

*OMG CORBAmed DTF*

## *Person Identification Service (PIDS)*

*(a.k.a. Patient Identification Service)*

### *Adopted Submission*

*2AB*

*Care Data Systems, Inc.*

*CareFlow/Net, Inc.*

*HBO & Company*

*HealthMagic, Inc.*

*HUBlink, Inc.*

*IDX Systems Corporation*

*IONA Technologies PLC*

*Oacis Healthcare Systems*

*Protocol Systems, Inc.*

*Sholink Corporation*

*OMG TC Document corbamed/98-02-29*

*12 February 1998*

- © Copyright 1997-1998 by 2AB
- © Copyright 1997-1998 by Care Data Systems, Inc.
- © Copyright 1997-1998 by CareFlow|Net, Inc.
- © Copyright 1997-1998 by HBO & Company
- © Copyright 1997-1998 by HealthMagic, Inc.
- © Copyright 1997-1998 by HUBlink, Inc.
- © Copyright 1997-1998 by IDX Systems Corporation
- © Copyright 1997-1998 by IONA Technologies PLC
- © Copyright 1997-1998 by Oacis Healthcare Systems
- © Copyright 1997-1998 by Protocol Systems, Inc.
- © Copyright 1997-1998 by Sholink Corporation

The submitting companies listed above have all contributed to this "merged" submission. These companies recognize that this draft joint submission is the joint intellectual property of all the submitters, and may be used by any of them in the future, regardless of whether they ultimately participate in a final joint submission.

The companies listed above hereby grant a royalty-free license to the Object Management Group, Inc. (OMG) for worldwide distribution of this document or any derivative works thereof, so long as the OMG reproduces the copyright notices and the below paragraphs on all distributed copies.

The material in this document is submitted to the OMG for evaluation. Submission of this document does not represent a commitment to implement any portion of this specification in the products of the submitters.

**WHILE THE INFORMATION IN THIS PUBLICATION IS BELIEVED TO BE ACCURATE, THE COMPANIES LISTED ABOVE MAKE NO WARRANTY OF ANY KIND WITH REGARD TO THIS MATERIAL INCLUDING BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.** The companies listed above shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance or use of this material. The information contained in this document is subject to change without notice.

This document contains information which is protected by copyright. All Rights Reserved. Except as otherwise provided herein, no part of this work may be reproduced or used in any form or by any means graphic, electronic, or mechanical, including photocopying, recording, taping, or information storage and retrieval systems without the permission of one of the copyright owners. All copies of this document must include the copyright and other information contained on this page.

The copyright owners grant member companies of the OMG permission to make a limited number of copies of this document (up to fifty copies) for their internal use as part of the OMG evaluation process.

**RESTRICTED RIGHTS LEGEND.** Use, duplication, or disclosure by government is subject to restrictions as set forth in subdivision (c) (1) (ii) of the Right in Technical Data and Computer Software Clause at DFARS 252.227.7013.

CORBA, OMG and Object Request Broker are trademarks of Object Management Group.

# History

<b>Initial Submission Draft 1</b>	<b>20 May 1997</b>
<p>Integrated ideas from many sources. Some of these were the issues documents created by Jon Farmer, the IDL document created by Fred Druseikis, many conference calls with half dozen or so people, and personal discussions with many people.</p> <p>This will be distributed to a wide audience to review and suggest edits.</p>	

<b>Initial Submission Draft 2</b>	<b>30 May 1997</b>
<p>Incorporated major sections for the Authority interface, security discussions and some scenarios. Updated the IDL as per many email discussions among the group. Corrected wording, misspellings, etc. Did NOT update the descriptions to the IDL.</p>	

<b>Initial Submission</b>	<b>30 May 1997</b>
<p>A few IDL descriptions were updated. Removed revision marks.</p>	

<b>Initial Submission Revision 2</b>	<b>17 June 1997</b>
<p>Corrected many typos, grammatical mistakes, etc.</p> <p>Converted all exceptions from upper case to upper/lower format to be consistent with other OMG specs.</p> <p>Converted trait numbers to an enum and added some more traits.</p> <p>Broke Pids module into four modules.</p> <p>Added support for subscription to notifications via CmPidEvents module.</p> <p>Added CmTimeBase modules.</p> <p>Replaced Validation and Permanence concepts with IdState.</p> <p>Made many small edits to the IDL.</p> <p>Corrected the descriptions of the IDL to correspond to the current version of IDL.</p> <p>Added Hublink to the list of submitters.</p> <p>Added British National Health Service to the list of supporters.</p> <p>Rearranged and reorganized many sections of the chapter 3.</p> <p>Added consistent diagramming notation.</p> <p>Solved many of the many open Issues and deleted them.</p> <p>Added diagrams for all major IDL interfaces.</p> <p>Added better defined semantics for matching.</p>	

<b>Initial Submission Revision 3</b>	<b>11 July 1997</b>
<p>Added IDX and Sholink as submitters and added a few more supporters.</p> <p>Replace CmNaming and CmTimeBase modules with PidBase module.</p> <p>Removed IdDomainClerk and put its operations and attribute directly on PersonLookUp.</p> <p>Changed name of CmEvents module and types to use the term 'Change' in place of 'Event'.</p> <p>Removed 'Cm' from the front of all module names.</p> <p>Changed TraitSelector::weight to be a float.</p> <p>Added 'Unknown' IdState.</p> <p>Added a replace_profile() operation to IdDomainMgr and Identity.</p> <p>Added get_profile() and some attributes to Identity.</p> <p>Added CannotDelete exception to delete_temporary_id() in case the service does not allow this.</p> <p>Removed the profile passed into the make_id_permanent() operation.</p> <p>Changed CorrelatingMgr::add_profile() to take a list of profiles.</p> <p>Added the Appendix - Interaction Patterns chapter.</p> <p>Modified the definitions and some type names for the trait concept and how traits are named.</p> <p>Other changes may be found throughout the spec as well.</p>	

<b>Initial Submission Revision 4</b>	<b>19 October 1997</b>
<p>Changed name on cover from Patient Identification Service to Person Identification Service.</p> <p>Made major changes through out the IDL section of the document.</p>	

<b>Submission Revision 5</b>	<b>7 November 1997</b>
<p>Updated with changes from submitters meeting in Atlanta and changes discussed over email.</p>	

<b>Submission Revision 6</b>	<b>13 January 1998</b>
<p>corbamed/98-01-02</p> <p>Final Submission.</p> <p>Started updating with changes discussed and agreed to since last revision. This includes some feedback from the PIDS Evaluation Committee. Not completed due to technical difficulties (ice storm and file system failure). Sent to the OMG any way as the last changes that met the Final Submission deadline.</p>	

<b>Submission Revision 7</b>	<b>15 January 1998</b>
<p>corbamed/98-01-09</p> <p>Updated final Submission.</p> <p>Finished including the updates that were started for Revision 6.</p> <p>Changes since Revision 5 include adding:</p> <ul style="list-style-type: none"> <li>• Conceptual Model</li> <li>• Identification Model</li> <li>• Component and Inheritance diagram</li> <li>• Use Case examples</li> <li>• Security Guidelines</li> <li>• Notification module</li> <li>• Event type descriptions</li> <li>• Trader Service type description</li> </ul> <p>There were reorganizations of parts of the docuemnt as well as minor IDL and semantic description corrections and changes such as:</p> <ul style="list-style-type: none"> <li>• Terminology update</li> <li>• Moved all non-interface type definitions in PersonIdService module to the top of that module.</li> <li>• Changed 'split' to 'unmerge'</li> <li>• Changed 'Unique...' to 'Qulaified...'</li> <li>• Updated Interaction Diagrams</li> </ul>	

<b>Initial Submission Revision 8</b>	<b>6 February 1998</b>
<p>Minor changes with revision marks. Copy to be discussed in Salt Lake City before voting.</p> <p>NOTE: When this is voted in Salt Lake City it is possible other changes could be proposed as well.</p>	

<b>Initial Submission Revision 9</b>	<b>11 February 1998</b>
<p>Minor changes as determined by submitters on Monday 9 February. Contains revision marks since Revision 7.</p>	

<b>Adopted Submission</b>	<b>12 February 1998</b>
<p>This includes all changes that were adopted by CORBAmEd and the Architecture Board on 12 February 1998.</p> <p>Incorporated errata sheets corbamed/98-02-14 and corbamed/98-02-15 as these were part of the CORBAmEd and Architecture Board vote.</p> <p>Updated the copyright statements to include 1998.</p> <p>Changed document number, date, etc.</p> <p>Made minor format changes.</p>	

# Table of Contents

<b>1. Preface</b>	<b>8</b>
1.1 Submission Contact Points	8
1.2 Supporting Organizations	9
1.3 Conventions	10
1.4 Terminology	10
1.5 Proof of Concept	12
1.6 Changes to Adopted OMG Specifications	12
1.7 Response to RFP Requirements	13
<b>2. Overview of Response</b>	<b>15</b>
2.1 Introduction	15
2.2 Problems Being Addressed	16
2.3 Problems Not Being Addressed	16
2.4 Domain Reference Model	18
2.5 Determination of RFP Scope	19
2.6 Design Goals	20
2.7 PIDS Conceptual Model	21
2.8 PIDS Identification Model	21
<b>3. PersonIdentificationService module</b>	<b>24</b>
3.1 Basic Types	26
3.2 IdentificationComponent interface	34
3.3 IdentifyPerson interface	37
3.4 ProfileAccess interface	39
3.5 SequentialAccess interface	40
3.6 IdentityAccess interface	42
3.7 IdMgr interface	44
3.8 CorrelationMgr interface	47
<b>4. NamingAuthority module</b>	<b>48</b>
<b>5. Traits</b>	<b>53</b>
5.1 PersonIdTraits module	54
5.2 HL7Version2_3 module	55
5.3 vCardTraits module	56
<b>6. Naming/Trader Interoperation</b>	<b>58</b>
<b>7. Conformance Classes</b>	<b>61</b>

<b>8. Appendix - Use Case Examples</b>	<b>63</b>
<b>9. Appendix - Interaction Patterns</b>	<b>72</b>
<b>10. Appendix - Event Descriptions</b>	<b>83</b>
<b>11. Appendix - Security Guidelines</b>	<b>87</b>
11.1 Security Requirements .....	87
11.2 CORBA Security Overview .....	88
11.3 Secure Interoperability Concerns .....	88
11.4 Trust Models .....	89
11.5 CORBA Credentials .....	90
11.6 CORBA Security Domain Access Policy .....	91
11.7 Request Content Based Policy .....	92
<b>12. Appendix - Complete IDL</b>	<b>95</b>

# 1. Preface

---

This submission is a response to CORBAmed RFP1, Patient Identification Service (PIDS), Object Management Group (OMG) document number corbamed/96-11-08.

## 1.1 Submission Contact Points

<p>Carol Burt 2AB 110 12<sup>th</sup> Street North Birmingham, AL 35203 1 205 250 6392 cburt@2ab.com</p>	<p>Jon Farmer Care Data Systems, Inc. 3500 Deepwoon Lane Lambertville, MI 48144 1 313 856 8181 jfarmer@caredatasystems.com</p>
<p>V. "Juggy" Jagannathan CareFlow Net, Inc. 235 High Street, Suite 225 Morgantown, WV 26505 1 304 293 7535 juggy@careflow.com</p>	<p>Mary Richards HBO &amp; Company 301 Perimeter Center North Mail Stop GAC3 420 Atlanta, GA 30346 1 770 804 2968 mary.richards@hboc.com</p>
<p>Frederick Druseikis HealthMagic, Inc. 1901 Main Street, Suite 1570 Mail Stop 504 Columbia, SC 29201 1 803 748 9444 x101 fredd@healthmagic.com</p>	<p>James M. Gabler HUBLink, Inc. 100 E. Campus View Blvd. Columbus, OH 43235 1 614 841 9877 x190 jgabler@hublink.com</p>
<p>Mary Rettig IDX Systems Corporation 1400 Shelburne Road Burlington, VT 05402-1070 1 802 864 1758 x6254 rettig@idx.com</p>	<p>Martin Chapman IONA Technologies PLC 8-10 Lower Pembroke Street Dublin 2, Ireland 353 1 662 5255 mchapman@iona.com</p>
<p>Mark Shafarman Oacis Healthcare Systems 100 Drake's Landing Road, Suite 200 Greenbrae, CA 94904 1 415 925 4570 mshafarm@oacis.com</p>	<p>Tim Brinson Protocol Systems, Inc. 8500 SW Creekside Place Beaverton, OR 97008-7107 1 503 526 4392 tim@protocol.com</p>

Kinji Yamasaki  
Sholink Corporation  
2672 Bayshore Parkway, Suite 1001  
Mountain View, CA 94043  
1 650 969 5597  
kinji@sholink.com

## 1.2 *Supporting Organizations*

The following organizations have been involved in the process of developing, prototyping and/or reviewing this submission. The submitters of this response thank them for participating and giving their valuable input. Many ideas considered during the development of the PIDS specification were gained by participating in the MPI Workshop series held by HOST, CPRI and Los Alamos National Laboratory. A special thank you goes out to those organizations.

- 3M HIS
- Actamed
- AGFA
- Baptist Health Systems of South Florida (BHSSF)
- Blackwatch Technologies
- British National Health Service (NHS)
- CEDITI
- CERC, West Virginia University
- Context Consulting
- Computerized Patient Record Institute (CPRI)
- Federal University of São Paulo
- GE Medical Systems
- Integrated Medical Systems (IMS)
- International Business Machines (IBM)
- LinkSoft Technologies, Inc.
- Los Alamos National Laboratory (LANL)
- Medical Records Institute (MRI)
- OSF Healthcare System
- Per-se Technologies
- Philips Corporation
- Shared Medical Systems Corporation (SMS)
- Siemens Nixdorf
- Software Technologies Corporation (STC)
- STAR Consortium (EU project)

- Theragraphics, Inc.
- University of Michigan Health Systems (UMHS)
- University of Missouri

### 1.3 Conventions

IDL appears using this font and in a border.

### 1.4 Terminology

**Ancillary System** - A subordinate, secondary, or auxiliary system within an organization. An ancillary system could have its own simple PIDS and/or be sending requests to other PIDS. Examples: Vital Signs Monitoring Systems, Laboratory Information Systems (LIS), Scheduling Systems, Pharmacy Systems, Radiology Information Systems (RIS), etc.

**Attended Matching** - Matching that occurs as a result of human interaction. *See* **Matching**.

**Bind** - To logically attach or associate information about a person to an ID. Within an ID Domain, it can be said that an ID is a short hand representation of a real person or a key to more information about a person. For example, if a person is admitted to a hospital for the first time, a registration clerk enters identifying information into a system. All or part of this information is used to build a profile for the person. The system assigns an ID for the person and associates, or *binds* the profile to the ID.

**Candidate** - A person returned in the matching process that meets matching criteria. For example, if a clerk enters identifying information about a person and searches a PIDS system for a match, the PIDS system returns one or more candidates along with indications of how well each candidate matches the entered information.

**Client** - Any system or application that accesses or requests service from a Person Identification Service.

**Collision** - Within an ID Domain, a situation in which an ID is suspected to have been used by more than one person. Contrast with **Duplicate**.

**Component** - A cohesive set of software services. In this specification, a PIDS implementation is referred to as a component. A PIDS component implements varying PIDS interfaces as defined by PIDS conformance levels. For example, a *Simple PIDS* supports IdentifyPerson and ProfileAccess interfaces.

**Confidence or Confidence Level** - A matching algorithm's measure of probability that a candidate is a match. When a matching process returns a list of candidates it also returns a confidence level value with each entry in the list of candidates. The range of values for the confidence indicator is 0.0-1.0 with 1.0 being the higher confidence (e.g. 100%).

**Correlating ID Domain** - An ID Domain that correlates one or more other ID Domains. For example, a PIDS in a health care setting can be set up to correlate ID's from multiple providers (hospitals, clinics, Physicians offices, etc.) and multiple ancillary systems (lab, pharmacy, registration, etc.), where each participating system implements a different ID Domain.

**Correlation** - The creation of a cross-reference or mapping between Person ID's within a single ID Domain or across multiple ID Domains. For example, a PIDS set up to correlate ID's from a hospital and a lab stores both the hospital's ID and the lab's ID for any person with an ID in both ID Domains.

**CPR** - Computerized Patient Record

**Deprecate** - To indicate that an ID is not valid any more, within an ID Domain.

**Deactivate** - To *deactivate* an ID. No new information may be recorded under a deactivated ID, but it is allowed and accepted that some already exists. For example, when merging duplicate ID's in a PIDS, one of the ID's remains active. The other is deactivated, or marked as inactive. Also when an ID is deprecated. Contrast with **Collision**.

**Domain** - *See ID Domain*.

**Domain Name** - The name of an ID Domain in which an ID has meaning. That is, IDs are only relevant in a particular ID Domain. Each ID Domain has a Domain Name that is unique and different from all other ID Domain Names.

**Duplicate or Duplicate ID** - An ID is deemed to be a *duplicate* when it refers to the same person as one or more other ID's within an ID Domain. This results in the person "being in the ID Domain more than once". ID's that are known to be duplicates should be merged such that the associated person has one and only one unique ID and profile in the corresponding ID Domain.

**Federation** - As it relates to PIDS *federation* is the ability to structure ID Domains into hierarchies where the higher level ID Domains contain IDs for a superset of the persons with IDs in the lower level ID Domains. Operations such as searches for persons can be performed the same on all levels but applies only to the IDs known at that level and below. PIDS implementations can manage each of these ID Domains in which case the PIDS is said to be *federated*. PIDS provides a CorrelationMgr, NamingContext, TraderComponents, and EventComponent specifically to facilitate federation.

**ID or Identifier** - A sequence of characters that one or more systems in an ID Domain use to represent a person and bind related information. This could be numeric, alpha, and may include punctuation, etc.

**Identity** - The distinct real world person an ID and profile represents. In a PIDS an ID is established that represents a person's *identity* where each ID corresponds with one real world person. As an IDL interface, 'Identity' instances correspond one for one with a particular real world person that has been represented by an ID in the ID Domain.

**Identification** - The process of assigning an ID or finding an ID based on knowing some traits about the person.

**ID Domain or Domain** - A set of person ID's among which there is to be one unique person ID value per person or entity represented. For example, a hospital Admission, Discharge & Transfer computer system creates ID's for people as they are entered into the system. The set of ID's it manages is an ID Domain. ID Domains have a Domain Name which uniquely identifies it from other ID Domains.

People can have an ID from many ID Domains. Therefore, a person ID value has meaning for identification only if the correct ID Domain qualifies the ID value. For example, in the USA, the ID value 123-45-6789 can be used to identify a person if it is prefaced with the Social Security Number acronym, SSN, and it was assigned by the Social Security Administration.

Multiple systems can 'reside' in an ID Domain if they utilize/reference person ID's from the same ID Domain. For example, a lab system and a billing system can use the same medical record numbers to identify people. Each system can be said to 'reside' in the same ID Domain.

**MPI and EMPI** - Master Patient Index and Enterprise Master Patient Index.

**Matching** - The process that determines from a set of traits whether a person may already be known to a PIDS. The matching operation may return zero to many persons depending on the algorithm, weights on traits and threshold parameters used in the matching process.

**Merge** - To apply an operation on two or more ID's representing the same person in an ID Domain, which then results in one active ID for that person. All except one of the ID's are

**deactivated.** In other words, Merge operations are used to rectify the discovery of **Duplicate ID's**. *Contrast with Correlation*, where ID's are not deactivated.

**Naming Authority** - Any organization that assigns names determines the scope of uniqueness of the names and takes the responsibility for making sure the names are unique within its name space. In the same way that ID values are meaningful only within the context of their ID Domains, names are unique only within the context of their *naming authority*.

**PersonID or Person ID** - Same as **ID** or **Identifier**.

**PIDS** - The term *PIDS* is used in two ways: 1) to represent this Person Identification Service specification; and 2) to represent conformant implementations to this specification.

**Profile** - A set of information about a person that can be used to identify them. A profile consists of one or more **Traits**.

**Unmerge** - To take a person ID that has been merged with one or more other ID's and undo the merge, resulting in two or more person ID's. The person profiles are bound back to their original ID's before the merge. For example, suppose Person A has ID 1 and Person B with ID 2 had been merged into ID 1. A successful unmerge operation would restore Person A with ID 1 and Person B with ID 2.

**Subdomain** - If an ID Domain is being correlated over by a Correlating ID Domain then it can be called a Subdomain of the Correlating ID Domain.

**System** - An application or set of applications that interact with each other, interact with the PIDS or implement PIDS. *System* in this context is synonymous with *application*. Examples of systems might include a hospital or clinical information system, an ancillary system such as a lab or radiology system, or a financial/administrative system such as an ADT.

**Trait** - An attribute (i.e., information) that can be used to help identify a person. Traits are grouped to create a profile. Examples of a trait include *name, date of birth, sex, address, etc.*

**Unattended Matching** - A matching process that occurs without human intervention. *See Matching*. For example, an automated process may be configured to run once a night to scan an ID Domain, searching for potential duplicate person entries. Also when profiles are added to an ID Domain and the PIDS automatically determines if an ID already exists for the person.

## 1.5 *Proof of Concept*

Prototypes have been created for all major functionality defined in this specification except those in the Notification module.

## 1.6 *Changes to Adopted OMG Specifications*

No changes to the existing OMG specifications are needed by this specification.

## 1.7 *Response to RFP Requirements*

### MANDATORY RFP REQUIREMENTS

*Storage and retrieval of person ID and profile information.*

The PersonIdService::ProfileAccess, PersonIdService::IdentityAccess and PersonIdService::IdentifyPerson interfaces provide the storage and retrieval of person IDs and profiles.

*Access to candidate matches.*

The PersonIdService::IdentifyPerson.find\_candidates() operation provides access to candidate matches.

*Indicator of accuracy of candidate matches.*

The PersonIdService::CandidateSeq parameter returned from PersonIdService::IdentifyPerson.find\_candidates() contains the candidate matches and their confidence levels.

*Update person profile information.*

The PersonIdService::ProfileAccess.update\_and\_clear\_profiles() operation updates person profile information.

*Merge and Split of PersonIDs and profiles.*

The PersonIdService::IdMgr.merge\_ids() and PersonIdService::IdMgr.unmerge\_ids() provide merge and unmerge operations on person IDs and profiles.

*Assign person IDs automatically.*

The PersonIdService::IdMgr.find\_or\_register\_ids() operation automatically assigns person IDs.

*Handle mapping of multiple person IDs for the same person.*

The PersonIdService::CorrelationMgr.load\_profiles() operation handles mapping or cross-referencing multiple IDs for the same person.

*Flexibility to deal with arbitrary number of person ID types and traits.*

The PersonIdService::Trait is defined as a name/value pair which is a robust mechanism for dealing with arbitrary numbers of person IDs and traits.

### OPTIONAL RFP REQUIREMENTS

*Ability to support the interchange of information between two PIDSs and ID Domains.*

The PersonIdService::CorrelationMgr bridges multiple ID Domains.

*Ability to support federation of PIDS and ID Domains.*

The PersonIdService::IdentificationComponent provides the federation of PIDS and ID Domains by referencing local Naming and Trader services. The PersonIdService::CorrelationMgr and Notification::EventComponent interfaces are used to maintain consistency between PIDS that are federated.

*Ability to subscribe to notifications of updates to person profiles.*

The Notification::EventComponent provides the ability to subscribe to change events.

*Ability to use a standard set of traits data for matching parameters.*

The PersonIdTraits, HL7Version2\_3, and vCardVersion2\_1 modules provide sets of traits for matching parameters that an implementation can choose to use. This specification does not

require or recommend the use of any particular traits except for the PersonIdService::CorrelationMgr where the minimum traits supported are the patient's name, date of birth and sex. The HL7 Version 2.3 format is to be used for these three traits.

*Ability to detect and report suspected source identifier conflicts.*

The PersonIdService::CorrelationMgr detects source ID conflicts and reports them by listing multiple IDs from the same ID Domain when using the get\_corresponding\_ids() or retrieving the PersonIdTraits::CorrelatedIds trait. The PersonIdDuplicate and PersonIdCollision event types are defined for systems that may want to be notified of suspected conflicts.

**DISCUSSION POINTS REQUESTED**

*Protections for person confidentiality and privacy.*

Implementation of the PIDS service will use the CORBA Security service for those times when confidentiality or other security issues are a concern. The simplest model for access control would give client applications an all or nothing access to the service depending on their rights. The PersonIdService::IdentityAccess provides support for person oriented access control. Many implementations of the PIDS will not support these interfaces but those that do can control access depending on the person IDs being accessed. The Appendix gives an in depth discussion on using CORBA Security with PIDS to implement security policies that relate to confidentiality.

*Provisions for matching on non-unique keys.*

There are many person traits that can have more than one value. This submission treats the common cases of such values as sequences so the PIDS can store multiple values for a trait and a client can pass multiple values in for these traits. A trait may also have historical values that are no longer valid. A PIDS could store these historical values and use them for matching but the interfaces do not provide any special access mechanisms for them. Whether a PIDS does store historical values or not is considered a quality of service issue, which is outside the scope of this RFP.

*Capabilities for Administration and Configuration.*

No direct support is given for administration or configuration of a PIDS. These are considered implementation specific capabilities.

*Consistency of Identifiers across multiple PIDS.*

This submission assumes the person ID values used in two different ID Domains for the same person have nothing in common.

*Audit trail and system data maintenance.*

Any implementation of the PIDS specification needs to determine its own requirements for audit trails and system data maintenance. Since this submission addresses the person identification requirements for a wide variety of systems in the healthcare setting it is not possible to address this issue directly. Two types of audit trails need to be considered: 1) recording and reviewing changes to ID states and profiles; 2) as well as recording and reviewing of accesses to IDs and profiles by individuals. An implementation could log this information and provide their own mechanism to view the log. Alternatively they could use the EventComponent as a way to log information to an external service. There are discussions of creating an OMG specification for a standard Logging Service that may be used in the future. In addition the CORBA Security Service can be utilized for creating audit trails of type 2).

## 2. Overview of Response

---

### 2.1 Introduction

Throughout a person's lifetime he or she may have episodes of care provided by dozens or hundreds of healthcare providers, most of which may be assigning and maintaining patient IDs autonomously. In this arrangement each organization simply assigns IDs that uniquely identify patients within its local ID Domain of ID values, with the result that these ID values are meaningless outside that system or organization. These autonomously managed IDs suit the purposes of recording and retrieval of service records for the local organization. However, there is no basis for efficient collection or correlation of health records among multiple venues.

A typical healthcare information system will permit the user to submit a search for a person's record using some combination of identifying parameters for the person. When the user must collect a patient's healthcare information from a different organization or from a disparately-keyed system in the same organization, he or she typically must perform a new search in that other system - or ask a medical records person in the other organization to perform the search - in order to identify the person and retrieve the needed information.

In recent years, changes in the business of healthcare have made it both increasingly important and increasingly difficult to access the continuum-complete record of care for an individual. Risk-shared and capitation-based reimbursement policies have made it absolutely necessary to avoid redundant treatments. Increased specialization of providers has caused increased fragmentation and distribution of patient records.

Finally, organizational consolidation, growth and flux are exacerbating the problems associated with managing IDs. Most large integrated delivery systems are now competing on the basis of population share.

This RFP response defines the interfaces of a CORBA Person Identification Service (PIDS) that organizes person ID management functionality to meet healthcare needs. The PIDS is designed to:

- support both the assignment of IDs within a particular ID Domain and the correlation of IDs among multiple ID Domain
- support searching and matching of people in both attended-interactive and message-driven-unattended modes, independent of matching algorithm
- support federation of PIDS services in a topology-independent fashion
- permit PIDS implementations to protect person confidentiality under the broadest variety of confidentiality policies and security mechanisms
- enable plug-and-play PIDS interoperability by means of a "core" set of profile elements, yet still support site-specific and implementation-specific extensions and customization of profile elements
- define the appropriate meaningful compliance levels for several degrees of sophistication, ranging from small, query-only single ID Domains to large federated correlating ID Domains

## 2.2 *Problems Being Addressed*

The PIDS functionality is under severe demands for integration with other clinical and financial information systems, *but it must integrate with these systems without coupling with them.* Therefore it is critical that the functional scope of the PIDS be drawn with extreme care. This section and the section that follows delineate the scope of the PIDS by explicitly distinguishing those problems that are addressed by the PIDS from those that are not. The PIDS addresses the following specific problems.

### Identification

The PIDS will directly support the identification of people currently receiving care in a specific venue and ID Domain, and will support identification in the face of highly incomplete identifying information.

### ID Correlation

The PIDS will support both manual and automated correlation of IDs and records associated with health care consumers that have received care in different settings, and will address the problems of correlating IDs among the ID Domains of highly autonomous and frequently-reorganizing entities.

### Patient Confidentiality

While the PIDS itself is not required to enforce confidentiality, its interfaces are delineated so that "request interceptors" (implemented by CORBA Security Services or otherwise) can enforce *any* policy that is defined in terms of:

- the user's identity
- the person identity that is the target of the information request
- the ID Domain(s) involved, and
- the person traits requested.

Thus it will become reasonable to expect and demand that PIDS implementations compete on the basis of their abilities to enforce complex or individualized confidentiality policy and to protect person information from inferential analysis.

## 2.3 *Problems Not Being Addressed*

The following problems are specifically **NOT** addressed by this response:

### Confidentiality Policy

In the face of various and evolving confidentiality policy at private and governmental levels, the PIDS must not make any assumptions on confidentiality policy. Rather, it must, wherever possible, permit and support the broadest range of policies. See the Appendix for a discussion of how the OMA specifications, such as CORBA Security, can be used to implement some of the common characteristics that may be needed by security policies for PIDS implementations.

### Locating Episodic Patient Data (limited support)

The PIDS will support retrieval of any information that has been recorded in it as traits. However, its interfaces do not attempt to support the set-query and instance navigation paths that are needed in a serious CPR application. The PIDS IDL will support simple searches of directly recorded person attributes, but no attempt was made to support complex retrievals of visits or treatment and observation structures. Future CORBAMED RFPs could address these other functions.

Since the PIDS is an ID management service rather than a clinical application, it will support CPR information retrieval only to the extent of recording and retrieving object references to storage locations. It is then the sole prerogative of those objects to support retrieval operations and to enforce their respective security policies.

The practical impact of this distinction is that useful commercial applications will combine:

- 1• The PIDS functionality for ID management and profile retrieval with
- 2• Richer content support that is external to the PIDS - such as insurance, visit histories, observations, and treatments.

### HIPAA PL 104-191 (formerly Kassebaum-Kennedy Bill) and National Healthcare Identifiers

Some nations already have national healthcare IDs. At the time of this writing it is clear that the following conditions will likely apply to U.S. national healthcare IDs:

- ID assignment will require humans-in-the-loop, possibly for both initial assignment and eventual certification
- the ID will be more than a simple assigned next-sequence value
- there must be provisions for designations of permanent vs. temporary values, and support for correction operations (merging and unmerging)
- the ID will need to coexist and correlate with numerous legacy ID Domains in all but the newest organizations and systems

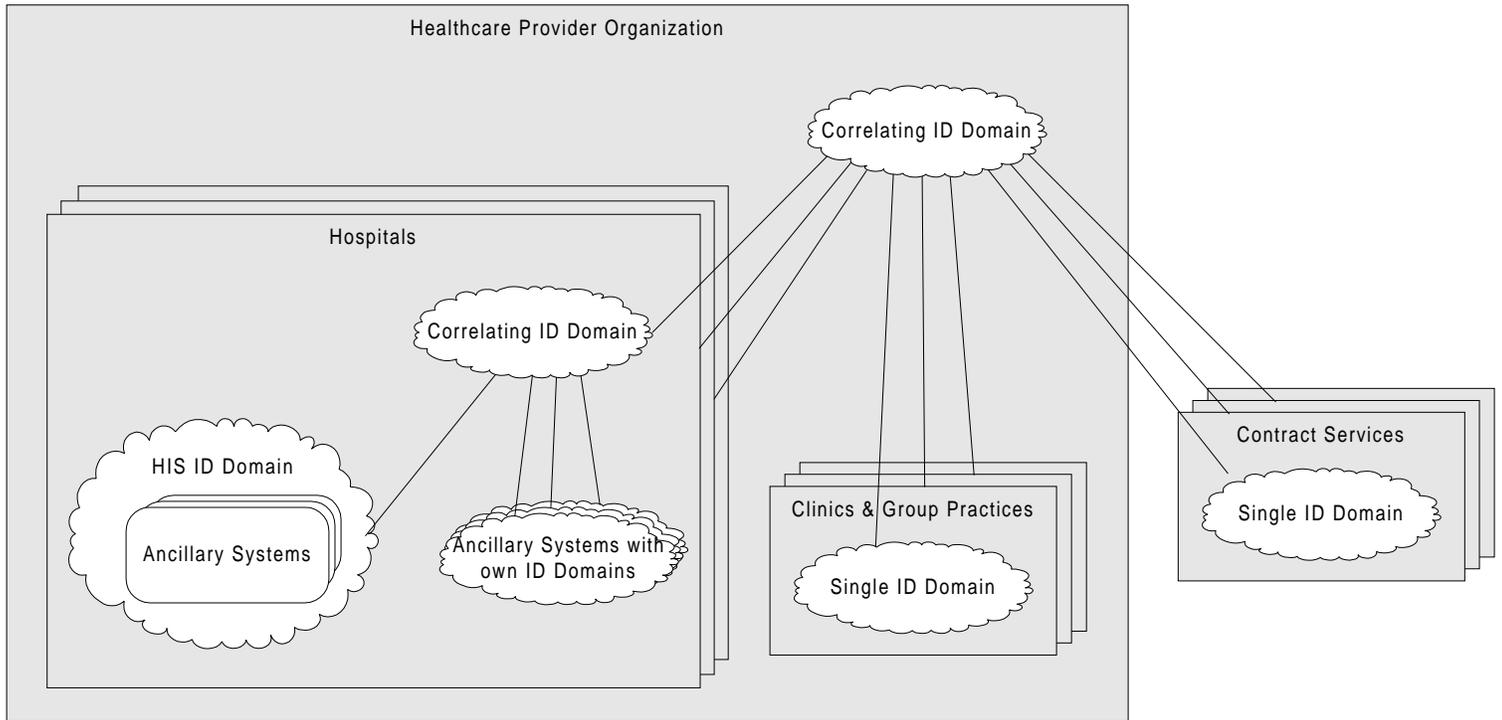
None of these requirements is in any way unique to national healthcare IDs. Any legally mandated ID, even though it has far reaching systemic and organizational impacts, is still an ID. As such it carries no unique *functionality or IDL requirements*. Furthermore, the PIDS solution must be generally applicable well beyond the U.S. Therefore such new ID Domains could be treated as "just another ID Domain" to the PIDS. On the basis of the importance and anticipated frequency of access of national healthcare IDs (such as to support government program claim processing) we include an explicit trait "NationalHealthId" to make it trivial to access these IDs, subject to security mechanisms.

While the issuance of the policy for national healthcare IDs is unlikely to affect the PIDS IDL, it is very likely to drive vendor implementations and applications to compete on the basis of:

- How effectively they can interface with related government deployed systems and administrative processes for ID assignment and certification
- How well they exploit the trustedness of digital certificates. For example, they may support the configuration of one or more certification authorities as ID Domains or support matching on certificates as Trait values.
- Their abilities to protect information integrity and patient confidentiality as the nationally-registered patient base grows and as private organizations merge. As patient empowerment grows in public awareness, vendors may compete by their abilities to support extreme cases of patient-administered access rights.
- How soon their products bring these strengths to the marketplace.

To summarize, the requirements of U.S. national healthcare IDs have served as a completeness check for IDL definition, and will raise the standard for scalability and flexibility of PIDS implementations. Its only PIDS IDL artifact is a predefined but optional trait to carry the national healthcare ID.

## 2.4 Domain Reference Model



*Domain Reference Model for PIDS.*

The above diagram represents a reference model for healthcare as it relates to this submission. This is not an attempt to define a globally useful reference model for healthcare as this model was chosen to highlight the aspects of healthcare related to person identity.

A hospital is likely to have one primary ID Domain which would typically be defined by the Healthcare Information System (HIS) or Admit Discharge Transfer (ADT) system. Many ancillary or departmental systems would be binding person related data (demographic, clinical, financial) to IDs in that ID Domain. There may be other ancillary systems that use their own identification make them each a separate ID Domain. The hospital may manually mirror these ID Domains against a 'main' ID Domain or electronically connect these systems to a common registration process to actually consolidate them into a single ID Domain. Today they use an MPI system to do this.

Today's Healthcare Provider Organizations may consist of multiple hospitals and multiple clinics forming an Integrated Delivery System. Each clinic that is automated will likely have one practice management system, which manages its own ID Domain.

A Healthcare Provider Organization also has contract services such as reference labs. Each of these would have their own ID Domain. The Healthcare Provider Organization may need a higher level of correlation that consolidates the various ID Domains.

There could be yet higher levels of correlation among Healthcare Provider Organizations. These do not add new relationships as they are just correlating over correlating ID Domains as we have here. Their population sizes would potentially be much larger. Even though it is not shown in the diagram an ID Domain could have multiple correlating ID Domains that correlate its IDs with other ID Domains.

## 2.5 Determination of RFP Scope

When identifying a person there is a variety of identifying information that may be used including demographics (address, place of birth, etc.), biometric information (finger print, photograph, blood type, etc.) and IDs they may have from other sources (SSN, drivers license number, insurance number, etc.). Some biometric information (e.g. blood type or weight) may also be considered clinical information.

**Question** *Can't PIDS be used for accessing clinical information about a patient such as electronic medical records?*

**Answer** *No. There is a fuzzy line between biometric information and clinical information but the distinguishing trait is their stability over time. The predefined names of traits do not cover clinical information and PIDS is not meant to store clinical information.*

*The information about a patient that is useful for identifying them are those traits that remain constant or change very slowly over time. The expectation in using them is that they have not changed so they can be stored and then used to identify the person the next time they are encountered. The information being gathered as part of the clinical diagnosis is expected to have changed and are in general unstable over time.*

The key word in Patient Identification Service is 'Identification'. Identity is definitely not unique to patients. When looking at the need for identification in many fields and looking at the interfaces being proposed, we envisioned this being applicable to much broader entities than just patients.

**Question** *Should the thing being identified in this specification be 'patients', 'persons', or 'subjects'?*

**Answer** *Persons!*

*The concept of identifying a subject by its characteristics is applicable to more than just persons but there are a couple problems proposing this specification to cover them. We need to define a default set of identifying information that can be used with this service for identifying patients. This information is specific to identifying persons (not just patients) but not general entities or subjects.*

*Applying this specification to general entities that need to be identified is beyond the scope of this RFP and even the CORBAmed DTF.*

**Question** *Isn't a Person Identification Service needed in other domains than just healthcare? Why is CORBAmed standardizing it?*

**Answer** *Yes it is needed in other domains but as an ancillary part of work flow as their real area of concern is finance, telecommunications, manufacturing, etc. In healthcare the person (we call them patients) is what the business is all about.*

*Other domains need to identify a person in order to make sure a financial transaction or other business contract is applied to the correct person. In healthcare it is important when diagnosing a problem that you have the same 'physical' being that a patient record refers to. Therefore incorrectly identifying the person can have more severe consequences.*

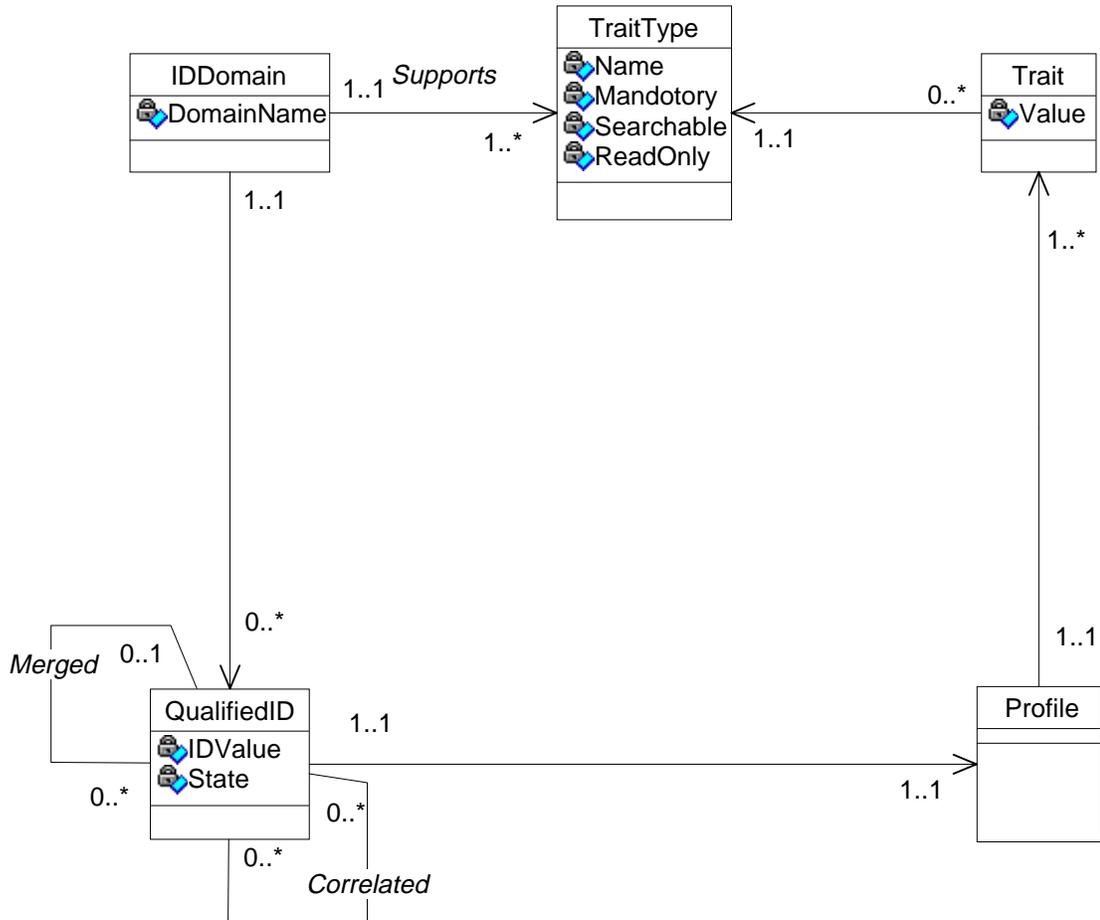
*For these reasons it does make sense that CORBAmed standardize interfaces for identifying persons. These interfaces should be useful outside healthcare as well.*

## 2.6 *Design Goals*

Besides meeting the requirements spelled out in the RFP there are a number of design goals that helped us determine solutions.

- The specification should be useful for a variety of systems and applications in the healthcare environment.
- The specification should be simple and easy to implement by vendors wrapping current products.
- The specification should be easy to deploy in healthcare settings in useful but manageable increments.
- The specification should afford interoperability when multiple systems implementing it need to communicate.
- The specification should provide as much out of box (plug-and-play) interoperability.
- The specification should provide a basis for future CORBAmed specifications to build upon.
- The specification should provide standard mechanisms to solving the problems without requiring a particular architecture.
- The specification should be confidentiality policy independent.
- The specification should be matching policy independent in order to provide various qualities of service and to not limit the use of future matching engine advancements.
- The specification should mandate only minimal, widely agreed upon content.
- The specification should include a standard set of person traits for interoperability but be extensible to meet local requirements.

## 2.7 PIDS Conceptual Model

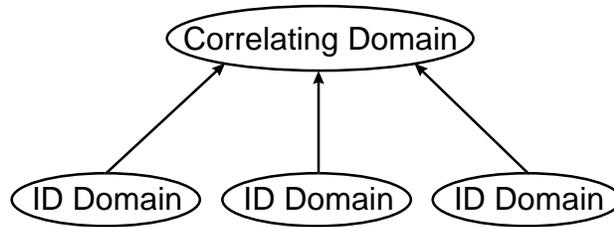


The PIDS Conceptual Model shows the core data elements that PIDS implementations must deal with. It also shows how these relate to a real world person. The following sections will describe this data and show how they are used by PIDS.

In PIDS implementations, the **QualifiedID** associations might carry administrative and auditing attributes such as, time stamp, user stamp, source system, and specific operation types for merge and correlation. For example, specific operation types could be added to correspond to the HL7 2.3 merge and link events.

## 2.8 PIDS Identification Model

The following figure provides the basic structural elements of our identification model:



The ID Domain is the basic building block of our PIDS model. An ID Domain maintains a unique Identifier (ID) for each person identity to be represented in the ID Domain. Ideally there is one and only one ID per person, but in reality there may be duplicates where a given person has been assigned an additional ID in the same ID Domain. For internal consistency, the ID Domain cannot assign two persons the same ID or it would have no means to distinguish between the two entries. The ID is an internal control mechanism and may or may not be used externally. Thus the ID and its ID Domain together create a unique ID for the person.

The ID Domain provides a framework for representing existing systems. For example, the MPI in a registration system is an ID Domain as is the patient index in an ancillary system. In other words, ID Domains already exist in any patient oriented system. Our terminology simply provides descriptive unit for existing identification processes.

Within the PIDS specification, several interfaces are detailed. The two most directly used with an ID Domain are IdentifyPerson interface and ProfileAccess interface. The IdentifyPerson interface is basically a query used to send traits to be matched in the ID Domain and to receive the matching candidate(s). The ProfileAccess interface can be a query or an update used to send an ID for a specific person and to receive that person's profile. A profile is a set of traits containing the person's values for their respective traits. Two additional interfaces are specialized forms of the ProfileAccess interface, namely the SequentialAccess interface and the IdentityAccess interface.

The coordinating structural unit for the PIDS model is the Correlating Domain. The Correlating Domain allows access to the correlated profiles of the IDs in all of the participating ID Domains. This building block provides the framework for cross-referencing IDs across the participating ID Domains. The Correlating Domain is a new building block which complements existing ID Domains.

Two interfaces illustrate the role of the Correlating Domain. First, profiles are added and correlated via the CorrelationMgr interface (PersonIdSeq register). Second, cross-referenced IDs are obtained by the CorrelationMgr interface (UniquePersonIdSeq get\_corresponding\_ids).

Integrity of each domain is maintained by their respective managers – the IDMgr interface and the CorrelationMgr interface.

To illustrate the use of this model in actual implementations, consider the following examples. The MPI in a registration system or the patient index in an ancillary system is an ID Domain. When a registration system is interfaced to ancillary systems, the registration systems MPI would be an ID Domain and the ancillary systems using the ID assigned by the registration system would be codependent domains. When multiple provider systems are interfaced, such as for an integrated delivery system, the enterprise MPI would be the Correlating Domain for the participating ID Domains of the provider systems. If the enterprise were to issue the IdentifyPerson interface and the ProfileAccess interface to the Correlating Domain instead of to each ID Domain, then the enterprise MPI would be both a Correlating Domain and an ID Domain.

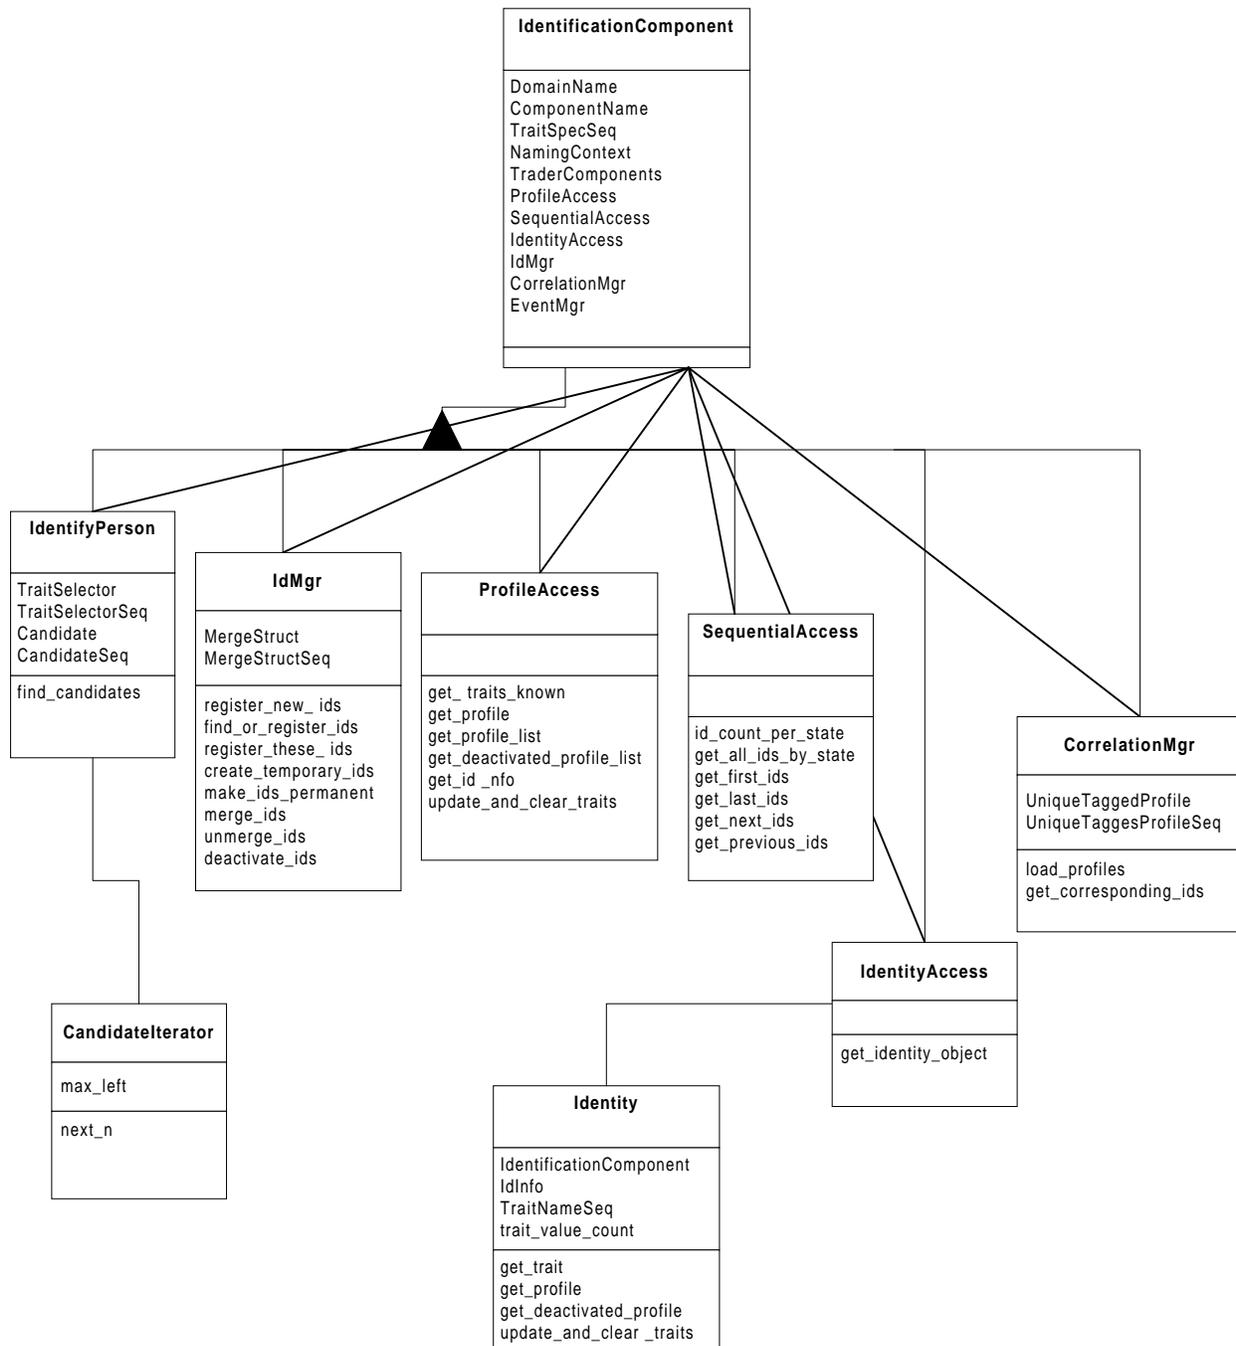
The PIDS model provides a tool to manage identification which provides flexibility for both the short-term (existing systems) and for the long-term (quickly accommodate business requirements). The PIDS model also allows both hierarchical and peer structures to evolve. Since a Correlating Domain can correlate both ID Domains and other Correlating Domains, hierarchical structures can be built for enterprises. In addition, an ID Domain can participate in multiple Correlating Domains, allowing peer structures to be built.

Peer structures offer significant potential for identification outside of enterprises. For example, any need to track a particular population could fund its own Correlating Domain and establish relationships with the appropriate ID Domains. The cost would be minimal for the Correlating Domain since it could take advantage of the accessibility to other Correlating and ID Domains. Furthermore, accessibility would be determined by the participating Domain, e.g., each participating Domain would only allow access by authenticated, authorized requesters.

One additional point involves the likelihood of a National Healthcare ID. The PIDS model views the National Healthcare ID as its own ID Domain since assignment and control would not be under the control of local ID Domains. Thus this national ID would be a trait in ID Domains (and Correlating Domains) creating an automatic cross-reference for searches.

### 3. *PersonIdentificationService module*

PIDS Components and Inheritance Diagram



The above diagram shows the interfaces defined in the PersonIdService module. The main functional interfaces inherit from the IdentificationComponent interface. They are also referenced by the IdentificationComponent. This section describes these interfaces and their data types in detail.

```
//File: PersonIdService.idl
#ifndef _PERSON_ID_SERVICE_IDL_
#define _PERSON_ID_SERVICE_IDL_

#include <orb.idl>
#include <NamingAuthority.idl>
#include <Naming.idl>
#include <Trading.idl>
#include <Notification.idl>

#pragma prefix "omg.org"

module PersonIdService
{
    // . . .
};

#endif // _PERSON_ID_SERVICE_IDL_
```

The core Person Identification Service (PIDS) consists of the type declarations and interfaces in this module. PIDS is structured as a component that has multiple interfaces that may be implemented by any particular instance of the service. Each of the interfaces represent a particular piece of functionality and are optional. Even though an implementation of PIDS is free to implement which interfaces they need, in the conformance section of this specification specific interfaces are called out that must be implemented in order to claim conformance to the various conformance classes.

**#include "NamingAuthority.idl"**

The types declared within the NamingAuthority module are used for defining the names of ID Domains, IDs, components and traits. The names of ID Domains, IDs, and components must be defined by implementations of the service or other parties. This specification defines some of the names for traits but they may be added independently by other parties. The global uniqueness properties of the NamingAuthority provide a mechanism for names to be created independently with low likelihood of name clashing.

**#include "CosNaming.idl"**

The CosNaming module is used for a couple reasons. This specification defines some standard names for PIDS to use within a naming service. PIDS components also include a NamingContext that is used to provide federation.

**#include "CosTrading.idl"**

The CosTrading module is used by PIDS in various ways. This PIDS specification defines Trader service types that PIDS components are to use when offered to a Trader service. This specification also uses many of the design patterns found in the Trader specification. In addition PIDS components include a TraderComponents that is used to provide federation.

**#include "Notification.idl"**

The Notification module is used for defining change event notifications on the component.

```
#pragma prefix "org/omg"
```

In order to prevent name pollution and name clashing of IDL types this module (and all modules defined in this specification) uses the pragma prefix that is the reverse of the OMG's DNS name.

### 3.1 Basic Types

```
// -----  
// Common Data Types  
//  
typedef NamingAuthority::AuthorityId DomainName;  
typedef sequence< DomainName > DomainNameSeq;  
  
typedef NamingAuthority::LocalName PersonId;  
typedef sequence< PersonId > PersonIdSeq;  
  
struct QualifiedPersonId {  
    DomainName domain;  
    PersonId id;  
};  
typedef sequence< QualifiedPersonId > QualifiedPersonIdSeq;  
  
typedef NamingAuthority::QualifiedNameStr TraitName;  
typedef sequence< TraitName > TraitNameSeq;  
typedef any TraitValue;  
struct Trait {  
    TraitName name;  
    TraitValue value;  
};  
typedef sequence< Trait > TraitSeq;  
typedef TraitSeq Profile;  
typedef sequence< Profile > ProfileSeq;  
  
enum IdState { UNKNOWN, INVALID, TEMPORARY,  
              PERMANENT, DEACTIVATED };  
typedef sequence<IdState> IdStateSeq;  
struct IdInfo {  
    PersonId id;  
    IdState state;  
    PersonId preferred_id;  
};  
typedef sequence<IdInfo> IdInfoSeq;  
  
// -----  
// Miscellaneous Data Types  
//  
typedef string ComponentVersion;  
struct ComponentName {  
    NamingAuthority::QualifiedName name;  
    ComponentVersion version;  
};  
  
struct TraitSpec {  
    TraitName trait;  
    boolean mandatory;  
    boolean read_only;  
    boolean searchable;  
};  
typedef sequence< TraitSpec > TraitSpecSeq;  
  
enum HowManyTraits { NO_TRAITS, SOME_TRAITS, ALL_TRAITS };  
union SpecifiedTraits switch ( HowManyTraits )  
{  
    case SOME_TRAITS: TraitNameSeq traits;  
};
```

```

struct TaggedProfile {
    PersonId id;
    Profile profile;
};
typedef sequence<TaggedProfile> TaggedProfileSeq;

struct QualifiedTaggedProfile {
    QualifiedPersonId id;
    Profile profile;
};
typedef sequence<QualifiedTaggedProfile> QualifiedTaggedProfileSeq;

struct ProfileUpdate {
    PersonId id;
    TraitNameSeq del_list;
    TraitSeq modify_list;
};
typedef sequence< ProfileUpdate > ProfileUpdateSeq;

struct MergeStruct {
    PersonId id;
    PersonId preferred_id;
};
typedef sequence< MergeStruct > MergeStructSeq;

struct TraitSelector {
    Trait trait;
    float weight;
};
typedef sequence<TraitSelector> TraitSelectorSeq;

struct Candidate {
    PersonId id;
    float confidence;
    Profile profile;
};
typedef sequence<Candidate> CandidateSeq;

interface CandidateIterator {
    unsigned long max_left();

    boolean next_n(
        in unsigned long n,
        out CandidateSeq ids );

    void destroy();
};

typedef unsigned long Index;
typedef sequence< Index > IndexSeq;

enum ExceptionReason {
    UNKNOWN_TRAITS,
    DUPLICATE_TRAITS,
    WRONG_TRAIT_FORMAT,
    REQUIRED_TRAITS,
    READONLY_TRAITS,
    CANNOT_REMOVE,
    MODIFY_OR_DELETE
};

struct MultipleFailure {
    Index the_index;
    ExceptionReason reason;
    TraitNameSeq traits;
};
typedef sequence< MultipleFailure > MultipleFailureSeq;

interface Identity;
typedef sequence< Identity > IdentitySeq;

```

```

// -----
// Exceptions
//

exception InvalidId { IdInfo id_info; };
exception InvalidIds { IdInfoSeq id_info; };
exception DuplicateIds { PersonIdSeq ids; };
exception UnknownTraits { TraitNameSeq traits; };
exception DuplicateTraits { TraitNameSeq traits; };
exception WrongTraitFormat { TraitNameSeq traits; };
exception InvalidStates {};
exception TooMany { unsigned long estimated_max; };
exception MultipleTraits { MultipleFailureSeq failures; };

exception ReadOnlyTraits { TraitNameSeq traits; };
exception CannotRemove { TraitNameSeq traits; };
exception ModifyOrDelete { MultipleFailureSeq failures; };
exception NotImplemented {};

exception InvalidWeight {};
exception CannotSearchOn { TraitNameSeq traits; };

exception IdsExist { IndexSeq indices; };
exception RequiredTraits { TraitNameSeq traits; };
exception ProfilesExist { IndexSeq indices; };
exception DuplicateProfiles { IndexSeq indices; };

exception DomainsNotKnown { DomainNameSeq domain_names; };
exception IdsNotKnown { QualifiedPersonIdSeq ids; };

```

There are a number of structured types used widely through out the PersonIdService module. These characterize ID Domains, IDs, traits of a person, etc.

### Common Data Types

#### **DomainName, DomainNameSeq**

This structure can be used to identify a fully qualified ID Domain name following the specified rules of the NamingAuthority module.

IDs for people are always relative to some ID Domain. An ID Domain is the extent of meaningfulness of the IDs within that ID Domain. The DomainName is an AuthorityId which makes it possible to define these independently without two ID Domains being named the same.

The DomainNameSeq is a sequence of DomainNames.

#### **PersonId, PersonIdSeq**

The PersonId is the simplest representation of an ID and is defined as a string. It can only be used in contexts where the ID Domain of the ID is either specified or is already known. The ID for a person relative to an ID Domain is represented as a name within that naming authority. A PersonId can only be interpreted in the context of an ID Domain.

A PersonId is the concrete representation of ID in an ID Domain. The representation of characters forming an ID is dependent on the specifics of the ID Domain. Given IDs from different ID Domains it is not possible to tell reliably what ID Domain they came from solely based on the syntax.

The PersonIdSeq is a sequence of PersonIds.

#### **QualifiedPersonId, QualifiedPersonIdSeq**

The QualifiedPersonId is a fully qualified ID since it contains the name of the ID Domain and the PersonId itself. If two QualifiedPersonIds are found to be identical they can reasonably be assumed to represent the same person (within the limits of the NamingEntity algorithms of the NamingAuthority module if they are followed).

#### **TraitName, TraitNameSeq**

The names given to the traits of a person are characterized by the TraitName type. The trait name is a string in the form of the QualifiedNameStr so that additional traits can be created by users of the specification without having name clashes or having to register their trait names with a centralized authority.

The TraitNameSeq is a sequence of TraitNames. It is very useful for specifying the set of traits a client is interested in.

#### **TraitValue**

The TraitValue represents the value of a person's trait. Since the set of traits is virtually endless and could be of any type imaginable (including multimedia) the 'any' data type is used.

#### **Trait, TraitSeq, Profile, ProfileSeq**

The characterization of a specific trait is given as a name value pair and is called 'Trait'. A sequence of these can be put together to create a TraitSeq.

A Profile is used when referring to the traits stored by a PIDS and bound to a PersonId. It is also used in the matching process for looking up persons. The term 'Profile' is given to these sequences of traits since they are used a lot. At times a sequence of Profiles are needed (ProfileSeq).

#### **IdState, IdStateSeq, IdInfo, IdInfoSeq**

The IdState represents the set of states that an ID may take on. Each state has some well defined semantics. The IdState is a characteristic of an **ID** where as Traits are characteristics of the **person** represented by an ID. For this reason there is a separate operation to access it. The IdState can also be used for change event audit trails and change event notification. The operations on the IdMgr interface affect the state of an ID.

The clients to the interface see a different behavior for each ID state.

- UNKNOWN - This state indicates the service does not know if the ID exists or not and if so what the actual state might be. This is used by components that do not manage an ID Domain but are a resident of an ID Domain where they may only know about a small subset of the IDs that exist.
- INVALID - An ID in this state can only be used in operations that create an ID. If a PIDS component knows all IDs in its ID Domain every other possible ID has the Invalid IdState.
- TEMPORARY - An ID can be created as temporary without indicating any mandatory traits. A common usage is to create an ID that data can be bound to a patient before that patient is identified with an appropriate confidence. A temporary ID can be made permanent, merged, or deprecated. A temporary ID is made permanent explicitly - just updating the profile to contain all required traits does not change the state.
- PERMANENT - When an ID is created as permanent all mandatory traits must be provided. A permanent ID can be merged or deprecated but not made temporary.
- DEACTIVATED - Once an ID is expected not to be needed any more it can be deactivated (merged or deprecated) which keeps it around for historical purposes. A deactivated ID is in its final state and can not be transitioned to any other state by PIDS operations except unmerging.

The IdStateSeq is used for passing in the states of interest in a query. In a query, an IdStateSeq having a value of zero length should be interpreted to mean, return all valid states for the operation except UNKNOWN and INVALID.

The IdInfo structure contains the ID and the state of the ID. It also contains a preferred ID that is blank if in any state other than DEACTIVATED or may contain the ID this ID has been superseded by (merged with) if the state is deactivated. The IdInfoSeq is a sequence of IdInfo structures.

#### **ComponentVersion, ComponentName**

The ComponentVersion is a numeric string with major and minor version indications. The part of the string up to the first period (‘.’) is the major version indication which must contain only numeric digits. The rest of the string after the first period is the minor version indication which may contain any printable characters. When a component is changed significantly enough that clients may need to react differently the major version should be rolled. The minor version part of the string is used for vendor specific purposes. If only the minor version number of a component changes clients can assume behavior from the service that is compatible with what came before. For example a change to the ID Domain name, supported traits, or supported interfaces would require changing the major version number for a component since they are considered static over time.

The ComponentName includes a globally unique name and the version for a component. The name is used to uniquely identify instances of components. There is no inherent meaning implied by the component name.

#### ***Miscellaneous Data Types***

##### **TraitSpec, TraitSpecSeq**

The TraitSpec characterizes a supported trait by listing its name and indicating if it can be changed, whether it is required for permanent IDs, and if it can be searched on. If the read\_only field is true then operations that attempt to modify it will throw an exception. If the mandatory field is true then it can be assumed that any IDs in the Permanent ID state will have a value for that trait. This may be helpful in knowing what traits to ask for. In addition if the IdMgr interface is implemented by the component an exception may be raised if there is an attempt to create a permanent ID without a value supplied for the trait. Also an attempt to clear the value of the trait on a permanent ID raises an exception. If the searchable field is true then this trait can be used for searches, assuming the component implements the IdentifyPerson interface. The find\_candidates() operation will raise an exception for traits that have this field set to false.

A TraitSpecSeq is a sequence of TraitSpec structures and is used for specifying a complete set of traits supported by a component.

##### **HowManyTraits, SpecifiedTraits**

The SpecifiedTraits union is used by clients to indicate the traits they are interested in obtaining values for. The selector being HowManyTraits could indicate all of them or just some. If it indicates some the union body has a list of the ones they are interested in.

##### **TaggedProfile, TaggedProfileSeq**

A TaggedProfile structure includes a profile and the ID associated with it. The TaggedProfileSeq is a sequence of TaggedProfiles.

##### **QualifiedTaggedProfile, QualifiedTaggedProfileSeq**

A QualifiedTaggedProfile contains the globally unique ID for a person and a profile for that ID.

The QualifiedTaggedProfileSeq is a sequence of QualifiedTaggedProfiles.

**ProfileUpdate, ProfileUpdateSeq**

The ProfileUpdateSeq is used to update (modify) a list of profiles in a single call. Each entry (ProfileUpdate) contains the ID, the names of traits that need to be cleared and the set of new trait values that need modified.

**MergeStruct, MergeStructSeq**

The MergeStruct is used to represent a request to deprecate one ID in preference to another one. The MergedStructSeq is used to batch a group of merges in a single call. The preferred\_id field contains the ID that the merged\_id is merged to.

**TraitSelector, TraitSelectorSeq**

The TraitSelector is the matching parameter used to tell the service identifying information about a person (or persons) that a client is interested in during a search. The structure contains a Trait plus the weight field as an indication of the relative weights to put on the trait. Since the matching engine is not being standardized by the PIDS this weighting hint may be ignored by the service.

**Index, IndexSeq**

The IndexSeq is needed for some exceptions to return the indices of passed in sequences that caused the exception.

**ExceptionReason, MultipleFailure, MultipleFailuerSeq**

The ExceptionReason enumeration is used to specify what the cause of the exception is in a MultipleFailure structure.

The MultipleFailure structure is used to report information back in an exception that can have problems with multiple traits each on multiple IDs or profiles passed in. The 'the\_index' field contains the index from the original sequence passed to the operation that contained profiles or IDs. The 'traits' field contains the names of the traits that had a problem at the index.

The MultipleFailureSeq is used to pass back problems that were found on multiple passed in profiles or IDs.

**IdentitySeq**

The IdentitySeq contains a sequence of object references to the Identity interface. It is used for retrieving multiple Identity object references from the IdentityAccess interface with one call.

**Candidate, CandidateSeq**

The Candidate is returned from looking up persons when there may be partial matches. It contains the ID of the person, a confidence of how well that person's profile matches the profile selector and the set of traits requested in the operation.

CandidateSeq is a sequence of candidates. These are returned from looking up a person in the service.

**CandidateIterator interface**

CandidateIterator is an iterator for candidates. It is used when more IDs match a look up request than the client wants at one time.

**max\_left()**

Returns the count of candidates left on the iterator.

**next\_n()**

Returns the number of candidates asked for or all left on the iterator, whichever is less. Removes the returned candidates from the iterator before returning. If all candidates on the iterator are returned the object is deleted upon returning the results.

**destroy()**

Destroys the iterator instance while there are still candidates on the iterator. Clients must always call this operation if they are finished accessing candidates from the iterator when more are left. It is a server implementation decision when to auto delete the iterator if the destroy is not called for a long time.

### Exceptions

The following exceptions are generally useful by most or all of the interfaces of this module.

#### **InvalidId**

This exception is generated when a single ID is passed into an operation but the ID is not valid for the operation. There are a number of reasons this may occur. The ID may not be known and hence may not be relevant to the service being called. The ID may not exist in the ID Domain that a service exists in. The operation being called may require the ID to be in a particular state. The IdInfo for the ID is passed back to the caller so they may inspect the reason of the exception.

If the service knows of all IDs in an ID Domain (e.g. manages the IDs) then IDs that don't exist are returned as Invalid. If the service resides in an ID Domain but is not aware all possible IDs the ID state is considered Unknown. The difference in semantics to the client is that if they receive the exception with Unknown IDs they could try the manager of the ID Domain and to see if the ID exists.

#### **InvalidIds**

This exception is similar to InvalidId but is used with operations that take more than one ID as part of the parameter(s). The complete set of violating IDs must be returned. A client could remove the violating IDs and repeat the operation knowing this exception will not be thrown again (unless the server state changed between calls).

#### **DuplicateIds**

This exception is raised when the same ID is passed more than once in a sequence. The complete set of duplicate IDs are returned.

#### **UnknownTraits**

This exception may be thrown for operations that take one or more trait names or traits as parameters (typically as sequences). The service throws the exception for traits it does not support. The complete set of violating traits must be returned. A client could remove the violating traits and repeat the operation knowing this exception will not be thrown again.

#### **DuplicateTraits**

This exception is raised when the same trait or trait name is passed more than once in a sequence. The complete set of duplicate trait names are returned.

#### **WrongTraitFormat**

This exception may be thrown for operations that take one or more traits as parameters (typically as sequences). The service throws the exception if the TraitValue is of the wrong IDL type. The complete set of violating traits must be returned. A client could remove the violating traits and repeat the operation knowing this exception will not be thrown again.

**InvalidStates**

This exception is thrown when duplicate states are passed in or the states are passed in to operation that are not allowed.

**TooMany**

Many operations return a sequence of data that could be very large. If the size is larger than is feasible for the service to return (and the service can detect the fact) it will throw this exception. Typically the size of the returned result is determined partly by passed in parameters. The service may return the estimated maximum size for a passed in sequence. If that does not make sense for an operation or the service does not know what the maximum size would be it can return zero. Alternatively the service could return a size that it is confident will work even though it may not be close to the maximum. In either case the client should be able to pass in a sequence of the estimated\_max and this exception would not be thrown again.

**MultipleTraits**

This exception is similar to the DuplicateTraits exception except it applies to operations that could have problems with traits on multiple IDs or profiles that were passed in. Each entry in the sequence returned contains the index of the original profile or ID that had a problem and the trait names that had a problem at the index.

**ReadOnlyTraits**

This exception is thrown if a client tries to modify or set a trait that is read only. The complete set of violating traits is returned to the client so they may remove them from the list and try again.

**CannotRemove**

This exception is raised when a client tries to clear a trait that is read only or mandatory and the ID is in the Permanent IdState. The complete set of violating traits is returned to the client so they may remove them from the list and try again.

**ModifyOrDelete**

This exception may be raised on operations that allow both clearing of traits and modifying/setting other traits. If a client passes in the same trait to be cleared and modified this exception is raised.

**NotImplemented**

This exception is raised by the few optional operations on the interfaces. The operations that may raise this exception are logically grouped with other operations on the same interface. That is why these optional operations are not separate interfaces on IdentificationComponent.

**InvalidWeight**

This exception is thrown if a weight is passed in for a trait selector that is less than 0.0 or greater than 1.0.

**CannotSearchOn**

This exception is raised if traits are passed in to be matched against when their searchable field is false.

**IdsExist**

If a client tries to create new IDs and they already exist this exception will be raised. The sequence indices for all violating IDs are returned.

#### RequiredTraits

If a client tries to create a permanent ID without giving all the mandatory traits the service raises this exception or returns a temporary ID instead.

#### ProfilesExist

If a client tries to create new IDs and the profiles passed in correspond to IDs that exist this exception will be raised. The sequence indices for all violating profiles are returned.

#### DomainsNotKnown

This exception is thrown when an operation on the CorrelationMgr interface receives a parameter specifying an ID Domain name that is not part of the source domains it knows about. All violating DomainNames are returned.

#### IdsNotKnown

The IDs referenced are not known by the service. A complete list of violating IDs are returned.

### 3.2 *IdentificationComponent interface*

```
// -----  
// IdentificationComponent  
//  
interface ProfileAccess;  
interface SequentialAccess;  
interface IdentityAccess;  
interface IdentifyPerson;  
interface IdMgr;  
interface CorrelationMgr;  
  
interface IdentificationComponent  
{  
    readonly attribute DomainName                domain_name;  
    readonly attribute ComponentName             component_name;  
    readonly attribute TraitSpecSeq              supported_traits;  
  
    readonly attribute IdentifyPerson            identify_person;  
    readonly attribute ProfileAccess             profile_access;  
    readonly attribute SequentialAccess          sequential_access;  
    readonly attribute IdentityAccess            identity_access;  
  
    readonly attribute IdMgr                     id_mgr;  
  
    readonly attribute CorrelationMgr            correlation_mgr;  
    readonly attribute Notification::EventComponent event_component;  
    readonly attribute CosNaming::NamingContext naming_context;  
    readonly attribute CosTrading::TraderComponents trader_components;  
};
```

The Person Identification Service is based on a component model similar in effect to the pattern used by the Trader service. All identification interfaces are inherited from a core IdentificationComponent. The IdentificationComponent also has references to each of the other interfaces that are implemented. This makes it possible for a client to obtain an object reference to any of the identification interfaces and be able to easily find out what other functionality is implemented and to navigate to those interfaces.

The IdentificationComponent service encapsulates a logical table with person characteristics (traits) that is keyed by an ID, a matching engine that uses that table to map traits back to their ID and a table of ID information (e.g. ID states).

This component can also be used by systems and applications (such as ancillary systems) that use IDs from an ID Domain but are not the manager of the IDs in the ID Domain. The IDs known by a system implementing this interface are likely to be a subset (possibly very small subset) of the IDs that exist for that ID Domain. Due to the fact that they may only be dealing with a few patients' data at any one time. The sequential access operations can be useful for systems with only a few IDs where a client application can request a page of data at a time and allow a user to scroll through the data.

Multiple systems in an ID Domain may implement the PIDS component interface and be binding identifying data (as well as other information beyond the scope of this RFP) to IDs in that ID Domain. These may (or may not) also be exposed in the naming context and Trader service of the ID Domain manager.

### IdentificationComponent interface

An IdentificationComponent has a number of optional interfaces that it may implement. There are a variety of systems and applications that may implement the interface for different reasons. They may pick and choose the functionality that is right for them. A component may be implemented by a single object (e.g. inheriting all the interfaces into one implementation dependent interface) or a different actual object for each interface supported (potentially distributed from each other) or any combination of the two. If multiple objects are used to implement the component they must all maintain consistency so the client can treat them as one. That is all the attributes on the component must return functionally identical results.

The interfaces implemented by a single component must operate over the same set of person IDs. Also the trait values must be consistent. Semantically they are using the same data base. If they are actually using separate data bases it is up to the implementation to maintain consistency.

There are a wide variety of ways an IdentificationComponent can be used. The following lists demonstrate the wide variety of possible uses.

- Master Patient Index Systems - These systems correlate over multiple ID Domains and represent the fullest use of the various interfaces.
- Registration Systems - These may manage a single ID Domain and not know of other ID Domains. The registration, identification, and demographics services could be created as a single component even though they may be separate systems and possibly from separate vendors. Alternatively the three separate services could be implemented as separate components.
- Ancillary Systems - Some ancillary systems may use IDs from the registration system and not manage an ID Domain at all. They would typically only know about a small population of the IDs in the ID Domain at any point in time. Examples of these could be Laboratory Information Systems (LIS), Radiology Information Systems (RIS), scheduling systems, monitoring systems, financial systems, pharmacy systems, etc. This represents the largest number of systems in a healthcare enterprise as they are doing specific functions that use the IDs as opposed to managing the IDs and ID information.

The first four interfaces (ProfilesAccess, SequentialAccess, IdentityAccess, IdentifyPerson) on this component are of general use that you may find on any type of component. They are the only interfaces that an ancillary system (that uses the medical record number) would typically implement (i.e. unless they maintained their own ID Domain). These basic interfaces give query access to lookup persons by matching against a passed in profile, and query access to find the profile information about persons, given an ID.

There is at most one component per ID Domain that implements the IdMgr interface. Other components in the ID Domain use notification events from the ID Domain manager to keep their cached values of the traits they are interested in kept up to date.

A PIDS implementation may implement security restrictions that prevent access to IDs and/or certain traits for an ID by not returning information on it. This can make it appear the person is not known by the service or certain traits about the person are not known to the service. Alternatively the service may raise security restriction exceptions.

#### **domain\_name**

This is the name of the ID Domain the component resides in. Even stand alone systems that implement their own ID Domain must create an ID Domain name and set this value. The entity responsible for naming the ID Domain must follow the rules defined in the NamingAuthority module for selecting names uniquely for each ID Domain created. The DomainName is invariant over time.

This is a globally unique, permanent name over the space of PIDS instances. It is expected that over time organizations will need to federate in ways they do not presently anticipate. It must never be necessary to modify a DomainName once it has been put into service.

#### **component\_name**

Each implementation instance of the component must create a unique name for the component. The name chosen does not need any particular meaning. The unique names make it possible for clients traversing a graph of components to recognize components they have encountered before. This way they can detect cycles in the graph and not visit a component multiple times when they don't need to. This part of the component\_name is invariant over time. The version within the component name can change over time according to the rules set out by the ComponentVersion definition.

If there are two or more objects with the same component name they must be replicas of each other with identical functionality. The mechanism used to maintain consistency between the replicas is implementation dependent and is not exposed as standard interfaces.

#### **supported\_traits**

This indicates the set of traits supported by a component. Each trait is indicated whether it is mandatory, read only and searchable. There is no ordering of the traits assumed.

The set of traits are static over time (see ComponentVersion).

#### **naming\_context, trader\_components**

These attributes are used to provide federation of PIDS components. Either one or both may be NULL as they are optional capabilities for any PIDS component. See the section on Naming/Trader Interoperability for more details.

#### **identify\_person**

This interface provides a way to identify a person (find a potential ID) from the traits known about them.

#### **profile\_access**

This is the main interface for accessing the traits associated with an ID.

#### **sequential\_access**

This interface provides mechanisms for scrolling forward and backward through the set of IDs the component knows about. This interface is mostly useful for components that have a small number of IDs that can easily be scanned by a person.

#### **identity\_access**

This interface provides access similar to `profile_access` but with the client first accessing a separate object per ID. This is needed to simplify certain security constraints by some implementations.

#### **id\_mgr**

Only one component within an ID Domain implements this interface. If there are multiple actual systems implementing this interface they should look like a single component. This is the interface implemented by a registration system. Other components within an ID Domain sets this attribute to point to the ID manager of the ID Domain, if one exists and they have a reference to it.

#### **correlation\_mgr**

This interface is only implemented by components that correlate over other ID Domains.

#### **event\_component**

The event component gives the ability to connect components together such that changes on one component are communicated to the other. For example IDs being created or deprecated and modifications of profiles.

### 3.3 *IdentifyPerson interface*

```
// -----  
// IdentifyPerson  
//  
interface IdentifyPerson :  
    IdentificationComponent  
{  
    void find_candidates(  
        in TraitSelectorSeq profile_selector,  
        in IdStateSeq states_of_interest,  
        in float confidence_threshold,  
        in unsigned long sequence_max,  
        in unsigned long iterator_max,  
        in SpecifiedTraits traits_requested,  
        out CandidateSeq returned_sequence,  
        out CandidateIterator returned_iterator )  
    raises (  
        TooMany,  
        UnknownTraits,  
        WrongTraitFormat,  
        CannotSearchOn,  
        DuplicateTraits,  
        InvalidStates,  
        InvalidWeight );  
};
```

This service defines the functionality for querying an ID Domain or individual system with a specified set of trait values and their weights.

#### **find\_candidates()**

Knowing some identifying information about a person (or group of people with common traits) a client can ask the service to find the candidate persons the service thinks may match those traits. The valid states that can be passed in by the client are Temporary, Permanent, and Deprecated.

The client can indicate the maximum number of candidates it wants passed back in the call using the `sequence_max` parameter. If the service matches more than that an iterator is created containing the rest. It is the responsibility of the client to either retrieve all candidates from the iterator or to call the `destroy()` method. The client can also specify the maximum number that should be returned in the iterator via the `iterator_max` parameter.

Candidate iterators require transient objects that could have significant memory management impact on a server. If a server does not support candidate iterators they return a NULL object reference for every invocation and clients can only obtain candidates that are returned as a sequence.

Using the `sequence_max` and `iterator_max` a client has the flexibility to do things like:

- get all candidates as a sequence,
- get all candidates via an iterator,
- get up to 30 candidates as a sequence and the rest up to 50 on an iterator,
- get up to 15 candidates as a sequence and all the rest via an iterator,
- etc.

The candidates returned by the operation have a confidence indication where the larger the number the better the match. No candidate is returned that has a confidence indicator that is less than the confidence threshold. The interpretation of the confidence threshold value is consistent with that of the confidence indicator. This indicates how well the stored profile for that person matched the passed in profile selector. The number of candidates returned could be zero if the matching engine does not find anything it feels matches close enough to return. Since there are so many algorithms this becomes a quality of service issue that is not standardized.

The range of values for the confidence indicator is 0.0-1.0 with 1.0 being the higher confidence (e.g. 100%). The exact semantics of the confidence indicator is determined by the service but some guidelines may help. The client should be able to compare two returned candidates and determine if they are of equal confidence or determine which is a higher confidence. If the returned candidates have different confidence values then they are returned in confidence order with the highest being returned first. If a service does not provide confidence determination they may return 0.0 as the confidence value the same for all candidates. In this way a client can know whether the service uses confidence values by whether the first candidate has a confidence of 0.0 or not.

The confidence value only has meaning relative to the single call it was returned from. There is no standard way to compare confidences returned from different services or from multiple calls to the same service.

Some matching engines may use discrete matching (such as the UNIX 'grep' utility) and others may use fuzzy matching (such as spell checkers or the use of phonic similarities). Even for discrete matching it is difficult to define exact semantics that could be applied to confidences since there are multiple traits and some traits contain multiple fields and the definitions would have to take into account how to combine the results from each field and the results from each trait to determine the confidence.

Matching engines map their matching capabilities to the defined confidence semantics as they see fit. How close they meet the expectations of the user is determined a Quality of Service (QoS) issue that is not standardized.

The 'weight' field is a **hint** from the client on how much preference it thinks the server should give to each profile passed in for determining matching confidences. It may be thought of as the confidence the client has in each trait of the profile selector.

The weight must be between 0.0 and 1.0. The weights are relative measures such that an exact match on a trait with weight = 0.5 results in twice the increase in confidence than an exact match on a trait with weight = 0.25.

A server implementation can ignore the weight field **hint** if it chooses. Using it is considered a Quality of Service (QOS) issue that is not standardized. The semantics here are just a way to define the measuring stick but does not require implementations to prove adherence. This is in part because fuzzy semantics can not be measured.

The `traits_requested` parameter indicates the traits to be returned from for every candidate, if the candidate has values for that trait.

### 3.4 ProfileAccess interface

```
// -----  
// ProfileAccess  
//  
interface ProfileAccess :  
    IdentificationComponent  
{  
    TraitNameSeq get_traits_known(  
        in PersonId id )  
        raises (  
            InvalidId );  
  
    Profile get_profile(  
        in PersonId id,  
        in SpecifiedTraits traits_requested )  
        raises (  
            InvalidId,  
            UnknownTraits,  
            DuplicateTraits );  
  
    TaggedProfileSeq get_profile_list(  
        in PersonIdSeq ids,  
        in SpecifiedTraits traits_requested )  
        raises (  
            TooMany,  
            InvalidIds,  
            DuplicateIds,  
            UnknownTraits,  
            DuplicateTraits );  
  
    TaggedProfileSeq get_deactivated_profile_list(  
        in PersonIdSeq ids,  
        in SpecifiedTraits traits_requested )  
        raises (  
            NotImplemented,  
            InvalidIds,  
            DuplicateIds,  
            UnknownTraits,  
            DuplicateTraits );  
  
    void update_and_clear_traits(  
        in ProfileUpdateSeq profile_update_spec )  
        raises (  
            InvalidIds,  
            DuplicateIds,  
            NotImplemented,  
            MultipleTraits );  
  
    IdInfoSeq get_id_info(  
        in PersonIdSeq ids )  
        raises (  

```

```

        TooMany,
        DuplicateIds );
};

```

ProfilesAccess is the most basic interface used to get identity information for persons. This service provides the simplest set of functionality for any PIDS service.

#### `get_traits_known()`

This operation returns the set of Traits known about a person by the service.

#### `get_profile(), get_profile_list()`

The `get_profile()` operation returns the profile or subset of the profile that the component knows about the person. The passed in traits indicate what subset of the profile that is being sought by the client.

The `get_profile_list()` is a shorthand mechanism for getting profiles for more than one ID at a time. This can be much more efficient than `get_profile()` for getting profiles on a lot of IDs since only one network round trip is required for all IDs as opposed to one per ID. The results are returned with exactly one value for each ID passed in. The results are tagged with the ID.

Both operations raise the `InvalidId(s)` exceptions if the IDs are not Permanent or Temporary.

Access rights to profiles are determined by the CORBAsec mechanisms (i.e., the user may not be able to access them).

#### `get_deprecated_profile()`

This is a special operation to get the profile of deprecated IDs. The ID passed in must be in the Deprecated state and known by the component or the `InvalidId` exception is raised. A service may choose to not keep the old profiles around in which case they raise the `NotImplemented` exception. The results are returned with exactly one value for each ID passed in. The results are tagged with the ID.

#### `update_and_clear_traits()`

This operation is used to modify the profile of already existing IDs. The structures passed in specify which traits in the profile to be cleared and which to change or add. Traits not mentioned to be cleared or changed remain the value they were before the call. If a PIDS component logs an audit trail of profile changes this operation will cause an event to be logged.

If a trait for an ID is listed to be cleared and changed the `ModifyOrDelete` exception is raised. If the passed in ID is not in the Temporary or Permanent `IdState` the `InvalidIds` exception is raised. If all supported traits are readonly the service raises the `NotImplemented` exception.

#### `get_id_info()`

This operation returns the current `IdStateInfo` for each ID passed in. The results are returned with exactly one value for each ID passed in. The results are tagged with the ID.

### 3.5 *SequentialAccess interface*

```

// -----
// SequentialAccess
//
interface SequentialAccess :
    IdentificationComponent

```

```

{
    unsigned long id_count_per_state(
        in IdStateSeq states_of_interest )
        raises (
            InvalidStates );

    TaggedProfileSeq get_all_ids_by_state(
        in SpecifiedTraits traits_requested,
        in IdStateSeq states_of_interest )
        raises (
            TooMany,
            UnknownTraits,
            DuplicateTraits,
            InvalidStates );

    TaggedProfileSeq get_first_ids(
        in unsigned long how_many,
        in IdStateSeq states_of_interest,
        in SpecifiedTraits traits_requested )
        raises (
            TooMany,
            UnknownTraits,
            DuplicateTraits,
            InvalidStates );

    TaggedProfileSeq get_last_ids(
        in unsigned long how_many,
        in IdStateSeq states_of_interest,
        in SpecifiedTraits traits_requested )
        raises (
            TooMany,
            UnknownTraits,
            DuplicateTraits,
            InvalidStates );

    TaggedProfileSeq get_next_ids(
        in PersonId reference_id,
        in unsigned long how_many,
        in IdStateSeq states_of_interest,
        in SpecifiedTraits traits_requested )
        raises (
            TooMany,
            InvalidId,
            UnknownTraits,
            DuplicateTraits,
            InvalidStates );

    TaggedProfileSeq get_previous_ids(
        in PersonId reference_id,
        in unsigned long how_many,
        in IdStateSeq states_of_interest,
        in SpecifiedTraits traits_requested )
        raises (
            TooMany,
            InvalidId,
            UnknownTraits,
            DuplicateTraits,
            InvalidStates );
};

```

#### `id_count_per_state()`

This operation indicates the number of IDs, having one of the ID states passed in, that the component knows of. The value could be zero at any point in time or it could be very large. Hospital ancillary systems would likely have from less than a dozen up to a hundred or many thousands. The valid states that can be passed in are Temporary, Permanent, and Deprecated.

`get_all_ids_by_state()`

This operation returns profiles for all patients the service knows about that match one of the states passed in. The returned profiles only contain the traits indicated by the passed in parameter, if they exist for the ID. The valid states that can be passed in are Temporary, Permanent, and Deprecated.

`get_first_ids(), get_last_ids(), get_next_ids(), get_previous_ids()`

These are the operations that provide sequential access to all the IDs known by the system/service. The service must have a consistent way to order the IDs. The ordering may be different for each set of traits asked for but for a particular set of traits the ordering is the same for all these operations. The ordering mechanism is implementation dependent and hidden behind the interface. The valid states that can be passed in are Temporary, Permanent, and Deprecated.

A client can request the profiles for a number of persons at the beginning and end of the ordered list by using the `get_first_ids()` and `get_last_ids()` respectively. The client passes in the number of IDs and profiles wanted.

A client can also request a number of profiles that are either after or before a particular ID via the `get_next_ids()` and `get_previous_ids()`. The ID and profile for the ID passed in is not returned.

Using these four operations a client can scroll forward or backward through the set of IDs, known by the service, a page at a time. The number of profiles returned may be smaller than that requested if the number of profiles held by the service is smaller than that requested.

### 3.6 *IdentityAccess interface*

```
// -----  
// IdentityAccess  
//  
interface IdentityAccess :  
    IdentificationComponent  
{  
    Identity get_identity_object(  
        in PersonId id )  
        raises (  
            InvalidId );  
  
    IdentitySeq get_identity_objects(  
        in PersonIdSeq ids )  
        raises (  
            InvalidIds );  
};  
  
interface Identity  
{  
    readonly attribute IdentificationComponent source_component;  
    readonly attribute IdInfo id_info;  
    readonly attribute TraitNameSeq traits_with_values;  
    readonly attribute long trait_value_count;  
  
    Trait get_trait(  
        in TraitName trait_requested )  
        raises (  
            UnknownTraits );  
  
    Profile get_profile(  
        in SpecifiedTraits traits_requested )  
        raises (  
            UnknownTraits,  
            DuplicateTraits );  
};
```

```

Profile get_deactivated_profile(
    in SpecifiedTraits traits_requested )
    raises (
        NotImplemented,
        UnknownTraits,
        DuplicateTraits );

void update_and_clear_traits(
    in ProfileUpdate profile_update_spec )
    raises (
        NotImplemented,
        UnknownTraits,
        WrongTraitFormat,
        ModifyOrDelete,
        ReadOnlyTraits,
        CannotRemove,
        DuplicateTraits );

void done();
};

```

The Identity object can provide individual ID level access control via CORBA Security. Specific access control parameters might be available in some security policies applied to an Identity object with the specific object being security unaware. These policies can be set up to prevent access on an operation/attribute basis. The service must be security aware in order to implement security policies that vary by the parameters passed in (e.g. separate access control for each trait).

Policies governing access to the interfaces is determined by administrative controls out of the scope of this specification.

#### **get\_identity\_object()**

Returns an identity object that represents the ID passed in, assuming it is a valid ID and known by the component.

#### **get\_identity\_objects()**

Returns an identity object for each ID passed in, assuming they are valid IDs and known by the component.

#### **Identity interface**

The Identity interfaces provide a way access control can be applied at the ID. Instances of Identities are accessed from the IdentityAccess interface.

#### **source\_component**

This is an object reference back to the component that created the Identity object. This may be useful if the reference was passed to a third party that may need to get to other functionality of the component.

#### **id\_info**

The id\_info attribute contains the ID by which this person is known. This consists of the PersonId (simple name) and the IdState of the ID. If it has been merged it has the preferred ID as well.

#### **traits\_with\_values**

This read only attribute returns the list of trait names for traits with values for the ID.

#### **trait\_value\_count**

This attribute is the count of the number of traits with values for the ID.

`get_trait()`

This operation returns the value set on the specified trait or raises an exception if the trait is not known.

`get_profile()`

This operation returns the traits requested that have values set.

`get_deactivated_profile()`

This is a special operation to get the profile of deprecated IDs. The ID must be in the Deprecated state or the InvalidId exception is raised. A service may choose to not keep the old profiles around in which case they raise the NotImplemented exception.

`update_and_clear_traits()`

This operation updates the profile by clearing (deleting the values) for some traits and setting the values for others. The ones not mentioned are left the same as before the call. If all supported traits are readonly the service raises the NotImplemented exception.

`done()`

This is called when the client no longer needs to access the object. Since there could be many Identity objects created for a single PIDS component the proper use of this operation can help the service in its own memory management. The server may delete the object reference if the object is transient or may delete the object from memory for persistent objects. This does not preclude Identity objects from being implemented persistently but it does not guarantee they are persistent. It is a service implementation issue as to when transient Identity objects are deleted if done() is not called. Clients should be prepared to get another object reference from the IdentityAccess interface if the Identity object is no longer valid while they still need to use it.

### 3.7 *IdMgr interface*

```
// -----  
// IdMgr  
//  
  
interface IdMgr :  
    IdentificationComponent  
{  
    PersonIdSeq register_new_ids(  
        in ProfileSeq profiles_to_register )  
        raises (  
            ProfilesExist,  
            DuplicateProfiles,  
            MultipleTraits );  
  
    PersonIdSeq find_or_register_ids(  
        in ProfileSeq profiles_to_register )  
        raises (  
            DuplicateProfiles,  
            MultipleTraits );  
  
    void register_these_ids(  
        in TaggedProfileSeq profiles_to_register )  
        raises (  
            NotImplemented,  
            IdsExist,  
            DuplicateIds,  
            ProfilesExist,  
            DuplicateProfiles,  
            MultipleTraits );  
}
```

```

PersonIdSeq create_temporary_ids(
    in ProfileSeq profiles_to_register )
raises (
    MultipleTraits );

PersonIdSeq make_ids_permanent(
    in PersonIdSeq ids_to_modify )
raises (
    InvalidIds,
    DuplicateIds,
    RequiredTraits );

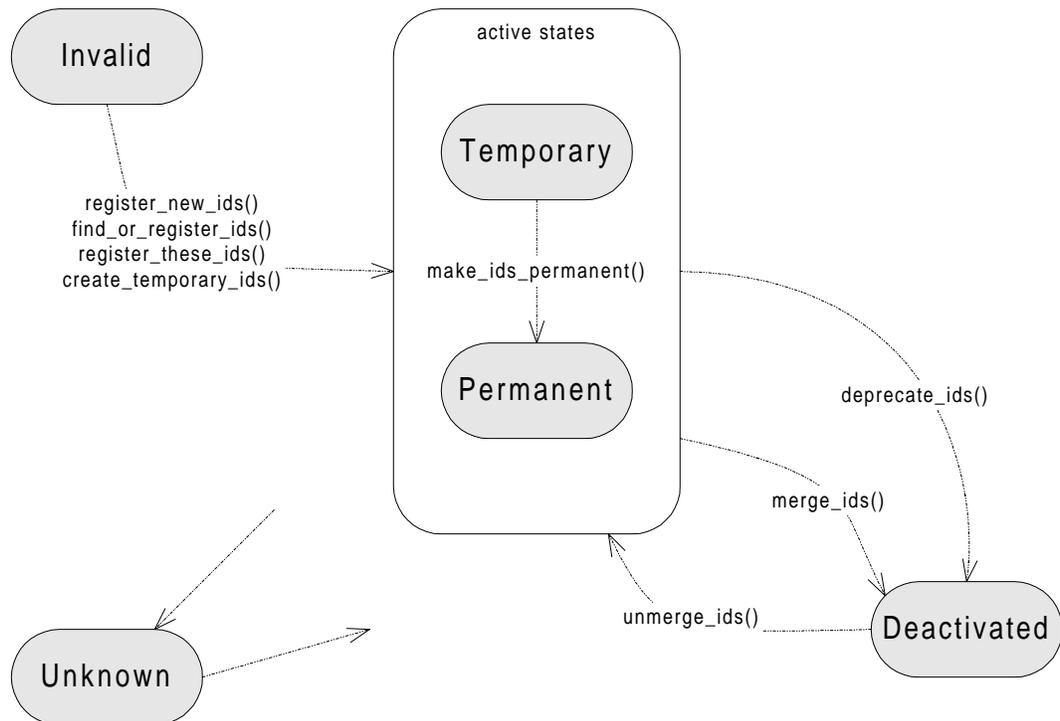
IdInfoSeq merge_ids(
    in MergeStructSeq ids_to_merge )
raises (
    InvalidIds,
    DuplicateIds );

IdInfoSeq unmerge_ids(
    in PersonIdSeq ids_to_unmerge )
raises (
    InvalidIds,
    DuplicateIds );

IdInfoSeq deprecate_ids(
    in PersonIdSeq ids_to_deprecate )
raises (
    InvalidIds,
    DuplicateIds );
};

```

IdMgr is an interface providing the core set of functionality for managing IDs in a single ID Domain. Specifically it can allocate a unique ID to a specific profile. All the operations on IdMgr are 'write' commands as opposed to 'read' commands. For this reason there may be more secure access control to prevent unwanted changes. The access control is managed by the CORBA security service and the implementation of the IdMgr service.



The transitions between IdStates are defined by operations on the IdMgr interface. Using the IdState as the values for a state machine pulls these operations together into a coherent set of semantics. The transitions shown are only suggestive as the policies of the service are implementation dependent. If a client calls a particular operation they are suggesting to the service that certain transitions should occur. The service may make the transition immediately, may wait for human review before making the transition or possibly never make the transition.

When a client needs to create a new ID the service returns it but does not guarantee whether it is Temporary or Permanent. Services are not required to use the Temporary ID state.

#### `register_new_ids()`

Generates new IDs in this ID Domain and binds the passed in profiles to them. The `register_new_ids()` operation establishes an association between IDs and profiles making them available for other component operations. Creates distinct, new IDs according to the rules and conventions for the ID Domain. These rules and conventions are implicit to the ID Domain and are not accessible via standard CORBA IDL.

By calling this operation the client is indicating it needs the IDs generated immediately and they are expected to become permanent. It is up to the service whether they are returned as permanent or as temporary IDs and require human interaction before becoming permanent.

#### `find_or_register_ids()`

This operation is used for generating IDs in unattended operation. An internal (hidden) value may indicate the confidence threshold that matching must attain in order to consider the person already has an ID in this Domain. Below this threshold a new ID is generated and the profile is bound to it.

This operation could produce two different resulting IdStates. It only changes the state of an ID if it creates a new ID. The resulting state of a newly created ID may be Temporary or Permanent which is implementation and site dependent. Some sites may choose to only produce temporary IDs automatically so they may be verified by a person before making them permanent.

#### `register_these_ids()`

This operation works similar to `register_new_ids()` except the client also indicates the value for the created IDs. The policy for generating the ID values may be hidden within the service and it would not allow clients to set the ID values. In this case the service raises the `NotImplemented` exception.

#### `create_temporary_ids()`

The `create_temporary_ids()` operation creates new IDs and indicates the client needs IDs that may be Temporary. A PIDS implementation may return a Permanent ID or Temporary ID. The mandatory traits are **not** required to create a temporary ID. This operation will still create a new ID even if an existing profile matches the one passed in.

#### `make_ids_permanent()`

Temporary IDs may be made permanent by the `make_ids_permanent()` operation. For some services the permanent IDs returned may be the same as the one passed in. Other services may want to use a separate part of the ID name space for temporary IDs so they can tell they are temporary by looking at them. For example a service may precede every temporary ID by the letters 'TEMP-'.

#### `merge_ids()`

If a person is found to have more than one ID in an ID Domain all but one of them can be merged into the preferred ID by calling the `merge_ids()` operation with each duplicate ID. The preferred ID will be unchanged except if it supports the `PersonIdTraits::MergedIds` trait, it will be modified

to reference the ID being merged to it. The merged ID will have its `IdState` set to `Deactivated` (possibly after human review) and have the `preferred_id` field on its `IdInfo` set to the preferred ID. The preferred ID will have its `MergedIds` trait set with the merged ID in it.

The `IdInfoSeq` returned has the `IdInfo` for each of the target IDs passed in to the operation.

#### `unmerge_ids()`

If an ID that has been merged with another is found later to represent a different person it may be unmerge with this operation. If the `MergedIds` trait is supported the profile for the ID this was merged with is changed to no longer reference this ID.

If the passed in ID(s) is(are) not merged (in the `Deactivated` state and `preferred_id` set to something other than "") the `InvalidIds` exception is raised.

The `IdInfoSeq` returned has the `IdInfo` for each of the target IDs passed in to the operation.

#### `deprecate_ids()`

Once an ID is expected to never be used again it may be retired from service by the `deprecate_ids()` operation which is a request to change their `IdState` to `Deactivated`. A service may either remove the profile for the ID or leave it intact for historical purposes but it can never be changed. The profile can be accessed only with the special operation `get_deprecated_profiles()`.

The `IdInfoSeq` returned has the `IdInfo` for each of the target IDs passed in to the operation.

### 3.8 *CorrelationMgr interface*

```
// -----  
// CorrelationMgr  
//  
interface CorrelationMgr :  
    IdentificationComponent  
{  
    readonly attribute DomainNameSeq source_domains;  
  
    void load_profiles(  
        in QualifiedTaggedProfileSeq tagged_profiles )  
        raises (   
            UnknownTraits,  
            WrongTraitFormat,  
            DomainsNotKnown );  
  
    QualifiedPersonIdSeq get_corresponding_ids(  
        in QualifiedPersonId from_id,  
        in DomainNameSeq to_domains )  
        raises (   
            DomainsNotKnown,  
            IdsNotKnown );  
};
```

The `PersonIdTraits::CorrelatedIds` trait is a special trait that is only used by components that implement the `CorrelationMgr` interface. The `CorrelationMgr` is responsible for setting this trait to reference all source IDs that are correlated to the same ID in the correlation Domain. This must be consistent with the information obtained via the `get_corresponding_ids()` operation.

The `CorrelationMgr` interface does not provide operations to support manual correlation or retrospective verification of unattended correlation. These capabilities may be addressed in a later RFP.

**source\_domains**

This read only attribute contains a list of source ID Domains that profiles may be loaded from and IDs may be mapped to or from.

**load\_profiles()**

This operation causes the profiles to be loaded into the Correlating ID Domain, from the specified source ID Domains. It is an implementation decision whether correlation is performed immediately.

**get\_corresponding\_ids()**

This operation returns the IDs in the destination ID Domains that correspond to the ID passed in.

## 4. NamingAuthority module

---

```
//File: NamingAuthority.idl

#ifndef _NAMING_AUTHORITY_IDL_
#define _NAMING_AUTHORITY_IDL_

#include <orb.idl>

#pragma prefix "omg.org "

module NamingAuthority
{
    enum RegistrationAuthority {
        OTHER,
        ISO,
        DNS,
        IDL,
        DCE };

    typedef string NamingEntity;

    struct AuthorityId {
        RegistrationAuthority authority;
        NamingEntity naming_entity;
    };
    typedef string AuthorityIdStr;

    typedef string LocalName;
    struct QualifiedName {
        AuthorityId authority_id;
        LocalName local_name;
    };
    typedef string QualifiedNameStr;

    exception InvalidInput {};

    interface translation_library
    {
        AuthorityIdStr authority_to_str(
            in AuthorityId authority )
            raises(
                InvalidInput );

        AuthorityId str_to_authority(
            in AuthorityIdStr authority_str )
            raises(
                InvalidInput );

        QualifiedNameStr qualified_name_to_str(
            in QualifiedName qualified_name )
            raises(
                InvalidInput );

        QualifiedName str_to_qualified_name(
            in QualifiedNameStr qualified_name_str )
            raises(
                InvalidInput );
    };
};

#endif // _NAMING_AUTHORITY_IDL_
```

The NamingAuthority module provides a means of giving globally unique names to name spaces and hence the names within those name spaces. The fundamental need is the ability to compare two names for equality. If they are equal they are known to represent the same entity, concept, or thing. This is needed when independent entities are generating names that may get compared for equality. However the reverse is not guaranteed to be true. That is an entity may have several names.

The authority for the name space may derive from several different types of roots, the choice of which depends upon the user requirements as each root has different qualities of management and uniqueness. The various root types are defined below.

**#pragma prefix "org/omg"**

In order to prevent name pollution and name clashing of IDL types this module (and all modules defined in this specification) uses the pragma prefix that is the reverse of the OMG's DNS name.

#### **RegistrationAuthority**

Identifies the root of the name space authority. An entity (e.g. person or organization) may be registered with multiple different roots (RegistrationAuthorities) and be able to assign names and other name spaces within each root. These may be used for the same or for different needs. For this reason there is no guarantee of any equality in the different name spaces managed by an entity. There are currently no means available to determine whether a given authority in an ISO hierarchy is the same authority as one specified in a DNS hierarchy.

**OTHER** This form of a naming authority should be used sparingly, and only in experimental or localized situations or special purposes. It is the responsibility of the implementing institution to guarantee uniqueness within the names themselves, and there is no uniqueness guarantee outside of the source institution. Services that define default naming authorities (and possibly also names) may also use the Other root to fore go long AuthorityIds. In this case the specification of the service must name AuthorityIds that may be expected with the Other root and still maintain name space integrity for that service.

**ISO** International Standards Organization<sup>A</sup>. The ISO specifies a registration hierarchy, identified by a series of named/numbered nodes. Many of the coding schemes used in the medical environment are or can be registered within the ISO naming tree. The ISO root form is one of the recommended forms when the naming authority is internationally recognized, such as international coding schemes, or when the authority is to be used across two or more different enterprises. ISO provides for the recording of a responsible person and address for each node in the authority hierarchy.

**DNS** Domain Name Services<sup>B</sup>. Internet domains are recorded with a central, global registration authority. Sub hierarchies within the domains are then maintained locally by the registered organization or person. The DNS form is recommended as an alternative to the ISO naming tree when the specific naming authority needs identity and uniqueness, but is not in an ISO registration. By using this common characteristic of many organizations it gives the ability to create globally unique name spaces and names without the need to register as an ISO name authority. It is up to the organization itself to maintain the integrity of the name space(s) (e.g. not reusing names or name spaces).

---

<sup>A</sup> ISO/IEC 8824-1 (1994) Information Technology - Abstract Syntax Notation One (ASN.1) - Specification of Basic Notation.

<sup>B</sup> P. Mockapetris, " Domain Names - Concepts and Facilities", RFC 1034, Information Sciences Institute, November 1987.

**IDL** The OMG Interface Repository<sup>C</sup>. The CORBA Architecture specifies a means of uniquely identifying entities within the interface repository, via the use of a *RepositoryId*. CORBA repository id's may be in either the OMG IDL format, the DCE UUID format or the LOCAL format. Within this specification, the "IDL" root refers only to the IDL format. The DCE format may be represented within the DCE root and the Local format within the Other root. The IDL authority may prove very useful when registering CORBA/IDL specific objects such as value sets, interface specifications, etc. It should be noted that OMG does not currently manage the repository name space in any rigorous fashion, and it is quite possible that two different developers may arrive at exactly the same repository ID for entirely different entities. For this reason some people give the repository ID a prefix that consists of their reverse DNS that is '/' separated instead of '.' separated. This root type may be very useful when the names within the name space are defined in IDL. For example it could be the RepositoryId for an enumerated type or a module that has constant integers or strings defined for each name within the name space.

**DCE** The Distributed Computing Environment<sup>D</sup>. While they don't actually register coding schemes or other entities, they do provide a means of generating a globally unique 128 bit ID, called a Universally Unique ID (UUID). This UUID may be used to guarantee the uniqueness of a name space in situations where it is not necessary for the identity of the authority to be known outside of the specific implementation.

### **NamingEntity**

Identifies a specific name in the syntax and format specified by the corresponding registration authority. The various naming authorities tend to provide a fair amount of leeway as far as the actual format of the registered names. As there may be situations where the full semantics of a specific authority's name comparison will not be available to an application, we have chosen to select a specific subset of the syntax of each representation. The intention is to be able to determine whether two registered entities are identical or not solely through the use of string comparison. The specific name formats are described below:

**OTHER** An arbitrary string, syntax undefined except locally by a specific service specification and/or by particular implementations and installations. The colon ":" character is illegal to use as it is reserved as a separator of components in the stringified version of QualifiedName.

**ISO** The name should be represented using the *NameForm* of the *ObjectIdentifierValue* as specified in ISO/IEC Recommendation 8824-1. Each name component should be separated by a single space.

Example: "joint-iso-ccitt specification characterString"

**DNS** The domain name and path in the form mandated in RFC 1034. The path name is represented as a dot separated tree which traverses up the hierarchy. Since DNS names are not case sensitive only lower case letter should be used such that simple string comparisons can determine equality. However it is OK to use case insensitive comparisons as well.

Example: "pidsserv.slc.mmm.com"

**IDL** The OMG RepositoryId format specified in the CORBA Architecture V2.0 manual, in the form: "<node>/<node>/.../<node>". The "IDL:" prefix and the version number suffix should NOT be used for the NamingEntity. The "IDL:" prefix is prepended to create the AuthorityIdStr.

Example: "CosNaming/NamingContext/NotFoundReason" is the NamingEntity for:

```
module CosNaming {
```

---

<sup>C</sup> The Common Object Request Broker: Architecture and Specification. Revision 2.0. OMG, July 1995.

<sup>D</sup> DCE 1.1 : Remote Procedure Call. OpenGroup Document Number C706, August 1997.

```

...
interface NamingContext {
    ...
    enum NotFoundReason { ... };
    ...
};
};

```

DCE The UUID in the external form <nnnnnnnn-nnnn-nnnn-nnnn-nnnnnnnnnnn>, where <n> represents one of the digits 0-9 and the characters A-F. The alpha characters should all be upper case.

Example: “6132A880-9A34-1182-A20A-AF30CF7A0000”

**AuthorityId, AuthorityIdStr**

The combination of a Registration Authority and Naming Entity, which identifies a specific naming authority. In situations where a given naming entity may have more than one naming authority, it should be agreed upon in advance which of the specific names for the entity is to be used. This specification makes no guarantees about the ability to recognize, for example, that an authority in the ISO structure is identical to an authority within the IDL structure.

The string version (AuthorityIdStr) is useful for situations where unique names are required in a string format. The string is created as <stringified RegistrationAuthority>:<NamingEntity>.

**LocalName, QualifiedName, QualifiedNameStr**

A local name is a name within (relative to) a namespace. It is simply a string representation.

A QualifiedName is a globally unique name for an entity by the fact that it carries the naming AuthorityId of the name space and the LocalName within that name space.

The QualifiedNameStr is a stringified version of the QualifiedName. The format of the string is <stringified RegistrationAuthority>:<NamingEntity>:<LocalName>. Notice that even though the colon character “/” can not be used within the name of a NamingEntity it can be used within the LocalName. The following table summarizes the format for QualifiedNameStr. Columns 1-3 give the format for an AuthorityIdStr.

Registration Authority	(1) Stringified Registration Authority	(2) RA-NE Delimiter	(3) NamingEntity Format	(4) NE-LN Delimiter	(5) LocalName Format
<b>OTHER</b>	“”	“:” optional	<no ‘/’>	“/” optional	<no ‘/’>
<b>ISO</b>	“ISO”	“:”	<use ISO rules>	“/”	<any characters>
<b>DNS</b>	“DNS”	“:”	<use DNS rules>	“/”	<any characters>
<b>IDL</b>	“IDL”	“:”	<use IDL rules>	“/”	<no ‘/’>
<b>DCE</b>	“DCE”	“:”	<use DCE rules>	“/”	<any characters>

The definitions for type OTHER are defined to allow using a QualifiedNameStr format in contexts where an IDL ‘string’ is currently used. A normal IDL string is a QualifiedNameStr with no RegistrationAuthority and no NamingEntity. The limitation is that any normal IDL strings that

start with one of the RegistrationAuthority strings can not be mapped into the QualifiedNameStr since they would be interpreted by the rules in this module.

The string for the 'OTHER' type of RegistrationAuthority being a blank string ("") makes it easier for locally defined names to be usable with no requirements on the format except they can not start with one of the strings reserved for the other RegistrationAuthority types. The ':' delimiter is optional for type OTHER. If the NamingEntity is "" for type OTHER then the '/' delimiter is also optional.

### Exceptions

#### **InvalidInput**

This exception is raised when the input parameter for the TranslationLibrary interface operations is determined to be of an invalid format.

### TranslationLibrary interface

This interface is meant to be a local library for translating between the structured version and stringified version of AuthorityIds and QualifiedNames.

`authority_to_str, str_to_authority, qualified_name_to_str, str_to_qualified_name()`

Each of these operations take either a structured version or stringified version of a type and return the opposite. The data content returned is the same as that passed in. Only the representation of the data is changed.

## 5. Traits

---

The definitions in the PersonIdService module define traits as name/value pairs with the name being a string (in the QualifiedNameStr format) and the value being an 'any'. This is done to provide complete flexibility in defining traits. The traits used to help identify a person may be of many types and formats. Many of them are represented as strings but others can be structured types and even multimedia types. In addition there are many inconsistent standards even for the more common traits like a person's name.

This specification is written to allow the use of different traits and different formats of traits to be used as needed. A specific implementation of PIDS may use the traits that meet their needs. The formats of even the common traits may need to be different depending on the technical and political environment in which it is deployed. It may need to work with other information services that have a predefined data format that the implementation decides to use. The data formats used for traits may be required by legal means as well.

While this specification does not limit what traits can be used with PIDS it does specify a set that can be well known and used by clients and servers if they choose. The set consists of data formats from a couple of well known industry standards as well as some special traits defined just for PIDS. The well known standards are the Health Level Seven (HL7) and vCard™.

If there are environments where other standard sets of traits are needed they could be standardized via the OMG's processes. These could be areas that require interoperable solutions to existing standards or legacy systems for which PIDS interfaces need to be provided.

The following sections define the three modules that contain the default set of traits defined by PIDS.

### References:

Health Level Seven (HL7) Standard, version 2.3, 1997.

Versit Consortium. VCard - The Electronic Business Card, version 2.1, September 18, 1996.

American Dental Association, "Proposed ANSI/ADA Specification No 1000 Standard Clinical Data Architecture for the Structure and Content of a Computer-based Patient Record. Part 1000.1 Individual Identification", ASC MD156, August 1997.

ASTM E1239, Standard Guide for Description of Reservation/Registration-Admission, Discharge, Transfer (R-ADT) Systems for Automated Patient Care Information Systems, Committee E-31 on Computerized Systems, Subcommittee E31.19 on Vocabulary for Computer Based Patient Records-Content and Structure, West Conshohocken, PA: ASTM, January 15, 1994.

ASTM E1714. Guide for the Properties of a Universal Healthcare Identifier. Committee E-31 on Computerized Systems, Subcommittee E31.12 on Medical Records. West Conshohocken, PA: ASTM, Aug. 15, 1995.

ASTM E1385. Guide for global environments for RADT.

ASTM E1715. RADT data model.

CPRI, 1996a. Action Plan for Implementing a Universal Patient Identifier, Draft Version 1.0. Schaumburg, IL: Computer-based Patient Record Institute, May.

ICSI, 1995. Data Communication Standard: Patient Identifier, September 1995.

HIN, 1997. The Essential Medical Data Set (EMDS), January 8, 1997.

UK National Health System Common Data Model.

## 5.1 *PersonIdTraits module*

```
//File: PersonIdTraits.idl

#ifndef _PERSON_ID_TRAITS_IDL_
#define _PERSON_ID_TRAITS_IDL_

#include <orb.idl>
#include <PersonIdService.idl>

#pragma prefix "omg.org"

module PersonIdTraits
{
    const PersonIdService::TraitName NULL_TRAIT = "";
    typedef any NullTraitType; // set to tk_null

    const PersonIdService::TraitName INTERNAL_ID = "PIDS/InternalId";
    typedef PersonIdService::PersonId InternalIdType;

    const PersonIdService::TraitName MERGED_IDS = "PIDS/MergedIds";
    typedef PersonIdService::PersonIdSeq MergedIdsType;

    const PersonIdService::TraitName DUPLICATE_IDS = "PIDS/DuplicateIds";
    typedef PersonIdService::PersonIdSeq DuplicateIdsType;

    const PersonIdService::TraitName CORRELATED_IDS = "PIDS/CorrelatedIds";
    typedef PersonIdService::QualifiedPersonIdSeq CorrelatedIdsType;

    const PersonIdService::TraitName EXTERNAL_IDS = "PIDS/ExternalIds";
    typedef PersonIdService::QualifiedPersonIdSeq ExternalIdsType;

    const PersonIdService::TraitName NATIONAL_HEALTH_IDS =
"PIDS/NationalHealthId";
    struct NationalHealthIdStruct {
        string country_code;
        PersonIdService::PersonId id;
    };
    typedef sequence< NationalHealthIdStruct > NationalHealthIdStructSeq;
    typedef NationalHealthIdStructSeq NationalHealthIdsType;
};

#endif // _PERSON_ID_TRAITS_IDL_
```

The `PersonIdTraits` module contains definitions for traits that are needed but are not defined by the HL7 Version 2.3 or vCard Version 2.1 standards. Most of these are for defining other IDs a person may have where these IDs can be characterized in the format of this specification.

### **NULL\_TRAIT, NullTraitType**

This trait is used in places where a trait or trait name must be passed but no valid value exists.

### **INTERNAL\_ID, InternalIdType**

The `InternalId`, if used, is always set to the `PersonID` itself. It is specified in the `supported_traits` as read only and searchable but not mandatory. The reason for including it as a trait is to allow matching on it. This is the only way the ID can be matched on when only part of the ID is known. For example this could occur if a person's ID card has a tear or only part of the digits are remembered by the person.

#### **MERGED\_IDS, MergedIdsType**

This trait indicates the set of other IDs a person may have. The IDs listed are from the same ID Domain as the ID this trait is bound to. It has the special semantics that the IDs listed are in the Deprecated IdState and the preferred\_id on their IdInfo references the ID this trait is bound to. The setting of this trait is controlled by the PIDS implementation in order to maintain consistency with the back references from the IDs listed. There for this trait is always read only when supported by a PIDS.

#### **DUPLICATE\_IDS, DuplicateIdsType**

This trait indicates the set of other IDs a person may have. The IDs listed are from the same ID Domain as the ID this trait is bound to. This is more general purpose than MergedIds. The IDs listed in this trait do not have to be deprecated or merged.

#### **CORRELATED\_IDS, CorrelatedIdsType**

This trait indicates the set of other IDs a person may have. The IDs listed are from a different ID Domain than the ID Domain of the ID this trait is bound to. This trait has special semantics in that it refers to IDs that have been correlated by the CorrelationMgr. There for this trait is always read only when supported by a PIDS.

#### **EXTERNAL\_IDS, ExternalIdsType**

This trait indicates the set of other IDs a person may have. The IDs listed may be from a different ID Domain than the ID Domain of the ID this trait is bound to. This is more general purpose than CorrelatedIds. The IDs listed in this trait may be set by other mechanisms than automatic correlation.

#### **NATIONAL\_HEALTH\_IDS, NationalHealthIdsType**

The NationalHealthId trait is added in this module because the HL7 2.3 and vCard 2.1 standards do not have a corresponding trait. This trait is important in countries that have a national healthcare ID. The 'country\_code' field indicates the name of the country issuing the IDs. The ISO 3166 two letter country codes or the telephone numeric country code are to be used. The value for this trait consists of a sequence since it is possible for a person to have IDs in multiple countries or accidentally get multiple IDs in the same country.

## 5.2 *HL7Version2\_3 module*

```
//File: HL7Version2_3.idl

#ifndef _HL7_VERSION_2_3_IDL_
#define _HL7_VERSION_2_3_IDL_

#include <orb.idl>
#include <PersonIdService.idl>

#pragma prefix "omg.org"

module HL7Version2_3
{
    typedef PersonIdService::TraitName PIDFieldName;
    typedef string PIDFieldValue;

    const PIDFieldName PATIENT_NAME = "HL7/PatientName";
    const PIDFieldName MOTHERS_MAIDEN_NAME = "HL7/MothersMaidenName";
    const PIDFieldName DATE_TIME_OF_BIRTH = "HL7/DateTimeOfBirth";
    const PIDFieldName SEX = "HL7/Sex";
    const PIDFieldName PATIENT_ALIAS = "HL7/PatientAlias";
    const PIDFieldName RACE = "HL7/Race";
    const PIDFieldName PATIENT_ADDRESS = "HL7/PatientAddress";
}
```

```

const PIDFieldName COUNTY_CODE = "HL7/CountyCode";
const PIDFieldName PHONE_NUMER_HOME = "HL7/PhoneNumber_Home";
const PIDFieldName PHONE_NUMBER_BUSINESS = "HL7/PhoneNumber_Business";
const PIDFieldName PRIMARY_LANGUAGE = "HL7/PrimaryLanguage";
const PIDFieldName MARITAL_STATUS = "HL7/MaritalStatus";
const PIDFieldName RELIGION = "HL7/Religion";
const PIDFieldName PATIENT_ACCOUNT_NUMBER = "HL7/PatientAccountNumber";
const PIDFieldName SSN_NUMBER = "HL7/SSNNumber";
const PIDFieldName DRIVERS_LICENSE_NUMBER = "HL7/DriversLicenseNumber";
const PIDFieldName MOTHERS_IDENTIFIER = "HL7/MothersIdentifier";
const PIDFieldName ETHNIC_GROUP = "HL7/EthnicGroup";
const PIDFieldName BIRTH_PLACE = "HL7/BirthPlace";
const PIDFieldName MULTIPLE_BIRTH_INDICATOR =
"HL7/MultipleBirthIndicator";
const PIDFieldName BIRTH_ORDER = "HL7/BirthOrder";
const PIDFieldName CITIZENSHIP = "HL7/Citizenship";
const PIDFieldName VETERANS_MILITARY_STATUS =
"HL7/VeteransMilitaryStatus";
const PIDFieldName NATIONALITY = "HL7/Nationality";
const PIDFieldName PATIENT_DEATH_DATE_AND_TIME =
"HL7/PatientDeathDateandTime";
const PIDFieldName PATIENT_DEATH_INDICATOR = "HL7/PatientDeathIndicator";
};

#endif // _HL7_VERSION_2_3_IDL_

```

The HL7Version2\_3 module defines the standard trait names for using trait information from the Health Level Seven (HL7) standard. The traits correspond to fields 5-30 in the PID segment as defined in Version 2.3 of the HL7 standard. The values for the traits are strings (HL7Version2\_3::PIDFieldValue) formatted according to the rules specified by the HL7 standard.

The trait names (HL7Version2\_3::PIDFieldName) include the complete QualifiedNameStr which has a registration authority of OTHER and the authority ID of 'HL7'.

Many of these traits from the HL7 PID segment are demographic traits instead of identifying traits. The difference between demographic traits and identifying traits is often dependent on the environment in which they are being used. It was decided to include all of these PID segment traits and let the PIDS implementers choose the ones that are considered identifying traits for their service.

**Note** *HL7 is working on version 3.0 of its standard. Once it is finished the list of traits may be updated or a new list created via the standard OMG processes.*

### 5.3 vCardTraits module

```

//File: vCardVersion2_1.idl

#ifndef _V_CARD_VERSION_2_1_IDL_
#define _V_CARD_VERSION_2_1_IDL_

#include <orb.idl>
#include <PersonIdService.idl>

#pragma prefix "omg.org"

module vCardVersion2_1
{
    typedef PersonIdService::TraitName PropertyName;
    typedef string PropertyValue;

    const PropertyName FORMATTED_NAME = "vCard/FN";
    const PropertyName NAME = "vCard/N";

```

```

const PropertyName PHOTOGRAPH = "vCard/PHOTO";
const PropertyName BIRTHDAY = "vCard/BDAY";
const PropertyName ADDRESS = "vCard/ADR";
const PropertyName HOME_ADDRESS = "vCard/ADR;HOME";
const PropertyName WORK_ADDRESS = "vCard/ADR;WORK";
const PropertyName TELEPHONE = "vCard/TEL";
const PropertyName PREFERRED_TELEPHONE = "vCard/TEL;PREF";
const PropertyName HOME_TELEPHONE = "vCard/TEL;HOME";
const PropertyName WORK_TELEPHONE = "vCard/TEL;WORK";
const PropertyName VOICE_TELEPHONE = "vCard/TEL;VOICE";
const PropertyName FAX_TELEPHONE = "vCard/TEL;FAX";
const PropertyName MESSAGE_TELEPHONE = "vCard/TEL;MSG";
const PropertyName CELLULAR_TELEPHONE = "vCard/TEL;CELL";
const PropertyName BULLETIN_BOARD_TELEPHONE = "vCard/TEL;BBS";
const PropertyName MODEM_TELEPHONE = "vCard/TEL;MODEM";
const PropertyName CAR_TELEPHONE = "vCard/TEL;CAR";
const PropertyName ELECTRONIC_MAIL = "vCard/EMAIL";
const PropertyName GEOGRAPHIC_POSITION = "vCard/GEO";
const PropertyName TITLE = "vCard/TITLE";
const PropertyName ORGANIZATION = "vCard/ORG";
const PropertyName SOUND_ANNOTATION = "vCard/SOUND";
const PropertyName UNIFORM_RESOURCE_LOCATOR = "vCard/URL";
};

#endif // _V_CARD_VERSION_2_1_IDL_

```

The vCardVersion2\_1 module defines the standard trait names for using trait information from the vCard standard. The traits correspond to the properties as defined in Version 2.1 of the vCard standard. The values for the traits are strings (vCardVersion2\_1::PropertyValue) formatted according to the rules specified by the vCard standard.

Only the properties in vCard that relate to this specification are included. The appropriate properties with parameters are also included.

The trait names (vCardVersion2\_1::PropertyName) include the complete QualifiedNameStr which has a registration authority of OTHER and the authority ID of 'vCard'.

## 6. Naming/Trader Interoperation

---

The Naming and Trader Services are used in multiple ways by the PIDS specification. One way is to define Trader Service Types explicitly for PIDS and some naming conventions for PIDS related objects. The other is that every PIDS component has the ability to expose Naming and Trader services that are local to the component in order to provide explicate federation of PIDS implementations.

### Naming Service

The following names are to be used for the 'kind' field in CosNaming::NameComponent.

- **'IdentificationComponent'** - Generic type for a PIDS component.
- **'Simple PIDS'** - A PIDS component that meets the conformance class of the same name.
- **'Sequential Access PIDS'** - A PIDS component that meets the conformance class of the same name.
- **'Identity Access PIDS'** - A PIDS component that meets the conformance class of the same name.
- **'ID Domain Manager PIDS'** - A PIDS component that meets the conformance class of the same name.
- **'Correlation PIDS'** - A PIDS component that meets the conformance class of the same name.

Naming contexts are freestanding nodes in the directory hierarchy unless they are associated with a particular system or component in some way. The use of the naming context as referenced from an IdentificationComponent has specific semantics. The naming context referenced from a PIDS component is the root of a naming tree relative to that component. It is a way for the component to publish references to other objects it is associated with or knows about. The naming directory it is the root for can be used for as a general naming service as well. The following directory names have special meaning for PIDS components.

- **'Source ID Domains'** - Components that do correlation over other ID Domains will have object references to PIDS components that manage those ID Domains (if they exist) in this directory. It is advised the sub ID Domains be referenced by their stringified ID Domain name as it provides a simple way to search for ID Domains when a trader component is not available. They may also be referenced by a logical name for easier identification by users and by their component name.
- **'Correlating ID Domains'** - Components that are being correlated by other components can publish references to these correlating components in this directory.
- **'ID Domain Components'** - PIDS components that manage an ID Domain it will have references to other PIDS components within the ID Domain in this directory. It is advised the ID Domain residents be referenced by their stringified component name as it provides a simple way to search for them when a trader component is not available. They may also be referenced by a logical name for easier identification by users.
- **'ID Using Services'** - This directory is a place for a component to have references to services it provides that use IDs from its ID Domain.

The following names for objects located in the root of the tree referenced by a PIDS component are also reserved.

- **'Trait Information'** - If this object exists it is an implementation of the Lexicon Query Service (LQS) that contains the traits supported by this PIDS component as concepts.

Trader components are typically stand alone services installed in an enterprise. The reference to the trader component from a PIDS component has special semantics. The trader component is used for searching for other PIDS components as well as other object services. The Trader referenced by a PIDS component knows of the PIDS service types defined in this specification.

### Trader Service

The following definitions are Service Types defined for PIDS components for use by the Trader Service.

```
service IdentificationComponent {
    interface IdentificationComponent;
    mandatory readonly property string domain_name;
    mandatory readonly property StringSeq interfaces_implemented;
    mandatory readonly property StringSeq conformance_classes;
    mandatory readonly property string component_name;
    mandatory readonly property string component_version;
    mandatory readonly property StringSeq supported_traits;
    mandatory readonly property StringSeq read_only_traits;
    mandatory readonly property StringSeq mandatory_traits;
    mandatory readonly property StringSeq searchable_traits;
    readonly property StringSeq source_domains;
};
```

Since all PIDS implement the IdentificationComponent only one Trader Service type is needed, which is also called 'IdentificationComponent'. The IdentificationComponent interface has attributes for the common characteristics for all PIDS. These are used as properties for the IdentificationComponent service type. One additional property is specified aswell which comes from an attribute from one of the derived interfaces. The stringified versions of the attributes are used for properties since the standard Trader constraint language does not provide a way to filter on user defined types.

### IdentificationComponent service

The interface type returned from the Trader Service for this service type is an IdentificationComponent. All except one of the properties are mandatory. These are found on all IdentificationComponent interface implementations.

#### **domain\_name**

The domain\_name property contains the information from the domain\_name attribute of the IdentificationComponent interface. It is formatted as specified for NamingAuthority::AuthorityIdStr.

#### **interfaces\_implemented**

This sequence contains the names of the interfaces the component has references to. This includes PIDS specific interfaces such as 'PersonIdService::ProfileAccess' and other interfaces such as 'CosNaming::NamingContext'. The names are fully qualified name which includes the module name.

#### **conformance\_classes**

This sequence contains the conformance classes the implementation supports. The strings are identical to the way they are spelled and capitalized in the definition of the conformance classes for PIDS.

#### **component\_name**

This property contains the 'the\_name' field of the component\_name attribute. It is a stringified version of the name in the format as a NamingAuthority::QualifiedNameStr.

**component\_version**

This property contains the 'the\_version' field from the component\_name attribute.

**supported\_traits**

This property contains a sequence with all of the TraitNames that are in the supported\_traits attribute.

**read\_only\_traits**

This property contains a sequence with all of the TraitNames that are in the supported\_traits attribute that have the read\_only field set to true.

**mandatory\_traits**

This property contains a sequence with all of the TraitNames that are in the supported\_traits attribute that have the mandatory field set to true.

**searchable\_traits**

This property contains a sequence with all of the TraitNames that are in the supported\_traits attribute that have the searchable field set to true.

**source\_domains**

This is the only optional property. It only applies to PIDS that implement the CorrelationMgr interface. The property contains the stringified DomainName for the source ID Domains being correlated.

## 7. Conformance Classes

---

The following interfaces are programmatic reference points for testing conformance. Conformance indicates implementing all of the attributes and operations for those interfaces and the specified behavior and raising exceptions as specified. For the operations that may raise the `PersonIdService::NotImplemented` exception a valid implementation of the operation may raise this exception. All other operations must perform as specified within this specification

- IdentificationComponent
- ProfileAccess
- IdentityAccess
- SequentialAccess
- IdentifyPerson
- IdMgr
- CorrelationMgr
- EventComponent

The following taxonomy is defined for specific conformance classes of PIDS service implementations. An implementation claiming conformance to any of these classes must conform to all of the interfaces specified for that class. An implementation may claim conformance to multiple conformance classes as long as it is conformant to each one it claims.

An implementation claiming conformance must also follow the rules in the NamingAuthority module for creating names that derive from those types.

All conformant implementations must support at least one trait as specified on their `supported_traits` attribute.

Implementations that use the traits defined in the `PersonIdTraits`, `HL7Version2_3` and `vCardVersion2_1` modules must maintain the semantics defined in those modules in order to be considered conformant.

An implementation claiming conformance must implement the semantics defined for IdentificationComponent such as consistency between the all interfaces implemented. They must also maintain the semantic mapping between attributes on the IdentificationComponent and the other interfaces implemented.

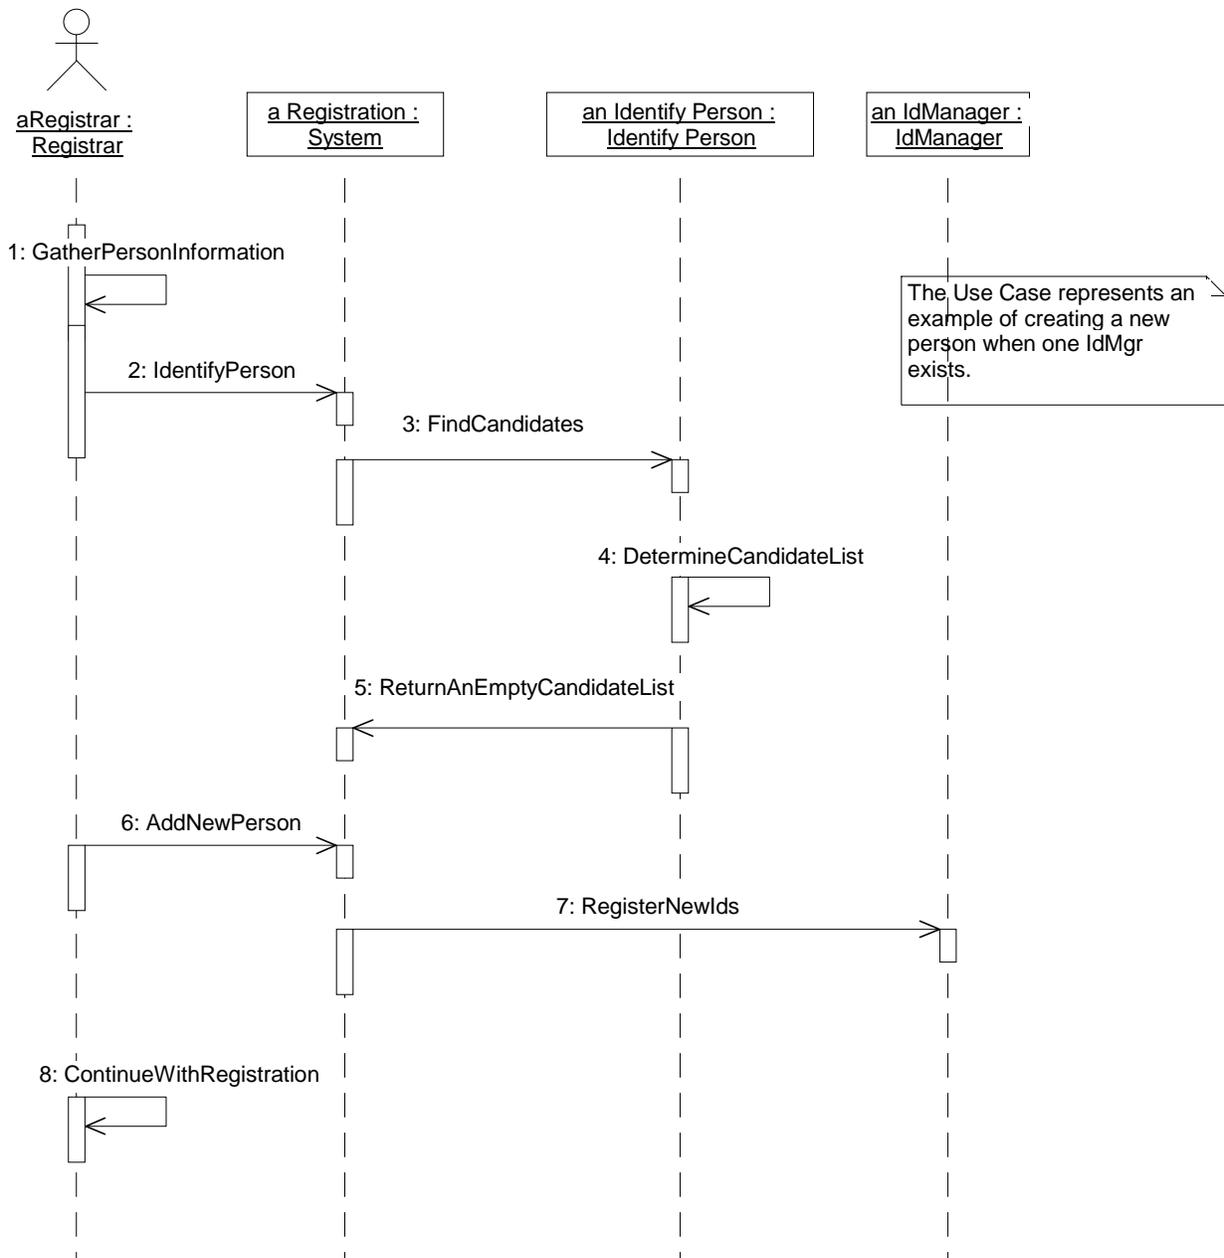
Each row in the following table includes the specification for a different conformance class. The columns represent the interfaces on the IdentificationComponent. A star '\*' in a column indicates the conformance class in that row includes the interface of that column.

Conformance Class	IdentifyPerson	ProfileAccess	SequentialAccess	IDMgr.	IdentityAccess	CorrelationMgr
Simple PIDS	*	*				
Sequential Access PIDS	*	*	*			
ID Domain Mgr PIDS	*	*		*		
Identity Access PIDS	*				*	
Correlation PIDS						*

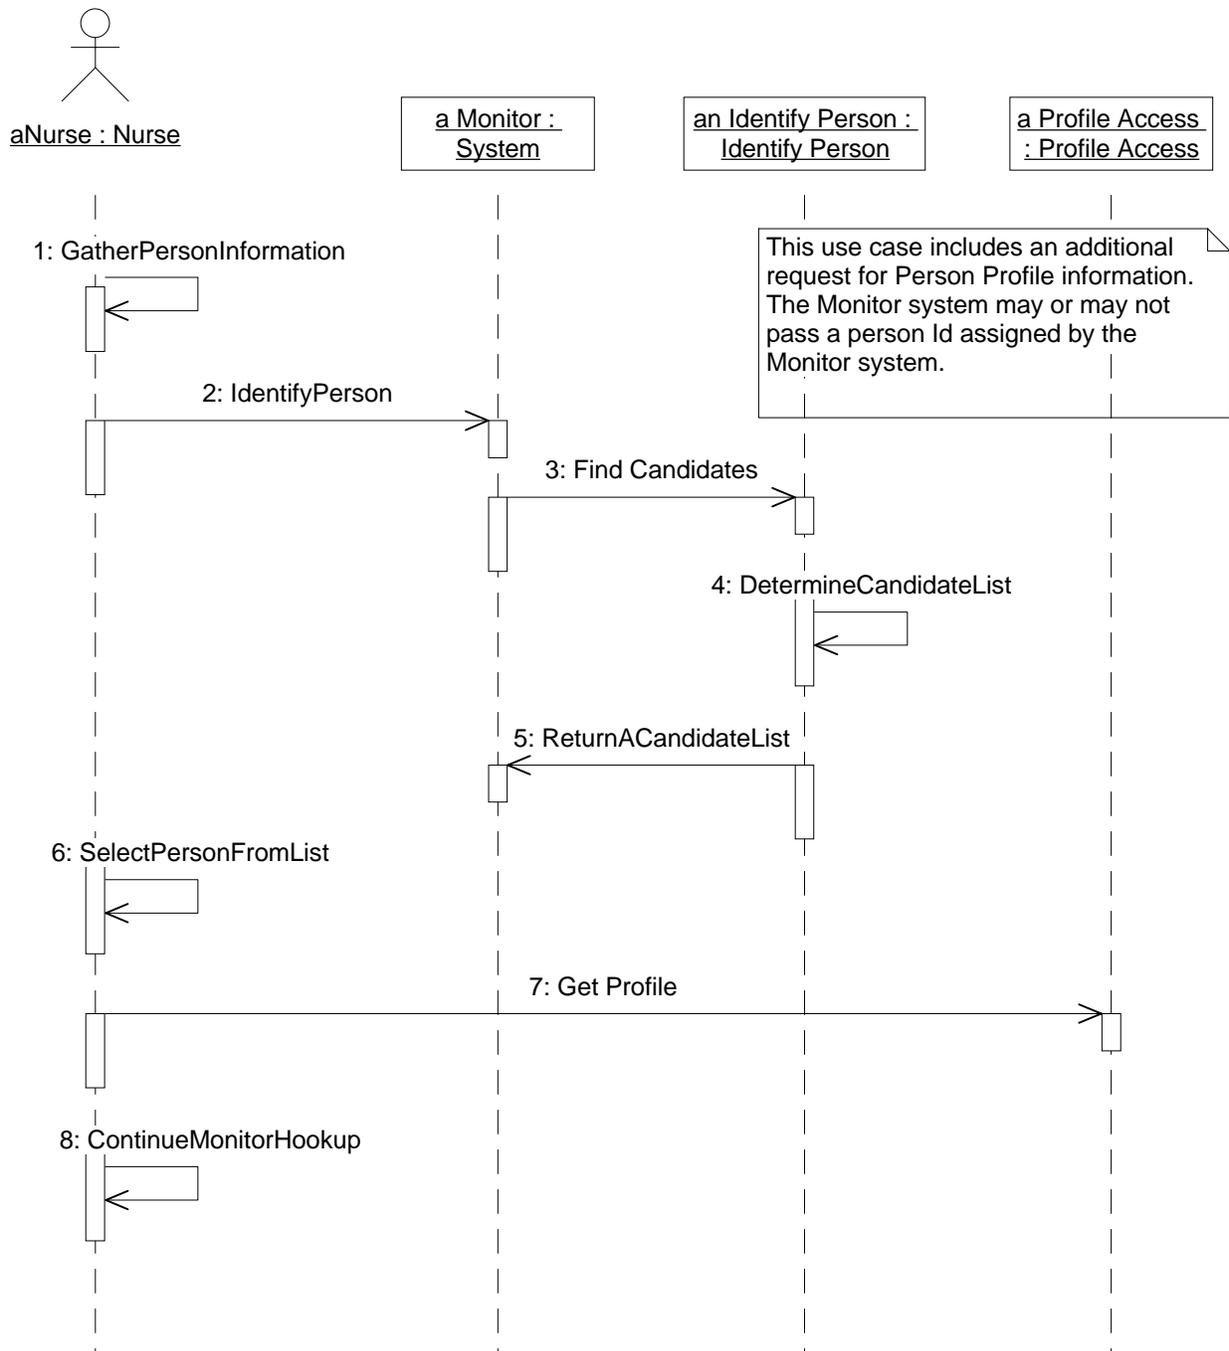
- **'Simple PIDS'** - This class provides the basic operations to access profiles from an ID and to match potential IDs given some traits.
- **'Sequential Access PIDS'** - This conformance class adds the ability to scroll through the set of IDs sequentially.

- **'ID Domain Mgr PIDS'** - This conformance class adds the ability to create new IDs and modify the states of already created IDs.
- **'Identity Access PIDS'** - This conformance class is an alternative to the Simple PIDS in that it provides the same functionality but uses the IdentityAccess interface instead of the ProfileAccess. It is intended for implementations that have different access policies for each ID, for example allowing people to control access to their own profile.
- **'Correlation PIDS'** - The supported\_traits attribute contains at least the following three traits: HL7:PatientName; HL7:DateTimeOfBirth; HL7:Sex. This conformance class contains a single interface as it provides functionality by its self. However it can be mixed in with the other conformance classes for additional functionality.

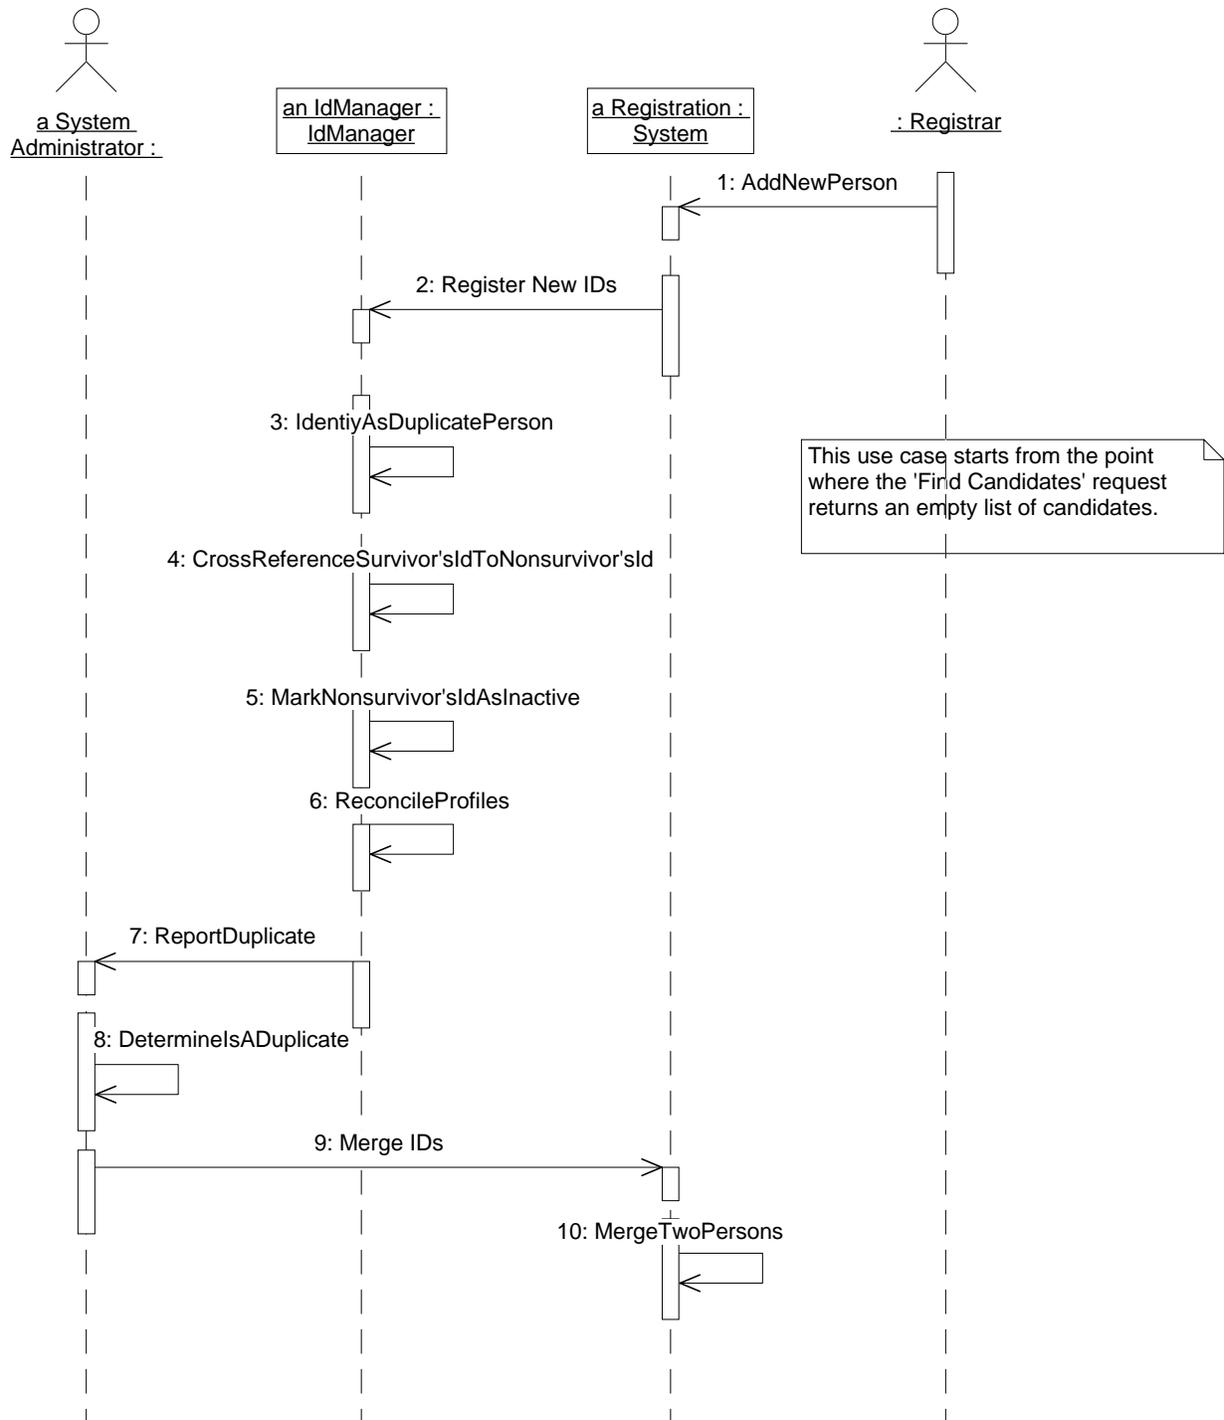
## 8. Appendix - Use Case Examples



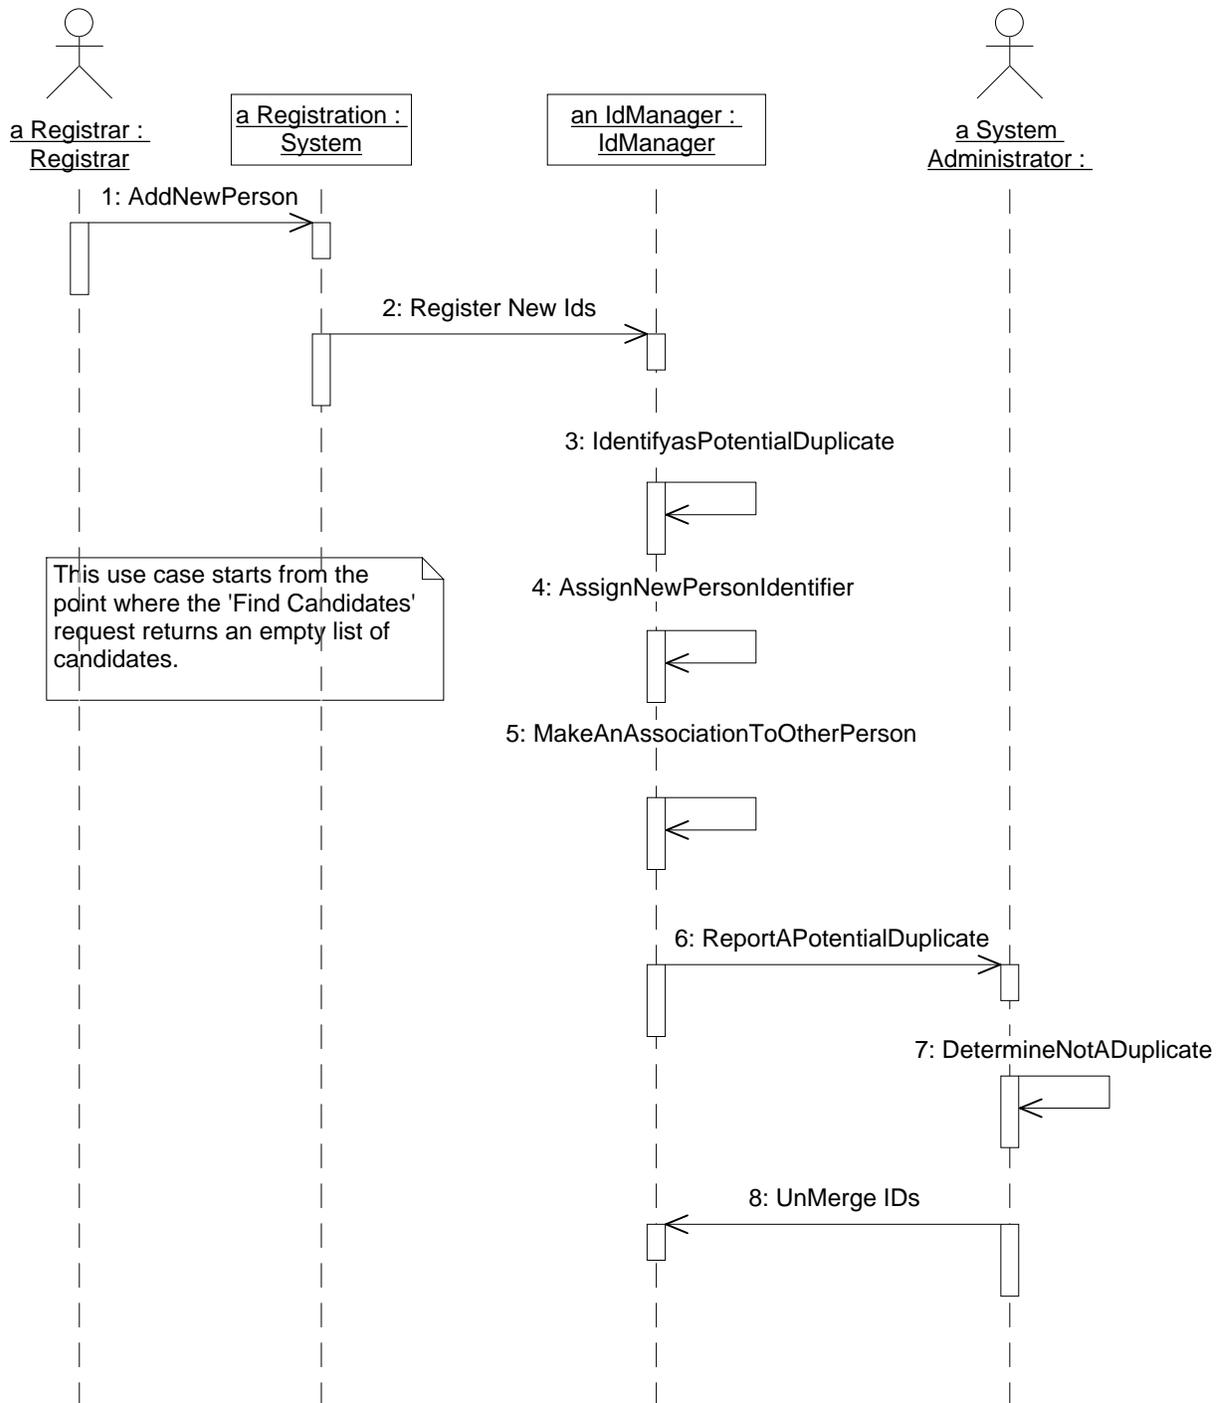
*Find Candidates and Register New IDs.*



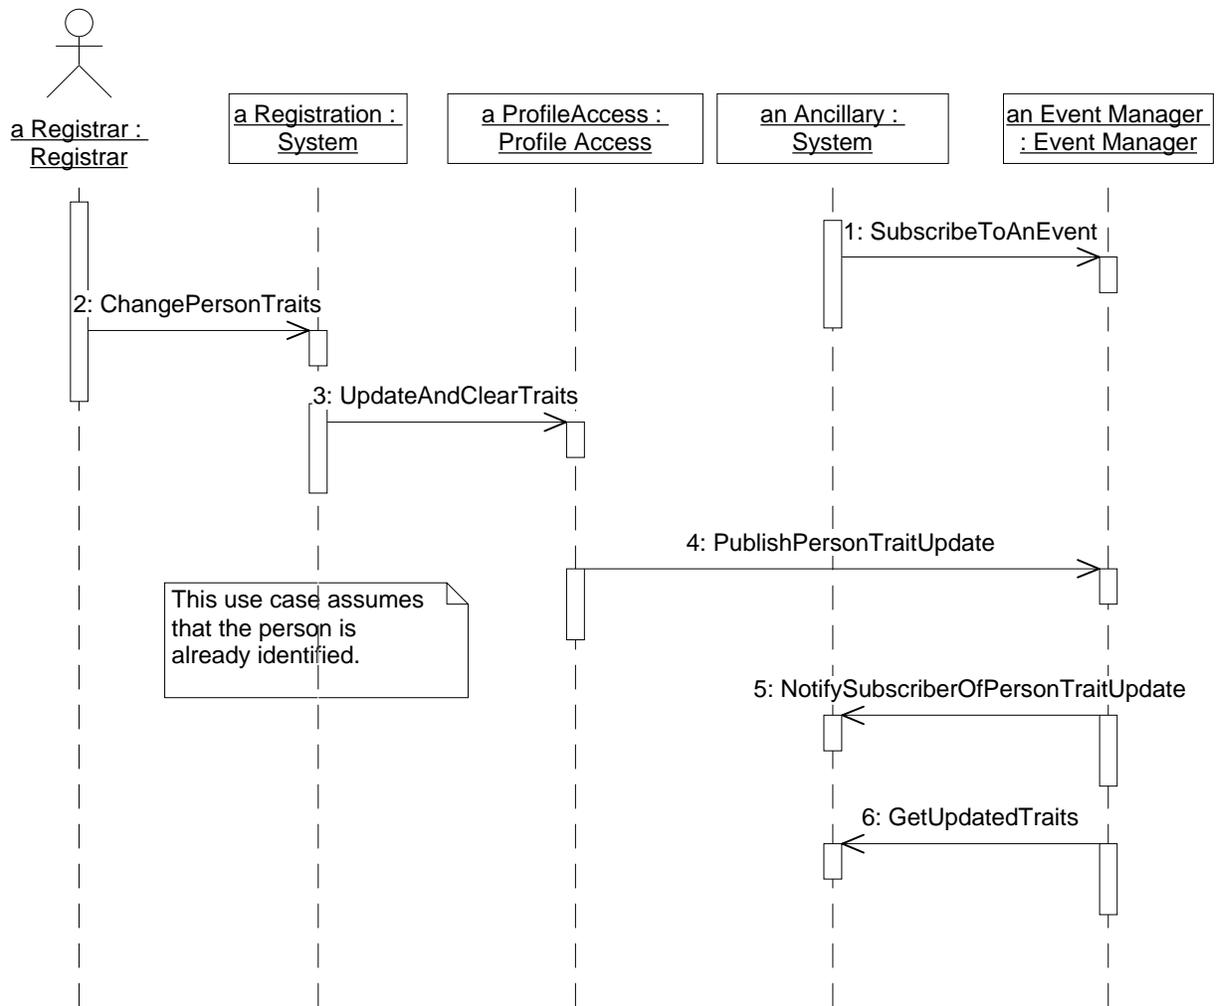
*Find Candidate and Get Candidate Profile.*



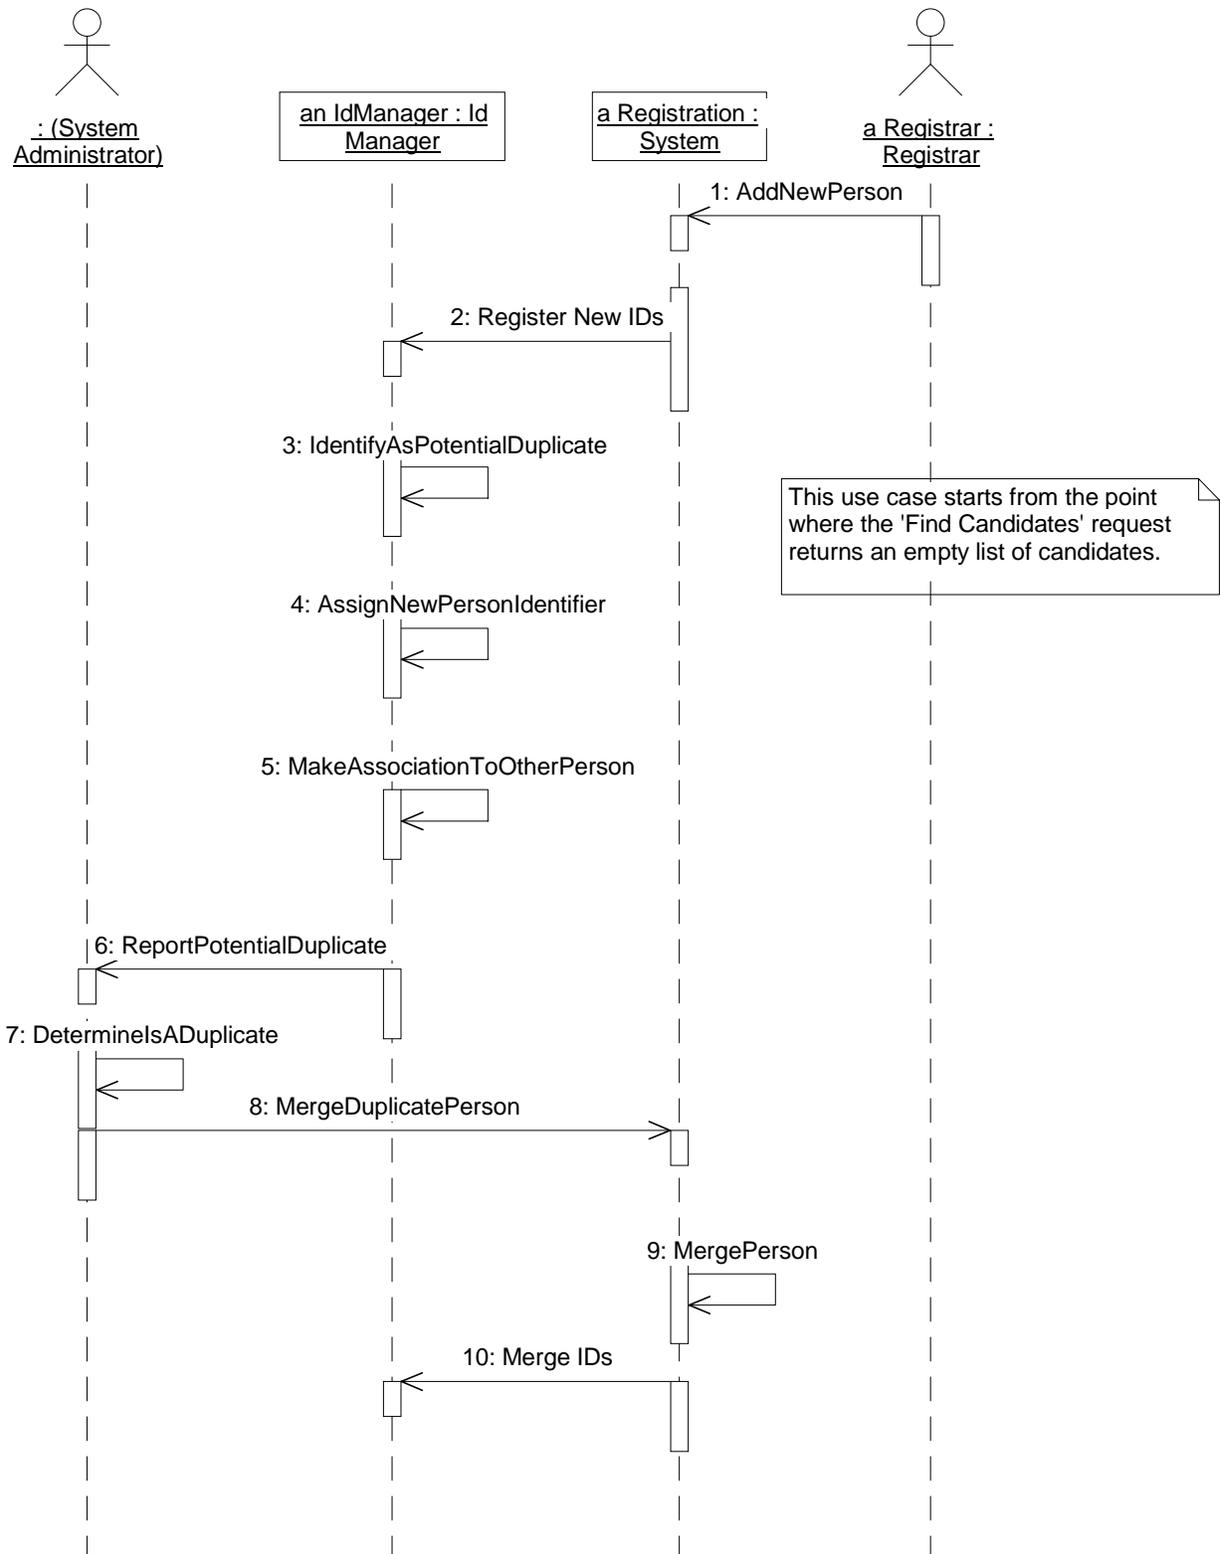
*Merge IDs.*



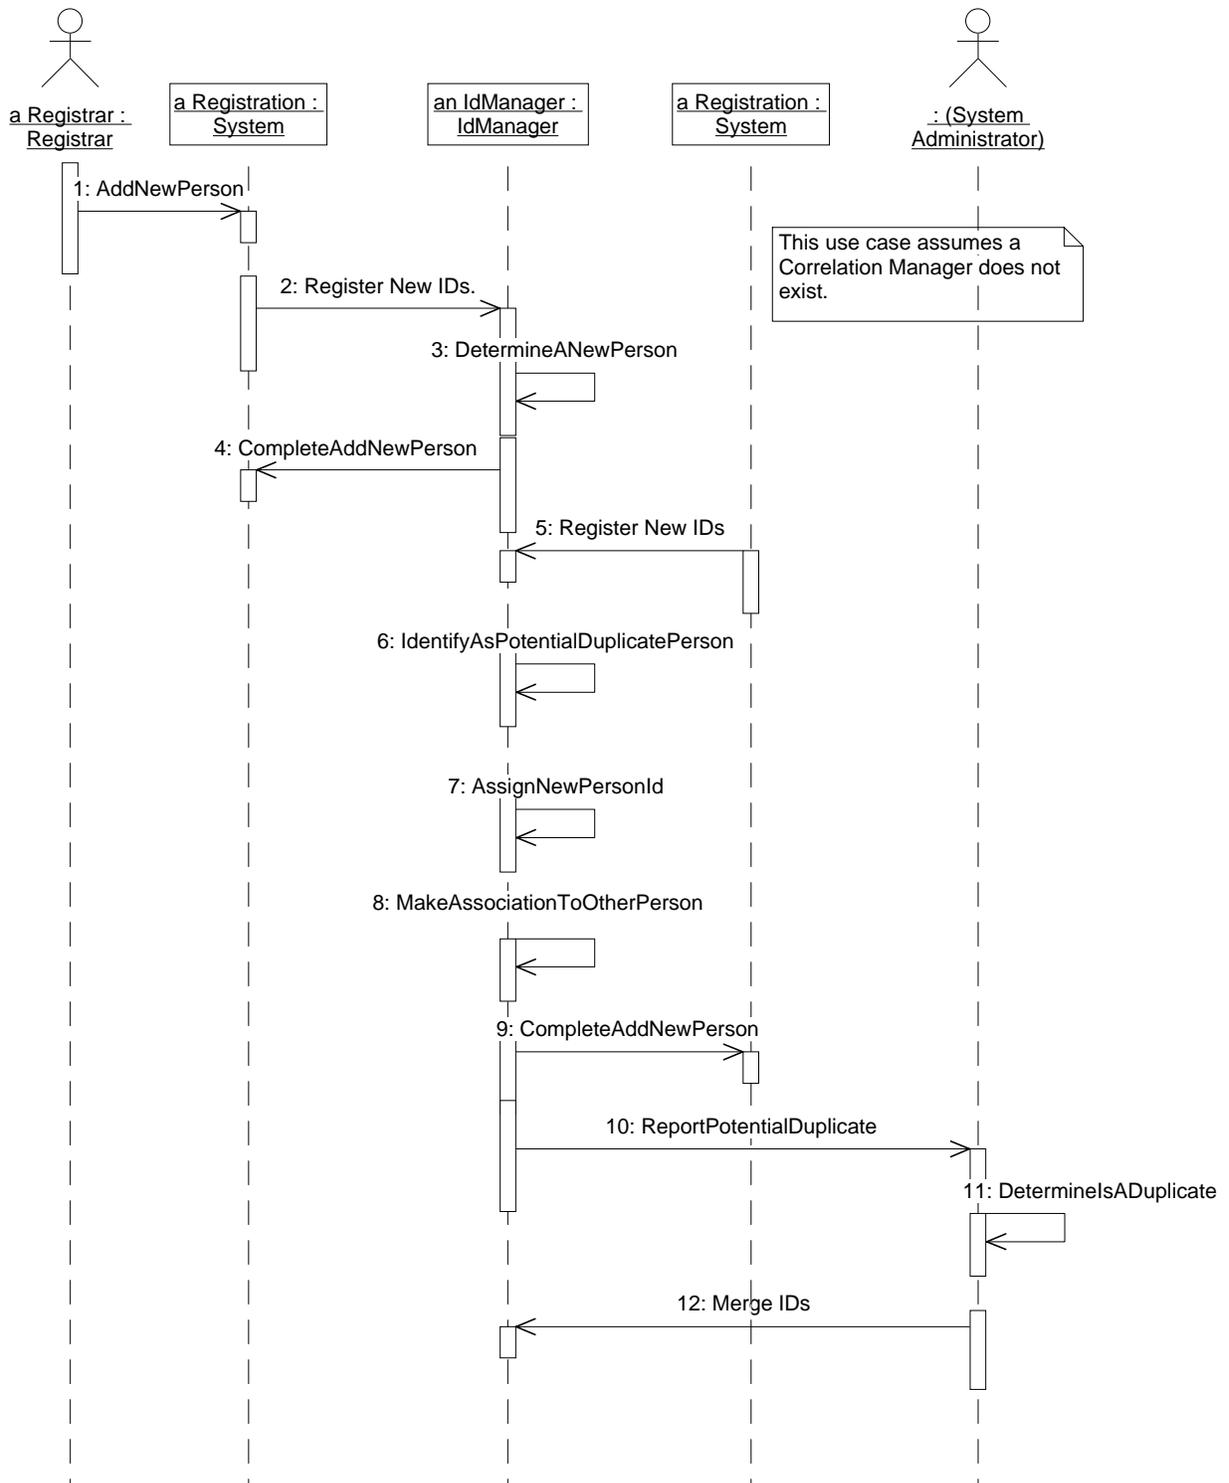
*Unmerge IDs.*



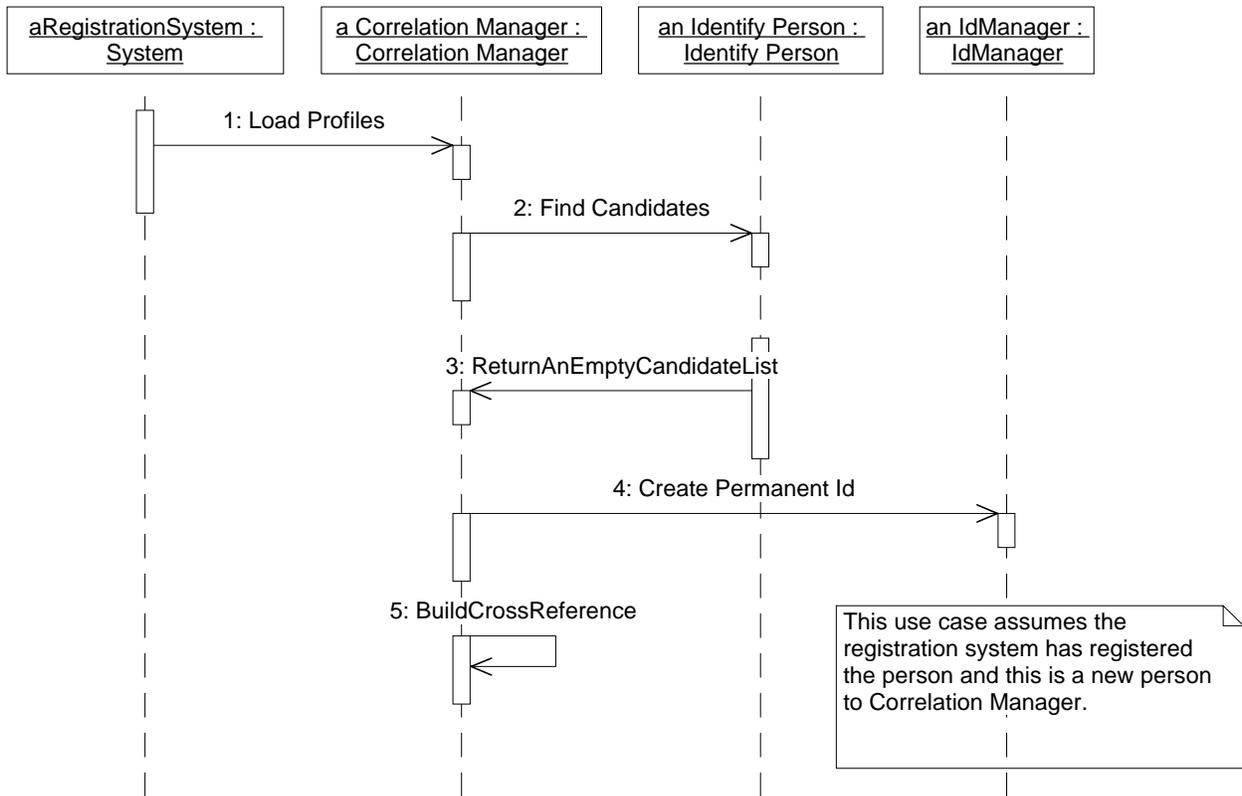
*Update Person Traits and Notify Subscribed Parties.*



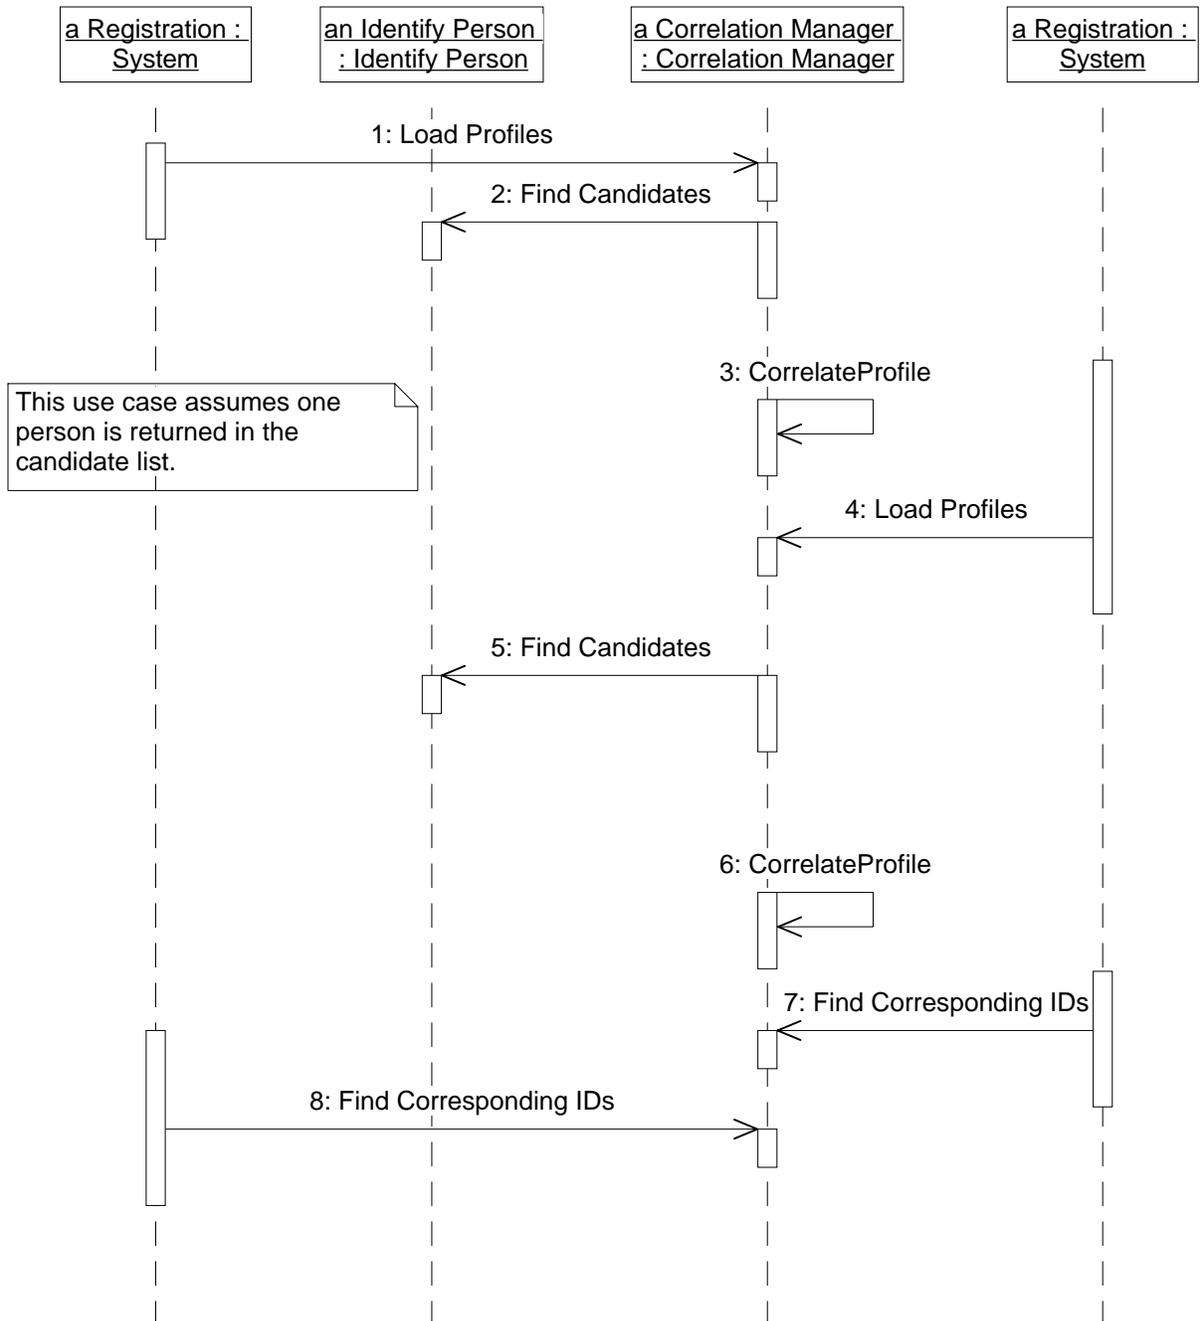
*Potential Merge Person in a Single ID Domain.*



*Merge Person without a Correlating Manager.*



*Correlate Person Profile When Person Does Not Exist.*



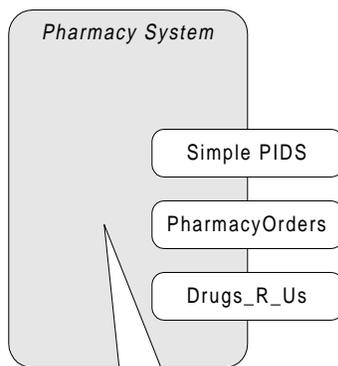
*Correlate Profile When Person Already Exists.*

## 9. Appendix - Interaction Patterns

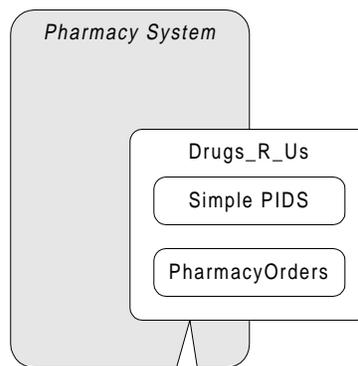
---

The following diagrams indicate a variety of ways that PIDS can be used. Obviously this is not a complete list but it does show there is extreme flexibility in the way it can fit into an enterprises architecture.

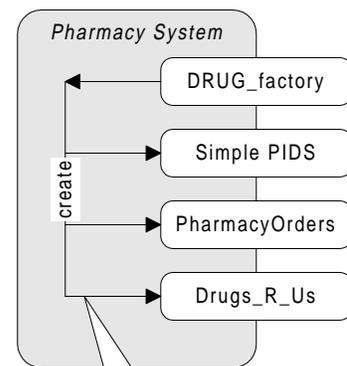
There are multiple ways to integrate the PIDS IdentificationComponent with other interfaces implemented by the same system as shown in the examples below.



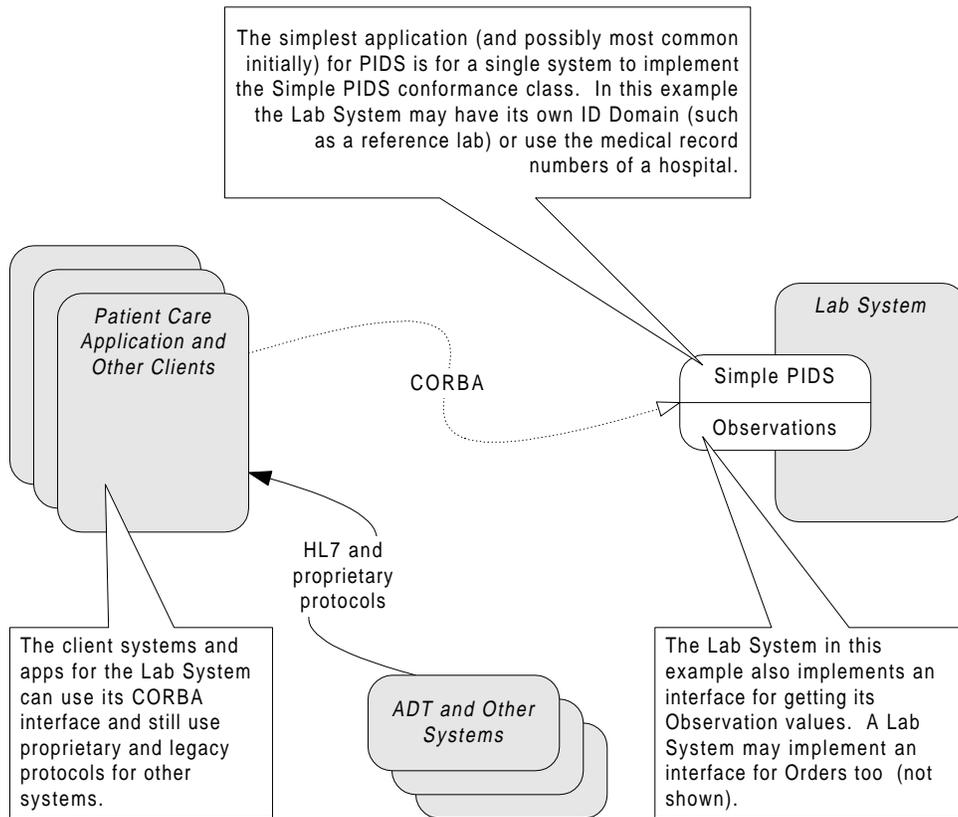
Implementing standard and proprietary interfaces separately that are independently used.

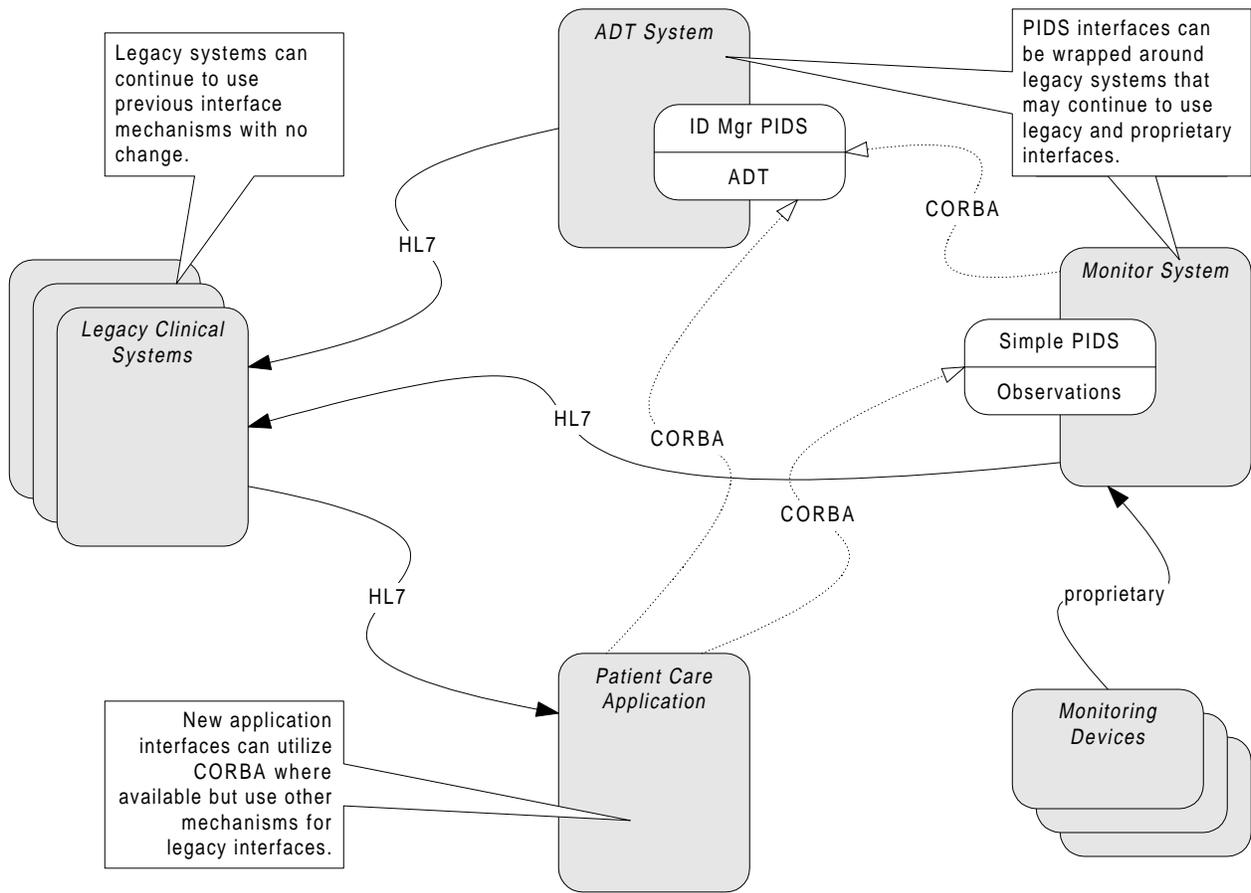


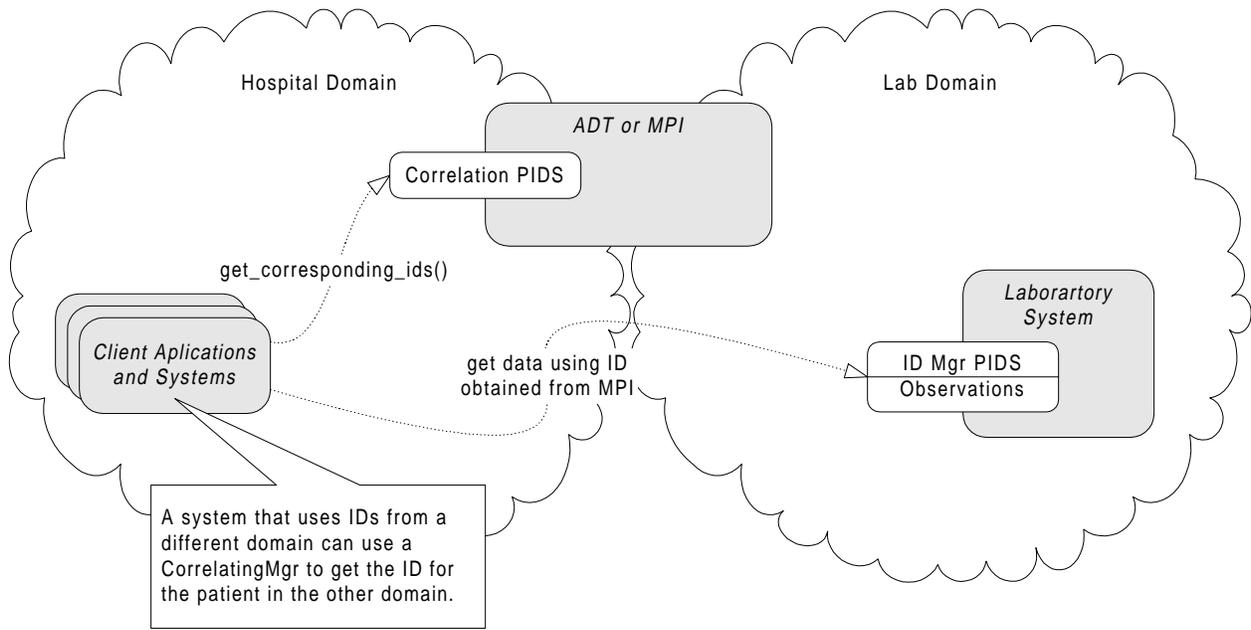
Mixing two standard interfaces into a proprietary one via multiple inheritance.

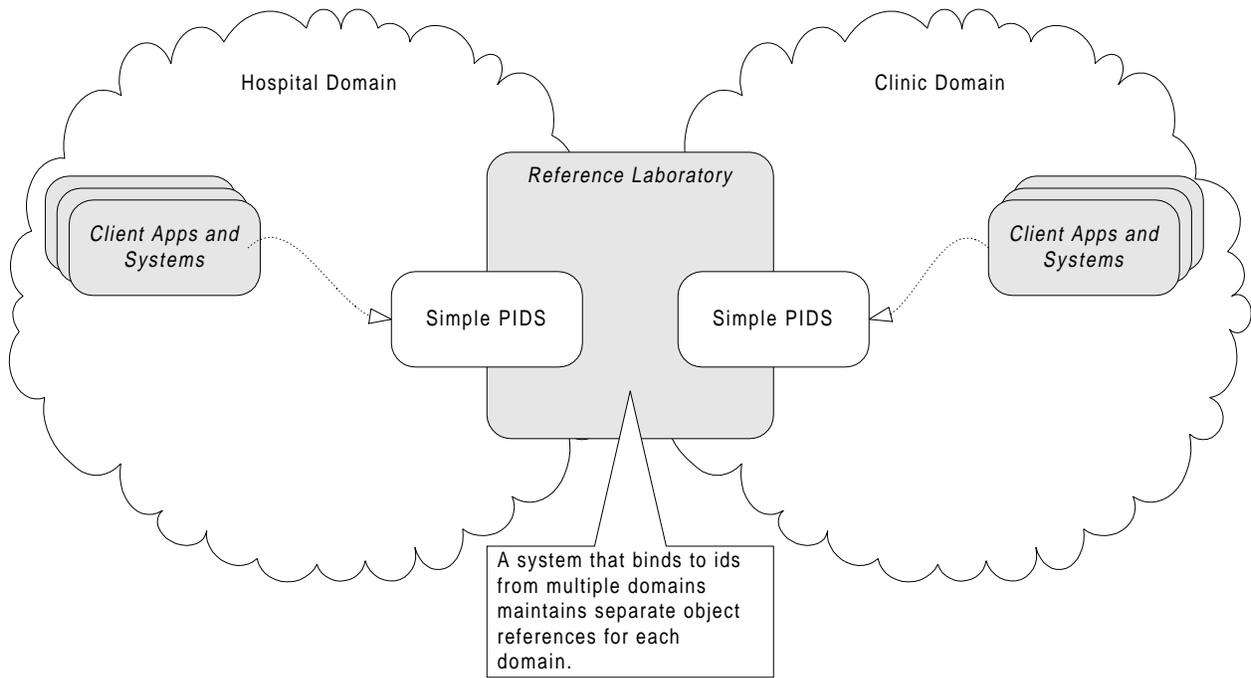


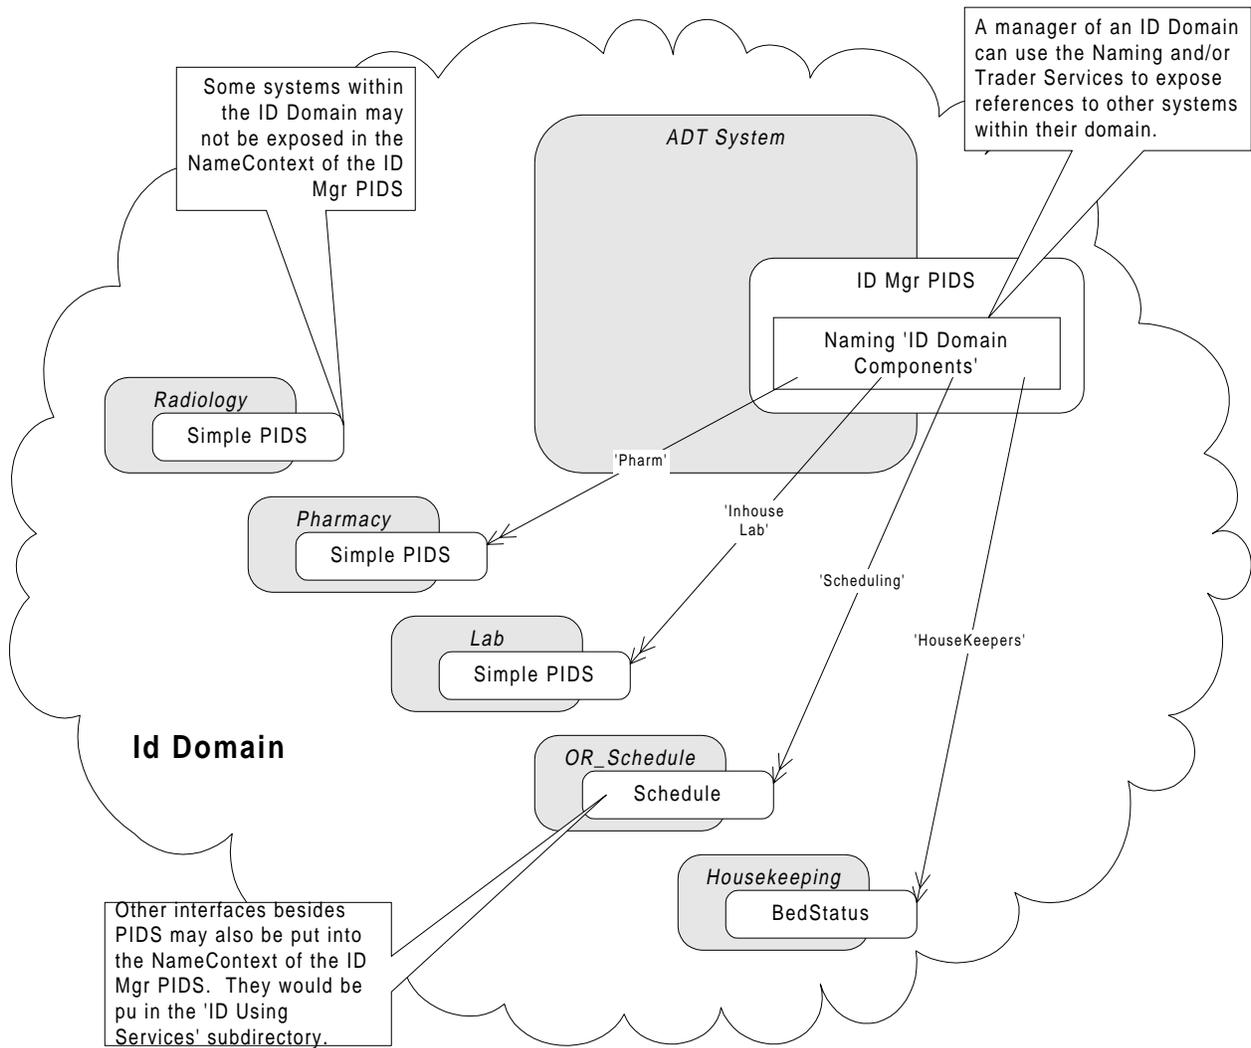
Separate interfaces pulled together (created) via a factory interface.

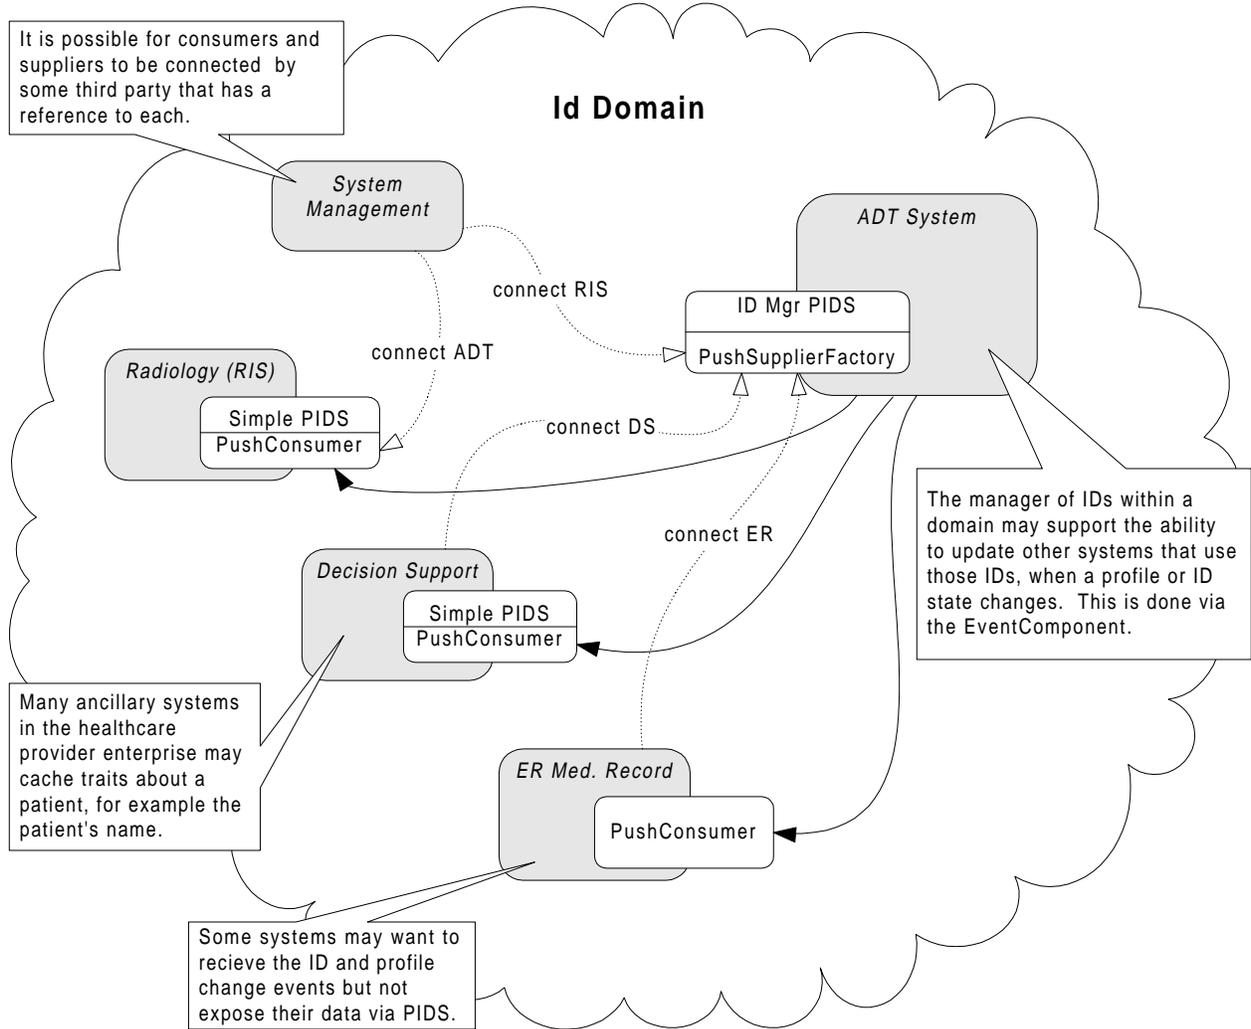


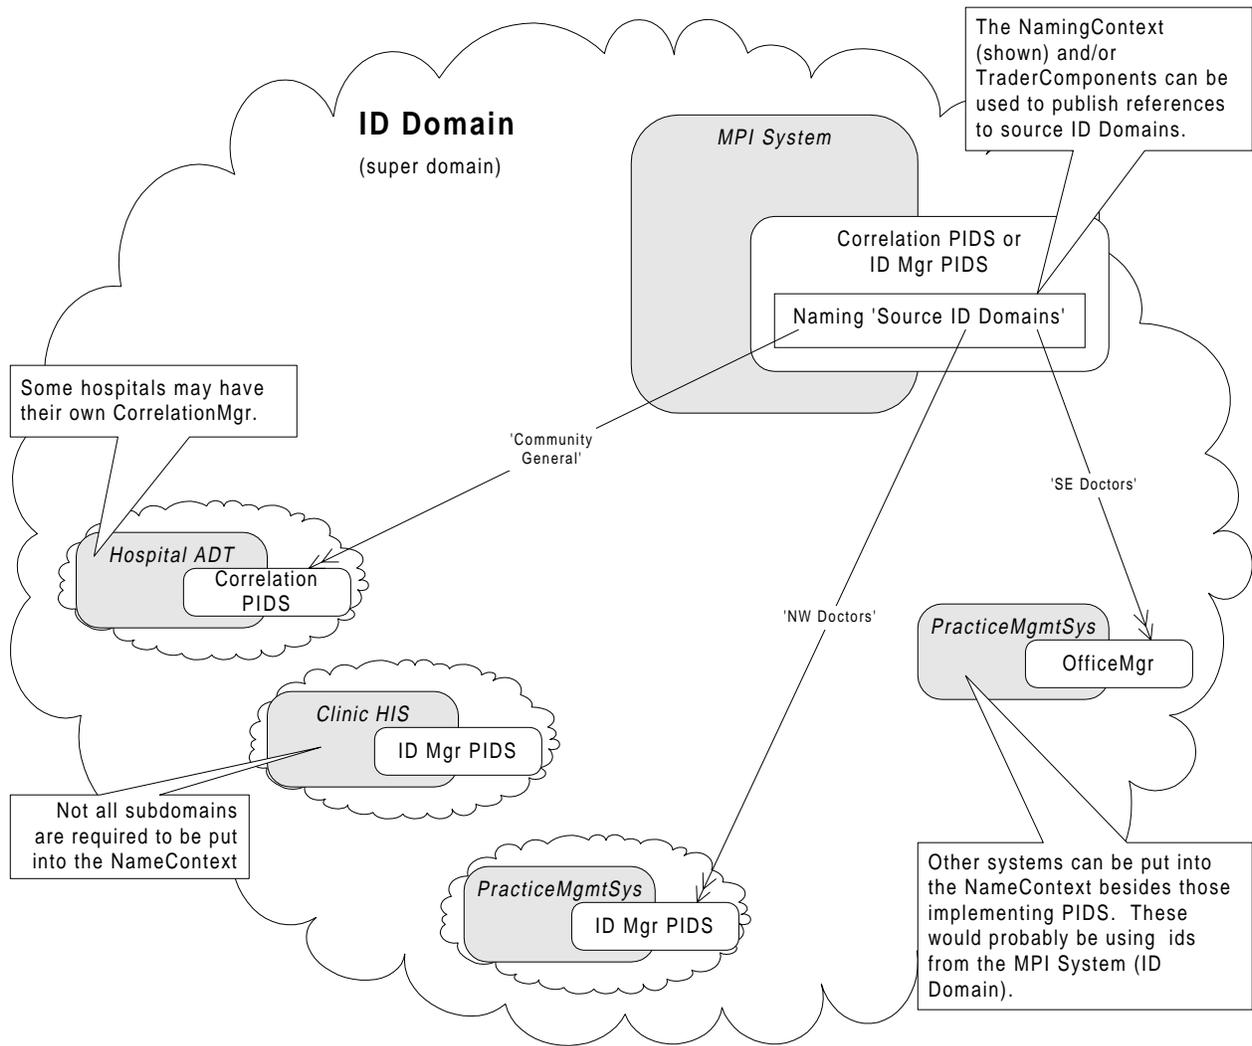


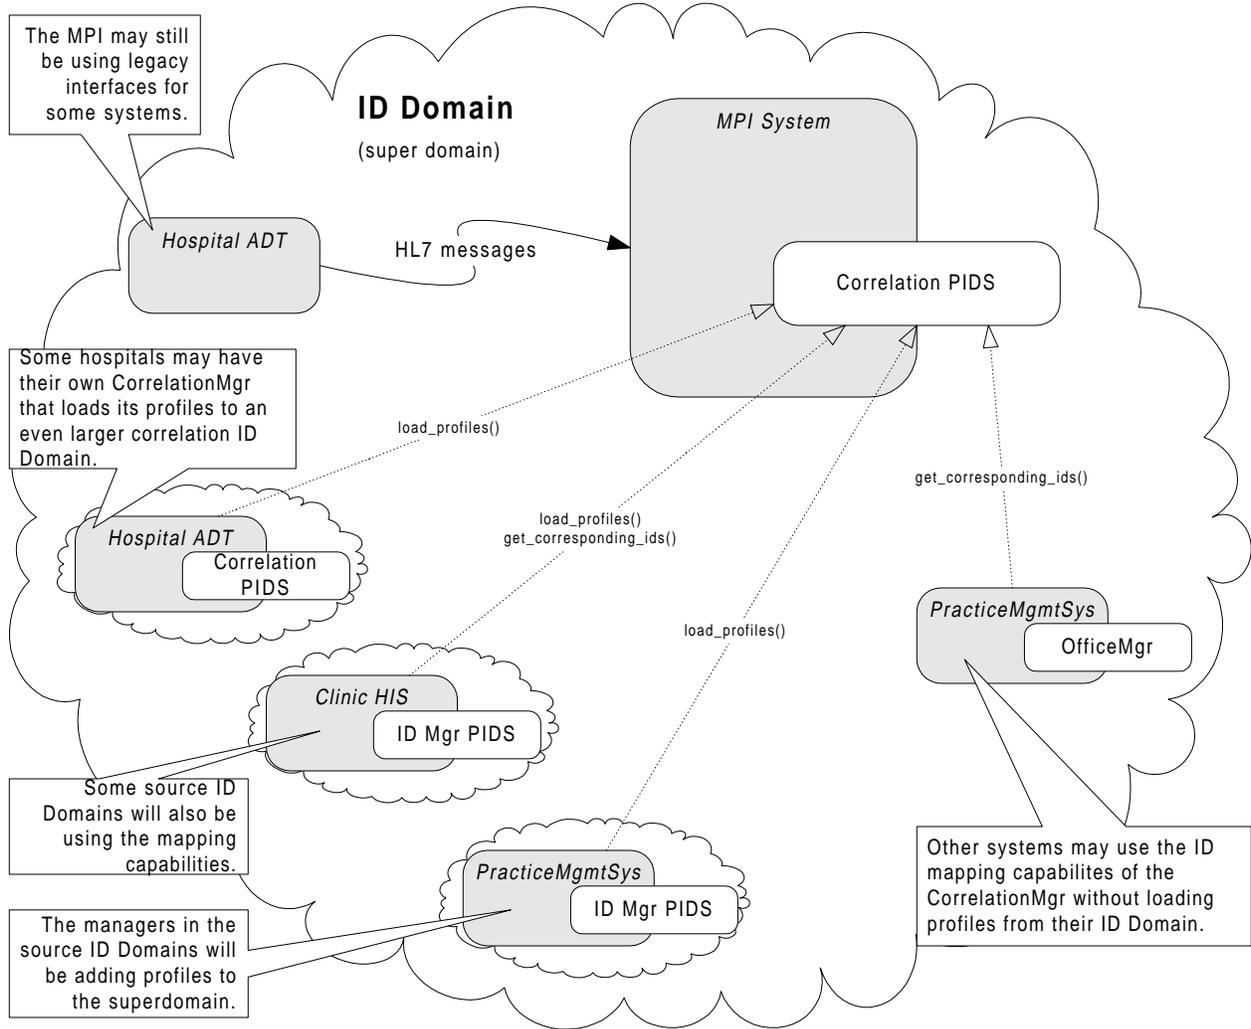


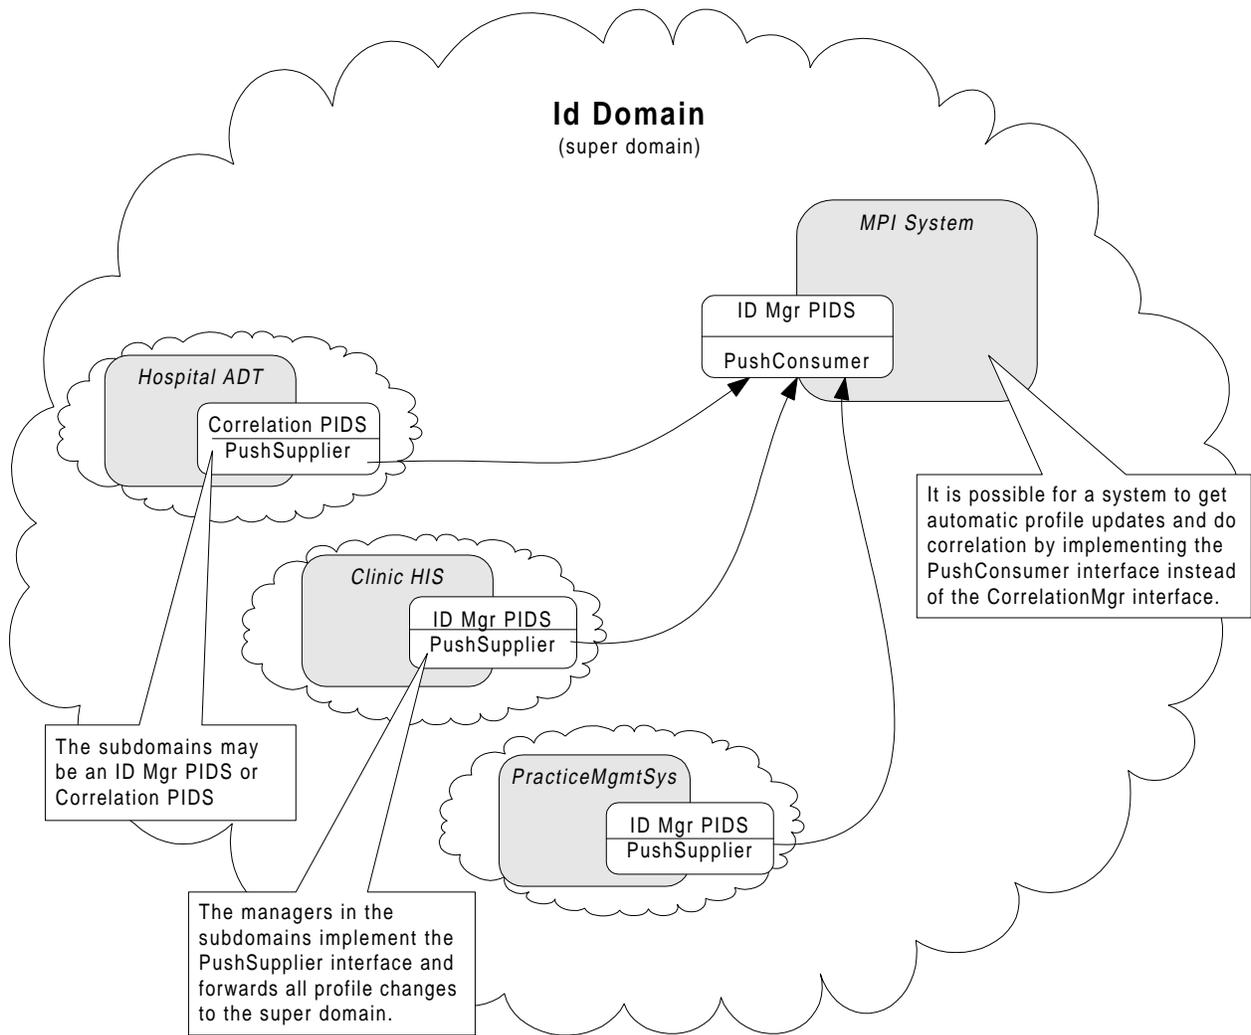


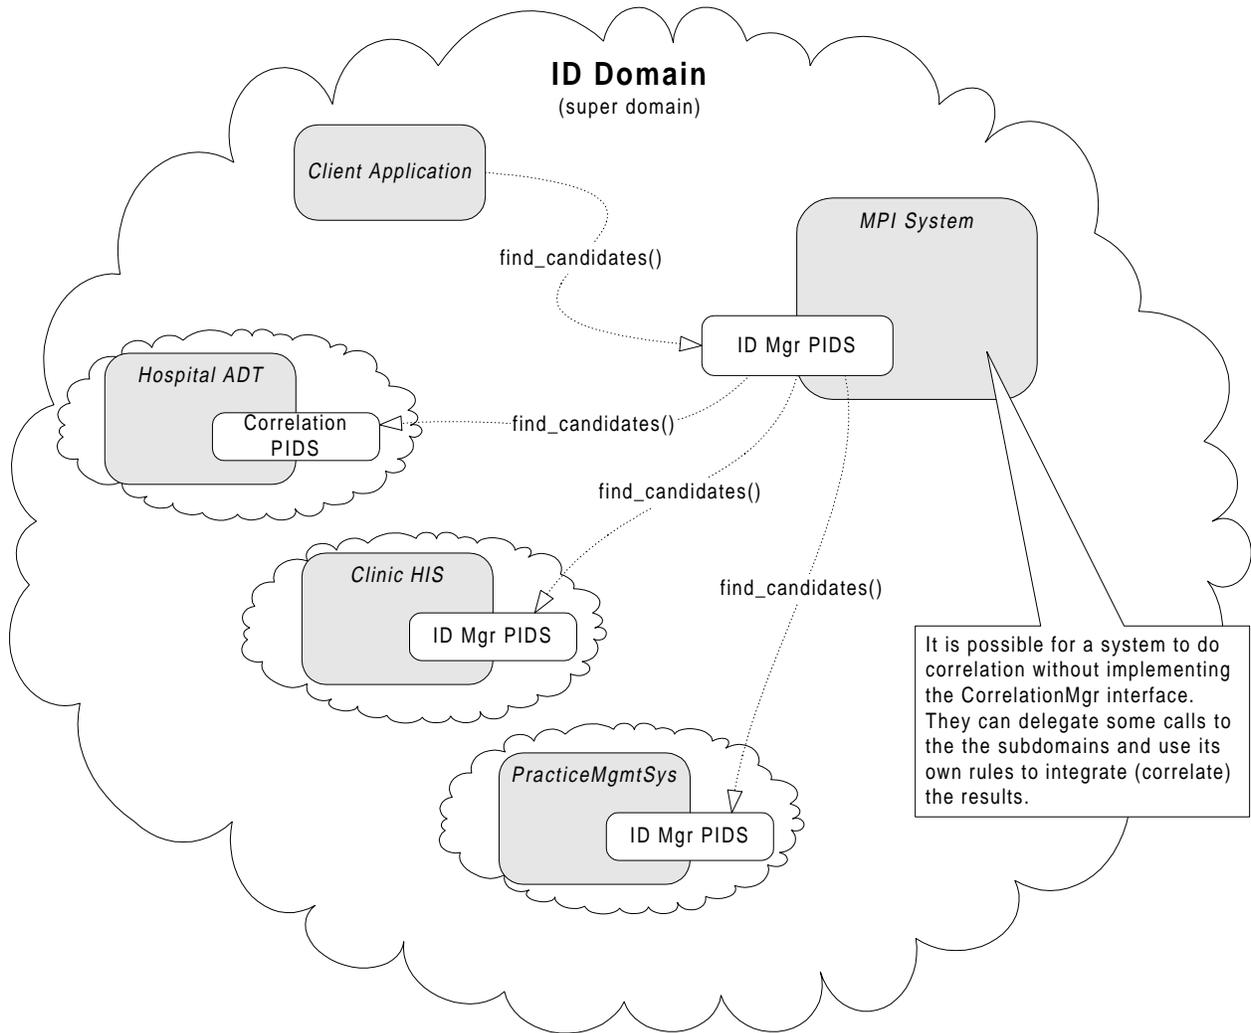












## 10. Appendix - Event Descriptions

---

The information in this appendix is informative and is not a normative part of the specification. It was prepared by the submitters under the expectation that there would be an adopted Notification Service.

Two of the initial submissions to the Notification Service propose the creation of event types similar to the service types defined by the Trader Service. They also propose using the constraint mechanism for event consumers to specify the filters to apply to event channels. The aspects of service types that apply to events are the type names and property descriptions as described below.

- **Event Type** - Event types have a name and a definition. The event may inherit from previously defined events. An inherited event has all the same properties defined by the parent plus any the new type adds. The definition of the type consists of a set of properties.
- **Properties** - A property is a name/value pair where the value has a specific IDL type. The property is also characterized by a boolean attribute indicating if it is mandatory.

The event types and associated property descriptions needed for PIDS are described below. The syntax used is similar to the Trader service type descriptions except they do not include an interface name or have the concept of 'readonly' properties.

```
event PersonIdChange {
    mandatory property string qualified_person_id;
};

event PersonIdStateChange : PersonIdChange {
    mandatory property octet new_state;
    property octet old_state;
};

event PersonIdStateMerged : PersonIdStateChange {
    mandatory property string preferred_id;
};

event PersonIdStateUnmerged : PersonIdStateChange {
    mandatory property string old_preferred_id;
};

event PersonIdProfileChange : PersonIdChange {
    mandatory property StringSeq trait_names_changed;
    mandatory property string component_name;
    property string component_version;
    property PersonIdService::TraitSeq changed_traits;
    property PersonIdService::TraitSeq new_profile;
};

event PersonIdDuplicate {
    mandatory property StringSeq qualified_person_ids;
};

event PersonIdCollision {
    mandatory property String qualified_person_id;
};

event IdentificationComponentChange {
    mandatory property string domain_name;
    mandatory property string component_name;
};

event IdentificationComponentVersionChange : IdentificationComponentChange {
    mandatory property string new_version;
    property string old_version;
};
```

```

};

event IdentificationComponentTraitChange : IdentificationComponentChange {
    property StringSeq new_trait_names;
    property StringSeq old_trait_names;
    property boolean trait_spec_changed;
};

event CorrelationSourceChange : IdentificationComponentChange {
    property StringSeq new_source_domain_names;
    property StringSeq old_source_domain_names;
};

```

### PersonIdChange event

This is a general event for any changes on an ID.

#### **qualified\_person\_id**

This is the ID that had something change on it. It is a stringified version of the QualifiedPersonId. The stringification is per the rules defined in the NamingAuthority for QualifiedNameStr.

### PersonIdStateChange event

This event type is a subtype of the general PersonIdChange event. It applies only to state changes on the ID.

#### **new\_state**

The new\_state is a mandatory property which has the new IdState for the ID. The IdState enum values are coded into the octet starting with 0 and increasing sequentially in the order they are defined in the IDL.

#### **old\_state**

The old\_state is optional and represents the IdState value before the change.

### PersonIdStateMerged event

This event is a special case of a PersonIdStateChange event where the new state is Deactivated and the ID has a preferred ID set (e.g. the merge\_ids() operation was called).

#### **prefered\_id**

This is the ID that the qualified\_person\_id is merged into.

### PersonIdStateUnmerged event

This event is a special case of a PersonIdStateChange event where the old state is Deactivated and had a preferred ID. The ID has been reactivated (e.g. the unmerge\_ids() operation was called).

#### **old\_prefered\_id**

This is the ID that the qualified\_person\_id was merged with.

### PersonIdProfileChange event

This event is a specialization of the PersonIdChange event where the ID profile bound to the qualified\_person\_id has changed.

#### **trait\_names\_changed**

This mandatory property indicates the TraitName for each trait that has changed in the profile.

**component\_name**

This property contains 'the\_name' part of the ComponentName where the change occurred. This may be used by a client to filter events except those from a particular component.

**component\_version**

This property contains 'the\_version' part of the ComponentName where the change occurred.

**changed\_traits**

This property contains the traits (including values) that changed. It is an optional trait since some trait values may be large (e.g. photographs) and the client will have to query for their value.

**new\_profile**

This property contains the whole new profile. It contains a superset compared to the changed\_traits property.

*PersonIdDuplicate event*

This event indicates the PIDS has detected suspected duplicate IDs for the same person.

**qualified\_person\_ids**

These are the set of IDs the service thinks may represent the same person.

*PersonIdCollision event*

This event indicates the PIDS has detected a suspected collision on an ID. That is two or more people may be represented by (or be using) the same ID.

**qualified\_person\_id**

This is the stringified QualifiedPersonId that a suspected collision has occurred on.

*IdentificationComponentChange event*

This is a general event representing some of the pseudo static information on an IdentificationComponent has changed.

**domain\_name**

This is stringified ID DomainName where the component resides in.

**component\_name**

This is the \_name from the ComponentName of the component the change occurred on.

*IdentificationComponentVersionChange event*

This event indicates the component's version has changed. That is the component may have had a configuration change, software update, etc. Whether there is any outward functionality change on the component cannot be determined for this event.

**new\_version**

This is the new the\_version from the ComponentName on the component that changed.

**old\_version**

This is the old the\_version from the ComponentName on the component that changed.

*IdentificationComponentTraitChange event*

This event indicates some aspect of the supported\_traits attribute has changed on the component.

**new\_trait\_names**

This sequence contains the trait names for any new traits that are now supported by the component.

**old\_trait\_names**

This sequence contains the trait names for any old traits that are no longer supported by the component.

**trait\_spec\_changed**

This boolean attribute indicates that some of the values on the TraitSpec (mandatory, read\_only, and/or searchable) has changed on some of the traits that were and still are supported.

**CorrelationSourceChange event**

This events occurs when the source ID Domains for which a PIDS is correlating over changes.

**new\_domain\_name**

These are the DomainNames for any new source ID Domains that is being correlated over.

**old\_domain\_name**

These are the DomainNames for any old source ID Domains that are no longer being correlated over.

# *11. Appendix - Security Guidelines*

---

The PIDS interfaces may be used in many different environments with widely varying security requirements that range from no security to extreme security. For this reason the PIDS interfaces do not expose any security information. PIDS relies on the underlying CORBA infrastructure and services which provides all the security mechanisms needed without exposing it in the interfaces.

An attribute of security that of concern to many people is to maintain confidentiality of certain (sensitive) information about them. For PIDS this implies being able to filter requests by:

- who is accessing the information,
- who the information is about,
- what information is being accessed.

Other common security concerns could be preventing unauthorized modification of data, tapping into communications to acquire sensitive information, and causing loss of service by overburdening a service.

CORBA Security provides robust mechanisms to address these and other concerns. Some of the security properties it does deal with includes authentication, authorization, encryption, audit trails, non-repudiation, etc. CORBA Security, in its default mode allows these security concerns to be addressed without the client and server software being aware of it. This is a powerful notion, allowing security policies to be created and enforced after applications and systems have been created and installed.

Other CORBA and CORBA Security features provide mechanisms for applications to extend these security capabilities. For example they can obtain credentials from the ORB and implement filters that can look at specific data passed to and returned from operations.

It is a requirement of the PIDS to provide confidentiality of information that is stored about an individual. This requirement fuels the need for fine grained access control on trait information that is associated with a PIDS ID.

PIDS provides two interfaces to access information about a person, given their ID. These are the ProfileAccess and IdentityAccess. The functional capabilities of these two interfaces are identical, but they have different strengths. The IdentityAccess interface can simplify the internal implementation of some types of security policies - those where there is a different policy for each ID. The ProfileAccess interface's strength is that it only requires a single call to access data for multiple IDs. The IdentityAccess requires separate calls for each ID information is accessed by.

The IdentityAccess interface allows a client to acquire an object reference to an Identity interface, containing information pertaining to only one particular ID. Creating a CORBA object reference as a single accessor to the information yields benefits in controlling the access to that information. It is a single point to which access of this information must flow. Since CORBA security services can automatically deny requests as a result of an access policy, access to information behind this interface can be controlled by that mechanism.

## *11.1 Security Requirements*

For the PIDS to be secure in its possible dissemination of information it needs to adhere to several requirements:

- The PIDS needs to authenticate a client's principal identity, role, and sensitivity level.
- The PIDS needs to transmit information confidentially and with integrity.

The first requirement states that the entire PIDS interface implementations must be able to identify a potential client. If it cannot authenticate a client, then the client may be severely limited in the particular requests that the PIDS can service. The CORBA Security Service provides the mechanisms for a server to authenticate a client.

The second requirement provides for the confidentiality of the information. The client must communicate with the PIDS using not only encryption to protect data, but signature as well, so as not to have data tampered with during communication. There is no sense in putting a Sensitivity level of "OwnerOnly" on a trait and have its value transmitted to the owner in the clear. The CORBA Security Service provides these capabilities, including SSL.

The problem is, How does one get CORBA to support this access policy model?

## *11.2 CORBA Security Overview*

In an effort to keep the PIDS interfaces security unaware, i.e. no extra visible security relevant parameters in methods, access policy must be adhered to from behind the interfaces. The CORBA security model offers several ways to apply security policy to method invocations.

The CORBA Security Specification (CORBAsec) is not a cookbook for using CORBA security in building applications. It is a general framework with which ORB vendors and application vendors can build a multitude of different security policy models. The CORBAsec also gives the interfaces for which implementations of applications can access those security services that are supplied with a secure ORB.

A secure PIDS implementation that can control access to specific traits must be aware of the security services offered by the ORB. This caveat also means that a client's ORB may have to know the kind of ORB and the security services that is used by the PIDS.

The CORBA security specification outlines a general security policy model. Although the specification is vague about which approach should be taken, it is specific enough to be able to choose from a couple of models that can be supported.

The CORBA security model bases itself on credentials and security domains. Credentials are data objects that contain attributes such as privileges, capabilities, and sensitivity levels, amongst others. Security domains are mappings from credentials to access rights. Credentials can be encrypted and signed to prevent tampering and achieve a level of trust between client and server. CORBA credentials get passed with requests beneath the visible level of the interface. CORBA security services give the clients and servers the ability to authenticate/verify credentials to implement policies in security domains.

Many different schemes, algorithms, services, and vendor implementations exist to provide implementation of security policy, and many different implementations of those schemes may be integrated into a CORBA compliant ORB. It is not the purpose of this specification to dictate the specific implementation of an ORB and security services that should be used, but to outline the external requirements for the PIDS implementation. These requirements and guidelines aid in selecting a secure ORB and the level security functionality needed to implement the PIDS access policy model.

## *11.3 Secure Interoperability Concerns*

CORBA has built the communication bridge between distributed objects creating a interoperable environment that spans heterogenous platforms and implementations. However, security adds another layer of complexity to the issue of interoperability. ORB implementations are not required to include security services nor provide an interoperable mechanism of security services.

However, a specification does exist for the target object to advertise, via the IOR, the security services that it supports and the services it requires from the client. Both the client and server ORBs must use compatible mechanisms of the same security technology.

The CORBA Common Secure Interoperability (CSI) Specification [document number??] defines 3 levels of security functionality that ORBs may support. The levels are named, CSI Level 0, CSI Level 1, and CSI Level 2. Each level has increasing degrees of security functionality.

The CSI Level 0 supports identity based policies only and provides mechanisms for identity authentication and message protection with no privilege delegation. The CSI Level 1 adds unrestricted delegation. The CSI Level 2 can implement the entire CORBA Security Specification at Security Level 2.

Each CSI level is parameterized by mechanisms that can support the level of security functionality, such as SPKM for CSI Level 0, GSS Kerberos for CIS Level 0 or CIS Level 1, and CSI\_ECMA for CSI Level 2. Future developments in security functionality and mechanism are not restricted, and mechanisms can be added to each level.

The ORB implementations may use different security technology with differing capabilities and underlying mechanisms, such as SSL, DCE, Kerberos, Sesame, or other standards. Choosing the ORB and its underlying security services will be critical to protecting PIDS, and it will influence the implementation of the access policy that a secure PIDS implementation must support.

For example, an ORB that only supports SPKM, i.e. CSI Level 0, can only authenticate clients and provide confidentiality and integrity of communication. It cannot support definition and use of security attributes beyond an access ID. Support for security attributes beyond an access ID require CSI Level 2. Therefore, using an ORB that only provides CSI Level 0 will require the PIDS to maintain its own information on the credentials of clients.

Even if an ORB's security technology supports the definition of security attributes that can be delivered to the PIDS, i.e. CSI Level 2, there are still concerns involving the trust between the client and the PIDS.

## 11.4 Trust Models

The available trust models for the PIDS is simplistic. Since the PIDS is a communications end point and does not require to make requests on other services on a client's behalf, a delegation trust model is not needed. This simplifies the model and eliminates an absolute need for a CSI Level 1 or CSI Level 2 secure ORB (although they may use them).

There are two basic trust models for the PIDS. If the PIDS and its client are implemented using CSI Level 0 or CSI Level 1 ORBs, only the first trust model can be supported. If a CSI Level 2 ORB is used, both trust models can be supported. The trust models are:

- 1• The client's identity can and is trusted to be authenticated. However, the client is unable or untrusted to deliver the valid credentials.
- 2• The client is trusted to deliver the correct credentials.

In the first model, the client ORB is required to authenticate its principal (the user) and provide authentication information to the server ORB. The methods used to accomplish principal authentication is specific to the mechanisms (e.g. DCE or Kerberos) that the selected ORB supports. Management of those identities is also specific to the mechanism. The server ORB must have a compatible mechanism that verifies the authentication information and carries out mutual authentication of the client.

With this trust model, a secure PIDS implementation must maintain and manage a map of identities to privilege attributes. CSI Level 0, 1, and 2 ORBs are able to support this trust model.

Even if the ORB has CSI Level 2 functionality, it may be a local policy that a PIDS does not trust the credentials brought forth from an authenticated client. In that case, the PIDS must maintain the map or use a default set of security attributes for requests from clients it does not trust.

In the second model, the client ORB is required to authenticate its principal and acquire its valid credentials. The methods used to accomplish principal authentication and acquisition of privilege attributes are specific to the mechanism that the selected ORB supports, such as DCE and Sesame. Management of those identities and attributes are also mechanism specific. A secure PIDS installation using this trust model must take a careful look at that management scheme and operation, evaluate it, and decide to trust it. In such a scenario, the server ORB, which has CSI Level 2 functionality, automatically verifies the credentials on invocation.

A secure PIDS built to the second model leaves management of identities and their attributes to the security services policy management system used by the ORB. The PIDS may manage security attributes for the data itself.

A secure PIDS built to the first model will have some scheme to manage trusted identities and their credentials. There is no interface or plan in the PIDS to specify this kind of management.

## 11.5 CORBA Credentials

To adhere to the credential model that supports trait specific access policy, a set of credentials must contain privilege attributes such as the identity of the client, the role in which the client is actively represented, and the sensitivity level of information to which the client is allowed access. It will be the responsibility of a PIDS implementation to advertise to potential client vendors the specifics of these attributes and how to represent them externally. A client ORB needs to ascertain certain credentials about the user and must pass them to the PIDS. An external representation of those credentials is needed so that credentials can be passed between client and server within the CORBA security services. The CORBA Security module defines the structure for this representation.

```
module Security {  
  
    const SecurityAttributeType  AccessId  = 2;  
    const SecurityAttributeType  Role      = 5;  
    const SecurityAttributeType  Clearance = 7;  
  
    struct SecAttribute {  
        AttributeType  attribute_type;  
        Opaque         defining_authority;  
        Opaque         value;  
    };  
    typedef sequence<SecAttribute> AttributeList;  
}
```

Listed above are the relevant pieces of the specification from the Security module that apply to externalizing credential information.

The Security::AccessId security attribute type could represent the person for which the ID and hence Identity object reference relates to. In constructing the value of a Security::SecAttribute of this type, the defining authority part could be the name of the PIDS ID Domain manager, and the value part can be the ID within that domain. However, if the ORB uses an underlying scheme where the value of the AccessId security attribute is supplied by some security services, such as a DCE name, a map to the PIDS ID may be needed.

The Security::Role security attribute type should represent the mandatory role. The defining authority part could take the name of the PIDS ID Domain that specifies the role, and the value can be the ID within that domain.

The Security::Clearance security attribute type can be used to represent the Sensitivity Level. For example the values could be represented by the strings, "OwnerOnly", "LevelA", "LevelB", "LevelC", "None", and "Undefined".

## 11.6 CORBA Security Domain Access Policy

In addition to a credential based scheme, CORBA defines security domains. The purpose of this section is to explain and illustrate the use of the standard CORBA security policy domain and the way in which it may be used to implement a security policy for the PIDS. This section offers a recommendation to a PIDS implementor in order to give a feel for the kinds of security policy a PIDS implementation may need to support. It should also guide the implementer in evaluating a secure ORB and available security services.

A security domain governs security (access) policy for objects that are managed within that domain. In order to make scalable administration of security policy, these domains map sets of security credentials to certain sets of rights. A right is a sort of an internal security credential.

CORBA defines a standard set of rights that are granted to principals within a security domain. A security domain administrator manages that map through the SecurityAdministration module's DomainAccessPolicy interface. Access decision then can be based on a set of required rights and the rights granted to the client by the domain access policy, by virtue of the client's credentials.

ORB security service vendors will supply a security policy management infrastructure that implements the standard CORBA rights scheme. The PIDS must use security services that can place different required rights on the PIDS interfaces. Some ORB security services may allow a security domain to create special rights. However, CORBA defines a standard set of rights: get, set, and manage. This right set will suffice to handle the PIDS.

In the model just described there is one security domain for all of the PIDS components. The CORBA rights families scheme within a single security policy domain suffices to differentiate the security nature of the methods. More generally any number of domain models can be used, such as a separate security domain for each PIDS component.

The PIDS interfaces are divided up so that for most of the interfaces one right can apply to all methods of each interface. The table below recommends the required rights for each of the PIDS interfaces. An asterisk implies that the listed right applies to all methods in the interface that are not list separately.

<b>Interface</b>	<b>Required Rights</b>
PersonIdService module	
IdentificationComponent::*	corba:g
ProfileAccess::*	corba:g
ProfileAccess::update_and_clear_traits()	corba:s
IdentityAccess::*	corba:g
IdentityAccess:: update_and_clear_traits()	corba:s
SequentialAccess::*	corba:g
IdentifyPerson::*	corba:g
IdMgr::*	corba:s
CorrelationMgr::*	corba:g
CorrelationMgr::load_profiles()	corba:s
Notification module	
Filter::*	corba:s

EventComponent::*	corba:g
Push/PullSupplier/ConsumerFactory::*	corba:m
Push/PullSupplier/Consumer::*	corba:s

Most methods on the IdentificationComponent, its subtypes and the EventComponent can be considered "get" methods. The domain access policy for the security domain should grant authenticated clients with the proper access credentials, i.e. access ID and role, with the get (corba:g) right.

All the methods in the IdMgr, Filter and consumer/supplier interfaces as well as the update\_and\_clear\_traits() and load\_profiles() operations can be considered "set" methods. They change information, and therefore these methods have a different security function other than the other methods. A client that is granted the right "get" should not necessarily be allowed access to methods that can change information. Clients that are allowed to change information in the PIDS should be granted the set (corba:s) right.

The Factory interfaces performs management of event end points, and therefore it is recommended that access to these objects should be more limited. The manage (corba:m) right may be sufficient to separate this duty from the others in a single security domain.

## 11.7 Request Content Based Policy

The CORBA standard domain access policy scheme only protects methods from invocation at the target and without regard to content of the request. The PIDS needs a more fine grained access control in order to implement the content based access policy required (e.g. access policies for different traits). The PIDS implementations must be made security aware to implement an access policy based on the value of arguments in a request. There are multiple ways to implement this policy using a secure CORBA implementation.

The CORBA Security Specification supplies two different schemes represented by an interface hierarchy, which are Security Level 1 and Security Level 2 (these should not be confused with CSI Levels 0, 1, and 2). These interfaces describe the level of security functionality that is available to security aware implementations.

### Security Level 1

For the PIDS to take advantage of CORBA security in order to implement its access policy model, the ORB must at least implement the CORBA Security Level 1 interfaces. A Security Level 1 compliant ORB supplies an interface to access the attributes of the credentials received from the client.

Using the SecurityLevel1 interfaces, which is simplistic, gives the implementation of the PIDS interfaces the ability to examine the client's credentials and compare them to the access control information that is managed as the access policy. However, the implementation of the PIDS must be security aware.

```

module SecurityLevel1 {
    Current get_current();

    interface Current {
        Security::AttributeList get_attributes(
            in Security::AttributeTypeList attributes
        );
    };
}

```

Using the Security Level 1 interfaces, each implementation of a PIDS query interface must call the `get_attributes()` function on the Current pseudo object, examine the attributes, compare to the access policy information, and make the access decision. The implementation should raise an exception if access is determined to be denied.

It is the responsibility of the client's ORB to acquire the proper credentials securely. It is the responsibility of the server's ORB to authenticate credentials received from the client, extract the security attributes from them, and make them available to the implementation through the `Current::get_attributes()` method.

## Security Level 2

Using an ORB which supplies the Security Level 2 interfaces, the implementation can be somewhat free of making the access control decision in the implementation of the query interfaces. Having an implementation that is security unaware is attractive in CORBA, because security policy decisions can be made underneath the functionality, and they have the ability to be changed without retooling the application.

As with any framework, there are several ways in which to use the Security Level 2 interfaces. One approach could be to implement a replaceable security service for the access decision. Security Level 2 describes a method in which security can be enforced by the creation of an Access Decision object. The AccessDecision object would interact with a DomainAccessPolicy object to get effective rights and compare those to rights returned from the RequiredRights interface.

Some secure ORB implementations may allow the installation of specialized Access Decision objects to be used in conjunction with specialized DomainAccessPolicy objects. In the Security Level 2 interfaces, the specification implies access control only on the invocation of a method and not the contents of the request.

```
module SecurityReplaceable {  
  
    interface AccessDecision {  
        boolean access_allowed (  
            in SecuirytLvel2::CredentialList    red_list,  
            in CORBA::Object                    target,  
            in CORBA::Identifier                operation_name,  
            in CORBA::Identifier                interface_name  
        );  
    };  
};  
}
```

Currently, the AccessDecision object specified in the SecurityReplaceable module does not take the invocation Request as an argument. It only makes an access decision based on the credentials received from the client, the target object reference and operation name, and the target's interface name. This criteria is insufficient to implement the content based access policy, if needed by a PIDS implementation to be automatically performed by the ORB.

Since the PIDS requires access control on the contents of the method invocation (such as asking for the value of the HomePhone trait), this scheme of replacing these Security Level 2 components cannot be used. ORB security services that use the standard CORBA domain access policy may use third party implementations for these components. This standard domain access policy functionality gives a the PIDS a high level of invocation protection that is orthogonal to the content based access policy. Some PIDS need the standard domain access policy functionality in addition to providing content based access policy. Therefore, another approach must be taken.

A content based access policy can be implemented in a Security Level 2 ORB by using an interceptor. A request level interceptor takes the Request as an argument and therefore, it can examine the content of the invocation arguments.

```
module CORBA {  
  
    interface Interceptor { ... };  
};
```

```
interface RequestLevelInterceptor : Interceptor {
    void client_invoke( inout Request request );
    void target_invoke( inout Request request );
};
}
```

Installing an interceptor on an ORB is ORB implementation specific, and each ORB vendor may have their own flavor of interceptors. The ORB calls the request level interceptor just before the invocation gets passed to the server implementation by using the `target_invoke()` operation. The interceptor uses the Dynamic Skeleton Interface (DSI) to examine values of the arguments of the invocation and make access decisions. These access decisions are also based on the credentials received from the client and the access policy. The interceptor will deny access to the invocation by raising the an exception. The server's ORB will transmit this exception back to the client.

The use of the interceptor scheme frees the implementation of the PIDS interfaces from the implementation of the access decision policy. If the access policy model changes, then the interceptor can be changed out without retooling the PIDS implementation.

As awareness of the need for more powerful and flexible security policy management grows, more tools to create, manage, and analyze policy will come into existence. A PIDS implementation relying on interceptors to implement its security policy may be able, with relative ease, to switch to using more robust policy services as they become developed.

## 12. Appendix - Complete IDL

---

```
//File: PersonIdService.idl

#ifndef _PERSON_ID_SERVICE_IDL_
#define _PERSON_ID_SERVICE_IDL_

#include <orb.idl>
#include <NamingAuthority.idl>
#include <Naming.idl>
#include <Trading.idl>
#include <Notification.idl>

#pragma prefix "omg.org"

module PersonIdService
{
    // -----
    // Common Data Types
    //
    typedef NamingAuthority::AuthorityId DomainName;
    typedef sequence< DomainName > DomainNameSeq;

    typedef NamingAuthority::LocalName PersonId;
    typedef sequence< PersonId > PersonIdSeq;

    struct QualifiedPersonId {
        DomainName domain;
        PersonId id;
    };
    typedef sequence< QualifiedPersonId > QualifiedPersonIdSeq;

    typedef NamingAuthority::QualifiedNameStr TraitName;
    typedef sequence< TraitName > TraitNameSeq;
    typedef any TraitValue;
    struct Trait {
        TraitName name;
        TraitValue value;
    };
    typedef sequence< Trait > TraitSeq;
    typedef TraitSeq Profile;
    typedef sequence< Profile > ProfileSeq;

    enum IdState { UNKNOWN, INVALID, TEMPORARY,
        PERMANENT, DEACTIVATED };
    typedef sequence<IdState> IdStateSeq;
    struct IdInfo {
        PersonId id;
        IdState state;
        PersonId preferred_id;
    };
    typedef sequence<IdInfo> IdInfoSeq;

    // -----
    // Miscellaneous Data Types
    //
    typedef string ComponentVersion;
    struct ComponentName {
        NamingAuthority::QualifiedName name;
        ComponentVersion version;
    };

    struct TraitSpec {
        TraitName trait;
        boolean mandatory;
    };
};
```

```

        boolean read_only;
        boolean searchable;
    };
    typedef sequence< TraitSpec > TraitSpecSeq;

    enum HowManyTraits { NO_TRAITS, SOME_TRAITS, ALL_TRAITS };
    union SpecifiedTraits switch ( HowManyTraits )
    {
        case SOME_TRAITS: TraitNameSeq traits;
    };

    struct TaggedProfile {
        PersonId id;
        Profile profile;
    };
    typedef sequence<TaggedProfile> TaggedProfileSeq;

    struct QualifiedTaggedProfile {
        QualifiedPersonId id;
        Profile profile;
    };
    typedef sequence<QualifiedTaggedProfile> QualifiedTaggedProfileSeq;

    struct ProfileUpdate {
        PersonId id;
        TraitNameSeq del_list;
        TraitSeq modify_list;
    };
    typedef sequence< ProfileUpdate > ProfileUpdateSeq;

    struct MergeStruct {
        PersonId id;
        PersonId preferred_id;
    };
    typedef sequence< MergeStruct > MergeStructSeq;

    struct TraitSelector {
        Trait trait;
        float weight;
    };
    typedef sequence<TraitSelector> TraitSelectorSeq;

    struct Candidate {
        PersonId id;
        float confidence;
        Profile profile;
    };
    typedef sequence<Candidate> CandidateSeq;

    interface CandidateIterator {
        unsigned long max_left();

        boolean next_n(
            in unsigned long n,
            out CandidateSeq ids );

        void destroy();
    };

    typedef unsigned long Index;
    typedef sequence< Index > IndexSeq;

    enum ExceptionReason {
        UNKNOWN_TRAITS,
        DUPLICATE_TRAITS,
        WRONG_TRAIT_FORMAT,
        REQUIRED_TRAITS,
        READONLY_TRAITS,
        CANNOT_REMOVE,
        MODIFY_OR_DELETE
    };
};

```

```

struct MultipleFailure {
    Index the_index;
    ExceptionReason reason;
    TraitNameSeq traits;
};
typedef sequence< MultipleFailure > MultipleFailureSeq;

interface Identity;
typedef sequence< Identity > IdentitySeq;

// -----
// Exceptions
//

exception InvalidId { IdInfo id_info; };
exception InvalidIds { IdInfoSeq id_info; };
exception DuplicateIds { PersonIdSeq ids; };
exception UnknownTraits { TraitNameSeq traits; };
exception DuplicateTraits { TraitNameSeq traits; };
exception WrongTraitFormat { TraitNameSeq traits; };
exception InvalidStates {};
exception TooMany { unsigned long estimated_max; };
exception MultipleTraits { MultipleFailureSeq failures; };

exception ReadOnlyTraits { TraitNameSeq traits; };
exception CannotRemove { TraitNameSeq traits; };
exception ModifyOrDelete { MultipleFailureSeq failures; };
exception NotImplemented {};

exception InvalidWeight {};
exception CannotSearchOn { TraitNameSeq traits; };

exception IdsExist { IndexSeq indices; };
exception RequiredTraits { TraitNameSeq traits; };
exception ProfilesExist { IndexSeq indices; };
exception DuplicateProfiles { IndexSeq indices; };

exception DomainsNotKnown { DomainNameSeq domain_names; };
exception IdsNotKnown { QualifiedPersonIdSeq ids; };

// -----
// IdentificationComponent
//

interface ProfileAccess;
interface SequentialAccess;
interface IdentityAccess;
interface IdentifyPerson;
interface IdMgr;
interface CorrelationMgr;

interface IdentificationComponent
{
    readonly attribute DomainName                domain_name;
    readonly attribute ComponentName              component_name;
    readonly attribute TraitSpecSeq              supported_traits;

    readonly attribute IdentifyPerson             identify_person;
    readonly attribute ProfileAccess              profile_access;
    readonly attribute SequentialAccess           sequential_access;
    readonly attribute IdentityAccess             identity_access;

    readonly attribute IdMgr                      id_mgr;

    readonly attribute CorrelationMgr             correlation_mgr;
    readonly attribute Notification::EventComponent event_component;
    readonly attribute CosNaming::NamingContext  naming_context;
    readonly attribute CosTrading::TraderComponents trader_components;
};

```

```

// -----
// IdentifyPerson
//

interface IdentifyPerson :
  IdentificationComponent
{
  void find_candidates(
    in TraitSelectorSeq profile_selector,
    in IdStateSeq states_of_interest,
    in float confidence_threshold,
    in unsigned long sequence_max,
    in unsigned long iterator_max,
    in SpecifiedTraits traits_requested,
    out CandidateSeq returned_sequence,
    out CandidateIterator returned_iterator )
  raises (
    TooMany,
    UnknownTraits,
    WrongTraitFormat,
    CannotSearchOn,
    DuplicateTraits,
    InvalidStates,
    InvalidWeight );
};

// -----
// ProfileAccess
//

interface ProfileAccess :
  IdentificationComponent
{
  TraitNameSeq get_traits_known(
    in PersonId id )
  raises (
    InvalidId );

  Profile get_profile(
    in PersonId id,
    in SpecifiedTraits traits_requested )
  raises (
    InvalidId,
    UnknownTraits,
    DuplicateTraits );

  TaggedProfileSeq get_profile_list(
    in PersonIdSeq ids,
    in SpecifiedTraits traits_requested )
  raises (
    TooMany,
    InvalidIds,
    DuplicateIds,
    UnknownTraits,
    DuplicateTraits );

  TaggedProfileSeq get_deactivated_profile_list(
    in PersonIdSeq ids,
    in SpecifiedTraits traits_requested )
  raises (
    NotImplemented,
    InvalidIds,
    DuplicateIds,
    UnknownTraits,
    DuplicateTraits );

  void update_and_clear_traits(
    in ProfileUpdateSeq profile_update_spec )
  raises (
    InvalidIds,

```

```

        DuplicateIds,
        NotImplemented,
        MultipleTraits );

    IdInfoSeq get_id_info(
        in PersonIdSeq ids )
        raises (
            TooMany,
            DuplicateIds );
};

// -----
// SequentialAccess
//

interface SequentialAccess :
    IdentificationComponent
{
    unsigned long id_count_per_state(
        in IdStateSeq states_of_interest )
        raises (
            InvalidStates );

    TaggedProfileSeq get_all_ids_by_state(
        in SpecifiedTraits traits_requested,
        in IdStateSeq states_of_interest )
        raises (
            TooMany,
            UnknownTraits,
            DuplicateTraits,
            InvalidStates );

    TaggedProfileSeq get_first_ids(
        in unsigned long how_many,
        in IdStateSeq states_of_interest,
        in SpecifiedTraits traits_requested )
        raises (
            TooMany,
            UnknownTraits,
            DuplicateTraits,
            InvalidStates );

    TaggedProfileSeq get_last_ids(
        in unsigned long how_many,
        in IdStateSeq states_of_interest,
        in SpecifiedTraits traits_requested )
        raises (
            TooMany,
            UnknownTraits,
            DuplicateTraits,
            InvalidStates );

    TaggedProfileSeq get_next_ids(
        in PersonId reference_id,
        in unsigned long how_many,
        in IdStateSeq states_of_interest,
        in SpecifiedTraits traits_requested )
        raises (
            TooMany,
            InvalidId,
            UnknownTraits,
            DuplicateTraits,
            InvalidStates );

    TaggedProfileSeq get_previous_ids(
        in PersonId reference_id,
        in unsigned long how_many,
        in IdStateSeq states_of_interest,
        in SpecifiedTraits traits_requested )
        raises (
            TooMany,

```

```

        InvalidId,
        UnknownTraits,
        DuplicateTraits,
        InvalidStates );
};

// -----
// IdentityAccess
//

interface IdentityAccess :
    IdentificationComponent
{
    Identity get_identity_object(
        in PersonId id )
        raises (
            InvalidId );

    IdentitySeq get_identity_objects(
        in PersonIdSeq ids )
        raises (
            InvalidIds );
};

interface Identity
{
    readonly attribute IdentificationComponent source_component;
    readonly attribute IdInfo id_info;
    readonly attribute TraitNameSeq traits_with_values;
    readonly attribute long trait_value_count;

    Trait get_trait(
        in TraitName trait_requested )
        raises (
            UnknownTraits );

    Profile get_profile(
        in SpecifiedTraits traits_requested )
        raises (
            UnknownTraits,
            DuplicateTraits );

    Profile get_deactivated_profile(
        in SpecifiedTraits traits_requested )
        raises (
            NotImplemented,
            UnknownTraits,
            DuplicateTraits );

    void update_and_clear_traits(
        in ProfileUpdate profile_update_spec )
        raises (
            NotImplemented,
            UnknownTraits,
            WrongTraitFormat,
            ModifyOrDelete,
            ReadOnlyTraits,
            CannotRemove,
            DuplicateTraits );

    void done();
};

// -----
// IdMgr
//

interface IdMgr :
    IdentificationComponent
{
    PersonIdSeq register_new_ids(

```

```

        in ProfileSeq profiles_to_register )
    raises (
        ProfilesExist,
        DuplicateProfiles,
        MultipleTraits );

PersonIdSeq find_or_register_ids(
    in ProfileSeq profiles_to_register )
    raises (
        DuplicateProfiles,
        MultipleTraits );

void register_these_ids(
    in TaggedProfileSeq profiles_to_register )
    raises (
        NotImplemented,
        IdsExist,
        DuplicateIds,
        ProfilesExist,
        DuplicateProfiles,
        MultipleTraits );

PersonIdSeq create_temporary_ids(
    in ProfileSeq profiles_to_register )
    raises (
        MultipleTraits );

PersonIdSeq make_ids_permanent(
    in PersonIdSeq ids_to_modify )
    raises (
        InvalidIds,
        DuplicateIds,
        RequiredTraits );

IdInfoSeq merge_ids(
    in MergeStructSeq ids_to_merge )
    raises (
        InvalidIds,
        DuplicateIds );

IdInfoSeq unmerge_ids(
    in PersonIdSeq ids_to_unmerge )
    raises (
        InvalidIds,
        DuplicateIds );

IdInfoSeq deprecate_ids(
    in PersonIdSeq ids_to_deprecate )
    raises (
        InvalidIds,
        DuplicateIds );
};

// -----
// CorrelationMgr
//

interface CorrelationMgr :
    IdentificationComponent
{
    readonly attribute DomainNameSeq source_domains;

    void load_profiles(
        in QualifiedTaggedProfileSeq tagged_profiles )
        raises (
            UnknownTraits,
            WrongTraitFormat,
            DomainsNotKnown );

    QualifiedPersonIdSeq get_corresponding_ids(
        in QualifiedPersonId from_id,

```

```

        in DomainNameSeq to_domains )
        raises (
            DomainsNotKnown,
            IdsNotKnown );
};

#endif // _PERSON_ID_SERVICE_IDL_

```

```

//File: NamingAuthority.idl

#ifndef _NAMING_AUTHORITY_IDL_
#define _NAMING_AUTHORITY_IDL_

#include <orb.idl>

#pragma prefix "omg.org "

module NamingAuthority
{
    enum RegistrationAuthority {
        OTHER,
        ISO,
        DNS,
        IDL,
        DCE };

    typedef string NamingEntity;

    struct AuthorityId {
        RegistrationAuthority authority;
        NamingEntity naming_entity;
    };
    typedef string AuthorityIdStr;

    typedef string LocalName;
    struct QualifiedName {
        AuthorityId authority_id;
        LocalName local_name;
    };
    typedef string QualifiedNameStr;

    exception InvalidInput {};

    interface translation_library
    {
        AuthorityIdStr authority_to_str(
            in AuthorityId authority )
            raises(
                InvalidInput );

        AuthorityId str_to_authority(
            in AuthorityIdStr authority_str )
            raises(
                InvalidInput );

        QualifiedNameStr qualified_name_to_str(
            in QualifiedName qualified_name )
            raises(
                InvalidInput );

        QualifiedName str_to_qualified_name(
            in QualifiedNameStr qualified_name_str )
            raises(
                InvalidInput );
    };
};

```

```
#endif // _NAMING_AUTHORITY_IDL_
```

```
//File: PersonIdTraits.idl

#ifndef _PERSON_ID_TRAITS_IDL_
#define _PERSON_ID_TRAITS_IDL_

#include <orb.idl>
#include <PersonIdService.idl>

#pragma prefix "omg.org"

module PersonIdTraits
{
    const PersonIdService::TraitName NULL_TRAIT = "";
    typedef any NullTraitType; // set to tk_null

    const PersonIdService::TraitName INTERNAL_ID = "PIDS/InternalId";
    typedef PersonIdService::PersonId InternalIdType;

    const PersonIdService::TraitName MERGED_IDS = "PIDS/MergedIds";
    typedef PersonIdService::PersonIdSeq MergedIdsType;

    const PersonIdService::TraitName DUPLICATE_IDS = "PIDS/DuplicateIds";
    typedef PersonIdService::PersonIdSeq DuplicateIdsType;

    const PersonIdService::TraitName CORRELATED_IDS = "PIDS/CorrelatedIds";
    typedef PersonIdService::QualifiedPersonIdSeq CorrelatedIdsType;

    const PersonIdService::TraitName EXTERNAL_IDS = "PIDS/ExternalIds";
    typedef PersonIdService::QualifiedPersonIdSeq ExternalIdsType;

    const PersonIdService::TraitName NATIONAL_HEALTH_IDS =
"PIDS/NationalHealthId";
    struct NationalHealthIdStruct {
        string country_code;
        PersonIdService::PersonId id;
    };
    typedef sequence< NationalHealthIdStruct > NationalHealthIdStructSeq;
    typedef NationalHealthIdStructSeq NationalHealthIdsType;
};

#endif // _PERSON_ID_TRAITS_IDL_
```

```
//File: HL7Version2_3.idl

#ifndef _HL7_VERSION_2_3_IDL_
#define _HL7_VERSION_2_3_IDL_

#include <orb.idl>
#include <PersonIdService.idl>

#pragma prefix "omg.org"

module HL7Version2_3
{
    typedef PersonIdService::TraitName PIDFieldName;
    typedef string PIDFieldValue;

    const PIDFieldName PATIENT_NAME = "HL7/PatientName";
    const PIDFieldName MOTHERS_MAIDEN_NAME = "HL7/MothersMaidenName";
};
```

```

const PIDFieldName DATE_TIME_OF_BIRTH = "HL7/DateTimeOfBirth";
const PIDFieldName SEX = "HL7/Sex";
const PIDFieldName PATIENT_ALIAS = "HL7/PatientAlias";
const PIDFieldName RACE = "HL7/Race";
const PIDFieldName PATIENT_ADDRESS = "HL7/PatientAddress";
const PIDFieldName COUNTY_CODE = "HL7/CountyCode";
const PIDFieldName PHONE_NUMBER_HOME = "HL7/PhoneNumber_Home";
const PIDFieldName PHONE_NUMBER_BUSINESS = "HL7/PhoneNumber_Business";
const PIDFieldName PRIMARY_LANGUAGE = "HL7/PrimaryLanguage";
const PIDFieldName MARITAL_STATUS = "HL7/MaritalStatus";
const PIDFieldName RELIGION = "HL7/Religion";
const PIDFieldName PATIENT_ACCOUNT_NUMBER = "HL7/PatientAccountNumber";
const PIDFieldName SSN_NUMBER = "HL7/SSNNumber";
const PIDFieldName DRIVERS_LICENSE_NUMBER = "HL7/DriversLicenseNumber";
const PIDFieldName MOTHERS_IDENTIFIER = "HL7/MothersIdentifier";
const PIDFieldName ETHNIC_GROUP = "HL7/EthnicGroup";
const PIDFieldName BIRTH_PLACE = "HL7/BirthPlace";
const PIDFieldName MULTIPLE_BIRTH_INDICATOR =
"HL7/MultipleBirthIndicator";
const PIDFieldName BIRTH_ORDER = "HL7/BirthOrder";
const PIDFieldName CITIZENSHIP = "HL7/Citizenship";
const PIDFieldName VETERANS_MILITARY_STATUS =
"HL7/VeteransMilitaryStatus";
const PIDFieldName NATIONALITY = "HL7/Nationality";
const PIDFieldName PATIENT_DEATH_DATE_AND_TIME =
"HL7/PatientDeathDateandTime";
const PIDFieldName PATIENT_DEATH_INDICATOR = "HL7/PatientDeathIndicator";
};

#endif // _HL7_VERSION_2_3_IDL_

```

```

//File: vCardVersion2_1.idl

#ifndef _V_CARD_VERSION_2_1_IDL_
#define _V_CARD_VERSION_2_1_IDL_

#include <orb.idl>
#include <PersonIdService.idl>

#pragma prefix "omg.org"

module vCardVersion2_1
{
    typedef PersonIdService::TraitName PropertyName;
    typedef string PropertyValue;

    const PropertyName FORMATTED_NAME = "vCard/FN";
    const PropertyName NAME = "vCard/N";
    const PropertyName PHOTOGRAPH = "vCard/PHOTO";
    const PropertyName BIRTHDAY = "vCard/BDAY";
    const PropertyName ADDRESS = "vCard/ADR";
    const PropertyName HOME_ADDRESS = "vCard/ADR;HOME";
    const PropertyName WORK_ADDRESS = "vCard/ADR;WORK";
    const PropertyName TELEPHONE = "vCard/TEL";
    const PropertyName PREFERRED_TELEPHONE = "vCard/TEL;PREF";
    const PropertyName HOME_TELEPHONE = "vCard/TEL;HOME";
    const PropertyName WORK_TELEPHONE = "vCard/TEL;WORK";
    const PropertyName VOICE_TELEPHONE = "vCard/TEL;VOICE";
    const PropertyName FAX_TELEPHONE = "vCard/TEL;FAX";
    const PropertyName MESSAGE_TELEPHONE = "vCard/TEL;MSG";
    const PropertyName CELLULAR_TELEPHONE = "vCard/TEL;CELL";
    const PropertyName BULLETIN_BOARD_TELEPHONE = "vCard/TEL;BBS";
    const PropertyName MODEM_TELEPHONE = "vCard/TEL;MODEM";
    const PropertyName CAR_TELEPHONE = "vCard/TEL;CAR";
    const PropertyName ELECTRONIC_MAIL = "vCard/EMAIL";
    const PropertyName GEOGRAPHIC_POSITION = "vCard/GEO";
    const PropertyName TITLE = "vCard/TITLE";

```

```
const PropertyName ORGANIZATION          = "vCard/ORG";
const PropertyName SOUND_ANNOTATION      = "vCard/SOUND";
const PropertyName UNIFORM_RESOURCE_LOCATOR = "vCard/URL";
};

#endif //_V_CARD_VERSION_2_1_IDL_
```