Alternative Approaches for Supporting Lattice-based Access Control (LBAC) in the Fast Healthcare Interoperability Resources (FHIR) Standard

Steven Demurjian[1], Thomas Agresta[2], Eugene Sanzi[1] and John DeStefano[3]
[1]*Department of Computer Science & Engineering, University of Connecticut, 371 Fairfield Way, Storrs, Connecticut, U.S.A.*
[2]*Department of Family Medicine, University of Connecticut Health Center, 263 Farmington Avenue, Farmington, Connecticut, U.S.A.*
[3]*SMC Partners, LLC, 10 Columbus Boulevard,, Hartford, Connecticut, U.S.A.*

Keywords: Healthcare, Multi-level Security, Lattice based Access Control, FHIR, Sensitivity Level.

Abstract: A major challenge in the healthcare industry is the selective availability, at a fine-grained level of detail, of a patient's data to the various clinicians, nurses, specialists, home health aides, family members, etc. where the decision of who can see which information at which times is controlled by a patient. The information includes: contact and demographics, current conditions, medications, test results, past medical history, history of substance abuse and treatment, mental health information, sexual health information, records relating to domestic violence, reproductive health records, and genetic information. To control sensitivity, multi-level security (MLS) using lattice-based access control (LBAC) can be used to extend the traditional linear sensitivity levels of mandatory access control with the ability to define a complex lattice of sensitivity categorizations suitable for the wide variety of the aforementioned information types. This paper applies and extends our prior work on multi-level security for healthcare using LBAC by exploring alternative approaches to integrate this approach into the *Fast Healthcare Interoperability Resources (FHIR)* standard at the specification level of the standard.

1 INTRODUCTION

One of the major challenges in the healthcare industry is to ensure that a patient's healthcare information is securely accessible to a wide range of stakeholders (e.g., physicians, clinicians, medical specialists, nurses, non-medical staff, home health care providers, pharmacists, patients, family members, etc.) to administer patient care in a variety of settings such as physician offices, hospitals, rehab facilities, emergency rooms, home based care settings, etc. The challenge from a data security perspective is to provide fine-grained access control to a patient's healthcare information that is able to precisely define which portions of the information should be available to which stakeholders at what time. Granular sharing of medical, health, and fitness data is becoming an increasingly important aspect of patient care, considering new government initiatives which aim to broaden the sharing of a person's health data beyond traditional boundaries.

In this situation, the information that needs to be controlled has many different levels of permissions due to its sensitivity and confidentiality: controlled by the various types and granularity of information to which patients want to control access (Caine & Hanania, 2013); providing fine grained access control to allow a patient to define: who may view/modify what (Sujansky, et al., 2010); and, providing a way for patients to provide their data to an emergency physician in time critical situations (Peleg, et al., 2008). Of particular relevance, the work of (Caine & Hanania, 2013) identified 11 medical information items that are partitioned into 5 protected items such as: contact information and demographics, information relevant to current conditions, medications (prescribed and over-the-counter), test results (blood pressure, blood tests, imaging tests, etc.), and past medical history; and, 6 sensitive items such as: history of substance abuse and treatment,

Demurjian, S., Agresta, T., Sanzi, E. and DeStefano, J.
Alternative Approaches for Supporting Lattice-based Access Control (LBAC) in the Fast Healthcare Interoperability Resources (FHIR) Standard.
DOI: 10.5220/0010150800930104
In *Proceedings of the 16th International Conference on Web Information Systems and Technologies (WEBIST 2020)*, pages 93-104
ISBN: 978-989-758-478-7

mental health information, sexual health information, records relating to domestic violence, reproductive health records, and genetic information.

Given the variety and scope of a patient's healthcare information, one security approach that may address the granularity issue for fine-grained control is multi-level security (MLS), which has its origins in the lattice-based access control model (LBAC) (Denning, 1976) and the mandatory access control model (MAC) (Bell & LaPadula, 1976). Both models rely on sensitivity levels (e.g., unclassified U, confidential C, secret S, top secret TS, etc.) that are assigned to objects (termed classifications) and users (termed clearances). Access to objects depends on a comparison of a user's clearance against an object's classification based on the type of operation (read, write, etc.). MAC utilizes a strict linear order while LBAC utilizes a lattice.

In fact, our recent work (Demurjian et al., 2017) explored the use of multi-level security for healthcare using an LBAC approach to define five sensitivity levels from least secure to most secure, further subdivided by different categories to replace the four traditional ones (TS, S, C, U):

- Level 0: *Basic Demographic Data* such as city, state, general health condition, fitness data.
- Level 1: *Medical History Data* such as patient name, address, day/month of birth, weight, height, next of kin, history, immunization, and a separate mental health history.
- Level 2: *Summary Clinical Data* such as R_x and OTCs, allergies, diagnoses, treatment plan, and, for mental health, separate diagnoses, and treatment plan.
- Level 3: *Detailed Clinical Data* such as imaging studies, laboratory tests, mental health encounters (excluding psychotherapy notes), and clinical data (e.g., heart rate, blood oxygen level, blood pressure, etc.).
- Level 4: *Sensitive Clinical Data* used by medical specialists on genetics, substance abuse, mental health psychotherapy notes, reproductive health, and domestic violence.

Different categories of information at the same sensitivity level can be authorized to users based on required Level/Category combinations.

The main focus of this paper is to apply our prior work (Demurjian et al., 2017) on multi-level security for healthcare that we developed and explore alternative approaches to integrate this approach into *Fast Healthcare Interoperability Resources (FHIR)* (HL7 International, 2020). FHIR provides structures for sharing EHR data between healthcare providers. Data is accessed through *resources* utilizing a location URL as part of a REST API in conjunction with a logical ID. This allows data that resources describe to sync between separate FHIR systems. We focus on the FHIR *base* resources (e.g., patients, practitioners, and family relationships; organizations, services, appointments, and encounters) and *clinical resources* which are for a patient's health history. Note that these alternative approaches are described at the specification level of the FHIR resources and not from an implementation perspective.

The main objective of this paper is to explore the utilization of our prior work (Demurjian et al., 2017) on multi-level security for healthcare to incorporate the lattice and our security approach into the resources of FHIR. Section 2 provides background on healthcare, multi-level security, and FHIR. Section 3 reviews our prior work in LBAC in the health care domain. Section 4 applies our prior LBAC work to FHIR by discussing the way that Level/Categories can be assigned at the schema level to a resource and its components. This includes exploring two approaches that utilize various FHIR capabilities and features and reviewing our LBAC implementation strategy. Finally, Section 5 concludes our paper.

2 BACKGROUND

This section provides background material on concepts used in the rest of this paper. Section 2.1 reviews the different kinds of healthcare information and systems. Section 2.2 briefly summarizes the history of MLS as realized by the MAC and LBAC access control models. Section 2.3 briefly reviews the FHIR specification. Section 2.4 briefly discusses multi-level security in healthcare.

2.1 Healthcare Information & Systems

Caine and Hanania (Caine & Hanania, 2013) organized the recipients and information in the context of patients managing and sharing their medical data into 11 data items: contact information and demographics, information relevant to current conditions, medications, test results, past medical history, history of substance abuse and treatment, mental health information, sexual health information, records relating to domestic violence, reproductive health records, and genetic information. These data categories have parallels to what we have loosely characterized as medical/health/fitness data in this paper. However, there are three categories of data that we believe are missing from this list: fitness data

collected by patients utilizing fitness devices and mobile apps; medical data collected by patients using their own medical devices and mobile apps; and, medical data collected by patients at the direction of their physician (e.g., Holter Cardiac Monitor) that may record data or feed data to the physician via a phone link, the web, or a mobile app. Thus, we propose adding Patient-Supplied Fitness, Patient-Supplied Medical, and Patient/Physician Directed Medical data to this list.

This challenge of patients sharing information is further complicated by the fact that a patient's healthcare information is stored in multiple locations in a variety of health information technology (HIT) that includes: electronic health records (EHRs), practice management systems (PMS), e-prescribing systems, personal health records (PHRs), etc. The majority of these systems must adhere to various laws such as the Health Insurance Portability and Accountability Act (HIPAA) (HIPAA, 2017) for the security, availability, transmission, and release of a patient's medical information. The sharing of information among multiple HITs is being facilitated in part by the Fast Healthcare Interoperability Resources (FHIR) (HL7 International, 2020) standard, a health information exchange (HIE) standard created by HL7 to promote secure sharing of healthcare data among multiple HIT systems.

2.2 Multi-level Security and LBAC

LBAC (Denning, 1976) and MAC (Bell & LaPadula, 1976) share the approach of security sensitivity levels that are assigned to subjects (clearance) and objects (classification) with the permissions for the subject to read and/or write an object based on the relationship between clearance and classifications. MAC typically is modelled using four sensitivity levels which are hierarchically ordered from most to least secure: top secret (TS) < secret (S) < classified (C) < unclassified (U). LBAC generalizes this approach by ordering the sensitivity levels in a lattice that determines the relative ranking of each sensitivity level vs. the others. Security policies in LBAC and MAC are defined by a security administrator to control information flow in computer systems where users are prohibited from changing their security attributes. In LBAC and MAC, access to objects (e.g., segments of an XML document, tables in a database, etc.) by subjects (e.g., users, processes in a system, etc.) is granted based on the security definitions on the targeted object (exhibited via tags) and the credentials granted to the user.

From a definition and management perspective, a security administrator would set the clearance level of users following the predefined sensitivity levels (e.g., TS, S, C, and U) to establish the levels for both subjects and objects. These levels are then augmented on a user-by-user basis by assigning the ability to read and/or to write an object. Once this has all been established for an application, definition of permissions and levels (e.g., the elements of a patient's health record) can be used to maintain confidentiality by preventing an unauthorized provider to access sensitive information (e.g., not all medical providers are able to access mental health history) and to prohibit a patient from changing their own record. The use of multi-level security in the traditional military context is directly analogous to its application in healthcare. The major difference is that the "sensitivity" of information in the traditional military context relates to the risk to national security of improper disclosure, while the "sensitivity" level of patients' medical information relates to the privacy risk associated with release to unauthorized users.

In LBAC, the work of Denning (Denning, 1976) on lattice-based access control defines a set of security classes SC (analogous to security classifications) that is then organized into a universally bounded lattice that defines a partial order across the set SC. Such a representation expands the traditional DoD version of MLS so that a set of security classification levels can be defined in a complex lattice to represent a richer set of relationships among various types of information from a sensitivity perspective that are more conducive to the complex and rich sensitivity of healthcare data. The work of Landwehr (Landwehr, 1981) is a comprehensive review of the formal models for computer security circa 1981 and discusses both MLS and lattice-based access control.

The lattice approach expands MLS by defining a finite set of elements (security levels) that are augmented with a partial ordering in order to define, for each pair of elements, a least upper bound and a greatest lower bound. The "compartment sets" can be partially ordered with one another via subset relationships, so that given two sets, one compartment set is greater than or equal to another compartment set. As a result, security classifications in MLS can include not only a sensitivity level (top secret, secret, confidential, etc.) but also a compartment set.

2.3 FHIR

FHIR enables the retrieval of healthcare data by providing a common API to locate and exchange

healthcare records. FHIR's data exchange structure is built on the concept of resources, which provide a meaningful set of healthcare related data concepts. FHIR provides over 145 different resources for: patients, observations, medications, patient consent, etc. Requests for a specific resource are available through a REST API that supports instance level interactions such as: read, vread (version read), update, patch (update a portion of a resource), delete, and history interactions. FHIR resources are organized in categories: *foundation resources*, *base resources*, *clinical resources*, *financial resources*, and *specialized resources*. We highlight only a subset relevant for the paper. The *base* resources describe: patients, practitioners, and family relationships; organizations, services, appointments, and encounters. The *clinical resources* are for a patient's health history, including: diagnostic data, medications, care provision, and request/response communication. *HAPI FHIR* (HAPI FHIR, 2020) is a Java implementation of the FHIR specification. HAPI FHIR provides resource models for all resources defined in the current FHIR r4 specification. Interactions with FHIR resources are defined by the FHIR standard's REST API.

2.4 MLS in Healthcare

Despite the long history of MLS (since 1976) and its wide usage in governmental and commercial settings, there has been limited attention paid to the usage of MLS for health care. A recent review of access control models deployed by EHRs (Fernández-Alemán, Señor, Lozoya, & Toval, 2013) found that out of 35 articles, 27 specifically utilized RBAC. Our own attempts to identify MAC examples in health care found only three references. (Alhaqbani & Fidge, 2008; Gajanayake, Iannella, & Sahama, 2014; Hafner, Memon, & Alam, 2007). In terms of the use of MLS and MAC for security and privacy in health care, a post in the Healthcare Exchange Standards Blog (Moehrke, 2010) discussed the usage of the traditional military classification scheme in a health care setting. This work points to the definition of confidentiality labels in HL7 standards that are part of the vocabulary for FHIR HL7 (FHIR Confidentiality, 2020). Specifically, in the HL7 FHIR standard, the confidentiality labels are: U – unrestricted, L – low, M – moderate, N – normal, R – restricted, and V – very restricted. Note that the usage of confidentiality levels in this standard denotes the type of data to protect and the conditions under which to protect that data; they are not the same as sensitivity levels in MAC/MLS.

3 LBAC CLASSIFICATION FOR HEALTHCARE

This section presents our work (Demurjian et al., 2017) on an appropriate set of sensitivity labels for the healthcare domain that can be utilized for both classifications and clearances. As part of the process, we demonstrate that the rich semantics of health/medical/fitness data along with the varied requirements of stakeholders, necessitates that we move beyond a traditional linear-based MLS scheme to one that is lattice-based. In the rest of this section, a three-part approach is presented. First, we explain and review lattice-based access control in detail through a discussion of three key efforts (Denning, 1976; Landwehr, 1981; Sandhu, 1993). In the process, we transition to an MLS schema that has sensitivity levels, within each of which there may be multiple different categories of data that are related to one another in a lattice-based context. Second, we propose a set of security levels and categories for healthcare data to achieve a fine-grained security of medical/health/fitness data per the items from Table 1 of Caine and Hanania (Caine & Hanania, 2013) reviewed in Section 1. Once defined, these levels are then organized into a lattice whose structure is impacted by the way that medical stakeholders utilize different categories of data within each level. The sensitivity level lattice that is proposed is one example of the way healthcare data could be classified, but is not the only possible way to characterize such data. The third part of this section illustrates several alternative characterizations, which we term *Sensitivity Profiles*. Each Sensitivity Profile includes sensitivity labels that categorize medical/health/fitness data in different ways that are consistent with how information is utilized by different stakeholders/HIT systems in different contexts, easily understood by stakeholders, and clearly connote the confidentiality of the different types of healthcare information. Our premise is that it is highly unlikely that one single, universal set of sensitivity labels could be defined that would be suitable for all of the possible use cases in healthcare; as a result, we present alternate Sensitivity Profiles and discuss the situation under which each would be relevant in terms of the involved stakeholders and/or HIT systems/health information exchange (HIE).

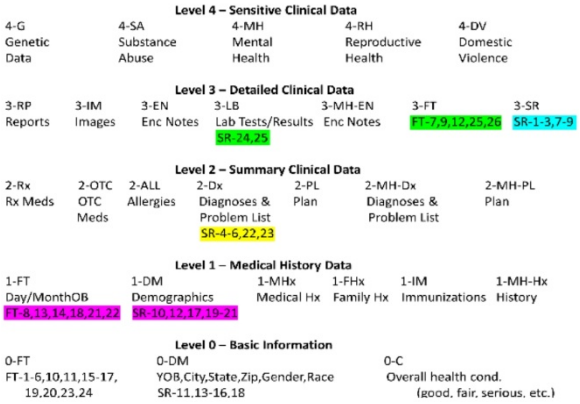

Figure 1: Sample Healthcare Sensitivity Levels.

The objective of this section is to propose and discuss a set of sensitivity levels for labelling fine-grained security of medical/health/fitness data per the items in Table 1 of Caine and Hanania (Caine & Hanania, 2013). The lattice to be presented in this section is intended for use by healthcare organizations (e.g., hospitals, clinics, medical specialist offices, etc.) to securely share healthcare data via HIE with an agreed upon set of security levels that are represented by a lattice. To begin, Figure 1 defines a set of five different sensitivity Levels (0 to 4) for healthcare, and within each level there are different categories of data that will be given to different users based on their need. The five levels replace the four traditional ones (TS, S, C, U) and are defined as:

Level 0: *Basic Information* contains data that is freely available to anyone: basic demographics such as city and state of residence and surveillance data from (11,13-16,18) (0-DM); general health condition (0-C); and information related to tracking fitness data (1-6,10,11,15-17,19,20,23,24) (0-FT) such as date, time, type, and duration of activity, etc.

Level 1: *Medical History Data* contains data that has some restrictions: detailed demographic data such as the patient name, address, day/month of birth, weight, height, next of kin, medical record ID of the patient, surveillance data (10,12,17,19-21) (1-DM); more sensitive patient-collected fitness data (8,13,14,18,21,22); history of the patient and his/her family, immunizations (1-MHx, 1-FHx, 1-IM respectively); and mental health history (1-MH-Hx).

Level 2: *Summary Clinical Data* including prescription (2-Rx) and over-the-counter medications (2-OTC), allergies (2-ALL), medical diagnoses and problem list that includes the provider name and ID and surveillance data (4-6,22,23) (2-Dx), plan for treatment or other related instructions (2-PL), and, for mental health, separate diagnoses (2-MH-Dx), and treatment plan (2-MH-PL).

Level 3: *Detailed Clinical Data* contains reports from imaging studies (CT Scans, MRIs, X-Rays, etc.) (3-RP), the images from the studies (3-IM), detailed information on each medical visit (encounter notes, 3-EN), laboratory tests ordered, dates, and results including surveillance data from (24,25) (3-LB), information about mental health encounters (excluding psychotherapy notes) (3-MH-EN), surveillance data (1-3,7-9) (3-SR), and clinical data (e.g., heart rate, blood oxygen level, blood pressure, etc.) from fitness devices (7,9,12,25,26) (3-FT).

Level 4: *Detailed Clinical Data* contains reports from imaging studies (CT Scans, MRIs, X-Rays, etc.) (3-RP), the images from the studies (3-IM), detailed information on each medical visit (encounter notes, 3-EN), laboratory tests ordered, dates, and results including surveillance data from (24,25) (3-LB), information about mental health encounters (excluding psychotherapy notes) (3-MH-EN), surveillance data (1-3,7-9) (3-SR), and clinical data (e.g., heart rate, blood oxygen level, blood pressure, etc.) from fitness device (7,9,12,25,26) (3-FT).

Level 5: *Sensitive Clinical Data* contains sensitive information used by specialists including data on genetics (4-G), substance abuse (4-SA), mental health psychotherapy notes (4-MH), reproductive health (4-RH), and domestic violence (4-DV).

Level 0 is the least secure, while Level 4 is the most secure. Each of the levels have different categories of information, that while at the same sensitivity level as one another, have the ability to be authorized to different users based on the combination of Level/Category. This Level/Category combination corresponds to the security level/compartment as defined in Landwehr (Landwehr, 1981). This was also shown in the example of Figure 5, where S-LW represents the combination of the S-L Level/Category and the S-W category, combining the two categories of data (L and W) within one level (S). In Figure 1, examples of Level 2 categories are: prescription (2-Rx) and over-

the-counter medications (2-OTC), allergies (2-ALL) and diagnoses/problems (2-Dx). In general terms, Level 0 is public data available to anyone without control, Level 1 is for use by administrative staff, Level 2 is for use by clinical staff (RNs, PAs, etc.), Level 3 is for use by medical providers, and Level 4 for is for use by specific medical specialists. A patient would have access to all of the levels.

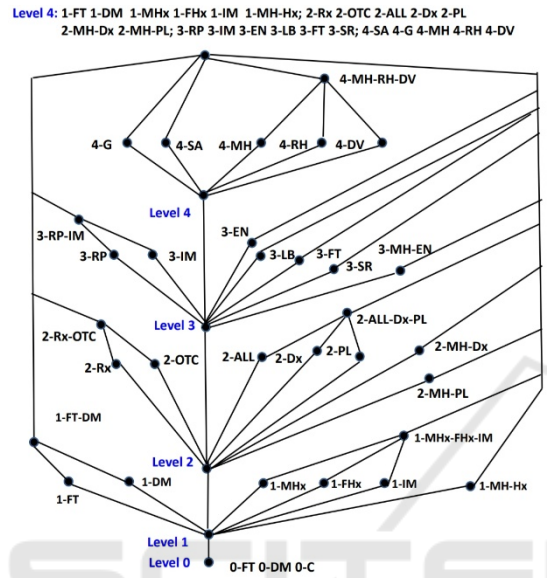

Figure 2: Corresponding Lattice based on Figure 1.

Figure 2 shows one possible lattice utilizing the sensitivity levels and categories as given in Figure 1, which is based on the work of Smith (Smith, 1990) which defines a product lattice that combines ordered security levels (i.e., TS, S, C, U) with eight different categories of data (A, K, L, Q, W, X, Y, Z); note that a category in this case is akin to a compartment as described by Landwehr (Landwehr, 1981) and this similar to our approach of sensitivity levels with categories. In Figure 2, there are different combinations of categories at each level that represent the likely usage of that medical data by a particular stakeholder. Starting from the bottom up in the figure, notice that on the lower right-hand side for Level 1, medical and family histories and immunizations are grouped as: 1-MHx-FHx-IM. On the lower left hand side, patient supplied fitness data and demographic categories are grouped as: 1-FT-DM. For Level 2, prescription and over-the-counter medications are put together into the group 2-Rx-OTC with medical plans, diagnoses, and allergies in a separate group 2-ALL-Dx-PL. Some medical providers might have access to medications (2-Rx-OTC) while others might need access to both and would be assigned 2-Rx-OTC and 2-ALL-Dx-PL. For Level 3, there is a linking of imaging and the associated reports into the group 3-RP-IM while information on encounter notes and laboratory tests/results can be separately assigned to a medical provider.

Access to summary mental health encounter information (3-MH-EN) can also be separately assigned. Finally, at Level 4, categories for mental health psychotherapy notes, reproductive health, and domestic violence are grouped into 4-MH-RH-DV since a medical provider treating one of those categories likely needs to know about the information in the other two, but may not require access to genetic or substance abuse data. Genetic (4-G) and substance abuse (4-SA) categories can be separately assigned. Note that in some cases, there may be a medical provider that needs all five of the categories in Level 4. The top level collects all categories into one logical unit. Two or more healthcare organizations (e.g., hospital A, hospital B, clinic C) that wish to securely share information on patients could agree to use the same set of sensitivity levels/categories (Figure 1) and corresponding lattice (Figure 2). This is analogous to DoD and federal organizations that agree to (TS, S, C, U) in that setting. One final note is that a recent article (Gajanayake, Iannella, & Sahama, 2014) on privacy for Electronic Health Records has utilized the mandatory access control approach to define an object sensitivity tree with allowed and prohibited sensitivity labels where a user would receive an aggregation of multiple allowed and prohibited levels. While their approach is related to our work it differs in that they are limited to a linear ordering of MAC (not the lattice of LBAC) and didn't demonstrate as comprehensive a treatment of medical data as given in Figure 1.

Given the lattice as presented in Figures 1 and 2, the various Level/Category combinations can be assigned to different users/stakeholders based on individual needs. For example, all information in level 0 is essentially public and freely available. Administrative users such as office staff would have access to level 0 as well as all of level 1 that includes 1-FT, 1-DM, 1-MHx, 1-FHx, 1-IM. Stakeholders that are clinical staff (RN, LPN, etc.) for a given patient would have access to levels 0 and 1, as well as 2-Rx, 2-OTC, 2-ALL, 2-Dx, and 2-PL. Stakeholders that are part of the medical provider team (MDs) for a given patient, would have access to levels 0, 1, and 2, as well as 3-RP, 3-IM, 3-EN, 3-LB, 3-FT, and 3-SR. Lastly, a specialist medical provider would have one or more of: 4-SA, 4-G, 4-MH, 4-RH, and 4-DV. In the last category, it may be necessary to protect specific information by specialist, for example: Protect mental health information: Eliminate 4-MH;

Protect HIV information: Eliminate 4-RH; Protect Genetic information: Eliminate 4-G; and, Protect Substance Abuse information: Eliminate 4-SA.

In addition, the different mental health information categories that exist from Level 1 through Level 4 provide the ability to tailor access to sensitive mental health information for different stakeholders. A psychiatrist would have access to: 1-MH-Hx, 2-MH-Dx, 2-MH-PL, 3-MH-EN, and 4-MH. This subdivision allows some of the mental health information on lower levels to be available to stakeholders that need access to lower level mental health such as 1-MH-Hx and 2-MH-Dx but would not be allowed access to other levels. The scenarios to be presented in Section 5 include more detailed examples of stakeholders and their necessary permissions in regards to Level/Category combinations of the Figure 2 lattice.

4 ALTERNATIVE APPROACHES FOR LBAC AND FHIR

This section explores alternative approaches for incorporating LBAC with Level/Category in Section 3 into FHIR. These approaches are at the level of the specification with specific examples of FHIR resources. Section 4.1 discusses the way that the Level/Category can be assigned at the schema level to a resource's attributes, resources, and references to other resources. The next two sections explore two different approaches for including LBAC with Level/Category, namely, FHIR Security Levels in Section 4.2 and FHIR Extensions in Section 4.3. Section 4.4 explores the various implementation strategies that can be utilized in order to realize LBAC within the FHIR framework. Note that all of the information in Figures 3, 4 and 7, as well as other FHIR examples are from publicly available freely citable web pages of (FHIR resources, 2020).

4.1 LBAC & Resources Concepts

In this section, we explore the way that the different FHIR resources can be labelled and classified using the sensitivity levels in the categories in Figure 1. To assist us in the process, Figure 3 contains the 145 currently defined FHIR resources. There are a number of resources that are particularly relevant to demonstrate five different sensitivity levels in Figure 1 and the way to position each of those resources in one primary sensitivity level. FHIR Resources of interest to us are:

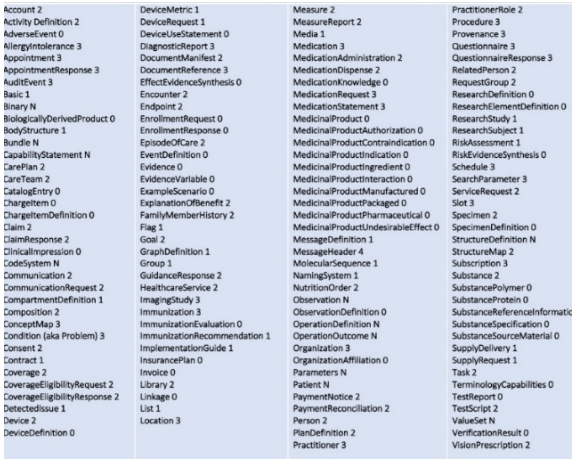

Account 2	DeviceMetric 1	Measure 2	PractitionerRole 2
Activity Definition 2	DeviceRequest 1	MeasureReport 2	Procedure 3
AdverseEvent 0	DeviceUseStatement 0	Media 1	Provenance 3
AllergyIntolerance 3	DiagnosticReport 3	Medication 3	Questionnaire 3
Appointment 3	DocumentManifest 2	MedicationAdministration 2	QuestionnaireResponse 3
AppointmentResponse 3	DocumentReference 3	MedicationDispense 2	RelatedPerson 2
AuditEvent 3	EffectEvidenceSynthesis 0	MedicationKnowledge 0	RequestGroup 2
Basic 1	Encounter 2	MedicationRequest 3	ResearchDefinition 0
Binary N	Endpoint 2	MedicationStatement 3	ResearchElementDefinition 0
BiologicallyDerivedProduct 0	EnrollmentRequest 0	MedicinalProduct 0	ResearchStudy 1
BodyStructure 1	EnrollmentResponse 0	MedicinalProductAuthorization 0	ResearchSubject 1
Bundle N	EpisodeOfCare 2	MedicinalProductContraindication 0	RiskAssessment 1
CapabilityStatement N	EventDefinition 0	MedicinalProductIndication 0	RiskEvidenceSynthesis 0
CarePlan 2	Evidence 0	MedicinalProductIngredient 0	Schedule 3
CareTeam 2	EvidenceVariable 0	MedicinalProductInteraction 0	SearchParameter 3
CatalogEntry 0	ExampleScenario 0	MedicinalProductManufactured 0	ServiceRequest 2
ChargeItem 0	ExplanationOfBenefit 2	MedicinalProductPackaged 0	Slot 3
ChargeItemDefinition 0	FamilyMemberHistory 2	MedicinalProductPharmaceutical 0	Specimen 2
Claim 2	Flag 1	MedicinalProductUndesirableEffect 0	SpecimenDefinition 0
ClaimResponse 2	Goal 2	MessageDefinition 1	StructureDefinition N
ClinicalImpression 0	GraphDefinition 1	MessageHeader 4	StructureMap 2
CodeSystem N	Group 1	MolecularSequence 1	Subscription 3
Communication 2	GuidanceResponse 2	NamingSystem 1	Substance 2
CommunicationRequest 2	HealthcareService 2	NutritionOrder 2	SubstancePolymer 0
CompartmentDefinition 1	ImagingStudy 3	Observation N	SubstanceProtein 0
Composition 2	Immunization 3	ObservationDefinition 0	SubstanceReferenceInformatic
ConceptMap 3	ImmunizationEvaluation 0	OperationDefinition N	SubstanceSpecification 0
Condition (aka Problem) 3	ImmunizationRecommendation 1	OperationOutcome N	SubstanceSourceMaterial 0
Consent 2	ImplementationGuide 1	Organization 3	SupplyDelivery 1
Contract 1	InsurancePlan 0	OrganizationAffiliation 0	SupplyRequest 1
Coverage 2	Invoice 0	Parameters N	Task 2
CoverageEligibilityRequest 2	Library 2	Patient N	TerminologyCapabilities 0
CoverageEligibilityResponse 2	Linkage 0	PaymentNotice 2	TestReport 0
DetectedIssue 1	List 1	PaymentReconciliation 2	TestScript 2
Device 2	Location 3	Person 2	ValueSet N
DeviceDefinition 0		PlanDefinition 2	VerificationResult 0
		Practitioner 3	VisionPrescription 2

Figure 3: Alphabetical List of Resources.

- Related to individuals: *Person* who is patient or medical stakeholder; *Patient* who receives medical services; Practitioner who is a physician, visiting nurse, home health aide, etc.; and, *Organization* that administers or provides medical care.
- Related to a patient's health record: *Medication* tracks medications that a patient is taking or has taken; *AllergyIntolerance* keeps tracks of any allergies; *FamilyMemberHistory* for personal and family medical history; and, Immunizations which tracks vaccines.
- Summary data: *MedicationRequest* to record a prescription for a patient; *Condition* for the different diagnosis for a patient; *Observation* that keeps track of actual results of vital signs, different types of tests, social history, etc.; and, *CarePlan* to track the different plans among medical stakeholders to manage care.
- Detailed clinical data on a patient: *ImagingStudy* for the actual test results of an imaging study; and, *DiagnosticReport* that contains information on a patient's laboratory or other medical tests.

The remainder of this section explores the relevant sensitivity level for a subset of the FHIR resources shown in the previous bulleted list, and in the way that the sensitivity of the resource itself, in terms of the actual data stored for the resource, is interpreted.

To begin, we start with the first basic resource that underlies all healthcare applications that are developed using FHIR, the Person resource as given in Figure 4 in the XML format. Note that in addition to XML, a resource can also have a json and Turtle format. The information in the Person resource has basic name and demographic information, including attributes for: identifier, name, gender, telcom, birthdate, and active. The address attribute represents

the inclusion of another resource, Address. The managingOrganization attribute is a reference to 0 or 1 instances of an Organization resource and the attribute target is a reference to one or more Practitioner, RelatedPerson, or Person resource instances who are involved with the Person. The type of information that is included in the Person resource is represented as sensitivity levels 1-DM and 0-DM in Figure 2, representing the demographics from levels 1 and 0 respectively. This means that a person resource is primarily demographics information.

```
<?xml version="1.0" encoding="UTF-8"?>
<Person xmlns="http://hl7.org/fhir">doco
    <!-- from Resource: id, meta, implicitRules, and language -->
    <!-- from DomainResource: text, contained, extension, and modifierExtension -->
  <identifier></identifier>
  <name></name>
  <telecom></telecom>
  <gender value="[code]"></gender>
    <!-- 0..1 male | female | other | unknown -->
  <birthDate value="[date]"></birthDate>
    <!-- 0..1 The date on which the person was born -->
  <address></address>
  <photo></photo>
  <managingOrganization></managingOrganization>
  <active value="[boolean]"></active>
    <!-- 0..1 This person's record is in active use -->
  <link>
      <!-- 0..* Link to a resource that concerns the same actual person -->
    <target></target>
    <assurance value="[code]"></assurance>
      <!-- 0..1 level1 | level2 | level3 | level4 -->
  </link>
</Person>
```

Figure 4: Person Resources in XML Format.

From a sensitivity Level and Category perspective, we assign a level of 1-DM for the entire Person resource, shown in Figure 5. This is a resource level assignment of sensitivity coupled with the demographics category. The 1-DM is utilized since there is some information in the Person resource that cannot be released to the general public. Within the resource itself, each of the individual attributes can have the same or lower sensitivity levels. Within the Person resource, the embedded address resource's attributes for city, district (county), state, country, and postal code, would be tagged at sensitivity level 0-DM, since this is public information that could be utilized for statistical analysis of state and country wide healthcare data. The sensitivity of the managingOrganization attribute will be governed by the sensitivity of the Organization resource, also 1-DM, to allow information about an organization to be protected. For the other Person attributes: date of birth and gender are 0-DM, with all remaining attributes defaulting to level 1-DM, which includes references via the target attribute to the involved Practitioner, RelatedPerson, or Person resource. Figure 6 summarizes the sensitivity levels.

Resource	SL				
Person	1-DM				
<identifier>	1-DM				
<name>	1-DM				
<telecom>	0-DM				
<gender>	0-DM				
<birthDate>	0-DM	Resource	SL		SL
<address>	1-DM	Address	1-DM	city	0-DM
				district	0-DM
				state	0-DM
				country	0-DM
				postal code	0-DM
<photo>	1-DM	Resource			
<managing Organization>		Organization	1-DM		
<active>	1-DM				
<target>	1-DM	Practitioner	1-DM		
		Related Person	1-DM		
		Person	1-DM		
<assurance>	1-DM				

Figure 5: Sensitivity Levels (SL) for Person.

Resource	SL	Resource	SL
Person	1-DM	Medication	2-Rx
Patient	1-DM	AllergyIntolerance	2-ALL
Organization	1-DM	MedicationRequest	2-Rx
Practitioner	1-DM	Condition	2-Dx
RelatedPerson	1-DM	CarePlan	2-PL
Address	1-DM	Observation	3-EN
Immunizations	1-Im	DiagnosticReport	3-RP
FamilyMemberHistory	1-MHx	ImagingStudy	3-IM

Figure 6: Sensitivity Levels (SL) of Select FHIR Resources.

4.2 Approach A: FHIR Security Labels

Approach A leverages the capabilities of the (FHIR Security Labels, 2020) in an attempt to represent the levels and categories of our LBAC approach using the coding capabilities available in FHIR. One of the capabilities of the FHIR security levels is the ability to define confidentiality levels U – unrestricted, L – low, M – moderate, N – normal, R – restricted, and V – very restricted as discussed in Section 2.4 (FHIR Confidentiality, 2020). Consider the example from (FHIR Security Labels, 2020) on a patient resource with a tag added with a confidentiality level of R. While this could support the Level of our LBAC approach, and we could map our five Levels it would not be sufficient to support the Categories.

```
<Patient xmlns="http://hl7.org/fhir">
  <meta>
    <security>
      <system
      value="http://terminology.hl7.org/Code
      System/v3-Confidentiality"/>
```

```
      <code value="R"/>
      <display value="Restricted"/>
    </security>
  </meta>
...  [snip] ...
</Patient>
```

One of the other capabilities that has the potential to be leveraged as part of the standard is the HL7 Healthcare Privacy and Security Classification System (HL7 HCS, 2020). The FHIR version of has five different security labels: confidentiality classification, sensitivity category, compartment category, integrity category, and handling caveat; we will focus on the first two that have been excerpted from (FHIR Security Labels, 2020):

- "Confidentiality Classification: 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. Example Uses: Unrestricted, Normal, Very restricted."
- "Sensitivity Category: Security label metadata that "segments" an IT resource by categorizing the value, importance, and vulnerability of an IT resource perceived as undesirable to share. Example Uses: STDs, Psychiatric care, Celebrity status."

For confidential classification, the security label has a cardinality of 0 or 1 which means it may or may not be present; so a resource may have a classification or may not. These are the confidentiality levels U, N, R, etc. For the sensitivity category, the label has a cardinality of 0 or multiple, may not be present, or one or more. This gives the impression that the sensitivity category potentially could be applied to not only the resource but to different parts of the resource, namely the attributes. This makes sense as shown in our examples in Section 4.1, a given resource might have different Level/Category combinations at the attribute level.

From a sensitivity Level and Category perspective, we assign a level of 1-DM for the entire Person resource, shown in Figure 5. The information sensitivity labels have been excerpted from (FHIR Information Sensitivity, 2020) and are shown in Table 1. Notice that there are 4 sensitivity levels that range from least sensitive 1 to most sensitive 4. Notice that many of these information sensitivities have parallels to our categories in Figure 1. So, it would be possible to map the codes in Table 1 to the categories in Figure 1. The issue is whether this mapping will allow us to establish different codes for each resource or even the attributes of each resource. If it only works on the resource level this would not be sufficient to support LBAC. To illustrate the potential correlation between our LBAC and the information sensitivity levels, Table 2 contains a mapping between Figure 1 and a subset of Table 1. Note that there was no obvious way to map all of our Level/Categories in Figure 1 to all of the different information sensitivity levels in Table 1; it not feasible to realize FHIR Security Labels.

Table 1: Information Sensitivity Levels.

Level	Code	Display
2	ETH	substance abuse
2	GDIS	genetic disease
2	SDV	sexual assault, abuse, or domestic violence
3	BH	behavioral health
3	MH	mental health
3	PSY	psychiatry disorder
2	VIO	violence
2	DEMO	all demographic
2	DOB	date of birth
2	GENDER	gender and sexual orientation
2	MARST	marital status
2	RACE	race
2	REL	religion
2	LOCIS	location
1	DIA	diagnosis
1	DRGIS	drug

Table 2: Correlating LBAC with Information Sensitivity.

LBAC Level/ Category	Sensitivity Level	Code	
4-SA	2	ETH	substance abuse
4-G	2	GDIS	genetic disease
4-DV	2	SDV	sexual assault, abuse, or domestic violence
4-RH	2	SEX	sexuality and reproductive health
4-MH	3	MH	mental health
3-MH-EN	3	PSYTHPN	psychotherapy note
4-SA	3	SUD	substance use disorder
1-DM	2	DEMO	all demographic
1-DM	2	DOB	date of birth
1-DM	2	GENDER	gender and sexual orientation
2-Dx	1	DIA	diagnosis
2-Rx	1	DRGIS	drug

4.3 Approach B: FHIR Extensions

Approach B to integrating LBAC into FHIR utilizes the extensions capability (FHIR Extensions, 2020) which allows a resource to be modified with additional features that then must be published with a

formal definition call and the structure definition to be shared within the FHIR community. This allows the extensions to be used by app developers and their applications. In this section, we provide an extension for supporting LBAC as presented in Section 3 in Figure 7. The extension is meant to represent Level and categories as given in Figure 1 such as 1-DM, 2-Rx, 2-MH-Dx, etc. These strings of characters are decomposed to be stored in the extension of the resource. The top portion of Figure 7 is the structure definition of the LBAC extension. You can see there are two attributes in the extension: an integer level with the value of 0 to 4 which represents the levels as shown in Figure 1; and, a character string that can be multiple characters and corresponds to the category such as DM, Rx, etc.

```
<!-- LBAC FHIR Extension-->
<extension xmlns="http://hl7.org/fhir"
url="http://hl7.org/fhir/StructureDefinition/lbac" >
 <extension url="level">
  <valueInteger value=integer>
       <!--value ranges from 0 to 4 -->
 </extension>
 <extension url="category">
  <valueString=”[string]”/>
        <!--  1..* multiple characters  -->
 </extension>
</extension>

<!-- Person instance with extension-->and l
<?xml version="1.0" encoding="UTF-8"?>
<Person xmlns="http://hl7.org/fhir">
<extension xmlns="http://hl7.org/fhir"
url="http://hl7.org/fhir/StructureDefinition/lbac" >
 <extension url="level">
  <valueInteger value=1>
    <!—Level 1 portion of 1-DM -->
 </extension>
 <extension url="category">
  <valueString=”DM”/>
     <!— Category DM portion of 1-DM -->
 </extension>
</extension>
<id value="example"/>
  <text>
   <status value="generated"/>
   <div xmlns="http://www.w3.org/1999/xhtml">
    <table>
     <tbody>
      <tr>
       <td> Name</td>
       <td> Peter James <b> Chalmers</b>
<!—other data for person-->
</Person>
```

Figure 7: LBAC Extension and Person Instance.

The bottom portion of Figure 7 contains the extension as applied to the Person resource instance, which has the sensitivity level of 1-DM as was shown in Figure 5. The extension goes at the top of the resource instance, and a portion of the person instance has been shown with the rest omitted. This LBAC FHIR extension would have to be integrated into all of the resources that you wish to control using LBAC in your healthcare application. Of course, when this is included in all of the different instances for your application, the remaining issue is the required implementation infrastructure that is necessary to enforce LBAC for any given application, which is the subject of Section 4.4.

4.4 LBAC Implementation Strategy

In this section, we report on the implementation strategy for Approach B FHIR Extensions in Section 4.3 in two related areas:

1. An LBAC database component that is intended to store the complete lattice for a given application and specific user permissions to access the resources.
2. An LBAC access control server that will be the piece of software that interacts with HAPI FHIR (HAPI, 2020), an open-source Java-based library of the FHIR standard.

We review each of these components.

The LBAC database component will be the repository that stores:

- The complete set of levels in all categories as shown in Figure 1. This will require recoding all of the different Level/Category combinations such as 1-DM, 2-Rx, 3-MH-Dx, etc. into a computer compatible form. For instance, the five different levels from 0 to 4 can be given the unique identifiers L0 to L4. Also, all of the different strings that represent the categories can be put into a table that maps each string to unique category ID, e.g., table entry such as <C1, DM>, <C2, Rx>, and <C3, MH-Dx>, etc. Finally, there will be a separate mapping table with entries such as [S1, L1, C1], [S2, L2, C2], and [S3, L3, C3], where each of these combinations has been given a unique identifier for sensitivity S, of which we assume that there are *n* of them.
- The complete lattice as given in Figure 2 which has the relationships among all of the Level/Category combination sensitivity levels from Figure 1. The application-specific lattice in Figure 2 contains pairwise comparisons among the S1, S2, ..., Sn sensitivities.

- All of the complete permissions to the specific Level/Categories from the lattice of Figure 2 that have been authorized to the different users. As previously stated this would include: Office Staff 1-FT, 1-DM, 1-MHx, 1-FHx, 1-IM; RN, levels 0 and 1, as well as 2-Rx, 2-OTC, 2-ALL, 2-Dx, and 2-PL; MDs) levels 0, 1, and 2, as well as 3-RP, 3-IM, 3-EN, 3-LB, 3-FT, and 3-SR; and, psychiatrist 1-MH-Hx, 2-MH-Dx, 2-MH-PL, 3-MH-EN, and 4-MH.

The LBAC access control server leverages the capabilities of HAPI FHIR (HAPI, 2020) with a HAPI-FHIR server that can be used in front of an HIT. The HAPI-FHIR server architecture has two main functionalities. First, a HAPI RestfulServer is a Servlet that a developer utilizes to: create instances of user-defined resource provider and, specify the Servlet path. Second, a Resource Provider is a class that represents one FHIR resource (e.g., Patient) that has a number of empty annotated methods for CRUD verbs that a developer needs to implement. These empty annotated methods are utilized to parse HTTP requests and convert the transferred data to/from FHIR format/Back-end System format.

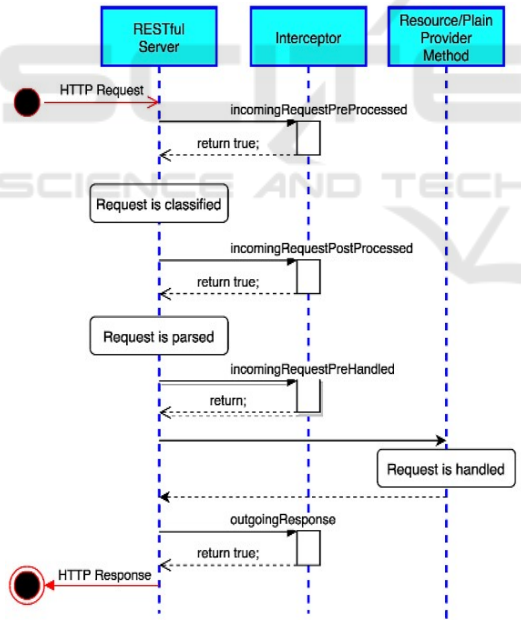

Figure 8: HAPI Interceptor Methods.

The HAPI-FHIR library provides a key capability for supporting LBAC, specifically, a general HAPI server interceptor in Figure 8 (HAPI Interceptor, 2016) which is a programmatic approach that allows a developer to examine each incoming HTTP request to add useful features to the HAPI RestfulServer such as authentication, authorization, auditing, logging, etc. The general HAPI interceptor, the InterceptorAdapter class, defines a number of methods that enable a developer to interact with the incoming HTTP requests at different points of the request lifetime. As Figure 8 shows, these methods are: incomingRequestPreProcessed that is invoked before performing any action to the request; incomingRequestPostProcessed that is invoked after determining the request type which classifies the request; incomingRequestPreHandled which is invoked before sending the request to the Resource provider; and, outgoingResponse which is invoked after the request is handled by the appropriate Resource provider. Each of these methods must returns either true, to continue processing the request, or false, to abort and reject the request. Moreover, a developer may extend the InterceptorAdapter class and implement the needed methods and register the extended class in the HAPI RestfulServer.

All of the capabilities of the interceptor can be utilized to do the required LBAC permission checks whenever a user is attempting to access a resource that has been extended using our Levels and Categories as discussed in Section 4.3. Specifically, whenever there is an attempt to access a resource by a health care application, the intercept capability performs a LBAC security check to determine if the user has the correct authorization to a subset of the lattice to access the requested instance of the resource. Recall the example from Section 3, an MD for a given patient would have access to levels 0, 1, and 2, as well as 3-RP, 3-IM, 3-EN, 3-LB, 3-FT, and 3-SR. That individual would have access to the resources through LBAC authorization to those Level/Categories. If the user attempted to access a resource at the domestic violence 4-DV, they would be denied. The interceptors can be utilized to ensure that only those resources that have been authorized by the lattice permissions can be returned to the user. Note that we have utilized interceptors in our prior research and have significant experience in their usage (Rivera Sánchez, et al. 2019).

5 CONCLUSIONS

The paper has applied our prior work on multi-level security using lattice-based access control (LBAC) for healthcare (Demurjian et al., 2017) to the Fast Healthcare Interoperability Resource (FHIR) standard. Specifically, this paper: explored the integration of our ideas into the resources of FHIR; presented and discussed two different approaches, FHIR Security Labels and FHIR Extensions, as

potential solutions for LBAC in FHIR; and, discussed an LBAC Implementation Strategy that could be actually utilized to realize a feasible solution. In support of this work: Section 2 reviewed background on healthcare, multi-level security, and FHIR; Section 3 reviewed our prior LBAC approach for healthcare; and, Section 4 discussed assigning Level/Categories to a resource, explored the two aforementioned approaches, and discussed our current implementation strategy. The alternative approaches presented in section 4 were described at the specification level of the FHIR resources and not from an implementation perspective.

REFERENCES

Alhaqbani B, Fidge C. Access Control Requirements for Processing Electronic Health Records, Business Process Management Workshops, AHM ter Hofstede, B Benatallah, and H-Y Paik, (eds.), Springer, LNCS, 4928, 2008;371-382.

Bell, D. E., & La Padula, L. J. (1976). Secure Computer Systems: Unified Exposition and Multics Interpretation. Bedford, Mass.: MITRE Corp.

Caine K, Hanania R. Patients want granular privacy control over health information in electronic medical records. J Am Med Inform Assoc 2013;20(1):7-15.

Denning DE. A Lattice Model of Secure Information Flow, Communications of the ACM 1976;19(5):236-243.

Department of Defense. Department of Defense Trusted Computer System Evaluation Criteria 15 August 1983;23. Available at: http://csrc.nist.gov/publications/history/dod85.pdf

Fernández-Alemán JL, Señor IC, Lozoya PÁO, and Tonal A. Security and privacy in electronic health records: A systematic literature review. J Biomed. Inform 2013;46(3):541–562.

Ferraiolo et al. (2001). Proposed NIST standard for role-based access control. ACM Transactions on Information and System Security (TISSEC), 4(3), 224–274.

FHIR Confidentiality (2020). HL7 v3 Value Set Confidentiality Classification. https://www.hl7.org/fhir/v3/ConfidentialityClassification/vs.html

FHIR Information Sensitivity (2020). HL7 v3 Value Set Information Sensitivity Policy. https://www.hl7.org/fhir/v3/InformationSensitivityPolicy/vs.html

FHIR Resources (2020). https://www.hl7.org/fhir/resourcelist.html

FHIR Security Labels (2020). Retrieved July 15, 2020. https://www.hl7.org/fhir/security-labels.html

Gajanayake R., Iannella R., and Sahama T. (2014). Privacy Oriented Access Control for Electronic Health Records, Special Issue on e-Health Informatics and Security, Electronic Journal for Health Informatics, 8(2),e15.

HAPI FHIR (2020). HAPI FHIR - The Open Source FHIR API for Java. https://hapifhir.io/

HAPI interceptor (2016). University Health Network. HAPI Server Interceptors. http://hapifhir.io/doc_rest_server_interceptor.html

HIPAA (2017). Health Information Portability and Accountability Act. http://www.hhs.gov/ocr/hipaa

HL7 HCS (2020). Healthcare Privacy and Security Classification System. Retrieved July 15, 2020. http://www.hl7.org/implement/standards/product_brief.cfm?product_id=345

HL7 International. (2020). Index - FHIR v4.0.1. http://hl7.org/fhir/

Landwehr CE, Formal Models for Computer Security, ACM Computing Surveys 1981;13(3):247-278.

Moehrke J. Data Classification - a key vector enabling rich Security and Privacy Controls, Healthcare Exchange Standards Blog, August 10, 2010. Available at: https://healthcaresecprivacy.blogspot.com/2010/08/data-classification-key-vector-through.html

Peleg, M, et al. Situation-Based Access Control: privacy management via modeling of patient data access scenarios. J Biomed Inform 2008;41(6):1028-40.

Rivera Sánchez, Y.K., Demurjian, S., and Baihan, M. (2019). A service-based RBAC & MAC approach incorporated into the FHIR *standard. Digital Communications and Networks*, 5(4):214-225, Elsevier.

Ryutov et al. (2005). Adaptive Trust Negotiation and Access Control. SACMAT '05 Proceedings of the tenth ACM symposium on Access control models and technologies (pp. 139-146).

Sandhu R. Lattice Based Access Control Model, IEEE Computer 1993;26(11):9-19.

Smith, G. (1990). The Modeling and Representations of Security Semantics for Database Applications. Doctoral Dissertation, George Mason UUniversit

Sujansky WV, et al. A method to implement fine-grained access control for personal health records through standard relational database queries. J Biomed Inform 2010;43(5 Suppl):S46-50.

The Office of the National Coordinator for Health Information Technology. (2018, September 19). Meaningful Consent Overview | HealthIT.gov. https://www.healthit.gov/topic/meaningful-consent-overview

The Office of the National Coordinator for Health Information Technology. (2019). Patient Consent for Electronic Health Information Exchange | HealthIT.gov. Retrieved January 24, 2020,

U.S. Department of Veterans Affairs - Office of Public and Intergovernmental Affairs. (2019). VA achieves critical milestone in its Electronic Health Record Modernization Program. Retrieved January 4, 2020, t

Yasnoff, W. A. (2015). A Feasible and Sustainable Approach to Health Information Infrastructure Via Mobile Devices. http://mediasite.uchc.edu/mediasite41/Play/b409b6fea70b4ec5b3fc34355340ac521