

*Prepared for the*  
**Administration for Children and Families (ACF)**

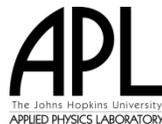
**National Human Services Interoperability  
Architecture**

**Security**

**DRAFT Version D0.2**

**June 2012**

*Prepared by:*  
**The Johns Hopkins University  
Applied Physics Laboratory (JHU/APL)**



## Draft Issue

It is important to note that this is a draft document. The document is incomplete and may contain sections that have not been completely reviewed internally. The material presented herein will undergo several iterations of review and comment before a baseline version is published.

This document is disseminated in the interest of information exchange. The Johns Hopkins University Applied Physics Laboratory (JHU/APL) assumes no liability for its contents or use thereof. This report does not constitute a standard, specification, or regulation. JHU/APL does not endorse products or manufacturers. Trade and manufacturer's names appear in this report only because they are considered essential to the object of this document.

Note: This document and other NHSIA-related documentation are available for review from the NHSIA SharePoint site. Updates and any additional documents will be published on that site. The URL for the site is <https://partners.jhuapl.edu/sites/HSNIA>. The version D0.1 and D0.2 documents may be viewed or downloaded from the document library named [NHSIA Drafts](#).

Review and comments to this document are welcome. To comment, either post your feedback in the [NHSIA Drafts Comments](#) library or send comments to [NHSIAArchitectureTeam@jhuapl.edu](mailto:NHSIAArchitectureTeam@jhuapl.edu).

John J. Tamer  
The Johns Hopkins University Applied Physics Laboratory  
11100 Johns Hopkins Road  
Laurel, MD 20723  
Phone: 443-778-8248  
E-Mail: [john.tamer@jhuapl.edu](mailto:john.tamer@jhuapl.edu)

## Table of Contents

List of Figures .....	ii
1 Introduction.....	1
2 Identity Management and Access Control .....	1
2.1 Authentication.....	2
2.2 Strong Authentication & Multi-Factor Authentication.....	2
2.3 Access Control .....	2
2.4 Types of Access Control.....	3
2.5 Identity Management .....	4
2.6 GFIPM .....	4
2.6.1 Security Assertion Markup Language (SAML).....	5
2.6.2 GFIPM in the NHSIA Context .....	5
2.6.3 Creating a NHSIA Identity Federation .....	6
2.6.4 Federation Managers.....	6
2.6.5 Identity Providers.....	7
2.6.6 Service Providers .....	8
2.6.7 Identity Attributes .....	8
2.7 Implementation Considerations .....	9
2.7.1 How Many Federations are Needed?.....	9
2.7.2 What Attributes are Used to Control Access? .....	10
2.7.3 Connecting Federal, State and Local Hubs.....	10
2.7.4 How to Control Human Services Client Access .....	10
2.8 GFIPM References and Resources .....	11
3 Network and Infrastructure Security.....	12
3.1 Encryption.....	13
3.2 Transmission Security.....	13
3.3 Security of Data at Rest .....	14
3.4 Remote Access.....	14
4 Privacy and Confidentiality .....	15
5 Contingency Planning and Disaster Recovery.....	16
5.1 Data Backup Methods and Offsite Storage.....	17
5.2 Emergency Mode Operations Plan .....	17
Appendix A: NHSIA Architecture Principles Related to Information Security.....	19

## List of Figures

Figure 1: GFIPM Document Map.....	11
Figure 2: GFIPM Home Page .....	12

# 1 Introduction

The ability to securely share and protect information is a key aspect of NHSIA. While organizations are likely to have invested significantly in securing their own environments, NHSIA complicates an already complicated area by bringing users and data together in a new, shared environment. Because of this, NHSIA must provide mechanisms to verify the identities of individuals who will access the environment, must authorize their admittance into the environment, and must control the applications and information to which those individuals have access. NHSIA must ensure that data is communicated into and out of the shared environment securely and that it is adequately secured and protected while in the NHSIA environment.

This document is intended to augment the NHSIA viewpoints by providing additional guidance on security-related topics relevant to jurisdictions planning to adopt the NHSIA framework. This document highlights some of the major security concerns and best practices that can be adopted by organizations as they implement interoperable systems that cross organizational boundaries. The material presented here will be of interest to both technology leaders, who must design and implement security solutions, as well as administrators who must ensure that applicable policies and regulations related to security are adhered to.

These security concerns are applicable to NHSIA regardless of the deployment approach used by the individual jurisdiction. In addition, the concepts discussed here are not intended to supersede the security practices implemented by individual organizations. Rather, they are intended to provide appropriate security in an environment that includes collaboration and information sharing across organization boundaries. Appendix A contains a set of basic security architecture principles. These principles establish the foundation upon which this document is based as well as provide a more general starting point for an organization to build its own security architecture.

The primary security concepts presented in this paper revolve around four key areas. These are:

- Identity Management and Access Control
- Network and Infrastructure Security
- Privacy and Confidentiality
- Contingency Planning and Disaster Recovery

## 2 Identity Management and Access Control

While the HIPAA (Health Insurance Portability and Accountability Act) Security Rule (45 CFR §164) was primarily intended to apply to protected health information that is maintained or transmitted in electronic form, the principles and guidelines established by the security rule are applicable to the sharing of any type of information where confidentiality is of utmost importance. The rule sets forth both required and addressable standards related to administrative, physical, and technical safeguards. This paper primarily addresses the technical safeguards (45 CFR §164.312), beginning with Identity Management and Access Control.

## 2.1 Authentication

Authentication is the process of determining that a person really is who they say they are. Authentication typically precedes identity management and is typically the responsibility of the organization that employs the person. When a person begins employment, he or she must present documentation that establishes one's identity and ability to work in the United States. Many organizations will also conduct credit checks or full background investigations on new employees. Employees are typically assigned user identifiers (IDs) and passwords that provide them access to the company's computer resources.

## 2.2 Strong Authentication & Multi-Factor Authentication

Strong authentication is a general term used to describe an authentication process that is more stringent than use of a simple ID and password. It can refer to the use of forms of proof, such as random password tokens, or smart cards, or it can refer to multi-factor authentication. Multi-factor authentication is the requirement for more than one form of proof of identity, from more than one type, or factor of proof.

The three main types of factors are:

- Human Factors, including biometrics such as retina scans or fingerprint readers.
- Personal Factors, such as passwords or personal identification numbers (PINs).
- Technical Factors, such as a smart card or token.

A multi-factor authentication process must include at least one form of proof from at least two of the above factor types. Multi-factor authentication greatly reduces the risk of establishing fraudulent identity over a scheme that uses only one factor. It takes away the ability to fraudulently authenticate by obtaining any single piece of technology or password secret.

## 2.3 Access Control

Access control refers to the technical means used to control who can access an information technology (IT) resource. The simplest form of access control involves logging on to a computer system with a user ID and password. The user ID assigned to an individual is unique to that person and is typically only assigned after the person has verified his or her identity to the company, typically the employer, that is issuing the user ID. Therefore, the combination of the user ID and password is used to authenticate the person, (e.g., verify that they are whom they say they are) and to control access to the resources the person is allowed to access.

Unfortunately, the situation becomes a bit more complicated when more than one organization is involved. Each organization must protect its information, but each must also facilitate the sharing of that information through the use of IT. Thus, true access control for a multi-organizational enterprise requires more robust authentication, authorization, and access control. Access control should determine what resources are authorized to be accessed by a user or process and prevent resources from being accessed by unauthorized users. This can be accomplished through the use of federated single sign-on.

Single sign-on (SSO) addresses the cumbersome situation of logging on multiple times to access different resources or different systems. In most cases, users should not be required to maintain separate sets of logon credentials to access both local and shared resources. When users must remember numerous passwords and IDs, they are more likely to take shortcuts in creating them that could leave them open to exploitation.

Federated SSO provides a secure, standard way to share user identities among multiple organizations. Users sign on once (the SSO) using their standard network login, typically assigned by their home organization. Their identity is then transparently and securely shared with the requested system or resource.

Use of federated SSO begins with the creation of a federation. A federation is a group of two or more trusted partners with business and technical agreements that allow a user from one federation partner (participating agency, or organization) to seamlessly access resources from another partner in a secure and trustworthy manner. The federation provides a standardized means for allowing agencies to directly provide services for trusted users that they do not directly employ or manage. Essentially, the users from one organization are granted access to the resources of another. A well-defined set of attributes about users is securely exchanged between the two organizations. This allows access decisions to be made by each participating organization in accordance with its local policies and business practices.

For example, an attribute could define the role of the individual as a case worker from Organization A. When this individual wishes to access information from Organization B, Organization A electronically informs Organization B that they have authenticated the individual and that the individual is a case worker. In advance, Organization B has determined and identified which information it is willing to share with caseworkers from Organization A. Using this approach, the technology not only facilitates the sharing of information between Organization A and Organization B, it also controls what information can be shared between the two and by whom.

## 2.4 Types of Access Control

Access control can be implemented in a number of ways depending on the nature of the information being accessed and the capabilities of the information technology infrastructure. These include:

- **Discretionary Access Controls (DAC)** are the controls placed on resources at the owner's discretion. Discretionary access controls are not subject to pre-defined rules but are rather granted at the discretion of the information owner.
- **Mandatory Access Controls (MAC)** are controls based on policies. Policies define rules by which a user, or program is permitted access to a resource. Unlike DAC, access is universally applied based on pre-defined rules, in other words, it is non-discretionary.
- **Content Dependent Access Control** involves restricting access to content, such as documents and emails, based on embedded keywords or metadata. It works by inspecting the content and applying rules to determine if access is permitted. It is also possible to

combine content dependent access control with role-based access control in order to limit access to content by established roles.

- **Role-Based Access Control (RBAC)** Given the potentially large number of users of a system, access privileges are generally not assigned at the user level. Instead, users are assigned to groups (mimicking the organizational structure of a company), or roles (defined based on job functions that users perform), or some combination of the two. Access privileges are then assigned to groups and/or roles. The most natural case is that they are assigned to roles, since roles align more closely with operations users naturally perform to accomplish their job. The industry term for this is Role-Based Access Control (RBAC). RBAC is more flexible than defining access rights based on usernames or static groups and enables an organization to be more versatile when allocating resources. With RBAC the system must determine if the subject (user or client) is associated with a role that has been granted access to a resource. This process of user to role ascertainment is called role mapping.
- **Attribute-Based Access Control (ABAC)** There are times when access should be based on characteristics the user has rather than the organization or roles to which the user belongs. Attribute-based access control offers a more dynamic method of controlling access by basing decisions on attributes assigned to users, which may in fact change as business events unfold. Access policies define the attributes and values a user must have, and access decisions are evaluated against the current values assigned to the user. Attributes can be used to support both course-grained and fine-grained authorization.

## 2.5 Identity Management

Identity Management refers to the capability to store, manage, provision, administer, and audit security data related to user identity. It is typically the responsibility of each organization to perform these functions for its own employees. Identity information is usually stored in some type of system directory, such as Microsoft's Active Directory, that manages usernames and passwords and can provide access to company resources. In addition, many organizations employ systems and applications that have their own identity information. These systems will usually record the functions and privileges for a user as well.

## 2.6 GFIPM

The Global Federated Identity and Privilege Management (GFIPM) is a program sponsored jointly by the U.S Department of Justice and the U.S. Department of Homeland Security. It is a part of the Global Reference Architecture (GRA), a service-oriented reference architecture for justice and public safety information sharing. The GFIPM program has been developing information sharing solutions based on the concept of federated identity and privilege management.

According to GFIPM, a federation is a group of two or more trusted partners with business and technical agreements that allow a user from one federation partner (participating agency, or organization) to seamlessly access resources from another partner in a secure and trustworthy manner. The federation provides a standardized means for allowing agencies to directly provide services for trusted users that they do not directly manage. The identities from one organization

are granted access to the services of another. A well-defined set of trusted attributes about locally authenticated users is securely exchanged between the two organizations, allowing for identification and fine-grained access decisions to be made by each participating organization in accordance with its local policies and business practices.

At the highest level within the GFIPM model, there are three vital components that must interact between multiple systems:

- Identity Provider (IDP)
- Service Provider (SP)<sup>1</sup>
- User Profile Assertion (Metadata)

Within a federation, one organization may be an identity provider, a service provider, or both. The identity provider (IDP) is the authoritative entity responsible for authenticating an end user and asserting an identity for that user in a trusted fashion to trusted partners. The identity provider is responsible for account creation, provisioning, password management, and general account management. These are typically performed as a part of the identity management process within the individual organization.

The service provider (SP) is the organization that has systems, information, or web services that it is willing to share. The service provider relies on the identity provider to assert information about a user, leaving the service provider to manage access control based on these trusted sets of attributes. The communication of these attributes to the service provider is provided in the User Profile Assertion and is implemented using an industry standard called Security Assertion Markup Language, or SAML.

### **2.6.1 Security Assertion Markup Language (SAML)**

SAML is an XML (eXtensible Markup Language) -based standard for exchanging authentication and authorization data between identity providers and service providers. SAML is a product of the OASIS (Organization for the Advancement of Structured Information Standards) Security Services Technical Committee. GFIPM uses SAML 2.0 to address single sign-on and secure exchange of user attributes to support both user to system interactions and web services security.

### **2.6.2 GFIPM in the NHSIA Context**

The concepts developed as a part of GFIPM provide a proven approach that is applicable to NHSIA. The mission of the Global Reference Architecture, the parent of GFIPM, is “to enhance justice and public safety through a service-oriented approach to information sharing<sup>2</sup>.” NHSIA has a similar mission with respect to health and human services. Just as GFIPM facilitates service-oriented information sharing by providing identity and access management to the justice community, so too can it provide guidance for the same approach for the health and human services community.

---

<sup>1</sup> The “Identity Provider” and “Service Provider” terms are those used by GFIPM. A service provider in this context refers to an IT service provider, not a human services provider.

<sup>2</sup> Justice Information Sharing Global Reference Architecture,  
<http://www.it.ojp.gov/default.aspx?area=nationalInitiatives&page=1015>

### 2.6.3 Creating a NHSIA Identity Federation

One objective of the GFIPM standards and specifications is to provide a framework for securely connecting justice and public safety personnel to interagency applications and data over the Internet. Federation is a fundamental concept within the GFIPM framework. As such, federation is also a fundamental concept of NHSIA, providing a means to allow human services workers and clients to securely connect and share information.

The goal of a federation is to provide participating organizations with the following benefits:

- Provide single sign-on capabilities to end users for accessing online services.
- Eliminate the requirement to register user identity information in multiple external systems.
- Retain identity management and user authentication responsibility at the local organization level.
- Provide an interoperable standard vocabulary of identity access attributes. Identity access attributes are information items about a person that serve to determine the resources that person is allowed to access.
- Support informed access and authorization decisions based on a trusted set of identity access attributes thereby improving the security controls and scalability of electronic information sharing.

The federated approach to identity and privilege management provides a standards-based means for human services entities to locally authenticate their organizations' users and provide accurate and current user identity attributes to shared, external systems which in turn utilize the trusted attributes to make authorization and system access decisions.

Formation of a federation represents a trust model that enables multiple independent but related organizations to access online services based on the federation attributes issued by trusted identity providers (IDPs).

In practice, federations may be formed at any level. In NHSIA, it may be practical to form federations at the state level to permit sharing of resources across the state. In other situations, a federation could be formed at the county level or at a level that includes multiple neighboring states. The decision will be made based on the number and locations of organizations participating as well as the ability of those organizations to agree on common standards for assertion of identity attributes (metadata).

### 2.6.4 Federation Managers

As previously stated, a federation is a trust model formed among a collection of Identity Providers and Service Providers spanning multiple agency organizational boundaries. Identity Providers manage the identity and authentication of their local users. Service Providers rely on Identity Provider assertions of attributes to make authorization and access control decisions for sharing their information or services.

Under the GFIPM model, a federation includes a federation manager. This may be the responsibility of one of the federation members, or it may be performed by an independent,

third-party. In either case, the federation manager will be responsible for the day-to-day operations of the federation. This includes such things as:

- Developing policies and guidelines pertaining to the definition and usage of the metadata standard for end-user attributes.
- Implementing approved processes for determining the membership of any new party in the federation.
- Developing technical architecture and providing documents, including interface specifications, for technical interoperability within the federation.
- Conducting day-to-day operational services, i.e., audits.

In addition to the above responsibilities, the manager of a federation typically operates a federation certificate authority (CA) to provide trust and security to the federation. The purpose of this CA is to sign the federation's Cryptographic Trust Fabric document. The Cryptographic Trust Fabric document defines the most current cryptographic security context of the federation.

The Cryptographic Trust Fabric is essentially a mechanism to authenticate federation members through the use of a Public Key Infrastructure (PKI). The federation members trust the CA to authenticate the federation members and to issue certificates that can be used to relay that authentication between members.

A Cryptographic Trust Fabric document contains public keys and system entity metadata for each trusted endpoint in the federation. The federation endpoints include identity providers (IDPs), service providers (SPs), Web Service consumers (WSCs), Web Service providers (WSPs), and others. The federation manager maintains the Cryptographic Trust Fabric document and makes a new version of it available to federation members whenever the membership of the federation changes.

The GFIPM standards include a number of documents that detail the responsibilities of the federation manager.

### **2.6.5 Identity Providers**

An identity provider is the organizational entity that manages the identities for a particular set of users. These users will typically be employees of the organization and that organization is responsible for vetting and authenticating those employees. In other words, the IDP assumes responsibility for establishing that a person is who he or she claims to be. Within an organizational context, employees will likely have a user id and password that are provided by their organization and are used to gain access to the organizations systems and resources. Organizations typically manage this user ID and password information in an identity management or directory system. For example, Microsoft Active Directory (AD) is a common system used to manage this information.

In addition to the basic information about an employee, the organization must also be able to provide attributes that the federation has agreed are necessary to grant access to shared systems and resources. In a NHSIA federation, these attributes might include the type of worker (e.g., case worker) and licenses that worker possesses (e.g., Licensed Social Worker). The IDP uses a

SAML assertion, which is an XML document formatted in a standard way to communicate the identity attributes to the organization providing the shared system or resource.

### **2.6.6 Service Providers**

A service provider is an organization that manages some set of resources (e.g., applications, systems, or web services) that it has agreed to share with other members of the federation. Within the boundaries of that organization, access to those resources is controlled by some means. In many cases, access is controlled by a user ID and a password. This approach may be used to restrict who may access which resources. Once an organization agrees to make its resources available outside of its organizational boundaries to members of the federation, it is agreeing to grant some level of access to its resources to an entity external to the organization based on the identity provided by the entity's own organization (the IDP).

In one example, one organization has a web-based application that displays case information about its clients. A user from another organization accesses that web-based application via the Internet. Because the external user belongs to a federation partner, the web-based application verifies his or her identity with the IDP, and, based on the SAML assertion that contains the user's identity attributes, decides whether to grant access to the external user. The specifics of this type of interaction are detailed in the "GFIPM Web User-to-System Profile," a specification that defines a set of protocols and bindings for web browser-based interaction between users and resources across trust domains within a federation.

Because there are likely to be multiple organizations within a federation that are providing identities on behalf of users, service providers will typically maintain what is known as a discovery service to find the appropriate IDP for a particular user. This is also known as a "where are you from" service.

In another example, an organization has developed web services that it has exposed to members of the federation. This organization is considered a web service provider (WSP). Another organization, the web service consumer (WSC) has developed software and wishes to make use of the shared web service to perform some process or return some information. The security credentials of the user as well as the WSC must be communicated to the WSP. The GFIPM Web Services System-to-System Profile specification defines a complete web services protocol stack for basic system-to-system use cases. It addresses relatively low-level details such as the proper use of the WS-Security standard for building SOAP (Simple Object Access Protocol) messages that can be trusted within the context of the federation's Cryptographic Trust Model. It also describes how to properly compose and constrain web services industry standards for use within a GFIPM-based federation.

### **2.6.7 Identity Attributes**

The metadata that is used to exchange identity information defines attributes about users, system entities, information resources, information-sharing actions, and environmental conditions within the information-sharing federation. The information takes the form of trusted statements, or assertions, about subjects and is structured according to the SAML specification. Subjects include end users, organizations, resources (systems, web services, etc.) An assertion makes one or more statements about a subject that is based on attributes that have been predefined and

agreed upon by the federation members. Every assertion includes an assertion ID (a unique identifier), an issuer identification string and a creation time stamp. Assertions will also contain the relevant attributes about the entity. Assertions may also contain additional data, such as conditions that define when the assertion is valid. Finally, all assertions will contain a digital signature to ensure the integrity of the assertion data.

The following types of attributes will become key to the assertions and need to be defined and agreed upon by federation member organizations:

- **User Attributes.** These contain information about end users and may include basic identifying information, employment information, and other attributes that will ultimately control the types of access the user has to which shared resources.
- **Entity Attributes.** These contain information about the entities that are part of the federation and include basic identifying information as well as the role the entity plays within the federation (e.g., IDP, SP, WSC, WSP, CA).
- **Resource Attributes.** These attributes indicate the type of resource. Resource attributes may be used to categorize data in various ways that are then used to determine whether a user has the required privilege necessary to access the data. For example, a resource may be identified as client case information. A user who is a case manager may have a corresponding user attribute that indicates that he or she has permission to access client case information.
- **Action Attributes.** These attributes govern what actions are permitted by a user against a resource. For example, action attributes will control whether an individual user is permitted to update data or merely to read that data.

## 2.7 Implementation Considerations

### 2.7.1 How Many Federations are Needed?

A preliminary question that needs to be answered is what the scope of a NHSIA federation will include. This will be determined based on several questions, including:

- Who are the potential information sharing partners?
- Are some or all of the potential information sharing partners already part of an identity federation?

If the answer is “yes” to the second question, an organization should investigate the possibility of joining an existing federation. Otherwise, the potential information sharing partners will organize to form a federation following the guidelines established by GFIPM. However, because of the time and effort as well as the number of agreements to be attained in creating one, federations should include set of organizations that will most likely be sharing information regularly and on a real-time or near real-time basis. Otherwise, the cost to join the federation, which includes enabling the local identity management or directory system to generate SAML assertions as well as enabling target systems or web services to respond accordingly, could be cost prohibitive. Further, while it will depend on the nature of a jurisdiction, it is likely that federations will be established to include at least one jurisdiction and neighboring jurisdictions as well.

### **2.7.2 What Attributes are Used to Control Access?**

Another important consideration will be the attributes that are used to control access to resources. For example, federations must agree on the characteristics of an individual user that will determine whether he or she has access to any particular resources. Similarly, resources must be identified with the corresponding attribute to either allow or deny access. These decisions will be made by the federation member organization, agreed upon, and codified in the federation's trust agreement and implemented in a metadata specification. Entities in the federation will then use this metadata to create the appropriate assertions to identify and control entities and resources.

### **2.7.3 Connecting Federal, State and Local Hubs**

One of the key concepts that make up the NHSIA framework is the development and implementation of hubs. A hub is an information technology environment, within a particular jurisdiction, that is used to host shared services and shared data. It is likely that within a particular jurisdiction, (e.g., a county) the hub for that jurisdiction will be a member of the federation. This will provide county workers with single sign-on to resources in the hub and can also be used to provide web service security for services hosted in the hub and called by external software applications. It may also be beneficial to include neighboring jurisdictions in the federation if workers in that jurisdiction will require regular access to the hub.

Hubs may also exist at the State level as well. These hubs will contain services and data that are intended to be shared at the state level and by lower-level or neighboring jurisdictions. While it will certainly depend on the specific implementation, it is unlikely that a federation will extend very far beyond the borders of any one jurisdiction. In other words, it is unlikely that a state-wide federation will be required. Instead, it may be possible to control access to the State hub via the county hubs. This will require the creation of PKI certificates for the servers in the county hub, but will provide a means to control and protect the interactions between the state and county hubs without the need to create a state-wide identity management process or federation. Since the interaction is likely to be between a web service client and a web service, certificates can be used to authenticate clients using well defined standards. A similar approach may be adopted for a Federal hub if one is created.

### **2.7.4 How to Control Human Services Client Access**

One remaining issue that must be addressed is how access to resources will be granted to individual human services clients, acting on behalf of themselves or on behalf of their family members rather than on behalf of an employer who can verify the client's identity.

Some jurisdictions may have an existing system to provide user IDs and passwords to individual citizens. For example, the State of Wisconsin has a system that allows its citizens to obtain a "Wisconsin User ID" that provides access to state-provided Internet resources. Where these types of systems exist, they could be extended to provide citizen access to appropriate human services information system resources as well.

Alternatively, a COTS Web Access Management (WAM) solution could be deployed. This type of solution would provide identity management and authentication for individual citizens. Most employ a self-service approach in which a person chooses his or her own user name and password as well as security questions and answers to be used to recover lost passwords. WAM solutions are successfully deployed across the Internet to provide secure access to any number of resources. One drawback to this approach is that it is unlikely to provide a means for federated identity across organizational boundaries. In other words, depending upon the specific implementation, a citizen may be required to obtain a different user ID and password for each resource being accessed.

A final alternative is to create an Identity Provider as a part of a federation that will serve as the identity “home” for individual citizens. Under this alternative, a citizen’s identity could be verified by a human services worker during an intake processes. As a part of this process, data about the citizen is entered into a local system, and then, as a function of the NHSIA core, an entry is created in the Master Person Index. An identity record could be created in a directory as a part of creating the Master Person Index. The resulting directory entry, perhaps along with information in the Master Person Index, could then be used to create the SAML assertion to establish the citizen’s identity to SPs in the federation.

## 2.8 GFIPM References and Resources

Figure 1: GFIPM Document Map illustrates the available GFIPM documentation. More information about GFIPM guidelines and standards can be found at [www.gfipm.net](http://www.gfipm.net) (see Figure 2).

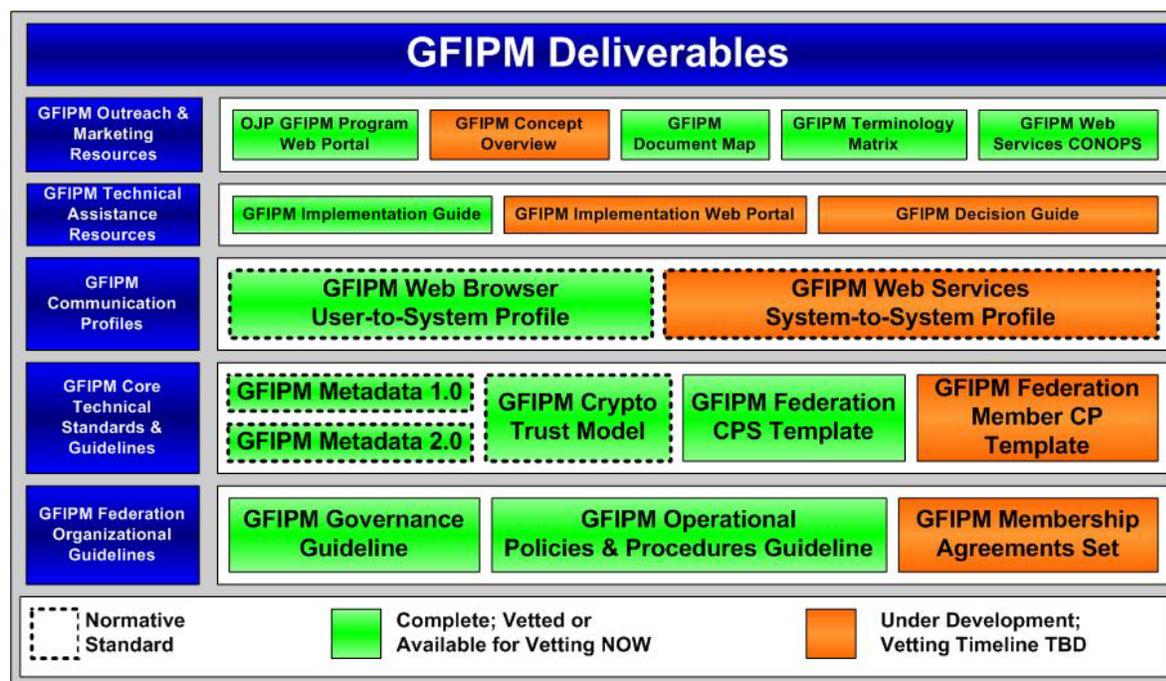
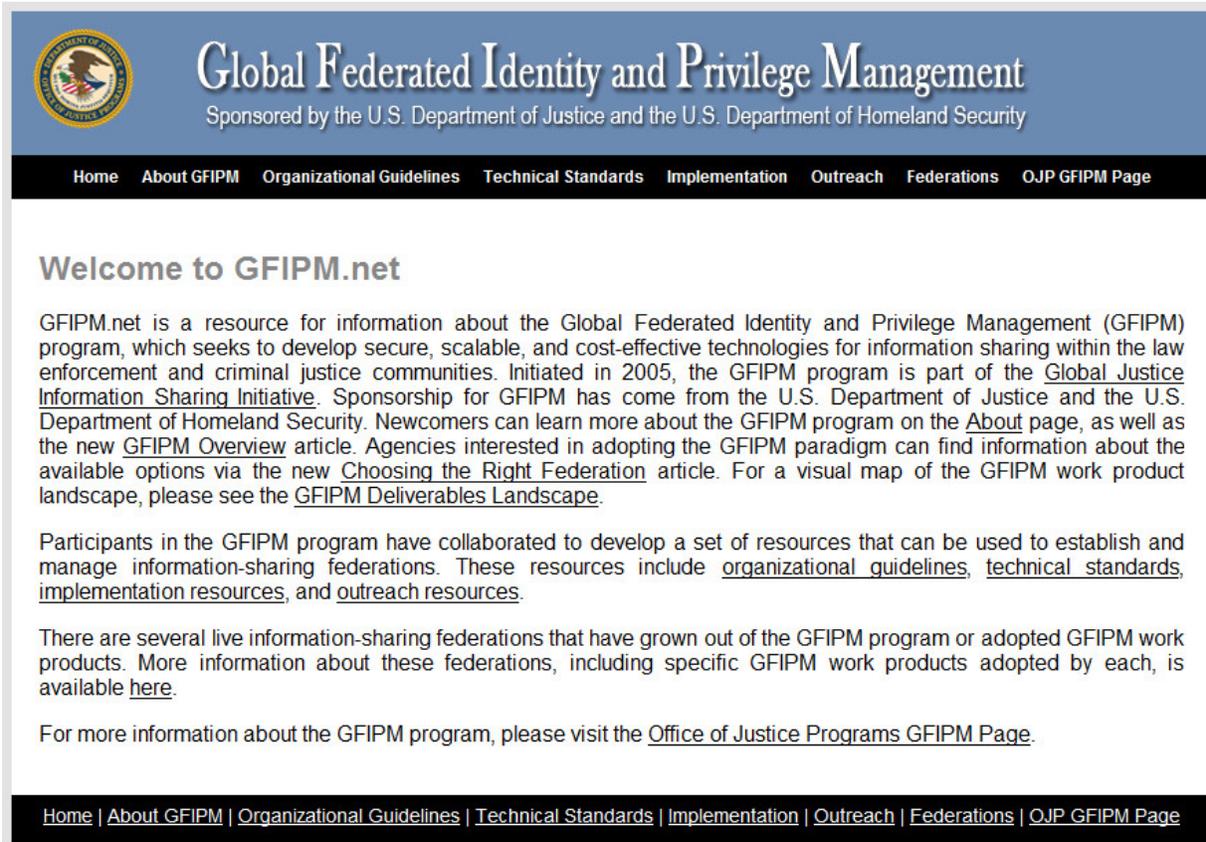


Figure 1: GFIPM Document Map



**Global Federated Identity and Privilege Management**  
Sponsored by the U.S. Department of Justice and the U.S. Department of Homeland Security

Home About GFIPM Organizational Guidelines Technical Standards Implementation Outreach Federations OJP GFIPM Page

## Welcome to GFIPM.net

GFIPM.net is a resource for information about the Global Federated Identity and Privilege Management (GFIPM) program, which seeks to develop secure, scalable, and cost-effective technologies for information sharing within the law enforcement and criminal justice communities. Initiated in 2005, the GFIPM program is part of the [Global Justice Information Sharing Initiative](#). Sponsorship for GFIPM has come from the U.S. Department of Justice and the U.S. Department of Homeland Security. Newcomers can learn more about the GFIPM program on the [About](#) page, as well as the new [GFIPM Overview](#) article. Agencies interested in adopting the GFIPM paradigm can find information about the available options via the new [Choosing the Right Federation](#) article. For a visual map of the GFIPM work product landscape, please see the [GFIPM Deliverables Landscape](#).

Participants in the GFIPM program have collaborated to develop a set of resources that can be used to establish and manage information-sharing federations. These resources include [organizational guidelines](#), [technical standards](#), [implementation resources](#), and [outreach resources](#).

There are several live information-sharing federations that have grown out of the GFIPM program or adopted GFIPM work products. More information about these federations, including specific GFIPM work products adopted by each, is available [here](#).

For more information about the GFIPM program, please visit the [Office of Justice Programs GFIPM Page](#).

Home | [About GFIPM](#) | [Organizational Guidelines](#) | [Technical Standards](#) | [Implementation](#) | [Outreach](#) | [Federations](#) | [OJP GFIPM Page](#)

Figure 2: GFIPM Home Page

### 3 Network and Infrastructure Security

Within the HIPAA Security Rule, the integrity (45 CFR §164.312(c)) and transmission security (45 CFR §164.312(e)) technical safeguards are related to one another in that a secure transmission also protects the integrity of the data. Also, both are intended to prevent an unauthorized party from reading or modifying information.

The integrity of information means that it has not been changed or altered in an unintended way. Compromising the integrity of information could happen through intentional means, for example, when an individual with malicious intent tries to destroy or falsify information. It could be through unintentional means, for example, when an employee makes a coding or transposition error while entering data. Or, it could happen through a system or media failure that causes some type of corruption in the data. Transmission security, on the other hand, is intended to safeguard the electronic transmission of information, through a network, from one system to another. A secure transmission implies that no one along the way was able to read or alter the data and that it reached its intended destination intact.

Consider the example of sharing a paper file folder with another person. If you send the folder through the mail, you have no guarantee that it reached its destination unaltered. Someone could have opened the envelope and changed the contents along the way. The mail carrier may have

left the envelope in the rain and the ink could have run and ruined the pages. When the other person is reading the files, you have no guarantee that someone else is not reading over his or her shoulder. Fortunately, a number of industry-standard security techniques can be implemented to ensure both integrity and transmission security.

For data that is stored in an IT system, proper use of access controls will ensure that no unauthorized person is able to access or modify data. Proper use of audit controls will ensure that mistakes are detectable and traceable back to a specific individual. In addition, data can be encrypted to ensure that it will not be understandable or even readable to anyone without the proper security keys to decrypt the data.

### **3.1 Encryption**

Encryption is a form of securing confidential or proprietary information as it is transmitted or while it is at rest within an IT environment. Encryption is based on mathematical formulae and comes in many forms, from secure e-mail to virtual private networks. The HIPAA Security Rule ((45 CFR §164.312(a)(1) and 45 CFR §164.312(e)(1)) specifically mentions encryption.

There are many forms of encryption suitable for different needs of an organization. The purpose of this section is to provide information about the use of encryption to secure confidential information in transit and at rest. It is not meant to provide a definitive standard because each organization will have varying needs for to securing data in transit and at rest.

### **3.2 Transmission Security**

Encryption plays an important role in transmission security. Secure transmission protocols are a part of most modern network infrastructures. These protocols automatically encrypt data as it is transmitted and automatically decrypt it as it is received. This ensures that someone “eavesdropping” on the transmission would be unable to understand the contents of the transmission. In addition, most transmission protocols also ensure the integrity of the data through built-in error checking and retransmission capabilities.

Many organizations are making confidential information available to clients or patients via company websites. Through the use of secure sockets layer (SSL), information transmitted between the end user and the website is secured. This type of security is commonly used to view bank statements, credit card statements, and order goods online. This requires deployment of a secure web server and appropriate software. Use of SSL will ensure that data exchanged between a human services organization and an end-user, via a web browser, will be encrypted during transmission. SSL may be used in addition to user authentication (e.g., user ID and passwords) to secure online access to confidential information.

Secure file transfer protocol (FTP) is another protocol often used to transmit large files between entities. It bulk-encrypts large files for transmission and allows for the secure transmission of data between computer systems. Use of Secure FTP will be appropriate for exchange of files between partner systems.

A virtual private network (VPN) can be used to secure the connection between two points. A VPN may be configured to require a user to authenticate before gaining access. In addition, a VPN will encrypt the exchange of data across the network. VPNs will allow users from one organization to securely access the network and resources of another organization assuming the proper credentials are presented. In the case of federated SSO, the credential will be the user ID and password defined by the user's home organization and shared within the federation.

### **3.3 Security of Data at Rest**

Data at rest includes data stored within an organization's IT infrastructure. The primary method to secure data at rest is through hard disk encryption. This can be accomplished through a variety of hardware or software approaches that encrypt stored data to prevent unauthorized access or misuse. Hard disk encryption is also important to consider if an organization uses laptops and hand-held devices to view or process data. Each of these is subject to theft, and use of hard disk encryption will prevent access to data stored on the device.

Of equal importance is the encryption of data contained in computer backups. Most commercial backup solutions provide a means to encrypt the contents of a backup. Organizations must apply the same levels of protect to data that is stored off-line as to data that is online. Encryption of backups provides this capability. In addition, it is critical that organizations also back up the keys used to encrypt the data and to store those keys separate from the backups themselves.

### **3.4 Remote Access**

Organizations have a number of methods to provide remote access to IT resources for users outside of the organization. Methods will vary based on the type of user and the access required. From the perspective of NHSIA, remote access is provided to both workers of the organization who need to connect from other locations and to workers who are members of other organizations.

The method used will vary based on the resources to be accessed, the means to access those resources and the capabilities of the organization. In many cases, an organization will provide access to any resource available to outside users through the use of a portal. A portal is a server that offers access to one or more applications through a single centralized interface. Most portals are accessed via a Web browser on the user's computer. The transmission of data between the user and the portal will usually be protected through the use of SSL.

In addition, access to the portal should be controlled via a user authentication credential. For jurisdictions that have implemented federated SSO, this will most likely be the user's ID and password from his or her home organizations. For jurisdictions that have not implemented federated SSO, each organization must provision credentials to external users in order to gain access to the portal.

Alternatively, external users may communicate through a secure communications tunnel which allows for secure transmission of information over public networks such as the Internet. Tunnels are typically established through virtual private network (VPN) technologies. Once a VPN tunnel

has been established between an external user and the organization's VPN gateway, the user can access the organization's computing resources through the tunnel. The VPN gateway will normally manage user authentication and access control for external users through the use of remote access services (RAS) or Remote-Access Dial-In User Service (RADIUS) servers.

## 4 Privacy and Confidentiality

Confidentiality is fundamentally about how we control information. It is not about hiddenness or concealment. Rather, it is about sharing the information we want to share and with whom.

The information technology can not only greatly facilitate the sharing of information; it can also greatly enhance both the security and the confidentiality of information in electronic form over that of paper-based information.

The HIPAA clearly defines how protected information is to be treated and what the exceptions are regarding release without patient authorization. State and federal laws other than HIPAA provide further guidance regarding the release or restrictions on release of information by identifying sensitive information. Legal requirements, though, represent the floor and not necessarily all of the privacy protections that an organization may choose to adopt or should adopt. In other words, organizations can go above and beyond the law providing greater protections than state or federal laws require.

Information may be shared for a number of reasons. First, information sharing may be expressly permitted (rather than prohibited) under federal or state laws. For example, HIPAA spells out the conditions and circumstances under which protected health information may be shared. Similarly, SAMHSA's 42 CFR, Part 2 specifies the conditions and circumstances, although more restrictive, under which substance abuse and mental health information can be shared. Finally, FERPA (Family Educational Rights and Privacy Act) spells out the conditions and circumstances under which educational information may be shared. These federal laws specify the minimum thresholds that must be complied with. States may add additional restrictions that go above and beyond these federal laws.

In addition to permitting the sharing of information, the law typically specifies the requirements for individual consent to the sharing of information. As a general rule, individuals must be able to deny consent, limit the consent to a particular period of time or a particular type of information, or withdraw consent altogether. For example, schools must generally have written permission from the parent or eligible student in order to release any information from a student's education record.

Finally, the sharing of information may be directed by court order. For example, FERPA allows schools to disclose records without consent in order to comply with a judicial order or lawfully issued subpoena.

While the specifics of client consent will vary from jurisdiction to jurisdiction, appropriate use of security mechanisms can be used to restrict access to client information in accordance with the client's consent decisions.

The first aspect is to electronically record information relative to an individual's consent to share his or her information. This should provide a means for an individual to consent to sharing information or to opt out. Further, it should provide a means for the individual to specify which types of information to share and with whom. Finally, consent must be set to expire as of a given date.

The second aspect is for the system where the information is stored to make appropriate information available based on the consent in place and the attributes of the person seeking to access the information. The attributes of the person should be a part of his or her federated single-sign on credentials.

## 5 Contingency Planning and Disaster Recovery

Organizations that are subject to HIPAA security and privacy laws will need to have a plan in place which specifies the steps to restore appropriate access to information after a major interruption in service. This is specified in 45 CFR §164.308, which, under the heading of contingency planning, requires organizations to have a data backup plan, a disaster recovery plan, and an emergency mode operations plan.

Because information system resources are so essential to an organization's success, the vast majority of organizations will undoubtedly have these plans in place. However, with the sharing of information and the use of common services across organization boundaries, contingency planning becomes even more important in that an outage in one organization may impact one or more other organizations.

According to the National Institute of Standards and Technology (NIST) Special Publication 800-34, Contingency Planning Guide for Federal Information Systems (2010), a disaster recovery plan typically applies to major physical disruptions to service that deny access to the primary facility infrastructure for an extended period. A disaster recovery plan is an information system-focused plan designed to restore operability of the target system, application, or computer facility infrastructure at an alternate site after an emergency. A disaster recovery plan is usually part of an overall business contingency plan that addresses the recovery of the business functions of the organization in the wake of a disaster.

A disaster recovery plan must begin with a business impact analysis that seeks to determine the impact of a system disruption to the functioning of the business. This will include the maximum downtime that an organization can tolerate while still maintaining the mission along with the resource requirements to resume operations.

NIST presents the following classification scheme for determining the recovery strategy relative to the importance of the resource to business operations (see Table 5-1):

**Table 5-1. FIPS 199 Category Backup and Strategy Examples<sup>3</sup>**

<b>Availability Impact Level</b>	<b>Information System Target Priority and Recovery</b>	<b>Backup / Recovery Strategy</b>
Low	Low priority - any outage with little impact, damage, or disruption to the organization.	Backup: Tape backup Strategy: Relocate or Cold site
Moderate	Important or moderate priority - any system that, if disrupted, would cause a moderate problem to the organization and possibly other networks or systems.	Backup: Optical backup, WAN/VLAN replication Strategy: Cold or Warm site
High	Mission-critical or high priority - the damage or disruption to these systems would cause the most impact on the organization, mission, and other networks and systems.	Backup: Mirrored systems and disc replication Strategy: Hot site

## 5.1 Data Backup Methods and Offsite Storage

System data should be backed up regularly. Policies should specify the minimum frequency of backups based on data criticality and the frequency that new information is introduced. Data backup policies should designate the location of stored data, file-naming conventions, media rotation frequency, and method for transporting data offsite. Data may be backed up on magnetic disk, tape, or optical disks, such as compact disks (CDs). The specific method chosen for conducting backups should be based on system and data availability and integrity requirements. In addition, given that backups are likely to contain confidential or otherwise sensitive information, backups should be encrypted.

It is good business practice to store backed-up data offsite and at a location sufficiently distant from the primary system location. Commercial data storage facilities are specially designed to archive media and protect data from threatening elements. Most commercial storage facilities will transport media in environmentally controlled vehicles from the primary site to the storage location.

## 5.2 Emergency Mode Operations Plan

Although major disruptions with long-term effects may be rare, they do in fact occur and organizations should develop a plan that includes a strategy to recover and perform system operations at an alternate facility for an extended period.

In general, three types of alternate sites are available. Many organizations have multiple, geographically diverse locations. It may be possible to for one location to assume the responsibilities of another during an extended outage. Of course, this should be detailed in the disaster recovery plan and should be tested periodically. Often times, organizations create

---

<sup>3</sup> NIST, Federal Information Processing Standards (FIPS) 199, Standards for Security Categorization of Federal Information and Information Systems (2004)

reciprocal agreements with other organizations and agree to provide recovery resources and facilities in the event that one or the other suffers an outage. Finally, commercially-leased facilities are available that will provide everything from an empty building in which to stage a recovery to a fully operational data center. Regardless of the type of alternate site chosen, the facility must be able to support system operations as defined in the plan. The three alternate site types commonly categorized in terms of their operational readiness are cold sites, warm sites, or hot sites.

## Appendix A: NHSIA Architecture Principles Related to Information Security

Architecture principles are general rules and guidelines, intended to be enduring, that inform and support the way in which an organization sets about fulfilling its mission. In the case of Information Security, the mission is to protect the confidentiality, integrity and available of information resources across the enterprise. The following principles are intended to establish the basic foundation for implementing information security within a particular organization, or in the case of NHSIA, across organizational boundaries, so as to facilitate secure collaboration and information sharing<sup>4</sup>.

### Defense in Depth

The principle of Defense in Depth states that the failure of any single component in the security environment must not compromise the entire environment. Defense in depth is typically implemented through the use of multiple security perimeters between public networks and internal, protected resources.

### Least Privilege

The principle of Least Privilege states that users and other consumers of resources must operate using the least set of privileges necessary to complete the job. This is because security risks increase with the amount of access a user or resource consumer is granted. Risks can stem from misuse of privilege, unintentional destructive actions, compromised accounts and systems, among others. Least privilege is typically implemented by granting access to system functions and data based on a user's role or attributes about that user.

### Security as a Service

The principle of Security as a Service states that technology solutions must be designed to consume common security services where possible as opposed to implementing custom security logic and replicating copies of security data. In other words, consistent, shared security services allow multiple solutions to share common security logic, features, policies, and identity information thereby eliminating redundancies and associated risks. It also enables more effective management of security in the IT environment.

### Secure Web Services

The principle of Secure Web Services states that adopting SOA and Web Services must not negatively impact system security, or negate the use of infrastructure-based security services. The provider organization's security infrastructure must protect Web Services end-to-end regardless of the number of intermediaries that exist between the consumer and provider.

---

<sup>4</sup> Readers should consult *NIST Special Publication 800-53*, "Recommended Security Controls for Federal Information Systems and Organizations," from the National Institute of Standards and Technology for additional information on implementing security within an individual organization (<http://csrc.nist.gov/publications/nistpubs/800-53-Rev3/sp800-53-rev3-final.pdf>)

Web Service security is typically provided through the use of open security and interoperability standards, such as those based on WS-Security and WS-I Basic Security Profile.

**Active Threat Detection and Analysis**

The principle of Active Threat Detection and Analysis states that the security infrastructure must be capable of detecting abnormal behavior and adapting accordingly to protect vulnerable resources. The security infrastructure must monitor for fraudulent use and abnormal behavior and take appropriate measures, such as sending alerts and suspending accounts.

**Complete Audit Trail**

The principle of Secure, Complete Audit Trail states that the security system must be able to identify when changes have been made to data within the organization, what changes have been made, and by whom. Many legal and regulatory concerns (e.g., HIPAA) require organizations to maintain a complete and secure audit trail. The infrastructure must be able to gather audit records from various sources into a secure repository where they can be collectively monitored and reviewed.

**Data Security**

The principle of Data Security states that the confidentiality, integrity, and availability of data must be ensured at all times. Data is a valuable resource and there are tremendous risks to an organization when security is compromised.

**System Availability**

The principle of System Availability states that systems, applications, services, etc., must be adequately protected to ensure their intended degree of availability, but not overly constrained by security measures to unnecessarily impede normal operations. This is particularly true when IT resources are shared across organizational boundaries. It is possible to very easily and completely secure a resource by preventing all to access it. However, too much security can have the effect of preventing legitimate access to systems and data. Security measures, and the overhead of using the system, should not outweigh their usefulness. Of course, risks must be assessed in order to properly determine the security measures necessary for a system.