



V3_PSAF_R1_N3_2019SEP

HL7 Version 3 Standard:
Privacy and Security Architecture Framework
Release 1

**Trust Framework for Federated Authorization
Supplemental Guidance**

HL7 Normative Ballot
September 2019

Sponsored by:
Security Work Group
Community Based Care and Privacy Work Group

Copyright © 2019 Health Level Seven International ® ALL RIGHTS RESERVED. The reproduction of this material in any form is strictly forbidden without the written permission of the publisher. HL7 and Health Level Seven are registered trademarks of Health Level Seven International. Reg. U.S. Pat & TM Off.

Use of this material is governed by HL7's [IP Compliance Policy](#).

IMPORTANT NOTES:

HL7 licenses its standards and select IP free of charge. **If you did not acquire a free license from HL7 for this document**, you are not authorized to access or make any use of it. To obtain a free license, please visit <http://www.HL7.org/implement/standards/index.cfm>.

If you are the individual that obtained the license for this HL7 Standard, specification or other freely licensed work (in each and every instance "Specified Material"), the following describes the permitted uses of the Material.

A. HL7 INDIVIDUAL, STUDENT AND HEALTH PROFESSIONAL MEMBERS, who register and agree to the terms of HL7's license, are authorized, without additional charge, to read, and to use Specified Material to develop and sell products and services that implement, but do not directly incorporate, the Specified Material in whole or in part without paying license fees to HL7.

INDIVIDUAL, STUDENT AND HEALTH PROFESSIONAL MEMBERS wishing to incorporate additional items of Special Material in whole or part, into products and services, or to enjoy additional authorizations granted to HL7 ORGANIZATIONAL MEMBERS as noted below, must become ORGANIZATIONAL MEMBERS of HL7.

B. HL7 ORGANIZATION MEMBERS, who register and agree to the terms of HL7's License, are authorized, without additional charge, on a perpetual (except as provided for in the full license terms governing the Material), non-exclusive and worldwide basis, the right to (a) download, copy (for internal purposes only) and share this Material with your employees and consultants for study purposes, and (b) utilize the Material for the purpose of developing, making, having made, using, marketing, importing, offering to sell or license, and selling or licensing, and to otherwise distribute, Compliant Products, in all cases subject to the conditions set forth in this Agreement and any relevant patent and other intellectual property rights of third parties (which may include members of HL7). No other license, sublicense, or other rights of any kind are granted under this Agreement.

C. NON-MEMBERS, who register and agree to the terms of HL7's IP policy for Specified Material, are authorized, without additional charge, to read and use the Specified Material for evaluating whether to implement, or in implementing, the Specified Material, and to use Specified Material to develop and sell products and services that implement, but do not directly incorporate, the Specified Material in whole or in part.

NON-MEMBERS wishing to incorporate additional items of Specified Material in whole or part, into products and services, or to enjoy the additional authorizations granted to HL7 ORGANIZATIONAL MEMBERS, as noted above, must become ORGANIZATIONAL MEMBERS of HL7.

Please see <http://www.HL7.org/legal/ippolicy.cfm> for the full license terms governing the Material.

Ownership. Licensee agrees and acknowledges that **HL7 owns** all right, title, and interest, in and to the Materials. Licensee shall **take no action contrary to, or inconsistent with**, the foregoing.

Licensee agrees and acknowledges that HL7 may not own all right, title, and interest, in and to the Materials and that the Materials may contain and/or reference intellectual property owned by third parties ("Third Party IP"). Acceptance of these License Terms does not grant Licensee any rights with respect to Third Party IP. Licensee alone is responsible for identifying and obtaining any necessary licenses or authorizations to utilize Third Party IP in connection with the Materials or otherwise. Any actions, claims or suits brought by a third party resulting from a breach of any Third Party IP right by the Licensee remains the Licensee's liability.

Following is a non-exhaustive list of third-party terminologies that may require a separate license:

Terminology	Owner/Contact
Current Procedures Terminology (CPT) code set	American Medical Association https://www.ama-assn.org/practice-management/cpt-licensing
SNOMED CT	SNOMED International http://www.snomed.org/snomed-ct/get-snomed-ct or info@ihtsdo.org
Logical Observation Identifiers Names & Codes (LOINC)	Regenstrief Institute
International Classification of Diseases (ICD) codes	World Health Organization (WHO)
NUCC Health Care Provider Taxonomy code set	American Medical Association. Please see www.nucc.org . AMA licensing contact: 312-464-5022 (AMA IP services)

Important Note to September 2019 Ballot Voters

The September 2019 Privacy and Security Framework (PSAF) ballot is a package containing all of the Volumes developed to date under the PSAF Project Scope Statement 914. See the September Ballot Announcement:

<https://confluence.hl7.org/display/HL7/2019SEP+Announcement+of+Formation+of+Consensus+Groups>

The Privacy and Security Architecture Framework (PSAF) is comprised of:

- Volumes 1 and 2, and the Informative Guidance document for Trust Framework for Federated Authorization conceptual and behavioral models (TF4FA), which passed normative ballot in May 2018. [Being normative, it is not in scope for September 2019 ballot comments.](#)
- Volume 3 Provenance, a conceptual model addressing topics needed for trustworthy information exchange, passed normative ballot in January 2019. It has been significantly restructured as a Domain Analysis Model (DAM) for the September 2019 ballot based on input from commenters and stakeholders. [Volume 3 Provenance is in scope for September 2019 ballot comments.](#)
- Volume 4 Audit, a conceptual model for the audit service interfaces. This document was approved as normative in January 2017 under the title HL7 Version 3 Standard: Privacy, Access and Security Services (PASS) - Healthcare Audit Services Conceptual Model, Release 1 (PI ID: 1264). However, the Security Work Group missed the publication deadline, so this volume was re-balloted and past normative during the May 2019 cycle. [Being normative, it is not in scope for September 2019 ballot comments.](#)
- The Security Work Group decided to combine all volumes into one ballot package to keep them moving in tandem through balloting, publication and potential reaffirmation.

[As stated, only Volume 3 Provenance, is in scope for comments for September.](#)

Inclusion of Volumes 1, 2, and the TF4FA Guide, and Volume 4 in the September PSAF ballot package also affords voters an opportunity to review the wider privacy and security context in which the Provenance DAM was developed, and to which it contributes a significant component.

Acknowledgements

TF4FA Contributor Table	
John “Mike” Davis, VHA Security Architect Project; Authoring Lead, Principal Contributor Publishing Facilitator	
Dave Silver, Electrosoft Inc. Contributor	Diana Proud-Madruga, Electrosoft Inc. Contributor
Sponsoring HL7 Security Work Group Co-chairs	
John Moehrke, By Light	Trish Williams Professor of Digital Health Systems Flinders University School of Computer
Alexander Mense, Fachhochschule Technikum Wien, Vienna	Kathleen Connor, Book Zurman Incorporated Contributor
Chris Shawn, VHA Project and Authoring Co-lead, Contributor	
Co-sponsoring HL7 Community Based Collaborative Care [CBCP] Work Group Co-chairs	
Suzanne Gonzales-Webb, Book Zurman Incorporated	Jim Kretz, Substance Abuse and Mental Health Services Administration [SAMHSA]
Johnathan Coleman, Security Risk Solutions	David Pyke, Ready Computing

Table of Contents

1	INTRODUCTION.....	1
2	DEFINITIONS	1
3	DOMAINS AND INFORMATION OBJECTS.....	5
3.1	Federated Domain Model	5
3.2	Domain Interactions within Multidomains	7
3.2.1	Base Domain	8
3.2.2	Multi-Domain Information Objects	8
4	DOMAINS UP CLOSE	9
4.1	Normal Domain Example	9
4.2	Restricted Domain Example	10
4.3	Very Restricted Domain Example	10
4.4	Multi-Domain Information Objects	11
4.5	A Practical Way Forward.....	12
4.6	A Note on Unrestricted Information	13
5	CONCLUSION	13

List of Figures

Figure 1: Domain Model.....	2
Figure 2: Policy Bridging.....	5
Figure 3: Sensitivity Layers in a Compound Federated Domain	7
Figure 4: Access Control Model	7
Figure 5: Normal Domain Example.....	9
Figure 6: Restricted Domain Example.....	10
Figure 7: Very Restricted Domain Example.....	11

List of Tables

Table 1: Problem List Example	12
-------------------------------------	----

List of Appendices

APPENDIX A: HL7 Classification Codes (Normative).....	14
APPENDIX B: Types of Domain Privacy and Security Policy.....	15
APPENDIX C: HIMSS Interoperability Definition.....	16

1 INTRODUCTION

It is often necessary to establish trust between partners in the exchange of protected health information. The exchange may involve a request from one party to another or a direct push. The parties may be health care organizations conducting business as well as patients directing exchanges among or requesting information from health care organizations.

This paper defines terms and concepts foundational to developing trust between parties in such exchanges. It draws upon international standards and the Health Level 7 (HL7) privacy and security standards for interoperability. These include standards for messaging, information classification, and terminology, as well as access control methods and services.

2 DEFINITIONS

Security Domain. A set of subjects, their information objects, and a common security policy (NIST Special Publication 800-33).

Security Policy Domains. A security policy domain is a set of objects to which a security policy applies for a set of security related activities and is administered by a security authority. (Note that this is often just called a security domain and are here treated as equivalent.) The

objects are the domain members. The policy represents the rules and criteria that constrain activities of the objects to make the domain secure. (OMG Security Services Specification (OMG SEC))

Security Authority: A security authority must be identifiable and responsible for defining the policies to be applied to the domain but may delegate that responsibility to a number of sub-authorities, forming subdomains where the subordinate authorities' policies are applied. Subdomains may reflect organizational subdivisions or the division of responsibility for different aspects of security. Typically, organization-related domains will form the higher-level superstructure, with the separation of different aspects of security forming a lower-level structure. (OMG SEC)

Domain Characterization. A domain is characterized by a domain identifier, domain name, domain authority, and domain qualifier (ISO/TS 22600-2:2006).

Subdomain: A domain might consist of sub-domains (which will inherit and might specialize policies from the parent domain). The smallest-scale domain might be an individual workplace or a specific component within an information system. (ISO 22600-2)

Superdomain: Domains can be extended into super-domains, by chaining a set of distinct domains and forming a common larger-scale domain for communication and co-operation. (ISO 22600-2)

Domain Attributes

- Within a security domain, all information objects exist at the same level of sensitivity (Note: this is synonymous with the “confidentiality classification” found in HL7 HCS.)
- Members of a domain may have different security attributes, such as read, write, or execute permissions on information objects.
- Security domains are not bound by systems or networks of systems.
- A security domain's objects may reside in multiple systems.




Figure 1: Domain Model

Figure 1 illustrates major characteristics and relationships between domain and their sub-domains

- Domains = N1, N2, N3, N4, R1, R2
- Subdomains = N2, N3, N4

Objects: are the domain members (OMG SEC)

Policy: The policy represents the rules and criteria that constrain activities of the objects to make the domain secure. (OMG SEC)

A policy is the formulation of the concept of requirements and conditions for trustworthy creation, collection, storage, processing, disclosure, retention, transmission, and use of sensitive information. (ISO 22600-2)

Policy: A set of *rules*, an identifier for the *rule-combining algorithm* and (optionally) a set of *obligations* or *advice*. May be a component of a *policy set*. (OASIS XACML v3.0)

Policy Set: A set of policies, other policy sets, a policy-combining algorithm and (optionally) a set of obligations or advice. May be a component of another policy set.

Advice: A supplementary piece of information in a *policy* or *policy set* which is provided to the *PEP* with the *decision* of the *PDP*.

Obligation: An operation specified in a *rule*, *policy* or *policy set* that should be performed by the *PEP* in conjunction with the enforcement of an *authorization decision*

Target: An element of an XACML *rule*, *policy*, or *policy set* which matches specified values of *resource*, *subject*, *environment*, *action*, or other custom attributes against those provided in the request context as a part of the process of determining whether the *rule*, *policy*, or *policy set* is applicable to the current decision.

Security Policy: A security policy is the complex of legal, ethical, social, organizational, psychological, functional, and technical rules for ensuring trustworthiness of health information systems. (ISO 22600-2)

The following concepts of Basic, Composite and Meta policy are adapted (additions in bold) from Ponder.¹

Basic Policy: The body of a basic policy consists of one or more policy elements. Several of these elements are common to all basic policy types: the subject, the target, the when-constraint, as well as import statements, constant definitions and external specifications. Other policy elements are specific to a particular policy type.

Policy elements can be specified in any order. The subject and the target for a basic policy are specified using domain scope expressions or by a formal identifier of type set. Actual parameters for subjects and targets are domain scope expressions. A subject or target keyword can be optionally followed by the Interface Definition Language (IDL) type of the objects specified. A name can also be assigned to subjects and targets in order to reuse it in expressions within the policy. The keywords subject and target themselves can also be used to refer to the current subject/target during the execution of the policy. Each basic policy can also optionally specify a when-constraint element that limits the applicability of the policy.

Basic policies cannot contain other policies. Although they usually need an explicit subject an exception is when a basic policy is specified as part of a Role, in which case the subject domain of the Role/**Clearance** is the implicit subject.

- Authorization policies: For both positive and negative authorization policies, the specification of the following policy elements is required. An authorization policy must contain the following policy elements:
 - subject (except in roles)
 - target
 - action (**roles only**)
 - rule (clearances only)²
- Obligation policies: An obligation policy must contain the following policy elements:
 - subject (except in roles)
 - action (**roles only**)
 - event
 - rule (clearances only)

¹ Ponder: A Language for Specifying Security and Management Policies for Distributed Systems, Version 1.11, 18 January 2000 <http://www.doc.ic.ac.uk/research/technicalreports/2000/DTR00-1.pdf>

² Rules are added to account for attribute-based access control which include the subject, target and a rule that links them

- **Refrain policies:** A refrain policy must contain the following policy elements:
 - subject (except in roles)
 - action (**roles only**)
 - rule (clearances only)
- **Delegation policies:** One or more positive authorization and/or delegation policies must always be associated with a delegation policy (both positive and negative). The only required policy element for a delegation policy is the specification of a grantee. Subjects and targets, if not specified, default to the aggregated subjects and targets of the associated authorization/delegation policies. If actions to be granted are not specified they default to those of the associated authorization/delegation policies.

Composite Policy: Used to group a set of related policy specifications within a syntactic scope with shared declarations in order to simplify the policy specification task for large distributed systems. **Five** types of composite policies are provided: groups, roles, relationships and management structures. Constraints can be specified to limit the applicability of policies based on time or values of the attributes of the objects to which the policy refers.

ALSO

There is a need to group a set of related policy specification within a syntactic scope with shared declarations in order to simplify the policy specification task for large distributed systems. This is a common concept in many programming environments and is the main motivation behind composite policy types in Ponder. At run-time, the set of policies defined in a composite policy, together with any constraints applying to the composite policy would be stored within a domain. All composite-policies can include types and instance definitions as well as nested groups. However, roles cannot include nested roles, relationships or management structures, and relationships cannot contain nested relationships or management structures. All composite-policies can be specified as types from which multiple instances can be created.

Multidomain Information Object (aka Compound Domain): A collection of objects from different security domains perceived by users as a single information object. In compound security domains, additional policies are written that apply to the newly created multidomain information objects. The multidomain information security policy states the privileges that a user must have to view, print, create, delete, or transfer multidomain information objects between information systems. It cannot be assumed that the Multidomain Information Object policies are simply inherited from the subdomains. [ASTM E2595]

Policy Bridging: The process used to derive (negotiate) the set of common, domain-specific security and privacy policies required for trustworthy co-operation between collaborating domains. (Derived from ISO 22600-1)

Management Structures: A management structure defines the configuration of roles and relationships in organizational units in terms of the required instances of the roles. For example, it would be used to define a management structure (type) for creating branches in a bank or departments in a university. Management structures can include any nested composite-policy.

Relationships: Relationships specify policies pertaining to the relationship rather than the individual participating roles.

Role: A role groups the policies specifying the duties and rights relating to a position within an organization. A role is thus a particular type of group in which all policies have the same subject domain. A role can contain basic policies and groups of basic policies but not nested roles, relationships or management structures. The role instantiation declaration may specify an optional path name, which is to be used as the subject domain for the role. This assumes the subject domain has already been created in the domain hierarchy. If the subject domain is not specified then a domain with the name of the role instance is implicitly created and used as the subject domain i.e. the subject for policies within the role.

Sensitivity: The characteristic of an IT resource which implies its value or importance and may include its vulnerability. (ISO 7492-2)

Privacy metadata for information perceived as undesirable to share. (HL7 Healthcare Classification System)

- Sensitive information is data that must be protected from unauthorized access and disclosure to safeguard the privacy or security of an individual or organization.
- Classification is the act or process by which information is determined to be sensitive or non-sensitive.
- The appropriate classification level is determined by the disclosure risks of the information, which usually are identified by the magnitude, amount or kind of damage that could be caused by disclosure.

3 DOMAINS AND INFORMATION OBJECTS

3.1 Federated Domain Model

The federated domain model describes the components of negotiated trust between two or more individual domains that provide a basis for assuring secure interchange of protected health information. Exchange occurs under the control of shared security and privacy policies managed by a common Federation Authority. The shared intersection of data, users and policy defines the elements of the Federated Domain.

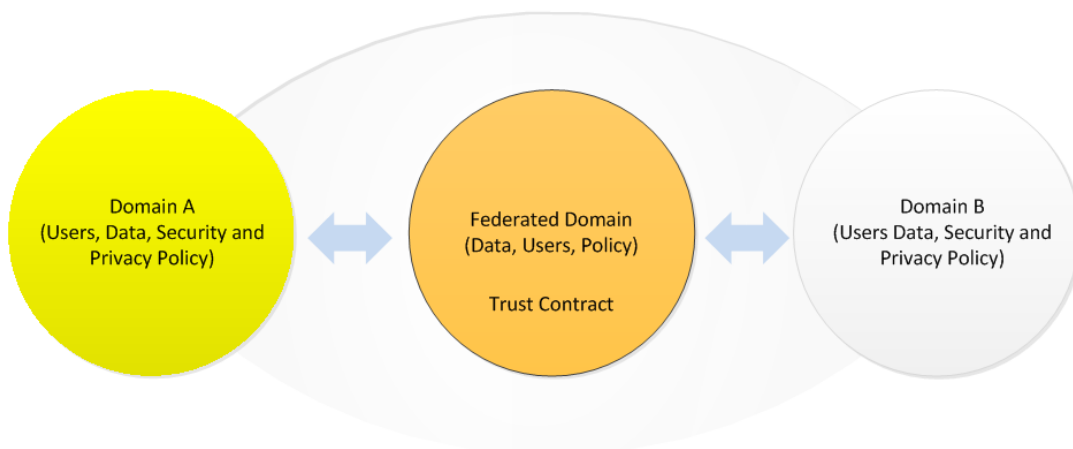


Figure 2: Policy Bridging

Figure 2 illustrates the result of a system where policy bridging has derived (negotiated) the set of common, domain-specific security and privacy policies required for trustworthy co-operation between collaborating domains (Federated Domain Composite Policy). (Derived from ISO 22600-1)

Domain Authorities agree to which users and what data are to make up the shared Federated Domain, and the rules governing information sharing. A Trust Contract (aka Federation Agreement) provides confidence that the mutual agreements will be honored. In a federation, each domain retains most of its authority while agreeing to afford the other limited rights.

- Sensitivity³ Singularity. Under domain rules, a domain may only contain single data sensitivity, however, to achieve real-world conditions, the full description of all desired interactions among cooperating partners involves chaining together of multiple individual federated subdomains representing all included sensitivities. The resulting extended domain forms a federated multidomain of communication and cooperation that is characterized by an agreed upon overall composite security and privacy policy.
- Federation agreement. The federation agreement records:
 - Rights given to both sides, such as the kind of access allowed,
 - Trust each has in the other,
 - An agreement as to how policy differences are handled, for example, the mapping of roles in one domain to roles in another.

Within the Federated Domain, sharing rules are specific to information sensitivity. Consequently, a complete description of sharing for all allowed sensitivities is provided by the aggregation of independent domains each at its own sensitivity level. For example, compound information objects such as a subject of care Medical History shared between two different organizations (Domains) might include Medications, Diagnosis, Allergies, and Immunizations. This information object inherits the top-level classifications of the most restrictive classifications of any of the instances of any of its included subordinate information objects.

Real world information objects may include multiple sensitivities which from the user's point of view, are perceived as layers within a Multidomain Information Object. Each layer represents a unique intersection of users, data and Federated Domain Sensitivity characterized by a unique domain sensitivity value. Combined together these layers define all possibilities within the Multidomain Information Object. The Compound Federated Domain is the resulting collection of all included subordinate information objects, users and merged policy. See Figure 3 Sensitivity Layers in a Compound Federated Domain.

³ For the purposes of this paper, "sensitivity" refers to the confidentiality classification of the data as defined in HL7 HCS: "Security label metadata classifying an IT resource (clinical fact, data, information object, service, or system capability) according to its level of sensitivity, which is based on an analysis of applicable privacy policies and the risk of financial, reputational, or other harm to an individual or entity that could result if made available or disclosed to unauthorized individuals, entities, or processes. Confidentiality classifications are hierarchical levels in a multilevel policy that permits a user with a clearance classification equal to the classification label assigned to an information resource to "read down," (i.e., to read less classified information objects, and to "write up", i.e., create information resources that are more highly classified, but does not permit the user to reclassify an information resource to a lower level of confidentiality).




Figure 3: Sensitivity Layers in a Compound Federated Domain

Within these domains, sensitive information may additionally be organized by particular attributes through labels that provide further segmentation into category subdomains.

3.2 Domain Interactions within Multidomains

Figure 4 illustrates subdomains (blue or purple circles) within an existing domain defined by their HL7 Confidentiality codes (Normal, Restricted and Very Restricted) and category values⁴. The full description of these and other domains are described in the HL7 Privacy and Security Healthcare Classification System (HCS) (summarized in Appendix A) and corresponding HL7 vocabulary. The model logically elaborates Figure 3's sensitivity layers.




Figure 4: Access Control Model

⁴ **HCS Security Category:** The HCS Security Category Named Tag Set may include multiple Tag Set Name fields for the different Security Categories such as Sensitivity, Compartment, Privacy Policies and Laws, Integrity, and Provenance. Each Security Category Name Tag Set field includes one or more security tags valuing the label field. [HCS]

Figure 4 also illustrates Domain Confidentiality and controlled access to classified information in accordance with the Bell-LaPadula model for enforcing access control. In this model, Users with lower authorization may “Write up” (indicated by the arrows) but not “Read up” to layers with higher sensitivity. In addition, users with higher authorizations (e.g., Very Restricted), may “Read Down” but may not “Write Down” to Domains of lower sensitivity. The domains in Figure 4 are separated by the grey border between domains. In a compound Federated Domain such as described here, Normal is not a subdomain of either Restricted or Very Restricted as these are different sensitivities.

Furthermore, each domain may include subdomains (additional domain organizational categories indicated by blue or purple lettered circles). The Bell-LaPadula rules apply to access control to these subdomains with one additional caveat regarding read down. This caveat is that while users with higher authorizations may read down to lower domains, they may or may not be able to access the lower subdomain categories of a Multidomain Information Objects depending on policy. For example, a user with Very Restricted clearance may read Restricted information but may or may not have access (clearance) for categories A, AB or B.

3.2.1 Base Domain

A Base Domain (Domain) consists of Users, Data, and a controlling security policy under control of a domain authority. The Base Domain information objects exist at a single level of sensitivity.

For example, consider a base information object that consists of an instance of a single Medication. At a minimum, the instance is assigned one and no more than one valid single classification of (VR, R, N, M, L or U). For example, a base domain information object attributes would look like:

Medication (Normal)		
Medication	Classification	Domain
Example Med 1	Normal	Domain 1

3.2.2 Multi-Domain Information Objects

Real-world information objects are complex, with properties that extend beyond those of the Base Domains. For example, the information object representing all Medications for a single subject of care may have multiple instances at various classifications. In this case, the Medication object acquires the overall classification of its most sensitive member instance perceived by its users as a single multi-domain information object.

Compound Healthcare Domains are characterized by their information object sensitivities. A Compound Healthcare Domain Sensitivity (DS) is a unique single valued function of each of six “HL7 Privacy and Security Classification System (HCS)” defined Classification values of Unrestricted, Low, Medium, Normal, Restricted and Very Restricted. For any Domain Sensitivity covered by this specification, there must exist at least one, non-null Classification value.

Each Classification function is further defined by its arguments. The arguments are the HCS defined categories of Sensitivity, Integrity and Compartment.⁵ The values of these arguments are specified by the HL7 Normative Privacy and Security Vocabulary.

4 DOMAINS UP CLOSE

This section provides expository examples of Normal, Restricted and Very Restricted Domains. Each domain has a corresponding set of users and data characterized by its classification. Each domain satisfies the core definition of domain and sub-domain.

4.1 Normal Domain Example

The common Normal domain is the first domain for which user clearances are appropriate. Normal domains may have subdomains further defined by categories; such as “Pharmacy Use Only”, Care Team etc. enforcing need to know and least privilege policies.

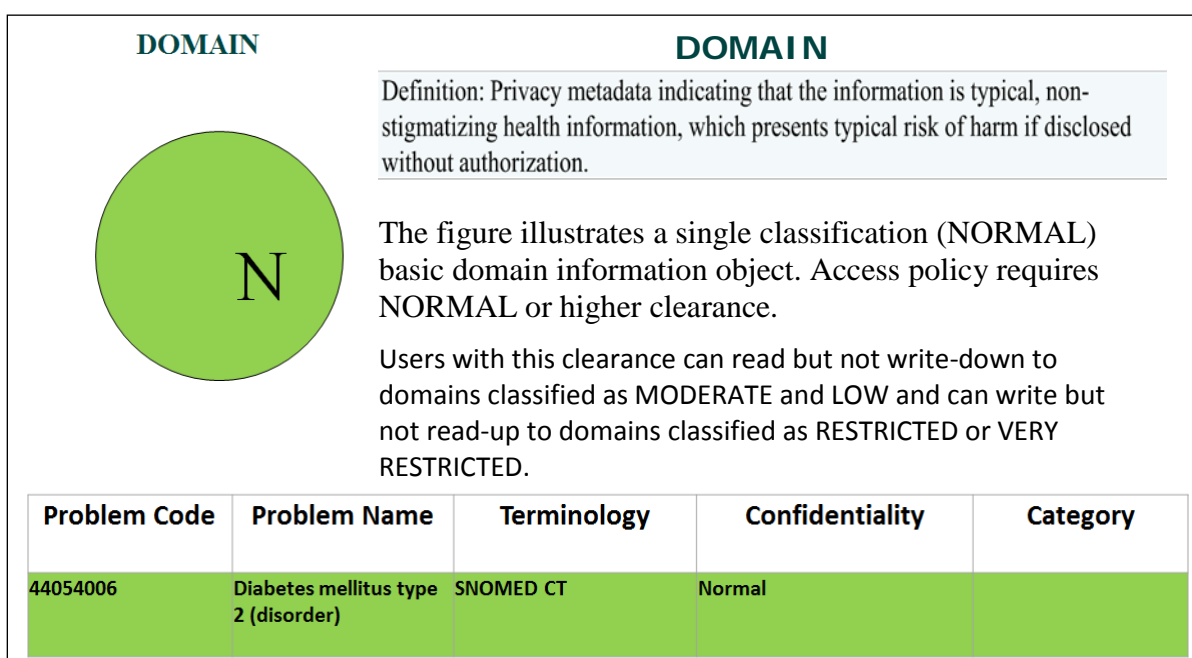


Figure 5: Normal Domain Example

⁵ Distinguish Domain Sensitivity (DS) assigned to a Domain from the HCS value of Sensitivity assigned to an information object. The HCS definition of “sensitivity” is “Privacy metadata for information perceived as undesirable to share.”

4.2 Restricted Domain Example

Restricted domains are characterized by rich user permissions and access to data further characterized by Sensitivity, Integrity and Compartment. Users with Restricted Permission can read down to all lower levels but cannot access (read up) to information classified as Very Restrictive.

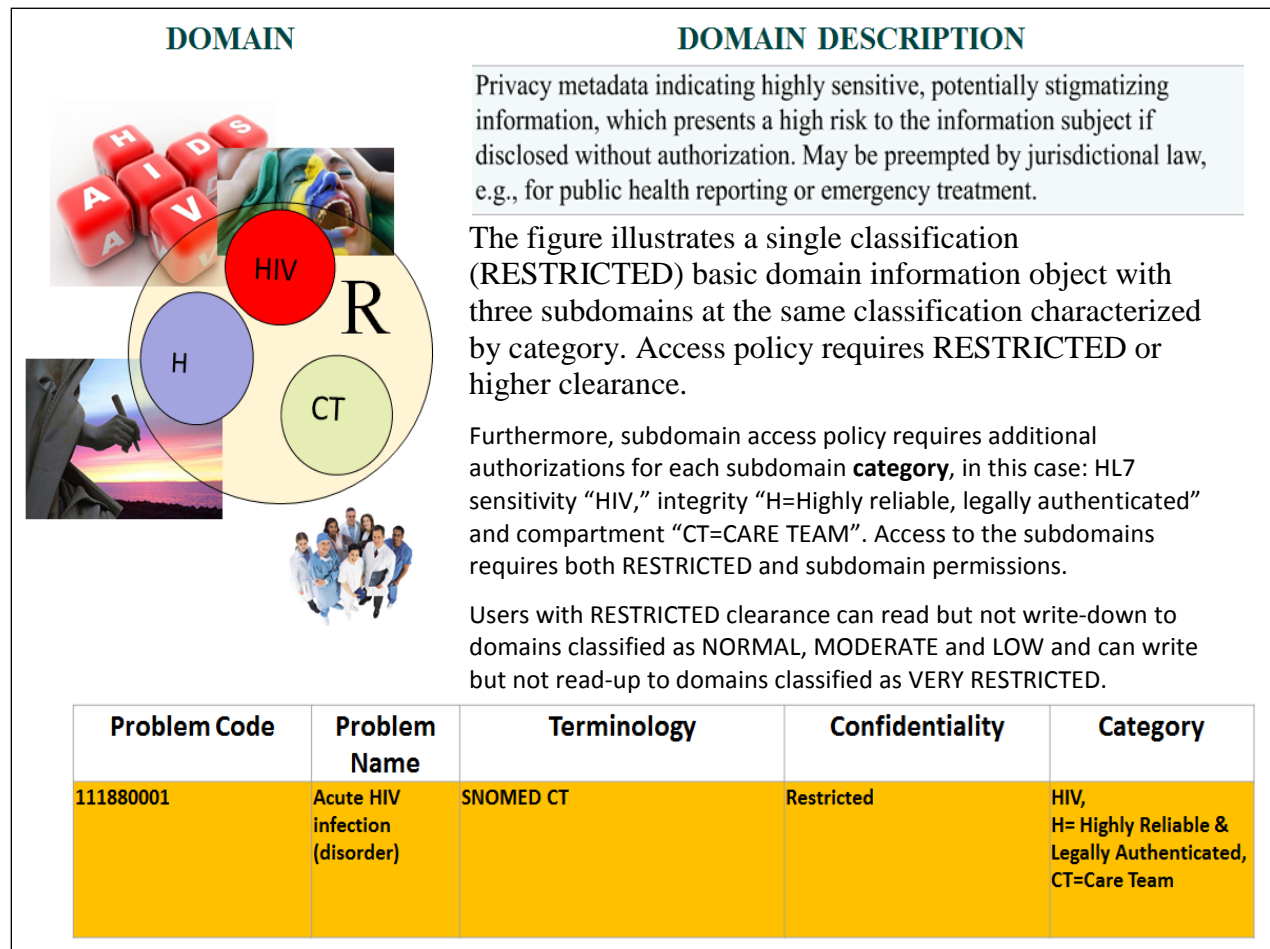


Figure 6: Restricted Domain Example

4.3 Very Restricted Domain Example

Very Restricted domains have fully defined subdomains of Sensitivity, Integrity and Compartment. Very Restricted domains are characterized by rich user permissions and access to data further characterized by Sensitivity, Integrity and Compartment. Users with Very Restricted Permission can read down to all levels but must in addition possess appropriate Sensitivity, Integrity and/or Compartment clearances to access this information, even though they can read down to other subordinate classifications.

Very Restricted domains have fully defined subdomains of Sensitivity, Integrity and Compartment. Very Restricted domains are characterized by rich user permissions and access to data further characterized by Sensitivity, Integrity and Compartment. Users with Very Restricted Permission can read down to all levels but must in addition possess appropriate Sensitivity, Integrity and/or Compartment clearances to access this information, even though they can read down to other subordinate classifications.

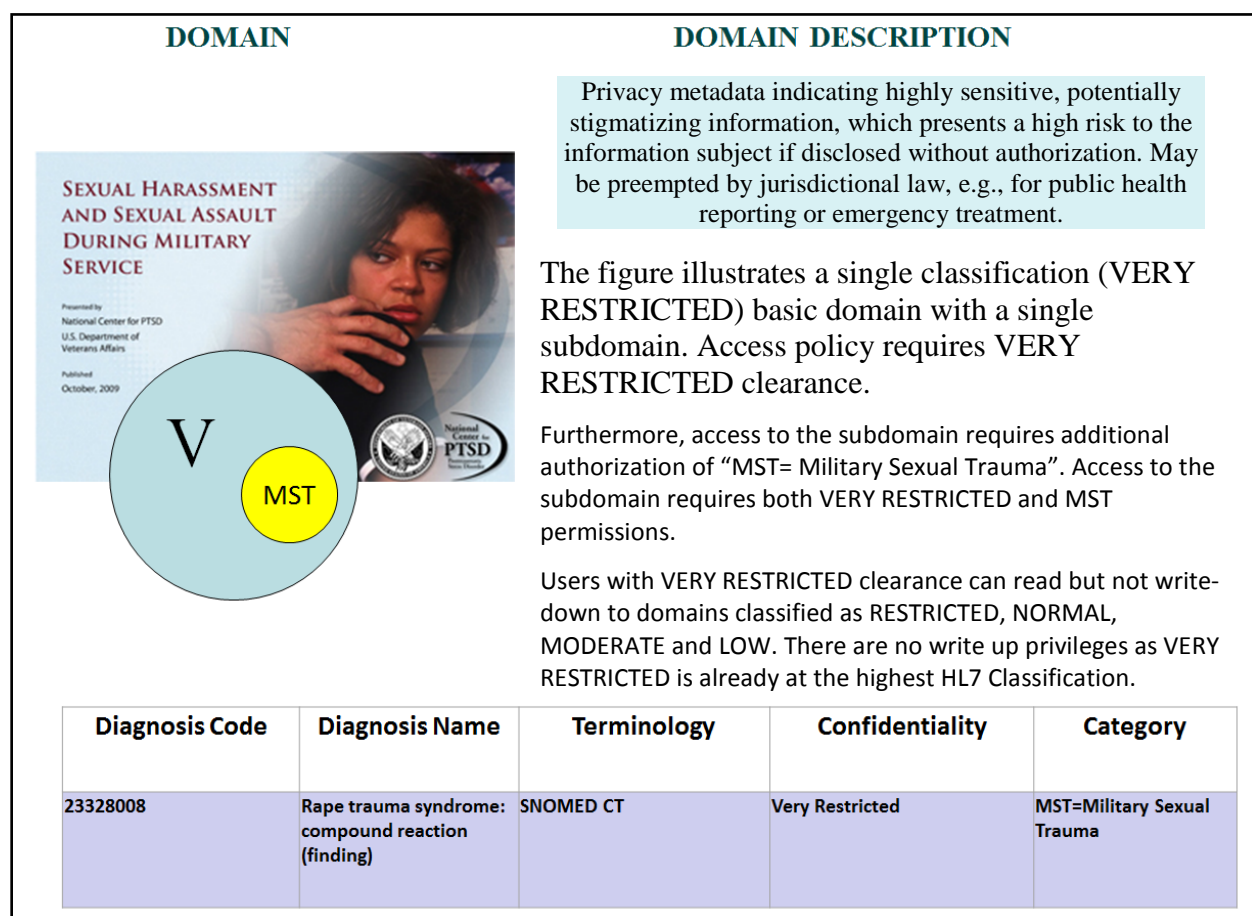


Figure 7: Very Restricted Domain Example

4.4 Multi-Domain Information Objects

Recall the definition of a multi-domain information object as: “A collection of objects from different security domains perceived by users as a single information object.”

Figure 4 illustrates a real-world information object consisting of a patient problem list. The problem list includes a number of entries which follow rules for individual domains. From the perspective of a user, this collection appears to be a single information object. The information object is classified overall “RESTRICTED//HIV” based upon the highest classifications and categories found.

Access policy for the Problem List information object requires that users hold both the RESTRICTED (or higher) clearance and HIV category. Furthermore, authorization for the HIV subdomain is required.

This information object itself is not a domain as it contains objects at both the RESTRICTED and NORMAL classification. However, using data segmentation with a Privacy Protective Service to redact or mask the RESTRICTED information, users with a clearance for NORMAL would have access, and the resulting single classification information object would be in one information domain.

In compound security domains, additional policies need to be written that apply to the newly created multi-domain information objects.

Table 1: Problem List Example

RESTRICTED//HIV Problem List			
Problem Name	Problem Code	Classification/ Category	Problem Status
Acute HIV infection (Disorder)	111880001	R, HIV	Active
Diabetes mellitus type 2 (disorder)	44054006	N	Resolved
Asthma (disorder)	19597001	N	Inactive
Coronary artery atheroma (disorder)	67682002	N	Inactive
Hyperlipidemia (disorder)	5582204	N	Active
Hypertension associated with transplantation (disorder)	427889009	N	Active

4.5 A Practical Way Forward

While security and privacy policies can be written in detail at both the classification and detailed category and sub-category level, in reality arbitrary policies for most health care organizations only exist with respect to patient defined policies (Don't share my weight, Don't share my HPV immunization, etc.). At the organization level, security classifications at the Restricted and Very Restricted level are few, most deriving from State and Federal law.

For example, HIV is widely covered under state law as protected, requiring patient authorization to disclose. Similarly, Federal law (42 CFR Part 2 and 38USC7332) require patient authorization to disclose drug and alcohol substance use disorder.

Such restrictions carry their own disadvantages. For example, a 42CFR patient who does not authorize disclosure of their substance use disorder information may receive a prescription for opioids from a clinician unaware of their addiction.

One attractive approach to dealing with this problem which is receiving attention is called "share with protections". In a share with protections scheme, information is always shared (except when an explicit patient request to opt-out or restrict information has been made and accepted). In this scheme, sensitive information is labeled appropriately (e.g. Restricted//SUD)

and shared with agreement among partners to only allow access by clinicians with appropriate clearances.

Furthermore, considering that the number of protected conditions under law is relatively limited, additional simplification is possible by only requiring a single label of “Restricted” to include all subordinate protected categories. As information is always shared, Clinical Decision Support can also be employed to evaluate and share protected information with clinicians on the basis of patient “Emergency” or “Patient Safety” conditions (e.g. Drug-Drug interactions), even though they may possess Restricted access clearances.

4.6 A Note on Unrestricted Information

The HL7 HCS Vocabulary also includes a classification type of unrestricted (U). This classification includes publicly available information that does not require the sender/receiver to consider additional policies when making access control decisions. This type of information includes for example, business name, phone, email and or physical address. The following considerations apply to unrestricted information:

- If information individually classified as U is mixed in the same block within text of higher classification, then the higher classification will apply to the entire block.
- When not associated with other labeled data, sources consisting entirely of Unrestricted information are often unlabeled in actual use, although they may be.
- Personally identifiable information (PII) can be classified U with certainty when obtained from a dedicated unrestricted source such as a Master Patient Index.
- If there is no law that specifies a classification requirement, then (MPI data) can be classified as “U” unclassified

5 CONCLUSION

Classification schemes leveraging attribute-based access control can provide detailed and fine-grained access to protected information. The close relation of policy-based labeling to the underlying data, should be a simplifying approach that provides for straight-forward insight into the functioning and understanding of security system operations. In addition, the relative flexibility of being able to change system response by way of managing policies contained in software provides for relatively direct and inexpensive management capable of change. Finally, the availability of extensive standardized vocabulary and codes sets within HL7 provides a platform for ready-made interoperability.

Data classification and ABAC seem to be ideal for FHIR as well, suitable for a single instance of a resource as well as for the entire medical history. Policy evaluation engines are mature and readily available (e.g., OASIS XACML vendors). No other security scheme offers the same level of simple interoperability.

Finally, ABAC systems can be implemented as services bound to an organizations security system. This means that changes to underlying EHRs themselves are unnecessary. In this way security and privacy solutions can be deployed at whatever scale is needed, simultaneously supporting both organizational (PHI Protection) and patient (Privacy Protection) requirements.

APPENDIX A: HL7 Classification Codes (Normative)

Confidentiality Code	Print Name	Definition	Business Use
V	Very Restricted	Definition: Privacy metadata indicating extremely sensitive, likely stigmatizing information, which presents a very high risk if disclosed without authorization. This information must be kept in the highest confidence.	Access to this information is allowed only for those with a corresponding Restricted clearance.
R	Restricted	Definition: Privacy metadata indicating highly sensitive, potentially stigmatizing information, which presents a high risk to the information subject if disclosed without authorization. May be preempted by jurisdictional law, e.g., for public health reporting or emergency treatment.	Access to this information is allowed only for those with a corresponding Restricted clearance.
N	Normal	Definition: Privacy metadata indicating that the information is typical, non-stigmatizing health information, which presents typical risk of harm if disclosed without authorization.	Access to this information is allowed only for those with a corresponding Normal clearance.
M	Medium	Definition: Privacy metadata indicating moderately sensitive information, which presents moderate risk of harm if disclosed without authorization.	Not a user permission. Provides staff guidance for use within health care environment. For electronic exchange, policy may be set by obligation/ terms of use
L	Low	Definition: Privacy metadata indicating that the information has been de-identified, and there are mitigating circumstances that prevent re-identification, which minimize risk of harm from unauthorized disclosure. The information requires protection to maintain low sensitivity.	Not a user permission. Provides staff guidance for use within health care environment. For electronic exchange, policy may be set by obligation to comply with a data use agreement.
U	Unrestricted	Definition: Privacy metadata indicating that the information is not classified as sensitive.	Publicly available information

APPENDIX B: Types of Domain Privacy and Security Policy

Domain policy is evaluated in terms of rules and rulesets and associated rule-combining algorithms. Individual rules are evaluated in terms of decision specific information attributes and conditions (ACI). Accordingly, distinct policy categories (types of rules), can be distinguished by ACI, may be evaluated under policy enforcement control of an HL7 Access Control System (ACS).

Policy Type 1: Environment ACI. Policies regarding environment ACI such as a patient authorization, location, time of day, relevant law, etc. Also, contextual policies regarding ACI associated with Purpose of Use, data use agreements, memorandum of understanding, and obligations where one or more parties must agree to accept certain terms and/or accept the responsibility to enforce obligations. This potentially includes all classifications and explicitly M, and L.

Policy Type 2: Request ACI. Policies regarding attributes of an initiators request (e.g., Policy distinctions for requests for information whose value is less than 5 million dollars as opposed to those more than, etc.)

Policy Type 3: Data ACI. Policies regarding data and types of data resources (e.g., Community Care, Department of Defense information, immunizations information, categories of data specifically protected by law, etc.),

Policy Type 4: User ACI.

- a. Policies regarding roles held by recipients (e.g. Care Team, workflow permissions, licensed clinician, organization, HIE, etc.),
- b. Policies regarding clearances held by recipients (e.g., N, R, VR as well as policies regarding data classifications for which a user must additionally hold corresponding sensitivity attributes (S, I and C),
- c. Policies regarding Access Control Lists (ACL).

Policy Type 5: Target ACI. Policies involving ACI of a target/target group including ACI of hierarchical and Functional Groups, Integrities, authorities, ownership, privacy and security data classifications.

Policy Type 6: Retained ACI. Policies regarding retained information ACI from previous decisions.

APPENDIX C: HIMSS Interoperability Definition

Definition of Interoperability © HIMSS 2013 - Approved by the HIMSS Board of Directors April 5, 2013 <http://www.himss.org/library/interoperability-standards/what-is-interoperability>

In health care, interoperability is the ability of different information technology systems and software applications to communicate, exchange data, and use the information that has been exchanged.⁶ Data exchange schema and standards should permit data to be shared across clinicians, lab, hospital, pharmacy, and patient regardless of the application or application vendor.⁷ Interoperability means the ability of health information systems to work together within and across organizational boundaries in order to advance the health status of, and the effective delivery of health care for, individuals and communities.⁸

There are three levels of health information technology interoperability:⁹ 1) Foundational; 2) Structural; and 3) Semantic.

1. “Foundational” interoperability allows data exchange from one information technology system to be received by another and does not require the ability for the receiving information technology system to interpret the data.
2. “Structural” interoperability is an intermediate level that defines the structure or format of data exchange (i.e., the message format standards) where there is uniform movement of health data from one system to another such that the clinical or operational purpose and meaning of the data is preserved and unaltered. Structural interoperability defines the syntax of the data exchange. It ensures that data exchanges between information technology systems can be interpreted at the data field level.
3. “Semantic” interoperability provides interoperability at the highest level, which is the ability of two or more systems or elements to exchange information and to use the information that has been exchanged.¹⁰ Semantic interoperability takes advantage of both the structuring of the data exchange and the codification of the data including vocabulary so that the receiving information technology systems can interpret the data. This level of interoperability supports the electronic exchange of health-related financial data, patient-created wellness data, and patient summary information among caregivers and other authorized parties. This level of interoperability is possible via potentially disparate electronic health record (EHR) systems, business-related information systems, medical devices, mobile technologies, and other systems to improve wellness, as well as the quality, safety, cost-effectiveness, and access to health care delivery.¹¹

⁶ HIMSS Dictionary of Healthcare Information Technology Terms, Acronyms and Organizations, 2nd Edition, 2010, Appendix B, p190, original source: Wikipedia

⁷ American Academy of Family Physicians (AAFP), Center for Health IT, 2013.

⁸ HIMSS Dictionary of Healthcare Information Technology Terms, Acronyms and Organizations, 3rd Edition, 2013, p. 75.

⁹ National Committee on Vital and Health Statistics (NCVHS) Report on Uniform Data Standards for Patient Medical Record Information, July 6, 2000, pp. 21-22.

¹⁰ Institute of Electrical and Electronics Engineers. IEEE Standard Computer Dictionary: A Compilation of IEEE Standard Computer Glossaries. New York, NY: 1990.

¹¹ HIMSS Dictionary of Healthcare Information Technology Terms, Acronyms and Organizations, 2nd Edition, 2010, Appendix B, p190, original source: HIMSS Electronic Health Record Association.