United States
Department of Justice

U.S. Department of Justice's Global

Federated Identity and Privilege Management (GFIPM)

Implementation Guide

Version 1.0

Global Security Working Group

September 14, 2010

This project was supported by Grant No. 2008-DD-BX-K520 awarded by the Bureau of Justice Assistance, Office of Justice Programs, in collaboration with the U.S. Department of Justice's Global Justice Information Sharing Initiative. The opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the U.S. Department of Justice.

About the Document

Justice organizations are looking for ways to provide secured access to multiple agency information systems with a single logon. The Global Federated Identity and Privilege Management (GFIPM) initiative, developed by the U.S. Department of Justice's Global Justice Information Sharing Initiative, provides the justice community with a security and information sharing architecture that is based on an electronic justice credential. This standards-based justice credential can be used to securely connect law enforcement and public safety personnel to interagency applications and data over the Internet.

Background: The GFIPM framework provides the justice community and partner organizations with a standards-based approach for implementing federated identity. Common use of these standards across federation systems is essential to their interoperability. Leveraging the Global Justice XML and National Information Exchange Model (NIEM), a standard set of XML-based elements and attributes (referred to collectively as GFIPM metadata) about a federation user's identities, privileges, and authentication can be universally communicated.

Value to the Justice Community:

1. **User Convenience:** Users can access multiple services using a common set of standardized security credentials, making it easier to sign on and access applications and to manage account information.
2. **Interoperability:** By specifying common security standards and framework, applications can adopt interoperable security specifications for authentication and authorization.
3. **Cost-Effectiveness:** GFIPM facilitates information sharing by using a standardized XML-based credential that includes information about each user's identity and privileges. This reduces the cost and complexity of identity administration required to access applications and vet users.
4. **Privacy:** GFIPM can reduce the propagation of personally identifiable information, reduce the redundant capture and storage of personal identity information, and depersonalize data exchanges across domains using privacy metadata.
5. **Security:** A federation model can improve the security of local identity information and data in applications by providing a standardized approach to online identities between agencies or applications.

Contents: The GFIPM Implementation Guide contains detailed instructions for implementers of identity providers (IDPs) and service providers (SPs), which are the two types of systems that participate in user-to-system transactions as specified in the

GFIPM Web Browser User-to-System profile. The document covers all aspects of IDP and SP implementation, from requirements analysis to system deployment.

Target Audience: The target audience for this document includes managers and technical representatives of prospective GFIPM participant organizations who are planning to implement an identity provider (IDP) and/or a service provider (SP) within a GFIPM federation. It also includes vendors, contractors, and consultants who are required to establish technical interoperability with GFIPM standards as part of their project or product implementation.

Table of Contents

Table of Tables

Table of Figures

1. Introduction

This document comprises a collection of accumulated expertise and insights about how to implement information systems that are interoperable with a Global Federated Identity and Privilege Management (GFIPM) federation. It addresses four system implementation scenarios: identity provider, service provider, Web service consumer, and Web service provider. The remainder of this section provides introductory information about the GFIPM program and the scope of this document. Sections 2 through 5 provide guidance on each of the implementation scenarios listed above, and Section 6 contains additional insights not covered in previous sections.

1.1 GFIPM Background

The Global Federated Identity and Privilege Management (GFIPM) Security Interoperability Demonstration project was initiated in 2005 by the Global Security Working Group (GSWG) to investigate the concept of federated identity and privilege management (FIPM) as a candidate solution for information sharing interoperability challenges that have arisen in the law enforcement and justice community. Jointly funded by the Department of Justice (DOJ) and the Department of Homeland Security (DHS), the demonstration project successfully met all of the initial objectives and resulted in the creation of several valuable products, including a set of draft interoperability specifications, a freely available implementation of GFIPM middleware, and an operational pilot GFIPM federation. As a result of this success, the Global Justice Information Sharing Initiative (Global) recognized GFIPM as "the recommended approach for development of interoperable security functions for authentication and privilege management for information exchange among cross-domain justice information sharing systems."[1]

The GSWG established a GFIPM Delivery Team (GFIPM DT) to evolve the initial GFIPM products and specifications into a fully vetted and production-quality capability that can be leveraged across federal, state, local, and tribal justice and public safety communities. The GFIPM DT currently acts as the governance body for a GFIPM pilot operational federation and for day-to-day decisions on developing the GFIPM implementation framework.

For purposes of this document, the term "operational federation" refers to a GFIPM-compliant federation that provides live data to authorized law enforcement users associated with a small number of "early adopter" organizations. This operational pilot federation is the National Information Exchange Federation (NIEF). NIEF allows law enforcement organizations to engage in the process of sharing controlled unclassified information (CUI) with each other in a cost-effective and scalable

[1] This recommendation was made by GAC at its November 2006 meeting.

manner. The Georgia Tech Research Institute (GTRI) is currently acting as the Federation Manager for NIEF with the GFIPM DT acting as the NIEF Board of Directors. (The roles of Federation Manager and Federation Board of Directors are defined in the GFIPM Governance Guideline document [GFIPM Gov].)

Additional information on Global and GFIPM can be found at http://it.ojp.gov/gfipm.

1.2 About This Document

The purpose of this document is to provide implementers with essential information and instructions on how to integrate information systems into a GFIPM federation such as NIEF. These systems may include existing user databases or directories, and/or various existing or planned databases, portals, or other mission information resources. The document provides an organized collection of knowledge and insights gained by the initial GFIPM implementers; it will evolve over time to capture additional best practices and insights as they are developed.[2]

Note that some of the key technologies used within GFIPM are relatively new open standards (e.g., OASIS SAML 2.0, WS-Security, and others). Given that most justice-community organizations will presumably be interested in integrating their existing, older legacy systems with these new technologies in a GFIPM federation, it is likely that they will encounter unique implementation challenges for which there is no established "best" solution. This document provides implementers with guidance based on initial experience, but it does not address all possible implementation scenarios. For the benefit of the entire GFIPM stakeholder community, implementers are encouraged to submit summaries of lessons learned during their GFIPM implementation process to gfipm-support@lists.gatech.edu, so that they can be incorporated into future versions of this document.

1.3 Target Audience

The target audience for this document includes managers and technical representatives of prospective GFIPM participant organizations that are planning to implement an identity provider (IDP) and/or a service provider (SP) within a GFIPM federation.[3] It also includes vendors, contractors, and consultants who are required to establish technical interoperability with GFIPM standards as part of their project or product implementation.

[2] In the near future, this document will be converted into a Web-based resource for implementers.

[3] Future versions of this document will also contain content about how to implement Web services components within a GFIPM federation.

1.4 GFIPM Normative Technical Standards

The implementation guidance in this document pertains to the use of the following GFIPM normative technical standards.

- ***GFIPM Metadata Specification 2.0 [GFIPM Meta]***—Defines a vocabulary of attributes that can be used to describe facts about federated users and federated service endpoints. The GFIPM Metadata standard is an essential part of the GFIPM concept of federated identity and privilege management.

- ***GFIPM Cryptographic Trust Model 1.0 [GFIPM Trust]***—Defines a cryptographic trust model that provides a technical basis for all trusted communications between service endpoints in a GFIPM federation.

- ***GFIPM Web Browser User-to-System Profile 1.0 [GFIPM U2S Profile]***—Specifies technical interoperability requirements for connection to a GFIPM federation as an IDP or SP.[4]

- ***GFIPM Web Services System-to-System Profile 1.0 [GFIPM S2S Profile]***—Specifies technical interoperability requirements for connection to a GFIPM federation as a Web service consumer (WSC) or Web service provider (WSP).[5]

Each of these standards is evolving. The maintainers of this Implementation Guide will make every attempt to keep its content up to date with respect to changes in the GFIPM technical standards, but in some cases this document may not reflect the most current implementation best practices.

Also note that the normative standards listed here do not cover non-IT topics such as governance, policy, or other nontechnical interoperability requirements.[6] Please review [GFIPM U2S Profile] and/or [GFIPM S2S Profile] in conjunction with this document prior to beginning the on-boarding process.

[4] Security Assertion Markup Language (SAML) 2.0 is a product of the Organization for the Advancement of Structured Information Standards (OASIS) Security Services Technical Committee (SSTC).

[5] The GFIPM Web Services System-to-System Profile is still in development, and this document does not yet contain any implementation guidance related to GFIPM Web services. The document will be revised to include Web services implementation guidance after GFIPM early adopters have gained enough experience and lessons learned to provide such guidance.

[6] GFIPM and NIEF stakeholders are in the process of developing and vetting a set of GFIPM federation governance and policy documents to address these topics.

1.5 Acknowledgements

The primary developer of this document was the Georgia Tech Research Institute (GTRI). Other contributors include the Criminal Information Sharing Alliance (CISA) and the Pennsylvania Justice Network (JNET). Specifically, the following individuals have made significant contributions to this document.

- John Wandelt—GFIPM Project Director, GTRI
- Stefan Roth—Research Scientist, GTRI
- Matt Moyer—GFIPM Technical Lead, GTRI
- Jeff Krug—Research Scientist, GTRI
- John Ruegg—Chair, GFIPM Delivery Team, and Global Security Working Group
- Glen Gillum—Executive Director, CISA
- James Dyche—Architecture Manager, JNET

1.6 References

The following documents or resources are referenced throughout this Implementation Guide. These documents describe companion standards or systems that may be used by GFIPM implementers.

Document ID	Document Name and URL
GFIPM Map	GFIPM Document Map Available at http://it.ojp.gov/gfipm.
GFIPM Terms	GFIPM Terminology Matrix Available at http://it.ojp.gov/gfipm.
GFIPM Gov	GFIPM Governance Guidelines Available at http://it.ojp.gov/gfipm.
GFIPM OPP	GFIPM Operational Policies and Procedures Available at http://it.ojp.gov/gfipm.
GFIPM Meta	GFIPM Metadata Standard 2.0 Available at http://it.ojp.gov/gfipm.
GFIPM Trust	GFIPM Cryptographic Trust Model 1.0 Available at http://it.ojp.gov/gfipm.
GFIPM U2S Profile	GFIPM Web Browser User-to-System Profile 1.0 Available at http://it.ojp.gov/gfipm.
GFIPM S2S Profile	GFIPM Web Services System-to-System Profile 1.0 Available at http://it.ojp.gov/gfipm.
SAML2	Security Assertion Markup Language (SAML) 2.0 http://wiki.oasis-open.org/security
GFIPM Demo	GFIPM Security Interoperability Demonstration Project Final Report, Georgia Tech Research Institute, 30 August 2007 Available at http://it.ojp.gov/gfipm.

Document ID	Document Name and URL
HTTPD	Apache HTTP Server Project The Apache Software Foundation http://httpd.apache.org/
Tomcat	Apache Tomcat Java Servlet and JavaServer Pages The Apache Software Foundation http://tomcat.apache.org/
Java JDK	Java Development Kit Sun Development Network http://java.sun.com/

Table 1: References for GFIPM-Specific Standards

2. Implementing an Identity Provider

This section outlines the steps necessary to implement a GFIPM Identity Provider (IDP).

A GFIPM IDP collects information (typically from an existing identity store) about a local user and generates corresponding user metadata when a user attempts to connect to a local or remote GFIPM Service Provider.

This section is organized into the following steps to help you implement an IDP. It covers topics from the early design to final deployment.

1. Developing a GFIPM information sharing plan
2. Submitting a request for federation membership
3. Choosing a product for building an IDP
4. Implementing your organization's IDP
5. Writing a test plan for an IDP
6. Deploying an IDP in a test environment
7. Executing the test plan for your IDP
8. Deploying your IDP in an operational federation

2.1 Developing a GFIPM Information Sharing Plan for an IDP

During this process, you will accomplish the following:

1. Develop a list of federation resources to which you want your users to have access.

2. Identity your local users and collect all sources of information about them.

3. Design the metadata to describe your users.

4. Fill out a Local Attribute Mapping Form to map local attributes about your users into GFIPM metadata attributes.

If your organization has multiple attribute stores and/or authentication systems, you may need to consider implementing multiple IDPs at your site. Alternatively, you may wish to consider using a virtual LDAP product to consolidate your multiple sources of user data into one source. This situation may be especially applicable if the information in the attribute stores cannot be merged or different sets of users must stay with separate authentication or authorization systems. In the case of multiple IDPs at your site, the steps in this section should be performed for each IDP.

2.1.1 Discover Federation Resources

This section guides you through the process of discovering available resources in your federation and determining which of them may be of value to your organization's users. Additionally, you need to consider the future as your federation adds more Service Providers with more and more resources. You will want your users to have access to the appropriate resources in the future, preferably without having to add capabilities to your IDP. At the end of this section, you should have a list of access control policies for which your users need to qualify. This list will be used in subsequent sections to help you implement an IDP that can assert the appropriate metadata.

The following steps will help you accomplish the above goals:

1. Survey the current landscape of available federation resources.

2. Determine which existing resources would be of value to your users. Also consider what other future resources you might want for your users and what the access requirements for those resources might be.

3. Review the access control polices for the resources you would like to use.

To explain these concepts and help you put them into action, resources in the NIEF federation are used as examples throughout this section.

2.1.1.1 Survey Federation Resources

This section describes how to examine the list of federation resources and generate a list of resources that you want your users to access.

To allow prospective federation members to determine what members' resources would be of value to their users, the federation should maintain a registry of existing resources, including a description and access control policy for each resource. The information should be available from your federation manager.

For NIEF, the resource list is available at http://gfipm.net/users/ with a section for each current federation member.

The following are a few examples of NIEF resources and their access control policies with implementation guidelines and explanations.

Resource Example Number 1:

As can be seen from the resource list, each resource has an access control policy. The policy may be as simple as for the Arizona Counter Terrorism Information Center resource, provided by the CISAnet portal:

> **Access Control Policy**
>
> Any user with a valid federation login may access this resource. In addition, sufficient audit data is required for all users.

In the above requirement, a "valid federation login" refers to a user who is able to log in to an IDP of any NIEF federation member. In effect, this is the absolute minimum requirement possible for being granted access to a resource. Some portals provide useful public resources in this manner and the users gain from seeing a collection of such resources in one location, but such public resources do not provide law enforcement data.

Because some Service Providers (SPs) may actually require that a user have permission to view public resources, we recommend that IDPs assert the "Public Data Self Search Home Privilege Indicator" in the user's metadata. Service Providers are free to assume that all users have public permission but are not required to and may not implement such an assumption.

The "sufficient audit data" includes information that can be used to uniquely identify a user and is stored in the audit log files of a service provider, portal, and/or application.

Resource Example Number 2:

Most resources have more complex policies; for example, the resource Texas Criminal Law Enforcement Online (CLEO) provided by the CISAnet Service Provider:

Access Control Policy

Any user who is a sworn law enforcement officer may access this resource. In addition, sufficient audit data is required for all users.

The definition of a sworn law enforcement officer (SLEO) is as follows. He/she is:

- A full-time employee of a state-recognized law enforcement agency.
- Authorized to make an arrest (has the authority).
- Certified by a state certifying authority (e.g., Peace Officer Standards and Training [POST]), or equivalent.

  OR

- A full-time employee of a state-recognized law enforcement agency, acting on behalf of a SLEO, in performance of his or her assigned duties.

An IDP indicates that a user is a sworn law enforcement officer by asserting the "Sworn Law Enforcement Officer Indicator" as "true" in that user's metadata. Many of the more useful resources available require that the user be a sworn law enforcement officer, so you as the IDP developer are strongly urged to implement the necessary code to determine whether your user qualifies.

The requirement for the audit data is described in example Number 1 above.

Resource Example Number 3:

Other resources may have very restrictive access policies, as demonstrated by the Texas Criminal Law Enforcement Reporting and Information System (CLERIS) resource:

- Be a sworn law enforcement officer.
- Have an agency ORI code.
- Have either the criminal investigative home data search privilege OR a combination of the criminal intelligence home data search privilege and 28 CFR certification.

- Identity-proofing assurance is NIST level 4, and electronic identity assurance is at least NIST level 3.

- In addition to the policy specified above, sufficient audit data is required for all users.

The IDP must assert the sworn law enforcement officer indicator as demonstrated in example Number 2 above.

To access CLERIS, a user must have an ORI code, which must be asserted by the user's IDP in the user's metadata in the attribute "Employer ORI." Many resources available from GFIPM participants require an ORI code, so IDP developers are strongly urged to supply them.

Further, a user's home IDP must assert the attribute "Criminal Investigative Data Self Search Home Privilege Indicator" or must assert the attributes "Criminal Intelligence Data Self Search Home Privilege Indicator" and "28 CFR Privilege Indicator."

The requirement for sufficient identity proofing (asserted as the attribute "Identity Proofing Assurance Level Code") refers to the level of assurance that the person assigned this electronic identity is actually the user. The levels of assurance vary from no requirements to appearing in person with certain photo IDs or supplying certain identification assurances from a remote location that are inspected or validated to a specified degree by the issuing authority.

The requirement for electronic identity authentication assurance level (asserted as the attribute "Electronic Authentication Assurance Level Code") refers to the level of confidence in the asserted digital identity. The asserted level depends on the IDP's authentication mechanism, which ranges from simple and relatively insecure methods (such as username/password authentication) to more secure methods (such as two-factor authentication with a hardware crypto token). The requirement for the audit data is described in example Number 1 above.

2.1.1.2 Determine Valuable Resources

At this point, you should have a list of resources for your users. After finishing the survey of federation resources in the previous section, you should have a feel for the types of resources that you might add to that list if they become available in the future.

If your federation plans to add new members, or if prospective members are already in the process of joining your federation, you may wish to inquire about the resources that may soon be added to the federation via the addition of those new federation members. If you determine that these additional resources would be

valuable for your users, you may wish to plan ahead for asserting their required attributes even before they come online in the federation.

For example, NIEF is a relatively new federation, with few members at this time. However, as of this writing, several new members are in the process of submitting their application packages, so it is likely that NIEF will increase significantly in size in the near future. A new IDP implementer should contact the NIEF federation manager for advice on expected new members and resources, so that the IDP's users can gain access to these new resources without the implementer having to add support for more attributes at a later date.

For example, on the CISAnet Service Provider, you saw that several resources (such as NM Law Enforcement Information Network with Corrections or NM Complete Arrest Information) require the Criminal History privilege. Suppose you determine that these would be useful resources for your users, except that your users are not necessarily interested in the New Mexico data but would be interested in other states' similar resources. You should add "Criminal History Resources" to your list of desired resources. Further, you see that some of these same resources also require the NCIC Criminal History Certification, so you add "NCIC Criminal History Certification Resources" to your list of desired resources.

2.1.1.3 Review Access Control Policies

Next, review your list of desired resources and determine which privileges, certifications, indicators, and electronic identifications you must support in your IDP to allow your users to access those resources. This list represents the privileges, certifications, and indicators that you would like to assert for your users, but not necessarily those that you can assert. At this point, it is better to build an over-inclusive list than a list that may be missing some desired permissions.

After you identify your users from a GFIPM perspective in the section below (Section 2.1.2), this list will be used in Section 2.1.3 to design the GFIPM user metadata to be built by your IDP.

2.1.2 Identify Local Users

This section guides you through a process to collect all known information about your organization's users and collect it for use as a basis for a GFIPM Identity Provider (IDP).

To implement a GFIPM IDP, you must gather existing sources of information about your local users. These sources may consist of a user directory, a database system, applications that manage user identities, and organizational policies and other documents.

A user directory may be implemented as LDAP or Active Directory or some other in-house or commercial system. A database may be implemented as a system such as Oracle or SQL Server or one of many other commercial or open-source systems. Other sources may include user applications such as criminal information systems, case management systems, or records management systems. These systems will then serve as the providers of local user information to the federation member's IDP.

Each source about local users should provide information such as the following:

- Name, address, phone number
- E-mail address
- Unique user ID
- Home organization, employer, assignment, job classification
- Certifications and clearances
- Permissions and privileges
- Electronic or digital identity

In addition, there may be other, more indirect sources of information about users. Organizations typically have documented security policies. Users may also be required to sign user agreements, which typically specify levels of training or qualifications for the user. These may specify conditions of employment such as background checks, user qualifications, certifications, or security clearances.

Three specific instances of these types of documents include the following:

1. ***Local Security Policy Document—***A document describing the security policy that is currently in place within your organization.

2. ***Local User Agreement Document—***A document describing the terms and conditions to which your users must agree as a prerequisite for using an electronic identity issued by your organization.

3. ***Local User Vetting Policies and Procedures Document—***A document describing the user vetting policies and procedures that are currently in place within your organization.

Implicit or derived information from the above documents can add to the knowledge base about your users, either individually or as a group. At this point, you should collect these documents from your organization and use them as a basis for additional knowledge about your users. In addition to serving as sources of information about users, the three documents listed above will be used during your organization's federation application process (see [GFIPM OPP]) described in Section 2.2.

An example of information derived from a security policy is the following derivation rule used by CISAnet, which is a federation member in NIEF:

> All CISAnet users have 28 CFR training as a documented organizational policy. These users have the CISAnet Criminal Intelligence permission. While the 28 CFR training information is not stored in the local identity management system, the policy is used as a basis for the CISAnet IDP to assert the "28 CFR Privilege Indicator" in the GFIPM user metadata. Furthermore, the CISAnet IDP also asserts the attribute "Criminal Intelligence Data Self Search Home Privilege Indicator" for its users.

After you finish this section, your GFIPM Information Sharing Plan should include details about all your sources of user information and also document details about which specific information is available for users from each source. This information will eventually appear in your Local Attribute Mapping Form; see Section 2.1.4 for more information.

2.1.3 Design User Metadata

This section presents instructions on how your organization should design its user metadata, based on your list of desired federation resources (from Section 2.1.1) and the information about your users (from Section 2.1.2).

The GFIPM Metadata Specification includes a standard set of informational attributes that can be asserted for a user. Example attributes are a user's name, phone number, title, permissions, etc. These attributes are collected by an IDP, assembled into a SAML assertion, and securely transmitted to an SP on behalf of a user.

Your federation manager should provide you with advice on which metadata attributes are required or recommended for assertion by IDPs in your federation.

Taking into account the GFIPM definition of each metadata attribute you want to assert, determine whether and how you can truthfully assert it for your users based on your locally available user information. Assertions can be based either on explicit local attribute data (stored in a user repository) or on implicit assumptions about users based on local policies.

You may encounter a situation in which you want to assert an attribute but are unable to assert it based on locally available information. In this case, you have two choices—do not assert it or collect and store the data necessary to assert it.

Remember that the basis for asserting each user attribute must be documented in the Local Attribute Mapping Form (Section 2.1.4).

At the end of this process, you should have a list of GFIPM metadata attributes that your IDP will assert, along with the precise GFIPM definition for each of those attributes.

Each metadata attribute that your IDP is able to assert must be asserted with a valid value in the SAML assertion. The values must be either extracted from one of your local data sources or validly derived for each user. In addition, how you assert these attributes (i.e., the data source or reasoning you used) must be documented in your Local Attribute Mapping Form to be prepared in Section 2.1.4.

As defined by your federation manager, required metadata attributes are mandatory based on their use to uniquely identify a user and to audit transactions. Your IDP must assert these metadata attributes.

Strongly recommended attributes are those attributes used by many Service Providers or resources in their access control policies. Assertion of these attributes typically leads to more data access opportunities for users. Your IDP should assert these metadata attributes if possible.

The other listed attributes are recommended, which means that useful resources tend to use them in their access control policies. Your IDP should assert these metadata attributes if possible.

2.1.4 Fill Out a Local Attribute Mapping Form

This section will help you fill out a Local Attribute Mapping Form to map local attributes about your users into GFIPM metadata attributes. The Local Attribute Mapping Form for your IDP is later used as part of your request for federation membership.

The Local Attribute Mapping Form is briefly described in the Operational Policies and Procedures document [GFIPM OPP]:

> A document describing how the organization plans to map its local policies and locally stored user attributes into attributes conforming to the GFIPM Metadata standard.

When the federation manager reviews your application package, he or she will provide a copy of your Local Attribute Mapping Form (for an IDP) to all existing members for review and comment.

The Local Attribute Mapping Form should be written as a spreadsheet (i.e., in Microsoft Excel). A template of this form is included with the membership

application forms provided by the federation manager when you request to join the federation.

Before editing the file, you should rename it to include your IDP name in the file name.

Table 2 contains an example of the design of the spreadsheet, including the headers followed by five rows describing attribute mappings. Note that these examples are from different members, so their derivations are not related to each other.

GFIPM Attribute Map Identity Provider Name: <Your Organization>			
Semantic Intent Of Mapping	**Mapping Rule From Local Attribute/Policy to GFIPM Metadata**		
	GFIPM Metadata Attribute	**Mapping Method**	**Local Source Attribute**
First name of user	Given Name	Calculated from Local Attribute CN (Common Name) from ABCD Directory	Take substring to the first space in CN starting from the left.
The unique federation-wide identifier for this user	Federation ID	Fixed text plus Local Attribute (e-mail address) from the ABCD Directory for this user	"GFIPM:IDP:ABCD:USER:" + e-mail
ABCD does not have an attribute to indicate whether a user is a public safety officer. This derivation should yield a reliable indicator if the user is a public safety officer or working at the behest of one.	Public Safety Officer Indicator	Derived from Local Attributes in Directory	"true" if (departmentNumber contains 'Police' OR 'Patrol' OR 'Sheriff' OR '911') OR (title contains 'Officer' OR 'OFFICER' OR 'Dispatch' OR 'Sheriff' OR 'District' OR 'Patrol' OR 'Lieutenant' OR 'Sergeant') OR (postalAddress = 'police')
Derive if a user is legitimately a sworn law enforcement officer even though ABCD does not store this information in our directory	Sworn Law Enforcement Officer Indicator	Derived from Local Attribute Criminal Intelligence permission	All our SLEO users who go through 28 CFR training are given the Criminal Intelligence permission in our directory. If a user has this permission, our IDP will assert this indicator.
The contact e-mail for questions about ABCD or the identity information in the ABCD SAML assertion. This is the ABCD help desk e-mail address.	Identity Provider Organization Point of Contact E-mail Address Text	Fixed text	techsupport@abcd.gov

Table 2: Example of Local Attribute Mapping Form

In the above spreadsheet, you must have a row for every GFIPM Metadata attribute that your IDP asserts, explaining the source of the value and how you plan to map from the source to the GFIPM attribute.

If you would like additional examples of a Local Attribute Mapping Form, please contact gfipm-support@lists.gatech.edu to request them.

At this point, you should have completed the GFIPM Information Sharing Plan and the Local Attribute Mapping Form.

2.2 Submitting a Request for Federation Membership

This section serves as a supplemental aid to the membership application process by listing the membership documents that should be collected or generated during the IDP implementation process. The authoritative document for the process of submitting a membership request in NIEF is the Operational Policies and Procedures Document [GFIPM OPP].

The membership process is defined in [GFIPM OPP]. The process follows the following four phases:

1. Request-to-join process
2. Application process
3. On-boarding process
4. Ongoing membership

While going through the IDP implementation process, you should either collect or produce the following membership documents:

1. ***Authority-to-Operate Document—***A document attesting to the organization's authority to operate as an identity provider and provide access to the federation for the organization's users. It typically takes the form of a signed memorandum or letter from the organization's executive officer to the federation manager.

2. ***Local Security Policy Document—***A document describing the security policy that is currently in place within your organization. This document should already exist within your organization. It should be collected during Section 2.1.2.

3. ***Local User Agreement Document—***A document describing the terms and conditions to which your users must agree as a prerequisite for using a digital identity issued by your organization.

This document should already exist within your organization. It should be collected during Section 2.1.2.

4. ***Local User Vetting Policies and Procedures Document—***A document describing the user vetting policies and procedures that are currently in place within your organization. This document should already exist within your organization. It should be collected during Section 2.1.2.

5. ***Local Attribute Mapping Form for IDP—***A document describing how the organization plans to map its local policies and locally stored user attributes into attributes conforming to the GFIPM Metadata standard [GFIPM Metadata]. You develop this document by following the instructions in Section 2.1.4.

6. ***Implementation Documentation Form for IDP—***A document describing how your local federation-aware infrastructure is implemented. You develop this document by following the instructions in Section 2.8.

Other documents are required for the membership application process, but these are outside the scope of this document. For more complete documentation about the membership application and technical onboarding process for a GFIPM federation, please see the GFIPM Operational Policies and Procedures Guideline [GFIPM OPP].

2.3 Choosing an IDP Product

This section lists the requirements for products that may be considered for a GFIPM IDP. It also briefly describes the IDP products for which GFIPM implementers currently have some amount of knowledge and implementation experience.

As you work through the process of choosing an IDP product, consider which product best meets your organization's needs, and keep in mind that the best product for you may not necessarily be included in this document. For those organizations that have an existing enterprise identity management platform, the best choice may be to implement a GFIPM IDP via that existing platform – especially if the existing identity management platform conforms to the GFIPM IDP technical requirements (listed below).

An Identity Provider (IDP) is responsible for authenticating an end user and asserting a SAML assertion for that user in a trusted fashion to Service Providers. When a user attempts to access a Service Provider, the user's IDP collects local attribute information about the user and uses it to generate a SAML assertion for the user.

A GFIPM IDP must meet the following minimum requirements:

1. It must conform to the SAML 2.0 Web Single Sign-On (SSO) Profile [SAML2].

2. It must support SP-initiated Web Browser SSO.

3. It must be compliant with the IDP requirements in [GFIPM U2S Profile].

Typically, an IDP consists of several components, including:

- User authentication
- Local user repository
- SAML assertion generation

An IDP product may address one or more of these components, but in any case, it must perform the SAML assertion generation. It is likely that your organization already supports several of these components, including user authentication and a local user repository. Any IDP product must support interfaces to these existing systems.

While an IDP generates a SAML assertion that provides attributes about a user, an SP handles access to protected resources based on information given to it by an IDP. To perform their respective roles, an IDP and an SP need to communicate with each other, and the protocol through which this communication occurs in GFIPM is the Security Assertion Markup Language [SAML2].

SAML is a product of the OASIS Security Services Technical Committee (SSTC). It is an XML-based framework for communicating user authentication, entitlement, and attribute information. SAML allows business entities to make assertions regarding the identity, attributes, and entitlements of a subject (an entity that is often a human user, but may also be an application or system) to other entities, such as a partner company or another enterprise application.

Any IDP product chosen for a GFPIM federation must be SAML 2.0 compatible. The product must also have support for looking up GFIPM Metadata attributes in a local data source, so they can be assembled into a SAML assertion.[7]

The following is a non-exhaustive list of products that provide SAML-based identity provider capabilities. You should evaluate these and other products to determine which best meet your needs within your budget.

[7] This capability is a standard feature in most SAML 2.0 IDP products.

2.3.1 Shibboleth IDP

Shibboleth is a standards-based, freely available open-source software package for Web single sign-on across or within organizational boundaries. It allows sites to make informed authorization decisions for individual access of protected online resources in a privacy-preserving manner. It was developed by the Internet 2 project using the OpenSAML open-source implementation of SAML 2.0. It is being used by at least one participant in GFIPM.

The GFIPM federation extends the authorization functions to include privilege management for the justice community and partner organizations with a standards-based approach for implementing federated identity. Note that Shibboleth has separate components that act as an IDP and an SP.

Several components of the GFIPM Reference Federation[8] are implemented using Shibboleth, including the Reference IDP, the Reference SP, and the production CISA IDP and SP.

Existing federation members and technical support staff have extensive implementation experience with Shibboleth. The GFIPM reference IDP and SP use Shibboleth.

The product Web page is at https://spaces.internet2.edu/display/SHIB2.

See Section 6.4 for Shibboleth implementation and installation instructions.

2.3.2 Ping Identity PingFederate IDP

PingFederate is a commercial product marketed by Ping Identity (www.pingidentity.com). It supports Internet Single Sign-On, Internet User Account Management, and Identity-Enabled Web Services. PingFederate is SAML 2.0-compatible. Note that the PingFederate Internet Identity Security Platform can act as both an IDP and an SP.

2.3.3 CA Federation Manager IDP

Federation Manager is a commercial product marketed by CA[9] (formerly Computer Associates, www.ca.com). It provides standards-based identity federation capabilities that enable the users of one organization to easily and securely access the data and applications of another. It delivers support of federation standards, including SAML, enabling both Identity Providers and Service Providers. CA Federation Manager also provides for the administration of federation partnerships.

[8] The GFIPM Reference Federation is a GFIPM-conformant testing environment maintained by GTRI.
[9] CA, formerly known as Computer Associates, changed its name in 2006.

Note that CA Federation Manager can act as both an IDP and an SP. Its product Web page is at http://www.ca.com/us/products/Product.aspx?ID=8231.

2.3.4 Sun OpenSSO IDP

The Sun Open Web SSO project (OpenSSO, https://opensso.dev.java.net/) provides core identity services to simplify the implementation of transparent single sign-on (SSO) as a security component in a network infrastructure. OpenSSO provides the foundation for integrating diverse Web applications that might typically operate against a disparate set of identity repositories and is hosted on a variety of platforms such as Web and application servers. This project is based on the code base of Sun Java System Access Manager (http://www.oracle.com/us/products/middleware/identity-management/oracle-access-mgmt/index.html), a core identity infrastructure product offered by Oracle (formerly Sun Microsystems, www.oracle.com). Sun OpenSSO is available as an open-source solution or as a commercialized packaged product from Oracle with support. Note that Sun OpenSSO can act as both an IDP and an SP.

2.3.5 Oracle Identity Federation IDP

Identity Federation is a commercial product marketed by Oracle (www.oracle.com). It is part of the Oracle Identity & Access Management Suite. Note that the Oracle Identity Federation product can act as both an IDP and an SP. Its Web page is at http://www.oracle.com/products/middleware/identity-management/identity-federation.html.

2.4 Implementing a GFIPM IDP

This section provides an overview of the logical structure of a GFIPM IDP and also guides the reader through the process of implementing a GFIPM IDP at a high level. Note that it is possible for implementers to become confused about the difference between a GFIPM IDP and an existing local identity management system. A GFIPM IDP is a secure service that produces a SAML assertion for a user. The user information in the SAML assertion should be based on your organization's local identity management system, which may be a directory, a database, or another application to manage your user identities. GFIPM does not dictate how your user identities are managed, nor does it require modification to your identity management system beyond establishing a means to interface your GFIPM IDP to it.

After choosing an appropriate IDP product, the next step is implementing a GFIPM IDP using it. While this document cannot lead an implementer through all the details of installing a product and interfacing it to your organization's identity management system, it does outline broad steps and offer guidelines on how to overcome specific

IDP implementation issues that other GFIPM participant organizations have experienced.

2.4.1 IDP Components

An IDP consists of several logical components. Using a Shibboleth IDP as an example, the IDP components are illustrated in Figure 1 and discussed below. Other IDP products vary in their implementation details; however, all the components described here need to be present in some form for any IDP product.

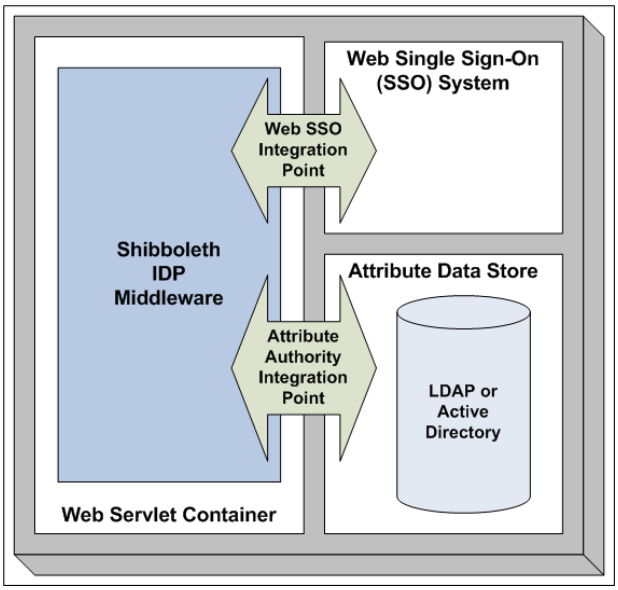

Figure 1: GFIPM Shibboleth Identity Provider Structure

2.4.2 IDP Core Software Module

The IDP software module, depicted by the blue box ("Shibboleth IDP Middleware") in Figure 1, consists of a set of interfaces, called *integration points*, which must be connected to other system components for the IDP to work. In the case of Shibboleth, it is implemented as a Java servlet and runs within a Web servlet container.

The IDP core software module handles the processing of incoming SAML messages from SPs, as well as the creation of outgoing SAML messages to SPs. In addition, it manages signing and encryption of all outgoing SAML messages, as well as signature verification and decryption of all incoming SAML messages. In the case of Shibboleth, the Web servlet container in which the software runs is responsible for handling the IDP's connection-level (TLS) encryption needs. Interested readers may refer to [GFIPM U2S Profile] for normative details about the SAML protocols and messages used within the GFIPM federation; however, note that detailed knowledge

of these protocols is not typically necessary for implementation of an IDP using COTS or open-source software.

2.4.3 Web Servlet Container

A Web servlet container is often required on a GFIPM IDP to run the IDP core software module. (It is required in the case of Shibboleth.) Many such Web servlet containers are available for use, but GFIPM participants have chosen to use the freely available Tomcat [Tomcat] open-source servlet container. Tomcat was chosen for several reasons. First, Tomcat is fully compatible with the Shibboleth IDP middleware; in other words, it runs the Shibboleth code without any problems. Second, Tomcat supports connection-level encryption using both SSL and TLS.[10] Third, Tomcat supports client certificate authentication of browsers with support for certificate revocation lists (CRLs). Fourth, Tomcat runs on both of the major operating system platforms (Microsoft Windows and Red Hat Enterprise Linux) that GFIPM participants use.

2.4.4 IDP Integration Points

Implementing an IDP in a federation requires that two integration issues be addressed. The first of these issues involves the integration of the IDP with the local site's user authentication system (the single sign-on integration point), and the other involves connecting the IDP to the local site's attribute repository (the attribute authority integration point).

At the single sign-on integration point, the user authentication system can be a username/password authentication system (though this form of authentication does not provide enough assurance for the use of most federation resources), a token-based authentication system, a PKI-based client certificate authentication system, or another two-factor authentication system. By having chosen Tomcat as the Web servlet container, we automatically gained support for client certificate authentication of browsers with support for certificate revocation lists (CRLs).

At the attribute authority integration point, the IDP should be connected to the existing attribute data store, which is your local identity management system (such as a LDAP repository).

2.4.5 Web Single Sign-On System

The Web single sign-on (SSO) system integrates with the IDP core software module via an SSO integration point and provides the basis for the IDP core software module to generate SAML authentication statements about users. To realize the

[10] The GFIPM User-to-System Profile [GFIPM U2S Profile] mandates the use of TLS and prohibits the use of SSL.

single sign-on benefit of federated identity management, an SSO system must already exist and must already be used to authenticate your users for other purposes.

2.4.6 Attribute Data Store

The attribute data store integrates with the IDP core software module via an attribute authority integration point. It provides a source of trusted data about users that can be used to construct SAML assertions. Any component that acts as an attribute data store is essentially a database. There are virtually no limitations on the attribute data store in terms of how it stores attributes; however, in the case of Shibboleth, it must store attributes in a fashion that allows for attribute queries based on a user ID or some other key that can be understood by the Shibboleth IDP and maps uniquely to a specific user. Typically, an organization will want to connect an IDP to an existing user attribute repository—often an LDAP repository or an Active Directory database. It is also possible to use an ODBC or SQL database, a flat file on the local machine's file system, or any other repository, via custom Java code. But these cases are rare. The most common implementation scenario involves connecting an IDP directly to a local LDAP or Active Directory server.

2.5 Writing an IDP Test Plan

This section describes how to write a plan for testing a new IDP in a test environment such as the GFIPM Reference Federation.

The new Identity Provider should be first deployed and tested in a test environment, followed by deployment and testing in the operational federation.

If you are also developing a new Service Provider, you can deploy and test it at the same time that you deploy and test your Identity Provider. Note that testing for interoperability between your own IDP and SP is insufficient in terms of full interoperability testing. Comprehensive interoperability testing requires that you test for interoperability between your IDP and other organizations' SPs, as well as between your SP and other organizations' IDPs.

There are a wide variety of possible test environments, each having different IDP implementations and security requirements. It is therefore not possible for this guide to provide you with a "complete" test plan for your IDP. But this section outlines the basic testing steps that you should perform, regardless of your test environment. Based on these steps, you should be able to generate a test plan that is sufficient and specific to your organization. Your test plan should address the following issues:

1. ***Test Goals—***Your plan must address the following questions:

   a. Will you test individual modules (IDP middleware, SSO system, attribute data store, assertion generator, etc.)?

   b. Will you perform integration testing? If so, how will you test the interfaces between the IDP's components (single sign-on integration point, attribute authority integration point, etc.)? For example, if you vary the data on one side of the interface, does the information on the other side reflect the change correctly?

   c. How will you perform interoperability testing of your IDP? Keep in mind that the ultimate purpose of GFIPM interoperability testing is to determine whether your IDP's user attributes are compatible with the SPs that consume and use them.[11] To ensure broad test coverage of user attributes, you should create multiple test identities with varying attributes (user data, permissions, privileges, etc.) to ensure that user attributes for your local identities are correctly translated into the appropriate GFIPM user attributes. Also, be sure to notify your testing partners (the federation manager, SP administrators, etc.) of your testing goals and plans.

   d. What are your organization's formal processes for acceptance testing? Be sure to notify your testing partners if you need results or other inputs from them.

2. ***Testing Schedule—***Your testing schedule can be as simple as an e-mail asking/notifying your testing partners of the days you will need to communicate with them. Note that before testing begins, all required resources must be available.

3. ***Resource Logistics—***You must identify all required resources, including all of the personnel, test machines, and network connections that will be required for testing.

   a. Test machines should be 100 percent available on test days.

   b. Identify and arrange for the required network connections (especially firewall and router configurations) well in advance of the test date. Connections (i.e., all specific port numbers)

[11] The GFIPM Reference Federation is a useful test bed for IDP integration testing. It includes both mock-up SPs and fully operational (test) SPs. See Appendix A for a full description.

should also be tested. If necessary connections are denied, your local network administrator may need to open firewalls, which can take several weeks in some organizations.

c. Notify all personnel who are required for testing well in advance of the testing date. This may include persons who are not part of your local organization.

4. ***Documented Testing Procedures—***Generate a list of test cases, each of which must pass. Test cases should include SAML assertions that include the following:

   a. Only the minimum attributes.

   b. An excessive number of attributes.

   c. A reasonable number of attributes.

   It is strongly recommended that you create a user identity in your identity management system that is not a valid identity (such as missing last name, missing e-mail, or malformed phone number) and perform tests to ensure that your IDP behaves properly when dealing with it.

5. ***Responsibility for Carrying Out the Test Plan—***Which person in your organization should execute the test plan? Experience has shown that a test plan works best when the person performing the testing is not the same person who implemented the software.

2.6 Deploying an IDP in a Test Environment

This section presents steps required to deploy an IDP in a test environment to ensure its connectivity and interoperability with SPs within the context of the test environment's GFIPM Trust Fabric.[12]

Any new IDP must be "connected" to the test environment by adding the IDP to the test environment's trust fabric file. The trust fabric file update process consists of these steps:

[12] The term "GFIPM Trust Fabric" refers to a cryptographically signed document containing SAML service endpoint information and X.509 certificate data for each endpoint (IDP, SP, etc.) in a GFIPM federation. In practice, "GFIPM Trust Fabric" or "trust fabric" refers to an XML file. It is generally signed by the federation manager and distributed to all federation member organizations as a cryptographic trust anchor upon which all cryptographic operations (signing and encryption) rely.

1. Provide your IDP's entity metadata to your federation manager (or the organization that manages the test environment you will use).

2. The federation manager adds the new entity to the trust fabric file.

3. All participants load the new trust fabric file into their IDPs and SPs.

You can test whether your IDP's entry into the test environment's trust fabric has succeeded by verifying its SAML interoperability with other members of the test environment, as described in Section 6.6.

The GFIPM Reference Federation is an excellent resource that is generally available for use as a test environment for GFIPM IDPs. It contains a large number of useful and necessary resources for implementing and testing your IDP. These topics are covered more thoroughly in Appendix A.

2.7 Executing an IDP Test Plan

After you have written your test plan and implemented your Identity Provider, you must execute the test plan.

While software issues almost always arise during testing, experience has shown that personnel scheduling problems and network configuration problems tend to cause the most delays during testing. The bottom line is that you must make sure all required personnel (especially required testing partners from other organizations) will be available as expected during the testing process, and that all machines used in the testing process have Internet connectivity on all required ports as needed.

Work through your test plan. Take thorough notes, and compile a test report to be distributed within your organization and to all testing partners.

2.8 Deploying an IDP in an Operational Federation

This section presents steps required to deploy an IDP in the operational federation to ensure its connection and interoperability to the GFIPM Trust Fabric.

During your deployment in the test environment, you were able to use all the test environment's resources. If you are now deploying your IDP in the NIEF operational federation, here are the equivalent production resources that you can leverage:

- Publicly accessible information about NIEF is available at http://nief.gfipm.net/trust-fabric/.

- The NIEF Trust Fabric document is located at **Error! Hyperlink reference not valid.**.

Note that there are no "test" IDPs or SPs in an operational federation such as NIEF. ***The operational federation contains live data, and test identities should never be used within it.***

Any new IDP must be "connected" to NIEF (or your own operational GFIPM federation) by adding the IDP to the federation's trust fabric. The trust fabric update process consists of these steps:

1. Provide your IDP's entity metadata to the federation manager.

2. The federation manager adds the new entity to the federation trust fabric.

3. All participants load the new federation trust fabric into their IDPs and SPs.

Before or during your deployment, you must also fill out an Implementation Documentation Form for your IDP and submit it to the federation manager as part of the membership application process described in Section 2.2. A template of this form is available from the federation manager. The form requests the following information:

- IDP software platform details (OS, Web server, SAML software)
- User authentication endpoint details
- User attribute endpoint details
- Network configuration notes

Your ability to test your IDP in the operational federation will be limited because of the lack of test SPs in the operational federation. According to the usage policies of most SPs in an operational federation, only real users using real identities (with valid user data, permissions, and privileges) are permitted to use (i.e., test) the operational systems. Therefore, when executing your IDP's test plan in the operational federation, you must perform the necessary tests with real users (from your organization and others). As before, write a test report to be distributed within your organization and to all testing partners.

3. Implementing a Service Provider

This section outlines the steps necessary to implement a GFIPM Service Provider (SP). The following steps are presented:

1. Developing a GFIPM information sharing plan
2. Submitting a request for federation membership
3. Choosing a product for building your SP
4. Implementing your organization's SP
5. Writing a test plan for your SP
6. Deploying your SP in a test environment
7. Executing the test plan for your SP
8. Deploying your SP in an operational federation

3.1 Developing a GFIPM Information Sharing Plan for an SP

During this process, you will accomplish the following:

1. Develop a list of your local resources to which you want to allow access for other organizations' users.
2. Determine your organization's business rules for granting access to your local resources.
3. Using your business rules, develop and codify the access control policies for your resources in terms of the GFIPM user metadata.
4. Fill out a Local Access Policy Mapping Form for your SP.

3.1.1 Identify Local Resources

This section will help you identify local resources that may be useful to other federation members and to determine whether they can be made available to other organizations' users under certain access rules.

Service Providers need to determine which production resources they should make available to other federation members. An equally important consideration is which of these resources can be permitted to be made available and how long the decision making process will take. Federation members typically "own" some resources; in other words, the federation member administers the resource or otherwise has control over access decisions. For these resources, the decision to make the resource available to federation users may be easy. But some federation members may be

distributing data that belongs to another organization. In this case, receiving permission to release the data may be more complex and time consuming. The required timelines for gaining permission to make resources available within a federation need to be built into an SP's implementation schedule.

Resources that might be useful to other law enforcement organizations include the following:

- Arrest records
- Incident reports
- Criminal history reports
- Criminal justice reports
- Criminal investigative reports
- Criminal intelligence reports
- Counterterrorism notifications
- Driver's license records
- Public safety messages
- Other federal/state/tribal/local law enforcement information

Types of resources that would be less useful but still of potential interest to federation users include the following:

- Amber Alerts (i.e., missing child alerts)
- Sex offender information
- Other publicly available information

As part of the process of determining which local resources are to be made available, you must also collect the appropriate authority-to-operate documents for providing those resources to federation users. These documents will be needed in Section 3.2.

3.1.2 Determine Business Rules for Resources

For each resource you identified in the previous section, you must now collect or document its access policy.

It is likely that all of the identified resources already have documented access policies. For example, a resource's access policy might be as simple as the following:

- Any user who is a sworn law enforcement officer may access this resource.

Or a resource's access policy might be more complicated, such as this:

- The user must be a sworn law enforcement officer.

- The user must also have NCIC certification and criminal history privileges.

Many resources with their access policies are listed on the NIEF federation Web site at http://nief.gfipm.net/ for NIEF participants.

In the unlikely event that your organization does not already have documented access control policies, you will need to work through the process of writing them and getting the appropriate approval(s) to provide those resources to federation members.

3.1.3 Develop Access Control Rules

This section uses the business rules from the previous section to guide you through the process of codifying the rules in terms of the GFIPM user metadata.

Note that this section uses the NIEF federation as an example for developing access control rules. Specifically, this section relies on the required and recommended attributes in a SAML assertion. If you are building an SP for a different federation, your set of metadata attributes may be different. However, it is likely that there will be many similarities to the NIEF requirements.

The access control rules are written in terms of attributes in the GFIPM user metadata. The minimal access requirements that a user identity must contain are the following fields:

- User's Last Name
- User's First Name
- User's Phone Number
- User's E-mail Address
- User's Federation ID
- User's Home Organization Name
- User's Identity Provider Name

The above fields are typically used for auditing purposes by a Service Provider to meet the "Federation Login" requirement. Because these attributes do not assert any permissions or privileges, a user identity that contains only the above attributes typically will not be granted access to any law enforcement resources.

3.1.4 Fill Out Local Access Policy Mapping Form

This section will help you fill out a Local Access Policy Mapping Form to translate your plain-English access policies into Boolean logic rules based on GFIPM metadata

attributes. The Local Access Policy Mapping Form for your SP is later used as part of your request for federation membership.

The Local Access Policy Mapping Form is a document describing how the organization plans to map its local access control policies into rules that can be expressed using attributes from the GFIPM Metadata standard.

When the federation manager reviews your application package, it will provide a copy of your Local Access Policy Mapping Form (for an SP) to all existing members for review and comment.

The Local Access Policy Mapping Form should be written as a spreadsheet (i.e., Microsoft Excel). A template of this form is included with the membership application forms provided by the federation manager when you request to join the federation.

Before editing the file, you should rename it to include your SP name in the file name.

Below is an example of the design of the spreadsheet, including the headers followed by several rows describing access policy mappings. Note that these examples are from different members, so their derivations are not related to each other.

GFIPM Access Control Policy Map Service Provider Name: <Your Organization>				
	Policy for Resource Discovery		**Policy for Resource Access**	
Service/ Resource Name	**Semantic Intent**	**GFIPM Boolean Logic**	**Semantic Intent**	**GFIPM Boolean Logic**
Arizona Counter-Terrorism Information Center	Any user with a valid federation login may discover this resource.	ALLOW if: (Given Name is present AND Surname is present AND Telephone Number is present AND Federation ID is present AND Employer Organization Name is present AND Identity Provider Name is present)	Any user with a valid federation login may access this resource. In addition, sufficient audit data is required for all users.	ALLOW if: (Given Name is present AND Surname is present AND Telephone Number is present AND Federation ID is present AND Employer Organization Name is present AND Identity Provider Name is present)
New Mexico Complete Arrest Information	Any user with a valid federation login may discover this resource.	ALLOW if: (Given Name is present AND Surname is present AND Telephone Number is present AND Federation ID is present AND Employer Organization Name is present AND Identity Provider Name is present)	To access this resource, a user must be a sworn law enforcement officer with NCIC criminal history certification and criminal history home data search privileges. In addition, sufficient audit data is required for all users.	ALLOW if: (Given Name is present AND Surname is present AND Telephone Number is present AND Federation ID is present AND Employer Organization Name is present AND Identity Provider Name is present) AND (Sworn Law Enforcement Officer Indicator = TRUE) AND (NCIC Certification Indicator is present AND (Criminal History Home Search Data Privilege Indicator = TRUE)

In the above table, you must have a row for every local resource that your SP will make available and then explain the policy for resource discovery and the policy for resource access and the corresponding access control rule expressed in GFIPM Metadata attributes.

If you would like additional examples of a Local Access Policy Mapping Form, please contact the federation manager at gfipm-support@lists.gatech.edu.

At this point, you should have completed the Local Access Policy Mapping Form.

3.2 Submitting a Request for Federation Membership

This section serves as a supplemental aid to the membership application process by listing the membership documents that must be collected or developed during the SP implementation process. The authoritative document for this process in NIEF is the Operational Policies and Procedures document [GFIPM OPP]. This section of this document is not a substitute for the OPP document.

The membership process follows these four phases:

1. Request-to-join process
2. Application process
3. On-boarding process
4. Ongoing membership

While working through the SP implementation process, you should collect or develop the following membership documents required above.

1. ***Authority-to-Operate (ATO) Document(s)—***A document attesting to the organization's authority to operate as a service provider and make available electronic resources belonging to, or under the legal control of, a specific legal jurisdiction. An ATO document typically takes the form of a signed memorandum or letter from the organization's executive officer to the federation manager.

2. ***Local Security Policy Document—***A document describing the security policy that is currently in place within your organization. This document should already exist within your organization.

3. ***Local User Agreement Document—***A document describing the terms and conditions to which your users must agree as a prerequisite for using a digital identity issued by your organization. This document should already exist within your organization.

4. ***Local Privacy Policy Document—***A document describing the policies that govern the practices of maintaining the privacy of users visiting the organization's service provider or portal. This document should already exist within your organization.

5. ***Local Access Policy Mapping Form for Your SP—***A document describing how the organization plans to map its local access control policies into rules that can be expressed using attributes from the GFIPM Metadata Specification [GFIPM

Metadata]. You develop this document by following the instructions in Section 3.1.4.

6. ***Implementation Documentation Form for SP—***A document describing how your local federation-aware infrastructure is implemented. You develop this document by following the instructions in Section 3.8.

Other documents are required for the membership application process, but these are outside the scope of this implementation guide.

3.3 Choosing an SP Product

This section outlines the requirements for products that may be considered for a GFIPM Service Provider (SP). We also present the list of products of which we have knowledge.

An SP is responsible for managing access to applications, services, and other resources used by federation users. To do this, it relies on Identity Providers (IDP) to assert information about users, leaving the SP to manage access control and dissemination based on the trusted set of attributes it receives for each user. There can be an arbitrary number of SPs in a federation, and each SP can manage an arbitrary number of resources.

An SP handles the management of access to protected resources based on information given to it by an IDP. To perform their respective roles, an IDP and an SP need to communicate with each other, and the standard through which this communication occurs in GFIPM is the Security Assertion Markup Language [SAML2].

SAML was developed by the Organization for the Advancement of Structured Information Standards (OASIS) Security Services Technical Committee (SSTC). It is an XML-based framework for communicating user authentication, entitlement, and attribute information. As its name suggests, SAML allows business entities to make assertions regarding the identity, attributes, and entitlements of a subject (an entity that is often a human user, but may be an application or system) to other entities, such as a partner company or another enterprise application.

A GFIPM SP must meet the following minimum requirements:

1. It must conform to the SAML 2.0 Web Single Sign-On (SSO) Profile [SAML2], including support for both SP-initiated SSO and IDP-initiated SSO.

2. It must be able to discover the user's IDP, by either:

    a. Supporting the OASIS Identity Provider Discovery Service Protocol and Profile.

    b. Implementing a local IDP Discovery Service at the SP.

See [GFIPM U2S Profile] for a thorough, normative specification of the technical requirements that a GFIPM SP must meet.[13]

Any SP product chosen for a GFPIM federation must be SAML 2.0-compatible. It must implement the Web Single Sign-On (SSO) Profile and support SP-initiated Web SSO. Further, it must support parsing and processing of SAML attributes containing GFIPM user metadata.

The following is a non-exhaustive list of products that provide SAML-based Service Provider capabilities. You should evaluate these and other products to determine which product best meets your needs and your budget.

3.3.1 Shibboleth SP

Shibboleth is a standards-based, open-source software package for Web single sign-on across or within organizational boundaries. It allows sites to make informed authorization decisions for individual access of protected online resources in a privacy-preserving manner. It was developed by the Internet 2 project using the OpenSAML open-source implementation of SAML.

The GFIPM federation extends the authorization functions to include privilege management for the justice community and partner organizations with a standards-based approach for implementing federated identity. Note that Shibboleth has separate components that act as an IDP and an SP.

Several components of the GFIPM Reference Federation are implemented using Shibboleth, including the Reference IDP, the Reference SP, and the production CISA IDP and SP.

Existing federation members and technical support staff have extensive implementation experience with Shibboleth. The GFIPM reference IDP and SP use

[13] In addition to the SAML 2.0 Web SSO Profile, [GFIPM U2S Profile] also specifies optional use of the SAML 2.0 Single Log-Out (SLO) Profile. In practice, however, support for SAML SLO is limited, and in addition, there are several known problems with the profile at the level of technical specification that render it unable to accomplish a true user logout operation in standards-compliant Web browsers. Therefore, it is not used in practice.

Shibboleth. The Shibboleth product Web page is available online at https://spaces.internet2.edu/display/SHIB2.

See Section 6.4 for Shibboleth implementation and installation instructions.

3.3.2 Ping Identify PingFederate SP

PingFederate is a commercial product marketed by Ping Identity (www.pingidentity.com). It supports Internet Single Sign-On, Internet User Account Management and Identity-Enabled Web Services. Note that the PingFederate Internet Identity Security Platform can act as both an IDP and an SP. See Section 6.5 for PingFederate implementation issues.

3.3.3 CA Federation Manager SP

Federation Manager is a commercial product marketed by CA (formerly Computer Associates, www.ca.com). It provides standards-based identity federation capabilities that enable the users of one organization to easily and securely access the data and applications of another. It delivers support of federation standards, including SAML, enabling both Identity Providers and Service Providers. CA Federation Manager also provides for the administration of federation partnerships. Note that CA Federation Manager can act as both an IDP and an SP. Its product Web page is at http://www.ca.com/us/products/Product.aspx?ID=8231.

3.3.4 Sun OpenSSO SP

The Sun Open Web SSO project (OpenSSO, https://opensso.dev.java.net/) provides core identity services to simplify the implementation of transparent single sign-on (SSO) as a security component in a network infrastructure. OpenSSO provides the foundation for integrating diverse Web applications that might typically operate against a disparate set of identity repositories and is hosted on a variety of platforms such as Web and application servers. This project is based on the code base of Sun Java System Access Manager (http://www.oracle.com/us/products/middleware/identity-management/oracle-access-mgmt/index.html), a core identity infrastructure product offered by Oracle (formerly Sun Microsystems, www.oracle.com). Sun OpenSSO is available as an open-source solution or as a commercialized packaged product from Oracle with support. Note that Sun OpenSSO can act as both an IDP and an SP.

3.3.5 Oracle Identity Federation SP

Identity Federation is a commercial product marketed by Oracle (www.oracle.com). It is part of the Oracle Identity & Access Management Suite. Note that Oracle Identity Federation can act as both an IDP and an SP. Its product Web page is at

http://www.oracle.com/products/middleware/identity-management/identity-federation.html.

3.4 Implementing a GFIPM SP

This section provides an overview of the logical structure of a GFIPM SP and also guides the reader through the process of implementing a GFIPM SP at a high level.

After choosing an appropriate SP product, the next step is to implement a GFIPM SP. This document cannot provide exhaustive details covering all possible implementation scenarios because of the inherent complexity of the topic. But the document does outline the broad steps you must take and offers guidelines on how to overcome common problems that other organizations have experienced.

3.4.1 SP Components

There are several basic components in a GFIPM SP that are noteworthy. Using a Shibboleth SP as an example, the SP components are illustrated in Figure 2 and discussed below. Other SP products may vary in their implementation details; however, the components described here will always be present in one form or another.

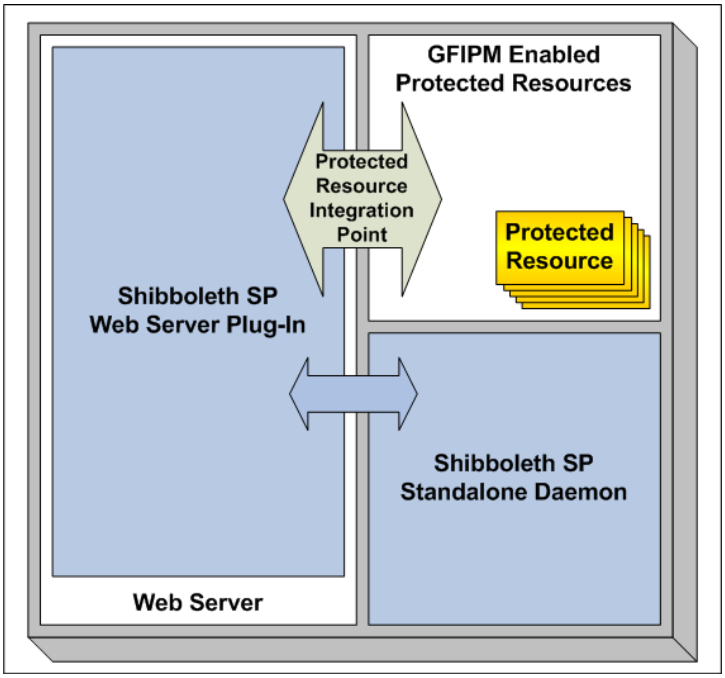

Figure 2: GFIPM Shibboleth Service Provider Structure

Note that the SP structure depicted in Figure 2 contains only the most basic components of a Shibboleth-based SP. In particular, the protected resource integration point in the diagram often contains custom code designed to provide an

interface between legacy resources (which are not natively GFIPM-enabled and do not support access control and auditing based on GFIPM Metadata) and the GFIPM-enabled federation components. Section 3.4.4 provides additional general-purpose information about the protected resource integration point within an SP. In addition, Section 6.3 provides guidance about how to define and implement a GFIPM-enablement strategy for a resource, based on the resource's technical characteristics.

3.4.2 Web Server

An SP enables basic federated identity management functionality for Web-based resources; therefore, it must contain a Web server. The Web server is responsible for serving sensitive Web-based resources to users who request them, subject to access controls and other usage policies that may exist. Resources may be served either directly or via a reverse-proxy-based architecture.

In addition to serving Web-based resources, the Web server is responsible for handling TLS encryption of HTTP traffic, including incoming SAML messages that are sent to the server from IDPs via the user's Web browser over HTTPS. The Web server component is typically integrated with SP core software module (see Section 3.4.3) and relies on the SP core software module to handle basic SAML operations.

GFIPM participants using Shibboleth have successfully used two different Web servers at their SPs: Microsoft Internet Information Services (IIS) on Windows and Apache on Linux machines. At this time, no other Web servers or host operating systems have been used within GFIPM.

3.4.3 SP Core Software Module

The SP core software module integrates with the Web server and handles basic SAML-level operations within the SP. It is depicted in Figure 2 by a pair of blue boxes ("Shibboleth SP Web Server Plug-In" and "Shibboleth SP Standalone Daemon"). In the case of Shibboleth, it is implemented in C++ and consists of several components (a Web server extension module and a standalone daemon) that work in tandem.

The SP core software module must be integrated with Web-based resources via an integration point that allows it to protect them. The core module handles the processing of incoming SAML messages from IDPs, as well as the creation of outgoing SAML messages to IDPs. In addition, it manages the signing and encryption of all outgoing SAML messages, as well as signature verification and decryption of all incoming SAML messages. Interested readers may refer to [GFIPM U2S Profile] for normative details about the SAML protocols and messages used within the GFIPM federation; however, note that detailed knowledge of these

protocols is not typically necessary for implementation of an SP using COTS or open-source software.

3.4.4 Protected Resource Integration Point

The protected resource integration point enables the SP to protect sensitive Web-based resources that are to be shared within the GFIPM federation. There are two ways in which the SP can be integrated with sensitive resources. Each method is discussed below.

- ***Integration Method Number 1—***The first integration method involves the use of access control logic that is built into the SP software itself. When using this integration method, it is possible to construct simple access control policies that permit or deny access to specific URLs served by the Web server based on logic involving the SAML attribute values presented for a user. At the implementation level, this integration method is as simple as configuring the appropriate access control logic in the SP core software module via its UI or a configuration file. This integration method works only when the following two conditions apply: First, the granularity of access control for a resource protected in this manner is limited to static URLs. In other words, it is not possible to make access control decisions that depend on HTTP query variables (such as CGI parameters) using this technique. Second, the SAML attribute values used to make access control decisions in this manner must be relatively simple (e.g., limited to simple values such as a name or an e-mail address).

- ***Integration Method Number 2—***The second integration method involves passing SAML attribute data directly from the SP core software module to a protected resource and allowing the resource itself to make its own access control decisions based on the data. SAML attribute data is passed to an application via the following mechanism. In the case of Shibboleth, the SP core software module inserts each SAML attribute value into an HTTP header that is passed to a resource at the same time that the resource receives notification of an HTTP/HTTPS request from the Web server. (At the implementation level, the resource typically receives all of its HTTP/HTTPS headers—including the headers constructed by the Shibboleth SP software—via environment values in its local environment.) After receiving the SAML attribute values, the resource can examine and process them as needed and make decisions accordingly.

SPs in NIEF typically use a combination of both integration methods, as follows. First, basic access to each SP is typically controlled using the SP's native access control facilities, with the following simple policy: "A user may access resources on this SP if and only if he has authenticated with an IDP in the federation." Then, the resources themselves (or an appropriate resource proxy—see Section 3.4.5) implement a more fine-grained access control policy using SAML attribute values passed from the SP core software module. As previously mentioned, there are important considerations that influence the decision to serve a sensitive resource in the GFIPM federation directly or via a reverse proxy arrangement.

3.4.5 Optional GFIPM-Enabled Proxy/Portal Service

A protected resource can interact with the SP core software module within a Web server through its resource integration point. It is through this integration point that a resource can receive GFIPM user attributes for processing. Of course, GFIPM user attributes are not useful to a resource unless the resource has the ability to understand them, process them, and use them as needed for purposes such as user identification, access control, and auditing. A resource that has this ability is typically called a *federation-enabled resource* or *GFIPM-enabled resource*.

Most of the resources that are currently useful to federation users tend to be legacy Web applications that contain dynamically generated Web pages, rather than simple static pages. And for most of these applications, it is not possible to modify the application's source code because of various business-level constraints.

Several of the participants in the GFIPM project act as umbrella organizations and manage networks containing resources that are owned by other organizations. For example, JNET manages access to data from the Pennsylvania Department of Corrections and Pennsylvania Department of Transportation; however, JNET does not own this data. This situation is likely to be the case for many other federation participants in the justice community.

If a resource is not already federation-enabled, and its source code cannot be modified for federation enablement, then the resource cannot exist natively within the GFIPM federation, so it must be made available to federation users via a proxy/portal service. Such a service exists to provide proxied access to a non-federation-enabled resource in a manner that allows the resource to exist in a GFIPM federation and be available to federation users without any modifications to the resource's source code or internal configuration. Additional details about federation enablement of resources are available in Section 6.3. The remainder of this section describes the basic internal structure of a proxy/portal service.

Figure 3 shows the GFIPM Shibboleth SP structure depicted in Figure 2 but also illustrates how a GFIPM-Enabled Proxy/Portal Service fits into the SP architecture.

Note that the proxy/portal service integrates directly with the Shibboleth SP middleware, providing a bridge between GFIPM-enabled service provider and non-GFIPM-enabled legacy resources.

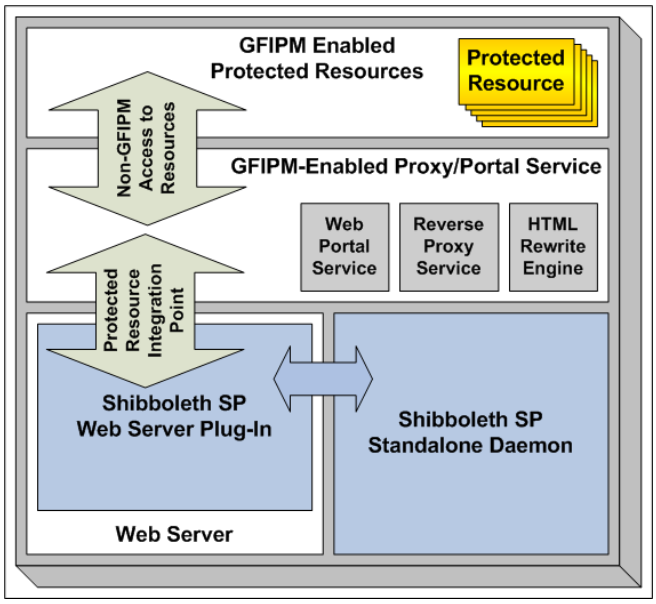

Figure 3: GFIPM Shibboleth Service Provider Structure With GFIPM-Enabled Proxy/Portal Service

A proxy/portal service typically includes the following logical components:

- ***Web Portal Service—***This is a federation-aware application that consumes GFIPM assertions and translates their content into access privileges for the proxied resources. Using these privileges, it defines network-level access control policies that are to be enforced by the reverse proxy service (see below). The Web Portal Service also presents a Web-based user interface for accessing the proxied resources. Finally, when necessary it may also handle the auditing of access attempts to the proxied resources.

- ***Reverse Proxy Service—***This is a network-level service that allows traffic to flow from a user's browser, through the SP's Web server, to protected resources. It enforces network-level access control based on instructions that it receives from the Web Portal Service. (The technical details concerning how access control policy is communicated from the Web Portal Service to the Reverse Proxy Service are implementation-dependent and outside the scope of this document.) Finally, when necessary, the Reverse Proxy Service handles authentication to protected resources in a manner that the resources understand. Please see Section 6.7 for

more detailed guidance about implementing a reverse proxy service.[14]

- ***HTML Rewrite Engine—***One of the challenging but necessary aspects of reverse-proxying Web content is rewriting the URLs and URL fragments in the HTML, JavaScript, and other content served by the proxied resource, so that to the user's browser the content appears to have come directly from the proxy. This process typically involves modifying URLs within the HTML and JavaScript content so that they refer to the reverse proxy server rather than referring directly to the proxied resource. Please see Section 6.8 for more detailed guidance about implementing an HTML rewrite engine.[15]

GTRI has developed a low-cost solution for Service Providers who need to deploy a GFIPM-Enabled Proxy/Portal Service capability. The solution leverages a relatively inexpensive COTS reverse proxy product and HTML rewrite engine, called EZproxy, and also contains custom code. CISA is currently using this GTRI-developed solution to serve non-federation-enabled resources to GFIPM users. JNET has also developed its own GFIPM-Enabled proxy/portal capability using entirely custom code, with which it is serving non-federation-aware resources.

In addition, GTRI plans to develop a freely available "Service Provider in a Box" solution for future participants who wish to implement the proxy-based Service Provider model described above. At the time of this writing, these plans are still tentative.

3.4.6 Protected Resources

A protected resource is the actual sensitive content that needs to be protected by the SP infrastructure. As implied by the discussion in the previous two sections, all protected resources fall into two categories: GFIPM-enabled resources, which natively understand GFIPM metadata, and non-GFIPM-enabled resources that must be proxied. Most of the protected resources in NIEF at this time are not natively GFIPM-enabled, and it is expected that nearly all resources that will be added to NIEF in the future will also be non-GFIPM-enabled.

Because the purpose of the Service Provider is to make resources available to federation users, GTRI has developed federation enablement techniques through which a wide range of legacy resources in the law enforcement domain can be made available to federation users. These techniques are not directly related to implementing an SP, so they are described in Section 6.3.

[14] Section 6.7 is currently a stub and does not yet contain any content.
[15] Section 6.8 is currently a stub and does not yet contain any content.

3.5 Writing an SP Test Plan

This section lists the necessary steps for a test plan for testing a new Service Provider in a test environment.

The new Service Provider should be deployed and tested in a test environment, followed by deployment and testing on an operational federation.

If you are also developing a new Identity Provider, you can deploy and test it at the same time that you deploy and test your Service Provider. Note that testing for interoperability between your own SP and IDP is insufficient in terms of full interoperability testing. Comprehensive interoperability testing requires that you test for interoperability between your SP and other organizations' IDPs, as well as between your IDP and other organizations' SPs.

Because there are a wide variety of possible test environments, with different Service Provider implementations and security requirements, it is not possible to give you a complete test plan for your SP. However, this section outlines the steps of the testing that should be performed, and you should be able to generate a test plan that is sufficient and specific to your organization.

Your test plan should address the following issues:

1. ***Test Goals—***Your plan must address the following questions:

    a. Will you test individual modules (SP middleware, various resources, proxy/portal service, HTML content rewriting service, etc.)?

    b. Will you perform integration testing? If so, how will you test the interfaces between the SP's components (protected resource integration point, etc.)? For example, if you vary the attributes in a SAML assertion on one side of the interface, does the other side of the interface behave correctly with respect to your local access control policy?

    c. How will you perform interoperability testing of your SP? Keep in mind that the ultimate purpose of GFIPM interoperability testing is to determine whether your SP's access control logic is compatible with the GFIPM assertions that it receives from IDPs.[16] To ensure broad test coverage of user attributes, you should use multiple test identities with varying attributes (user

[16] The GFIPM Reference Federation is a useful test bed for SP integration testing. It includes both mock-up IDPs and fully operational (test) IDPs. See Appendix A for a full description.

data, permissions, privileges, etc.) to ensure that your SP handles all conceivable cases properly. Also, be sure to notify your testing partners (the federation manager, IDP administrators, etc.) of your testing goals and plans.

d. What are your organization's formal processes for acceptance testing? Be sure to notify your testing partners if you need results or other inputs from them.

2. ***Testing Schedule—***Your testing schedule can be as simple as an e-mail asking/notifying your testing partners about the number of days you will need to communicate with them. Note that before testing begins, all required resources must be available.

3. ***Resource Logistics—***You must identify all required resources, including all of the personnel, test machines, and network connections that will be required for testing.

   a. Test machines should be 100 percent available on test days.

   b. Identify and arrange for the required network connections (especially firewall and router configurations) well in advance of the test date. Connections (i.e., all specific port numbers) should also be tested. If necessary connections are denied, your local network administrator may need to open firewalls, which can take several weeks in some organizations.

   c. Notify all personnel who are required for testing well in advance of the testing date. This may include persons who are not part of your local organization.

4. ***Documented Testing Procedures—***Generate a list of test cases, each of which must pass. Test cases should include SAML assertions that include the following:

   a. Only the minimum attributes

   b. An excessive number of attributes

   c. A reasonable number of attributes that allow access to some resources and not others. Carefully vary the values of attributes and check that the access permissions change as expected in accordance with the access control policy.

It is strongly recommended that you create a user identity in your identity management system that is not a valid identity (such as a missing last name, missing e-mail, or malformed phone number), and perform tests to ensure that your SP behaves properly.

5. ***Responsibility for Carrying Out the Test Plan—***Which person in your organization should execute the test plan? Experience has shown that a test plan works best when the person performing the testing is not the same person who implemented the software.

3.6 Deploying an SP in a Test Environment

This section presents steps required to deploy an SP in a test environment to ensure its connectivity and interoperability with GFIPM IDPs within the context of the GFIPM Trust Fabric.

Any new SP must be "connected" to a test environment by adding the SP to the test environment's trust fabric. The trust fabric update process consists of these steps:

1. Provide your SP's entity metadata to the federation manager.

2. The federation manager adds the new entity to the test environment trust fabric.

3. All participants in the test environment load the new test environment trust fabric into their IDPs and SPs.

You can test whether your SP's entry into the test environment's trust fabric has succeeded by verifying its SAML interoperability with other members of the test environment, as described in Section 6.6.

The GFIPM Reference Federation is an excellent resource that is generally available for use as a test environment for GFIPM SPs. It contains a large number of useful and necessary resources for implementing and testing your service provider. These topics are covered more thoroughly in Appendix A.

3.7 Executing an SP Test Plan

Once you have written your test plan and implemented your Service Provider, you must execute the test plan.

While software issues almost always arise during testing, experience has shown that personnel scheduling problems and network configuration problems tend to cause

the most delays during testing. The bottom line is that you must make sure all required personnel (especially required testing partners from other organizations) will be available as expected during the testing process, and that all machines used in the testing process have Internet connectivity on all required ports as needed.

Work through your test plan. Take thorough notes, and compile a test report to be distributed within your organization and to all testing partners.

3.8 Deploying an SP in an Operational Federation

This section presents steps required to deploy an SP on the operational federation to ensure its connection and interoperability to the GFIPM Trust Fabric.

During your deployment in the test environment, you were able to use all the test environment's resources. If you are now deploying your SP in the NIEF operational federation, here are the equivalent production resources that you can leverage:

- Publicly accessible information about NIEF is available at http://nief.gfipm.net/.
- The NIEF Trust Fabric document is located at http://nief.gfipm.net/trust-fabric/nief-trust-fabric.xml.

Note that there are no test IDPs or SPs in an operational federation such as NIEF. ***The operational federation contains live data, and test identities should never be used within it.***

Any new SP must be "connected" to the NIEF federation (or your own GFIPM federation) by adding the SP to the federation's trust fabric. The trust fabric update process consists of these steps:

1. Provide your SP's entity metadata to the federation manager.
2. The federation manager adds the new entity to the federation trust fabric.
3. All participants load the new federation trust fabric into their IDPs and SPs.

Before or during your deployment, you must also fill out an Implementation Documentation Form for SP and submit it to your federation manager as part of the membership application process described in Section 3.2. A template of this form is available from your federation manager. It requests the following information:

- SP software platform details (OS, Web Server, SAML Software, etc.)
- GFIPM Metadata enablement of resources
- Network configuration notes

Your ability to test your SP in the operational federation will be limited because of the lack of test IDPs in the operational federation. According to the usage policies of most SPs in an operational federation, only real users using real identities (with valid user data, permissions, and privileges) are permitted to use (i.e., test) the production systems. Therefore, when executing your SP's Test Plan in the operational federation, you must perform the necessary tests using real users (from your organization and others). As before, write a Test Report to be distributed within your organization and to all testing partners.

To publicize your organization's resources to federation users, you must supply a list of your GFIPM-available resources to the federation manager, including the following information about each resource:

- Resource name
- Resource description
- How to use the resource
- Access control policy
- Usage scenarios

Extensive examples for the above information are available at http://nief.gfipm.net/ for each of the existing participants.

4. Implementing a Web Service Consumer

This section serves as a placeholder for information on implementing a Web service consumer (WSC) within a GFIPM federation. Content for this section will be developed as GFIPM program participants and stakeholders gain knowledge about how to implement Web services in a GFIPM environment.

5. Implementing a Web Service Provider

This section serves as a placeholder for information on implementing a Web service provider (WSP) within a GFIPM federation. Content for this section will be developed as GFIPM program participants and stakeholders gain knowledge about how to implement Web services in a GFIPM environment.

6. Additional Implementation Guidance

This section highlights certain specific issues that you may experience during the implementation and testing processes. It also presents possible solutions and guidance based on experiences from previous implementers. The following issues are addressed:

1. IDP Discovery Service
2. Web browser choices and usage
3. GFIPM enablement of resources
4. Shibboleth implementation of Identity Provider and Service Provider
5. PingFederate implementation notes
6. Testing SAML interoperability
7. Implementing a reverse proxy solution
8. Implementing an HTML rewriting solution

6.1 IDP Discovery Service

An IDP Discovery Service (DS) performs the task of discovering a user's IDP and providing that information to an SP so that the SP knows which IDP to use in the subsequent single sign-on (SSO) process.

The DS provides a convenient means by which a user may specify which IDP he or she would like to use for SSO within the federation. The GFIPM Reference Federation currently uses a single DS, which is managed by GTRI; however, there is no inherent limitation on the number of discovery services that a federation can use.

All Service Providers must provide a method for users to discover their Identity Providers. You may elect to use the centralized DS provided by the federation, deploy your own DS, or implement/configure a custom discovery solution for your deployment. The NIEF federation also currently uses a single DS, but again there is no limitation on the number of discovery services in a federation. For convenience, your Service Providers may point to the central DS when a user tries to access a resource without a SAML assertion. But you are free to implement your own discovery service or an equivalent service to determine a user's IDP for SSO. In the event that a participant's Service Provider solution cannot interface with the DS, the SP must provide an equivalent functionality.

The GFIPM Reference Federation and the NIEF federation use the Switch PHP Discovery Service implementation. The Switch Discovery Service is described in detail at the following URL:

http://www.switch.ch/aai/support/tools/wayf.html

Note: The Discovery Service was previously named the "Where Are You From" (WAYF) service. Most switch documentation uses the two names interchangeably.

As an example, Figure 4 contains an image of a Web page generated by the Discovery Service in the GFIPM Reference Federation.

6.2 Web Browser Choices and Usage

The choice of Web browser is a personal and/or business decision made according to personal or corporate preferences and end-user device constraints (e.g., mobile access). In an ideal scenario, GFIPM would be able to extend this liberty to federation users, but in reality this is often not possible. During GFIPM implementations, participants have gained several important insights into the Web browser options that are available to federation users. These insights are described below.

1. Any Web browser used in GFIPM must be able to support HTTPS (HTTP over TLS), as well as HTTP redirection. All reasonably modern browsers (released in the last five years) can do this. However, as participants discovered during the project, not all modern browsers are configured by default to use TLS. While this issue is relatively easy to solve in the browser (by simply changing the browser's configuration), it can nevertheless cause usability problems, because the problem typically manifests as a cryptic Web server error that is not easily identifiable as a browser configuration problem. This problem specifically affects Internet Explorer (IE) version 6, which has TLS turned off by default. The simplest work-around for the problem is for the user to upgrade to IE version 7, which is configured to use TLS by default. If an upgrade to IE 7 is impossible, users must be given brief instructions on how to modify their IE 6 settings to enable TLS.

2. The only other constraints imposed by the GFIPM federation are specific to the limitations of certain applications within the federation. For example, if a specific application requires the use of Internet Explorer (IE) by its users prior to becoming federation-enabled, then it will almost certainly require IE for federation users as well. In these (admittedly rare) instances, users who use the Firefox browser would not be able to access the application.

3. Participants may place further constraints on browsers for their own users. For example, CISA requires that CISA users must use IE for authenticating with the CISA client certificate SSO system at

CISA's IDP. Therefore, since CISA users cannot use Firefox to authenticate to their IDPs, they also cannot use Firefox to access federation resources.

4. Interoperability problems may arise if the browser required by a user's IDP is Browser X (e.g., IE) and the browser required by an SP that the user wishes to access is Browser Y (e.g., Mozilla Firefox). However, this scenario has not happened yet. Typically, if an application requires a specific browser, the browser required is IE. There are no known instances in the federation in which an application has a browser-specific requirement for a browser other than IE.

6.3 GFIPM Enablement of Resources

This section covers common methods of enabling applications to be used with a GFIPM Service Provider. Distinctions are made between new and legacy applications and how their enablement may differ.

From the user's perspective, the implementer's primary job is to make a large set of resources available on the Service Provider. From the implementer's perspective, the implementer's primary job is to retain and enforce the usage requirements of his resources; in other words, ensuring that access control and security requirements are met by all users. Techniques exist through which a wide range of legacy resources in the law enforcement domain can be made available to federation users. The process of making a resource available to federation users is called federation enablement.

As a federation grows, resource owners will look to the GFIPM model to help them realize the following value propositions of federated identity and privilege management.

- Achieve resource sharing with a large base of established users and partners who would normally not have access to their resources, while keeping costs low.

- Provide a simplified and improved user experience (via single-sign-on access to all federation resources, subject to access control policies).

- Provide better security and privacy protection for users' personal data (via the reduction or elimination of redundant data capture and storage processes).

But achieving federation enablement for a wide range of legacy resources can be challenging in that they can be very diverse in many aspects, including application architecture, implementation platforms and vendor products, type and structure of resource, application functionality, support model, security and access policies, etc. Many insights and lessons about federation enablement of resources have been gained from current federation members during the process of federation-enabling several existing resources.

Applications and resources tend to have usage requirements that must be met by all of their users. Most usage requirements fall into the following categories:

- ***Terms of Use—***The application may require that a user agree to specific terms of use prior to using it.

- ***Provisioning—***The application may require that a user register a local account with it before using it.

- ***Intersession Persistence—***The application may need to maintain status about the user from one session to another.

- ***Identification—***The application may need to know the user's identity at all times while the user is using it.

- ***Access Control—***The application may impose certain access restrictions based on some combination of the user's rank, certifications, role, or some other important personal characteristics.

- ***Auditing—***The application may log all actions performed by a user in an audit log for review, compliance, etc.

- ***Personalization—***The application may need to maintain miscellaneous personal data about a user for the purpose of delivering certain features. For example, locality information would help the application deliver a list of alerts or bulletins that specifically pertain to a user's region or locality.

One of the fundamental tenets of the GFIPM concept is that resource owners must be able to maintain control over the usage requirements of their resources within the federation and are not forced to modify the requirements in a manner they find unacceptable. GFIPM allows a wide range of existing resources and applications to be federation-enabled and made available to federation users in a manner that fulfills the usage requirements of those applications.

The GFIPM concept provides many valuable tools that help to simplify federation enablement of resources, while still allowing those resources to meet their usage requirements.

- Federation-wide policy-level agreements and memoranda of understanding can form the basis of interagency trust, which can be layered with additional bilateral or community agreements as required.

- The basic SAML-based infrastructure provides a standard means of authentication/identification of users and the convenience of SSO.

- The GFIPM metadata provides detailed personal information about individual federation users, including identification, contact information, affiliations, memberships, certifications, and basic data access privileges within the user's home organization. This information can be trusted because it comes to the resource from a secure, trustworthy, authoritative source: the user's IDP.

The following information will help federation implementers better understand the basic federation enablement options that are available to them for various categories of resources.

6.3.1 Resource Integration Profiles

The following resource integration profiles are based on common categories of resources and applications. Their purpose is to help resource owners better understand the level of effort required to federation-enable specific types of resources. Note that a specific resource may or may not fit neatly into a specific integration profile. This section does not necessarily describe an exhaustive set of resource classes, but rather provides enough detail about the critical differences between resources to illuminate the important issues that must be addressed when federation-enabling them.

Profile 1: Read-Only Content Without Individual User Accounts

A resource in Profile 1 typically has the following characteristics:

- It is used for dissemination of information.

- It does not require a unique pre-provisioned user account for each user.

- It may require the user's identity and contact information for auditing purposes.
- It requires some basic information about the user for access control.
- It does not require personalization data.
- It has no persistence requirement.

Profile 2: Resource With Individual User Accounts and Dynamic Provisioning

A resource in Profile 2 typically has the following characteristics:

- Its provisioning requirement can be met by GFIPM Metadata and leverage the IDP user vetting without the need for any additional out-of-band communication or user vetting during the provisioning process.
- It requires the user's identity and contact information for auditing purposes.
- It requires information about the user at account provisioning time for provisioning the account's access control permissions.
- It may require personalization data.
- It has a requirement for persistence of user account information between sessions.

Profile 3: Resource With Individual User Accounts and Pre-Provisioning

A resource in Profile 3 has the following characteristics:

- It requires a unique pre-provisioned user account for each user.
- Its provisioning requirement cannot be met by GFIPM Metadata and IDP vetting alone, since it requires out-of-band communication to facilitate a direct relationship with the user during the provisioning process. However, GFIPM can provide single sign-on functionality with an account linking capability for it after the provisioning process is complete.

- It requires the user's identity and contact information for auditing purposes.
- It requires information about the user at account provisioning time for provisioning the account's access control permissions.
- It may require personalization data.
- It has a requirement for persistence of user account information between sessions.

The next section discusses the actual techniques that can be used for federation enablement of resources that fit these integration profiles.

6.3.2 Resource Integration Techniques

The following resource integration profiles are based on common categories of resources and applications. Their purpose is to help resource owners better understand the level of effort required to federation-enable specific types of resources. Note that a specific resource may or may not fit neatly into a specific integration profile. The intent of this section is not to prescribe an exhaustive set of resource integration techniques, but rather to provide enough detail about the critical differences between the techniques to illuminate the important issues that must be addressed when considering their use.

Technique 1: GFIPM-Aware Reverse Proxy and HTML Rewriting Service with No Secondary Authorization

Integration Technique 1 works best for a resource that meets the following criteria:

1. It is not natively GFIPM-aware.
2. It does not require local user accounts or user authentication as a prerequisite to access.
3. It does not require or have any notion of intersession persistence or personalization.
4. It cannot be modified at the source code level because of business/technical impracticalities.

In this integration technique, access to the resource is provided for federation members via network configuration by a reverse proxy service that is GFIPM-aware.

Also, the proxy service must implement access control and auditing if they are required by the application.

Technique 2: GFIPM-Aware Reverse Proxy and HTML Rewriting Service with Secondary Authorization

Integration Technique 2 applies to any resource that meets the following criteria:

1. It is not natively GFIPM-aware.

2. It does require local user accounts for access.

3. It cannot be modified at the source code level because of business/technical impracticalities.

4. Its access policy may require that each federated user have a unique account, or it may allow multiple users to share a group account.

5. It may or may not require intersession persistence for an account, and it may or may not have any personalization requirements.

In this integration technique, access to the resource is provided for federation members via a reverse proxy architecture in which the proxy maps each GFIPM user into a corresponding user account for the proxied resource. The mapping from GFIPM user to back-end resource account may be many-to-one (using local accounts at the resource as group accounts for GFIPM users) or one-to-one (using individual accounts). The proxy must authenticate the user to the proxied resource using the appropriate account.

If access to the resource is to be permitted via group accounts, a typical configuration would consist of relatively few local accounts on the resource (e.g., the number of group accounts in the application), with each account corresponding to an access class, group, or role that would be used by many federation users. In this configuration, since the resource knows only group accounts, it has no way to know which specific GFIPM user is accessing it at any given time. Therefore, to provide end-to-end auditing, it is necessary to combine the proxy's audit logs with the resource's audit logs. Two NIEF participants have successfully enabled federation resources using group accounts in this manner.

If access to the resource requires each GFIPM user to have a unique back-end account with the resource, it is necessary to provision individual user accounts for access. The most elegant approach to solving this problem is to allow user accounts to be provisioned via an account sign-up page on the back-end resource, and

configure the proxy to populate the resource's account provisioning (sign-up) page with GFIPM user attribute data extracted from a SAML assertion. If this type of dynamic account provisioning is not possible (for either business or technical reasons), the proxy can act as an account-linking bridge between GFIPM federated user accounts and local back-end accounts that have been pre-provisioned out-of-band.

In either scenario, the proxy must somehow know how to map each GFIPM user ID to a back-end user ID on the resource. This may be accomplished in several ways, but the most straightforward technique is for the proxy to maintain a database or map from one domain to the other. In some cases where accounts follow a regular form, a set of rules or algorithmic mapping may be possible.

Technique 3: Native GFIPM Enablement

If it is possible and practical to modify the source code of a resource during the federation enablement process, the resource can be configured so it natively understands GFIPM metadata and can exist in the GFIPM environment without the aid of a proxy.

If the resource is a newly developed application, it may be advantageous to design and build the application such that it fundamentally understands GFIPM Metadata as its primary internal model for fulfilling its requirements related to identifying the user, enforcing access control rules, and auditing access.

But if the application already exists, it may be necessary to take an alternate approach and develop a GFIPM-aware module that exists within the application and serves to translate information from GFIPM metadata into the native user model that the application uses.

Issues related to account provisioning (dynamic or out-of-band) are the same for this technique as they are in the discussion above for Technique 2.

6.3.3 Profiles and Techniques for Existing Resources

To provide a more concrete perspective on the discussion about integration profiles and integration techniques in the previous two sections, this section contains summary information[17] about how resources currently in the federation were federation-enabled.

Table 3 lists resources currently in the NIEF federation, along with their integration profiles and the integration technique used to integrate those profiles with the

[17] Note: This section was derived from content in [GFIPM Demo].

federation. It is important to note that while the table provides some insight into the breakdown of GFIPM resources for federation enablement purposes, it is not necessarily representative of the broader set of information sharing resources in the justice community.

For federation growth and outreach purposes, it is important that the federation gain some insight into how well the current set of integration techniques can accommodate the broad range of resources that are potentially applicable to GFIPM. Current and future implementers should be able to add their experience to the enablement knowledge base.

	Resource Name	Integration Profile	Integration Technique
CISA	Arizona Counter-Terrorism Information Center (ACTIC)	1	1
	Arizona Sex Offender Information Center	1	1
	Arizona Amber Alert	1	1
	Georgia Bureau of Investigation Sex Offender Registry	1	1
	Oklahoma State Bureau of Investigation Officer Safety Bulletin	3	2
	Texas Criminal Law Enforcement Online (CLEO)	3	2
	California Joint Regional Information Exchange System (JRIES)	3	2
	CISAnet Federated Query Tool • New Mexico Complete Arrest Information (CAI) • New Mexico Incident Based Reporting System (NMIBRS) • New Mexico Sex Offender Registration • New Mexico Missing Person & Unidentified Bodies • New Mexico Field Interview (FI) • New Mexico Law Enforcement Information Network with Corrections (LINC) • New Mexico Criminal Law Enforcement Reporting and Information System (CLERIS) • Arizona Criminal Investigative Database • Texas Criminal Law Enforcement Reporting and Information System (CLERIS)	2	3
JNET	Pennsylvania Department of Corrections Intake/Exit Photos	3	2
	Pennsylvania Arrest Warrants Outstanding for Parolees who Failed to Report (Absconders)	3	2
	Pennsylvania State Prisoner Locator	3	2
	Pennsylvania Criminal Trial Case Information	3	2
	Pennsylvania Arrest Warrants Outstanding for Failure to Pay Child Support	3	2
	Pennsylvania Amber Alert	3	2
	GFIPM Lessons Learned	1	3
RISS	HSIN Counterterrorism Briefs, Reports, and Documents	1	3
	RISS Counterterrorism Briefs, Reports, and Documents	1	3

Table 3: Integration Profiles and Integration Techniques for Resources in NIEF

6.4 Shibboleth Implementation

This section details technical information about current federation members' implementation strategies for possible use by prospective members.

GTRI has experience with implementing products for fielding Identity Providers and Service Providers. Most of the experience is with Shibboleth, which is a freely available, open-source SAML software package for Web single sign-on across or within organizational boundaries. It allows Web sites to make informed authorization decisions regarding user access to protected online resources in a privacy-preserving manner. Shibboleth was developed through the Internet2 initiative. The current version of Shibboleth is version 2.x.

6.4.1 Hardware Recommendations

This section describes the basic hardware, operating system, Web server, and network requirements for participants who choose to use Shibboleth in their GFIPM federation deployment.

IDPs and SPs each require server-class machines. Both machines must be of contemporary performance specifications. Specifically, the following minimum performance characteristics are recommended for each machine:

- 3.0 MHz Core2 Duo Pentium-4 (or equivalent) processor
- 2 GB of memory
- 160 GB of disk space

If a participant plans to implement multiple IDPs or SPs, additional servers are recommended. Some participants may also wish to implement their servers on virtual machines, in which case appropriate hardware should be chosen.

The only other hardware requirement is that the machines be able to run the participant's chosen operating system. Each participant must choose and install an OS platform on the machines. GTRI strongly recommends that participants choose one of the following OS platforms:

- Microsoft Windows Server 2003 or later
- Red Hat Enterprise Linux (AS or ES)

GTRI currently operates a reference implementation of an Identity Provider and a Service Provider on each of these two OS platforms. By using one of these platforms, participants can ensure that GTRI will be able to provide them with the best possible technical assistance.

Each participant must choose and install a Web server on the two machines. GTRI recommends that participants choose one of the following Web servers:

- Microsoft Internet Information Server (for Windows-based systems)
- Apache HTTP Server (for Linux-based systems)

GTRI operates a reference implementation of an Identity Provider and a Service Provider on each of these two Web servers. By using one of these Web servers, participants can ensure that GTRI will be able to provide them with the best possible technical assistance.

Participants must configure the Service Provider machine to have a static public IP address, a domain name, and Internet connectivity on port 443 (HTTPS). In addition, port 80 (HTTP) may be used for a public page of the portal or merely as a means to forward incoming connections to port 443. For security purposes, GTRI recommends that all other ports be blocked by a firewall. The Identity Provider machine must be accessible to allow local users to authenticate, but Internet access to the IDP may or may not be provided depending on local needs.

6.4.2 Install Identity Provider

The Shibboleth Identity Provider (IDP) is a standard Java Web application based on the Servlet 2.4 specification.

The official Shibboleth installation instructions are brief and lacking in some details. They are available from the Shibboleth Web site at:

https://spaces.internet2.edu/display/SHIB2/IdPInstall

You must follow the above instructions to install Shibboleth. Additional Shibboleth installation instructions with more descriptive details are provided below for your convenience. You should follow both sets of instructions simultaneously.

> *Note that the instructions below make certain assumptions about version numbers for software packages and file names. These version numbers reflect the most recent versions available at the time of this writing. You may need to make minor adjustments during this process, depending on the versions that are available as you work through the installation. UNLESS NOTED OTHERWISE, YOU SHOULD ALWAYS USE THE LATEST VERSION OF EACH SOFTWARE PACKAGE AT THE TIME OF YOUR INSTALLATION.*

Download file shibboleth-idp-2.1.2-bin.zip as directed and unzip the file (it will unzip into the directory `./identityprovider`). Follow the official Shibboleth installation instructions above and install Shibboleth IDP into the following directory:

/opt/shibboleth-idp-2.1.2 (or a directory of your choice).

The following is the output of running the ANT script to install the files:

```
# ./ant.sh
Buildfile: build.xml

install:
Is this a new installation? Answering yes will overwrite your current configuration.  [yes|no]
yes
Where should the Shibboleth Identity Provider software be installed? [default: /opt/shibboleth-idp-2.1.2]
What is the hostname of the Shibboleth Identity Provider server? [default: idp.example.org]
idp.cisanet.net
A keystore is about to be generated for you.  Please enter a password that will be used to protect it.
XXXXXXXX
Updating property file: /home/gtri/downloads2/identityprovider/install.properties
Created dir: /opt/shibboleth-idp-2.1.2/bin
Created dir: /opt/shibboleth-idp-2.1.2/conf
Created dir: /opt/shibboleth-idp-2.1.2/credentials
Created dir: /opt/shibboleth-idp-2.1.2/lib
Created dir: /opt/shibboleth-idp-2.1.2/lib/endorsed
Created dir: /opt/shibboleth-idp-2.1.2/logs
Created dir: /opt/shibboleth-idp-2.1.2/metadata
Created dir: /opt/shibboleth-idp-2.1.2/war
Generating signing and encryption key, certificate, and keystore.
Copying 5 files to /opt/shibboleth-idp-2.1.2/bin
Copying 8 files to /opt/shibboleth-idp-2.1.2/conf
Copying 1 file to /opt/shibboleth-idp-2.1.2/metadata
Copying 38 files to /opt/shibboleth-idp-2.1.2/lib
Copying 4 files to /opt/shibboleth-idp-2.1.2/lib/endorsed
Copying 1 file to /home/gtri/downloads2/identityprovider/build/WEB-INF
JARs are never empty, they contain at least a manifest file
Building jar: /opt/shibboleth-idp-2.1.2/war/idp.war

BUILD SUCCESSFUL
```

After the installation, further configuration operations are required as follows.

Note: Some file names contain a version number (2.0.2 or 1.5, etc.). In some cases, you may find files with a later version than listed here. Unless otherwise noted, you should use the later versions.

Create required link:

```
cd /opt
```

```
ln –s shibboleth-idp-2.1.2 shib-idp
```

Edit file **/opt/ shib-idp/conf/relying-party.xml**:

- Modify MetadataProvider of type FileBackedHTTPMetadataProvider to retrieve the metadata from http://gfipm.net/metadata/gfipm-signed-metadata.xml.
- Modify security:Credential to refer to the correct key and crt files.
- *Uncomment* and modify security:TrustEngine to refer to the correct GFIPM CA crt file.

Edit file **conf/attribute-filter.xml**:

- Add <AttributeFilterPolicy id="releaseGFIPM"> … for GFIPMAssertion-1.0.

Edit file **conf/attribute-resolver.xml**:

- Get a new, slimmer version from GTRI and edit it.
- Modify resolver:DataConnector to use CISA's LDAP directory.

Create file **gfipm/dsml2gfipm_cisanet.xsl**.

Create all files in **gfipm/xsd**.

Take the following files from **`Conn2.0.2b.zip`** and copy them to **/opt/shib-idp/lib**:

- castor-1.0.jar
- castor-1.0-srcgen-ant-task.jar
- castor-1.0-xml.jar

Take the following file (33298 bytes, dated Feb 20, 2008 15:56) and copy it to **/opt/shib-idp/lib**:

- gfipm-shib-trunk-jdk-1.5.jar

Note that the above four files also need to be copied to directory /opt/tomcat5/CISAIDP/Webapps/idp/WEB-INF/lib

Additionally, files in **./credentials** and **./metadata** need to be configured from certificate files set up during your Web server configuration. If you would like

information on how to install or configure an Apache Web server and/or Tomcat servlet engine with Shibboleth, please contact gfipm-support@lists.gatech.edu to request available instructions.

Create directory and change ownership to the **tomcat** user:

```
cd /opt/shib-idp
mkdir users
chown -R tomcat:tomcat.
```

Testing the IDP

The basic method for testing the Shibboleth IDP is to use the status URL, which is of the form:

https://FQDN/idp/profile/Status

where FQDN is your IDP machine's fully qualified domain name. If it is working, it will show a Web page with a simple

```
ok
```

in the Web browser. If it is not working, it will show one of several error messages (and also check the log files as described below).

There are two sets of log files that you can monitor:

- Tomcat servlet engine log files:
  These are typically located in the [Tomcat_Install_Dir]/logs/ directory. There are several log files in this directory, but the main file is usually named catalina.out.
  If a different Java servlet container than Tomcat is used, the log files should be in that particular container's log directory.

- Shibboleth IDP log files:
  These are typically located in the [IDP_Install_Directory]/logs/ directory. The file idp-process.log (or similar) is relevant for tracing operations in a minimal IDP installation. The idp-access.log and idp-audit.log files are generally relevant only on fully functional systems.

6.4.3 Install Service Provider

The Shibboleth Service Provider (SP) is a stand-alone daemon running as a background process.

To read the Shibboleth 2 release announcement, go to:

http://shibboleth.internet2.edu/shib-v2.0.html

To download the necessary software for Shibboleth 2 SP, go to the "Downloads" section. For the Linux RHEL RPMfiles, go down to the RPMS directory.

Note that the instructions below make certain assumptions about version numbers for software packages and file names. These version numbers reflect the most recent versions available at the time of this writing. You may need to make minor adjustments during this process, depending on the versions that are available as you work through the installation. UNLESS NOTED OTHERWISE, YOU SHOULD ALWAYS USE THE LATEST VERSION OF EACH SOFTWARE PACKAGE AT THE TIME OF YOUR INSTALLATION.

Download these files (version numbers may be slightly different):

1. log4shib-1.0-1.i386.rpm
2. xerces-c-2.8.0-1.i386.rpm
3. xml-security-c-1.4.0-1.i386.rpm
4. xmltooling-1.0-6.i386.rpm
5. opensaml-2.0-6.i386.rpm
6. shibboleth-2.0-6.i386.rpm

NOTE: To uninstall an old Shibboleth version, uninstall the above packages in reverse order, as given above.

Use the rpm package command to uninstall:

```
rpm -e -v <package-name>
```

Official Shibboleth 2 installation instructions may be found at:

https://spaces.internet2.edu/display/SHIB2/NativeSPLinuxInstall

Follow the above official Shibboleth installation instructions carefully and fully, and then follow the detailed configuration instructions given below.

Edit file `/etc/shibboleth/shibboleth2.xml`:

- **MetadataProvider—**specify the URI to the http/xml metadata file. This file might not yet be available because 1) the metadata doesn't exist until there is at least 1 SP and 1 IDP deployed or 2) the file is not yet globally available. In these cases, either use a local xml file of your IDP metadata (built during the installation of your IDP) or a copy (if available) of the federation metadata file.

  - URL is http://gfipm.net/metadata/gfipm-signed-metadata.xml

- Change all instances of **https://sp.example.org** to your SP URL.

- Change all instances of **https://idp.example.org** to your IDP URL.

- Change all instances of **https://ds.example.org/DS** to your Discovery Service URL (note that your DS host may use http).

- The **CredentialResolver** section needs the service provider's certificate and key files added.

- **SignatureMetadataFilter—**provide the root certificate for the GFIPM CA.

- **SessionInitiator** – remove 'relayState="cookie"' to work around a known bug in v 2.0.

Edit the file `conf/attribute-map.xml` so that it exports all GFIPM user attributes required by your Service Provider. For example:

```
    <!-- GFIPM 2.0 Attributes -->
<Attribute name="gfipm:2.0:user:LocalId" id="GfipmLocalId">
      <AttributeDecoder xsi:type="StringAttributeDecoder" caseSensitive="true"/>
    </Attribute>
    <Attribute name="gfipm:2.0:user:GivenName" id="GfipmGivenName">
      <AttributeDecoder xsi:type="StringAttributeDecoder" caseSensitive="true"/>
    </Attribute>
    <Attribute name="gfipm:2.0:user:SurName" id="GfipmSurName">
      <AttributeDecoder xsi:type="StringAttributeDecoder" caseSensitive="true"/>
    </Attribute>
    <Attribute name="gfipm:2.0:user:EmailAddressText" id="GfipmEmailAddress">
      <AttributeDecoder xsi:type="StringAttributeDecoder" caseSensitive="true"/>
    </Attribute>     …  <!—Snipped for brevity, but all relevant attributes would be included
typically -->
```

Not making this modification will cause the SP to receive no metadata, even though both the IDP and SP seem to be working (no error messages).[18]

Add participant SP public and private key certificates to **`/etc/shibboleth/ssl/.`**

- cisanetsp.crt
- cisanetsp.key
- Also specify these filenames in file **shibboleth2.xml** in section CredentialResolver.

Verify that file **/etc/init.d/shibd** looks correct (generally, it is correct as is).

Logging output:

- To get more details in the log output, modify file `shibd.logger`: (on line number 2)
  - change: log4j.rootCategory=INFO, shibd_log
  - to: log4j.rootCategory=DEBUG, shibd_log
- By default, log files are written to directory `/var/log/shibboleth.`

To get the Shibboleth SP metadata file (to send to GTRI), use a browser to go to the following:

- https://sp.cisanet.net/Shibboleth.sso/Metadata.
- Save to sp-metadata.xml and send to GTRI to incorporate into the GFIPM entities metadata.
- The GFIPM entities metadata will be available at http://gfipm.net/metadata/gfipm-signed-metadata.xml.

The installation should be complete at this point.

To start or stop the Shibboleth process, use the script /etc/init.d/shibd:

```
/etc/init.d/shibd start
/etc/init.d/shibd stop
```

[18] GTRI makes available a sample `attribute-map.xml` file with all GFIPM Metadata 2.0 user attributes in it.

6.4.4 Known Issues

The Shibboleth SP is typically installed on the Tomcat application server [Tomcat], which may be running behind the Apache httpd Web server. A connector is required to connect Apache and Tomcat.

6.5 PingFederate Implementation

GTRI does not have any direct experience with implementing PingFederate; however, one NIEF member is currently using PingFederate, and GTRI has provided some assistance during the implementation process. During our limited experience, we found only one known issue. Specifically, PingFederate is not able to parse a GFIPM Trust Fabric file, which includes entries for multiple entities (IDPs and SPs). PingFederate requires that each entity in the federation be defined individually. You should be able to extract individual entities out of the GFIPM Trust Fabric file and load them into PingFederate according to the PingFederate documentation.

6.6 Testing SAML Interoperability

The GFIPM federation implements the SAML Web Single Sign-On Profile. When an IDP and/or SP product is initially deployed, one of the first tests to be performed is the SAML interoperability test. This ensures that the IDP or SP is able to communicate with other SPs or IDPs at the SAML level.

This test is performed after the product has been fully deployed (for IDPs, see Section 2.6 and for SPs, see Section 3.6) by integrating them into the federation Trust Fabric.

The SAML-interoperability test consists of the following steps:

1. Protect a resource with your SAML product so it requires Single Sign-On (SSO).

2. Access that resource to initiate SSO.

3. The SP should generate an "Authn Request," which is sent to the IDP.

4. The IDP should process this request, authenticate the user, and generate a SAML response containing a SAML assertion containing GFIPM user attributes.

5. Verify that the SP can consume the SAML response and extract the GFIPM user attributes from the SAML assertion.

For this test, it is NOT required that the GFIPM user attributes be created from a local user directory, or that they be parsed at the SP. The purpose of the test is simply to verify the SAML level communications process between IDP and SP.

6.7 Implementing a Reverse Proxy Solution

This section serves as a placeholder for information on implementing a reverse proxy solution within a GFIPM federation. Content for this section will be developed as GFIPM program participants and stakeholders gain knowledge about how to implement a reverse proxy in a GFIPM environment.

6.8 Implementing an HTML Rewriting Solution

This section serves as a placeholder for information on implementing a HTML Rewrite Engine within a GFIPM federation. Content for this section will be developed as GFIPM program participants and stakeholders gain knowledge about how to implement a HTML Rewrite Engine in a GFIPM environment.

Appendix A: GFIPM Reference Federation

This appendix provides background information about the GFIPM Reference Federation. It also provides guidance on how to use the GFIPM Reference Federation to test IDPs and SPs without the concern of exposing sensitive live data.

The GFIPM Reference Federation contains a collection of Internet-accessible IDPs and SPs that are configured for testing purposes. It also contains an IDP Discovery Service (DS).

GTRI maintains the GFIPM Reference Federation and makes it available to prospective NIEF and other GFIPM program stakeholders members for proof-of-concept and interoperability testing. The federation consists of several core software components located in a laboratory at GTRI, as well as other IDP and SP systems that may be connected to the federation from time to time. The core components of the federation are accessible via the Internet and are generally available 24/7; however, GTRI and other GFIPM stakeholders also use it as a test bed, so some of its components may be unavailable at times. To ensure availability of the GFIPM Reference Federation, GTRI recommends that you arrange an appointment at which to perform your formal testing.

To use the GFIPM Reference Federation, you must meet the following configuration and networking requirements:

1. Set up and configure your organization's test IDP and/or SP in your test environment.

2. Your test user(s) should use a workstation (typically a Windows PC) with an industry-standard Web browser (typically Internet Explorer).

3. Your testing environment must have the following network connectivity:

    a. Your test users must have network connectivity to your IDP.

    b. Your test users should have network connectivity to the Reference IDP. While this is not required, it is a great convenience, because the Reference IDP contains a number of useful test identities that can be used by any test users.

    c. Your test users must have network connectivity to your SP.

d. Your test users must have network connectivity to the GFIPM Reference Federation SPs and Directory Service (DS) over the Internet.

e. The GFIPM Reference Federation SPs and DS are listening on ports 80 and 443 on the Internet, so your site must allow outbound traffic on these ports.

f. Other test users (such as at GTRI) must have network connectivity to your SP over the Internet. This means that your test environment must allow inbound traffic to your SP server, usually on ports 80 and 443 (or possibly other ports depending on your SP).

g. Note: Inbound traffic from the Internet to your IDP is not required for other test users or SPs unless your own test users need to come in from the Internet (and they are not using a VPN or other security layer). Due to security considerations, participants' IDPs are normally not exposed to the Internet.

Never disseminate "live" or real data via the GFIPM Reference Federation. It is imperative that your test IDPs and SPs contain no real data when testing in the GFIPM Reference Federation or any other nonproduction GFIPM testing environment.

Agencies that are interested in the GFIPM program are invited to join the GFIPM Reference Federation to learn more about operating within the federation. By joining the GFIPM Reference Federation, an agency can do all of the following:

- Verify proper generation and processing of GFIPM Metadata.
- Verify interoperability with other federation members.
- Learn how to deploy SAML2 software.
- Test new deployment strategies.
- Use test identities to test its SP.
- Test new services with reference IDPs.
- Test IDPs with reference SPs.
- Examine and analyze other agencies' reference SPs.

The early federation components were reference implementations of each major functional component. These reference components were deployed by GTRI, and they serve two purposes. First, the process of deploying and maintaining the reference components serves as a valuable learning experience and a source of documentation artifacts that may be used by other participants. Second, the reference components themselves provide a valuable testing platform for each

organization's IDP and/or SP deployments. Each reference component is discussed individually in the subsections below.

Reference Identity Provider (IDP)

GTRI deployed two reference IDPs in the pilot federation. Both IDPs are based on the Shibboleth 2.x implementation of SAML 2.0. One of the reference IDPs is deployed on a Microsoft Windows platform, the other on Red Hat Enterprise Linux (RHEL). There is no functional difference between a Shibboleth IDP running on Windows and one running on RHEL; however, the deployment processes for a Shibboleth IDP on each platform are different enough to merit the task of working through each and documenting them separately. During and after the deployment process, GTRI created a detailed set of instructions for deploying a Shibboleth IDP on each platform.

Both IDPs address two integration issues. The first is the Single Sign-On Integration Point, which involves the integration of the IDP with the local site's user authentication system, and the other is the Attribute Authority Integration Point, which involves connecting the IDP to the local site's attribute repository.

During the deployment of the reference IDPs, the Windows-based reference IDP was integrated with a username/password authentication system, and the RHEL-based reference IDP was integrated with a PKI-based client certificate authentication system. At the Attribute Authority Integration Point, both reference IDPs were connected to a reference LDAP repository.

Since their deployment, it has become clear that both reference IDPs are very useful during the process of deploying new SPs. They contain a large number of test identities that are designed to allow for testing a wide variety of metadata attributes. Additional credentials for test accounts on each reference IDP can be provided to new participants, and these test accounts have repeatedly proven to be a valuable resource for participants who have brought their SPs online.

Some of the current participants continue to maintain their reference IDPs online on a full- or part-time basis. For example, CISA maintains its reference IDP (a full test copy of its production IDP) for its testing use. On request, CISA administrators and users can test other participants' reference SPs with their reference IDPs.

Other reference IDPs are occasionally available in the GFIPM Reference Federation. These IDPs belong to other participants and can be used only by the users of those participants.

A representative list of reference IDPs can be seen in the drop-down menu in Figure 4.

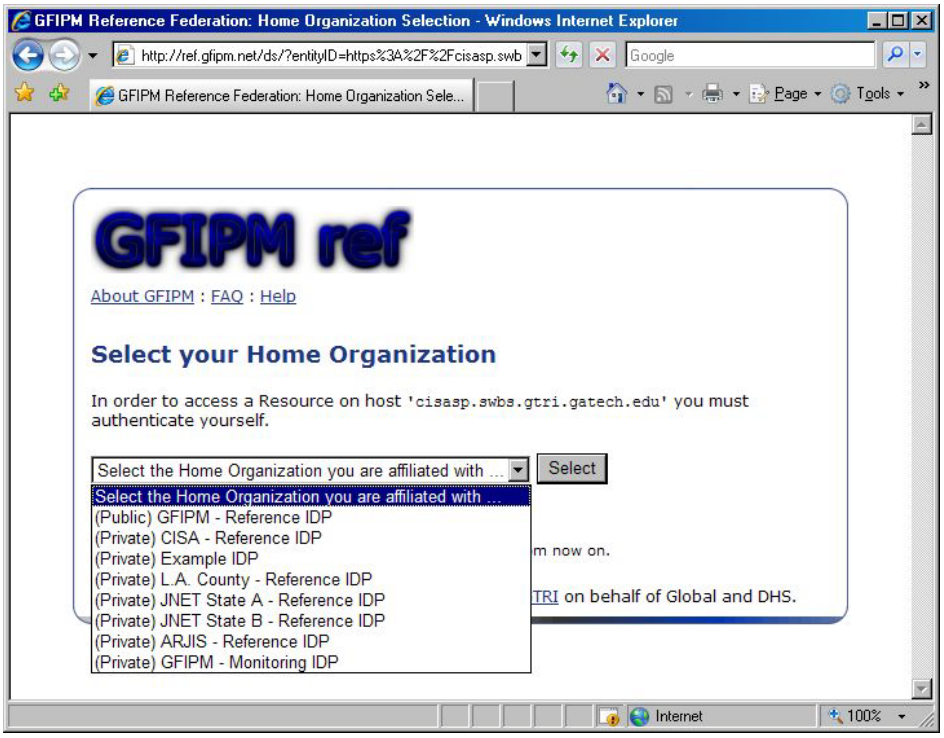

Figure 4: Screen Shot Containing a List of Available Reference IDPs

Reference Service Provider (SP)

In addition to deploying reference IDPs, GTRI has also deployed two reference SPs in the GFIPM Reference Federation. Both SPs are based on the Shibboleth 2.x implementation of SAML 2.0. As with the reference IDPs, one of the reference SPs was deployed on a Microsoft Windows platform, and the other on RHEL. Again, as with the reference IDPs, there is no functional difference between a Shibboleth SP running on Windows and one running on RHEL; however, the deployment processes for a Shibboleth SP on each platform are different enough to merit the task of working through each and documenting them separately. During and after the deployment process, GTRI created a detailed set of instructions for deploying a Shibboleth SP on each platform.

The integration work required for an SP involves setting up the SP to provide protected access to sensitive resources (see Section 6.3 for options on how to do this). During the deployment of the reference SPs, GTRI created some simple HTML and PHP pages to serve as protected resources. These pages serve two important purposes. First, they help participants debug various problems with their IDPs at the SAML configuration level. Second, they allow for careful inspection of the GFIPM user metadata that an IDP sends to a reference SP. This feature has been very valuable in helping participants identify and correct problems related to the generation of metadata by their IDPs.

As with the reference IDPs, participants have found the reference SPs to be useful during the deployment process for their infrastructure. Participants are able to test their IDPs by attempting to access resources on the reference SPs.

Some of the current participants continue to maintain Reference SPs online within the GFIPM Reference Federation on a full- or part-time basis. For example, CISA maintains a reference SP for its testing use. The CISA reference SP provides links to test resources for testing purposes and can also be used to test the user metadata from another participant's IDP. Other reference SPs from other participants may occasionally be available in the GFIPM Reference Federation. When available, these SPs can be used by the users of any participant with a referenced IDP or by using the GFIPM Reference IDP.

Reference IDP Discovery Service (DS)

The final reference component in the GFIPM Reference Federation is the IDP Discovery Service (DS). The DS allows a convenient means for a user to specify which IDP he or she would like to use for single sign-on within the federation. The GFIPM Reference Federation currently uses a single DS, which is managed by GTRI; however, there is no inherent limitation on the number of discovery services that a federation can use.

Figure 4 depicts a Web page generated by the Reference Discovery Service.

Participants in a GFIPM federation need not implement their own Discovery Service. Instead, your Service Providers can redirect to the central DS when a user tries to access a resource without a SAML assertion. If a participant's Service Provider solution cannot interface with the DS, the SP must provide an equivalent functionality.

Useful GFIPM Reference Federation Information

The GFIPM Reference Federation contains useful test documentation as well as reference SPs and IDPs, the Discovery Service, and the signed federation trust fabric document. These items are summarized below with their respective URLs.

- GFIPM Reference Federation Home: http://ref.gfipm.net/
  This Web site offers an introduction to current members and prospective members of a GFIPM federation for the purpose of getting started using the GFIPM Reference Federation. The GFIPM Reference Federation is a public federation that agencies interested in GFIPM are invited to join to learn more about operating within a federation. Topics covered on this Web site include the following:

- Overview and purpose
- Information for participating
- Members and their reference resources
- Downloads page
- FAQ
- How to get more help

- Reference SP: https://rhelsp.ref.gfipm.net
  This test Service Provider contains one Shibboleth Protected Resource, which acts as a protected resource that requires authentication at a GFIPM Reference Federation IDP. When you try to use the resource, you will be redirected to the GFIPM Reference Federation's Directory Service.

- Reference IDP: https://rhelidp.ref.gfipm.net
  This test Identity Provider contains multiple test GFIPM user attribute sets for use by federation members in SAML assertions for testing. These attribute sets are suitable for testing a new Service Provider in the GFIPM Reference Federation. The attribute sets represent identities with a wide variety of authentication and privilege information. There are also multiple similar user attribute sets that vary only slightly among themselves so that testers can observe small privilege changes on their SPs.

  Important: Because these user attribute sets do not represent real people, they must not be used to access live data.

- GFIPM Reference Federation IDP Discovery Service (DS): http://ref.gfipm.net/ds/
  The DS is a service that performs the task of discovering the user's IDP and providing that information to the SP so that the SP knows which IDP to use in the subsequent SSO process.

- Federation Trust Fabric File: http://ref.gfipm.net/gfipm-signed-ref-metadata.xml
  A document signed by the Federation Manager Organization, containing trusted information about each IDP and SP in the federation. It includes X.509 certificate data for each software entity, as well as a GFIPM Entity Assertion providing various informational attributes about each entity. This GFIPM Trust Fabric is the cryptographic trust anchor for all federation transactions. Before any new SP or IDP can join the GFIPM Reference Federation, the federation manager must first enter it into this file. All operational SPs and IDPs must download and use

this file. In addition, the providers must periodically check for new versions and download them (new versions are typically announced to the participant administrators by e-mail).

- CISA Reference SP: https://cisasp.swbs.gtri.gatech.edu
  This is a complete test version of the CISAnet production SP with test resources and access control rules suitable for testing IDPs and test user identities.

Appendix B: Acronyms

This section provides a list of acronyms used throughout this document.

CISA	Criminal Information Sharing Alliance
COTS	Commercial Off-The-Shelf Software
CRL	Certificate Revocation List
CSP	Credential Service Provider
DHS	U.S. Department of Homeland Security
DOJ	U.S. Department of Justice
DS	Discovery Service
FIPM	Federated Identity and Privilege Management
GFIPM	Global Federated Identity and Privilege Management
GFIPM DT	GFIPM Delivery Team
GSWG	Global Security Working Group
GTRI	Georgia Tech Research Institute
HSIN	Homeland Security Information Network
IDP	Identity Provider
IDPO	Identity Provider Organization
JNET	Pennsylvania Justice Network
NIEF	National Information Exchange Federation
PKI	Public Key Infrastructure
PMO	Program Management Office
RA	Registration Authority
RHEL	Red Hat Enterprise Linux
RISS	Regional Information Sharing Systems
SAML	Security Assertion Markup Language
SP	Service Provider
SPO	Service Provider Organization
SLO	Single Log-Out
SSL	Secure Sockets Layer
SSO	Single Sign-On

TLS	Transport Layer Security
URI	Uniform Resource Identifier

About Global

The U.S. Department of Justice's Global Justice Information Sharing Initiative (Global) serves as a Federal Advisory Committee to the U.S. Attorney General on critical justice information sharing initiatives. Global promotes standards-based electronic information exchange to provide justice and public safety communities with timely, accurate, complete, and accessible information in a secure and trusted environment. Global is administered by the U.S. Department of Justice, Office of Justice Programs, Bureau of Justice Assistance.

For more information on DOJ's Global and its products, including those referenced in this document, visit

www.it.ojp.gov/gfipm

or call

(850) 385-0600