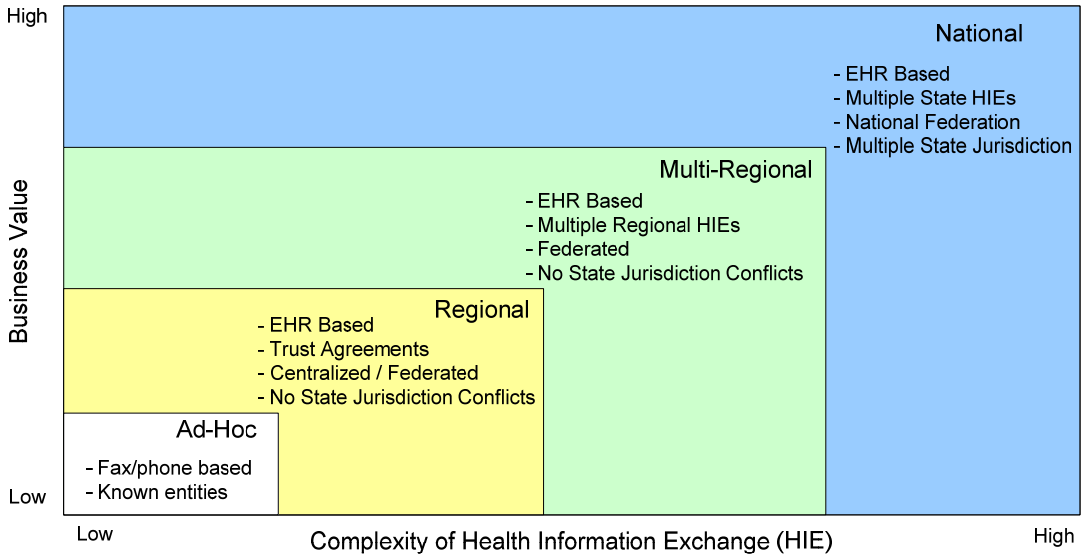

3.1 Ad Hoc HIEs

An ad hoc HIE occurs when two healthcare organizations exchange paper-based health information, usually under the precondition of familiarity and trust, using traditional mechanisms such as faxing and phone calls. Health organizations that currently practice paper-based ad hoc HIEs may find it impractical to justify the cost to migrate into electronic health record (EHR)-based HIEs unless there is a regional force behind it.

3.2 Regional HIEs

Regional HIEs are those that consist of two or more legally and commercially independent institutions that share EHRs, but where no state jurisdictional issues exist that prevent or impede the sharing of data.[3] The HIE network includes clinicians, hospitals, labs, pharmacies, insurance companies, and other key health domain players. Participating organizations will normally draft a trust agreement to govern the information exchange. Depending on the scale, the technical architecture might be centralized or federated. Regional HIEs do not have state jurisdictional conflicts and are large enough to justify self-sustained EHR-based HIEs. They are considered simpler than multi-regional and national HIEs because of their smaller scale and lack of state jurisdictional conflicts.

3.3 Multi-Regional HIEs

Multi-regional HIEs connect multiple regional HIEs. They may cross state lines or other physical boundaries. They are usually EHR-based. Since they connect multiple regional HIEs, they will have a federated technical architecture. For multi-regional HIEs, conflicts of laws may require complex solutions.

3.4 National HIEs

The national HIE connects many multi-regional HIEs. Hence it is based on EHRs, involves multiple state jurisdictions, and has a national federated technical architecture. Multi-regional and national HIEs have a different focus than regional HIEs. Regional HIEs are basic building blocks that focus on developing effective and localized solutions to meet specific HIE needs (research, clinical trials, patient transfer, etc). Multi-regional and national HIEs focus on building the backbone infrastructure needed to connect various regional HIEs.

This publication focuses on the needs of regional HIEs. Assuming that the security architectures and other system aspects of regional HIEs are interoperable, these HIEs can serve as the “building blocks” for larger multi-regional and national HIEs, and therefore represent a scalable solution for the ultimate emergence of a national HIE.

4.0 HIE Security Architecture Design Process

Technical solutions that facilitate the exchange of health information can be complex. With various policies and standards, and an ever-changing technical landscape, a systematic approach to designing an HIE security architecture can allow practitioners to analyze all policy requirements and ultimately refine them into a technology-neutral, vendor-neutral, standards-based architecture to drive technical solution decisions.

[3] Regional HIEs might include HIE members from different states. This document assumes the state jurisdiction conflicts, if any, have been reviewed and resolved.

The use of a systematic approach plays a significant role in a successful and secure HIE implementation. The HIE security architecture design process was developed to assist HIEs in meeting this need by providing a five-layered methodology for successful HIE security technology identification and selection as illustrated in the following figure:

Figure 2. HIE Security Architecture Design Process

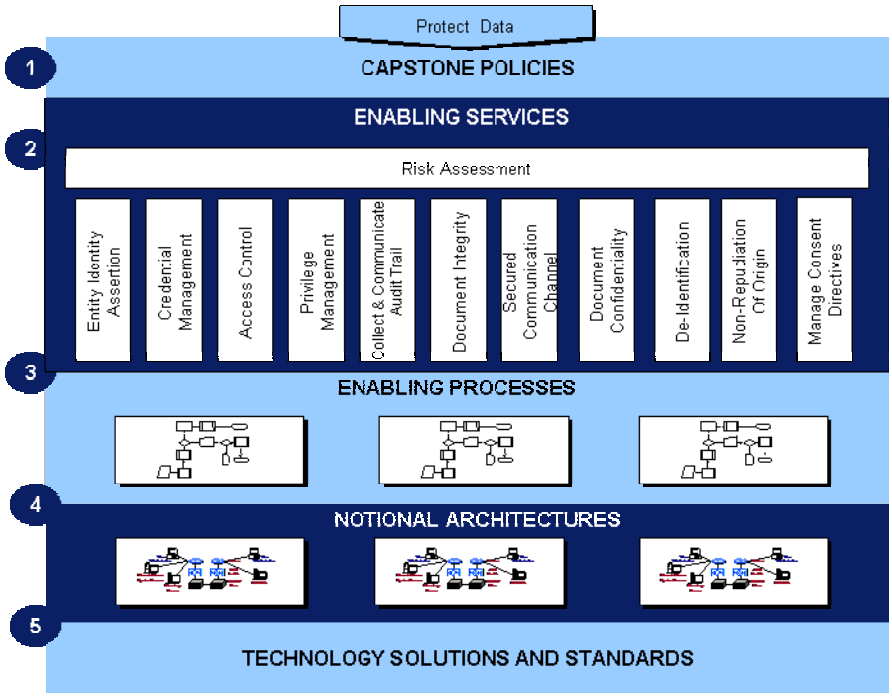

1) ***Capstone Policies:*** Capstone policies provide overall requirements and guidance for protecting health information within HIEs. They can be driven by national laws, regulations, and guidelines, state regulations, organizational policies, business needs, or policies developed for specific HIEs.

2) ***Enabling Services:*** Enabling services define the nomenclature of services required to implement capstone policies. Enabling services are designed to be HIE context-independent. Services presented in this publication are derived from common industry-wide data protection practices and then customized to specifically address the requirements of HIEs.

3) ***Enabling Processes:*** Enabling processes define the operational baseline via use cases and scenarios for enabling services. Enabling processes are HIE context-dependent. Two HIEs could, for example, have different enabling processes implementing the "Access Control" service.

4) ***Notional Architectures:*** Notional architectures define the technical constructs (e.g., role-based access control and directory services) and their relationships to implement enabling processes. Notional architecture is the blueprint to drive the selection of technical solutions and data standards. Notional architecture is standards-based, technology-neutral, and vendor-neutral.

5) ***Technology Solutions and Standards:*** Technical solutions and data standards represent the selected the technical solutions and the data standards needed to implement the notional architecture.

Each layer of the design methodology is described in the following sections.

5.0 Capstone Policies – Layer 1

Capstone policies are those policies that are developed by institutions participating in HIEs and that provide overall requirements and guidance for protecting health information within those HIEs. Ideally, capstone policies should address the requirements imposed by all laws, regulations, and guidelines at the national, state, and local levels; business needs; and policies at the institutional and HIE levels.

In developing capstone policies, organizations must identify the requirements that these laws, regulations, and other authorities impose on HIEs. One challenge in ascertaining that all such requirements have been identified is that these sources of requirements may not be specific to health information systems. For federally owned or operated systems, for example, other requirements such as the Federal Information Security Management Act (FISMA[4]) will also need to be considered. For this reason, organizations must consider the expert input of appropriate legal counsel in assembling these requirements.

Within this section, many major U.S. federal laws relevant to the development of HIE security and privacy architectures are identified. For virtually all U.S. entities, however, other federal and state laws will also need to be considered. These representative laws are identified as illustrative and as assistance to organizations. In particular, state laws may be significant. Under the Health Insurance Portability and Accountability Act (HIPAA), more stringent state laws that may require additional or greater protections for PHI must be followed. The existence of HIPAA does not negate such requirements, or excuse the covered entity from addressing them.

In many cases, relevant laws, regulations, and policies will impose other requirements aside from those that help identify capstone policies. These authorities may also establish broad goals or end states, without specifically defining enabling services, and may need to be interpreted based on industry best practices or reasonable safeguards. Addressing the text alone, therefore, may not be sufficient in order to ensure secure HIEs. In cases where authorities urge the institution of appropriate policies without proposing specific safeguards, practitioners should not confine themselves to developing capstone policies that merely satisfy compliance, but should view these authorities as setting only a minimum set of requirements.

5.1 Health Insurance Portability and Accountability Act

The Health Insurance Portability and Accountability Act of 1996 (HIPAA, Public Law 104-191) is the most well-known and influential law affecting the security and privacy practices of many healthcare organizations in the United States (those that are "covered entities" under the Act). HIPAA required the Secretary of Health and Human Services (HHS) to create sets of regulations on several topics related to electronic healthcare transactions, including the privacy of protected health information (PHI) and the security of electronic protected health information (ePHI). Any private and secure health information exchange must accommodate the functionality to support

[4] FISMA (P.L. 107-347, Dec. 2002) requires each federal agency to develop, document, and implement an agency-wide program to provide information security for the information and information systems that support the operations and assets of the agency, including those provided or managed by another agency, contractor, or other source.

the standards of both the HIPAA Privacy Rule[5] and the HIPAA Security Rule.[6] This publication reflects the standards and implementation specifications of the HIPAA Security and Privacy Rules that will drive the architectural framework and technical solutions of a mature Health Information Exchange. However, those requirements that do not necessarily create parameters or other requirements for health information exchange are not explored. These requirements include policies and practices that may be implemented and enforced at individual organizations. For a fuller discussion of the HIPAA Security Rule, including resources for understanding and addressing its requirements, see NIST Special Publication 800-66 (October 2008), *An Introductory Resource Guide for Implementing the Health Insurance Portability and Accountability Act (HIPAA) Security Rule*; for a discussion on the HIPAA Privacy Rule, see *Summary of the HIPAA Privacy Rule*, HHS Office for Civil Rights, (May 2003), available at www.hhs.gov.

5.2 Other Key Drivers for Capstone Policies

Other laws and regulations may also drive requirements for the functionality of security controls, depending on an HIE or its components' functions, activities, business partners, the types of information it handles, status as a government agency or private commercial entity, or its geography. This publication identifies many of the most common federal laws and regulations that create requirements for capstone policies for large numbers of organizations across the United States. However, all related federal laws and regulations, and relevant state and local policies may not be identified.

Table 1 lists a selection of the laws and regulations that may affect the healthcare transactions for some entities. In addition to these, there are numerous state and local laws and regulations that may impact technology selection and implementation.

Table 1. Capstone Policy Drivers

Entities Affected	Capstone Policy Drivers	Enabling Services Affected	Comment
Any institution conducting research involving human subjects conducted, supported or otherwise subject to regulation by any federal department or agency	Federal Regulations: Protection of Human Subjects, 45 CFR Part 46 ("The Common Rule" for Human Subjects Protection)	Risk Assessment; Document Confidentiality; Manage Consent Directives; De-Identification	Requires consent of research subjects, subject to some exemptions, including research on records where **if** the information is recorded by the investigator in such a manner that subjects cannot be identified, directly or through identifiers linked to the subjects. Consent forms must include notice of adequate provisions to protect the privacy of subjects and to maintain the confidentiality of data
All institutions conducting clinical investigations regulated by the Food and Drug Administration or supporting applications for research or marketing permits for products regulated by the Food and Drug Administration.	Federal Regulations: Protection of Human Subjects, 21 CFR Part 50 ("The FDA Rule" for Human Subjects Protection)	Risk Assessment; Document Confidentiality; Manage Consent Directives	Requires consent of research subjects, subject to some exemptions, especially for research conducted on investigative new drugs (INDs). Consent forms must include notice of adequate provisions to protect the privacy of subjects and to maintain the confidentiality of data

[5] Standards for Privacy of Individually Identifiable Health Information, Final Rule ("The HIPAA Privacy Rule"), 45 CFR Parts 160, 162, and 164.

[6] Health Insurance Reform: Security Standards; Final Rule ("The HIPAA Security Rule"), 45 CFR Parts 160, 162, and 164.

Entities Affected	Capstone Policy Drivers	Enabling Services Affected	Comment
Research facilities that transmit information to the Food and Drug Administration	Federal Regulations: Electronic Records; Electronic Signatures; Electronic Submissions; Establishment of Public Docket; Notice 28 CFR Part 11	Entity Identity Assertion (Authentication). Collect and Communicate Audit Trail; Document Integrity; Secured Communication Channel; Non-Repudiation of Origin	Especially relevant to transmission of information related to new drug trials
Organizations that disclose patient alcohol and drug abuse records	Federal Regulations: Confidentiality of Alcohol and Drug Abuse Patient Records 42 CFR Part 2	Manage Consent Directives; Document Confidentiality	Restrictions on disclosing patient alcohol and drug abuse records Intended to encourage patients to seek help for abuse and addiction
Public and private postsecondary educational institutions receiving federal funds	Family Education Rights and Privacy Act of 2000 (FERPA) 20 U.S.C. § 1232g; 34 CFR Part 99	Manage Consent Directives; Document Confidentiality	FERPA protects education records of students enrolled at covered institutions. Its effect on health records is complex, but in general health records are deemed to be "education records" if held by the covered institution.
Medicaid providers	Federal Medicaid Confidentiality Standards 42 CFR §431.300 et seq.	Manage Consent Directives; Document Confidentiality	Requires state Medicaid agencies to document rules for specifying the conditions for release and use of information about applicants and recipients.
Federal agencies and contractors	Privacy Act of 1974 5 U.S.C. § 552a	Manage Consent Directives; Document Integrity; Document Confidentiality; Risk Assessment	Requirements related to the collection, disclosure, and documentation of most personal information, including health information, held by federal agencies
Federal agencies and contractors	Federal Information Security Management Act (FISMA) 44 U.S.C. § 3541	Entity Identity Assertion (Authentication); Access Control (Authorization); Collect and Communicate Audit Trail; Document Integrity; Secured Communication Channel; Document Confidentiality; Non-Repudiation of Origin; Risk Assessment; Credential Management; Privilege Management	Among many other requirements, federal agencies must provide quarterly Privacy Management Report on privacy protections
Federal agencies and contractors	OMB Memoranda	All	Various requirements related to privacy in the system development life cycle, analysis, reporting, and risk reduction relevant to federal agencies

6.0 Enabling Services - Layer 2

Enabling services are those services that are necessary to implement capstone policies. These services are typically HIE "context-independent," meaning they will be included in all HIEs although the manner of implementation may be different for different systems. For example, two HIEs both providing "Access Control" services might have different implementation models for them.

The function of enabling services is to provide a standard set of minimum requirements across HIEs, but not to establish definitive methods for obtaining them. This means that every HIE will need to deploy enabling services using appropriate solutions that must be identified and selected.

Having a consistent, standards-based set of enabling services can benefit future interoperability between HIEs. This standardization provides a basic assurance level on the implementation of security and privacy controls, and it will be easier to determine and address discrepancies among HIEs.

Services presented in this publication are derived from the Healthcare Information Technology Standards Panel (HITSP)(www.hitsp.org) Security, Privacy, and Infrastructure constructs, which detail the selection of standards to meet Use Case requirements defined by the American Health Information Community (AHIC), and common established security principles. These are then distilled to specifically address the data protection requirements of HIEs.

Figure 3. Enabling Services

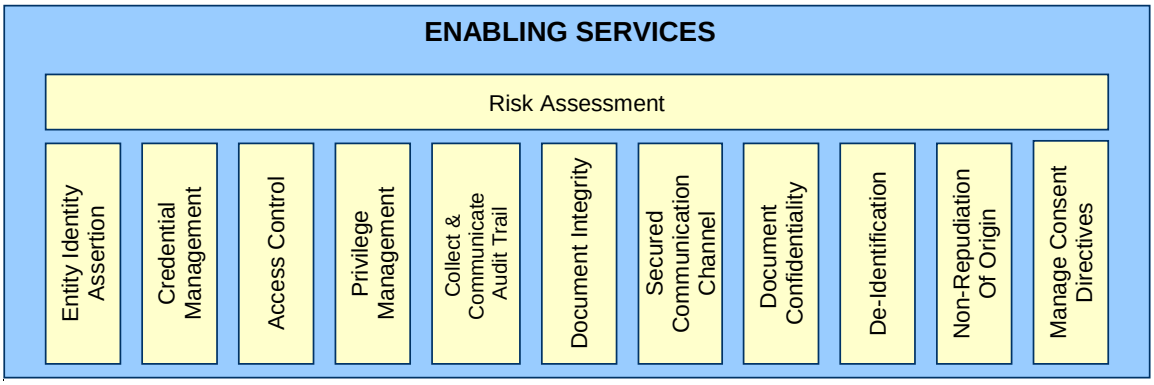

Table 2. Enabling Services and Definitions

Service Name	Source	Definition
Risk Assessment	Security and Privacy Principles	To identify risks to HIE operations based on threats, assets, vulnerabilities, and probabilities of threat success.
Entity Identity Assertion (Authentication)	HITSP Construct	To ensure that an entity is the person or application that claims the identity provided.
Credential Management	Security Principles	To manage the life cycle of entity credentials used for authentication and access control.
Access Control (Authorization)	HITSP Construct	To ensure that an entity can access protected resources if they are permitted to do so.
Privilege Management	Security Principles	To manage privileges (grant or deny) associated with entities for HIE transactions.
Collect and Communicate Audit Trail	HITSP Construct	To define and identify security relevant events and the data to be collected and communicated as determined by policy, regulation, or risk analysis.
Document Integrity	HITSP Construct	To ensure the integrity of a document that is exchanged or shared.
Secured Communica-tion Channel	HITSP Construct	To ensure the authenticity, the integrity, and the confidentiality of transactions, and the mutual trust between communicating parties.
Document Confidentiality	Security Principles	To ensure the confidentiality of a document that is exchanged or shared.
De-identification	Privacy Principles	To remove individual identifiers, so that it cannot be used to identify an individual.
Non-Repudiation	HITSP Construct	To ensure that information received can be confirmed as having been sent by the apparent sender, and that no reasonable basis exists for claiming that the information came from some other source; and to ensure that the sender can confirm that the intended recipient has received the information."

Service Name	Source	Definition
Manage Consent Directives	HITSP Construct	To ensure that individually identifiable health information is only accessed with an individual's consent.

6.1 Assumptions

Enabling services identified in this section focus on the "exchange" aspect of HIE operations. To truly create a secure HIE environment, additional services are required to protect the data of the participating entities' organization infrastructure (end points). Services such as contingency planning and configuration management that are used to secure a participating entity's infrastructure are not covered in this document, which focuses on the exchange of health information only. Internal services, such as Contingency Planning and Configuration Management, which are used to secure an entity's internal infrastructure, are not covered. Also, many of the managerial and operational security controls that are not directly part of a cross enterprise exchange are also not addressed but these measures may be critical for a complete security program for an organization. Organizations must ensure that these controls implemented in their HIEs are integrated and mutually supportive of the technology architecture derived from the design process outlined in this publication.

6.2 Enabling Services

The twelve enabling services identified below are derived from the HITSP Security, Privacy, and Infrastructure construct definitions, and common established security principles. Information is provided, where available, to consolidate work previously conducted in this area in order to support a standardized, common vocabulary for HIE concepts.

In the following section, a definition and an illustrating example are provided for each service. Also provided is a list of other documents with further information regarding the specific enabling service. These referenced documents offer information for further insight and clarification.

6.2.1 Risk Assessment

Definition: To identify risks to HIE operations that may compromise PHI information resulting in unauthorized disclosure, loss of integrity, or lack of availability.

Illustration: A county government decides to build a health information exchange network to research heart disease. HIE-participating entities perform a comprehensive risk assessment by examining the information to be exchanged over the network. They decide to categorize the information into three assurance levels (low, medium, high) based on the sensitivity of the information. The community then decides what measures are required for each assurance level. Specific threats are evaluated for their potential to exploit existing vulnerabilities and documented as threats. Existing measures are evaluated for their ability to mitigate these threats and additional measures are decided upon to ensure residual risks are acceptable. The complete set of security controls is documented and used in a trust agreement enforced by the HIE.

Other References:
HIPAA Security Rule references: 164.308(a)(1), Implementation Specification: Risk Analysis.

NIST 800-53 security control family: Risk Assessment (RA)

NIST Publications:
SP 800-30, *Risk Management Guide for Information Technology Systems*

SP 800-39, DRAFT *Managing Risk from Information Systems: An Organizational Perspective*

6.2.2 Entity Identity Assertion (Authentication)

Definition: To ensure that an entity is the person or application that claims the identity provided.

Illustration: A Doctor at Hospital One wishes to access Joan Taylor's records for the purposes of entering new data concerning her health status. This new data may later be accessed by other healthcare providers that are members of the HIE. The Doctor uses an approved computer terminal to access the Hospital's HIE system. Before accessing the patient's record, the Doctor is asked to provide his username and a password that he has chosen and that is known only to him, but that is recognized by the HIE information system. The enabling service then receives the input of the Doctor's asserted user name and identity, and compares that assertion with preexisting records to authenticate the Doctor.

Other References:
HITSP C 19, Entity Identity Assertion

CCHIT Requirement description: Security: Item numbers 10012-10023, 10031, "Authentication," 9081-9083 and 9144-9146, "Entity Identity Assertion."

HIPAA Security/Privacy Rule references: 164.308(a)(5), Implementation Specification: Password Management; 164.312(a)(1), Implementation Specification: Unique User Identification; 164.312(d), Standard: Person or Entity Authentication.

NIST 800-53 security control family: Identification and Authentication (IA)

NIST Publications:
Draft SP 800-63-1, *Electronic Authentication Guideline*

6.2.3 Credential Management

Definition: To manage the life cycle of entity credentials used for authentication and access control.

Illustration: Hospital One has three assurance levels for the information exchanged on its network. For information of each assurance level, accepting authentication credentials and the life cycle of those credentials are defined. The credential life cycle includes an identity proofing process to obtain, validate, renew, and revoke the credential.

Other References:
CCHIT Requirement description: See item numbers 10012-10023, 10031, "Authentication," 9081-9083 and 9144-9146, "Entity Identity Assertion," although CCHIT does not make a strong distinction between the fields of "Credential Management" and "Authentication."

HIPAA Security Rule references: 164.312(d), Standard: Person or Entity Authentication.

NIST 800-53 security control family: AC-2, Account Management; IA-4, Identifier Management; IA-5, Authenticator Management.

NIST Publications:

SP 800-63, *Electronic Authentication Guideline*

SP 800-57, *Part 1, Recommendation for Key Management - Part 1: General*

800-56A, *Recommendation for Pair-Wise Key Establishment Schemes Using Discrete Logarithm Cryptography*

6.2.4 Access Control (Authorization)

Definition: To ensure that an entity can access protected resources if they are permitted to do so.

Illustration: A Doctor joins the staff at Hospital and needs to have access to the information system supporting Hospital's participation in the HIE. The attending physician, having confirmed that such access would be appropriate, contacts the Hospital's HIE system administrator and requests that the Doctor receive access to the system. The system administrator creates an account for the Doctor and sets his permissions such that the Doctor is able to access PHI for all patients and services that he may need to access in order to perform his job. This process ensures that the Doctor is in the restricted group of individuals who, under the policies of the institution, may receive access to PHI, including PHI accessed at other healthcare providers within the HIE. The enabling service reflects that the Doctor has received permission to access these systems based on the organization's authorization policy and validates the permission every time an access request is made.

Other References:
HITSP TP 20 "Access Control"

CCHIT Requirement description: Security: Item numbers 9003 and 9004, "Access Controls," category descriptions relevant to management of security authorizations.

HIPAA Security/Privacy Rule references: 164.308(a)(4), Implementation Specification: Access Authorization, Implementation Specification: Access Establishment and Modification.

NIST 800-53 security control family: Access Controls (AC)

6.2.5 Privilege Management

Definition: To manage privileges (grant or deny) associated with entities for HIE transactions.

Illustration: Doctor from Hospital One participates in a research HIE. He has access to certain research projects within the network. The research HIE administrator provisions Doctor and grants him access privileges only to research projects in which he participates. When any research project is finished, the HIE administrator will remove Doctor's access privilege accordingly.

Other References:
CCHIT Requirement description: See Item numbers 10012-10023, 10031, "Authentication," 9081-9083 and 9144-9146, "Entity Identity Assertion," although CCHIT does not make a strong distinction between "Privilege management" and "Authorization."

HIPAA Security Rule references: 164.308(a)(4), Standard: Information Access Management.

NIST 800-53 security control family: AC-13, Supervision and Review-Account Management.

6.2.6 Collecting and Communicating Audit Trails

Definition: To define and identify security relevant events and the data to be collected and communicated as determined by policy, regulation, or risk analysis to support identification of those security relevant events.

Illustration: The System Administrator reviews a file that is generated by the HIE-enabling system on a daily basis. The Audit Trail enabling service generates a record of the users who have accessed what files and when. The enabling service also makes note of any attempts to access the system from an unauthorized terminal; the use of an expired username or password; unusual numbers of password attempts; and other potential attempted violations of security policies. The System Administrator may take appropriate action to ensure that future attempts at gaining unauthorized access are unsuccessful.

Other References:
HITSP T 15 "Collect and Communicate Security Audit"

CCHIT Requirement description: Security: Item numbers 8019-8021, "Audit logging and error handling for data access and exchange" and 1284 and 1285, "Audit trail."

HIPAA Security/Privacy Rule references: 164.308(a)(5), Implementation Specification: Log-In Monitoring, 164.312(b), Standard: Audit Controls

NIST 800-53 security control family: Audit and Accountability (AU)

NIST Publications:
SP 800-92, *Guide to Computer Security Log Management*

6.2.7 Ensuring Document Integrity

Definition: To ensure the integrity of a document that is exchanged or shared.

Illustration: Hospital One sends a record to Hospital Two using a one-way hash to confirm that the record has not been altered in transit.

Other References:
HITSP TP 13, Manage Sharing of Documents

CCHIT Requirement description: Security: Item number 8039, "Routing of consumer requests for data corrections," 8014-8018, "Data integrity and non-repudiation checking," 9047-9080, "Data integrity auditability."

HIPAA Security/Privacy Rule references: 164.308(a)(5), Implementation Specification: Protection from Malicious Software; 164.312(c)(1), Standard: Integrity; 164.312(e)(1), Implementation Specification: Integrity Controls.

NIST 800-53 security control family: Security Controls: Transmission Integrity (SC-8), System and Information Integrity (SI)

NIST Publications:
SP 800-106, DRAFT *Randomized Hashing Digital Signatures*

SP 800-25, *Federal Agency Use of Public Key Technology for Digital Signatures and Authentication*

6.2.8 Secure Communication Channel

Definition: To ensure the authenticity, the integrity, and the confidentiality of transactions, and the mutual trust between communicating parties.

Illustration: Hospital One sends records containing PHI to Hospital Two. All information is protected in transit by a virtual private network (VPN) ensuring that there are no opportunities to intercept the data in transit.

Other References:

HITSP T 17, Secure Communication Channel

CCHIT Requirement description: Security: Item numbers 1234-1236, "Inter-provider communication."

HIPAA Security/Privacy Rule references: 164.312(e)(1), Standard: Transmission Security; Implementation Specification: Encryption.

NIST 800-53 security control family: System and Communications Protection (SC)

NIST Publications:
SP 800-113, *Guide to SSL VPNs*

SP 800-77, *Guide to IPSec VPNs*

SP 800-58, *Security Considerations for Voice Over IP Systems*

SP 800-52, *Guidelines for the Selection and Use of Transport Layer Security (TLS) Implementations*

SP 800-45, Version 2, *Guidelines on Electronic Mail Security*

6.2.9 Preserving Document Confidentiality

Definition: To ensure that personal health information is not sent intentionally or unintentionally to a party that is not authorized to view it, either by the patient or by a provider that has received the patient's authorization or a waiver of the patient's authorization.

Illustration: Hospital One intends to send records containing PHI to Hospital Two. In selecting recipients, Hospital One is asked to provide the name of the receiving physician in an appropriate entry field, but the sender cannot type in any e-mail address, or create a new account for an unauthorized recipient. The enabling service that limits the choice of recipient helps to ensure that only those with access to the HIE are able to receive patient's PHI.

Other References:
HITSP TP 13, Manage Sharing of Documents

CCHIT Requirement description: Security: Item numbers 1276-1280, "Enforcement of confidentiality."

HIPAA Security Rule references: The HIPAA Privacy Rule contains significant information on conditions for maintaining confidentiality.

NIST 800-53 security control family: System and Communications Protection (SC), Transmission Confidentiality (SC-9)

NIST Publications:

SP 800-113, *Guide to SSL VPNs*

SP 800-111, *Guide to Storage Encryption Technologies for End User Devices*

SP 800-77, *Guide to IPSec VPNs*

SP 800-58, *Security Considerations for Voice Over IP Systems*

SP 800-52, *Guidelines for the Selection and Use of Transport Layer Security (TLS) Implementations*

SP 800-45, Version 2, *Guidelines on Electronic Mail Security*

FIPS 140-3, DRAFT *Security Requirements for Cryptographic Modules*

6.2.10 De-Identification

Definition: To ensure that individuals' records have all data elements removed before the data is shared for statistical, research, public health, or other reasons that do not benefit the data subject directly, and for which no authorization has been provided, such that there is no reasonable basis to believe that the information can be used to identify an individual.

Illustration: A Researcher at a Hospital wants to study the records of all patients with a particular form of cancer within a certain age range. He contacts his organization's research review board to confirm that his protocol will be conducted ethically and within all state, federal, and local laws and guidelines. He then contacts all providers in the HIE network and asks them to help him populate a database of de-identified information. Providers contact all patients fitting the profile and secure their consent. Each provider then uses the De-Identification enabling service to remove all potentially identifying information from each consenting patient's record, and then sends the record to a database. Because ages are a relevant research parameter, birth years are retained in each record, although exact birth dates are removed.[7]

Other References:
CCHIT Requirement description: Item number 9150, "Secondary uses of data" ("The system shall provide the means to suppress data elements upon request") and 7076-7078, "Secondary uses of clinical data."

HIPAA Security Rule and Privacy Rule references: The requirements for de-identification under the HIPAA Privacy Rule are explicitly laid out in Section § 164.514, *Other requirements relating*

[7] This illustration addresses several issues of potential varying interpretations by Institutional Review Boards and/or Privacy Boards. For example, some HIPAA analysts interpret the Privacy Rule that patient consent must be received to send patient information into a database or similar repository for the purposes of being de-identified. Where possible varying interpretations exist, the scenario adopts the most stringent requirements; it also intentionally avoids complications such as the availability of limited data sets.

to uses and disclosures of protected health information, subsections (a) (Standard: de-identification of protected health information), (b) (Implementation specifications: requirements for de-identification of protected health information), and (c) (Implementation specifications: re-identification).

NIST 800-53 security control family: None. This construct is highly specific to the healthcare environment and to healthcare information exchange in particular.

6.2.11 Non-Repudiation

Definition: To ensure that information received can be confirmed as having been sent by the apparent sender and that no reasonable basis exists for claiming that the information came from some other source; and to ensure that the sender can confirm that the intended recipient has received the information.

Illustration: Patient had routine blood work performed at Hospital One. Patient then asks Hospital One to electronically transfer the results of her blood work to her primary care physician. When the results were available, Hospital One digitally signed and securely transmitted them to the primary care physician specified by Patient. Upon receipt, Patient's primary care physician was able to confirm that the results of the blood work were submitted by Hospital One.

Other References:
HITSP C26, Nonrepudiation of Origin

CCHIT Requirement description: Item numbers 9138-9140, "Nonrepudiation of Origin."

HIPAA Security Rule references: 164.312(d), Standard: Person or entity authentication.

NIST 800-53 security control family: Non-Repudiation (AU-10)

NIST Publications:
SP 800-25, *Federal Agency Use of Public Key Technology for Digital Signatures and Authentication*

6.2.12 Managing Consent Directives

Definition: To ensure that individually identifiable health information is collected, accessed, used or disclosed only with a consumer's consent.

Illustration: A Hospital and a Specialist are both entities within the same HIE. The Hospital sends the Patient's PHI to the Specialist. The Specialist will review the PHI and provide a medical opinion, but will not interact with Patient directly. The Specialist's own in-house rules require that the Hospital confirms that the Patient has received an appropriate notice of privacy practices (NOPP) for data to be shared. The Managing Consent Directives service would enable the Specialist to confirm that Patient has received the NOPP.

Later, the Specialist wishes to de-identify the Patient's data and share it with a Researcher, also an HIE participant. Under HIPAA, patients must provide adequate consent before their data is sent to a repository for de-identification, so the Specialist asks the Hospital to contact the Patient to provide the necessary consent. The Patient does provide the consent, and when it is reflected

via the enabling service, the Specialist de-identifies the record and submits it to the Researcher's repository.

Other References:
HITSP TP 30, Manage Consent Directives

CCHIT Requirement description: Item numbers 1167-1172, "Manage consents and authorizations," and item numbers 9035-9043, "Consent."

HIPAA Security/Privacy Rule references: See Subpart E generally, especially 164.506, Consent for uses or disclosures to carry out treatment, payment, and healthcare operations

NIST 800-53 security control family: None. This construct is highly specific to the healthcare environment and to healthcare information exchange in particular.

7.0 Enabling Processes – Layer 3

The enabling processes define business processes for enabling services. While enabling services define the nomenclature of HIE data protection requirements, enabling processes expand the enabling services into detailed requirements based on an HIE's business practices. Enabling processes are HIE context-dependent (e.g., treatment, public health). Hence, HIEs of different contexts could implement the same enabling services with different enabling processes. The following paragraph is an example of enabling processes for an "Access Control" service:

> Joan Taylor owns a protein database at a research institution. Her protein database is used in the ***Hope research project*** HIE with research scientists from a local university. Joan defines the following processes for the "Access Control" service:
>
> - Only the Hope research project manager has read/write privileges to the database. However, the project manager can delegate read/write privileges to research project members.
> - The database will be open for Hope research project use from 10:00 a.m. to 3:00 p.m. everyday. Joan wants to reserve the other time slots for the research institution scientists.
> - All accesses, internal and external, to the protein database need to be logged.

As illustrated in this example, enabling processes are written in plain English and are derived from an HIE participants' business practices. Enabling processes are detailed requirements for enabling services. They should be clearly defined and fully vetted within the HIE context for each enabling service.

8.0 Notional Architecture – Layer 4

Capstone policies (Layer 1), enabling services (Layer 2), and enabling processes (Layer 3) serve as the inputs to create the notional architecture which will be the blueprint to drive technical solution decisions. The notional architecture defines major architecture constructs and their relationships to implement enabling processes. It is standards-based, and technology- and vendor-neutral. The notional architecture is dependent on the enabling processes and will vary between HIE implementations. The following components for building an HIE notional architecture are presented:

- *Architecture design principles* are guiding principles identified from information-sharing implementations in the industries that apply to HIE contexts; and
- *Architecture constructs* can serve as the basic building blocks for a notional architecture.

The notional architecture development process is illustrated in Figure 4.

Figure 4. Notional Architecture Development Process

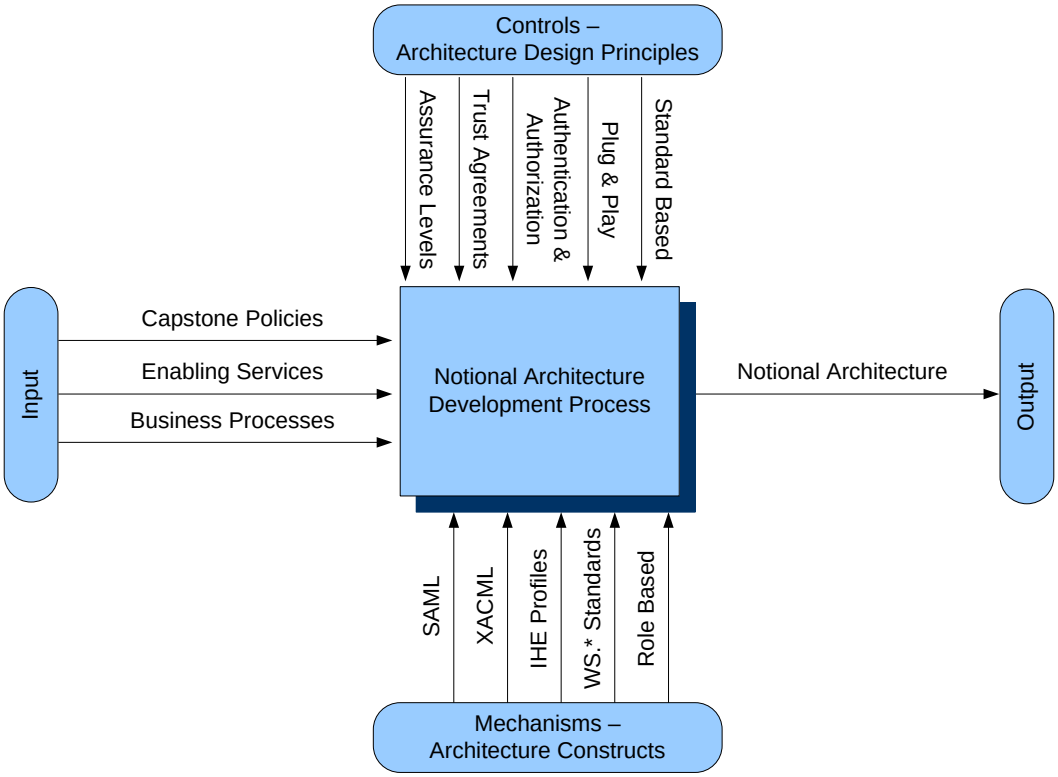

8.1 Architecture Design Principles

Architecture design principles are best practices derived from large-scale information-sharing implementations. Design principles serve as the overall guidance for building security and privacy services for HIEs. This publication identifies five design principles:

- Conduct a risk assessment to determine appropriate assurance levels for shared information;

- Create a “master” trust agreement describing requirements for a trust domain (trust domain is defined in Section 8.1.2);
- Separate credential management and privilege management;
- Develop data protection capabilities as plug-and-play services; and
- Maintain a standards-based, technology-neutral, and vendor-neutral architecture.

These design principles are described in more detail in the following sections.

8.1.1 Conduct a Risk Assessment to Determine Appropriate Assurance Levels for Shared Information

Conducting a risk assessment on the information exchanged in any HIE is fundamental and critical to the effective protection of the information. Organizations should be aware of the security and privacy risks with the exchanged information in order to design a proper architecture.

The results of a risk assessment can enable HIE transactions to be categorized into assurance levels. Assurance levels define the degree of confidence required to conduct a specific HIE transaction. The assurance levels reflect the sensitivity and criticality of the information. The following table lists several examples of HIE transactions with associated assurance levels.

Table 3. Illustrative Examples of Assurance Levels

Assurance level (Relative)	Example HIE Transactions	Information Classification
Low	Share public protein database	Public information
Medium	Share patient de-identified medical research information	Sensitive information
High	Share patient HIV information	Confidential information

Assurance levels are represented by a range (e.g., 1-2-3; high-medium-low, gold-silver-bronze) rather than absolute values due to their comparison nature. The representation of assurance levels helps an organization decide what kind of credential and what identity proofing process is needed for an HIE transaction (See Section 9.1.3 for details). The number of required assurance levels depends on the complexity of the information exchanged in HIEs.

8.1.2 Create a “Master” Trust Agreement Describing Security and Privacy Requirements for a Trust Domain

Once assurance levels are defined and risks are identified and mitigated, a trust domain can be created. A trust domain is a logical construct within which a single set of access control policies can be enforced for all HIE transactions. A master trust agreement can be created to enforce security requirements within a trust domain. The master trust agreement should be honored in every HIE transaction. For unique HIE transactions, specialized trust agreements might be created based on the master trust agreement. The organization that has a master trust agreement is able to provide every participating HIE entity a basic assurance, and avoid the complexity of requiring each participant in the HIE to execute a unique agreement with every other participant in the HIE.

8.1.3 Separate Credential Management and Privilege Management

Credential management governs the types of authentication credentials and their life cycle based on defined assurance levels. The following table lists examples of credentials of various assurance levels specific to credential management.

Table 4. Authentication assurance levels are mapped to application risks

Assurance level	Authentication Token and Required Process
Level 1 – Low	PIN # or password
Level 2 – Medium	Strong password, one-time password / ID proof
Level 3 – High	PKI credential / ID proof
Level 4 – Very High	Hard crypto token / ID proof

Credential management grants a HIE entity its identity within the HIE context. The HIE identity usually has a global effect within a specific HIE. Once a HIE credential is granted to a HIE entity, it will be recognized across the HIE context until it is revoked or expired.

Privilege management governs privileges of an HIE entity (i.e., what an entity can do after authentication). Granting privileges requires a trusted credential on a HIE entity who requests access to certain information. The decision to grant privileges is usually made locally by the HIE entity which guards the requested information.

Due to the very different nature of credential and privilege management (global authentication vs. local authorization), these two topics should be separated when an organization is developing the notional architecture. The trust agreement should identify what kinds of credentials will be accepted for each assurance level in a HIE. HIE entities which guard requested information need to use an interoperable authorization language to express authorization policies. Authorization decisions will be made locally at HIE entities. HIE entities which guard the information should assume full authority on granting access privileges.

The trust agreement can be easily created when credential management and privilege management are separated. Having the global authentication credential and local authorization authority allows HIE entities to better control what information is exposed to HIEs and what information should be protected inside their own boundary.

8.1.4 Develop Data Protection Capabilities as Plug-and-Play Services

As described under “enabling services,” the word “services” refers to the protections that a requester should be able to expect in each information exchange, regardless of whether the requester explicitly and knowingly makes such a request for these protections. Modeling data protection capabilities as services will have the following benefits:

- *Loose coupling:* scalable as requirements change;
- *Plug-and-play:* service users do not need to know the implementation details and interoperability is improved;
- *Efficiency:* instead of having every entity create its own services, the entities can use a common set of services; and
- *Effectiveness:* it is easier to enforce if every transaction goes through the same set of services.

Developing data protection capabilities as services also improves future interoperability with other HIEs.

8.1.5 Maintain a Standards-Based, Technology-Neutral, and Vendor-Neutral Architecture

Standards-based, technology-neutral and vendor-neutral characteristics are important for a notional architecture. These characteristics will aid in driving the selection of technical solutions and standards while maintaining forward compatibility as the solutions landscape evolves.

8.2 Architecture Constructs

Architecture constructs, usually originating from various industry standards, identify basic building blocks for a notional architecture. This section lists several important architecture constructs only as illustrative examples.

8.2.1 Security Assertion Markup Language (SAML)

SAML, developed by the Security Services Technical Committee of the Organization for the Advancement of Structured Information Standards (OASIS), is an Extensible Markup Language (XML)-based framework for communicating user authentication, entitlement, and attribute information. SAML allows business entities to make assertions regarding the identity, attributes, and entitlements of a subject (an entity that is often a human user) to other entities, such as a partner company or another enterprise application. SAML is a flexible and extensible protocol designed to be used – and customized if necessary – by other standards. The Liberty Alliance, the Internet2 Shibboleth project, and the OASIS Web Services Security (WS-Security) committee have all adopted SAML as a technological underpinning for various purposes.

For more information on SAML, visit www.oasis-open.org.

8.2.2 eXtensible Access Control Markup Language (XACML)

XACML was ratified as an OASIS standard in February 2003 (1.0 version). XACML defines a generic authorization architecture and the constructs for expressing and exchanging access control policy information using XML. Policy constructs include policies, rules, combining algorithms, etc. XACML complements SAML so that not only policy decisions, as well as the policies themselves, can be exchanged in a standard fashion.

For more information on XACML, visit www.oasis-open.org.

8.2.3 Integrating the Healthcare Enterprise (IHE) Profiles

IHE is a global initiative that creates the framework for passing vital health information seamlessly – from application to application, system to system, and setting to setting – across multiple healthcare enterprises. IHE brings together healthcare information technology stakeholders to implement standards for communicating patient information efficiently throughout and among healthcare enterprises by developing a framework for interoperability.

For more information on IHE profiles, visit www.himss.org.

8.2.4 Web Services Security Standards

Web services security standards represent various specifications defined to implement Web services security. Figure 7 identifies Web services security standards.

Figure 5. Web Service Security Standards

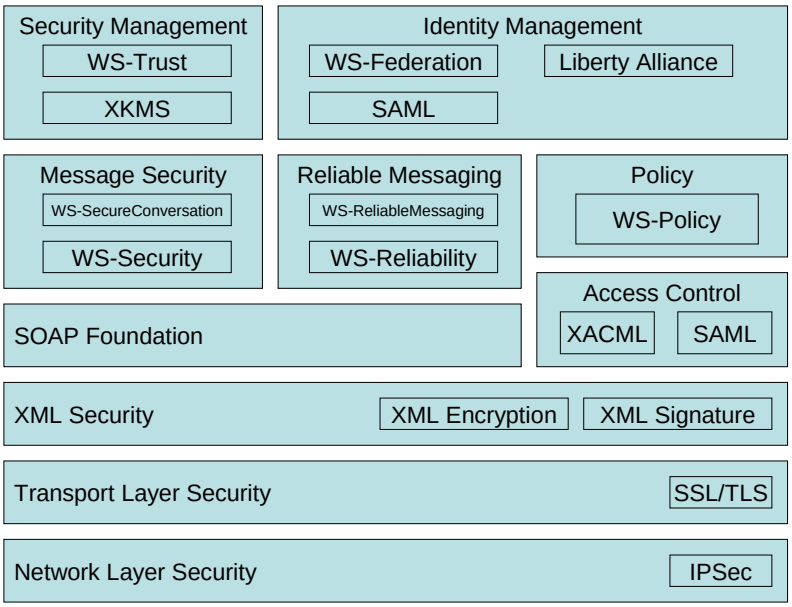

Web service security standards are composable standards. Depending on the notional architecture, an implementation might use only one or two standards from the entire Web services security stack.

For more information on Web services security, visit www.oasis-open.org.

8.2.5 Role-Based Access Control (RBAC)

With role-based access control, access decisions are based on the roles that individual users have as part of an organization. Users take on assigned roles (e.g., doctor, nurse, teller, manager). Access rights are grouped by role name, and the use of resources is restricted to individuals authorized to assume the associated role. For example, within a hospital system, the role of doctor can include operations to perform diagnosis, prescribe medication, and order laboratory tests, and the role of researcher can be limited to gathering anonymous clinical information for studies.

For more information on RBAC, visit csrc.nist.gov.

8.2.6 Attribute-Based Access Control (ABAC)

An attribute-based access control model recognizes that a flexible access control policy should address the evaluation of multiple dimensions of an entity, including identifiers, roles, and qualifications. Since roles can be viewed as nothing more than attributes of principals, RBAC can be wholly absorbed into an attribute-based mechanism.

Attribute-based authorization policies have some distinct advantages over other approaches. First, an attribute-based approach recognizes from its inception that a flexible access control policy cannot be locked into evaluating only one dimension of a principle (such as an identity or role). For example, in order to provide proper controls for accessing sensitive information, it may be necessary to consider various other principal attributes such as doctor qualifications, formal access approvals, or organization affiliation. Second, an attribute-based approach takes into consideration that there are other attributes that are relevant to authorization policies besides those associated with resources or environmental attributes.

9.0 Technology Solutions and Standards – Layer 5

Once the notional architecture is complete, the last phase is to select the technical solutions and data standards that will satisfy the requirements specified in the architecture. Technical solutions and data standards represent the implementation of the notional architecture. Technical solutions and data standards to implement secure and private HIE services are determined based on the notional architecture. The following figure shows illustrative steps taking an organization from notional architecture to the implementation of secure HIE services.

Figure 6. Illustrative Steps from Notional Architecture to Secure HIE Services

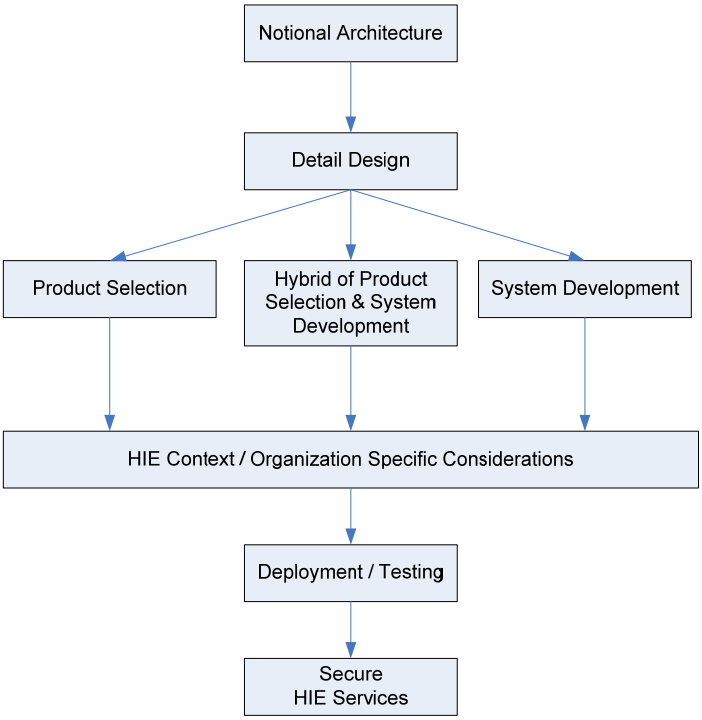

While a notional architecture defines architecture constructs and their relationships, detailed design transforms a notional architecture into detailed implementation specifications that are ready for product selection, system development, or a hybrid of product selection and system development. Many HIE context and organization-specific considerations will need to be evaluated. For example, if most of the HIE participants utilize Java development resources, Java-related technical solutions might be a better choice. Once technical solutions and data standards are selected, they go through deployment and testing cycles to assure that secure HIE services are provided.

10.0 Building a Nationwide HIE using Regional HIEs

As discussed in the previous section, this publication presents a five-layer development operating model for building security architectures for Regional HIEs. If Regional HIEs follow the five-layer operating model, there will be many Regional HIEs using a standard set of data protection services.

Figure 7. Regional HIEs with Standard Enabling Services

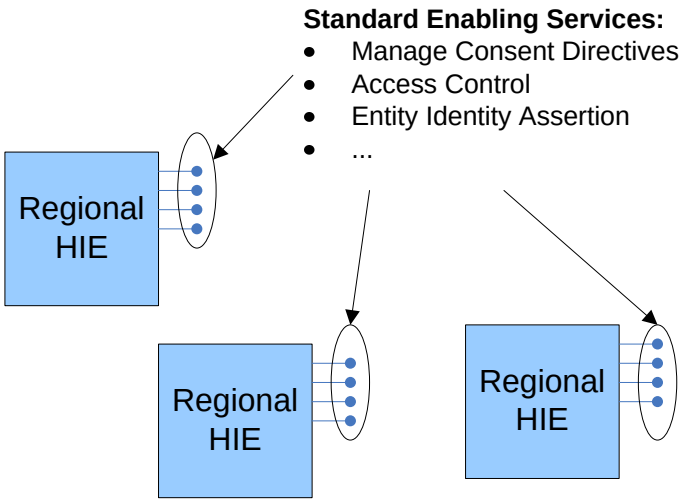

Although it is likely that each Regional HIE might implement the services differently based on its own HIE requirements, having a standard set of services allows for a common understanding of assurance levels that can allow for risk-based interconnection decisions. For example, while HIEs might have different access control policies and implementations, the existence of the common core access control service provides a foundation from which to further evaluate risk.

Using Regional HIEs as the building blocks, Multi-Regional HIEs can be built using a federated architecture as illustrated in the following figure. The federated architecture will centralize certain elements (e.g., trust agreements, assurance levels) while allowing the regional HIE to remain autonomous.

Figure 8. Multi-Regional HIE with Federated Enabling Services

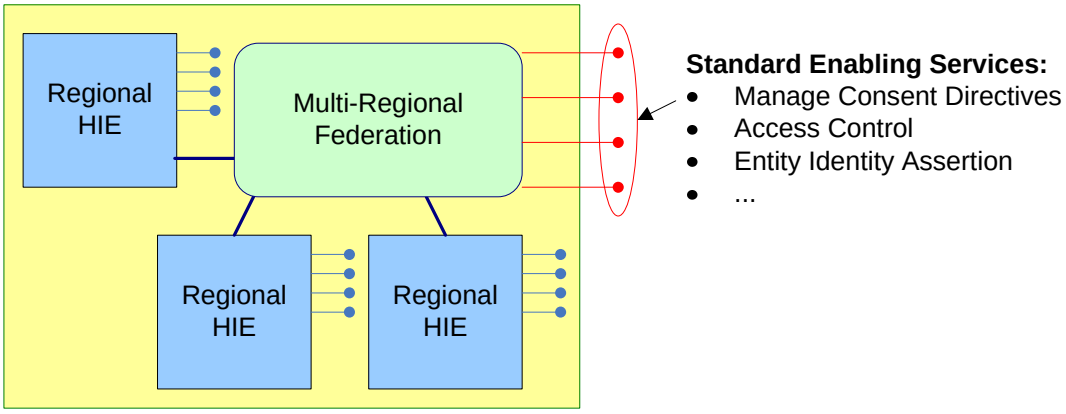

A nationwide HIE can be constructed in a similar way by connecting Multi-Regional HIEs using a federated architecture as illustrated in the following figure.

Figure 9. Nationwide HIE with Federated Data Protection Services

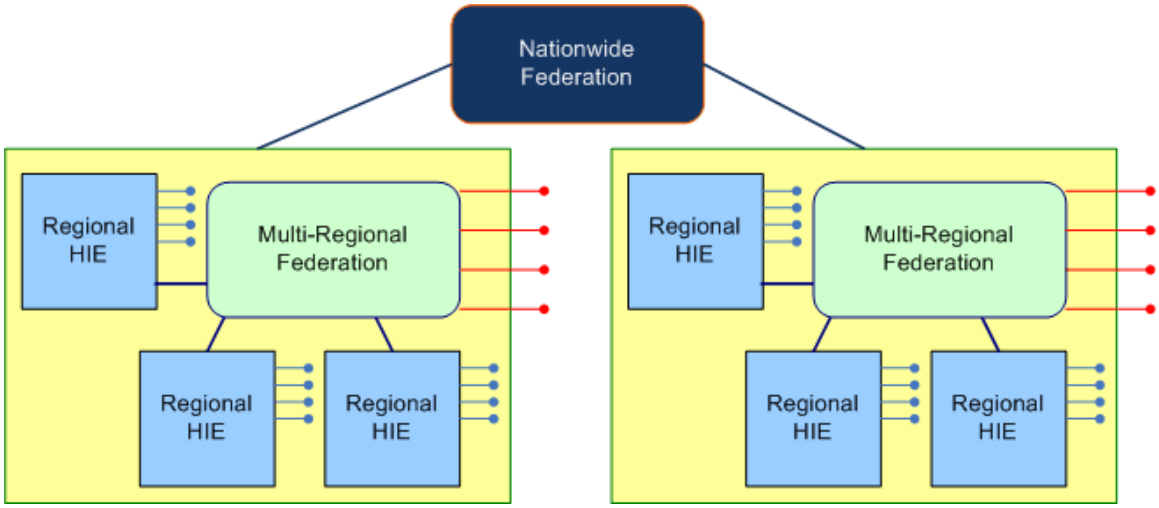

This publication provides a standardized systematic design methodology for developing a security architecture and technical solutions that support a core set of enabling services necessary for the secure exchange of health information. The overarching premise is that if entities engaged in the exchange of health information utilize a standard approach for the selection of security architectures, the probability and the ease of scalability would be dramatically increased.

Appendix A: Applying the Security Architecture Design Process

In this section, the security architecture design process is applied to a specific American Health Information Community (AHIC)[8] Use Case to illustrate the analyses and considerations that need to be made when applying this process to the exchange of health information. This scenario considers issues and data flows surrounding the implementation of information technology to enable the delivery of personalized healthcare. This case is supplied for illustrative purposes only, and may not consider all the complexities, requirements, and interdependencies that could be encountered in particular environments.

The Office of the National Coordinator for Health Information Technology (ONC) developed this use case in two stages. First, a "Prototype Use Case" was developed, which described the data flows of the use case at a high level. AHIC solicited public feedback on the Prototype Use Case in February 2008. Feedback was received and incorporated into the "Detailed Use Case," which comprehensively documents all of the events and actions within the use case at a detailed level.

Figure 10. Clinical Assessment Scenario of 2008 ONC Personalized Healthcare Use Case

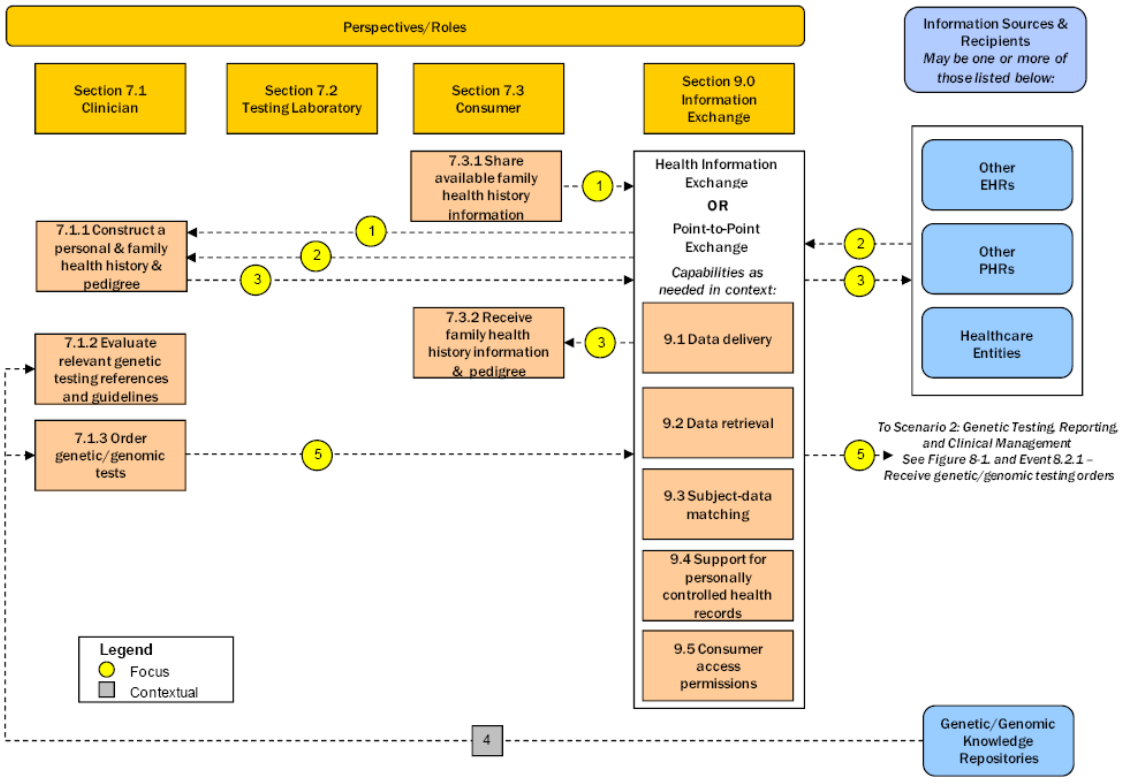

The detailed Personalized Healthcare use case was further broken out into two scenarios. This document uses the Clinical Assessment scenario[9] to develop a similar situation that is specific to the implementation of this architecture design process. Data protection issues are analyzed at each

[8] AHIC is a federally chartered advisory body assembled to make recommendations to the Secretary of the U.S. Department of Health and Human Services on how to accelerate the development and adoption of health information technology; http://www.hhs.gov/healthit/community/background/

[9] For detail definition of each perspectives/roles/actions in the Clinical Assessment scenario, please refer to the 2008 ONC Personalized Healthcare use case: http://www.hhs.gov/healthit/usecases/phc.html.

layer of the operating model. Only one enabling service is further defined in subsequent layers in this document.

A.1 Illustrative Clinical Assessment Scenario

Carol has, in the past, used a Web-based Personalized Health Record (PHR) system to store her personal medical history, including health conditions of her parents and her genetically-related relatives. Carol then begins seeing Dr. Alice. Before seeing Dr. Alice, Carol grants her access to some, but not all, of her PHR.

Dr. Alice retrieves Carol's PHR from the Web-based system based on Carol's authorization. To make a sound clinical assessment, Dr. Alice asks for Carol's permission to request additional information from her previous family doctor, Dr. Bob. She also asks Carol for more information on the health conditions and health history of her parents, and for "read-write-delete" level permissions in order to make appropriate entries or corrections to Carol's PHR. Having received Carol's authorization, Dr. Alice obtains the information and constructs a consolidated view of Carol's personal and family health history. While reviewing Carol's health records, Dr. Alice finds several duplications and eliminates them from the record.

Dr. Alice uses the consolidated information to conduct a clinical assessment, develop a diagnostic plan, and update Carol's PHR. Carol receives an updated PHR with Dr. Alice's inputs.

A.2 Identifying the Health Information Exchanges

Figure 11 is a graphical representation of the illustrative clinical assessment scenario.

Figure 11. Illustrative Clinical Assessment Scenario

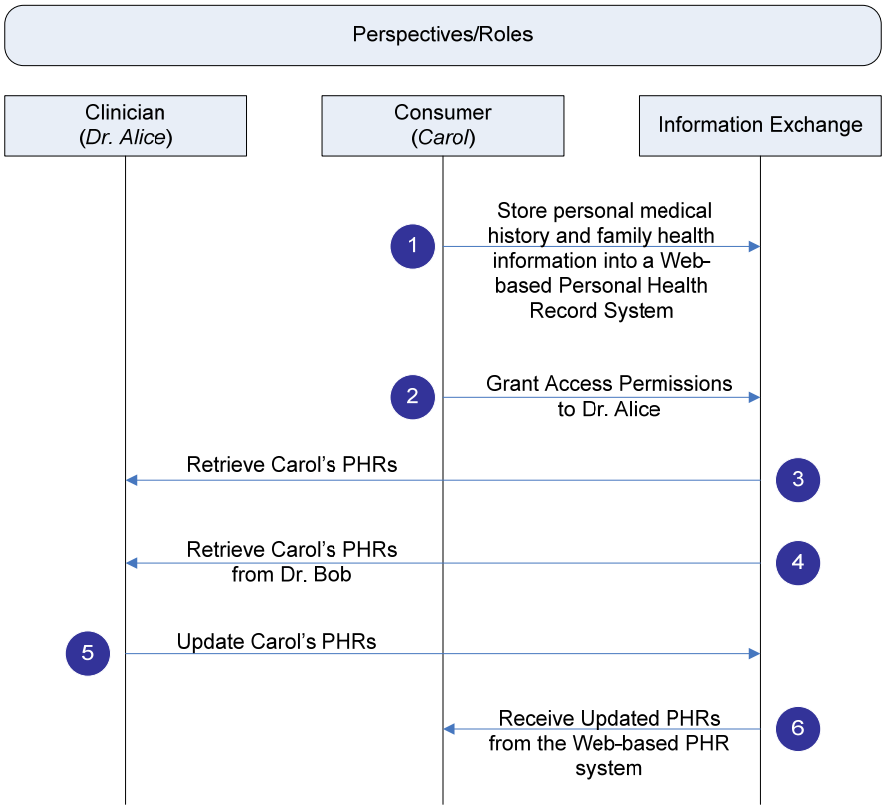

In the illustrative scenario, there are six health information exchanges:

1. Carol stores personal medical history and family health information into a Web-based Personal Health Record (PHR) system;
2. Carol grants access permissions to Dr. Alice;
3. Dr. Alice retrieves Carol's personal medical history and family health information from the Web-based PHR;
4. Dr. Alice retrieves Carol's personal medical information from her previous doctor, Dr. Bob;
5. Dr. Alice updates Carol's personal medical history; and
6. Carol receives updated PHR from the Web-based system.

The HIE security architecture design methodology will be applied to scenario to demonstrate a secure exchange of health information.

A.3 Identify Capstone Policies – Layer 1

Capstone policies include the national regulations, state and local regulations, and organizational policies that apply to the exchange of health information in this illustrative example.

A.3.1 National Regulations

The most significant data protection requirements governing the exchange described are the HIPAA Privacy and Security Rules. Other federal laws and regulations may govern other kinds of healthcare transactions, particularly those that involve exchanges of information with particular government agencies; healthcare research activities; exchanges of particularly sensitive healthcare information, such as information about substance abuse treatment; or sharing information for purposes other than healthcare treatment, payment, or operations (such as law enforcement, public health reporting, or marketing). Given this scenario, however, HIPAA is certainly the most significant capstone policy driver.

A.3.2 State and Local Regulations

Depending on the state and jurisdiction, other rules may govern the use, disclosure, or security of protected health information (PHI). For example, a majority of states require entities conducting business in the state to provide notice to all affected individuals in the event of a breach or loss of private data, including PHI.

In cases where state law conflicts with HIPAA, HIPAA explicitly allows state law to take precedence over HIPAA if the state law is "more stringent." That is, the state law supplies even greater protections to an individual's privacy, or requires additional or stronger security protections than those required by HIPAA.

In the current case, it is assumed that relevant state law:

- Requires the disclosure of PHI to any healthcare provider at the patient's written request. HIPAA merely allows covered entities to share this information with other healthcare providers for payment, treatment, or operations purposes. As disclosure is compelled under this hypothetical state law, it is "more stringent" and must be followed.
- Forbids the disclosure of PHI to a patient's otherwise authorized representative if, in the judgment of the healthcare provider, releasing that information could cause the patient harm. Many states have a provision such as this one such that healthcare providers would not, for example, be obliged to disclose information about mental or physical abuse to the patient's

possible abuser. As this measure allows the patient even more protection than HIPAA explicitly allows, it is "more stringent" and should be followed.

- Requires the provider to disclose to the patient or a patient's representative or guardian if there is a known or suspected breach of the patient's unencrypted information. As this notification is not explicitly required by HIPAA, it is "more stringent" and must be followed.

A.3.3 Organizational Policies

Further requirements may be imposed by the institutions at which Dr. Alice and Dr. Bob practice. While HIPAA sets parameters for data protections, its standards often require the institution to implement their own reasonable policies in certain areas. Other individual institutional rules may be driven by other laws, such as the Common Rule for Human Subjects Research; institutional accreditation standards; contractual obligations with business partners; or best practices.

This scenario assumes that certain appropriate rules apply to the Web-based PHR system used by Carol. Rules will be proposed only to the extent necessary to address one enabling service, Entity Identity Assertion. These rules are not intended to be complete, and no assertion is made as to their adequacy for any real-world entity or environment.

This scenario assumes that the following institutional (corporate) rules apply for Carol's access to her PHR:

- Carol, and anyone to whom she grants access to her account, must log in using a username and ID.
- Passwords must have a minimum "strength," as described below.
- Carol, and anyone to whom she grants access to her account, must use a digital certificate to access her account
- Carol, and anyone to whom she grants access to her account, must use a hardware token to assert their identity.
- Carol has unrestricted access to her own PHR.
- Carol has unlimited privileges to grant access and privileges to others, including privileges to read, write, edit, or delete her account.

In addition, other institutional-level restrictions may apply to Dr. Alice's and Dr. Bob's institutions. Dr. Alice and Dr. Bob may have to log on to their accounts using separate identity assertion controls.

A.4 Identify Enabling Services - Layer 2

The Security and Privacy Operating Model identified the 12 enabling services that every HIE should consider. Based on the six identified HIEs, the following table lists enabling services that should be used in each HIE:

Table 5. Enabling Services for each HIE

HIE #	HIE Description	Enabling Services	Enabling Services Description
1	Carol stores personal medical history and family health information into a Web-based Personal Health Record (PHR) system	Risk Assessment	Risk assessment is used to analyze the business risks of compromising the security and privacy of the health information exchanged.
		Entity Identity Assertion	The Web-based PHR requires Carol to identify herself using a registered credential every time she logs in.
		Access Control	The Web-based PHR grants access permissions based on privileges an

HIE #	HIE Description	Enabling Services	Enabling Services Description
			authenticated individual has.
		Credential Management	The Web-based PHR that Carol selects will require Carol to use certain types of credentials to register.
		Privilege Management	Carol has full access permissions to her PHR and she can assign access permissions to her doctors.
		Audit Trail	All accesses to Carol's Web-based PHR will be logged. Suspicious accesses will trigger warning messages that will be sent to Carol.
		Secure Communication Channel	All information transmitted is secured between Carol's terminal and the Web-based PHR.
2	Carol grants access permissions to Dr. Alice	Entity Identity Assertion	The Web-based PHR requires Carol to identify herself using registered credentials every time she logs in.
		Access Control	The Web-based PHR allows Carol to change access permissions associated with her PHR.
		Privilege Management	Carol assigns access permissions on certain portions of her personal medical history to Dr. Alice.
		Audit Trail	The Web-based PHR maintains a record of Carol's action of assigning access permissions to Dr. Alice.
		Secure Communication Channel	All information transmitted is secured between Carol's terminal and the Web-based PHR.
3	Dr. Alice retrieves Carol's personal medical history and family health information from the Web-based PHR	Entity Identity Assertion	The Web-based PHR requires Dr. Alice to identify herself with a registered credential.
		Access Control	The Web-based PHR allows Dr. Alice to read the portion of Carol's PHR to which she has access permission.
		Audit Trail	Dr. Alice's access to Carol's PHR is logged.
		Secure Communication Channel	All information transmitted is secured between Carol's terminal and the Web-based PHR.
4	Dr. Alice retrieves Carol's personal medical information from Dr. Bob	Entity Identity Assertion	Dr. Bob requires Dr. Alice to identify herself with registered credentials. Dr. Alice requires Dr. Bob to identify himself with registered credentials.
		Manage Consent Directives	Dr. Bob obtains Carol's consent to share Carol's personal medical information.
		Secure Communication Channel	All information transmitted is secured between Dr. Alice and Dr. Bob.
5	Dr. Alice updates Carol's personal medical history	Entity Identity Assertion	The Web-based PHR requires Dr. Alice to identify herself with registered credentials.
		Access Control	The Web-based PHR allows Dr. Alice to read and update the portion of Carol's PHR to which she has access permission.
		Audit Trail	Dr. Alice's access to Carol's PHR is logged.
		Secure Communication Channel	All information transmitted is secured between Dr. Alice's terminal and the Web-based PHR.
6	Carol receives updated PHRs from the Web-based	Entity Identity Assertion	The Web-based PHR requires Carol to identify herself using registered credentials every time she logs in.

HIE #	HIE Description	Enabling Services	Enabling Services Description
	PHR system	Access Control	The Web-based PHR allows Carol to read her personal medical history.
		Audit Trail	Carol's access to her PHR will be logged.
		Secure Communication Channel	All information transmitted is secured between Carol's terminal and the Web-based PHR.

The identified enabling services supporting the health information exchanges need to be in place to enable the information exchanges among Dr. Alice, Carol, Dr. Bob and the Web-based PHR system. Implementation of one enabling service, Entity Identity Assertion, is addressed in the following sections.

A.5 Develop Enabling Processes – Layer 3

The Entity Identity Assertion service of the Web-based PHR system has the following requirements based on the HIE activities indicated in the illustrative scenario above.

- The system shall accept three types of credentials to authenticate users (including service providers, consumers, and any others):
  - User created ID with Strong password;
  - Digital certificates; and
  - Hardware tokens.
- The system shall authenticate every transaction.
- The system shall accept credentials (any of the three types) issued from trusted third parties.

A.5.1 Authentication Credentials

The system has defined the following processes on how three types of credentials can be accepted:

Table 6. Processes for Credential Acceptance

Credentials	Processes
User ID and Password	• Passwords must be stored in irreversible encrypted form and the password file cannot be viewed in unencrypted form. • A password must not be displayed on the data entry/display device. • Passwords must be at least eight characters long. • Passwords must be composed of at least three of the following: English uppercase letters, English lowercase letters, numeric characters, and special characters. • Password lifetime will not exceed 60 days. • Users cannot use the previous six passwords. • The system will give the user a choice of alternative passwords from which to chose. • Passwords must be changed by the user after initial logon.
Digital Certificates	• The certificate must be an X.509v3 certificate. • The certificate must be within the valid period. • The certificate must be verified and validated through authentication. • The system will not issue digital certificates. Users will present trusted third party issued certificates that are valid and verifiable by the system.
Hardware Tokens	• The system will accept and support pre-approved types of hardware tokens as authentication credentials.

A.5.2 Accepting Trusted Third-Party-Issued Credentials

The system defines its processes and policies of accepting third-party authentication credentials as follows:

Table 7. Acceptance of Third-Party Authentication Credentials

Credentials	Processes
User ID and Password	• A trusted third party must comply with the system's User ID and Password rules (e.g., minimum password strength requirements must be met). • The system shall accept authentication claims from a third party authentication authority. • The third party authentication claim shall comply with the system's profile for authentication claims.
Digital Certificates	• Since the system will not issue digital certificates, all certificates will be issued by trusted third parties. • The system shall only accept digital certificates that issued by authorities comply with the system's X.509 profile. • The system's X.509 profiles defines requirements to be a trusted certificate authority and the certificate validation process.
Hardware Tokens	• User can only request hardware tokens from the system. No third party hardware tokens will be accepted.

As illustrated above, the enabling processes further refine the Entity Identity Assertion service for the Web-based PHR system. The enabling processes will vary for different HIEs. These processes will translate into part of the governance policies which could be part of the trust agreement between HIE entities.

A.6 Develop Notional Architecture – Layer 4

Based on the defined enabling processes, a notional architecture can be developed for the Entity Identity Assertion service. This notional architecture is a standards-based, platform-independent and vendor-neutral implementation blueprint for enabling services and processes, and it will drive the selection of technical solutions.

Figure 13 provides an illustrative example of the notional architecture for the Web-based PHR system's Entity Identity Assertion Service.

Figure 12. Illustrative Notional Architecture for Entity Identity Assertion Service

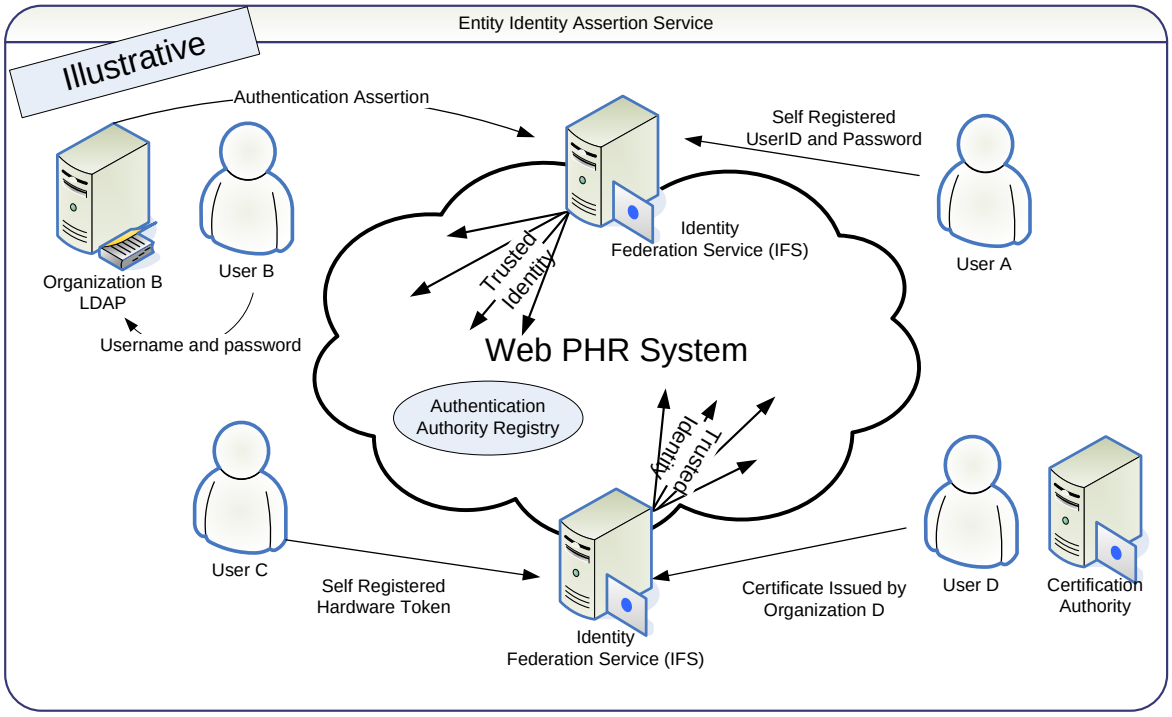

In the notional architecture, four different scenarios are described:

- Self-Registered UserID and Password (User A): Users register themselves with the Web-based PHR system by creating a User ID and password. Users must go through the identity proofing process defined by Web-based PHR system.
- Third-Party UserID and Password (User B): Users are redirected to their home organizations to perform the authentication. The home organization's (Organization B) authentication authority (e.g., LDAP) will issue a SAML assertion to the Web-based PHR system as the authentication credential.
- Hardware Token (User C): Users who request hardware tokens from Web-based PHR system can use the issued token as the authentication credential.
- Third-Party Certificates (User D): Users use third-party issued certificates as authentication credentials.

The Identity Federation Service (IFS) will serve as the authentication portal which accepts all types of credentials and creates a trusted identity for authenticated users into the Web-based PHR System. The trusted identity could be a digital certificate or a SAML assertion.

A.7 Select Technical Solutions – Layer 5

Once the notional architecture is determined, technical solutions can be selected and deployed to implement the architecture. An illustrative example of the deployment of possible technical solutions is depicted in Figure 14.

Figure 13. Illustrative Technical Solutions for Entity Identity Assertion Service

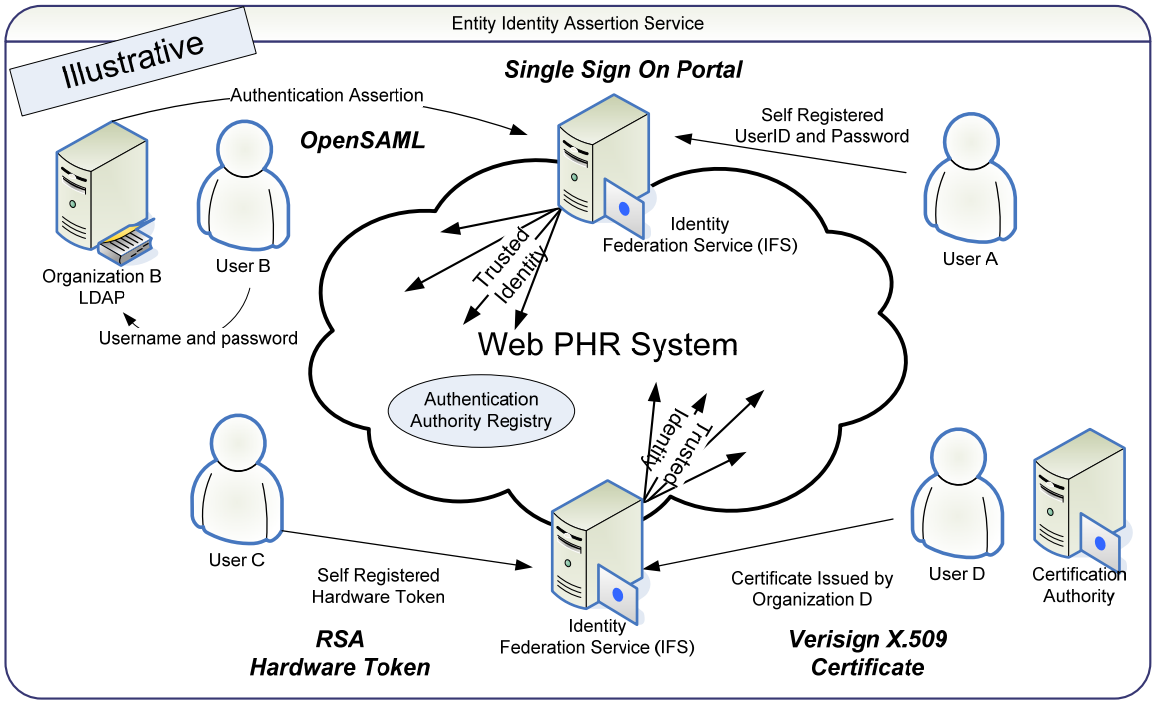

A.8 Considerations for Health Information Exchange

Following a similar process as illustrated above, an organization can implement all enabling services necessary to facilitate the secure and private exchange of health information in this scenario. The services described in this architecture design methodology focus only on the exchange portion; they do not focus on those services necessary to implement security and privacy within the involved organizations. To provide end-to-end protection of health information, the involved organizations need to implement relevant services that provide adequate protection for the information outside the bounds of exchange.

The actual data exchange -- that is, the act of transmitting and receiving the health information--is a point of particular vulnerability to the security and privacy of consumer information in the health information exchange, because it is usually done outside of the participating organizations' security boundaries. However, to ensure that a consumer's information is adequately protected, the "non-exchange" portions of the data usage, including collection, storage, modification, and destruction, must also receive security and privacy protections, which may include disaster recovery and contingency planning, configuration management, and other processes and technologies.

Appendix B: Acronyms

ABAC	Attribute-Based Access Control
AC	Access Control
AHIC	American Health Information Community
AU	Audit and Accountability
CCHIT	Certification Commission for Health Information Technology
CSRC	Computer Security Resource Center
EPHI	Electronic Protected Health Information
FISMA	Federal Information Security Management Act
EHR	Electronic Health Record
HHS	Health and Human Services
HIE	Health Information Exchange
HIPAA	Health Insurance Portability and Accountability Act
HIT	Health Information Technology
HITSP	Health Information Technology Standards Panel
IA	Identification and Authentication
ID	Identity
IHE	Integrating the Healthcare Enterprise
IPSec	Internet Protocol Security
IT	Information Technology
LDAP	Lightweight Directory Access Protocol
NHIN	Nationwide Health Information Network
NIST	National Institute of Standards and Technology
NISTIR	National Institute of Standards and Technology Interagency Report
NOPP	Notice of Privacy Practices
OASIS	Organization for the Advancement of Structured Information Standards
ONC	Office of the National Coordinator
PHI	Protected Health Information
PHR	Personal Health Record
PIN	Personal Identification Number
PKI	Public Key Infrastructure
RA	Risk Assessment
RBAC	Role-Based Access Control
SAML	Security Assertion Markup Language
SC	Systems and Communications (NIST SP 800-53 Security Control Family)
SP	Special Publication
TLS	Transport Layer Security
VPN	Virtual Private Network
WS	Web Services
XACML	Extensible Access Control Markup Language
XML	Extensible Markup Language

Appendix C: References

Public Laws

Public Law 107-347, E-Government Act of 2002 (Title III: Federal Information Security Management Act [FISMA] of 2002), December 17, 2002.

Public Law 104-191, Health Insurance Portability and Accountability Act (HIPAA) of 1996, August 21, 1996.

Federal Regulations

Health Insurance Reform: Security Standards; Final Rule ("The HIPAA Security Rule"), 68 FR 8334, February 20, 2003.

Standards for Privacy of Individually Identifiable Health Information, Final Rule ("The HIPAA Privacy Rule"), 45 CFR Parts 160, 162, and 164.

Federal Information Processing Standards (FIPS) Publications

FIPS 140-3, *DRAFT Security Requirements for Cryptographic Modules,* July 2007.

NIST Special Publications (SPs)

NIST SP 800-25, *Federal Agency Use of Public Key Technology for Digital Signatures and Authentication*, October 2000.

NIST SP 800-30, *Risk Management Guide for Information Technology Systems*, January 2002.

NIST SP 800-39, *DRAFT Managing Risk from Information Systems: An Organizational Perspective*, April 2008.

NIST SP 800-45, Version 2, *Guidelines on Electronic Mail Security*, February 2007.

NIST SP 800-52, *Guidelines for the Selection and Use of Transport Layer Security (TLS) Implementations*, June 2005.

NIST SP 800-53, Revision 2, *Recommended Security Controls for Federal Information Systems*, December 2007.

NIST SP 800-57, *Recommendation for Key Management*, March 2007.

NIST SP 800-58, *Security Considerations for Voice over IP Systems*, January 2005.

NIST SP 800-63-1, *DRAFT Electronic Authentication Guide*, December 2008.

NIST SP 800-66, *An Introductory Resource Guide for Implementing the Health Insurance Portability and Accountability Act (HIPAA) Security Rule*, October 2008.

NIST SP 800-77, *Guide to IPsec VPNs*, December 2005.

NIST SP 800-92, *Guide to Computer Security Log Management*, September 2006.

NIST SP 800-106, *Draft Randomized Hashing Digital Signatures*, July 2008.

NIST SP 800-111, *Guide to Storage Encryption Technologies for End User Devices*, November 2007.

NIST SP 800-113, *Guide to SSL VPNs*, July 2008.

Web sites and Other Resources

Certification Commission for Healthcare Information Technology (CCHIT): www.cchit.org

Department of Health and Human Services (DHHS): www.hhs.gov

Healthcare Information and Management Systems Society (HIMSS): www.himss.org

Healthcare Information Technology Standards Panel (HITSP): www.hitsp.org

NIST Computer Security Resource Center (CSRC): http://csrc.nist.gov/

Organization for the Advancement of Structured Information Standards (OASIS): www.oasis-open.org

Office of the National Coordinator for Health Information Technology (ONC): www.hhs.gov/healthit/onc/mission/