pennsylvania
OFFICE OF ADMINISTRATION
INFORMATION TECHNOLOGY

Information Technology Policy

Directory Services Architecture

ITP Number	***Effective Date***
GEN-SEC013B	June 22, 2006
Category	***Supersedes***
Recommended Policy	
Contact	***Scheduled Review***
RA-ITCentral@pa.gov	Annual

1. Introduction

The purpose of this document is to define the Enterprise Directory Services architectural policy that was established in ITP-SEC013 - *Identity Protection and Access Management (IPAM) Architectural Standard-Identity Management Services.*

Numerous commonwealth applications store and retrieve identity-related information, including identity credentials, using various data repositories generically referred to as identity stores. Many of these identity stores are directories, which are specialized databases optimized for reading and searching. Examples of identity stores include e-mail or Network Operating System (NOS) directories, HR databases, and flat files of data such as XML or delimited text. They are used by identity-centric applications such as for human resource management, commonwealth access management systems, subscriber services management, emergency response management and law enforcement.

A proliferation of identity-centric applications has resulted in the creation and administration of a substantial number of disparate identity stores across the commonwealth enterprise. Such incompatible data sources often have separate security protocols, complicating authorization and authentication tasks, and limiting interoperability between commonwealth agencies, business partners, other states, and federal agencies like Homeland Security or law enforcement.

The IPAM initiative establishes a shared identity information store, the Commonwealth of Pennsylvania Enterprise Directory (CoPED), to provide agencies with shared identity services and technologies to support their identity-centric applications. This policy ITP establishes the architecture of that directory and those services. These services incorporate virtual directory technology to provide consolidated and cross-sectional views of the three separate identity domains: *Employee*, *Business Partner*, and *Subscriber*. Metadirectory technology is used for synchronization. Implementation strategies and best practices for migration from legacy systems and processes are provided in BPD-SEC013H -*Directory Services Implementation Guide*. The directory technologies are named in SEC014A - *IPAM Technical Architectural Standards – Identity Management Services.*

1.1 Organization

This document provides the following information:

- Section 2 CoPED Directory Services
- Section 3 Directory Standards
- Section 4 Synchronization Services
- Section 5 Identity Data Management.
- Section 6 Security
- Section 7 Governance and Administration

Note: References and acronym definitions are provided in APP-SEC013A - *IPAM Glossary*.

2 CoPED Directory Services

The architectural model for the Commonwealth of Pennsylvania Enterprise Directory (CoPED) has several key characteristics. The first characteristic is that the directory conforms to a distributed model; it is comprised of three logically separate domains, *Employee*, *Business Partner*, and *Subscriber*, which together describe all identities that are to be stored in CoPED. Directory, virtual directory, and metadirectory technologies are available for agencies to access and use data from one domain at a time, or combine data from multiple domains with their own internal identity stores to create single consolidated views of data in response to specific application and reporting needs.

The second characteristic is that agencies are to use the IPAM directory technologies to maintain their own "master" directories. An agency's master directory represents a consolidated, multifunctional view of its own internal array of identity stores, useful such as internal application development and reporting. Master directories are also important for synchronizing with CoPED because agencies are responsible for populating the attributes in CoPED, and agency master directories can reduce the number of necessary CoPED connections to one per agency.

The remainder of this Section is devoted to further detailing this model. The directory architecture is laid out in Section *3 Directory Standards,* and directory synchronization is explained in *Section 4, Synchronization Services.*

An intermediate model that provides agencies with a jumpstart to develop their master directories, as well as the implementation architecture to be used for CoPED in general, is detailed in BPD-SEC013H - *Directory Services Implementation Guide*.

2.1 Strategic Purpose

CoPED serves two strategic purposes. The first is as the repository for identity authentication credentials for the Shared Authentication Service described in the supporting document, GEN-SEC013C, *Access Management and Control*. These include:

- The CoPED Global Unique Identifier (GUID)
- The identity's UserID and password (where applicable)
- The X.509 certificate holding the identity's public key

- The proofing level assigned to the agency process that was used to vet the identity before granting the credentials (see GEN-SEC013D, E*nrollment, Identity Proofing and Vetting*)
- Various IDs as needed to index into other agency and CoP identity stores

The CoPED second strategic purpose is to store identity attributes that are shared across many agencies and CoPs to support services and applications, as described in Section Section *2.3, Directory Integration*. These include:

- The identity's history of agency proofing levels and vetting processes.
- General authorization role information such as organization affiliations and associated IDs for duplicate accounts.
- Contact information.
- Attributes leveraged by portal and other enterprise applications.

Note: CoPED is not a general attribute directory providing a repository for hosting all identity attributes. The core identity attributes maintained in CoPED are listed in OPD-SEC014F - *Commonwealth of Pennsylvania Enterprise Directory (CoPED) Schema.* Policy governing the selection of additional attributes is discussed in Section *3.1.1 Attribute Selection Policy*.

2.2 Distributed Model Overview

CoPED consists of three logically separate domains: *Employee*, *Business Partner*, and *Subscriber*. Generally, a given identity does not exist more than once in a single domain, and agencies are directed to apply due diligence to prevent such occurrences from happening. However, an individual may well have a separate identity such as with separate user IDs and credentials in each of two or more of the domains. For example, a commonwealth employee would have an identity in the *Employee* domain for performing job-related functions and an identity in the *Subscriber* domain to use for non-work related functions such as driver's license renewal.

The *Employee* domain contains a comprehensive list of employees, and is to be used for employee-centric applications.

The *Subscriber* domain includes individuals who access the commonwealth's Web sites to receive services, or who are enrolled in one or more commonwealth service program. Although primarily composed of state residents, this domain may also include individuals who are not residents of Pennsylvania.

The *Business Partner* domain is reserved for all persons providing non-employee services to or on behalf of the commonwealth. These persons may also be subscribers or employees, but are tracked for something other than employee or subscriber-related purposes.

Note: Within the *Business Partner* domain, significant hurdles exist for establishing a single account for each unique entity, due to the varying nature of how these identities are created and used by the different agencies:

- The same entity might be used for different business purposes.

- Organizations contain multiple individuals; while each of those individuals might exist as separate entities, agencies might also be identifying the organizations themselves as entities (IPAM is to be used only for personal identities, not organization identities).
- An individual might have multiple roles that make it difficult to record this person as a single business identity.
- A single sub-contractor may be a member of multiple teams or multiple businesses.

Therefore, although the *Business Partner* domain, like each of the other domains, is intended to store each individual as a unique identity, it is permissible to assign an individual multiple identities in the *Business Partner* directory when deemed necessary by the agency.

In order to distinguish and ensure integrity of these identities, a format was established for the GUID used to identify Business Partner Organizations within the CoPED Services Model.

The GUID will contain the first ten letters/numbers from the company name, and the last four digits of the Federal Employer Identification Number (FEIN) to create the Organizational Unit (OU) in the Active Directory structure. The entire GUID will then be encrypted and stored to avoid duplication. A Business Partner from the same company will be added to the existing OU based on the information they provided (Company name/FEIN).

To account for any instance where a Business Partner Organization's first ten letters/numbers from their name and the last four digits of the FEIN are identical to one another, the GUID will be appended with a sequential letter/number combination.

The letter/number combination format starts with X0 (zero). If additional combinations are needed, they will continue with X1 through X9. If more than ten unique identifiers are required, the sequence continues with Y0 (zero) though Y9, and finally Z0 (zero) through Z9. This format allows for an additional thirty unique identifiers for each GUID.

2.3 Directory Integration

The CoPED architecture incorporates virtual directory technology to provide integrated views of the three domains, allowing agencies to develop composite data sets of entities and attributes from multiple domains, while the domains themselves remain physically separate. Integrated security mechanisms prevent unauthorized persons from viewing or editing restricted identities or their attributes, as determined by the data's authoritative source, governing authority or hosting agency.

Agencies may also choose to use the virtual directory technology to create virtual instances of their master directories, combining CoPED and internally stored identities and attributes into a single composite view; or they may choose metadirectory technology to create and populate "actual" instantiations of their master directories. Because metadirectory instances persist even while disconnected from the source directories, they can be used to support applications that are disconnected from the commonwealth network, or which require new attributes that do not already exist in CoPED or the agency's existing identity

stores. The metadirectory technology is discussed thoroughly in Section *4, Synchronization Services.* The importance of agency master directories is pursued in Section *2.5 Agency Master Directory*.

2.4 Virtual Directory

Virtual directories provide a real-time or near real-time, customizable view of all or part of one or more identity stores, and include features such as caching, filtering, data transformations, and schema transformations. The CoPED virtual directory configuration includes:

- Redundant virtual directory engines, which:
  - Provide administration tools that allow user-customized configuration.
  - Interact with a configuration store to save and retrieve configurations.
  - Cache directory data to improve performance
  - Filter the directory tree so that only appropriate objects are published (retired employees are filtered so that only current employees may be seen in the *employee* directory branch).
  - Filter the schema to only publish appropriate attributes (if the source identity store includes SSN, do not publish SSN in CoPED).
  - Transform the published data. Set a value to the Proofing Level attribute based on which identity store sourced that person.
  - Transform the published schema. Read the eight-digit driver's license ID as numeric, but publish it as an eight-character alphanumeric.
- Back-end access to source identity stores via Lightweight Directory Access Protocol (LDAP) or Structured Query Language (SQL)/Open Database Connectivity (ODBC)

2.5 Agency Master Directory

As previously discussed, agencies may leverage the commonwealth's investment in enterprise directory technology by using it to develop their own master directories. An agency's master directory represents a consolidated, multifunctional view of all identities in all of its own internal identity stores. A master directory may also include identities from CoPED that are not otherwise stored by the agency, if the agency deems it useful to do so.

The master directory is useful as a common identity source (such as internal applications and reporting). The master directory is also important for synchronizing with CoPED because it reduces the number of necessary connections for each agency to one. Agency master directories can be used to update CoPED dynamically via virtual directory technology, or as a scheduled task with metadirectory synchronization technology according to agency needs and the impact on system performance (see *Section 4, Synchronization Services*). Thus, master directories simplify CoPED administration and provide greater flexibility to agencies and CoPs for control over their own data. The virtual directory technology can be used by agencies to create consolidated views of their own internal identity directories to simplify internal application development and reporting, or to create completely new views that combine their internal data with attributes from CoPED and other agencies. The following sections explain these technologies and strategies in detail.

BPD-SEC013H, *Directory Services Implementation Guide,* provides guidance for the implementation of these technologies.

Agencies may leverage either the virtual directory technology described above or the metadirectory technology described in *Section 4, Synchronization Services,* to build their master directories. Each agency also has the option of creating a separate directory for each domain, or a single directory with an attribute that explicitly identifies which domain each identity belongs to. Agencies that have not yet created their master directories are directed to BPD-SEC013H, which provides the Community Model Hybrid to assist in transitioning to the Shared Directory Services architecture.

3 Directory Standards

CoPED follows all standards as named in the appropriate commonwealth ITPs. In particular, all access to CoPED (read or write, for authentication or attribute gathering) is to use the LDAP v3. All Web Service (XML) access uses the Directory Services Markup Language (DSML) v2 or later. All Web Services interfaces are to support Simple Object Access Protocol (SOAP) and WS-Security Version 1.1 or later.

3.1 Schema Architecture

The CoPED schema adheres to the appropriate Internet Engineering Task Force (IETF) Request for Comments (RFC). In particular, the attribute syntax definitions follow RFC 2252 and the standard base object class for user identities is *inetOrgPerson* as defined in RFCs 2256 and 2798. This defines a number of the attributes used for CoPED. A complete list of the core CoPED attributes, along with the attribute characteristics most relevant to IPAM, are listed in OPD-SEC014F - *CoPED Schema*.

Although *inetOrgPerson* provides many of the required attributes, the architecture provides additional attributes. These attributes are defined as belonging to custom object classes that are structural under *inetOrgPerson*. Each domain is to have one initial custom object class; based on the naming conventions defined in Section 3.1.3 below, these would be *copedEmplPerson*, *copedBusPartPerson*, and *copedSubscriberPerson*. If additional custom object classes are needed, the governance process defined below in Section 3.1.1 also provides the ability to create them. Custom object classes would group new custom attributes defined for a specific purpose that are not applicable to all identities within a given domain. For example, the custom object class *copedPIVcardPerson* might be created to hold custom attributes related to a PIV card that do not apply to all business partner identities.

In general, only a minimum set of attributes are to be included in CoPED. The inclusion of personally identifiable attributes is to be reviewed for compliance with existing laws, regulations, and standards governing their use. This limitation, however, only means that such attributes (if belonging to a standard object class) are not to be populated with data. The standard *inetOrgPerson* attributes from OPD-SEC014F - *CoPED Schema,* or any of its superiors <u>are not to be removed</u>.

3.1.1 Attribute Selection Policy

This section defines the governance model for managing the schema. In Section 7, *Governance and Administration*, governance for CoPED rests primarily with the Enterprise Architecture group. Governance includes setting and modifying policy as needed, ruling on exceptions to those policies when appropriate, and policy enforcement.

Schema governance flows from the strategic directory usage outlined in Section *2.1 Strategic Purpose*, which states that CoPED serves as the repository for users' authentication credentials for the Shared Authentication Service, and stores those additional information attributes that are shared across many agencies. The next guiding principle is legal compliance, especially to meet the privacy requirements of existing laws, regulations, and standards governing the use of personally identifiable attributes.

Additional policies may be added by Enterprise Architecture as circumstances dictate. Policies do not need to be created or modified to encompass every requested exception; an exception will be approved and implemented as an exception without changing the underlying policy.

OPD-SEC014F - *CoPED Schema,* lists the core set of inetOrgPerson attributes included in CoPED. All custom attributes are to be documented in a similarly structured table, providing name, description, data type and format, the primary object class to which the attribute belongs (could be more than one), indication of whether the attribute is mandatory for that object class, and whether it is single or multi-valued. Additional information if found useful may be included, such as indexing, sample syntax, data for complex attributes, and access permissions if different from the object as a whole.

3.1.2 Object Identifiers

To simplify data sharing and ensure greater interoperability with external systems, CoPED is registered with the Internet Assigned Numbers Authority (IANA) for a universally unique Private Enterprise Number (PEN). The IANA PEN is recognized internationally as the de *facto* standard for registered enterprises. The PEN is a dot-delimited numbering scheme, which not only uniquely identifies the Commonwealth of Pennsylvania Active Directory forest (CWOPA), but is also used to uniquely identify a wide variety of custom objects within the Enterprise, including the directory attributes and object classes described in Section *3.1.1, Attribute Selection Policy*. This is accomplished by adding additional dot-delimited values to the chain. A complete chain of values for any given object is referred to as an Object Identifier (OID). The OID conforms to a strict hierarchical pattern where each value in the chain is a sub-component of the preceding value. For the most part, the objects defined will be specific to the various agencies and business units using them, with governance provided by Enterprise Shared Services. The commonwealth has been assigned the PEN 1.3.6.1.4.1.29245 as its base identifier.

The CWOPA OIDs are associated with the custom attributes and object classes in the directory schema definition. Each CWOPA OID can be described as 1.3.6.1.4.1.29245**. A.B.C.D** where:

1.3.6.1.4.1.29245 is the unique IANA PEN for the CWOPA.

A represents a top-level technology classification. The following values are reserved:

- 1 = SNMP MIB
- 2 = Directory Services

The value 1 is reserved for Network Management using SNMP and 2 is reserved for Directory Services, which is the subject of this document. Additional values are assigned as needed for other technologies.

The definition of the remaining letters is specific to the selected technology classification; in other words, each subsequent value in the OID chain to the right of **A** is specific to the

technology class defined by **A.** The values described below are all specific to Directory Services.

B represents the source agency or business unit such as office, department, or CoP. Since all subsequent values in the OID chain pertain to custom objects developed specifically for a particular business unit, it follows that the organization is to be listed in the OID chain. The following values have been reserved:

- 1 through 9999 = Business Area number. Agencies are to use the Business Area number assigned for use by their fiscal offices.
- 10001 = Enterprise Shared Services

C represents the artifact type; for Directory Services these are schema element types. The following values have been reserved:

- 1 = Directory Custom Object Class
- 2 = Directory Custom Attribute

D represents the unique number for each schema element. For example, if the attribute *copedGUID* were assigned the **D** value of 15, then its complete OID would be 1.3.6.1.4.1.29245.2.10001.2.15 where:

1.3.6.1.4.1.29245 - is the base identifier obtained from IANA
2 - represents Directory Services
10001 - represents "Enterprise Shared Services"
2 - indicates a Directory Custom Attribute
15 - specifies the attribute *copedGUID*

3.1.3 Naming Conventions

Any CoPED custom attributes or object classes added are to be named in accordance with the following conventions. All such custom names are to begin with the prefix *coped* (all lower case) to indicate they are added for CoPED. They are to have as-meaningful-as-possible names appended, with the first letter of each word capitalized and with all words running together. For example, the custom object classes for the three domains will be: *copedEmplPerson, copedBusPartPerson*, and *copedSubscriberPerson*. As seen, structural object classes are to also give an indication of their superiors.

Section 5.1 details the policy surrounding the selection of the various authoritative sources.

3.2 CoPED GUID

Each entity in CoPED is to have one Global Unique Identifier (GUID) for each domain to which he or she belongs. These GUIDs are unique across all domains and directories; in other words, each GUID is specific to a single domain and is not to be used as an identifier in any other domain or directory within the Commonwealth enterprise, even if for the same person. The algorithms used for generating CoPED GUIDs are described below in *Section 4.1.2, Globally Unique ID*. The primary uses of the CoPED GUID are as follows:

- **CoPED RDN** – For the overall CoPED virtual directory interface, the CoPED GUID will form the identity's Relative Distinguished Name (RDN), which is the identity's unique ID within the bottom-most container. Since CoPED GUID will always be unique across all three domains, it will also always be unique within any single container. Also, since it will always be invariable (no matter what other changes an identity undergoes), it makes an excellent choice as the base key for the directory.

- **Business Partner and Subscriber RDN** – For the same reasons as above for using CoPED GUID for the overall CoPED RDN, the CoPED GUID will also be used as the RDN in the new directories for both the Business Partner and Subscriber domains.

- **Synchronization Link** – The CoPED GUID will serve as the unique ID that is keyed for connecting CoPED with each of the agency and CoP identity stores and master directories. This is discussed thoroughly in Section *4.1.2, Globally Unique ID*.

3.3 CoPED Relationship to PKI Model

GEN-SEC013G - *Public Key Infrastructure,* and the *X.509 Certificate Policy* (CP) *for the Commonwealth of Pennsylvania Enterprise Public Key Infrastructure,* define in detail the commonwealth's PKI model. In addition to what is supplied by those documents, this section defines the relationship and interactions of that PKI model with CoPED.

The primary objective of the commonwealth PKI is to issue X.509 digital certificates to users as needed for: user authentication; encryption; and digital signature. The commonwealth's federally certified Shared Service Provider (SSP), as named in STD-SEC014C - *Product Standards for Public Key Infrastructure/Shared Service Provider,* ensures *FIPS 201* and Federal Bridge compliance. Upon user registration, the SSP will issue an X.509 certificate containing the user's public key and will store it in CoPED, in the *userCertificate* attribute included in the *inetOrgPerson* object class for that user. This certificate is used by various applications as appropriate:

- By the Shared Authentication Service for authentication.
- By the e-mail client (Microsoft Outlook or equivalent) for encrypting or verifying digitally signed e-mail.

The SSP serves as the primary Certificate Authority (CA) for the Commonwealth, to issue certificates and manage the Certificate Revocation List, as explained thoroughly in GEN-SEC013G - *Public Key Infrastructure.* CoPED stores a copy of the *certificateRevocationList* of the certificates the SSP has revoked, via an attribute in the *certificationAuthority* object class. The SSP Certification Practice Statement (CPS) defines how (and whether) users' revoked certificates are disabled.

3.3.1 X.509 Certificate Storage and Retrieval

As noted above, the SSP stores the X.509 certificate containing the user's public key in CoPED, in the *userCertificate* attribute (part of the *inetOrgPerson* object class), at the same time that the SSP delivers the user's private key to the user (generally stored in the

browser or the user's PIV card, and protected by PIN or biometric). The X.509 certificate is retrieved by various applications as needed, via LDAP v3 access. The three primary certificate uses are for encryption of data to be sent to (or used by) the user, validation of a digital signature from this user, and user authentication. These are described below:

1. **Encryption** – To encrypt data or e-mail messages intended to be read by a given user, the application will retrieve that user's certificate from CoPED via LDAP and use the public key to encrypt the data (or message). When the user later attempts to read the data (or message), the user's application will retrieve the user's private key stored in the Windows certificate store on the user's system or in the user's PIV card, confirm the user's identity via the PIN or biometric protection, use the private key to decrypt the data (or message), and return it to the user in clear text.

2. **Digital Signature** – To digitally sign an e-mail message or other document, the application will perform a hash function against the message to create a relatively short string called a message digest. The application will encrypt the digest with the user's private key (the user will be required to enter the PIN or provide the biometric protecting the private key), and store the resulting string as the signature. The receiver's application will retrieve the user's certificate from CoPED via LDAP and use the public key to decrypt the signature, and then apply the same hash function to the original message. If the decrypted signature matches the locally computed hash, it indicates that the signature is authentic and that the "signed" document has not been altered.

3. **Authentication** – The "user" interactions described in this paragraph are typically performed by the user's browser with no user interaction required. When an application (such as the Shared Authentication Service described in GEN-SEC013C - *Access Management and Control*) wants to perform certificate authentication for this user, it will generally send the user a challenge string. The browser will digitally sign the challenge as described above: it will encrypt the challenge string with the user's private key (if protected, the user would be required to enter the PIN or provide the biometric) and return the resulting string. The application will retrieve the user's certificate from CoPED via LDAP and use the public key to decrypt the string. If the decrypted string matches the challenge, the user is authenticated.

3.3.2 User Naming

The X.509 certificates include the user's distinguished name (DN) as the subject of the certificate. The SSP is to retrieve the user's DN from CoPED when issuing a certificate, and store it as the subject. The user's DN is formed from the branch of the Directory Information Tree (DIT) containing the user's entry, and either the user's Common Name (cn) for *Employee* or CoPED GUID for *Business Partner* and *Subscriber*. See BPD-SEC013H, *Directory Services Implementation Guide,* for details on user identification best practices within the directory.

4 Synchronization Services

When metadirectory products first appeared, they actually provided a "meta" directory that is a single directory that consisted of a copy of all the data in the various lower-level directories and other identity stores. The value of the metadirectory was derived from the power of its synchronization engine. As the market matured, however, the vendors realized that while the synchronization services of their products were useful, most large enterprises found their aggregated identity data stores lacking when compared to industrial-strength directory server products. Although they have kept the name, current metadirectory products now consist almost entirely of synchronization-related services, focused on publishing aggregations of data into an enterprise directory server like CoPED.

The IPAM initiative defines the deployment of a metadirectory service to provide identity data synchronization at the enterprise level for commonwealth agencies. This section details the architecture of that service, as shown in *Figure 1 – Identity Data Synchronization Architecture*.

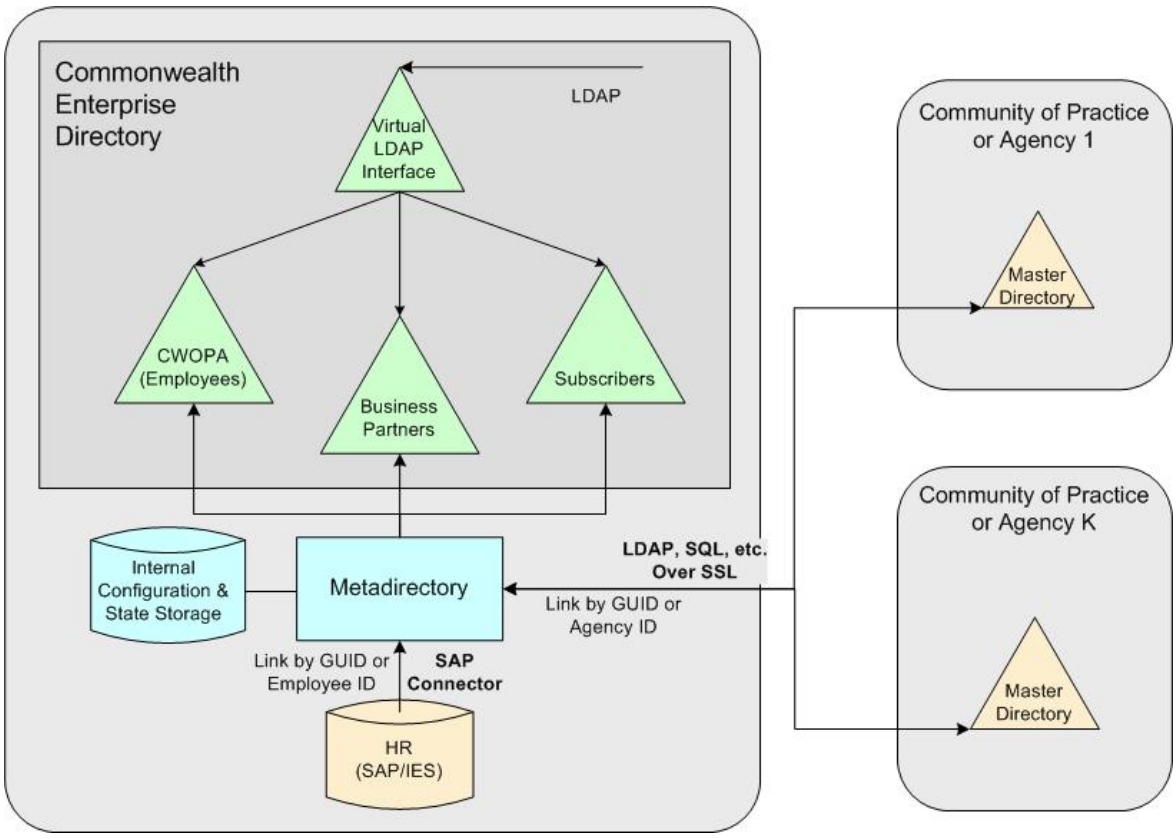

Figure 1 – Identity Data Synchronization Architecture

This architecture is based on:

- The metadirectory service (shown in Figure 1) reading employee data from Human Resources (HR) Department (from the SAP/IES HR system) and populating CWOPA (only for employees at this time);
- *Subscriber* and *Business Partner* data being read from the various agency and CoP master directories; and,
- The population of the *Business Partner* and *Subscriber* directory domains in CoPED

As shown in Figure 1, the metadirectory usually contains an internal data store for holding configuration information and its current view of the state of all aggregated identities. This internal store is not normally accessible as an identity store. such as by LDAP, but is where the metadirectory first writes any identity object updates before publishing them to the various connected identity stores, including CoPED and agency/CoP master directories.

This section describes the identity synchronization model while discussing:

- Use of a centralized metadirectory to aggregate identity data
- Requirements for and use of a globally unique identifier (the CoPED GUID) for each identity to link identity objects between identity stores
- Type and structure of connectors between the metadirectory and the connected identity stores

4.1.1 Central and Distributed Metadirectories

The full commonwealth synchronization architecture includes both centralized and distributed metadirectories to aggregate identity data. As shown in *Figure 1 – Identity Data Synchronization Architecture*, the centralized metadirectory server described in this document is used to aggregate the agencies' and the CoPs local identity data from their master directories into CoPED (as well as writing a defined subset of that data back into the connected agency and CoP master directories). In addition to that metadirectory, however, many agencies and CoPs leverage separate instances of the same metadirectory products to aggregate their own distributed identity data into their own master directory. This latter use of the metadirectory is not detailed further in this document.

4.1.2 Globally Unique ID

To synchronize identity data between identity stores, the metadirectory is to be able to link an identity in one store with the same identity in the other store. For the IPAM architecture, the metadirectory creates a globally unique identifier called the CoPED GUID, which is assigned to an identity in CoPED when that identity object is first created. It will never change or be deleted as long as that identity record exists.

The metadirectory service stores the CoPED GUID in CoPED when it creates a new identity object. It also writes the CoPED GUID back to each of the connected identity stores that links to this identity, assuming the agency or CoP controlling that identity store allows the

write-back. The metadirectory also stores in CoPED the identifier or key from the agency/CoP master directory for this identity *(AgencyIDs)* to provide an additional link (the primary or only link when the agency doesn't allow write-back) between the agency/CoP identity object and CoPED.

The CoPED GUID is structured by domain in the following format:

1. Initial (prefix) single character indicating the identity's domain (E, B, or S for Employee, Business Partner or Subscriber)

2. An arbitrary twelve-digit number assigned and managed by the metadirectory. This number is unique across all domains, to ease any future domain reconfiguration.

3. Repeat the single prefix character as a single suffix character (to better support those applications that prefer a suffix to a prefix).

For example, an employee might be assigned the CoPED GUID "E100000000001E" (without the quotes). A subscriber might then get "S100000000002S" but would never get "S10000000001S" since that number was already used.

When a user authenticates to the Shared Authentication Service described in the supporting document GEN-SEC013C - *Access Management and Control*, that service informs the application being accessed (using the Security Assertion Markup Language [SAML] federation standard, version 2.0) that the user has been authenticated. The service sends a SAML assertion to the application, providing the CoPED GUID as the user's main identifier in the SAML authentication statement. The assertion also contains attribute statements that provide additional user data, including local agency/CoP IDs (such as Employee ID, Citizen ID – either a state driver's license number or non-driver number – or others as appropriate) that applications can use to link to the local identity store for additional user information. Each such attribute statement is null if the user does not have a corresponding value for that ID in CoPED.

4.2 Metadirectory Architecture

This section describes the architecture for the metadirectory.

4.2.1 Physical Topology

The metadirectory resides in the data layer of an n-tier architecture, co-located on a high-speed backbone with those protected commonwealth data stores with which it will communicate, including:

- Underlying directory servers of CoPED
- A CWOPA Active Directory Domain Controller
- HR SAP (IES) system
- Agency and CoP identity stores, where possible

Agency and CoP identity stores are not required to be co-located in the data layer with CoPED.

4.2.2 Disaster Recovery

Even though the metadirectory and its processes have no time-critical responses, they are mission-critical processes. This infrastructure is to therefore, adhere to existing commonwealth protocols and policy for mission-critical data systems. Data to be backed up includes:

- Metadirectory server configuration
- Metadirectory server internal data store

4.2.3 Agent-less Remote Connectors

Some metadirectory products connect to remote identity stores by installing an agent into the identity store product. Although this can sometimes provide some security or performance benefits, the IPAM synchronization does not use the agents to connect to the remote identity stores, but instead works with industry standard access protocols, primarily LDAP and SQL, as well as typical proprietary protocols of the connected identity stores (primarily ADSI for Microsoft Active Directory).

4.3 Identity Synchronization Scenarios

This section presents three key scenarios that illustrate the identity synchronization process. The first describes the process for when a new user is added as an identity in an agency's master directory after having been vetted by that agency's identity proofing process when that user does not yet have an identity in the CoPED domain of interest (Business Partner or Subscriber). The second scenario is similar to the first, except that the user already has an identity in the CoPED domain of interest; this scenario describes the procedure for linking agency identities to existing CoPED identities. The final scenario describes the process followed when a synchronized attribute has its value changed.

In the first two cases, the metadirectory assigns a proofing level value to the user based on the value assigned to that agency's vetting process. It also assigns a unique code identifying the approving agency. These vetting processes and their proofing levels are explained in the supporting document GEN-SEC013D - *Enrollment, Identity Proofing and Vetting*. The vetting process is Web-based and user interaction with the Web interface is required.

4.3.1 New Identity Vetted at an Agency

The first scenario describes the process used when a user without an identity in CoPED is added to an agency's master directory after having been vetted under that agency's identity proofing process. This process begins with determining that the user is in fact new, and proceeds through the various metadirectory functions to synchronize that identity with CoPED as required. The process is shown in Figure 2, New Identity Scenario Flow. The steps in the figure correspond to the numbered paragraphs.

1. At the end of the agency vetting process, the user is at some Web page, either the last page in the process or a screen informing the user of the completion of the off-

line vetting process. That screen is to inform the user that the next step is to procure login credentials. This step is not shown in Figure 2.

2. The user selects the link to procure login credentials, and is taken to an application that will first look to see if the user exists.

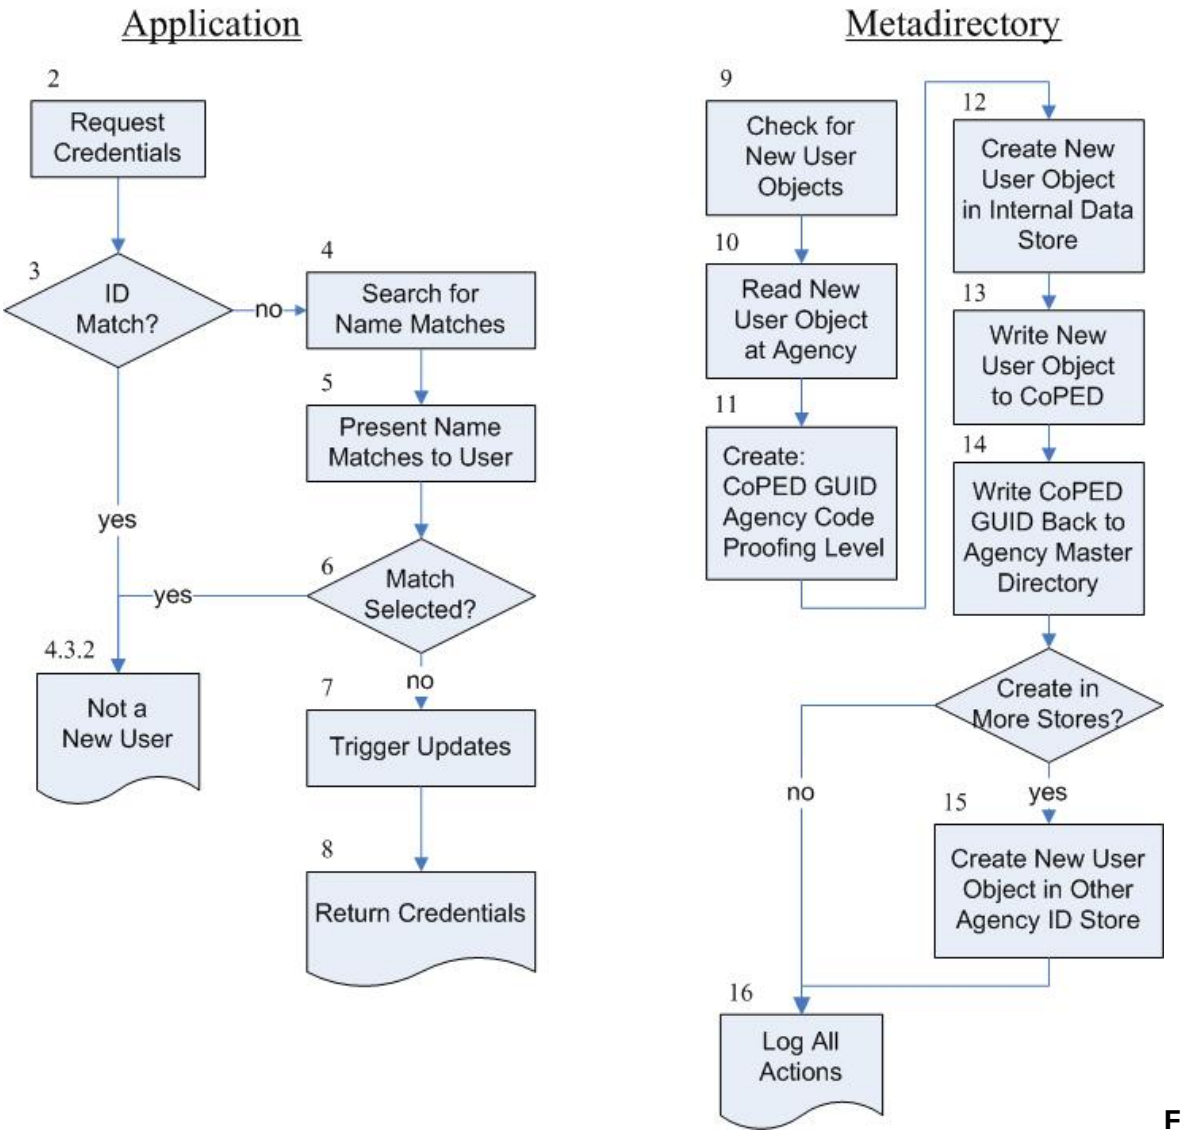

Figure 2 – New Identity Scenario Flow

3. This application looks for potential matches in the domain of interest (Business Partner or Subscriber) using any agency-internal ID that is stored in CoPED. This scenario assumes that none is found; continue with the next step. If there is a

match, that would constitute the second scenario; see Section 4.3.2, *Existing CoPED User Added at an Agency*.

4. If there is no ID match (or no ID exists to search on), the application then looks at the user's last name (surname) and searches for potential matches or close calls (Soundex or other approximate matching functions may be appropriate) in that domain. The agency can customize this application to allow users to narrow or expand the search criteria to find themselves among the existing users.

5. The application presents the user with a choice of the potential matches.

6. If the user chooses a match, that would constitute the second scenario; see Section 4.3.2, *Existing CoPED User Added at an Agency*. If there is no match, the user selects a "No Match" link indicating that this is a new user for CoPED, and not just for the agency.

7. The selection of "No Match" triggers the metadirectory to discover the new identity object in the agency or CoP master directory on its next update. Depending on which directory product is used, this may be:

    a. A specific activation identifying the particular master directory with the new object (or even the specific object);
    b. A general activation asking the metadirectory to scan all agency/CoP master directories to look for changes;or
    c. Only a trigger in the sense that on the metadirectory's next regular scan it will recognize that an identity object has been added to this master directory.

8. Finally, the application returns to the user the appropriate credentials for future logins.

9. Once the metadirectory is triggered, it checks the agency/CoP master directory for new identity objects. Depending on the product selected, this may be all master directories or only the specific directory that received the addition.

10. The metadirectory finds the new object and reads it, including its attribute values.

11. The metadirectory creates a CoPED GUID for this user and attaches it to the new identity as another value. The metadirectory also determines the agency code and proofing level value corresponding to the vetting process that was used.

12. The metadirectory creates a new identity object in its internal data store, writing the data from the attribute values with the proofing level and agency code, and using the CoPED GUID as its primary search key.

13. The metadirectory creates a new identity object in CoPED, writing the data from the attribute values and using the CoPED GUID as its primary search key.

14. The metadirectory writes the CoPED GUID back to the agency/CoP master directory to maintain the link between the objects in the two directories, if allowed.

15. If there are any other agencies or CoPs whose master directories have been configured to have identity objects created automatically for a set of criteria that this user meets, the metadirectory also creates new identity objects for those agencies and CoPs, and writes all configured and populated attribute values for that identity to those master directories.

16. The metadirectory logs all of its actions.

4.3.2 Existing CoPED User Added at an Agency

The second scenario is similar to the first, except that the user already has an account in CoPED. This scenario describes the process for when a new identity is added to an agency's (or CoP) master directory after having been vetted by that agency's identity proofing process and that identity already exists in CoPED and is found. This process begins with matching the identity, and proceeds through the various metadirectory functions to synchronize that identities attributes as required. This process is shown below in *Figure 3 – Existing Identity Scenario Flow*. The steps in this figure correspond to the numbered paragraphs below.

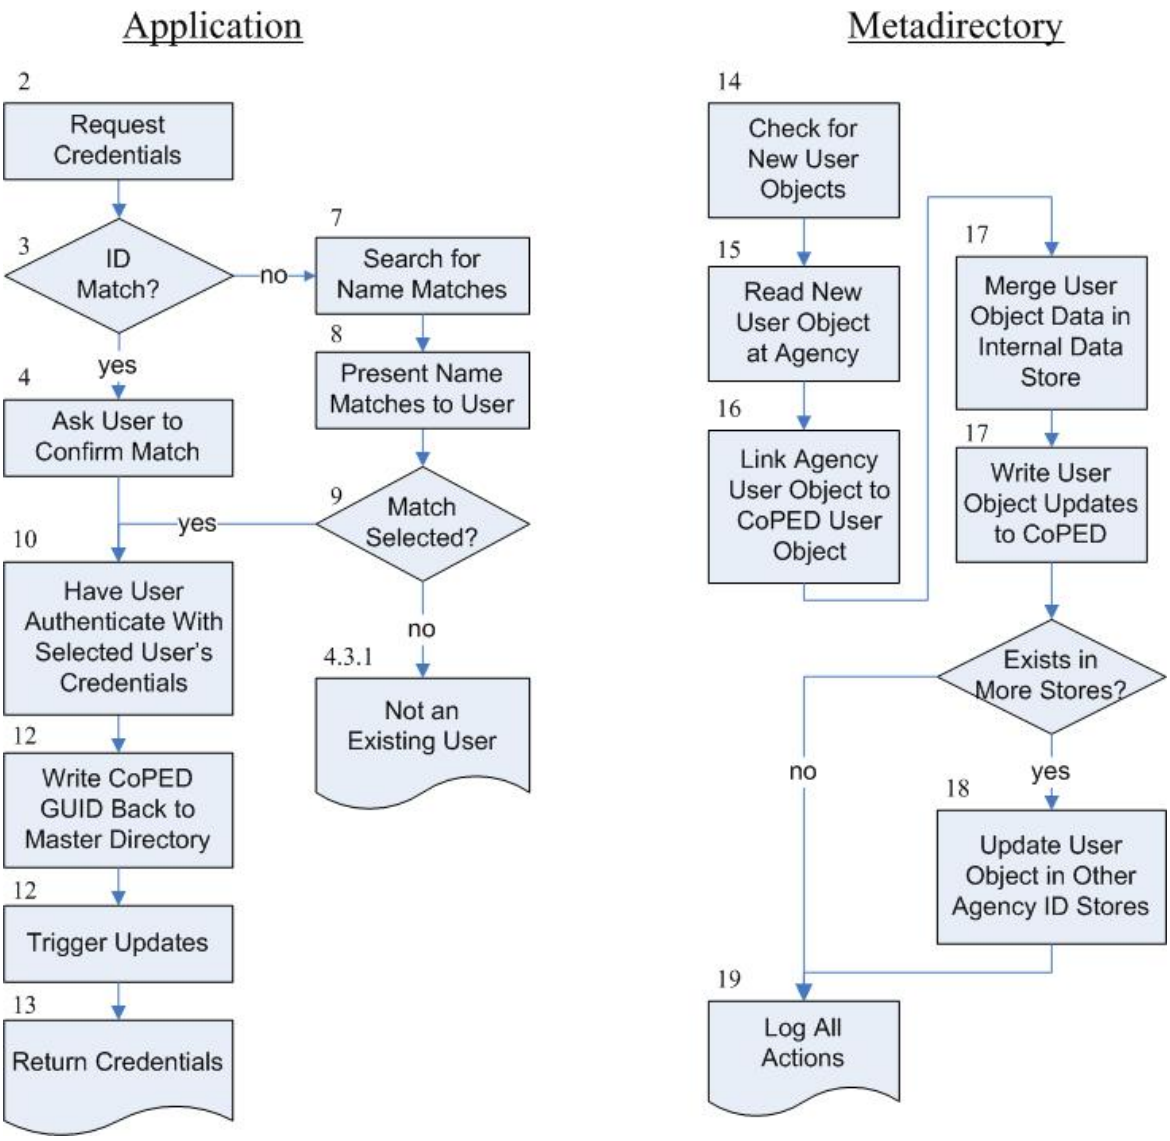

Figure 3 – Existing Identity Scenario Flow

1. At the end of the agency/CoP vetting process, the user is at some Web page, either the last page in the process or a screen informing the user of the completion of the off-line vetting process. That screen is to inform the user that the next step is to procure login credentials. This step is not shown in Figure 3.

2. The user selects the link to procure login credentials, and is taken to an application that first looks to see if the user exists as a CoPED identity.

3. This application looks for potential matches in the domain of interest (Business Partner or Subscriber) using any agency-internal ID that is stored in CoPED. If none is found, continue with Step 7.

4. If a unique ID match is found, the user is asked to confirm the matching identity.

5. If confirmed, continue with Step 10.

6. If unconfirmed, the user is told to call a help number for manual intervention. The status is logged.

7. If there is no ID match (or no ID exists to search on), the application then looks at the user's last name (surname) and searches for potential matches or close calls (Soundex or other approximate matching functions may be appropriate) in the domain of interest. The agency can customize this application to allow users to narrow or expand the search criteria to find themselves among the existing identities.

8. The application presents the user with a choice of the potential matches.

9. If the user chooses no match, that would constitute the first scenario; see Section 4.3.1. The user selects a link identifying the match.

10. The user is asked to authenticate with the discovered identity's credentials.

11. If authentication is unsuccessful, the user is told to call a help number for manual intervention. The status is logged.

12. If authentication is successful, the application retrieves and stores the CoPED GUID in the agency's or CoP master directory and triggers the metadirectory for update.

13. Finally, the application returns to the user the appropriate credentials for future logins.

14. Once the metadirectory is triggered, it checks the agency/CoP master directory for new identity objects.

15. The metadirectory finds the new object and reads it, including its attribute values which now includes the CoPED GUID written by the application.

16. Seeing the CoPED GUID, the metadirectory links this new object to the existing CoPED identity object.

17. The metadirectory merges the data from the new object's attribute values with the existing CoPED object. The metadirectory also writes to the CoPED identity object an updated agency code and proofing level value corresponding to the vetting process

that was used. Proofing levels are thoroughly explained in SEC013D, *Enrollment, Identity Proofing and Vetting.*

18. The metadirectory updates the new attributes to other connected directories as appropriate; see the following Section 4.3.3.

19. The metadirectory logs all of its actions.

4.3.3 Existing CoPED Identity Attribute Change

The final scenario describes the process followed when a synchronized attribute has its value changed, whether in CoPED directly or in one of the agency or CoP master directories. This process begins with changing the value of an attribute for a linked user in one of the connected directories. The process is shown in Figure 4.

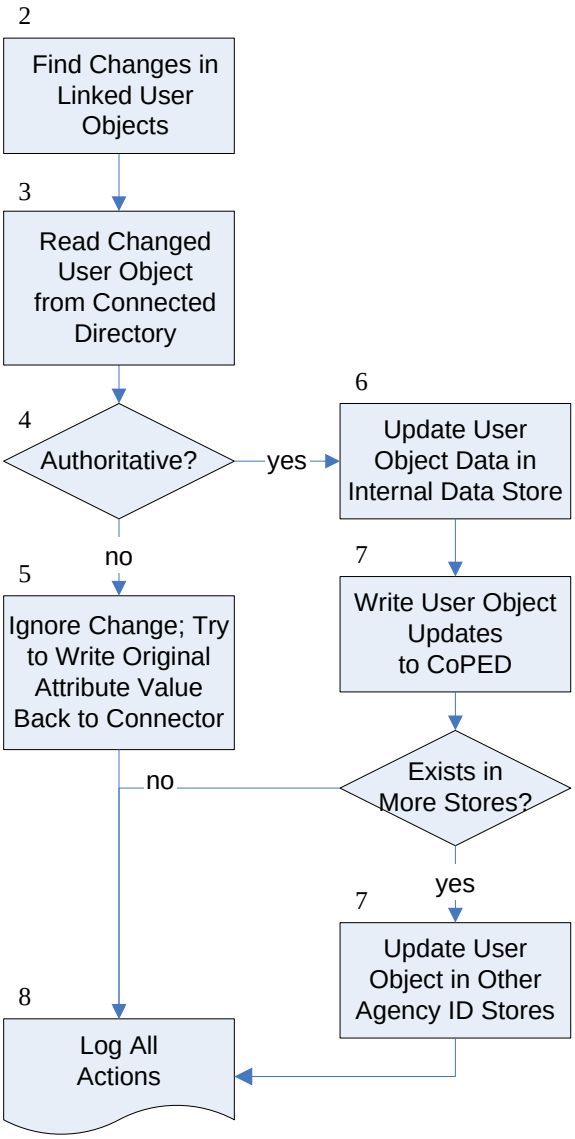

Figure 4 – User Attribute Change Scenario Flow

1. When an attribute value changes in a linked identity object in a connected directory, the metadirectory is triggered for update. Connected directory includes both CoPED and the various agency and CoP master directories. As before, the meaning of trigger depends greatly on the product used. It may be a specific activation identifying the particular master directory with the changed object (or even the specific object or attribute), a general activation asking the metadirectory to scan all agency/CoP master directories to look for changes, or only a trigger in the sense that on the metadirectory's next regular scan it recognizes that an object has an attribute that has changed in this master directory. This step is not shown in Figure 4.

2. Once the metadirectory is triggered, it checks the connected directories for changes in linked identity objects.

3. The metadirectory finds the change in the linked identity object and reads it, including all of its linked attribute values.

4. The metadirectory determines whether the change was read from its authoritative source for making this change to this attribute; see Section *5.1.5 Modify Attribute,* for details on authoritative sources for attribute modifications.

5. If not authoritative, the metadirectory ignores the change and attempts to write the original value back to the changed attribute.

Note: The agency or CoP may opt to disallow updates to its master directory from the metadirectory.

6. If authoritative, the metadirectory updates its internal data store, writing the new value to the attribute and marking it for update.

7. The metadirectory writes the attribute's new value to all connected directories that are accepting writes of that attribute, including CoPED and any agency or CoP master directory.

8. The metadirectory logs all of its actions.

4.4 Sharing Agency Data

The identity data synchronization service described in this policy involves sharing identity data between an agency and OA/OIT, as well as potentially sharing that data among the agencies themselves. Agencies may also leverage commonwealth directory technologies to share data with approved business partners, other states and federal organizations, but these will be on a case-by-case basis and is beyond the scope of this documentation. Each agency and CoP is to be aware of commonwealth data privacy standards for any and all data that it allows to be synchronized. Section 7, *Governance and Administration* , discusses the governance model used to enable agencies to manage the release of any such data.

Different directories often contain conflicting identity information about the same person. In addition, the agency or CoP that owns and manages the data in a specific connected identity store typically believes that its data is authoritative when compared to similar data that resides in a different connected identity store. In these cases, data owners are often reluctant to relinquish control of their identity data.

To resolve problems that result from conflicting identity information, rules are established in the metadirectory to determine attribute flow precedence for each specific attribute of the identity objects. The metadirectory then updates the other connected identity stores with

that authoritative value. Attribute precedence and authoritative sources are detailed in Section *5.1 Authoritative Sources*.

4.5 Extraction, Transformation and Loading

One of the greatest challenges for most metadirectory projects is the extraction, transformation and loading (ETL) of existing identity data into the new directory, especially when there is not already a widely disseminated globally unique identifier to link the various identity accounts. CoPED largely sidestepped this by identifying only a small subset of identity data for its initial rollout.

Most uses of employee identity data in the commonwealth already leverage the existing employee ID number. Since CoPED uses its virtual directory technology to incorporate CWOPA directly, keyed on the Employee ID, no additional work is required to ensure consistency for employees. See Section 4.6 for details.

For the remaining, non-employee identities (*Business Partner* and *Subscriber*), CoPED does not bulk load the identity data from the existing agency directories. Rather, as described in the scenarios listed previously in Section 4.3, identities are entered into the new CoPED identity store when they first access the Shared Authentication Service, at which time they are linked to one or more appropriate agency/CoP identities.

4.6 Transitional Model for Synchronization Services

The IPAM Identity Data Synchronization service, as with most IT initiatives, is deployed in phases to provide a smooth transition to the architecture described previously in *Figure 1 – Identity Data Synchronization Architecture*. See Section 5, *Transitional Model,* in BPD-SEC013H for a detailed explanation of the Transitional Model and its phases.

4.7 Standards

The Identity Data Synchronization service follows all standards listed in the appropriate commonwealth ITPs. In particular, all access to CoPED (to publish users or attribute modification or gathering) uses the LDAP v3 protocol. Access to the connected agency master directories use the standard access protocol most appropriate for each store; for example, ADSI for Active Directory, LDAP v3 for other directories, or SQL for databases.

5 Identity Data Management

This section defines the IPAM policy for creating and managing CoPED identity data.

Figure 5 – Authoritative Attribute Sources

5.1 Authoritative Sources

In addition to defining synchronization processes, this policy also describes how to determine the authoritative sources for:

- Creation of each class of user object
- "Deletion" (actually disabling) of each class of user object
- Creation of each attribute, for each class of user object
- Modification of each attribute, for each class of user object
- Deletion of each attribute, for each class of user object

Figure 5 shows an example of a metadirectory combining identity data from several sources, and in particular HR/IES and PDE. In this example, HR/IES is the authoritative source for the user object and therefore the identifying attribute, *"Name."* It is also authoritative for the attributes *Employee ID* and *Address*, while PDE is authoritative for *e-mail* and *Access* (a role attribute used to define entitlements). The metadirectory creates the *CoPED GUID* and feeds it back to PDE together with *Name*.

5.1.1 Create User Object

Each domain has its own authoritative source to create identity objects. The authoritative source for all employee identities is the SAP HR database (IES). No employee identity is to exist in CoPED (or CWOPA) without a record in IES.

Every agency can be the authoritative source for identities in *Business Partner* and *Subscriber*. When any agency (or any Shared Provisioning Service provided by Shared Services) vets a user through its defined proofing process, it creates an identity account for that user in its master directory, which is then discovered by the metadirectory and created in CoPED with a Proofing Level correlating to that particular proofing process. The proofing agency code is also stored as an attribute in CoPED.

In addition, agencies may delegate the ability to create and manage various identity types (especially business partners) to external entities whose defined processes for provisioning and managing have been approved by the appropriate commonwealth authorities. These delegation rights are further explained in SEC013D, *Enrollment, Identity Proofing and Vetting.* These processes also have set proofing levels defined.

See Section 4.3.1, *New Identity Vetted at an Agency*, for the scenario describing the creation of new users.

5.1.2 Link User Object

In addition to creating new identities in CoPED, an agency may vet a user who already exists in one of the other CoPED domains (*Employee, Business Partner* or *Subscriber*). Once the sponsoring agency has verified the existence of the CoPED identity, the agency can simplify its data entry task by linking the user to the corresponding CoPED GUID. When the user object is discovered by the metadirectory, it can automatically update those attributes that the agency is authoritative for (see Sections 5.1.4, *Create Attribute*, and 5.1.5*, Modify Attribute*, on attribute management).

In particular, it updates the agency code stored for the user to include its own agency identifier, indicating that that agency also has vetted the user. If its assigned Proofing Level is now higher than the user's currently stored level, CoPED updates that value to the higher level.

As with new users, agencies are permitted to delegate the ability to create and manage sets of users (especially business partners) to external entities who have defined processes for provisioning and managing those users. These processes also have set Proofing Levels defined. Regardless of which administrator manages the user account, once the account is created in the agency master directory the processing is the same.

5.1.3 Disable and Archive User Object

Identities are not to be deleted from CoPED. If there is a security reason for denying user access (such as intrusion attempt or fraud), the user's account is to be disabled to remove all authorization that the account had provided. Any appropriate security tool or process

can set a user's account to be disabled, and appropriate administrative processes are to be followed to re-enable accounts when needed.

If a decision is made to completely remove an identity from CoPED, the identity account is to be disabled and archived first. Archiving in CoPED can be arranged by setting an attribute *(copedArchive),* by moving the identity to an Archive branch of the directory, or by other means as appropriate. For employees, removal and archiving is also determined by IES, or by the tools and process used to manage CWOPA.

Archiving ensures that deactivated users can be referenced when tracing historical information, especially audit trails and other forensics. By maintaining the archived users within each domain's directory, the directory server's internal functionality for forcing ID uniqueness can be leveraged to ensure that CoPED GUIDs assigned to that directory remain unique forever, and are never reassigned.

For users in *Business Partner* and *Subscriber* domains, an automated process periodically traverses CoPED to find inactive users. These are defined as users who have not authenticated for some predefined period of time. Although adjustable through the governance process (see Section 7, *Governance and Administration*), this "period of time" is initially defined as three (3) years for *Business Partners* and seven (7) years for *Subscribers*.

As with all other attribute changes in CoPED, disable and archive events change attribute values and are discovered by the metadirectory and sent to agencies' connected master directories. The agencies determine whether and how to process these events in their local identity stores.

5.1.4 Create Attribute

IES and the existing CWOPA administration tools continue to be authoritative for creating attribute values for employee users.

As with creating user objects, each agency can create new attribute values for any existing user in *Business Partner* and *Subscriber*. The governance process described in Section 5 discusses whether such new attributes are actually read by the metadirectory and written into CoPED.

5.1.5 Modify Attribute

IES and the existing CWOPA administration tools continue to be authoritative for modifying attribute values for employee users.

Unlike creating user objects or even creating values for attributes which are presently null, when modifying existing values for *Business Partner* or *Subscriber* attributes it is possible that valid data may be overwritten with less valid data. The definition of the authoritative source and an enforcement mechanism work to minimize that risk. In

general, any agency may change most attributes, especially common attributes such as contact information. Local IDs (such as PennDOT ID) are changed only by their owning agency. Authentication credentials may only be changed by one of the agencies that originally vetted the user.

When additional attributes are added through the governance process of CoPED, the parties involved in requesting the new attributes work with Shared Services to determine which agencies are allowed to change the attribute's value once it has been set.

5.1.6 Delete Attribute

IES and the existing CWOPA administration tools continue to be authoritative for deleting attribute values for employee users.

In general, any agency may change most attributes, especially common attributes such as contact information. Local IDs (such as PennDOT ID) are changed only by their owning agency. Authentication credentials may only be changed by one of the agencies that originally vetted the user.

When deleting existing values for *Business Partner* or *Subscriber* attributes, it is possible that valid data may be lost. The definition of the authoritative source and an enforcement mechanism work to minimize that risk. In general, agencies or CoPs are allowed to delete few if any attributes; most attributes cannot be deleted, including common attributes (such as contact information), local IDs (such as PennDOT ID), or authentication credentials. They will be maintained in CoPED for historical purposes.

When additional attributes are added through the governance process of CoPED, the parties involved in requesting the new attributes work with Shared Services to determine which agencies (if any) are allowed to delete the attribute's value once it has been set.

5.2 Directionality

Even after the authoritative sources are determined, there is another policy issue: whether the identity data collected by the metadirectory should be written back to the various agencies identity stores in addition to writing to CoPED. To provide the greatest value of this shared synchronization service, the metadirectory provides bi-directional synchronization, allowing data updates back to the agency master directories.

During the detailed design and implementation of their connection to the metadirectory, each agency and CoP is able to define which attributes (if any) accept write-back from the metadirectory. The agencies and CoPs are also able to specify whether creation of new user objects by other agencies or CoPs are to trigger the metadirectory to create those objects in their local master directories, or only write to attributes in user objects that already exist locally. This is described in greater detail in BPD-SEC013H - *Directory Services Implementation Guide.*

6 Security

The security for any government directory warrants careful examination. This is especially true for CoPED because it is a core component of commonwealth access control systems. The CoPED system configuration and encryption are to adhere to existing commonwealth protocols and policy for securing mission critical data systems and for physical access security restrictions. This section discusses the policy governing these issues.

6.1 Administrative Access Control

Access by administrators to the CoPED virtual directory server configuration, the *Business Partner* and *Subscriber* directory server configuration and data, and CWOPA is strictly limited. One of the guiding principles for attribute selection is legal compliance, especially meeting the privacy requirements of existing laws, regulations, and standards governing the use of personally identifiable attributes. Therefore, storage of private data is limited, although not eliminated. Limiting access is also important to ensure the integrity of the identity data contained in CoPED.

CoPED is located in the Data Tier of the network infrastructure, with firewall protection from casual access (both Internet and internal). Once accessed, the internal administrative access control policy is rigorously enforced in accordance with *FIPS 201* and *NIST SP 800-73* stipulations, federal and commonwealth legislation pertaining to the storage and handling of private information, and existing commonwealth network and information security standards and protocols.

In general, all access to CoPED is to be programmatic by approved commonwealth (including agency) applications. Administrative access directly to the CoPED virtual directory server configuration or the *Business Partner* and *Subscriber* directory server configuration is restricted to only appropriate maintenance personnel. Direct administrative access to the *Business Partner* and *Subscriber* directory server data is restricted to an approved data administrators group and protected by appropriate access control lists (ACLs) or equivalent, based on product.

6.2 Data Storage

To ensure the security of the data in place, the following standards are to be followed:

6.2.1 Classification Standards

At the highest level, CoPED data can be classified as data accessible via CoPED and data stored by CoPED. Due to the virtual directory technology and leveraging of CWOPA and local agency identity stores, only a subset of the data accessible via CoPED is actually stored by CoPED, namely that data stored in the *Business Partner* and *Subscriber* directory servers themselves. This is the only data of interest in this section on data storage security.

The stored data can be further classified into private data (also referred to in *FIPS 201-1* as personal information in identifiable form) and public (non-private) data. Some data are

universally private and, if stored, are to be managed as indicated below under Privacy Standards. Some other data are only considered private within some context, either for a specified group of individuals or when used by certain applications.

6.2.2 Privacy Standards

The best method to meet privacy standards is not to store any private data. While CoPED will minimize its use of private data (for example, CoPED will not store SSN), it cannot totally avoid all personally identifiable identity data. Such data is to be stored in compliance with existing commonwealth and relevant national standards.

As noted above, CoPED access is restricted to minimize the release of any stored private data. Private and sensitive data is to be encrypted in place, with an encryption level meeting commonwealth protocols and policy for securing mission critical and sensitive data as well as the 3DES or AES encryption standard.

7 Governance and Administration

The Office of Administration/Office for Information Technology/Bureau of Enterprise Architecture (OA/OIT/BEA), rather than a separate dedicated organization, provides ongoing governance for suggested modifications to the shared CoPED service. This is provided through the review and approval of the Enterprise Architecture Standards Committee (EASC). Once the EASC approves modifications, including new functionality or integrations, they are to be designed and implemented by a designated administrative group. This is to be the same group that performs day-to-day operational oversight and management of CoPED.

As noted in *FIPS-201-1* (version 5, page 17), "To ensure the privacy of applicants, [the commonwealth] shall… [a]ssign an individual to the role of senior [commonwealth] official for privacy…. The individual serving in this role may not assume any other operational role in the PIV system." This senior privacy official is responsible for reviewing all requests for additional attribute data to be stored in CoPED, to ensure that inappropriate private data is not added. This official is also responsible for identifying stored attributes that require encryption or special (ACLs or equivalent) to protect their sensitivity.

The senior privacy official and other identity management roles, along with their associate responsibilities, are described completely in GEN-SEC013F - *Identity Card Production, Personalization and Issuance.*

8. Related ITPs/Other References

- APP-SEC013A - *Identity Protection and Access Management Glossary*
- GEN-SEC013C - *Access Management and Control*
- GEN-SEC013D - *Enrollment, Identity Proofing and Vetting*
- GEN-SEC013F - *Identity Card Production, Personalization and Issuance*
- GEN-SEC013G - *Public Key Infrastructure (PKI)*
- BPD-SEC013H - *Directory Services Implementation Guide*
- STD-SEC014C - *Product Standards for Public Key Infrastructure/Shared Service Provider*
- OPD-SEC014F - *Commonwealth of Pennsylvania Enterprise Directory Schema*

9. Authority

- Executive Order 2011-05, Enterprise Information Technology Governance

10. Publication Version Control

It is the user's responsibility to ensure they have the latest version of this publication. Questions regarding this publication are to be directed to RA-itcentral@pa.gov.

This chart contains a history of this publication's revisions:

Version	Date	Purpose of Revision
Original	9/7/2006	Base Policy
Revision	9/25/2009	Update format
	4/2/2014	ITP Reformat