Nationwide Health Information Network (NHIN) Exchange

Architecture Overview

DRAFT v.0.9
4/21/2010

The Nationwide Health Information Network (NHIN) comprises standards, services and a trust fabric that enables the secure exchange of health information over the Internet. This critical part of the national health IT agenda will enable health information to follow the consumer, be available for clinical decision-making and support appropriate use of healthcare information beyond direct patient care to improve population health.

One instance of the NHIN standards, services and trust fabric has been in pilot testing through the NHIN cooperative and is now ready for a limited production pilot to a broader community of Federally sponsored entities. The NHIN Exchange includes the robust technology and trust fabric necessary to support health information exchange among large nationwide organizations and federal agencies. **The information presented in this draft Architectural Overview document pertains only to the NHIN Exchange.** During the course of 2010, technical and policy initiatives will expand the value of NHIN standards, services and trust fabric and extend the ability to securely exchange health information to a larger audience.

Table of Contents

1 Introduction

This document is intended to provide an architectural overview of the Nationwide Health Information Network (NHIN). It is not intended to replace the individual NHIN specifications and documents available publicly at http://healthit.hhs.gov/nhin. Rather, it offers a contextual synopsis of how those pieces of the NHIN puzzle fit together, in order to facilitate wide participation in health information exchange.

This document is intended to help technical architects, developers and project managers to understand:

- The problems that the NHIN seeks to address
- The underlying principles that have informed the development of the NHIN architecture and the corresponding specifications
- The overall architecture framework, including the scope of the architecture and the individual "layers" that describe specific capabilities of the network
- How the NHIN enables health organizations to implement several common scenarios for information exchange
- The underlying logic and important details about many service specifications
- The role the NHIN plays in the larger national health IT agenda, and its relationship to other standards and governance organizations

Although this is the first release of the document, we are already thinking of ways in which it can be expanded to further meet future needs. In short order, this document and others like it will be hosted on a NHIN wiki in order to maximize feedback and collaboration.

2 Defining the NHIN

This section of the document describes the NHIN and provides context for its role within the broader national healthcare agenda.

2.1 Introduction to the NHIN

The Nationwide Health Information Network (NHIN) is a collection of standards, protocols, legal agreements, specifications, and services that enables the secure exchange of health information over the internet. The NHIN is a key component of the nationwide health information technology strategy and will provide a common platform for health information exchange across diverse entities, within communities and across the country, helping to achieve the goals of the Health Information Technology for Economic and Clinical Health (HITECH) Act. This critical part of the national health IT agenda will enable health information to follow the consumer, be available for clinical decision making, and support appropriate use of healthcare information beyond direct patient care, such as to improve public health.

The NHIN will enter production beginning in Q1 2010. Because several mid to large health information organizations (HIOs), including Federal agencies, have already committed to implementing NHIN technical and policy infrastructure for their multi-point information exchange needs, the NHIN will initially operate as a network of networks. HIOs joining the NHIN are termed NHIN Nodes (further description of NHIN Nodes can be found in section 3.4 “NHIN Nodes”). During this period, smaller providers seeking to use the NHIN will do so through association with a NHIN Node. The NHIN will evolve as changes in technology, users, the purposes users seek to achieve, and the functionalities users will require to accomplish these purposes.

2.2 Rationale for the NHIN

Healthcare IT (HIT) plays an important role in healthcare improvement and reform. The national agenda for HIT is twofold: increase adoption of Electronic Health Records (EHRs) and build a framework that enables these EHRs to be sharable and interoperable.

The HITECH Act calls for the Office of the National Coordinator for Health Information Technology (ONC) to develop “a nationwide health information technology infrastructure that allows for the electronic use and exchange of information and that…promotes a more effective marketplace, greater competition…[and] increased consumer choice” among other goals. (Section 3001(b)). The NHIN is a critical part of that technology infrastructure and plays an important role in the National health agenda.

The role of the NHIN is to provide means by which health and healthcare entities are able to securely exchange interoperable health information. In order to serve in this capacity, NHIN has specified a set of standards-based Web Service Interfaces which will allow disparate systems to securely communicate with each other over the Internet.

2.3 Current State

The NHIN consists of very minimal infrastructure (an UDDI Web Services Registry and the NHIN Security Infrastructure, based on a managed PKI Service) and a relatively few number of Web Service Interfaces and foundational specifications. The NHIN begins to take shape only when its specifications are implemented in physical systems and used to support business functions. NHIN specifies the use of standardized Internet protocols and Health IT standards to allow NHIN Nodes, which initially will include private and public, academic, business and government networks, to securely exchange health information. The “network of networks” terminology used to describe the NHIN is not intended to imply that the NHIN consists only of Health Information Exchanges or Integrated Delivery Networks. A NHIN Node may be any type of Health Information Organization – large or small, and may soon include Health Information service Providers

Throughout 2009, the NHIN was in limited production as a handful of NHIN Nodes, then termed NHIN Participants, used the services developed for the 2008 Trial Implementations in order to exchange live data in various pilots and projects. During this time, the NHIN Specifications Factory worked closely with HITSP and a number of key Standards Defining Organizations (SDOs) to advance the Trial Implementation service specifications in preparation for production. In January 2010, ONC published the initial set of NHIN Exchange specifications. The future direction of the NHIN is currently being shaped by the activities of the NHIN Work Group, which falls under the ONC HIT Policy Committee. The NHIN Work Group is charged with offering recommendations to drive the secure and standards-based exchange of health information in alignment with national healthcare priorities.

3 NHIN Architectural Concepts

This section introduces key principles and concepts which characterize NHIN architecture. Later sections of the document will discuss these principles and concepts in further detail.

3.1 The Network Approach

At its most fundamental level the NHIN is a network. Networks may be modeled as graphs of nodes and the links between them. In the context of the NHIN, a node is an HIO that participates in the exchange of health information with other nodes on the NHIN via a NHIN Gateway. A NHIN Gateway is an instantiation (implementation) of the NHIN technical specifications which supports secure, interoperable health information exchange across the NHIN. CONNECT, developed by the Federal Health Architecture (FHA), is one example of a NHIN Gateway. Implementation of a NHIN Gateway is only one aspect of exchanging information via the NHIN; other requirements are described separately in NHIN On-Boarding Process documentation.

Regardless of its internal structure, an HIO's implementation of a NHIN Gateway enables each NHIN Node to maintain autonomy inside their domain, while adhering to NHIN specifications for inter-node communication. This flexibility is achieved by the set of architectural principles, described in section 3.2, which inform the NHIN's design.

Figure 1: The Network Approach

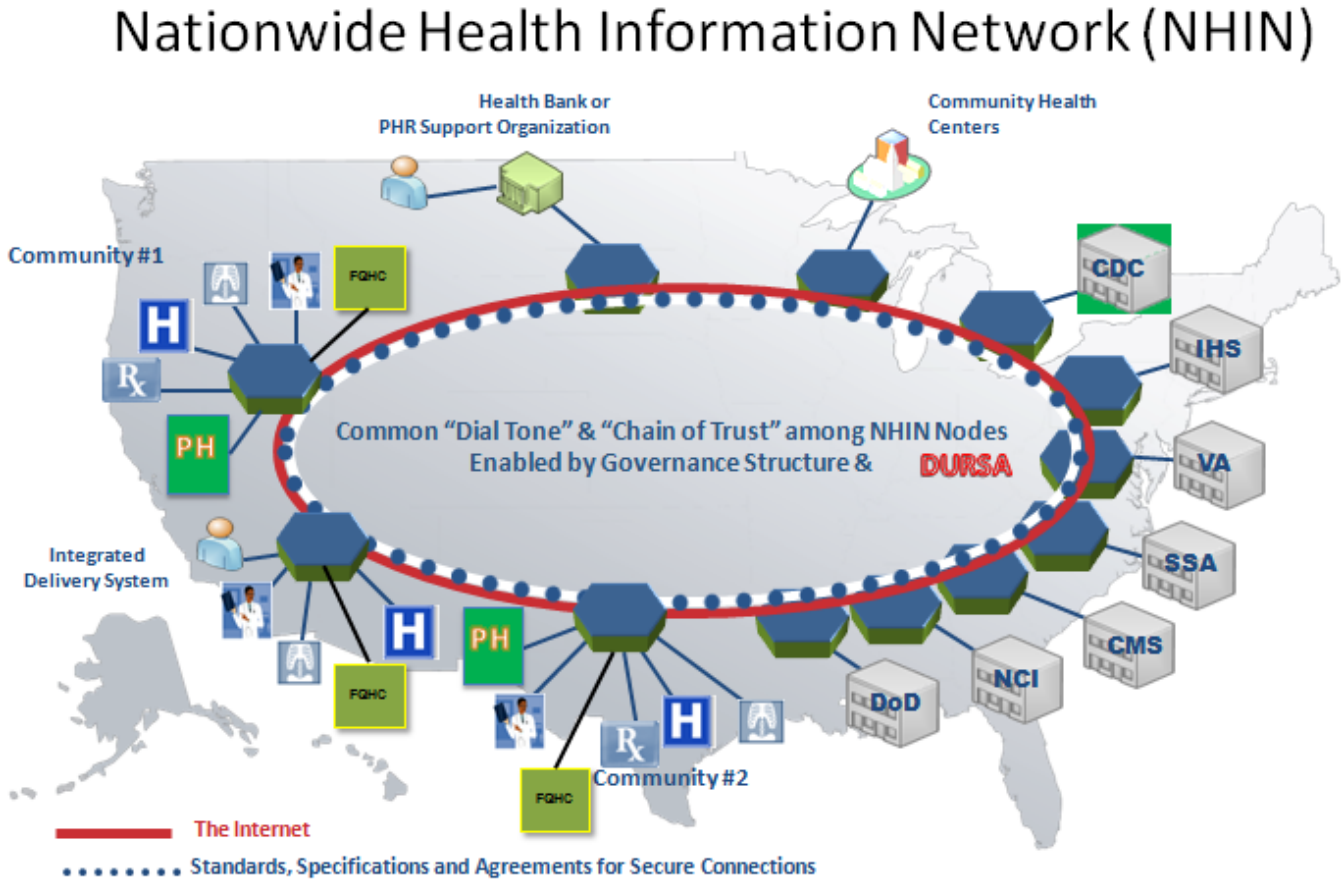

3.2 Architectural Principles

1) **De-centralization:** The NHIN architecture is highly decentralized; there is neither a central Master Patient Index (MPI) nor Central Data Repository (CDR). Patient information is retained locally by the NHIN Nodes.

2) **Local Autonomy** – Acknowledges that the decision to release information from one NHIN Node to another is a local decision, governed by Federal and State regulations and local policies and permissions. Given this principle, NHIN transactions must include enough information about the originating NHIN Node (requestor/sender depending on whether it is a push or pull transaction) for the target NHIN Node to make a decision about whether to participate in the information exchange.

3) **Local Accountability** - Each NHIN Node is accountable for the accuracy and truth of the information it provides to assist the decision making process, as embodied by the local autonomy principle.

4) **Adherence to standards:** The NHIN has taken the initiative to adopt a series of harmonized standards which have been developed by voluntary consensus standards bodies for exchange of health information among all such entities and networks.

5) **Service-Oriented, Layered Architecture:** There is a common messaging, security and privacy foundation which supports the NHIN discovery and information exchange services.

    - Cross-platform integration - Messages are the “universal translators” between different platforms and languages and permit each system to work with their native data types.
    - Reliable communication - Messages can use a “store-and-forward” style for delivery.
    - End-to-end security - Messages can transfer the complete security context using a combination of headers and tokens which increases the ability to improve control over the *authentication* of the personal identity of system users and the *authorization* of specific system users for the privileges to use specific IT system functions.

6) **Utilizes Web Services:** Web Services provide the basis for transport, discovery and exchange capabilities.

    - **Standard protocol:** Functionality is exposed via web services interfaces.
    - **Web service description**: This description is provided via an XML document called a Web Services Definition Language (WSDL) document.
    - **Finding web services:** The discovery capabilities are provided by a listing of web services implemented via the NHIN Web Services Registry. For more information regarding the registry, see section 4.3.1 “NHIN Operational Infrastructure Components”.

7) **Specifications Driven:** All web services have been designed based on specifications agreed upon by the NHIN Technical Committee. The current specifications include:

    - Authorization Framework
    - Messaging Platform
    - Patient Discovery
    - Retrieve Documents
    - Query for Documents
    - Health Information Event Messaging (HIEM)
    - Document Submission
    - Access Consent Policies

- Geocoded Interoperable Population Summary Exchange (GIPSE) Profile
- CARE (Continuity Assessment Record and Evaluation) Profile

8) Utilizes **public key infrastructure** as the basis for security.

3.3 Architectural Approach

3.3.1 NHIN Architecture Requirements

Requirements which drive the NHIN architecture include:

- Ability to discover and exchange healthcare information amongst participant entities
- Ability to match patients to their data without a universal or national patient identifier
- Ability to support patient preferences regarding their data exchange
- Support secure data exchange
- Support harmonized standards
- Support diverse set of organizations, technologies, and approaches
- Support a common trust agreement

Since the 2008 Trial Implementations, an additional significant requirement for the NHIN is to support the concept of "Meaningful Use". The American Recovery and Reinvestment Act of 2009 (ARRA) authorizes the Centers for Medicare & Medicaid Services (CMS) to provide reimbursement incentives for eligible professionals and hospitals who are successful in becoming "meaningful users" of certified electronic health record (EHR) technology. The NHIN Direct Project is working to expand current NHIN services to support Stage 1 meaningful use requirements.

3.3.2 NHIN Specification Development Process

NHIN specifications, which prescribe the technical and security requirements necessary to support information exchange among NHIN Nodes, are developed through a collaborative, multi-stakeholder process. This development work is carried out by the NHIN Specifications Factory, which is an ONC-facilitated work group comprised of representatives from a variety of public and private NHIN stakeholders. NHIN specification development efforts are guided by the NHIN governance structure, which includes oversight by the NHIN Technical Committee and input from the NHIN Coordinating Committee.

Updates or enhancements to NHIN functionality are driven by development requests submitted by NHIN stakeholders. The NTC prioritizes these requests and votes to approve any new or updated NHIN specifications for use in NHIN production. More detailed descriptions of the Specifications Factory, the NHIN Coordinating Committee, and the NTC can be found in section 6 – "Relationships between the NHIN and Other Efforts/Organizations".

3.4 NHIN Nodes

The NHIN is comprised of all of its interconnected nodes. NHIN Nodes currently consist of various Health Information Organizations who have implemented a NHIN Gateway and who have completed the NHIN On-Boarding Process. While initially, the NHIN will serve mid to large HIOs because of their early commitment to NHIN implementation for their information exchange needs; from a purely technical perspective, a NHIN Node can range from the simple to the complex:

- A single computer with an application that implements the core set of standard NHIN services and content could potentially be considered a node on the NHIN.
- A server that implements the core set of standard NHIN services and content and provides those services to the clients it serves may be considered a node on the NHIN.

- A network that has a component that implements the core set of standard NHIN services and content and provides those services to the other members in its network may be considered a node on the NHIN.
- Finally, a hierarchy of networks, with at least one component in one of the networks implementing the core standard NHIN services and content and providing those services to the other members of said hierarchy of networks, may be considered a node on the NHIN.

NHIN Nodes are autonomous. NHIN specifications apply at the communication boundaries between each NHIN Node, with the NHIN defining the standards and web services to be used to communicate with other NHIN Nodes.

The types of HIOs initially exchanging information via the NHIN and those envisioned to do so in the future include:

- Care Delivery Organizations (CDOs) that use Electronic Health Records (EHRs)
- Organizations that operate Personal Health Records (PHRs) and other consumer applications
- Organizations known as Health Information Exchanges (HIEs) that enable health related data exchange between state, regional or non-jurisdictional participant groups;
- Other participant organizations that operate for specific purposes, including secondary users of data such as public health, research, and quality assessment. Based on the specialized nature or purpose of use of these organizations, they may only require a subset of the shared architecture, processes, and procedures

3.5 DURSA Architectural Implications

The organizations currently committed to exchanging information via the NHIN have determined that they require a robust trust fabric, reflected in the Data Use and Reciprocal Support Agreement (DURSA).
The DURSA is a comprehensive, multi-party trust legal agreement and is based upon a set of policy assumptions that bridge varying state and federal laws and regulations, as well as various policies. This legal contract, signed by all entities currently exchanging information via the NHIN, provides a framework of trust assurance to support multi-point health information exchange across the NHIN. The DURSA signators agree to be governed by its provisions, and require other parties who wish to exchange information with them to also sign the DURSA.

As part of the NHIN's privacy, security, and messaging foundation, the NHIN Authorization Framework specification supports a multi-level approach to addressing security requirements in which NHIN Nodes may exchange (or reference) transaction-specific agreements, if that requirement is established by ONC. More details regarding the Authorization Framework may be found in section 4.3.2.2.

4 The NHIN Architectural Framework

4.1 Introduction

This section provides an overview of the architectural framework. It begins with an explanation of the "network of nodes" concept which has been used to describe the NHIN and then examines the various architectural infrastructure components.

NHIN has developed a collection of specifications which define how each Health Information Organization (HIO) uses a NHIN Gateway to communicate health information across organizations. The NHIN Gateway concept is critical to the effectiveness of the NHIN, because it encapsulates all the complexity of inter-organizational exchange (authentication, authorization, auditing, policy management, etc.) in one package. This document introduces the NHIN Gateway as the entity which encapsulates the NHIN interactions and defines the interactions between NHIN Nodes. In addition to interacting with each other, NHIN Gateways also interact with the NHIN Web Services Registry and NHIN Security Infrastructure.

4.2 The NHIN is a Network

At its most fundamental level, and as the acronym calls out, the NHIN is a network. The phrase "network of networks" which has been used to describe the NHIN is not intended to imply that nodes on the NHIN consist only of networks such as Health Information Exchanges or Integrated Delivery Networks. A NHIN Node may be any type of Health Information Organization – large or small, as discussed in section 3.4 "NHIN Nodes"

4.2.1 NHIN Network Zones

One way to illustrate the network concept is to view the NHIN as a series of software "zones", the HIO Zone, the NHIN Zone, and the NHIN Infrastructure Zone. Zones define the actions, relationships and responsibilities of NHIN Nodes, and the network itself. Each zone contains a set of architectural components, which execute based on the actions, relationships and responsibilities of the particular zone in which they exist. This concept can be seen in the figure below.

Figure 2: NHIN Network Zones

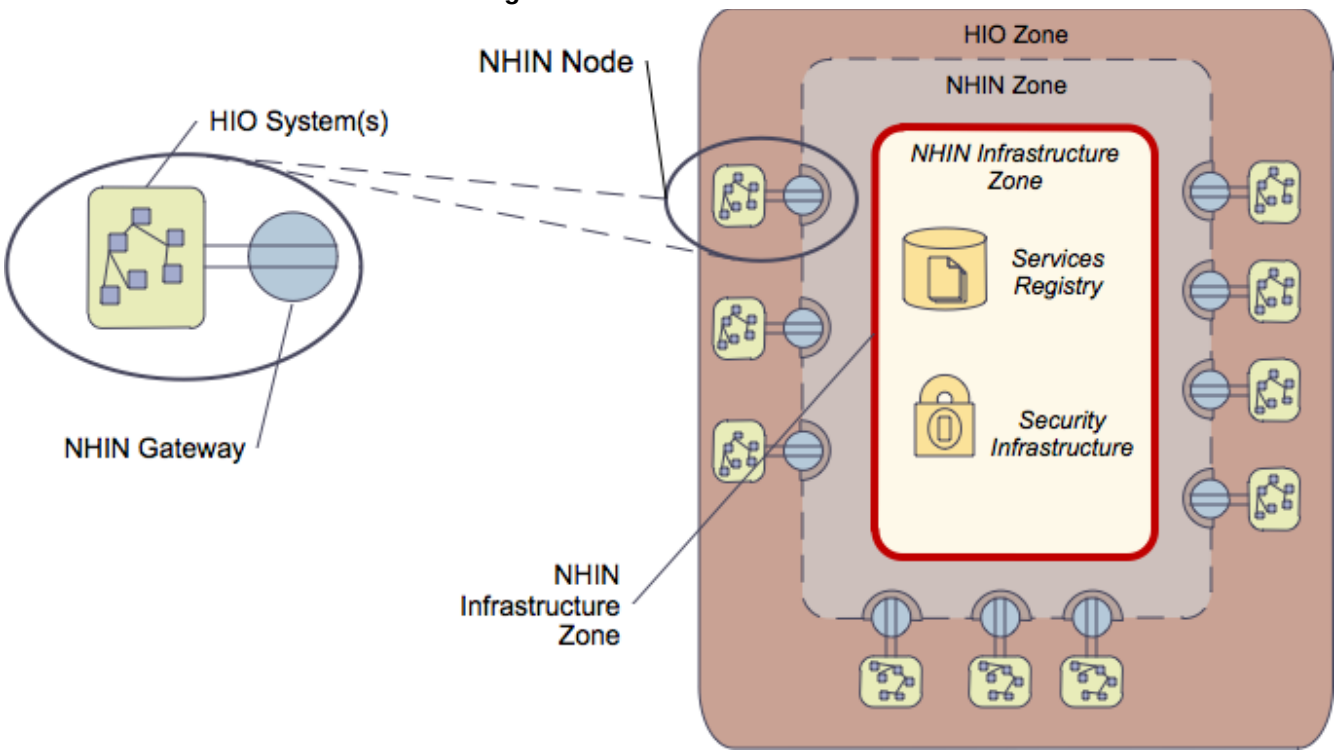

When taken as a whole these two concepts, zones and components within them, form a network view of the NHIN. The three zones which comprise the NHIN are described below:

1. The outer-most zone, the **HIO Zone**, contains the distinct HIO IT systems. As noted previously these systems can represent anything from a set of networks, such as an HIE, to a single clinicians portal running in an office. Additionally, according to the Local Autonomy principal of the NHIN, these systems are managed and governed by the policies and regulations particular to the respective owner organizations.

2. The **NHIN Zone** establishes the technical and governance trust fabric elements for secure exchange of information as well as the standard specifications required for reliable interoperability. Within this zone are the defined business rules and Web Service Interfaces that each NHIN Node must support in order to exchange information with other nodes. Those Web Service Interfaces are implemented by each NHIN Node through the use of a NHIN Gateway, which is required to connect nodes in the HIO Zone to the NHIN Zone.

3. The innermost zone is the **NHIN Infrastructure Zone**. This zone consists of the set of components and capabilities established by the NHIN for use by all NHIN Nodes.

4.2.2 NHIN Architectural Components

As described previously, each zone contains a number of architectural components. These are listed below:

HIO System(s) - This component represents the collection of Health IT systems that exchange information within an HIO – they are the nodes on the NHIN. These may be provider EMR systems, laboratory information systems, PHR services, etc According to the Local Autonomy architectural principle, each HIO is self managed and self governed.

NHIN Gateway - The NHIN Gateway is an HIO's implementation of the NHIN specified Web Service Interfaces. The NHIN Gateway enables HIOs to communicate with each other according to NHIN specifications, which are discussed in subsequent sections.

NHIN Operational Infrastructure - These represent the common components needed to support the NHIN Nodes in conducting distributed, secure communications between HIO NHIN Gateways. This Infrastructure consists of a Web Services Registry and the Security Infrastructure.

4.3 NHIN Architectural Layers

The NHIN Architecture can also be illustrated in a traditional layers model. In this model, each layer is defined by the distinct set of specifications, technologies, or mechanisms which define its capabilities. Each NHIN layer depends upon and extends the capabilities of the layers below.

Figure 3: NHIN Architectural Layers

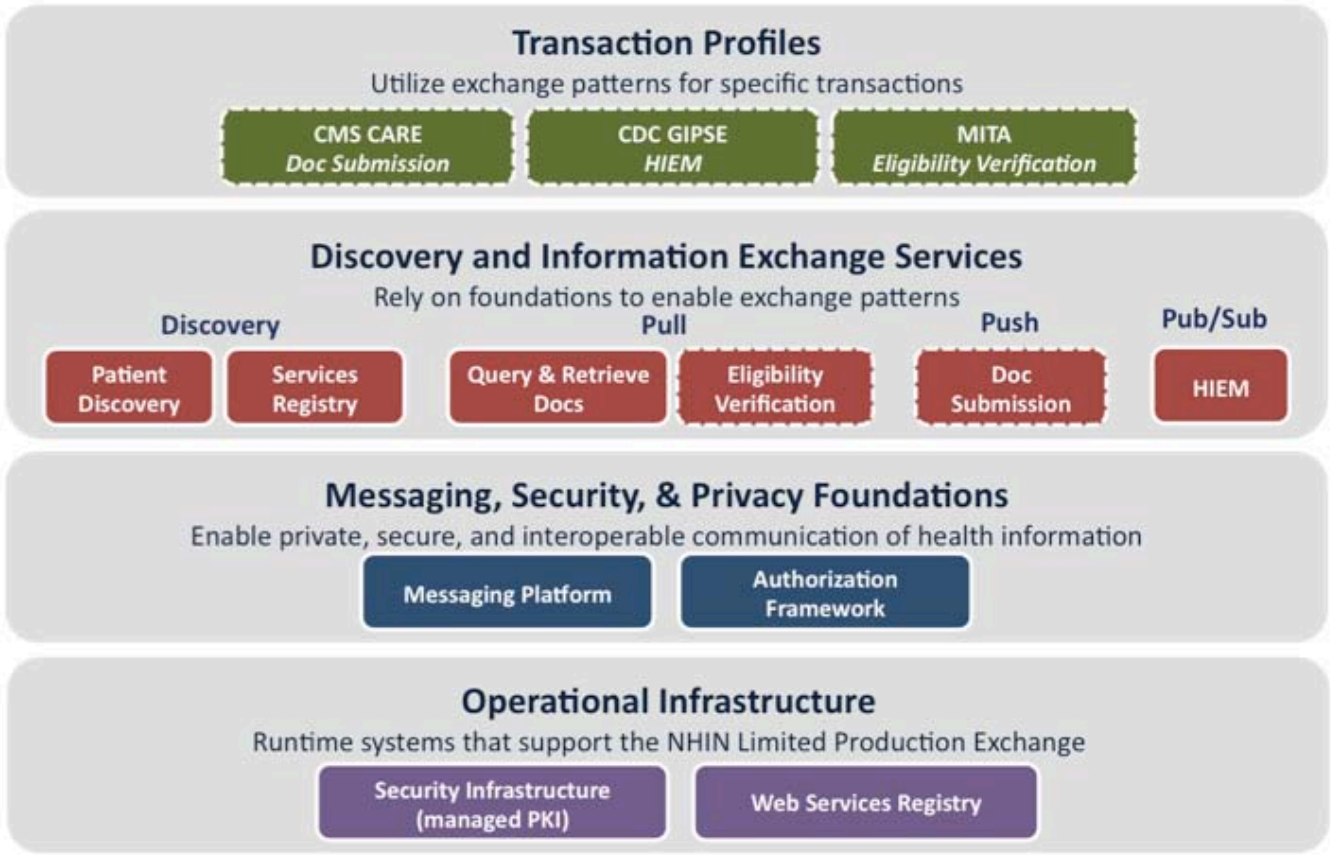

The lowest layer, which corresponds with the NHIN Infrastructure Zone, represents the infrastructure needed to support common network functions. The next two layers, which fall under the NHIN Zone, contain the specifications which define the fundamental mechanisms and standards to be used for used for information exchange with other NHIN Nodes.

The following sections take a more detailed look at the composition of each of the NHIN Architecture Layers.

4.3.1 NHIN Operational Infrastructure Components

The Operational Infrastructure Components, which comprise the NHIN Infrastructure Zone, represent the runtime systems that are used to support exchanges between NHIN Nodes. Unlike the other layers of the NHIN Architecture, the Operational Infrastructure components are implemented and maintained by the NHIN itself as a foundational utility. The table below describes each of the Operational Infrastructure Components.

Infrastructure System	System Description
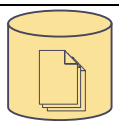 *Services Registry*	Based on a commercially available implementation of the Universal Description Discovery and Integration Registry[1] (UDDI) v3 specification, the NHIN Web Services Registry stores meta-data about the NHIN Gateway web services supported by each NHIN Node. Runtime system-to-system access to the NHIN Web Services Registry is restricted to NHIN Nodes through PKI secured interfaces. The information stored in the NHIN Web Services Registry will be managed by the NHIN program in coordination with the appropriate NHIN governance bodies. The NHIN Web Services Registry supports the use of sub-registries implemented and maintained locally by HIOs. The information in these registries is based on a central Web Services Registry operated and managed by the NHIN program.
 *Security Infrastructure*	The primary component of the Security Infrastructure is the Managed PKI (mPKI) Service. This service will be responsible for the management of X.509 based keys used to establish 2-way SSL for all exchanges between NHIN Nodes. Will be implemented by a commercial provider who will host services needed to support the issuance and revocation of PKI certificates. Issuance and revocation of security certificates will be performed by the NHIN program team as directed by the NHIN Governance bodies.

4.3.2 NHIN Messaging, Security, and Privacy Foundation

The NHIN Foundational Specifications, which fall within the NHIN Zone, define the base set of messaging and security standards which must be implemented by each NHIN Node, as well as the means to convey the metadata needed to inform information exchange decisions. In support of its architectural principles, the NHIN has defined a non-proprietary, implementation-agnostic foundation.

NHIN Nodes must authenticate each other and must encrypt the information exchanged between them. This mechanism is specified in the Messaging Platform and is accomplished by using a public key infrastructure (PKI) to create a secure channel between NHIN Nodes using TLS. Driven by the local autonomy and local accountability principles, NHIN Nodes provide each other with information needed to support the local authorization decisions about whether to engage in a requested exchange. This information is conveyed through the use of SAML assertions by the node attempting to initiate a NHIN transaction, as specified by the Authorization Framework.

[1] http://uddi.org/pubs/uddi_v3.htm

4.3.2.1 Messaging Platform Specification

The NHIN Messaging Platform is based on two interoperability profiles established by the Web Services Interoperability Organization (WS-I). These WS-I profiles provide interoperability guidance for core web service specifications created by organizations such as W3C and OASIS.

The two profiles contained in Messaging Platform, WS-I Basic v2.0 (*http://www.ws-i.org/Profiles/BasicProfile-2_0(WGD).html*) and WS-I Security v1.1 (*http://www.ws-i.org/Profiles/BasicSecurityProfile-1.1.html*) define the NHIN's common transport layer, specifying a base set of messaging standards and web service protocols which must be implemented by each NHIN Node and applied to all transactions. These profiles define the different specifications that make up the platform. For example the WS-I Basic profile details the use of SOAP 1.2, WSDL 1.1, WS-Addressing 1.0, WS-Policy 1.5 and MTOM 1.0. Together, these provide the basics for information exchange across the NHIN. Also contained within the WS-I Basic profile is the specification for UDDI v3.0.2 that is used as the implementation of the NHIN Web Services Registry.
The WS-I Security Profile provides the security mechanisms required for the NHIN through defining specifications for SOAP message security and various Security Token Profiles, as well as encryption and digital signatures. These profiles, as part of the NHIN Messaging Platform Specification, ensure that all NHIN inter-nodal messages conform to the same set of messaging standards.

4.3.2.2 Authorization Framework

Building upon the security related specifications found in the Messaging Platform, the Authorization Framework details NHIN specific constructs that are required for exchange of health information between NHIN Nodes. The Authorization Framework supports the Local Autonomy Principle by providing the responding node with the information needed to inform a response decision, and the Local Accountability Principle by requiring the initiating node to attest to the truth of the information provided.

There are four main areas covered:

- The first pertains to the authentication of the individual making a request for information. Based upon the meta-model defined in SAML 2.0, the responding node is provided with the method used to authenticate the requestor, the date and time they were authenticated, and optionally, the requestor's DNS, IP, and Session ID.
- The second area defines the information which must be provided to characterize the authenticated requestor and to identify the initiating NHIN Node and the subject of the message (patient). This information includes the requestor's name, role, and purpose for use; the entity to which they are associated; and the technical identifiers for the initiating node and the patient. This information is also conveyed via SAML 2.0 assertion.
- The third area defines the optional inclusion of Authorized Decision Statements among the information contained in the SAML assertion. These allow a requestor to assert that it holds an Access Consent Policy which specifically pertains to the request. The underlying assumption is that this assertion may be required to further inform access restrictions which may be in place at the responding NHIN Node. The initiator may either refer to a mutually known access consent policy, or provide information needed to allow the responder to retrieve a specific instance of such policy. The NHIN Access Consent Policy specification defines the format to be used.
- The final area covered by the Authorization Framework is the assertion signature, which supports assertion integrity, the authentication of the asserting party (e.g. the HIO), and may support non-repudiation.

4.3.3 NHIN Discovery and Information Exchange Services

NHIN information exchange patterns are defined by the NHIN Web Service Interface Specifications, which can be categorized as Discovery or Information services. When invoked by an initiating NHIN Node, these services are directed to specific service end points at one or more responding NHIN Nodes. The NHIN Discovery and Information Service specifications build upon the NHIN's Messaging, Security, and Privacy Foundation, as defined by the Messaging Platform and Authorization Framework specifications.

Each type of health information exchange transaction employs some combination of these services. Scenario-based examples of NHIN transactions and associated services are provided in section 5. More detail regarding these specifications can be found in the NHIN Implementation Guidance wiki[2].

4.3.3.1 NHIN Web Services Registry Interface

The NHIN Web Services Registry Interface specification defines the means by which a NHIN Node interacts with the UDDI registry in order to discover NHIN Nodes and the web service endpoints to which web services may be directed.

As described in section 4.3.1, the NHIN infrastructure includes the NHIN Web Services Registry, a UDDI registry. This registry contains the service end point information for each of the services offered by each NHIN Node. The means by which a HIO gains approval to be a node on the NHIN is detailed in the NHIN Validation Process, which is contained in a separate document.

4.3.3.2 Patient Discovery

Before one NHIN Node attempts to exchange patient specific information with another via the query and retrieve transactions, the two nodes must reciprocally establish that they are referring to the same person. Patient Discovery defines the mechanism by which one NHIN Node can query another to determine if it is a source of information for a specific patient. This query is intended to be directed to the most likely source nodes, as opposed to broadcast across the NHIN. Patient Discovery follows the following pattern:

1. The initiating NHIN Node sends a patient discovery request to one or more patient discovery web service end points for target (responding) NHIN Nodes obtained from the NHIN Web Services Registry. The request includes a set of demographic attributes associated with the patient for whom the requestor is seeking sources of information.

2. Each responding NHIN Node attempts to match the set of demographics provided by the requestor with those in its own master person index (MPI). If a single match (one per assigning authority) is achieved, the responding node responds with the Patient Identifier (PID) and the set of demographics locally associated with the matched person. If a single match cannot be achieved, the target may respond with an error condition which may either request additional attributes or indicate a match cannot be achieved.

3. The initiating NHIN Node may either accept the candidate match or evaluate the returned set of demographics to determine if agrees with the match.

A successful match allows the initiating NHIN Node to retrieve the PID used to identify the patient by the responding node. That PID is used for subsequent patient-specific transactions with the responding NHIN Node.

[2] Planned for release in Q2, 2010

4.3.3.3 Query for Documents and Retrieve Documents

This pair of specifications defines the first of the three NHIN information exchange transaction patterns. The Query for Documents service allows an initiating NHIN Node to request a patient-specific list of available documents from a responding NHIN Node. The query must include the PID used by the responding node to identify the patient for whom documents are sought. The initiating node presumably obtains the responder's PID via prior use of Patient Discovery. The Query for Documents service exhibits the following pattern:

1. The initiating NHIN Node directs a query to the responding NHIN Node containing the PID used by the responding node to identify the patient in question and, optionally, other constraining query metadata to refine the query.

2. Subject to an authorization decision, the responding returns a list of unique document IDs and other metadata used to describe any available documents.

The returned document IDs may be used by the initiating NHIN Node in a subsequent Retrieve Document request.

The Retrieve Documents service allows an initiating NHIN Node to retrieve one or more documents for a specific patient from a responding NHIN Node. The service requires the initiating node's use of the responding node Document IDs to specify the documents requested. Those Document IDs are presumably (but not necessarily), obtained by a prior Query for Documents. The Retrieve Documents service exhibits the following pattern:

1. The initiating NHIN Node sends a Retrieve Document request to the target HIO with the Document IDs used by the responding node to identify the documents sought by the requestor.

2. The responding NHIN Node evaluates the Document Retrieve request and, subject to an authorization decision, returns the requested documents to the initiating node.

It should be noted that the Document Retrieve request does not contain the PID used by the responder to identify the patient. PID is an attribute associated with the object identified by the Document ID.

4.3.3.4 Health Information Event Messaging

The Health Information Event Messaging (HIEM) service is the second of three information exchange transaction patterns. It allows an initiating NHIN Node to establish a subscription to information held by a responding node for periodic exchange. The service enables two transactions, subscription and notification.

1. Subscription – the initiating NHIN Node sends a subscription request to a responding node. The request contains the information needed to specify the type of events or information the requestor would like to receive. The responding node evaluates the request, and subject to an authorization decision, grants the request to establish the subscription.

2. Notification – the responding NHIN Node notifies the subscriber as events occur which match the approved subscription criteria.

The HIEM service specification defines the generic "pub/sub" mechanism. It does not include the information needed to request or fulfill any specific type of subscription. That information is to be contained in an HIEM Profile specification. GIPSE, which is described in section 4.3.4.1, is an example of an HIEM Profile specification.

4.3.3.5 Document Submission

The Document Submission service is another NHIN information exchange transaction type. It allows an initiating NHIN Node to send one or more documents for a given patient to a receiving node. Unlike Query/Retrieve and Pub/Sub, Document Submission does not require a prior request to retrieve a document or to subscribe to content. It is categorized as a "push" transaction that exhibits the following pattern:

1. The initiating NHIN Node submits one or more documents and associated Metadata that includes information used by the recipient to identify the patient. A Patient Identifier (PID) can be retrieved using the NHIN Patient Discovery service, or other similar attribute. This service does not require an initiating node to persist the document in its document repository upon document submission.

2. The receiving NHIN Node, upon receiving the submission, may evaluate the document(s) submitted and accept or reject them subject to local decision, and return a response with Accept or Reject status to the initiating node.

4.3.4 Information Service Profiles

Profile specifications further extend or constrain NHIN Discovery and Information services in order to define specific uses of that service. For example, an HIEM profile would include the information needed to request or fulfill a specific type of subscription.

4.3.4.1 HIEM GIPSE Profile

The Geocoded Interoperable Population Summary Exchange (GIPSE) profile defines how the HIEM service may be used to exchange GIPSE data. GIPSE is a data format created by the U.S. Centers for Disease Control and Prevention (CDC) to allow the electronic exchange of health condition/syndrome summary data that has been stratified by a number of variables, including geography.

4.3.4.2 Document Submission CARE Profile

The Continuity Assessment Record and Evaluation (CARE) profile defines how the Document Submission service may be used to exchange CARE data. The CARE data set defined by CMS includes discharge assessment summary data obtained prior to and after the discharge of a patient. The profile is part of the C-HIEP proof of concept being conducted by CMS.

5 NHIN Use Case Scenarios

5.1 Introduction

The intent of this section is to convey the context behind the NHIN architecture by illustrating how NHIN services are orchestrated. The following scenarios provide examples of how several NHIN services are used in concert to carry out a typical series of health information exchange transactions.

5.2 Scenario A: Retrieving Information for a New Patient (Query/Retrieve Pattern)

Dr. Smith, who is a member of an HIO which is a node on the NHIN, schedules an appointment for Ms. Jones, a new patient. As part of the workflow prior to seeing the patient, the doctor's office seeks medical information on that patient which may be available from other NHIN Nodes. In addition to the patient information held by the clinician's own system as well as that held by his local HIO, Dr. Smith is presented with information retrieved from other NHIN Nodes. The following is an explanation of the sequence of transactions needed to support this scenario:

Step 1 – Identify Target NHIN Nodes

As part of the appointment scheduling process, the patient, Ms. Jones, provides identifying information including name, birthdate, gender, address, and phone number. In addition, she is asked to indicate if she has received treatment in another city or state. In most cases, a NHIN Node will assume to search for health information within a geographically local area, unless there is reason to believe a wider search is necessary.

In order to use the NHIN to seek Ms. Jones' health information, Dr Smith's HIO must identify which other NHIN Nodes may hold her information. In order to promote a positive user experience, the local HIO's NHIN Gateway might be configured to target a predefined set of NHIN Nodes most likely to hold the patient's medical information (e.g. those in close proximity or in the same healthcare delivery area). Alternatively, target NHIN Nodes might be identified using information provided by the patient such as location of prior treatment or provider type (e.g. VHA). Using information gathered during scheduling, Dr. Smith's HIO communicates with the NHIN Web Services Registry to identify target nodes and discover their web service endpoints. The Web Services Registry, a component of the NHIN Infrastructure Zone, maintains a list of NHIN Nodes, as well as information about the web services they support.

Step 2 – Discover Sources of Patient Information

Dr. Smith's HIO uses the NHIN to determine if Ms Jones is known by the targeted NHIN Nodes in order to determine if they are potential sources of her health information.

Dr. Smith's HIO invokes the Patient Discovery Service, directing a Patient Discovery request to the appropriate target NHIN Node web service endpoints obtained from the NHIN Web Service Registry. The request contains the demographic information collected from the patient (name, birth date, gender at a minimum). The target node evaluates the request, and if the target node is able to achieve a single match, it responds with the set of demographics and Patient Identifier (PID) locally associated with the matched person. If the target node cannot achieve a single match, it responds with a request for additional attributes which may help achieve a match, or with an error indicating a match could not be achieved.

The initiating NHIN Node (Dr. Smith's HIO) may choose to accept the candidate match or it may choose to use its own criteria to verify the match. If the initiating node agrees on the match, it is accepted and the PID retrieved from the target node will be used for subsequent transactions.

Step 3 – Search for Relevant Documents

Once Dr. Smith's HIO has obtained a PID for Ms. Jones from at least one of the target NHIN Nodes, the Query for Documents service will be used to communicate with the target nodes to obtain a list of documents available for Ms. Jones.

The query sent by the initiating NHIN Node includes the target node PID and, optionally, other restrictions such as timeframe or document type. In this case Dr. Smith requests a list of all Ms Jones' documents, which may have the potential to result in a long list. The responding node includes only those documents which are available according to its local policies and permissions, based on the information Dr. Smith's HIO provided via the Authorization Framework.

Step 4 – Select Documents to Retrieve

Dr. Smith, via his HIO, is presented with the list of available documents and is able to select those to be retrieved. This selection may be automatic, based on some predetermined criteria ('retrieve all' being the simplest), or specified by a user. In this case, a member of Dr. Jones staff reviews the list of Ms Jones documents and selects those that contain the information Dr. Smith typically seeks. Each document that is selected is retrieved from the appropriate target node via the Retrieve Documents service.

5.3 Scenario B: Population Health - Biosurveillance (Pub/Sub Pattern)

A key capability supported by the NHIN is the ability to broadcast specific types of health information to a set of NHIN Nodes who have requested such information. Often referred to as Publish/Subscribe or Pub/Sub, this exchange capability is defined in the Health Information Event Messaging (HIEM) Web Service Interface specification. The GIPSE Profile defines a specific use of HIEM which includes hierarchically organized types of content that can be subscribed to.

The sample scenario outlined below describes how an organization, such as the U.S. Centers for Disease Control and Prevention (CDC), asks State Public Health Agencies to send the organization regular updates on cases involving the novel H1N1 Influenza (A) using HIEM and the GIPSE Profile.

Step 1 – Subscribe to Health Information

A State Public Health Agency agrees to send the CDC de-identified cases of novel H1N1 Influenza (A) using an HIEM service which supports the GIPSE profile. The CDC then uses its NHIN Gateway to subscribe to this data using the HIEM service and the GIPSE defined subscription defined by the GIPSE profile. This involves creating a subscription message that indicates a subscription request for DeidentifiedPopulationHealthData/Biosurveillance/Influenza content. The subscription message also contains the web service endpoint where such content should be sent when it is available for delivery. The subscription message is sent to the State Public Health Agency's NHIN Gateway.

The State Public Health Agency's NHIN Gateway receives a subscription message from the organization. An official at the State Public Health Agency reviews the subscription request and approves the subscription. A SubscribeResponse message is then sent from the State Public Health Agency's NHIN Gateway to the Subscriber NHIN Gateway.

Step 2 – Gather and Identify Relevant Health Information

AJ is a 20 y/o female who visits Dr. Matthews, a primary care provider who is not her regular doctor, on her way back from an overseas trip. AJ describes body aches that began 2 days ago and were followed this morning by fever, coughing, shortness of breath, and fatigue. Because she has never been seen at this clinic before, Dr. Matthews suspects that the case is influenza even though this would be a very early case (flu season has not yet officially begun). Using a swab from a rapid Influenza A kit, Dr. Matthews confirms the patient has Influenza A. Since Dr. Matthews is an influenza sentinel physician, he sends the swab to the state lab[3] for identification of H1N1.

The sample submitted by Dr. Matthews is analyzed at the state lab. The sample is identified as novel H1N1 Influenza (A). These results are shared with a regional health information exchange (HIE), which forwards them to Dr. Matthews as well as the State Public Health Agency.

An epidemiologist at the State Public Health Agency is able to look at displays of current influenza data from the HIE (gathered by regular polling of Laboratory Documents) and stratifies the data using appropriate demographic variables. A dramatic increase in the number of positive Influenza-A labs is noted by the epidemiologist, indicating that flu season has begun.

Step 3 – Publish and Notify Health Information

Influenza data from the State Public Health Agency is aggregated and de-identified every 6 hours into a GIPSE (Geocoded Interoperable Population Summary Exchange) file. When the GIPSE file has been created, it is placed on a server that is connected to the agency's NHIN Gateway.

The State Public Agency's NHIN Gateway periodically polls the server that contains GIPSE files for new data. That server forwards the NHIN Gateway the GIPSE file. The NHIN Gateway then initiates delivery of the GIPSE file to subscribers of DeidentifiedPopulationHealthData/Biosurveillance/Influenza content.
For each subscriber of DeidentifiedPopulationHealthData/Biosurveillance/Influenza content, an HIEM Notification message containing the GIPSE file is constructed and transmitted to the web service endpoint specified in the subscription message received from the subscriber. It contains the GIPSE file.

A subscriber to DeidentifiedPopulationHealthData/Biosurveillance/Influenza content receives a Notification message from the State Health Agency NHIN Gateway which contains the new GIPSE data. An acknowledgement message is returned.

The Subscriber's internal systems now process the GIPSE message and use it to carry out business functions. The CDC may use the data to track the incidence of novel H1N1 Influenza A in the state that transmitted the notification message and compare the data with other states in the region.

[3] The State Labs are the only labs confirming influenza samples submitted from influenza sentinel physicians or from flu-related deaths.

5.4 Scenario C: Document Submission (Push Pattern)

A clinician at an acute care hospital, which is a member of a NHIN Node, discharges a patient who is a Medicare beneficiary. As part of the discharge process, the clinician submits Continuity Assessment Record and Evaluation (CARE) patient assessment data through a NHIN Gateway to the Centers for Medicare and Medicaid Services (CMS) via the NHIN Node. As described in the NHIN CARE profile specification, CARE data contains clinical information about a Medicare beneficiary at a designated transition point in care. The "push" mechanism leveraged in this scenario is detailed in the NHIN Document Submission specification. The following steps outline this scenario:

Step 1 –Submit Health Information

A patient who is a Medicare beneficiary is discharged from an acute care hospital. The acute care hospital is a member of an NHIN Node. As part of the discharge process, the clinician performs a CARE assessment for the patient using its EMR and/or a combination of other HIT system. The data obtained from this assessment, referred to as CARE data, is transmitted to the NHIN Node. The NHIN Node creates a structured CARE Assessment document (using the HITSP C83 standard) that contains the de-identified CARE data and submits this document to CMS across the NHIN using the NHIN document submission service.

The web service end point location for the CMS NHIN Gateway has been obtained prior to initiating the document submission transaction. The submission includes an identifier for the patient that has been previously acquired through some verifiable means, such as a Patient Discovery transaction.

Step 2 – Receive Health Information

CMS receives the submitted document containing the CARE data. Receiving the CARE data, CMS first sends an HTTP acknowledgement to the submitting NHIN Node, providing a confirmation of message receipt.

At some later point in time after CMS evaluates the submitted document, it determines whether to accept or reject the document. CMS then sends an "Accepted", "Accepted with Warnings" or "Rejected" status response to the submitting NHIN Node describing appropriate indications of reasons for rejection and warnings, if any.

6 Glossary

Clinical Information – refers to all information associated with a patient over time. The information can contain health assessments (physical health, mental health, substance abuse, etc.), diagnosis, problems, progress notes, medications (prescription and over-the-counter), alerts, laboratory test results, demographics and any information relevant to the treatment of a patient.

Data Use and Reciprocal Support Agreement (DURSA) - The DURSA is a comprehensive, multi-party trust legal agreement and is based upon a set of policy assumptions that bridge varying state and federal laws and regulations, as well as various policies. This legal contract, signed by all entities currently exchanging information via the NHIN Exchange, provides a framework of trust assurance to support health information exchange across the NHIN.

Health Information Organization (HIO) - refers to an organization that maintains a common index of patients and their associated documents/data, encompasses participants that have established trust and business relationships, and have agreed to adhere to common standards. The scope of an HIO may range from a single provider practice, to large chains or delivery networks, existing jurisdictional HIEs/RHIOs, and state or county agencies.

Interoperability - The IEEE defines interoperability as: the ability of two or more systems or components to exchange information and to use the information that has been exchanged.[1]
James A. O'Brien and George M. Marakas define interoperability as: Being able to accomplish end-user applications using different types of computer systems, operating systems, and application software, interconnected by different types of local and wide area networks.

Message – A "message" in information technology is an electronic file that may be received and read either by a human, a computer, or both. In addition to the payload or "content" of the message, the message includes a "wrapper" or "header" that defines one or more characteristics of the message, especially the message destination. Messages often are used by software programs to transfer content to other software programs via programming interfaces.

Nationwide Health Information Network (NHIN) – ONC guides the technical and policy development of the Nationwide Health Information Network. The NHIN is a collection of standards, protocols, legal agreements, specifications, and services that enables the secure exchange of health information over the internet.

NHIN Gateway - An instantiation of the NHIN specifications which supports secure, interoperable health information exchange across the NHIN. NHIN CONNECT is an example of an NHIN Gateway.

NHIN Gateway Zone – The software architecture concept that includes all the NHIN Gateways instantiated on the NHIN. The NHIN Gateway Zone functions architecturally "between" the NHIN Infrastructure Zone, operated by the Federal Government, and the NHIN HIO Zone, which represents the "internal" IT systems of the HIO. In this sense, the NHIN Gateway Zone includes the architectural concepts of communications internally with the NHIN HIO Zone and communications externally over the Internet with the NHIN Infrastructure Zone and other NHIN Gateways in the NHIN Gateway Zone.

NHIN HIO Zone – The "internal" IT systems of a HIO that are protected from the Internet by a NHIN Gateway.

NHIN Infrastructure Zone – The NHIN Infrastructure Zone includes systems operated by the Federal Government including the NHIN Web Services Registry and the NHIN Security Infrastructure. NHIN Gateways from the NHIN Gateway Zone communicate directly with the systems of the NHIN Infrastructure Zone.

NHIN Node - "NHIN Node" is the preferred term in this document and refers to an HIO that is exchanging information via the NHIN.

NHIN On-Boarding Process – Process by which a qualified HIO becomes a NHIN participant. This process is described in separate NHIN on-boarding documentation.

NHIN Security Infrastructure – The primary component of the Security Infrastructure is the Managed PKI (mPKI) Service. This service will be responsible for the management of X.509 based keys used to establish 2-way SSL for all exchanges between NHIN Nodes. Further description of the NHIN Security Infrastructure can be found in section 4.3.1.

NHIN Specifications - Prescribe the data content, technical, and security requirements necessary to support information exchange among NHIN Nodes

NHIN Web Services Registry – A component of the NHIN Infrastructure Zone that enables a NHIN Node to discover the capabilities of other NHIN Nodes. It is a UDDI registry, which lists NHIN Nodes, the NHIN web services supported by each NHIN Node, and how to reach those web service end points.

Provider - Includes any entity that delivers healthcare to patients, whether in an emergency, ambulatory or inpatient setting. This encompasses behavioral as well as physical health. Providers may be an individual doctor's office or clinic, a hospital or chain of hospitals, an integrated delivery network, or a Federal agency that delivers care, such as the Departments of Defense or Veterans Affairs.

Semantic Interoperability - Beyond the ability of two or more computer systems to exchange information, semantic interoperability is the ability to automatically interpret the information exchanged meaningfully and accurately in order to produce useful results as defined by the end users of both systems. To achieve semantic interoperability, both sides must defer to a common information exchange reference model. The content of the information exchange requests are unambiguously defined: what is sent is the same as what is understood.

Service Interface – "Service Interface" is a type of information technology programming interface in which software program logic is separated from underlying software technologies through the use of software interfaces, each of which defines a contract between a software service consumer and a software service provider. Examples include web service interfaces, CORBA service interfaces, and other message-based services (based on Open Knowledge Initiative).

Syntactic Interoperability - If two or more systems are capable of communicating and exchanging data, they are exhibiting syntactic interoperability. Specified data formats, communication protocols and the like are fundamental. In general, XML or SQL standards provide syntactic interoperability. This is also true for lower-level data formats, such as ensuring alphabetical characters are stored in ASCII format in both of the communicating systems. Syntactical interoperability is required for any attempts of further interoperability.

Web Service – A "Web Service" represents the ability to manage data that may be created, read, updated, deleted or transformed via methods described by a Web Service Interface. A Web Service Interface is a type of information technology programming interface defined by W3C in which software program logic is separated from underlying software technologies through the use of software interfaces, each of which is defined by a Web Service Definition Language (WSDL) contract between a software service consumer and a software service provider. The presence of a "web service interface" implies the presence of a "web service." However, these two terms are not necessarily interchangeable.

Web Service Endpoint – A “Web service endpoint” is a (referenceable) entity, processor, or resource to which Web service messages can be addressed. (From W3C).

UDDI Registry – Universal Description Discovery and Integration Registry (UDDI) is a type of registry in the NHIN Infrastructure Zone known as the NHIN Web Services Registry. It lists the NHIN Nodes, the NHIN web services supported by each NHIN Node, and how to reach those web service end points.