

V3_PSAF_R1_N3_2019SEP



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

Volume 3: Provenance Domain Analysis Model

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	Mohammad Jafari, Book Zurman Incorporated Contributor
Dave Silver, Electrosoft Inc. Contributor	Diana Proud-Madruga, Electrosoft Inc. Contributor
Sponsoring HL7 Security Work Group Cochairs	
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 Cochairs	
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

Special acknowledgement goes to the International Virtual Observatory Alliance (IVOA) and their IVOA Provenance Data Model Version 1.0, which is an IVOA Proposed recommendation 2018-1015. Much of the material herein is adapted from IVOA description of provenance as a tool of science, as such equally applicable to healthcare.

Table of Contents

1	PREFACE	10
2	INTRODUCTION	11
2.1	PROVENANCE OVERVIEW	11
2.2	PROVENANCE AND METADATA	12
2.3	PROVENANCE AND TRUST	12
2.4	HEALTHCARE LIFECYCLE EVENTS (LCEs)	13
3	SCOPE	14
3.1	ASSUMPTIONS	14
3.2	LIMITATIONS	14
3.3	PRECONDITIONS WITHIN SCOPE	14
3.4	OUT OF SCOPE	15
4	FEDERATED PROVENANCE MODEL	16
4.1	ENTERPRISE VIEW	17
4.2	FEDERATED PROVENANCE DATA FLOWS	19
4.2.1	<i>Direct</i>	19
4.2.2	<i>Redirect</i>	19
4.2.3	<i>Query</i>	19
4.2.4	<i>Brokered Exchange</i>	20
5	STORYBOARDS	21
5.1	STORYBOARD 1 - TRACKING PRODUCTION HISTORY	21
5.1.1	<i>Claims Workflow Provenance</i>	21
5.1.2	<i>Clinical Documentation Workflow</i>	23
5.2	STORYBOARD 2 – DIRECTED AND FEDERATED PROVENANCE CHAIN	23
5.3	STORYBOARD 3 - LOCATE ERROR SOURCES AND SHARING WITH PROTECTIONS	25
5.4	PROVENANCE REPORTING WORKFLOW	26
5.5	PRIVACY PREFERENCES AND PROVENANCE	28
5.6	STORYBOARD PERSONA	28
6	USE CASES	30
7	ACTIVITIES (NORMATIVE)	34
7.1	LIFE-CYCLE EVENTS	34
7.2	FUNCTIONAL FLOWS	34
7.2.1	<i>Send Provenance</i>	34
7.2.2	<i>Send Provenance Redirect (Optional)</i>	34
7.2.3	<i>Request Provenance</i>	35
7.2.4	<i>Request Provenance Store Analysis</i>	36
7.2.5	<i>Request Provenance Store Notifications</i>	37
8	CLASS MODEL (NORMATIVE)	38
8.1	ENTITY	ERROR! BOOKMARK NOT DEFINED.
8.2	COLLECTIONS	39
8.3	ACTIVITY	ERROR! BOOKMARK NOT DEFINED.

8.4	AGENT	ERROR! BOOKMARK NOT DEFINED.
8.5	USED	41
8.6	WASGENERATEDBY	42
8.7	WASASSOCIATEDWITH.....	42
8.8	WASATTRIBUTEDTO.....	43
8.9	PROVENANCE CHAINING	43
8.9.1	<i>Entity Chains (wasDerivedFrom)</i>	45
8.9.2	<i>Activity-Based Chaining (wasInformedBy)</i>	45
8.9.3	<i>Agent-Based Chaining (actedOnBehalfOf)</i>	46
8.9.4	<i>Activity-Entity-Based Chaining (used, wasGeneratedBy)</i>	46
8.10	HEALTHCARE LIFE-CYCLE EVENTS AS PROVENANCE	48
8.10.1	<i>Access Or View</i>	49
8.10.2	<i>Add Legal Hold</i>	49
8.10.3	<i>Amend (Update)</i>	49
8.10.4	<i>Archive</i>	50
8.10.5	<i>Attest</i>	50
8.10.6	<i>Encrypt</i>	51
8.10.7	<i>Decrypt</i>	51
8.10.8	<i>De-Identify (Anonymize)</i>	51
8.10.9	<i>Deprecate</i>	52
8.10.10	<i>Destroy or Delete</i>	52
8.10.11	<i>Disclose</i>	53
8.10.12	<i>Extract</i>	53
8.10.13	<i>Link</i>	54
8.10.14	<i>Merge</i>	54
8.10.15	<i>Originate</i>	55
8.10.16	<i>Pseudonymize</i>	55
8.10.17	<i>Re-Activate</i>	56
8.10.18	<i>Receive</i>	56
8.10.19	<i>Re-Identify</i>	57
8.10.20	<i>Remove Legal Hold</i>	57
8.10.21	<i>Report (Output)</i>	58
8.10.22	<i>Restore</i>	58
8.10.23	<i>Translate or Transform</i>	59
8.10.24	<i>Transmit</i>	59
8.10.25	<i>Unlink</i>	60
8.10.26	<i>Un-Merge</i>	60
8.10.27	<i>Verify</i>	61
9	REQUIREMENTS.....	62
	<i>Data analytics services enable enhanced understanding of the</i>	79

List of Figures

FIGURE 1: ELEMENTS OF TRUSTWORTHY INTEROPERABILITY	10
FIGURE 2: W3C PROVENANCE MODEL.....	11
FIGURE 3: FEDERATED PROVENANCE SERVICE	17
FIGURE 4: FEDERATED PROVENANCE OVERVIEW.....	17
FIGURE 5: FEDERATED PROVENANCE DATA FLOWS.....	19
FIGURE 6: OPTIONAL BROKERED EXCHANGE FLOWS IN THE FEDERATED PROVENANCE MODEL.....	20
FIGURE 7: STORYBOARD 1 - CLAIMS PRODUCTION TRACKING.....	23
FIGURE 8: STORYBOARD 2 – CLAIMS AND CLINICAL DOCUMENTATION RESPONSE TO DR. BOB WITH HOP-TO-HOP PROVENANCE	24
FIGURE 9: STORYBOARD 2 ATTRIBUTION.....	25
FIGURE 10: PROVENANCE REPORTING WORKFLOW	27
FIGURE 11: FEDERATED PROVENANCE DOMAIN	30
FIGURE 12: ACTIVITY DIAGRAM FOR SEND PROVENANCE	34
FIGURE 13: ACTIVITY DIAGRAM FOR SEND REDIRECT	35
FIGURE 14: ACTIVITY DIAGRAM FOR RECEIVE PROVENANCE.....	35
FIGURE 15: ACTIVITY DIAGRAM FOR REQUEST PROVENANCE FROM AGENT	36
FIGURE 16: ACTIVITY DIAGRAM FOR REQUEST PROVENANCE FROM PROVENANCE STORE	36
FIGURE 17: ACTIVITY DIAGRAM FOR REQUEST ANALYSIS.....	36
FIGURE 18: ACTIVITY DIAGRAM FOR REQUEST NOTIFICATION.....	37
FIGURE 19: PROVENANCE DATA MODEL.....	38
FIGURE 20: PROVENANCE CHAINING.....	44
FIGURE 21: PROVENANCE CHAINING BASED ON THE WASDERIVEDFROM RELATION.	45
FIGURE 22: PROVENANCE CHAIN BASED ON THE WASINFORMEDBY RELATION.	46
FIGURE 23: PROVENANCE CHAIN BASED ON THE ACTEDONBEHALFOF RELATION	46
FIGURE 24: PROVENANCE CHAIN BASED ON USED AND WASGENERATEDBY RELATIONS.....	47
FIGURE 25: RELATIONSHIPS OF LIFECYCLE EVENTS TO CREATE, READ, UPDATE, DELETE, AND EXECUTE [ISO 21089]	48
FIGURE 26: ACCESS OR VIEW PROVENANCE EVENT.....	49
FIGURE 27: ADD LEGAL HOLD PROVENANCE EVENT.....	49
FIGURE 28: AMEND/UPDATE PROVENANCE EVENT.....	50
FIGURE 29: ARCHIVE PROVENANCE EVENT	50
FIGURE 30: ATTEST PROVENANCE EVENT.....	50
FIGURE 31: ENCRYPT PROVENANCE EVENT	51
FIGURE 32: DECRYPT PROVENANCE EVENT	51

FIGURE 33: DE-IDENTIFY (ANONYMIZE) PROVENANCE EVENT	52
FIGURE 34: DEPRECATE PROVENANCE EVENT	52
FIGURE 35: DESTROY OR DELETE PROVENANCE EVENT	53
FIGURE 36: DISCLOSE PROVENANCE EVENT	53
FIGURE 37: EXTRACT PROVENANCE EVENT	54
FIGURE 38: LINK PROVENANCE EVENT	54
FIGURE 39: MERGE PROVENANCE EVENT	55
FIGURE 40: ORIGINATE PROVENANCE EVENT	55
FIGURE 41: PSEUDONYMIZE PROVENANCE EVENT	56
FIGURE 42: RE-ACTIVATE PROVENANCE EVENT.....	56
FIGURE 43: RECEIVE PROVENANCE EVENT	57
FIGURE 44: RE-IDENTIFY PROVENANCE EVENT	57
FIGURE 45: REMOVE LEGAL HOLD PROVENANCE EVENT	58
FIGURE 46: OUTPUT (REPORT) PROVENANCE EVENT.....	58
FIGURE 47: RESTORE PROVENANCE EVENT	59
FIGURE 48: TRANSLATE OR TRANSFORM PROVENANCE EVENT	59
FIGURE 49: TRANSMIT PROVENANCE EVENT	60
FIGURE 50: UNLINK PROVENANCE EVENT	60
FIGURE 51: UN-MERGE PROVENANCE EVENT.....	61
FIGURE 52: VERIFY PROVENANCE EVENT	61
FIGURE 53: ISO DATA LINEAGE INFORMATION MODEL.....	77
FIGURE 54: IBM WATSON RESEARCH MODEL	78
FIGURE 55: INTRA-DOMAIN PROVENANCE MODEL	79

List of Tables

TABLE 1: FEDERATED PROVENANCE USE-CASE DESCRIPTIONS	31
TABLE 2: ATTRIBUTES OF THE ENTITY CLASS	39
TABLE 3: ATTRIBUTES OF THE HADMEMBER RELATION.	40
TABLE 4: ATTRIBUTES OF THE ACTIVITY CLASS	40
TABLE 5: ATTRIBUTES OF AGENT CLASS	41
TABLE 6: AGENT TYPES	41
TABLE 7: ATTRIBUTES OF THE USED RELATION	41
TABLE 8: ATTRIBUTES OF THE WASGENERATEDBY RELATION	42
TABLE 9: ATTRIBUTES OF THE WASASSOCIATEDWITH RELATION	42
TABLE 10: HL7AGENT ROLES.....	43

TABLE 11: ATTRIBUTES OF THE WASATTRIBUTEDTO RELATION.....	43
TABLE 12: PROVENANCE CHAINING METADATA ATTRIBUTES.....	44
TABLE 13: ATTRIBUTES OF THE WASDERIVEDFROM RELATION.....	45
TABLE 14: ATTRIBUTES OF THE WASINFORMEDBY RELATION.....	46
TABLE 15: ATTRIBUTES OF THE ACTEDONBEHALFOF RELATION.....	46
TABLE 16: HEALTHCARE PROVENANCE REQUIREMENTS (FPA: FEDERATED PROVENANCE AUTHORITY; PS: PROVENANCE STORE).....	62
TABLE 17: COMPLETE LIST OF HEALTHCARE LIFECYCLE EVENTS (LCEs)	65
TABLE 18: COMPLETE LIST OF PROVENANCE MODEL CLASSES	69
TABLE 19: COMPLETE LIST OF PROVENANCE MODEL PROPERTIES	71
TABLE 20: INVERSE PROVENANCE NAMES.....	75

List of Appendices

APPENDIX A – LIFE-CYCLE EVENTS DEFINITIONS	65
APPENDIX B – PROVENANCE MODEL LISTING	69
APPENDIX C – INVERSE PROVENANCE NAMES	75
APPENDIX D – PROVENANCE DATA LINEAGE	77
APPENDIX E – IBM WATSON RESEARCH MODEL.....	78
APPENDIX F – PROVENANCE INTRA-DOMAIN MODEL.....	79
APPENDIX G – ACRONYMS AND ABBREVIATIONS	81
APPENDIX H – GLOSSARY	82
APPENDIX I – REFERENCES.....	87

1 Preface

This document is part of a series of interrelated Privacy and Security Architecture Framework (PSAF) documents that address core security, policy, and traceability topics needed to enable trustworthy interoperability for information exchange. The series of documents are:

- *PSAF Volume 1, Trust Framework for Federated Authorization (TF4FA), Conceptual Model* [[HL7 PSAF TF4FA Vol. 1](#)]: presents a general architecture for creating a trusted relationship with a healthcare partner supporting policy derivation for security and privacy. This document provides a general conceptual overview of what defines interoperable authorized exchange and what is needed to achieve it.
- *PSAF Volume 2, Trust Framework for Federated Authorization (TF4FA), Behavioral Model* [[HL7 PSAF TF4FA Vol. 2](#)]: presents a more technical behavioral model describing logical interaction among Federated Authorization components.
- *PSAF TF4FA Guide* [[HL7 PSAF Guide](#)]: presents an informative supplement that amplifies information contained in Volumes 1 and 2.
- *This Volume 3, Federated Provenance*: presents a general conceptual overview of what defines resource lifecycle events and associated provenance events, and what is needed to process, share, and leverage that provenance data for resource trustworthiness decisions (i.e., “fitness for use” decisions by resource recipients).
- *PSAF Volume 4, Audit* [[HL7 PSAF Vol. 4 Audit](#)] – planned for May 2019 Ballot

Figure 1 illustrates the document series larger context of establishing trustworthy interoperability for information exchange.

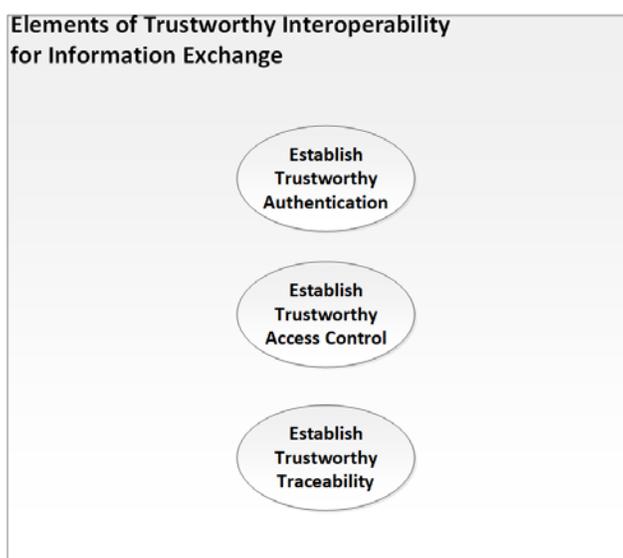


Figure 1: Elements of Trustworthy Interoperability

2 Introduction

This Domain Analysis Model (DAM) describes the conceptual-level artifacts for sharing standardized provenance information between independent security and privacy domains.

Provenance as used in this DAM follows the definition of provenance as proposed by the W3C, i.e. that “provenance is information about Entities, Activities, and people involved in producing a piece of data or thing, which can be used to form assessments about its quality, reliability or trustworthiness.” [W3C Prov Overview].

This DAM discusses provenance in terms of the business requirements governing healthcare information exchange, the parties involved in creating, sharing, and managing provenance data, and presents various models (static/informational and dynamic/behavioral) describing the relevant behaviors of those parties.

This DAM is implementation- and technology-agnostic. Nothing that follows implies or recommends a particular approach. Further, no current or emerging technologies are precluded by this conceptual model.

2.1 Provenance Overview

“Provenance of a resource is a record that describes Entities and processes involved in producing and delivering or otherwise influencing that resource. Provenance provides a critical foundation for assessing authenticity, enabling trust, and allowing reproducibility. Provenance assertions are a form of contextual metadata and can themselves become important records with their own provenance.” [W3C Prov XG FR].

At its core, provenance describes the use and production of *Entities* by *Activities*, which may be influenced in various ways by *Agents*. These core types and their relationships are illustrated in Figure 2.

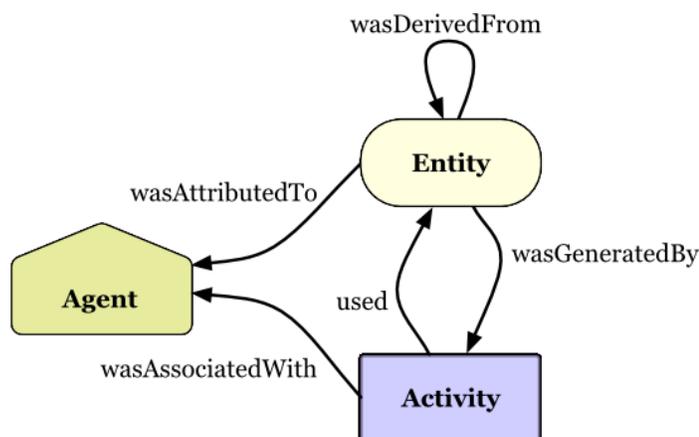


Figure 2: W3C Provenance Model

An Activity is something that occurs over a period of time and acts upon or with Entities. It may include consuming, processing, transforming, modifying, relocating, using, or generating Entities [W3C Prov DM]. Activities are how Entities come into existence and how their attributes change to become new Entities, often making use of previously-existing Entities to achieve this. They are dynamic aspects of the world, such as actions and processes. For example,

if the second version of document D was generated by a translation from the first version of the document in another language, then this translation is an Activity (adapted from [W3C Prov Primer]).

An Agent is something that bears some form of responsibility for an Activity taking place, for the existence of an Entity, or for another Agent's Activity [W3C Prov DM]. An Agent can be a person, software, process, inanimate object, organization, or other Entities that may be ascribed responsibility. Consider a lab report. To represent the provenance of that report, one could state that the person who created the report was an Agent involved in its creation, and that the software used to create the report was also an Agent involved in that Activity. (adapted from [W3C Prov Primer]).

An Agent may be acting on behalf of others and one can express such chains of responsibility in the provenance [W3C Prov Primer] (see [HL7 DPROV CDA IG] for discussion of Agents in various contexts).

2.2 Provenance and Metadata

“Metadata is used to represent properties of objects (e.g. an image). Many of those properties have to do with provenance, so the two are often. How does metadata relate to provenance? Descriptive metadata of a resource only becomes part of its provenance when one also specifies its relationship to deriving the resource. For example, a file can have a metadata property that states its size. But this is not typically considered provenance information since it does not relate to how the file was created. The same file can have metadata regarding its creation date, which would be considered provenance-relevant metadata. So even though a lot of metadata potentially has to do with provenance, both terms are not equivalent. In summary, provenance is often represented as metadata, but not all metadata is necessarily provenance.” [W3C Prov XG FR].

Core provenance metadata to be captured are:

- *Who* contributed to the generation of a resource (e.g., the participating authors, authenticators, legal authenticators, custodians, data enterers, performers, and other participants, including assembly and composing software, and scoping organizations at the document, section, and entry levels),
- *When* an information event recorded in a resource occurred (e.g., the Activity's start and stop times),
- *Where* an information event recorded in a resource occurred,
- *Why* an information event recorded in a resource occurred,
- *How* an information event recorded in a resource differs from a predecessor or successor information event, and the context surrounding that change including any privacy or security policies that influenced the manner in which the information event was changed.
- *What* provenance metadata about the information event that a recipient system may need to evaluate its authenticity, integrity, and trustworthiness, and to establish the receiver's confidence that this information is fit for use within its enterprise.

2.3 Provenance and Trust

Provenance is fundamental to trusted end-to-end flows of health data/records, capturing, retaining and rendering basic health data/record metadata – typically at the point of service/care

or separately at the point of record entry origination or update. Provenance information in healthcare can be used for the purpose tracing resources back to their origins (e.g., lab reports, entries in electronic health records).

It is essential that users know the provenance of healthcare information about to be relied upon in order to make trustworthiness (fitness for purpose) decisions before relying upon that information.¹

“Provenance data is the mechanism that memorializes and conveys provenance details. That data can then be used for many purposes such as understanding how data was collected so it can be meaningfully used, determining ownership and rights over an object, making judgements about information to determine whether to trust it, verifying that the process and steps used to obtain a result complies with given requirements, and reproducing how something was generated.” [W3C Prov Primer]

Trust is a term with many definitions and uses, but in many cases establishing trust in an object or an Entity involves analyzing its origins and authenticity. How does trust relate to provenance? Trust is often equated with provenance, and it is indeed related, but it is not the same. Trust is derived from provenance and from other data quality metrics, and typically is a subjective judgment that depends on context and use. With provenance, the focus is on how to represent, manage, and use information about resource origins, but not on detailed approaches as to how trust may be derived from it. In essence, provenance is a platform for trust algorithms and approaches [W3C Prov XG FR].

Note that authentication is often conflated with provenance because it leads to establishing trust. However, current mechanisms available for authentication address the verification of an identity or the access to a resource, such as digital signatures and access control. Provenance information may be used for authentication purposes, for example the creator of a document may provide a signature that can be verified by a third party but is only one component of authentication [W3C Prov XG FR].

In short, there are two key goals of provenance [ONC HIT S&I PI]:

- Improve the visibility of permutations of health information from creation to exchange, integration and use across multiple health information systems.
- Improve the confidence healthcare stakeholders have in the authenticity, reliability, and trustworthiness of shared data.

2.4 Healthcare Lifecycle Events (LCEs)

Underlying the Federated Provenance conceptual model are healthcare lifecycle events (LCEs) from which provenance data derives. The set of LCEs are adopted from [ISO/TS 21089] and consistent with [ISO/HL7 10781]. [Appendix A](#) provides the definitions for the LCEs.

¹ Fitness for use decisions are ultimately risk-based determinations.

3 Scope

The scope of federated provenance describes sharing of provenance data among signatories to a Federated Provenance agreement.²

3.1 Assumptions

This document assumes the following within scope:

1. Participants have coordinated and agreed to all elements of federated provenance policy.
2. All member provenance services and clients are turned on and fully operational in accordance with current configuration settings.
3. If present, participants are authorized to write to and read from a jointly shared and managed provenance records.
4. All participants are willing and capable of sharing their provenance information within the Federated Provenance context.

3.2 Limitations

Only provenance data already captured (as directed by provenance configuration) is available for sharing. It is not possible to obtain provenance data that has not been previously configured to be captured.

3.3 Preconditions within Scope

The following preconditions are required for a Federated Provenance:

1. Establishment of trustworthy identity and authorization have been successfully completed and a mutually agreed to contract provisioned into each member's local Access Control Service.
2. Use of HL7 vocabulary.
3. Establishment of Integrity policy and vocabulary.³
4. Use of the LCEs defined by [ISO/TS 21089] and noted in this document.
5. Use of the provenance events defined by W3C and noted in this document.
6. Configuration of provenance policy.
7. Mechanisms to ensure proper provenance chaining across domains are in place.

² An agreement is the participants' agreement to share provenance information for federation authority defined LCE and to use received provenance information according to federation rules.

³ This integrity field in the context of this document pertains to provenance (e.g., trustworthiness of a resource), not security (e.g., whether resource has been tampered with).

3.4 Out of Scope

The following are out of scope for this document:

1. Normative discussion of provenance processing within a participant's local operating environment is out of scope.
2. Though cited in this conceptual model, defining Audit Services is out of scope as it is detailed in [\[HL7 PSAF Vol. 4 Audit\]](#).
3. Retrieval of a resource and the Access Control Service (ACS) that supports retrieval requests are out of scope.
4. Determining and assessing provenance data lineage/quality through W3C LI/LE extensions is out of scope.
5. Determining the level of trust based on provenance data is out of scope.

4 Federated Provenance Model

The Federated Provenance model relies on the well-known notion of federation in information systems which is based on the two fundamental goals of a) maintaining autonomy for individual members while b) collaborating and sharing information among members [Heimbigner-McLeod 1985]. The federated provenance model leverages the federation model specified by NIST Special Publication 800-63-C [NIST SP 800-63-r3] as a model that allows for the conveyance of information across a set of networked systems. Based on this fundamental definition, in the provenance context, federation is an architecture model that allows conveyance of provenance information across a set of networked systems, known as Members.

Figure 3 shows a view of a Federated Provenance architecture, consisting of autonomous Members agreeing to share provenance information collected by each member. The semantics of the provenance information collected and shared in this model will be discussed in the later sections of this document.

A Federation Authority governs the federation in accordance with policies and oversees and facilitates members joining or leaving the federation in a dynamic fashion. Members who join the federation will agree to comply with the policies governing the federation, which in this case define what provenance information must be collected and shared with other members. There are no inherent constraints on the number of members who participate in a federation.

Members of a Federated Provenance Service could be individual healthcare systems, healthcare organizations, or in general any Domain in which a group of users and respective data are governed by some policy.

The core requirement for a Federated Provenance Service is for Members to capture, collect, and share provenance information, in accordance with the federation policies. This does not preclude members from retaining the same provenance information locally, or capture and collect other provenance information independently and for local use. In other words, while Members must comply with the federation provenance policies in capturing and sharing provenance information, they may also follow different local policies for local collection and organization of provenance information.

The goal of the Federated Provenance Service is to ensure access to a comprehensive cross-organization, and cross-domain provenance record corresponding to any given health information across its lifespan for all the Members of the Federated Provenance Service. This means that even when health information is transferred between different domains, there is a single source of truth for a complete provenance picture for that, so that, users in different domains can make judgments about trustworthiness, reliability, and “fitness for use” of a given piece of health information based on the provenance history, i.e., observing how the information has been generated, altered, and transformed at various domains in the federation and exchanged among them.

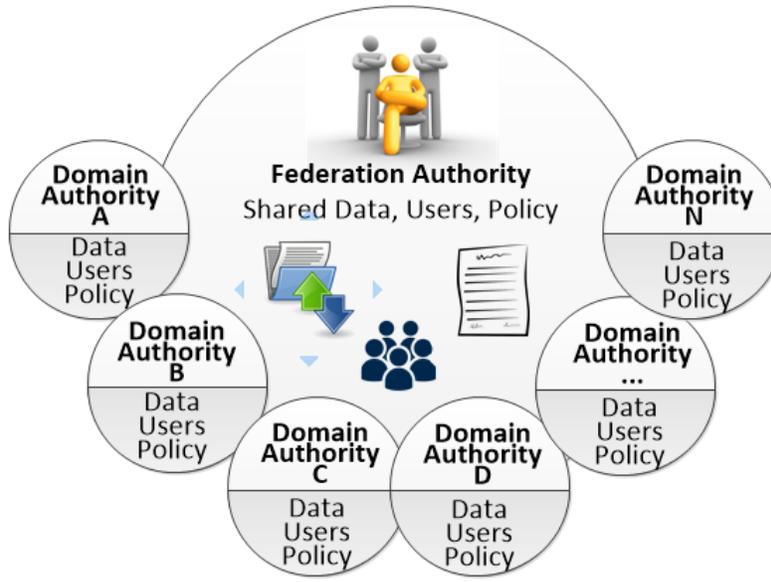


Figure 3: Federated Provenance Service

4.1 Enterprise View

Figure 4 provides a generalized business view of the Federated Provenance Service. The major components and actors are discussed below.

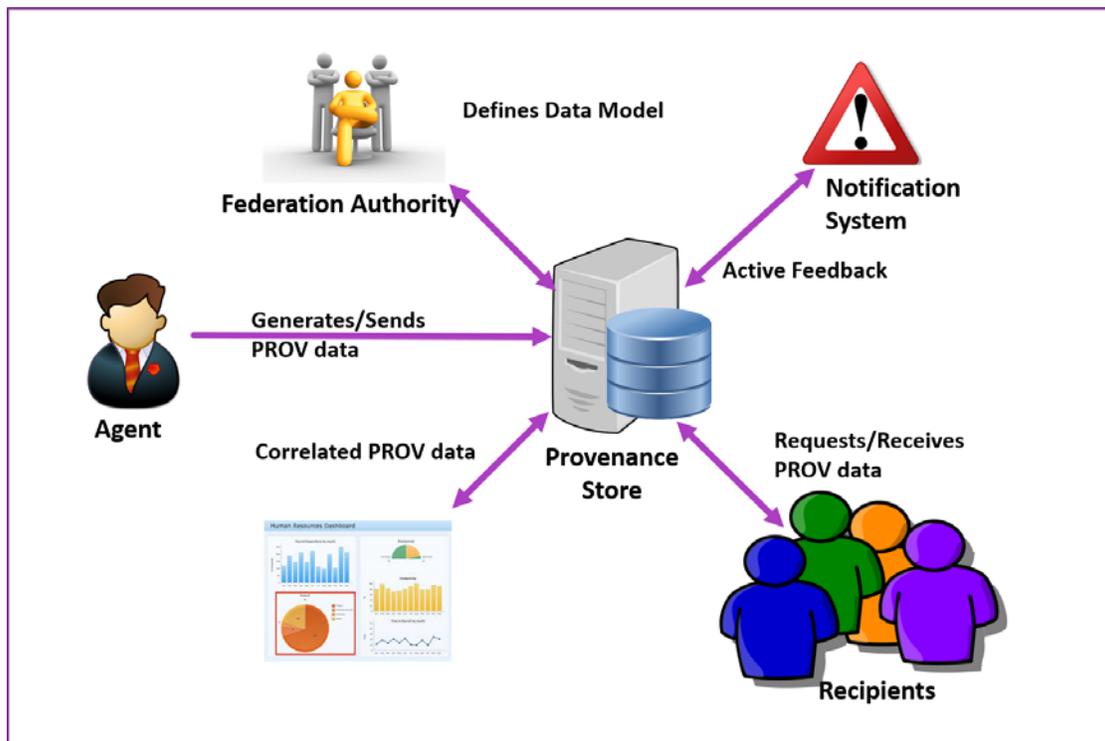


Figure 4: Federated Provenance Overview

Underpinning everything is Provenance Policy, as agreed to by all participants and managed and enforced by a Provenance Policy Authority. Provenance Policy governs all Activities and

processing within the Members of the Federated Provenance and is provisioned into all applicable contexts including the Provenance Store and all Member domains.

Provenance Stores are central storage or directories which receive and store instance of provenance information from the participating Members of the Federated Provenance Service. Provenance Stores may persist individual provenance instances or could act as a directory which only stores a pointer to the persisted provenance instance residing at one of the Member domains.

An Agent is the actor within each Member system in charge capturing, collecting, storing, and submitting provenance information to the Federated Provenance Service. This is usually a software service interfacing both the local system and the Federated Provenance Store.

Recipient is a Member system which receives provenance information from the Federated Provenance Store. This is usually a software service interfacing both the Federated Provenance Store and the local system.

Optionally, as shown in **Figure 6**, Brokers in the Federated Provenance Model can act as intermediaries between Recipient, Agents, and the Provenance Store to facilitate their communication or integration.

Provenance Analysis Service is a service provided within the Federated Provenance Service that enables queries for data analysis purposes, such as parameter-based queries, aggregations, and trends. Enhanced analytics may leverage additional and external data such as business data to support correlation queries.

Provenance Notification Service is a service provided within the Federated Provenance Service that enables invoking or calling external services (e.g., web services, emails, or text messages) or sending alerts based on provenance-related events, depending on policies and configurations. Notifications may be sent to various targets such as individuals, reporting systems, and dashboards. Advanced analytics can be leveraged to create notifications or alerts based on patterns and trends, such as suspicious pattern of updates to a resource. Notifications and alerts may be based on the global policies of the Federated Provenance Service, but it could also be configured by individual users interested in certain events and Activities, in the form of a subscription.

Members of a Federated Provenance Service often take the roles of an Agent, Recipient, and a User of the Analysis Service. and may also be recipient of notifications and alerts generated by the Notification Service. These different role may be implemented by different software components within the Members' Provenance Systems.

From a Provenance Store perspective, the Federation consists of the Agents, Recipients, and their Users which it serves, while the Provenance Store itself often takes the role of a Recipient. From a Recipient perspective, the Federation consists of the Provenance Store and the Agents that direct provenance information to the Provenance Store. From an Agent perspective, the Federation consists of Recipients including the Provenance Store and other Recipients that consume provenance data.

Provenance Federation Authority (PFA) provides governance for the Federated Provenance Service. Participating Members are under the governance of a Provenance Federation Authority and associated contracts.

4.2 Federated Provenance Data Flows

Figure 5 shows the Federated Provenance Service and its data flows as discussed in this section. Note that these data flows are focused on the exchange of provenance information among different Members. The initial capturing of provenance, which happens within an Agent's local operating environment, is often triggered by standard Life-Cycle Events that lead to creating corresponding provenance instance (see Section 8.10). The details of capturing provenance within a participating Member of the Federated Provenance Service is implementation-specific, and will not be discussed.

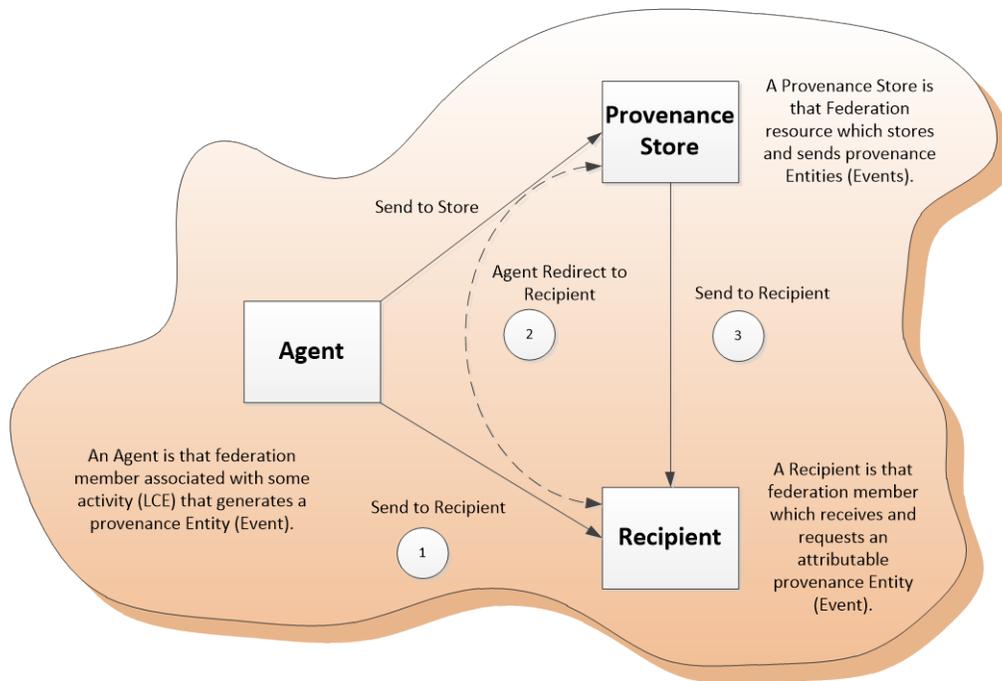


Figure 5: Federated Provenance Data Flows

4.2.1 Direct

This flow covers the direct sharing of provenance information with a Recipient by an Agent, in which an Agent directly provides provenance data to a Recipient. The Recipient can be the Federated Provenance Store, or another Member of the Federated Provenance Service.

4.2.2 Redirect

In this flow, the Agent shares the provenance information indirectly with a recipient via submitting it to the Federated Provenance Store. The Federated Provenance Store acts as an intermediary, which makes the provenance information available to the Recipient.

4.2.3 Query

In this flow, a Recipient queries the Federated Provenance Store to get access to a specific instance of provenance pertaining to some data. The query can be reference-based, seeking the provenance corresponding to a specifically identified data resource; or it can be parameter-based, by specifying a criteria, which could potentially cover a collection of provenance information [W3C PROV AQ].

4.2.4 Brokered Exchange

These optional flows cover the cases of exchanging provenance information using an intermediary, which acts as a Broker as shown in **Figure 6**. Brokered Exchange flows are essentially variations of the three flows discussed above in which the communication between an Agent and the Provenance Store and/or the communication between the Provenance Store and a Recipient that take place via a Broker, which facilitates or simplifies the integration by providing common communication or integration support as well as possible translation.

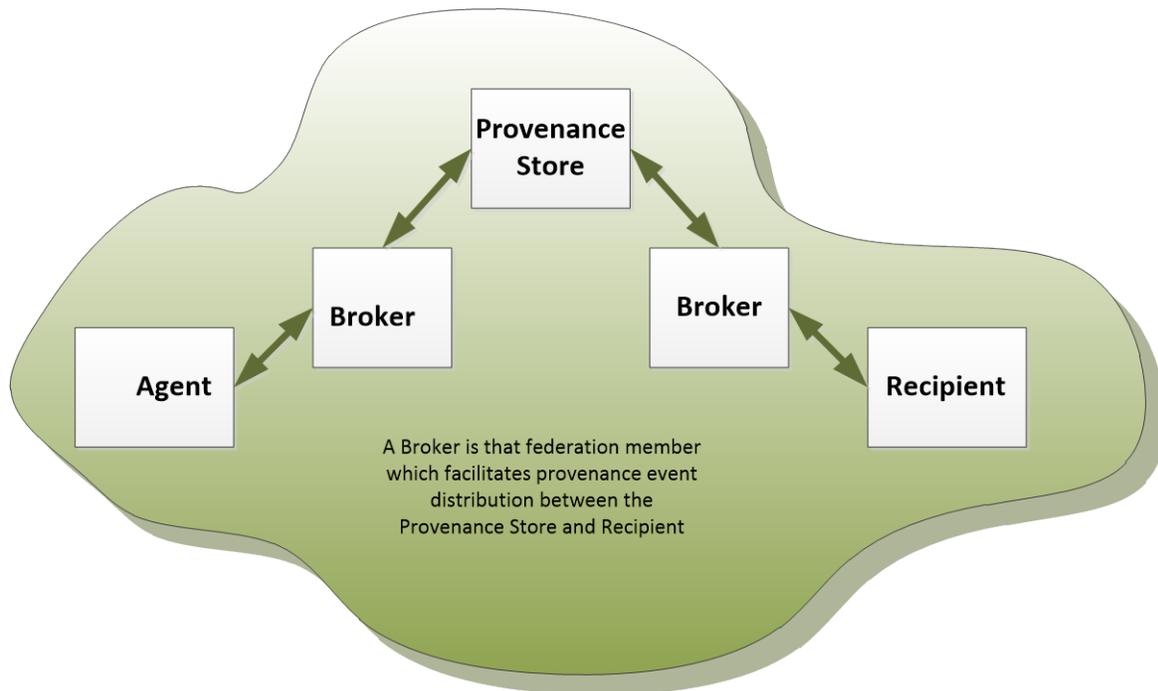


Figure 6: Optional Brokered Exchange flows in the Federated Provenance Model

5 Storyboards

These Storyboards are intended to be the basis for the [Activity Diagrams](#) in Section 7 and the [Class Models](#) in Section 8. These traceability links are Domain Analysis Model requirements from which those models are derived and to which those models are traceable as required by the HL7 Canonical Domain Analysis Model Guidance.^[1]

5.1 Storyboard 1 - Tracking Production History

Storyboard 1 explores how a Provider might use Provenance Records to establish confidence in a patient's records by identifying the sources, processes, and reliability of devices used to produce these records. This Storyboard is loosely based on the Ad-hoc PDex Member History Request Use Case from the Da Vinci Payer Data Exchange [HL7 Da Vinci PDex]. For details about persona characteristics, roles, and perspectives in the following storyboards, see [Section 5.6 Storyboard Persona below](#).

Alice enrolled in Good Health Plan six months ago after transitioning from her Medicaid Health Plan, in which she had been enrolled for many years.

Dr. Bob is her new primary care physician in Good Health Plan. Dr. Bob wants to check Alice's health history to prepare for her first appointment.

Using an application (App) in his EHR, Dr. Bob requests that Alice's Good Health Plan, send him claims history for services and medications, and any clinical documentation that the plan has received from providers caring for Alice since her enrollment. Dr. Bob clicks the App's button indicating that he also wants to review the provenance of the requested information from Good Health Plan.

5.1.1 Claims Workflow Provenance

Good Health Plan responds to Dr. Bob's requests by having its Clearinghouse compile the set of Alice's X12 and NCPDP Claims for services and medications from Good Health Plan providers, and transforming these into FHIR Claims.

As shown in **Figure 7**, Good Health Plan's Clearinghouse records the following provenance information to track the production history of the transformed claims, and will send a Provenance Report to Dr. Bob along with Alice's FHIR Claims. Within the FPD in which Dr. Bob participates, a health plan creates a Provenance Record documenting the source of the records, the identity of the health plan and the action taken to transform the data into FHIR Claims.

The following is an example of the Provenance Record information that might be required to be sent by a Health Plan, and is loosely based on the "Claim Information from Provider" example from Da Vinci [HL7 Da Vinci PDex].

- Agent 1 (claims receiver/retainer), which is the Clearinghouse claims adjudication software, that received and retained Entity t1 (set of Alice's X12 and NCPDP Claims)
- Agent 2 (claims transformer), which is the Clearinghouse FHIR transformation software, that acts on behalf of Agent 1 (claims receiver/retainer)

^[1] HL7 Specification: DAM Specifications and Requirements, Release 1
https://www.hl7.org/implement/standards/product_brief.cfm?product_id=463

- Agent 3 (verifier), which is the Clearinghouse transform algorithm, that calibrates the Plan Entity used by Activity 2 to transform X12 and NCPDP Claims into FHIR Claims by checking standards conformance and reliability of their map
- Agent attributes to be included in the Provenance record
 - Agent identifier, name, role, contact information, and affiliation. [See [Table 5 Attributes of Agent Class](#) prov:id, prov:label, prov:type]
 - Agent type including Software Agent and organization [See [Table 6 Agent Types](#) prov:Organization, prov:SoftwareAgent.]
- Entity t2 (FHIR Claims), which were derived from Entity t1 (set of Alice’s X12 and NCPDP Claims)
- Plan Entity is the ActivityDefinition, which is the algorithm used by Activity 2 to perform the transform of Entity t1 (X12 and NCPDP claims) into Entity t2 (FHIR claims)
 - Entity attributes in the Provenance record may include identifiers, names, direct representation of the ActivityDefinition instance for the transform and its location, creation time, and licenses for the map. [See [Table 2 Attributes of the Entity Class](#) prov:id, prov:name, prov:value, prov:location, prov:generatedAtTime, and rights.]
- Activity 2 (transform), which used Entity t1 to generate Entity t2
 - Activity 2 attributes such as name, start/end times, and status of complete. [See [Table 4 Attributes of the Activity Class](#) prov:label, prov:startTime, prov:endTime, prov:status.]
 - Activity 2’s relationship with the output type FHIR Claim Entity. [See [Table 8 Attributes of the wasGeneratedBy Relation](#) prov:time and prov:role.]
 - Activity 2 used Plan Entity, which has the role of being the transform algorithm. [See [Table 7 Attributes of the Used Relation](#) prov:role]
 - Activity 2’s relationship with Agent 2’s role. [See [Table 9 Attributes of the wasAssociatedWith Relation](#) prov:role.]
- Activity 3 (verity), which was informed by Activity 2 and used the Plan Entity description of the algorithm for Activity 2 transform to calibrate the correctness of the resulting (Entity t2) FHIR claim. The Provenance Record may include:
 - Activity 3 attributes include an annotation that Plan Entity is calibrated correctly for Activity 2’s FHIR Claim output, and status of complete. [See [Table 4 Attributes of the Activity Class](#) prov:label, prov:startTime, prov:endTime, prov:description, and prov:status.]
 - Activity 3’s wasInformedBy relationship with Activity 2 attributes an identifier of a link to the informing predecessor Activity. [See [Table 14 Attributes of the wasInformedBy Relation](#) prov:id, prov:Activity]
 - Activity 3’s wasAttributedTo relationship with the Agent attributes including an identifier [See [Table 11 Attributes of the wasAttributedTo Relation](#) prov:id, prov:role.]
 - Activity 3 used Plan Entity, which has the role of being the transform algorithm that the Agent verifies as being correctly calibrated. [See [Table 7 Attributes of the Used Relation](#) prov:role]

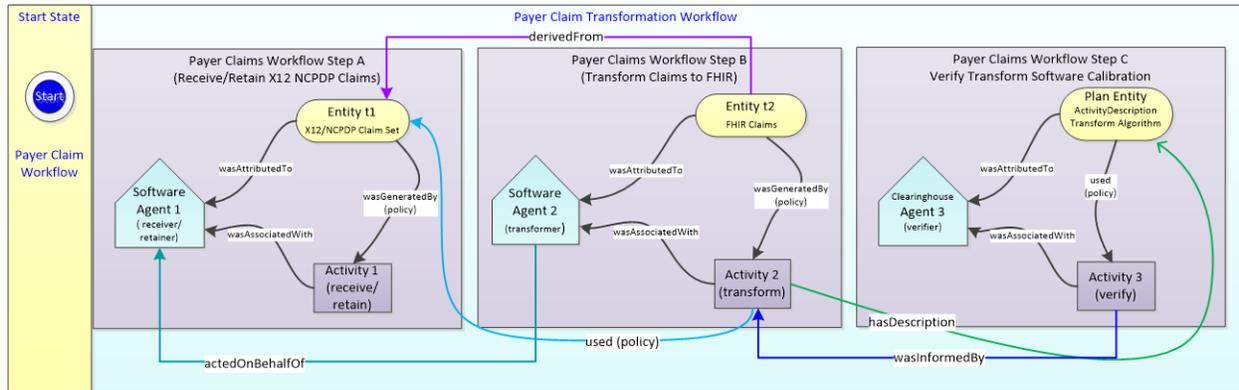


Figure 7: Storyboard 1 - Claims Production Tracking

Based on this Provenance Report, Dr. Bob surmises that he can trust the FHIR Claims information because Good Health Plan’s Clearinghouse regularly tracks the origins for the Claims and verifies the accuracy with which the transforms are produced.

5.1.2 Clinical Documentation Workflow

In response to Dr. Bob’s request, Good Health Plan’s Clearinghouse also compile all of the clinical documentation received from Alice’s providers for coordination of care, prior authorization and referral requests, risk adjustment, quality reporting, and for claims and medical necessity and appropriateness documentation [HL7 Da Vinci PDex]. This includes the Provenance Report of key Lifecycle Events for all documentation, which Good Health Plan requires providers to track. Dr. Bob can review the Provenance Report to determine his confidence in the FHIR Claim and FHIR DocRef information.

5.2 Storyboard 2 – Directed and Federated Provenance Chain

For details about persona characteristics, roles, and perspectives in the following storyboards, see [Section 5.6 Storyboard Persona below](#).

Storyboard 2 explores the differences from a Provider’s perspectives between a point-to-point “Last Hope” Provenance Chain and a Federated Provenance Chain when answering the questions:

Who are the persons or organizations involved in the production of a data set?

Who can provide answers to questions about this data set?

This Storyboard is loosely based on the Da Vinci Payer Data Exchange Implementation Guide [HL7 Da Vinci PDex] and contrasts it with the Federated Provenance approach described in this document.

According to Storyboard 1, Dr. Bob received the provenance information for FHIR Claims and FHIR DocRef Resources from Good Health since Alice enrolled as he requested.

But this only includes the provenance information from the “Last Hops” accumulated by Good Health. I.e., He only receives the “point-to-point” provenance information as it moved through the Payer and Provider Workflows a “Hop at a Time”.

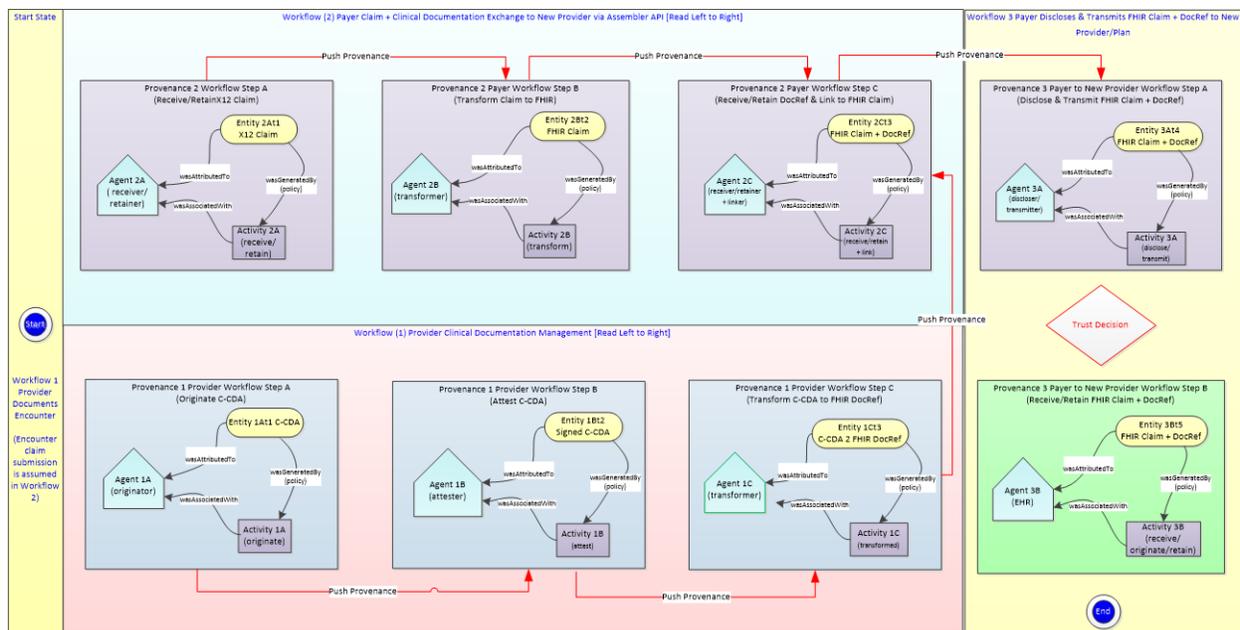


Figure 8: Storyboard 2 – Claims and Clinical Documentation Response to Dr. Bob with Hop-to-Hop Provenance

Based on the reason for Alice’s appointment, Dr. Bob is particularly interested in gathering information related to Alice’s previous PCP’s referral to a nephrologist, Dr. Carla; lab results for her diabetes; visits to her podiatrist, Dr. Foote, related to diabetic neuropathy pain; her DME orders to Ez-Supplies for a wheel chair, walker, and orthotic boots; and prior authorization requests from Karen Kind, her community care manager, for a ramp to be built to her front door.

While reviewing the C-CDAs referenced in the FHIR DocRef, Dr. Bob notices that Alice’s nephrologist, Dr. Carla did not attest to her C-CDA upon which she based her prior authorization request to refer Alice to a podiatrist, Dr. Foote for diabetic neuropathy. To evaluate the confidence he has in this Payer furnished information, Dr. Bob wants more information about the trustworthiness of the claims and clinical documentation on Alice sent by Good Health Plan.

Since he can’t back-track past the “Last Hop” provenance record from Good Health, he queries the Federated Provenance Store [FPS] for the Provenance Domain in which he is a participant.

His FPS query parameters include Alice’s identifiers, the date range for the target information Entities in scope, i.e., her last 2 years of encounters, observations, and orders and the associated claims for Alice’s providers. His query excludes the last 6 months because he already has that provenance information

The FPS store returns provenance information prior to Alice’s enrollment in Good Health Plan while she was a member of a State Medicaid Plan.

While reviewing the FPD Provenance Report for Alice’s State Medicaid Plan claims and C-CDAs, Dr. Bob discovers that Dr. Carla’s C-CDA sent with a prior authorization request for a podiatry services which was not attested to, was based on her consult notes about a call with Alice’s endocrinologist, Dr. Dan. See Storyboard 2.1 Attribution Diagram below.

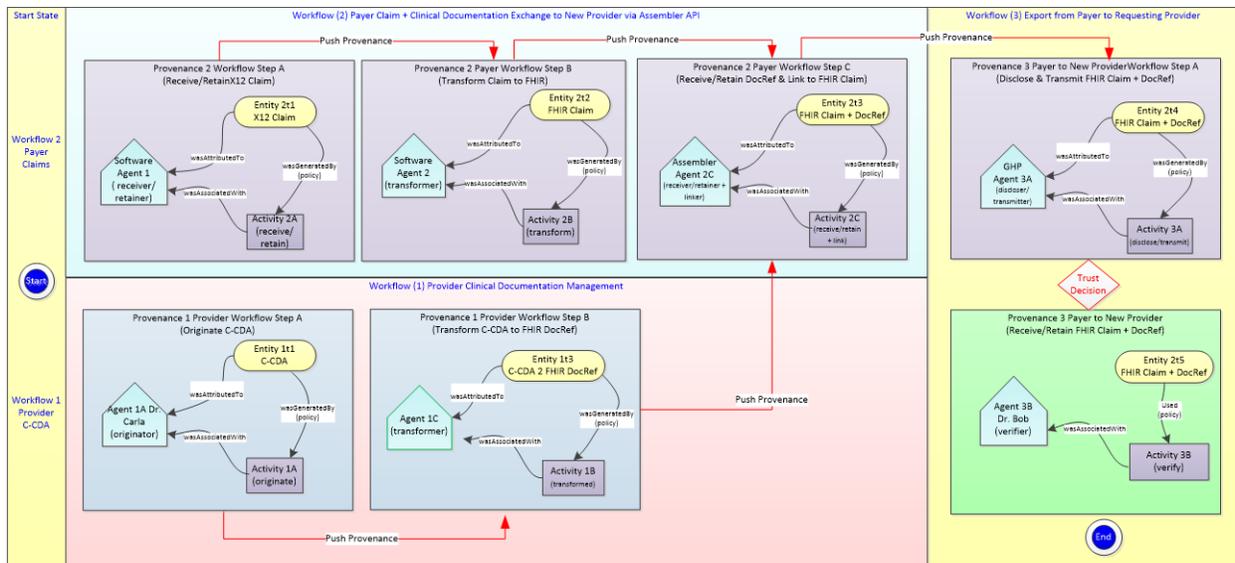


Figure 9: Storyboard 2 Attribution

Dr. Bob was not aware that Alice was seeing Dr. Dan, nor that he was the first to diagnose her diabetes condition. He also notes that the FPS sent Dr. Dan claims information,

Dr. Bob again queries the FPS for Alice’s provenance information where Dr. Dan is the Agent. The FPS returns a Provenance Report for Dr. Dan, which were created more than 6 months ago. Dr. Bob is able to retrieve these Entities using the Entity.location, and intends to follow up with Dr. Dan using the Agent.phone for more information about Dr. Carla’s authorization request to Good Health Plan.

While Dr. Bob appreciates getting Last Hop Chain of Provenance, he also wants the ability to query or subscribe to the FPD to get finer-tuned Provenance Reports of interest to him rather than rely on what each health plan wants to send him.

5.3 Storyboard 3 - Locate Error Sources and Sharing with Protections

Storyboard 3 features the ability of Provenance Reports to help providers uncover what seems to be missing clinical information with the potential downside of revealing confidential information by finding the location of possible error sources in a data set.

- Where does this data set come from?
- What process was used to create the data set, an old one or a new one?
- If a data set does not look right, how can one determine how it was obtained?

For details about persona characteristics, roles, and perspectives in the following storyboards, see [Section 5.6 Storyboard Persona below](#).

While reviewing the Provenance Report retrieved from FPD in Storyboard 2, Dr. Bob notes provenance information about a self-assessment of Alice’s socio-economic, emotional, and behavioral health, which Alice’s Care Manager, Karen Kind, received and retained. Using the assessment’s Entity.identifier and Entity.location, Dr. Bob is able to retrieve Alice’s self-assessment from the Care Coordination Platform he shares with Alice’s other providers. He notes that Alice includes information about her mental health about which Dr. Bob was unaware. He

was also unaware that Alice has been taking antidepressant medications and seeing Dr. Mind Urbiz, a psychiatrist.

Dr. Bob sees that the Provenance Record does not include claims or clinical documents from Dr. Urbiz. He concludes that Alice doesn't share information about her mental health information with anyone else, and Dr. Mind is very supportive of Alice's privacy preferences. For example Dr. Mind ensures that Alice can pay for her psychiatric services and medications out of pocket in accordance with HIPAA. Dr. Bob decides that he will broach the question of mental health with Alice but leave it up to her as to whether to disclose more. Since he inadvertently came across this sensitive information, which is specially protected under state law, he does not feel an obligation to incorporate it into his own records since he was not authorized by Alice to access it.

Dr. Bob is grateful that he now understands the discrepancy in the clinical information that he has obtained because of the chain of provenance, but he is also mindful of the confidential nature of this information, and will defer to his patient's privacy preferences about whether to incorporate it into his records.

He is also aware that any provider who is a member of the FPD, has a Clinical Decision Support systems (CDS), which is privileged to receive patient safety alerts from the FPS. So even if a patient pays for medications out of pocket, and no claim is sent to a payer, the pharmacy will still send provenance about the prescribed medication to the FPS.

As agreed by FPD governance, when the FPS records provenance related to prescriptions for medications with high likelihood of dangerous drug-drug interactions, that provenance will be sent as an alert to the CDS of Alice's providers but will be labels as "restricted".

Access is permitted only if the requester has clearance for restricted information, such as a provider authorized by Alice's consent directive, or if the provider is attempting to prescribe a dangerously contraindicated drug, which would trigger the provider's CDS to throw a "Break the Glass" warning so that the provider overrides the restriction to find out why the intended prescription is a patient safety issue.

This is exactly what happened when Alice's podiatrist, Dr. Foote, entered a prescription for an opioid to treat Alice's diabetic neuropathy pain. Dr. Foote's CDS alerted him to "Break the Glass". When he did, he realized that Alice has been taking an anti-anxiety medication. The CDS recommends non-opioid pain medications, so Dr. Foote changed the prescription upon learning of this error and the gap in information, which the Provenance Notification was able to provide.

5.4 Provenance Reporting Workflow

The following diagram shows an example of Provenance Chaining based loosely on the Da Vinci Payer Data Exchange use of point-to-point Provenance, which carries forward to the next "Hop" in a workflow. In this example, there are three workflows. It also shows that Provenance can also be sent to a FDS.

Walk-through by Workflows from Left to Right

- Provider Workflow
- Payer Workflow

- Payer Discloses/Transmits - New Provider/Plan Imports Patient's Provider's Claims/Clinical Documentation

The Provider Workflow 1 results in a FHIR DocRef for clinical information, such as C-CDAs, originated and attested to by a provider.

The Payer Workflow 2 results in a set of FHIR Claims based on transforms of a patient's X12 and NCPDP Claims, which are linked with the same patient's FHIR DocRef.

Workflow 3 is the export/disclosure of the FHIR Claims and FHIR DocRef to a new provider. For simplicity, this is a "push" transaction.

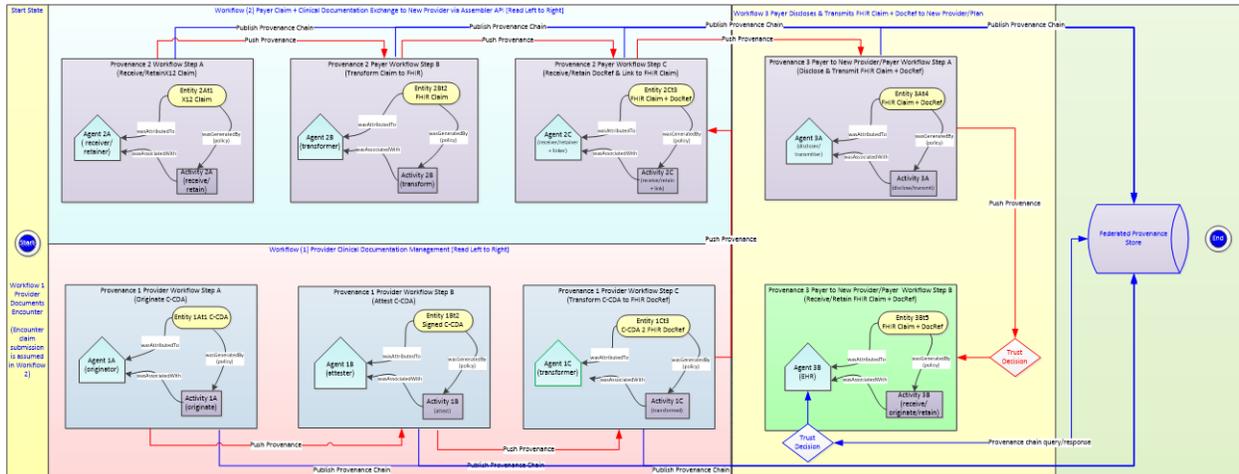


Figure 10: Provenance Reporting Workflow

The Provenance Reporting Workflow illustrates two compatible approaches to recording and tracking a chain of provenance, each one of which may result in a different type of Trust Decision: The Push and Publish Provenance Workflows. For detailed discussion see [Section 8.9 Provenance Chaining](#).

Push Provenance Workflow: This workflow requires that a Provenance Report be generated by the Provenance Event Agent after each Provenance Event related to an Entity during its Lifecycle. The Provenance Report is either the initial Provenance Report of the origination of an Entity, or a compilation of previous Provenance Reports with the one generated after the current Lifecycle Event involving an Entity derived from the original Entity.

In circumstances where the Agent of a Provenance Event did not receive or cannot retrieve the preceding Provenance Report, the Agent may be required as a term of participation in the FPD, to generate a Provenance Report of the preceding Provenance Event post facto.⁴ A post

⁴ Since Health Plans compile information from many sources to create a Member's Health History it is important that data traceability is maintained. The HL7 FHIR Provenance resource is used for this purpose. It is used to identify the source of information, the agents the data passed through and the actions the performed on the data.

Health Plans maintain provenance records that they receive as part of any exchange of FHIR data. Where a FHIR Provenance resource is not provided, such as when data is received from other non-FHIR sources, the Health Plan shall create FHIR Provenance record(s) to identify the source of the information being received and the actions that is taken on the data, such as converting from one format to another. Health Plans pass on Provenance records in any PDex information exchange. Provenance is covered in more detail in [Section 6-7 Handling Data Provenance](#) [HL7 Da Vinci PDex].

facto Provenance Report may be deemed less trustworthy than one generated by the Agent of the preceding Provenance Event.

Publish Provenance Workflow: This workflow requires that a Provenance Report or its Registry Location be published by the Provenance Event Agent after each Provenance Event to a Federated Provenance Store after each Provenance Event related to an Entity during its Lifecycle. There is no need to compile the current Provenance Report with preceding Provenance Reports related to an Entity during its Lifecycle.

Authorized FPD participants interested in the Provenance chain related to an Entity during its Lifecycle will be able to retrieve the full chain of Provenance Reports from the Federated Provenance Store (PS) as reported by the Agents of each Lifecycle Event. Depending on the methods by which the Provenance Report chain is published and logged, interested parties will likely have more trust in the retrieved Provenance chain from a FPS than they would have in the compilation of Provenance Reports about the same Entity because there is a higher level of confidence in the completeness, reliability, and authenticity of those reports as illustrated in the following Storyboards.

5.5 Privacy Preferences and Provenance

Additional privacy protection considerations are required when implementing the recording and exchange of provenance because this information is itself possibly confidential. Additional precautions are needed for governance of provenance information within a Federated Provenance Domain (FPD) as is also illustrated in the proceeding Storyboard 3.

5.6 Storyboard Persona

The following list of the Storyboard Persona describes their roles and perspectives in more detail.

- Alice: Patient with diabetes, and her condition has progressed to include diabetic neuropathy. She also have mental health conditions.
- Clearinghouse: An intermediary that runs:
 - A Claims Adjudication System, which receives, processes, and retains claims information for Good Health Plan and the State Medicaid Plan
 - Transformer Software, which converts X12 and NCPDP Claims into FHIR Claims.
- Dr. Bob: PCP recently assigned to Alice by her new health plan, Good Health Plan, in which she enrolled 6 months ago.
- Dr. Carla: A nephrologist to whom Alice was referred by Dr. Dan prior to her enrollment in Good Health Plan, and before she became a patient of Dr. Bob.
- Dr. Dan: Alice's long-time endocrinologist who is not in the Good Health plan provider network. Dr. Dan referred Alice to Dr. Carla after her diabetes diagnosis a year ago.
- Ez-Supplies: Alice's DME provider, to which Dr. Foote has sent diabetic neuropathy off-loading orders for a wheelchair and a walker.
- Dr. Foote: Alice's Podiatrist, to whom Dr. Carla referred Alice for off-loading orthotic care and orders.

- Good Health Plan: Alice's new health plan, in which she enrolled 6 months ago.
- Dr. Mind Urbiz Alice's psychiatrist for over most of her adult life. Alice doesn't share information about her mental health information with anyone else, and Dr. Mind is very supportive of Alice's privacy preferences. For example Dr. Mind ensures that Alice can pay for her psychiatric services and medications out of pocket in accordance with HIPAA.
- Karen Kind: Alice's Care in the Community Care Manager, to whom Dr. Foote has referred for social services including installation of a wheelchair/walker ramp to her house.
- State Medicaid Program: Alice's health plan for 4 years prior to enrolling in Good Health Plan 6 months ago.

6 Use Cases

Details of Federated Provenance Service use-cases are discussed in **Table 1**. The involved actors in these use-cases are Provenance Store, Agent, Recipient, Broker, Analysis Service, and Notification Service, as defined in Section 4.1. **Figure 11** provides a diagrammatic summary.

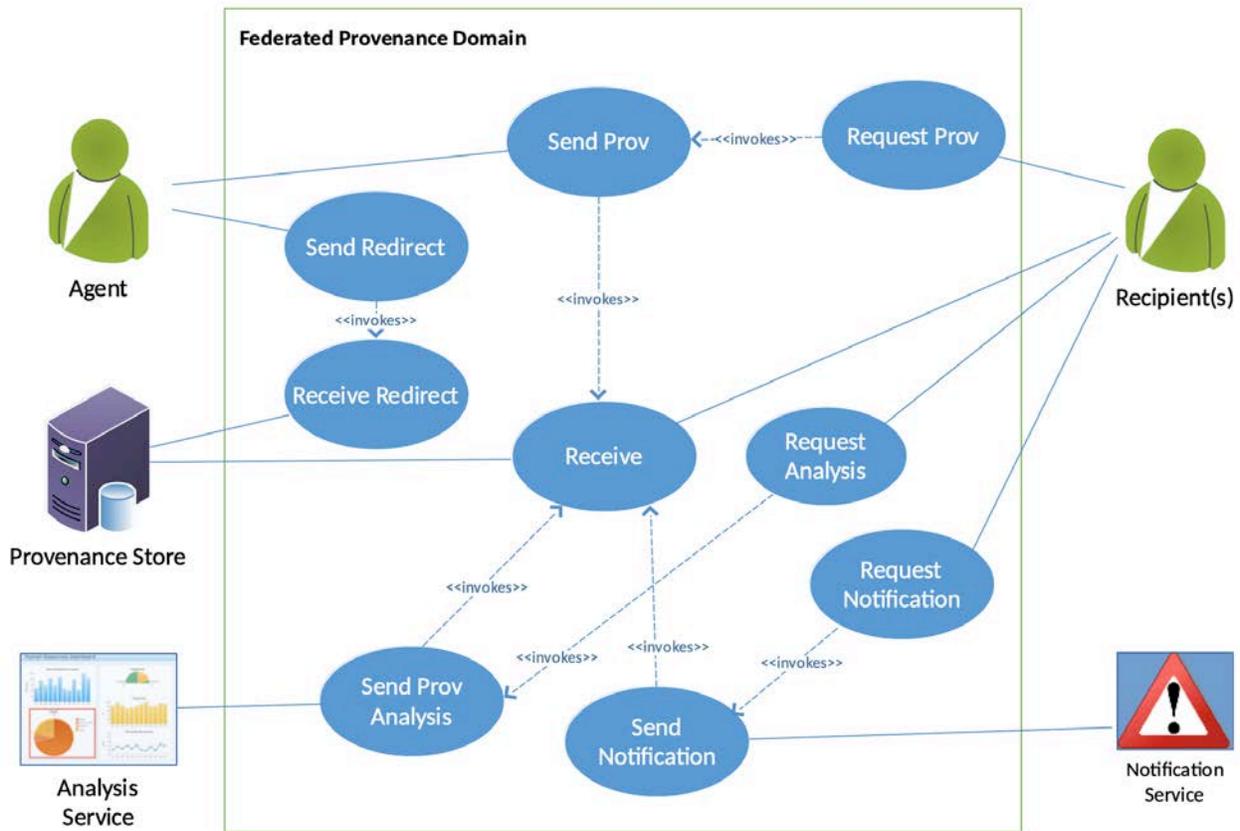


Figure 11: Federated Provenance Domain

Table 1: Federated Provenance Use-Case Descriptions

1. Send Provenance (Normative)	
Actors: Agent, Provenance Store, Recipient	Trigger Events: Designated Life-Cycle Events (LCEs) per policies.
Description: <ol style="list-style-type: none"> 1. Triggered by an applicable LCE, the Agent captures and generates a provenance instance. 2. The Agent sends the provenance instance via a push to the Recipient or Provenance Store. 	
Pre-Conditions: <ul style="list-style-type: none"> - The Agent, Provenance Store, and Recipient are on-boarded members of the Federated Provenance Service. - Through an unspecified out-of-band mechanism, the Agent is aware of the identifying address for the Recipient or the Provenance Store. 	
Post-Conditions: <ul style="list-style-type: none"> - The provenance instance is received, consumed, and/or persisted by the Recipient or the Provenance Store. 	
2. Send Redirect (Normative)	
Actors: Agent, Provenance Store, Recipient	Trigger Events: Designated Life-Cycle Events (LCEs) per policies.
Description: <ol style="list-style-type: none"> 1. Triggered by an applicable LCE, the Agent captures and generates a provenance instance. 2. The Agent sends the provenance instance via a push to the Provenance Store. 3. The Provenance Store relays the provenance instance to the Recipient. 	
Pre-Conditions: <ul style="list-style-type: none"> - The Agent, Provenance Store, and Recipient are on-boarded members of the Federated Provenance Service. - Through an unspecified out-of-band mechanism, the Agent is aware of the identifying address for the Provenance Store. - Through an unspecified out-of-band mechanism, the Provenance Store is aware of the identifying address for the Recipient. 	
Post-Conditions: <ul style="list-style-type: none"> - The provenance instance is received and persisted by the Provenance Store. - The provenance instance is received, consumed, and/or persisted by the Recipient. 	
3. Request Provenance (Normative)	
Actors: Provenance Store, Recipient	Trigger Events: N/A
Description: <ol style="list-style-type: none"> 1. The Recipient sends the Provenance Store a query requesting a specific provenance instance or a collection of provenance instances based on some criteria. 2. The Provenance Store consumes the Recipient's query and identifies the matching provenance instances. 3. The Provenance Store sends the provenance instance (or collection) to the Recipient. 	

Pre-Conditions:	
<ul style="list-style-type: none"> - The Provenance Store and Recipient are on-boarded members of the Federated Provenanance Service. - Through an unspecified out-of-band mechanism, the Recipient is aware of the identifying address for the Provenance Store. - Through an unspecified out-of-band mechanism, the Recipient is aware of the query language for requesting provenance information from the Provenance Store. 	
Post-Conditions:	
<ul style="list-style-type: none"> - The provenance instance is received, consumed, and/or persisted by the Recipient. 	
4. Request Analysis (Non-Normative)	
Actors: Recipient, Analysis Service	Trigger Events: N/A
Description:	
<ol style="list-style-type: none"> 1. The Recipient sends the Analysis Service a query requesting for analytical information regarding provenance instances within the the Federated Provenance Service domain. 2. The Analysis Service consumes the Recipient’s analysis query and invokes the required analytics processes to extract the information. 3. The Analysis Service sends the analysis results to the Recipient. 	
Alternatives:	
If running the analysis is time-consuming and process-heavy, the results may not be ready instantly. In such cases:	
<ol style="list-style-type: none"> 3. The Analysis Service sends the Recipient a follow-up ticket. 4. The Recipient returns to the Analysis Service and presents the follow-up ticket. 5. If the results are ready, the Analysis Service sends the analysis results to the Recipient. 	
Pre-Conditions:	
<ul style="list-style-type: none"> - The Analysis Service and Recipient are on-boarded members of the Federated Provenanance Service. - Through an unspecified out-of-band mechanism, the Recipient is aware of the identifying address for the Analysis Service. - Through an unspecified out-of-band mechanism, the Recipient is aware of the query language for requesting provenance analytics information from the Analysis Service. 	
Post-Conditions:	
<ul style="list-style-type: none"> - The provenance analytics is received, consumed, and/or persisted by the Recipient. - If the request is in the form of a subscription, the Recipient continues to receive the results until subscription is canceled. 	
5. Request Notification (Non-Normative)	
Actors: Agent, Notification Service, Analysis Service (optional)	Trigger Events: N/A
Description:	
<ol style="list-style-type: none"> 1. The Recipient sends the Notification Service a subscription request to be notified when provenance-related events matching a specific query occur. 2. The Notification Service consumes the Recipient’s request and registers a subscription for the Recipient. 	

Pre-Conditions:

- The Notification Service, Recipient, and Analysis Service are on-boarded members of the Federated Provenance Service.
- Through an unspecified out-of-band mechanism, the Recipient is aware of the identifying address for the Notification Service.
- Through an unspecified out-of-band mechanism, the Recipient is aware of the query language for requesting provenance information from the Provenance Store and optionally, the query language for requesting provenance analytics information from the Analysis Service, in order to specify a pattern of interest for the notification subscription.
- Through an unspecified out-of-band mechanism, the Notification Service is aware of the identifying address for the Analysis Service (optional).

Post-Conditions:

- An event watch is registered within the Notification Service to monitor for the events matching the subscription request.
- The Notification Service sends alert to the subscribed Recipient until the Recipient cancels the subscription.

7 Activity Diagrams (Normative)

An Activity diagram is used to show the different Activities that need to be carried out to accomplish the goals of a system (or higher-level Activity). It also shows the organization and sequencing of those Activities.

7.1 *Life-Cycle Events*

Life-Cycle Events start with the conceptual realization of an Entity (creation), and thereafter its instantiation in memory. The Entity did not exist before creation and exists afterwards as a realized Entity (LCE created) which can be called and used. An initialized Entity may be persisted (stored) and retained as a permanent object.

7.2 *Functional Flows*

This section describes system Activities specific to functional uses-cases and inter-relationships among the principal actors of Agent, Recipient and Provenance Store.

7.2.1 *Send Provenance*

Based on a request or a pre-established sharing agreement, an Agent sends provenance information meeting federated sharing rules are passed from one organization to another as depicted in **Figure 12**. The sending organization may be the recipient (pass through) or the author of the provenance information sent.

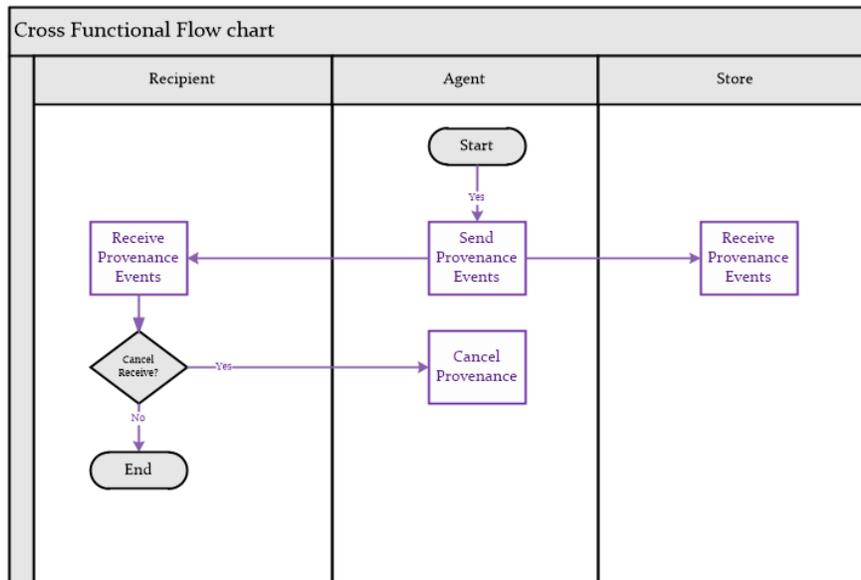


Figure 12: Activity Diagram for Send Provenance

7.2.2 *Send Provenance Redirect (Optional)*

Based upon policy or operational information, an Agent submits a redirect request to a Provenance Store to forward one or more provenance events to one or more federation member recipients as depicted in **Figure**. This is followed by receiving the provenance event by the Recipient as depicted in **Figure 14**.

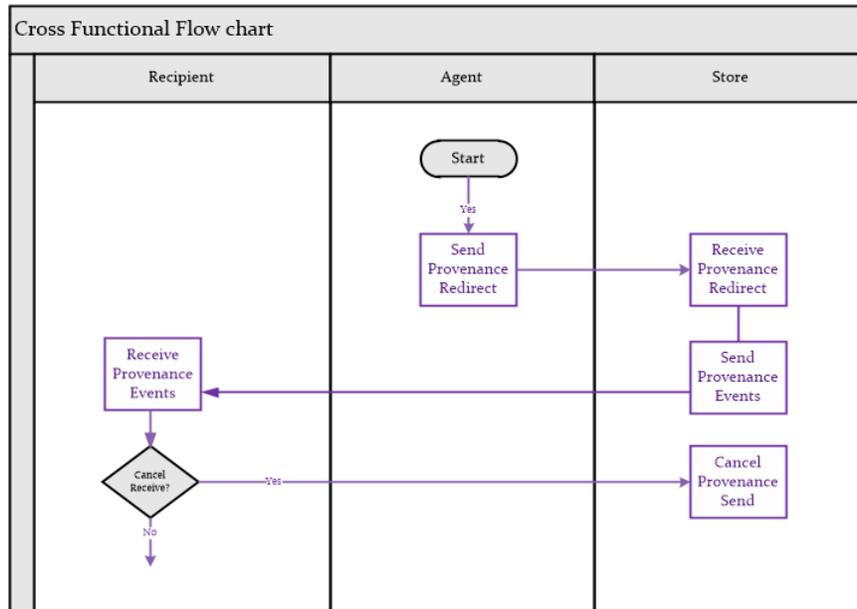


Figure 13: Activity Diagram for Send Redirect

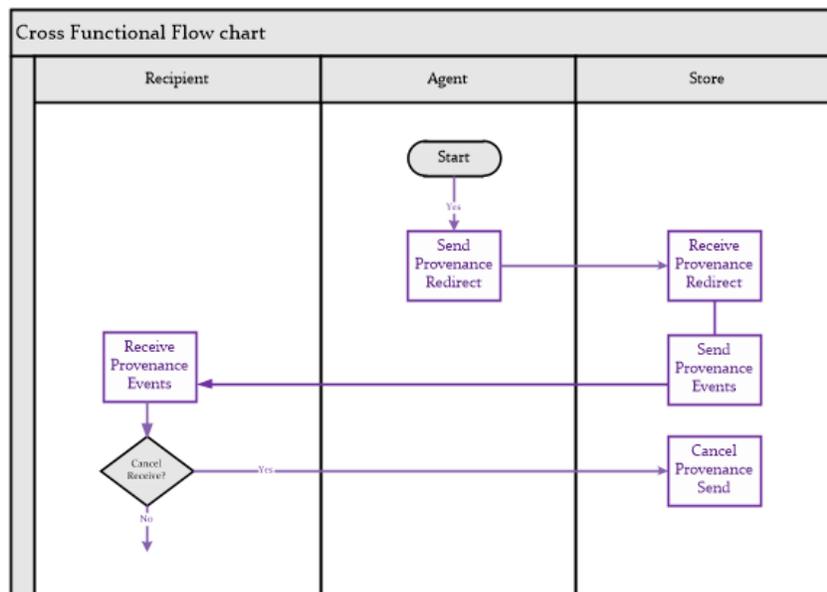


Figure 14: Activity Diagram for Receive Provenance

7.2.3 Request Provenance

A Recipient requests provenance information of some type meeting policy requirements directly from another Agent(s). Agents respond and continue to provide requested provenance until a cancellation request is received. Similarly, a Recipient requests to receive provenance information meeting policy requirements directly from the Provenance Store. The subscription remains in effect until altered or cancelled. These are depicted in **Figure 15** and **Figure 16**.

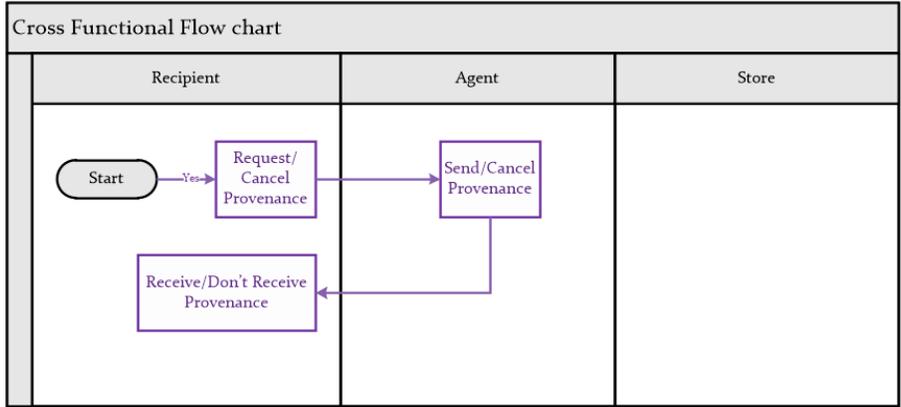


Figure 15: Activity Diagram for Request Provenance from Agent

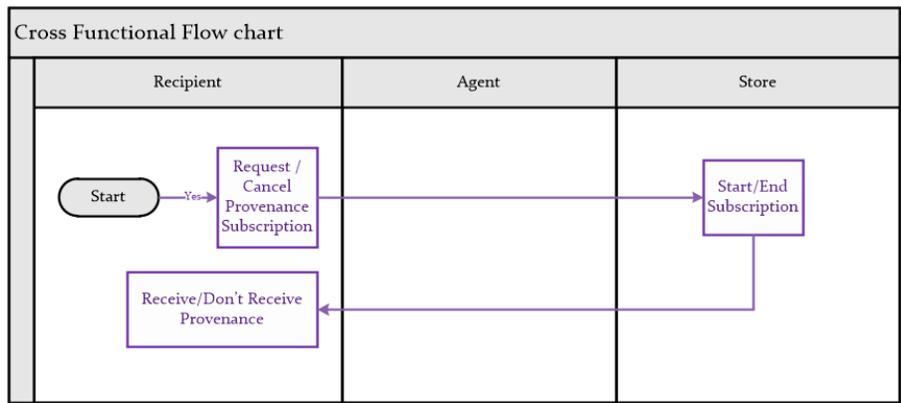


Figure 16: Activity Diagram for Request Provenance from Provenance Store

7.2.4 Request Provenance Store Analysis

A recipient requests access to Provenance Store Analysis (provided as a service) as depicted in **Figure 17**. Analysis can include multiple categories which can be requested in total or in part so long as they are available per federation agreement. The recipient continues to receive requested analysis until the subscription is cancelled.

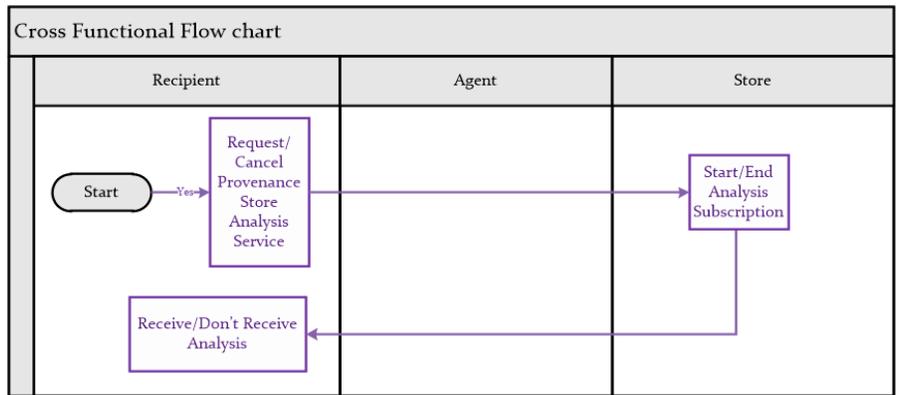


Figure 17: Activity Diagram for Request Analysis

7.2.5 Request Provenance Store Notifications

A recipient requests access to Provenance Store Notifications as depicted in **Figure 18**. Notifications may include warnings, system down-time alerts, etc. as established by policy. Custom notification can be specified by the Recipient who is interested in notifications triggered by certain events specified in the form of a provenance or analysis query. The recipient continues to receive notifications until the service is cancelled.

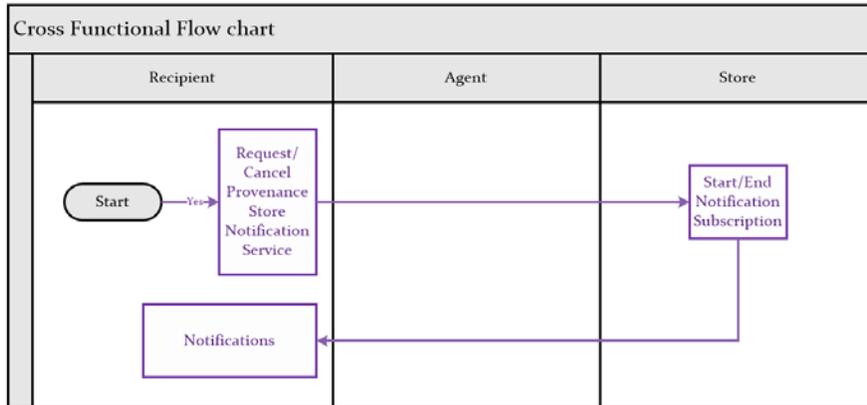


Figure 18: Activity Diagram for Request Notification

8 Class Models (Normative)

Provenance is captured by the relations between three core elements, Agents, Entities, and Activities. In the most general case, an Agent (potentially acting *on behalf of* a different Agent) invokes an Activity (which could potentially be *informed by* other Activities) that *uses* some Entities and could end up *generating* new Entities. The resulting Entities are said to be *derived from* the original Entities and could be *attributed to* one or all of the involved Agents. Moreover, Entities could be grouped together in the form of special type of Entity called a Collection. Based on this model, Provenance is defined as a class which captures in instance of the relationship between the classes and relations shown in **Figure 19**.

The core Entities and their relations which constitute the main classes of this model are discussed in this rest of section. Three relations, *wasDerivedFrom*, *actedOnBehalfOf*, and *wasInformedBy*, are specifically relevant to provenance chaining which will be discussed at the end of this section.

This model is based on the W3C Provenance Data Model [W3C Prov DM], the corresponding ontology [W3C Prov Ontology], and International Virtual Observatory Alliance (IVOA) Provenance Data Model [IVOA Prov DM]. A comprehensive informational list of W3C provenance classes and properties is given in Appendix B.

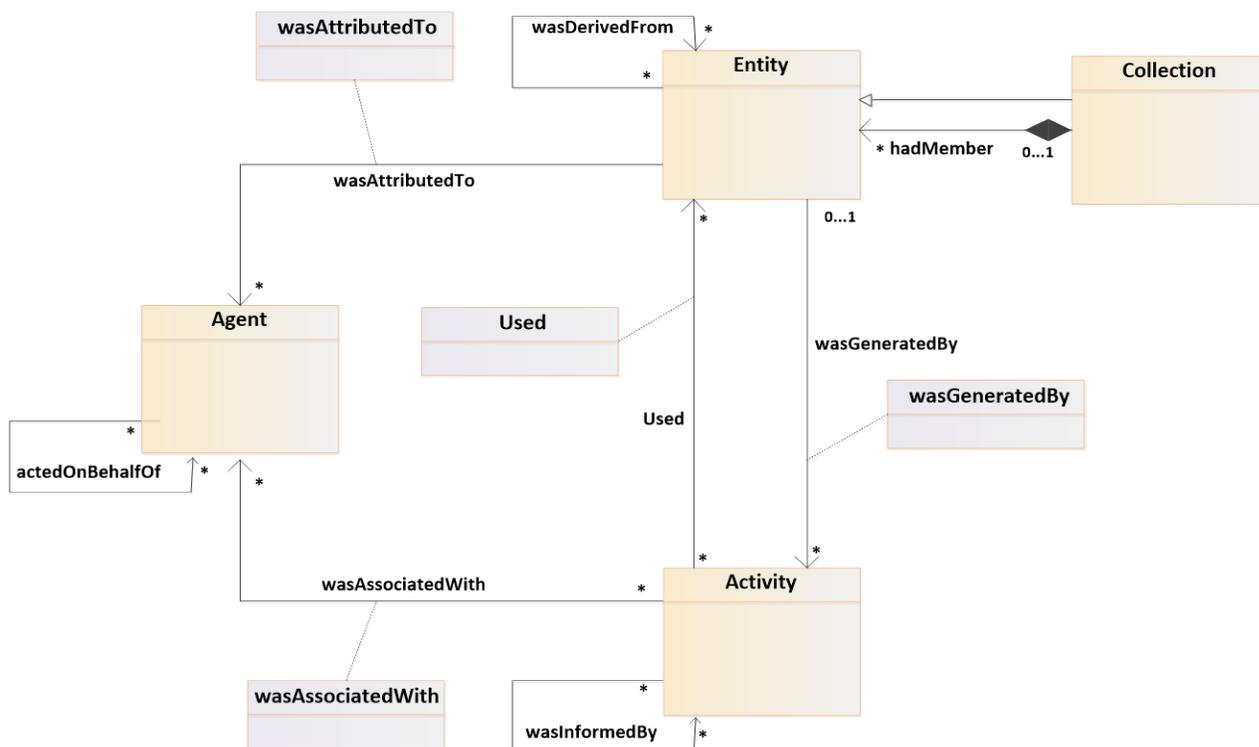


Figure 19: Provenance Data Model

8.1 Entity

Entity refers to a physical, digital, conceptual, or other kind of thing with some fixed aspects, such as medical records, medical images, Clinical Document Architecture (CDA), HL7 V2

message, Fast Healthcare Interoperability Resources (FHIR), Medical Device Data, calibration data, discharge summary, patient consent directive, etc.

What constitutes an Entity depends on the granularity requirements of the system; for example, depending on the application and use-cases, an Entity could be a FHIR resource, a CDA document, a section in a CDA document, a file, or a number in a table.

Table 2: Attributes of the Entity Class provides the attributes of the Entity class along side the corresponding attribute names from the W3C Provenance Data Model where it exists.

Table 2: Attributes of the Entity Class

Attribute	Type	Description	W3C PROV
id	(qualified) string	a unique id for this Entity (unique for its domain)	prov:id
name	string	a human-readable name for the Entity (to be displayed by clients)	prov:label
value		a direct representation of an Entity	prov:value
organization	string	a named legal institution such as a healthcare provider.	prov:Organization
location	string	an identifiable geographic or non-geographic place, such as coordinates, a URL, or a row-column pair.	prov:atLocation
type	string	a provenance type, i.e. one of: prov:collection, prov:bundle, prov:plan; or any specialized Entities	prov:type
annotation	string	text describing the Entity in more detail	prov:description
creationTime	datetime	date and time at which the Entity was created (e.g. timestamp of a file)	prov:generatedAtTime
destructionTime	datetime	date and time at which the Entity was erased or invalidated	prov:invalidatedAtTime
rights	string	access rights for the Entity, values: public, secure or proprietary.	

8.2 Collections

A Collection is a group of Entities which can be treated as one single Entity from the provenance perspective; for example, a Collection can include all the Entities produced by the same Activity. Collections can be used to collect Entities with the same provenance information together in order to hide complexity where necessary by creating another layer of abstraction with higher granularity. For example, a collection could be a FHIR Bundle containing all immunizations, allergies, or medications for a patient, or, a CDA document can be considered a collection containing various more fine-grained Entities.

As shown in **Figure 19**, the Entity-Collection relation can be modelled using the Composite design pattern: Collection is a subclass of Entity, but also an aggregation of one or many Entities, each of which could in turn be Collections.

Table 3: Attributes of the hadMember Relation.

Attribute	Type	Description	W3C PROV
↻collection	(qualified) string	the id of the Collection.	prov:id
↻Entity	(qualified) string	identifiers of the member Entities.	prov:id

8.3 Activity

Activity refers to an action, process, or a series thereof, occurring over a period of time, performed on, or caused by, Entities, often resulting in new Entities. For example, treatment in a general Activity performed by a care team which can result in an Entity such as a discharge summary document.

Activities start and end at particular points in time and during their lifespan can *use* and *generate* a variety of Entities as captured by the *used* and *wasGeneratedBy* relations which will be discussed further below. **Table 4: Attributes of the Activity Class** provides a summary of attributes of the Agent class.

Table 4: Attributes of the Activity Class

Attribute	Type	Description	W3C PROV
id	(qualified) string	a unique id for this Entity (unique for its domain)	prov:id
name	string	a human-readable name for the Entity (to be displayed by clients)	prov:label
startTime	datetime	start of an Activity	prov:startTime
endTime	datetime	end of an Activity	prov:endTime
annotation	string	additional explanations for the specific Activity instance	prov:description
status	string	can be used to describe the terminal status of the Activity (e.g. completed, aborted, error...)	

8.4 Agent

An Agent describes the party responsible for a certain Activity. It could be a person, a group, a software system, a team, a project, or an organization. For example, the physician who makes a diagnosis, the patient who signs and submits a consent directive, the administrative staff who recorded and entered the patient's consent into the system, or a software agent that encrypts a data element. If an Agent acts on another Agent's behalf, this is reflected by the *actedOnBehalfOf* attribute.

Table 5: Attributes of Agent Class. Agent types considered in this model are presented in **Table 6: Agent Types**.

The association of an Activity with an Agent is captured by the many-to-many relationship *wasAssociatedWith*. Also, an Entity can be associated with an Agent via the *wasAttributedTo* relation. These relations are discussed further below.

Table 5: Attributes of Agent Class

Attribute	Type	Description	W3C PROV
id	(qualified) string	a unique id for the Agent, unique for its domain.	prov:id
name	string	A common name for this Agent; e.g. first name and last name; project name, agency name.	prov:label
type	string	type of the Agent as given in Table 6: Agent Types .	prov:type
email	string	contact email of the Agent	
affiliation	string	Affiliation of the Agent	
address	string	Address of the Agent	
phone	string	Phone number	

Table 6: Agent Types

Type	Description	W3C PROV
Party		prov:Agent
Individual	a person, specified by name, email, address (though all these parts may change in time)	prov:Person
Organization	a healthcare provider, payer, affiliate, service organization, institute, standards body or scientific project	prov:Organization
SoftwareAgent	a software Agent is running software, e.g. a cron job or a trigger.	prov:SoftwareAgent

8.5 Used

An Entity is Used by an Activity, when it is used as an input in the course of that Activity. For example, the Activity *DNA Sequencing* used the Entity *DNA Data*. Aside from the participating Activity and Entity, this relationship also includes attributes to record the time when using the Entity by the Activity is started, as well as the role of the Entity in the Activity. **Table 7: Attributes of the Used Relation** summarizes the attributes of this relation.

The Used relation is closely coupled to the Activity, so the relationship is modeled as a composition (note the used of the filled diamond in **Figure 19**) to indicate that if an Activity is deleted, the corresponding Used relations need to be removed as well, while the Entities used still remain.

Table 7: Attributes of the Used Relation

Attribute	Type	Description	W3C PROV
id	string	an identifier for this relation	prov:id
role	string	role of the Entity, defined as what it is being used for	prov:role
time	datetime	Time at which the usage of an Entity started	prov:time
↻Activity	link	link to an Activity	prov:Activity
↻Entity	link	link to an Entity	prov:Entity

8.6 wasGeneratedBy

An Activity can result in the generation of a new Entity; this is captured by the *wasGeneratedBy* relation. For example, a discharge summary document can be generate as a result of an in-patient encounter.

As shown in **Table 8: Attributes of the wasGeneratedBy Relation**, aside from links to the Activity and Entity involved, the time of generation of the Entity and the role of the Entity in the relationship which deterimes the output type is acptured by additional attribuetes in the relation.

Table 8: Attributes of the wasGeneratedBy Relation

Attribute	Type	Description	W3C PROV
id	string	an identifier for this relation	prov:id
role	string	role of the Entity that is generated by an Activity, defines which output type it is	prov:role
time	datetime	Time at which the generation of an Entity is finished	prov:time
↻Activity	link	link to an Activity	prov:Activity
↻Entity	link	link to an Entity	prov:Activity

As shown in **Figure 19**, the *wasGeneratedBy* relation is closely coupled with the Entity via a composition, since without its Entitty a *wasGeneratedBy* relation cannot meaningfully exist. So, if an Entity is deleted, then all the *wasGeneratedBy* relation instances associated to that Entity must also be deleted. There is a one-to-many multiplicity between Activity and Entity in the *wasGeneratedBy* relation because an Activity can generate many Entities, but it is assumed that an Entity can be generated by only one Activity.

8.7 wasAssociatedWith

As shown in **Figure 19**, an Agent is responsibility for an Activity is captured using the *wasAssociatedWith* relation. For example, a physician can be associated with the Activity diagnosis. Aside from the links to the Activity and the Agent, the role of the Agent in the Activity can be recorded as an attribute of this relation as shown in **Table 9: Attributes of the wasAssociatedWith Relation**.

Table 9: Attributes of the wasAssociatedWith Relation

Attribute	Type	Description	W3C PROV
id	string	an identifier for this relation	prov:id
role	string	role of the Agent	prov:role
↻Activity	link	link to an Activity	prov:Activity
↻Agent	link	link to an Agent	prov:Agent

Table 10 shows a list of Agent roles adopted by this model based on the HL7 Data Provenance Implementation Guide for CDA Documents [HL7 DPROV CDA IG].

It is desired to have at least one Agent for each Activity but this is not a normative constraint. There can also be more than one Agent for each Activity with different Roles and one Agent can be responsible for more than one Activity. This many-to-many relationship is made explicit in

Figure 8 Storyboard – Claims and Clinical Documentation Response to Dr. Bob with Hop-to-Hop Provenance by adding the two following relation classes: *wasAssociatedWith*, that relates an Activity to an Agent, and *wasAttributedTo*, that relates an Entity to an Agent.

Note that the Agent associated with an Entity via a *wasGeneratedBy* relation may be different from the Agent associated with the Activity that created an Entity via *wasAssociatedWith*. An Agent performing a task is not necessarily given full attribution, especially if it acts on behalf of another Agent.

Table 10: HL7Agent Roles

Assembler	A provenance author which collates and repackages existing content (with existing provenance information intact) as an automated response to a query [HL7 DPROV CDA IG].
Composer	A provenance author which collates and repackages existing content in response to an author’s selection, and which may incorporate new content generated by the author in the process [HL7 DPROV CDA IG].
Patient	A provenance author in the role of patient or related person who acts as the patients advocate but is clearly not a member of the provider-related organizations [HL7 DPROV CDA IG].
Provider	A provenance author in the role of clinical staff [HL7 DPROV CDA IG].
Device	A provenance author which captures and creates new information independently of a human author [HL7 DPROV CDA IG].
Individual	A provenance author in the role of non-clinical human person, such as administrative, clerical, canteen, legal, police, volunteer.
Organization	A provenance author in the role of company, firm, institution, group, establishment, federation, society, etc.

8.8 *wasAttributedTo*

An Agent’s responsibility for an Activity or Entity, potentially in some role, is described using the properties *wasAttributedTo* as shown in **Figure 19**. For example, a diagnosis can be attributed to a physician using this relationship to indicate that the physician is responsible for making the diagnosis. **Table 11: Attributes of the *wasAttributedTo* Relation** shows the attributes of this relation.

Table 11: Attributes of the *wasAttributedTo* Relation

Attribute	Type	Description	W3C PROV
id	string	an identifier for this relation	prov:id
role	string	role of the Agent in this relation	prov:role
↻Agent	link	link to an Agent	prov:Agent
↻Entity	link	link to an Entity	prov: Entity

8.9 Provenance Chaining

The relationships between different instances of provenance are captured by *provenance chaining*. Depending on the type of chaining, the successor and predecessors in a chain sequence can be Agents, Activities, or Entities. Chaining is based on a generic relation that links a *predecessor* to a *successor* as shown in **Figure 20**.

Based on the W3C Provenance Ontology [W3C Prov Ontology], three of the relations defined by the core class model (**Figure 19**), *wasDerivedFrom*, *wasInformedBy*, and *actedOnBehalfOf* can form chains of, respectively, Entities, Activities, and Agents. A fourth type of chaining can be formed among Entities based on pairs of *used* and *wasGeneratedBy* relations associated with the same Activity, when an Activity uses some Entities to generate some new Entities. [Figure 7: Storyboard 1 – Claims Production Tracking](#) illustrates chaining with *wasDerivedFrom*, *wasInformedBy*, *actedOnBehalfOf*, and *used* relations. These chain-forming relations are discussed in the rest of this sub-section.

Note that provenance chaining is different from provenance data lineage, which is metadata primarily related to provenance data quality and enriched details of data transformations. As noted earlier, provenance data lineage is out of scope for this document; however, a brief informational overview is presented in Appendix D – Provenance Data Lineage.

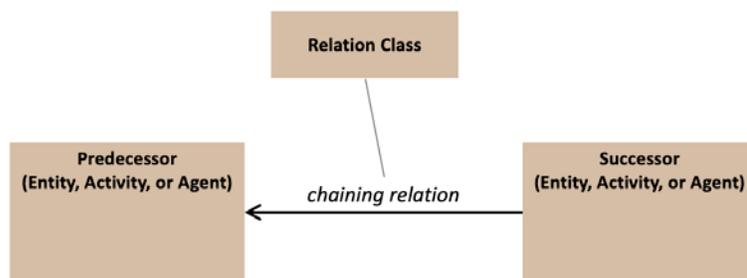


Figure 20: Provenance Chaining

Any additional information pertaining to the provenance chain can be modeled as metadata associated with the chaining relation. Any of the chaining relations discussed below could be augmented by adding such metadata attributes to capture further context about the chain. **Table 12** provides a summary of the metadata attributes for provenance chaining supported by this model.

Table 12: Provenance Chaining Metadata Attributes.

Attribute	Type	Description	W3C PROV
name	string	a human-readable name for the Entity , to be displayed by clients.	prov:label
↻chain_event	link	link to the event that changed predecessor to successor	prov:event
performer	string	performer of the change	prov:Agent
author	string	author of the provenance metadata	prove:Agent
facilitating software	string	link to facilitating software	prov:Activity
signature	string	whether there was a signature on the predecessor prior to incorporation	prov: signature
policy	string	applicable provenance policies for recording this provenance metadata	prov:description
security label	string	security labels, e.g. regarding integrity and confidentiality of the provenance metadata	prov:description

type	string	a provenance type	prov:type
annotation	string	text describing the Entity in more detail	prov:description
value		provides a value that is a direct representation of an Entity	prov:value

8.9.1 Entity Chains (*wasDerivedFrom*)

A chain can be formed when an Entity is connected to another Entity based on the fact that one was derived from the other, which is modeled by the *wasDerivedFrom* relation as shown in **Figure 21**. The attributes of this relation is summarized in **Table 13**.

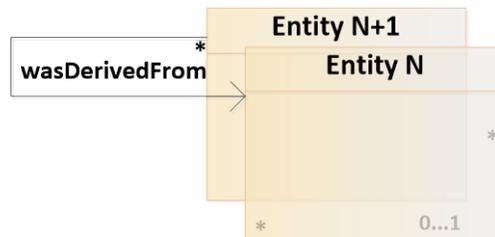


Figure 21: Provenance Chaining based on the *wasDerivedFrom* Relation.

Derivation chains between Entities capture provenance instances (e.g, transformation) between Entities when the Agent or the Activity involved in the derivation is not of interest or is unknown. For example, when the Agent or Activity involved in de-identifying a FHIR resource is not of interest, a provenance chain can be recorded simply by capturing that the de-identified FHIR resource *wasDerivedFrom* the original FHIR resource.

Note that although the Entity generated by an Activity is often derived from the Entity or entities used that Activity, the *wasDerivedFrom* relation cannot always automatically be inferred from following existing *wasGeneratedBy* and *used* relations alone. If there is more than one input and more than one output to an Activity, it is not clear which Entity was derived from which and *wasDerivedFrom* relation makes the derivation explicit. For this reason, the Entity chains resulting from pairs of *used* and *wasGeneratedBy* relations are discussed as a separate type of chaining below.

Table 13: Attributes of the *wasDerivedFrom* Relation.

Attribute	Type	Description	W3C PROV
id	(qualified) string	An identifier for this relation	prov:id
➤generatedEntity	link	link to the generated successor Entity	prov: Entity
➤usedEntity	link	link to the predecessor Entity from which the generated Entity was derived	prov: Entity

8.9.2 Activity-Based Chaining (*wasInformedBy*)

A chain can be formed when an Activity influences or provides input to another Activity, modeled by the *wasInformedBy* relation, as shown in **Figure 22**. For example, the Activity of rendering a FHIR consent resource can be informed by the Activity of filling a consent questionnaire by the patient. Attributes of this relation are summarized in **Table 14: Attributes of the *wasInformedBy* Relation.**

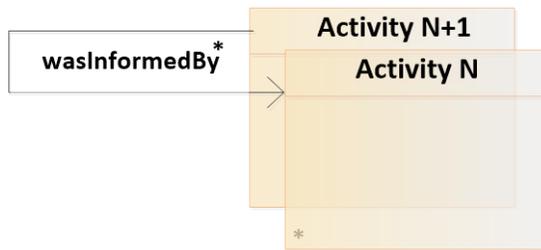


Figure 22: Provenance Chain based on the wasInformedBy Relation.

A *wasInformedBy* relation between two Activities often suggests that the informed Activity used an Entity that was generated by the informing Activity, however, since these Entities are not unknown or are not of interest, only the link between the Activities is recorded.

Table 14: Attributes of the wasInformedBy Relation.

Attribute	Type	Description	W3C PROV
id	(qualified) string	An identifier for this relation	prov:id
→informed	link	link to the Activity being informed by another (“second”) Activity	prov:Activity
→informant	link	link to the informing predecessor Activity (“first” Activity)	prov:Activity

8.9.3 Agent-Based Chaining (*actedOnBehalfOf*)

A chain can be formed based on the link between an Agent that acts on behalf of another Agent as shown in **Figure 23**, for example, when a patient’s legal guardian signs a consent form on behalf of the patient, or when an administrative assistant can assign a patient to a specialist on behalf of a general practitioner. Attributes of this relation are summarized in **Table 15**.

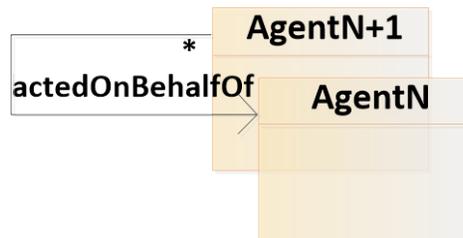


Figure 23: Provenance Chain based on the actedOnBehalfOf Relation

Table 15: Attributes of the actedOnBehalfOf Relation.

Attribute	Type	Description	W3C PROV
id	(qualified) string	an identifier for this relation	prov:id
→delegator	link	Link to initial Agent	prov:Agent
→delegate	link	link to the Agent that acted on behalf of the initial Agent.	prov:Agent

8.9.4 Activity-Entity-Based Chaining (*used, wasGeneratedBy*)

Activities can use and generate Entities during their lifespan, as captured by the *used* and the *wasGeneratedBy* relations. When the same Activity uses a predecessor Entity to generate a new

Entity, this indirect link can form a chain as shown in **Figure 24**. This chaining records the fact that some Entities were used as an input to an Activity which eventually generated certain other Entities as output. The use of an Entity as an input in generating another Entity implies that the information provided by the predecessor were used in generating the successor, and therefore, the chains of this type capture the information flow between Entities.

Note that this form of chaining is conceptually similar to the chaining based on the *wasDerivedFrom* relation, but in this form of chaining, the Activity is an essential component of the chaining which is not omitted. Moreover, this type of chaining captures a more general information flow link between Entities which may not be as strong as derivation, especially when multiple Entities are used as the input for generating the successor Entity. For example, when a physician uses a number of observations in a diagnosis which leads to generation of a careplan, the careplan is not directly *derived* from the observations, but its generation is linked to those observations via this broader form of chaining which captures the information flow from those observations into the resulting careplan.

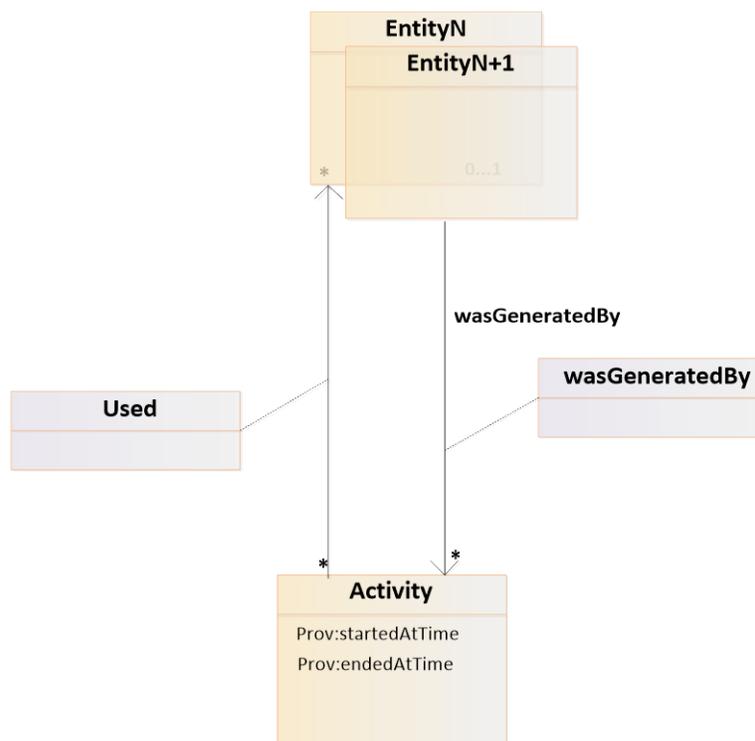


Figure 24: Provenance Chain based on used and wasGeneratedBy Relations

8.10 Healthcare Life-Cycle Events as Provenance

Life-Cycle Events (LCE) are standard events in a healthcare system as identified by various standards (as will be discussed below). These events capture the most common events in a healthcare information system.

This section presents models of these LCEs as instantiations of the provenance model, in the form of a provenance instance or a provenance chains, as presented and discussed in the previous subsections. These will be referred to as Provenance Events. Providing a provenance model for these events facilitates automatic capturing and generation of provenance information triggered by existing known events in a healthcare system.

As **Figure 25** illustrates, LCEs are specializations of the well-known set of Create, Read, Update, Delete (CRUD) security operations.

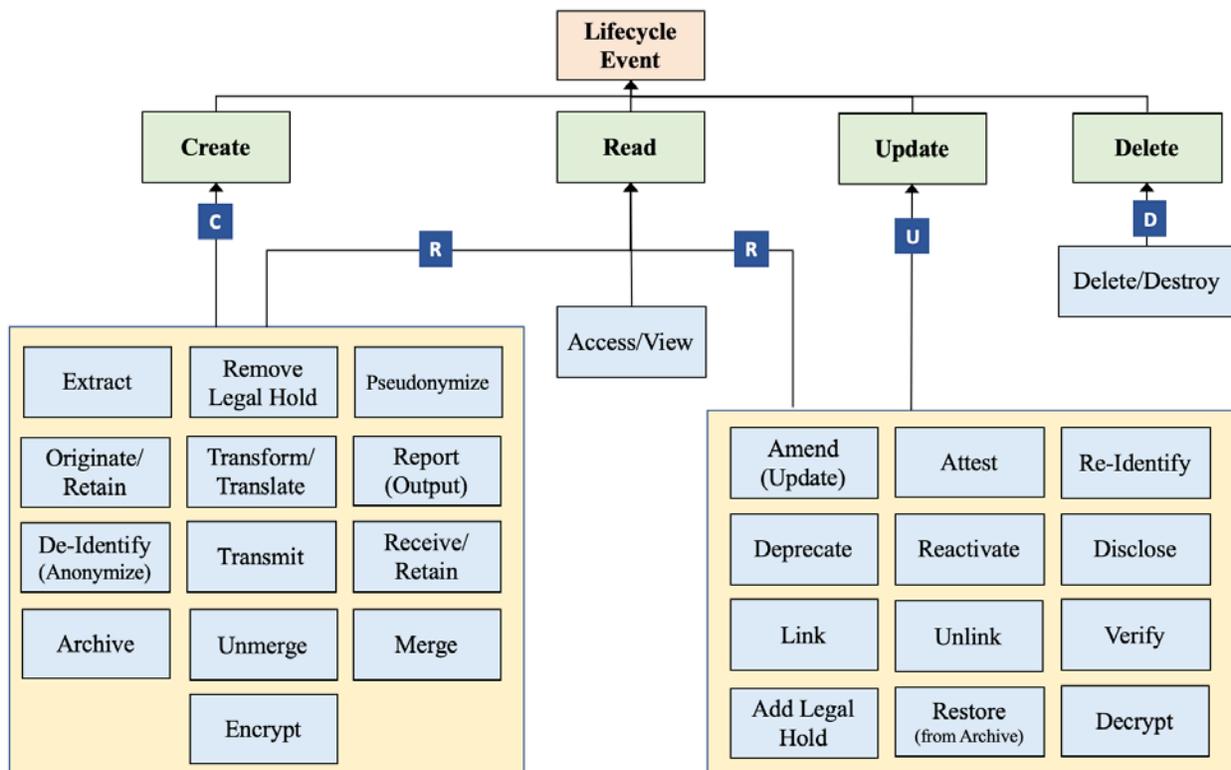


Figure 25: Relationships of Lifecycle Events to Create, Read, Update, Delete, and Execute [ISO 21089]

The LCEs discussed in this section are adapted from HL7 EHR LCE Definitions and Models [HL7 EHR LCE] and ISO/HL7 10781:2015 [ISO/HL7 10781]. This models are consistent with the healthcare LCE information model specified by [ISO/TS 21089]⁵ and the provenance information model specified by [W3C Prov DM] and [W3C Prov Ontology].

⁵ Except for Disclose, Merge, Remove Legal Hold, Unlink and Unmerge. See LCE diagrams below for details.

8.10.1 Access or View

This event captures the provenance when an Agent obtains the information content of an Entity.

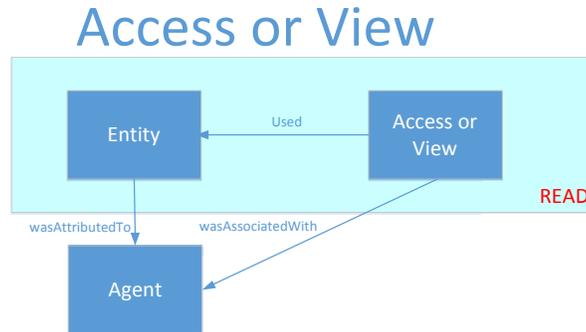


Figure 26: Access or View Provenance Event

8.10.2 Add Legal Hold

This event captures the provenance when an Agent places a tag or otherwise indicates special access management and suspension of destruction of an Entity when it is deemed relevant to a law suit, are reasonably anticipated to be relevant, or are consistent with organization policy under the legal doctrine of “duty to preserve.”

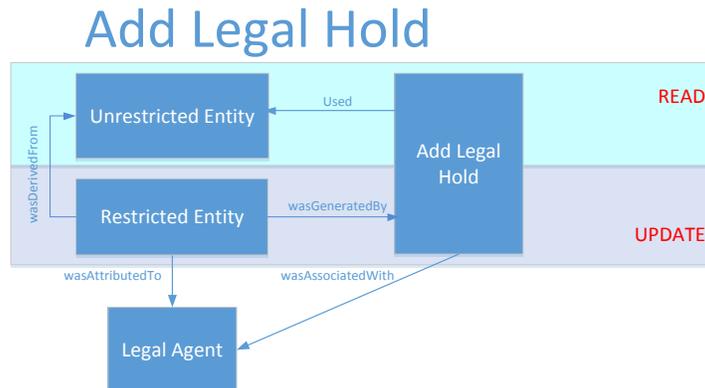


Figure 27: Add Legal Hold Provenance Event

8.10.3 Amend (Update)

This event captures the provenance when an Agent makes changes to the information content of an Entity. For the purposes this model, amend and update are considered synonymous.

Amend/Update

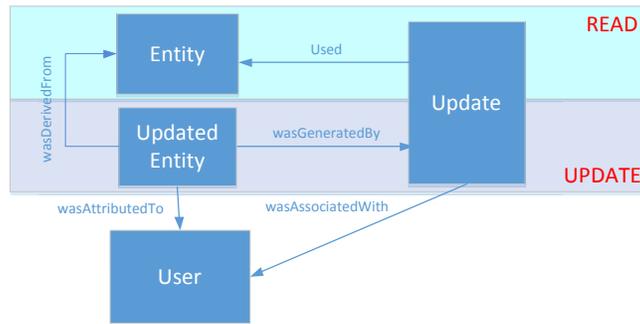


Figure 28: Amend/Update Provenance Event

8.10.4 Archive

This event captures the provenance when an Agent moves the contents of an Entity to long-term storage.

Archive

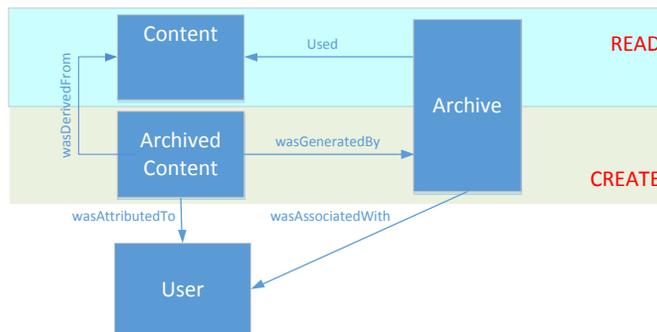


Figure 29: Archive Provenance Event

8.10.5 Attest

This event captures the provenance when an Agent performs a formal validation on the information content of an Entity.

Attest

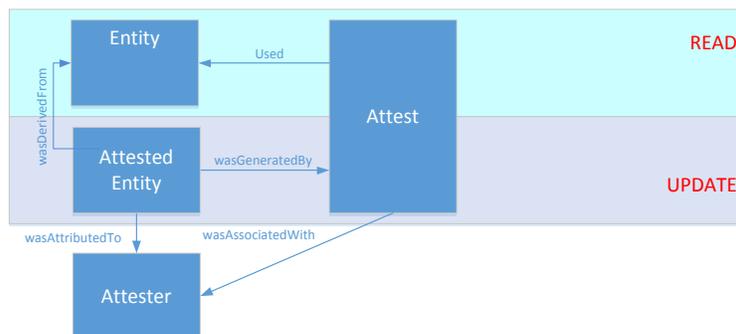


Figure 30: Attest Provenance Event

8.10.6 Encrypt

This event captures the provenance when an Agent performs an Activity that renders the information content of an Entity unreadable by algorithmically transforming plain text into ciphertext.

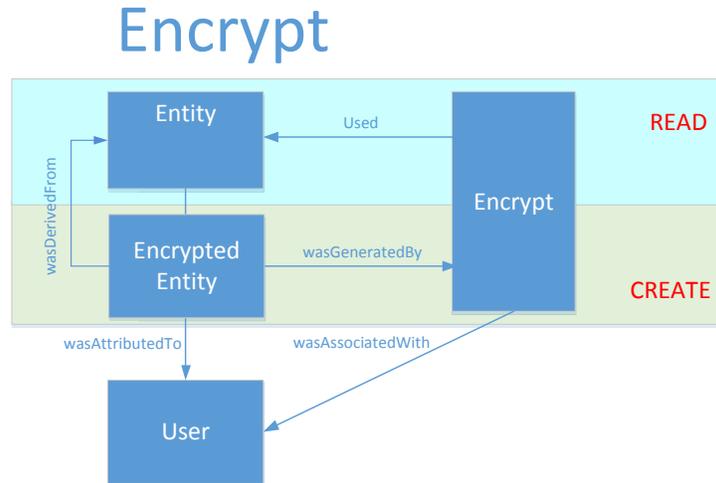


Figure 31: Encrypt Provenance Event

8.10.7 Decrypt

This event captures the provenance when an Agent performs an Activity which renders information content of an Entity readable by algorithmically transforming ciphertext into plaintext.

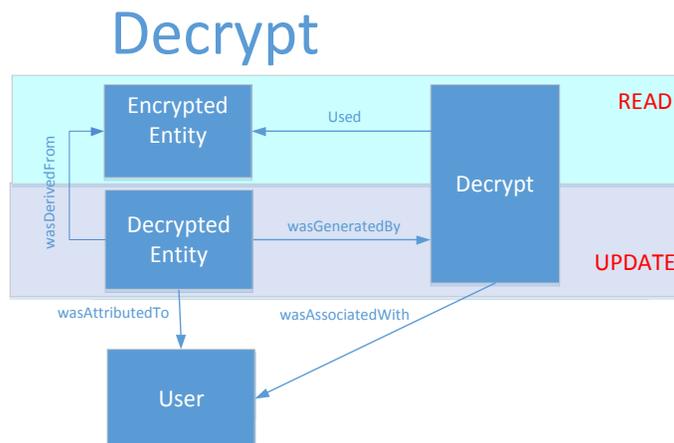


Figure 32: Decrypt Provenance Event

8.10.8 De-Identify (Anonymize)

This event captures the provenance when an Agent performs the process of reducing the association between an individual (patient) and the set of identifying information in the content of an Entity in a relatively unreversible way. For the purposes of this model, de-identify and anonymize are considered synonymous.

De-identify

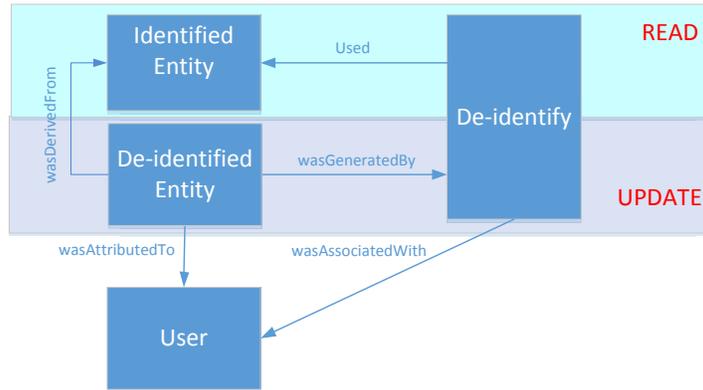


Figure 33: De-Identify (Anonymize) Provenance Event

8.10.9 Deprecate

This event captures the provenance when an Agent designates the information content of an Entity as obsolete, erroneous or untrustworthy in order to warn against its use in the future.

Deprecate

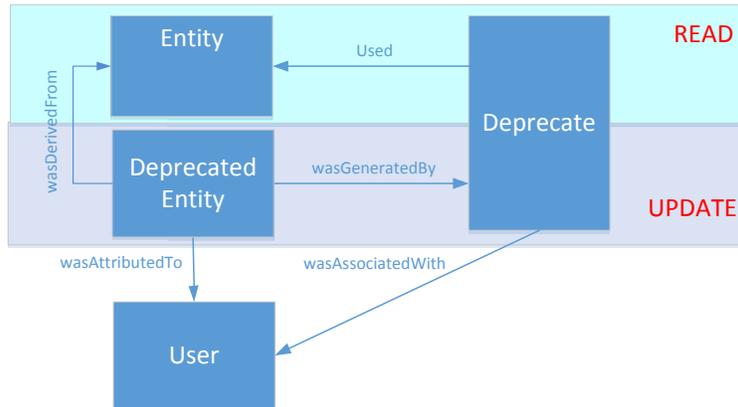


Figure 34: Deprecate Provenance Event

8.10.10 Destroy or Delete

This event captures the provenance when an Agent either permanently erases the information content of an Entity (destroy), or just makes the data inaccessible by removing the information about the Entity from memory or storage (delete).

Destroy or Delete

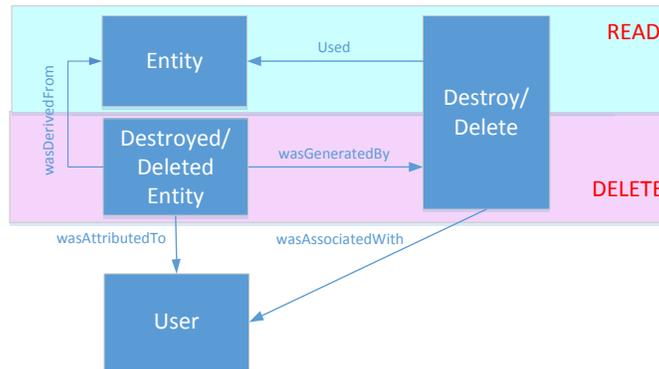


Figure 35: Destroy or Delete Provenance Event

8.10.11 Disclose

This event captures the provenance when an Agent releases, transfers, provisions access to, or divulges in any other manner, the information content of an Entity from an individual’s health record to third parties within or outside the healthcare organization.

Disclose

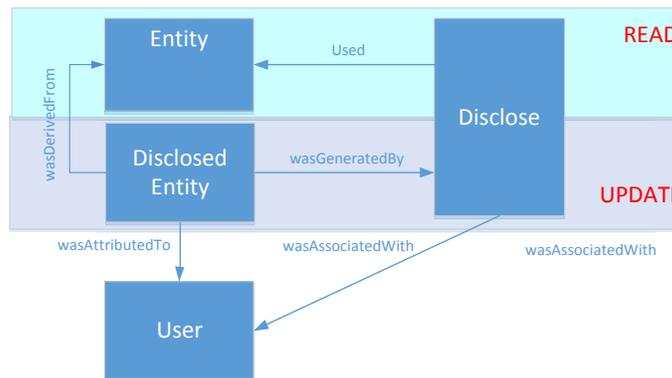


Figure 36: Disclose Provenance Event

8.10.12 Extract

This event captures the provenance when an Agent pulls a subset of the information content of an Entity, or a subset of the Entities included in a Collection based on an explicit criteria.

Extract

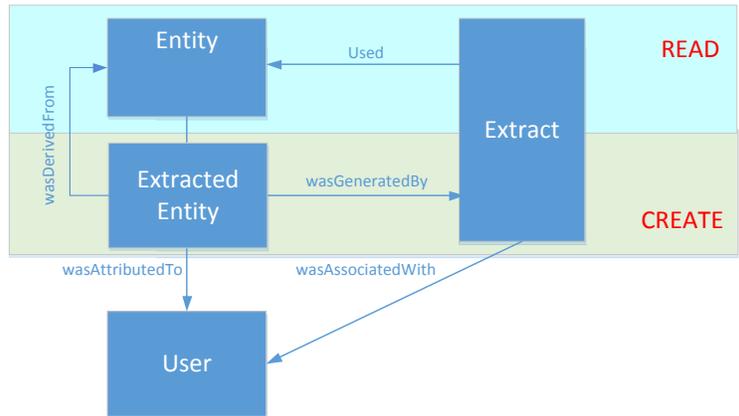


Figure 37: Extract Provenance Event

8.10.13 Link

This event captures the provenance when an Agent performs an Activity that connects two or more separate Entities so that access or use of one record entry implies equal access to, and ability to use, the connected record entries.

Link (Chaining)

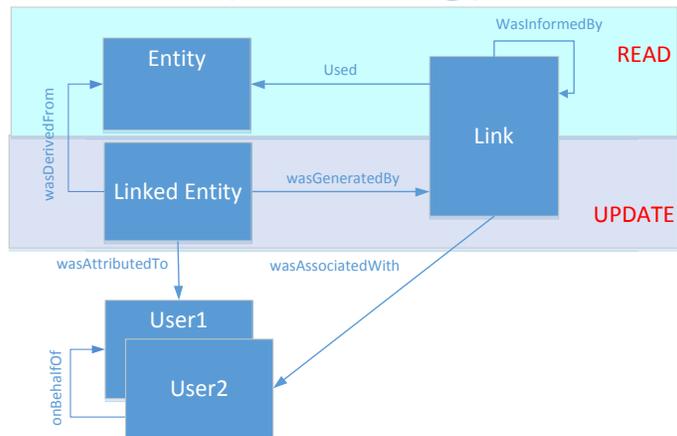


Figure 38: Link Provenance Event

8.10.14 Merge

This event captures the provenance when an Agent performs an Activity that combines the content of two or more Entities to generate one Entity.

Merge

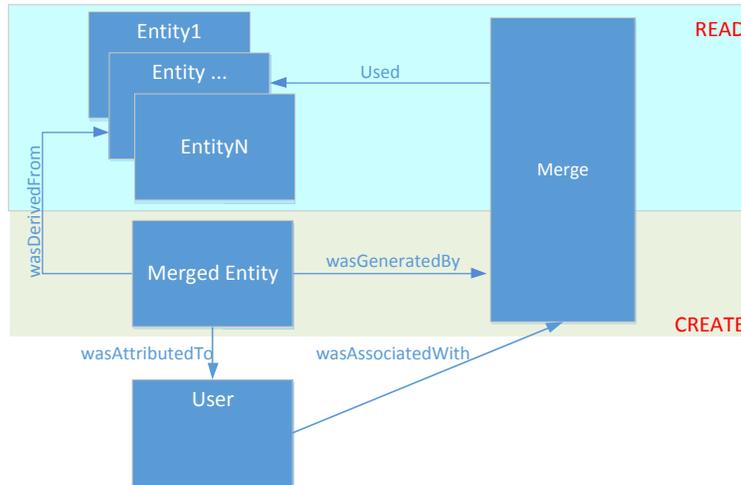


Figure 39: Merge Provenance Event

8.10.15 Originate

This event captures the provenance when an Agent initiates the entry of the information content of an Entity (originate) and enters that Entity into permanent storage (retain).

Originate/Retain (local)

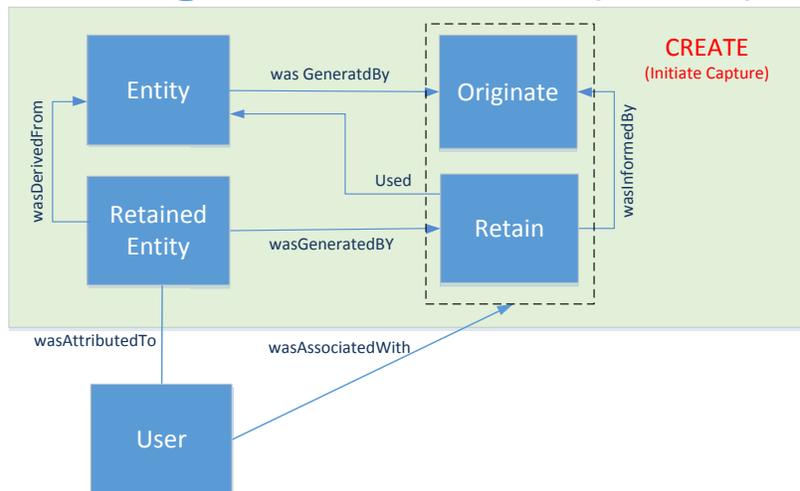


Figure 40: Originate Provenance Event

8.10.16 Pseudonymize

This event captures the provenance when an Agent performs potentially reversible de-identification by altering some of the identifying information content of an Entity that could link the Entity to an individual (e.g., patient).

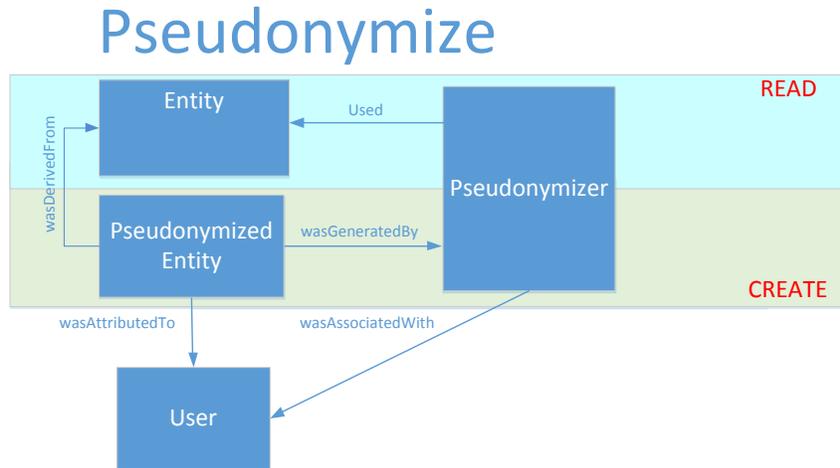


Figure 41: Pseudonymize Provenance Event

8.10.17 Re-Activate

This event captures the provenance when an Agent recreates previously deleted or deprecated Entities and restores them to full active status.

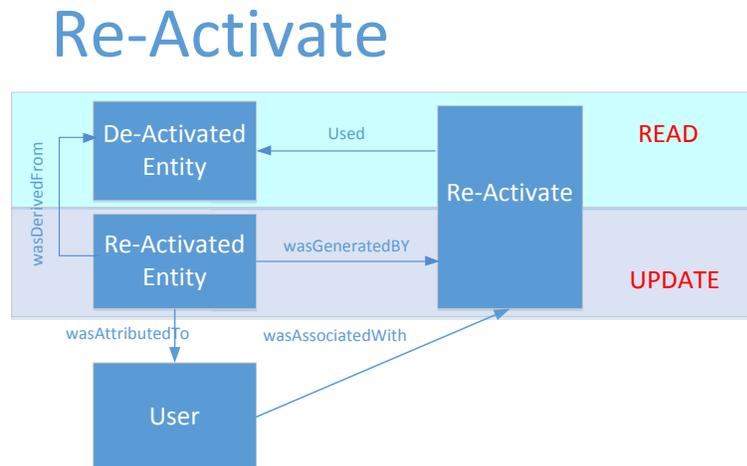


Figure 42: Re-Activate Provenance Event

8.10.18 Receive

This event captures the provenance when an Agent acquires data that exist elsewhere as potential content for generating an Entity (receive) and enters that Entity into permanent storage (retain).

Receive/Retain

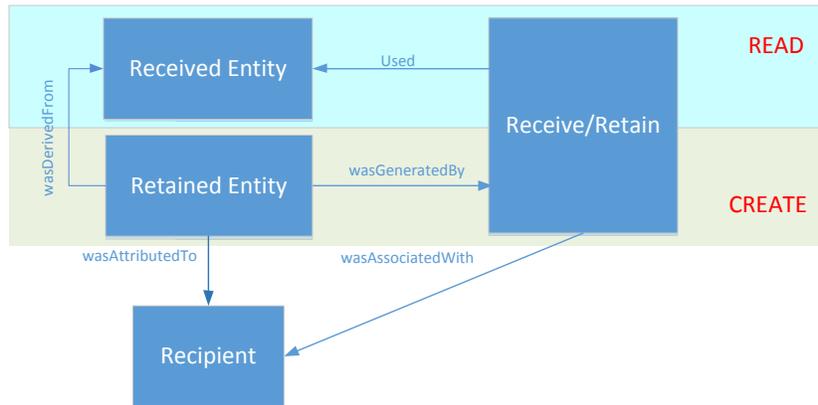


Figure 43: Receive Provenance Event

8.10.19 Re-Identify

This event captures the provenance when an Agent restores the identifying information content of an Entity, usually from a previously pseudonymized record, resulting in restoring the linkability of the Entity to and individual (e.g. patient).

Re-Identify

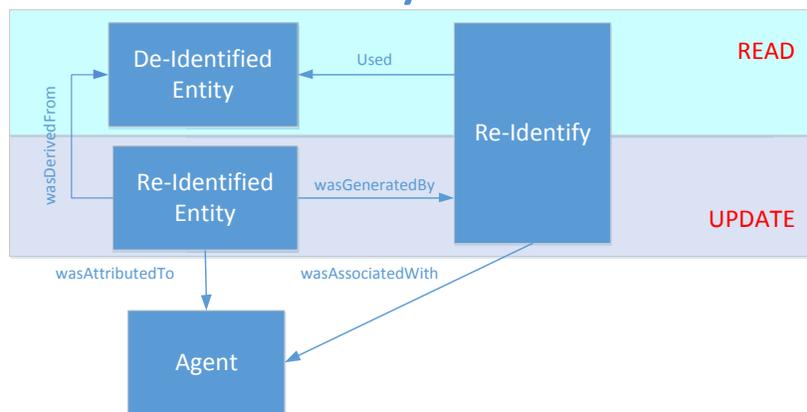


Figure 44: Re-Identify Provenance Event

8.10.20 Remove Legal Hold

This event captures the provenance when an Agent removes a tag or other cues for special access management and suspension of destruction from an Entity previously deemed relevant to a law suit, are reasonably anticipated to be relevant, or are consistent with organization policy under the legal doctrine of “duty to preserve.”

Remove Legal Hold

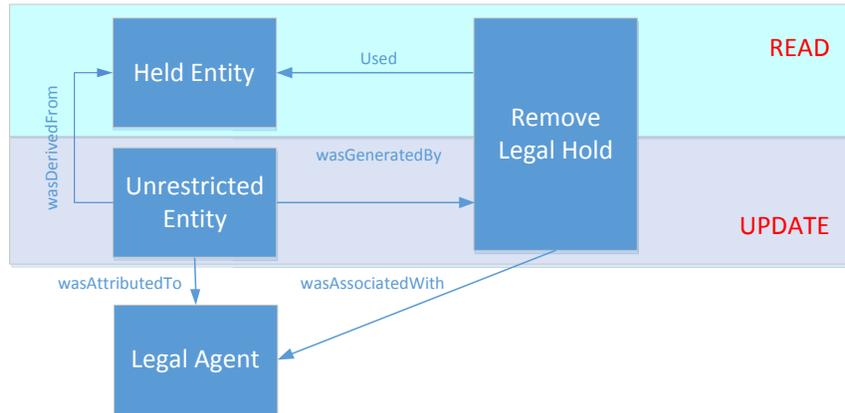


Figure 45: Remove Legal Hold Provenance Event

8.10.21 Report (Output)

This event captures the provenance when an Agent uses the information content of an Entity or a number of Entities to generate a new Entity in the form expected by a recipient. For the purposes of this model, report and output are considered synonymous.

Output/Report

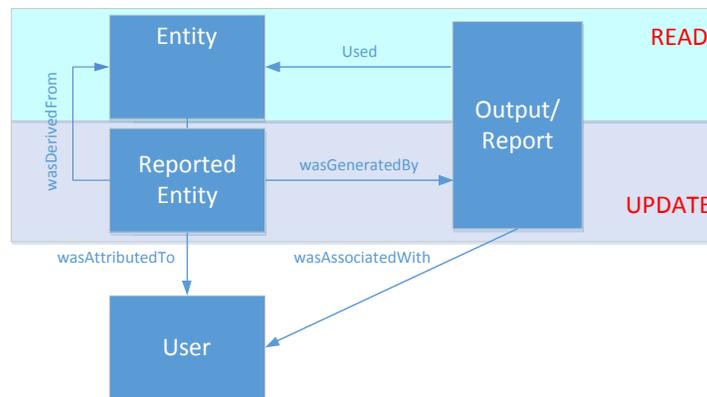


Figure 46: Output (Report) Provenance Event

8.10.22 Restore

This event captures the provenance when an Agent generates an Entity with active status based on a previously archived Entity.

Restore

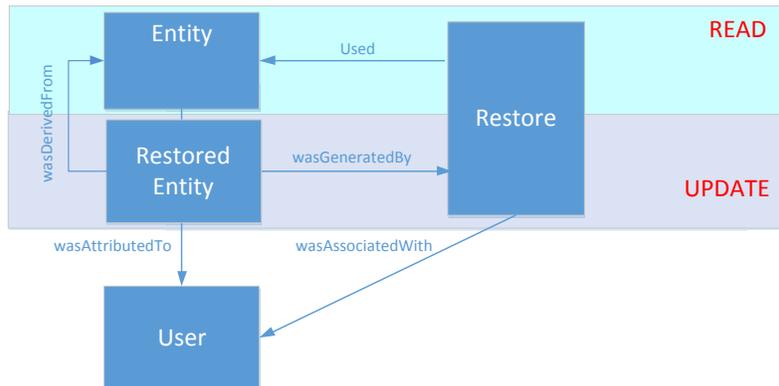


Figure 47: Restore Provenance Event

8.10.23 Translate or Transform

This event captures the provenance when an Agent makes changes to the content of an Entity in a way that semantics of the content stay the same but the form changes, e.g. by changing the form (transform), language, or coding system (translate) used to represent data.

Translate/Transform

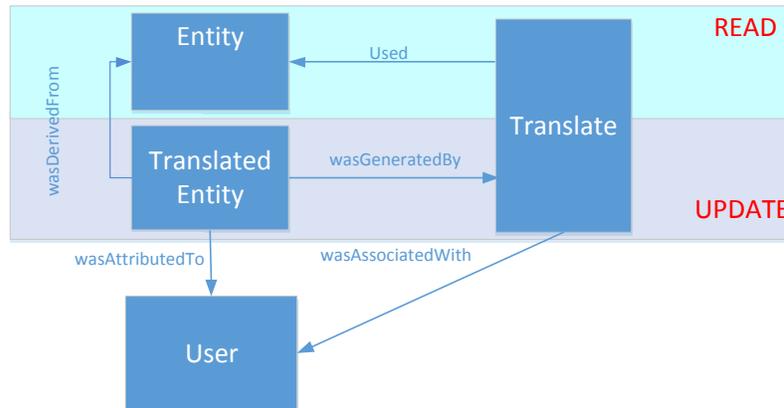


Figure 48: Translate or Transform Provenance Event

8.10.24 Transmit

This event captures the provenance when an Agent sends the information content of an Entity from one system to another.

Transmit

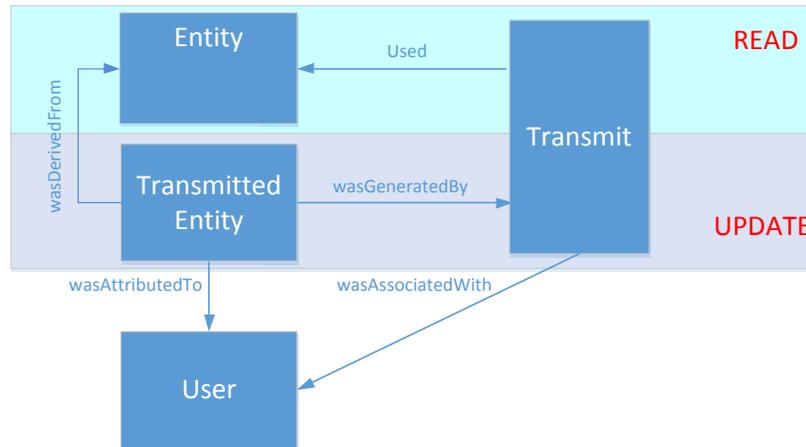


Figure 49: Transmit Provenance Event

8.10.25 Unlink

This event captures the provenance when an Agent performs an Activity which removes the linking between Entities previously established by a linking event.

Unlink

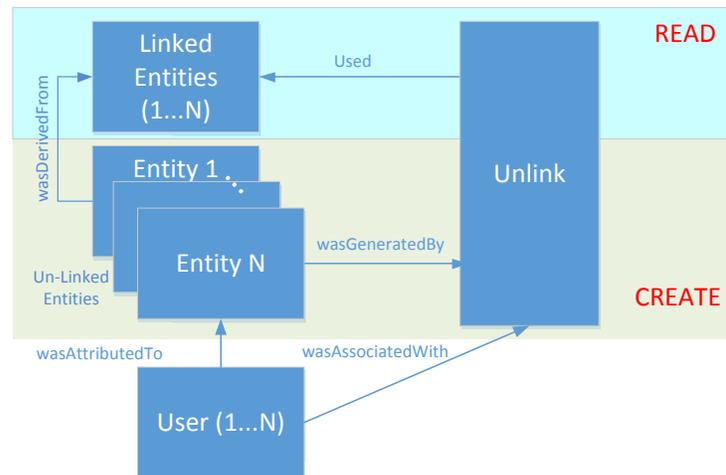


Figure 50: Unlink Provenance Event

8.10.26 Un-Merge

This event captures the provenance when an Agent performs an Activity which reverses a previous merge operation by generating a number of separate Entities based on a single Entity created by a previous merge operation.

Unmerge

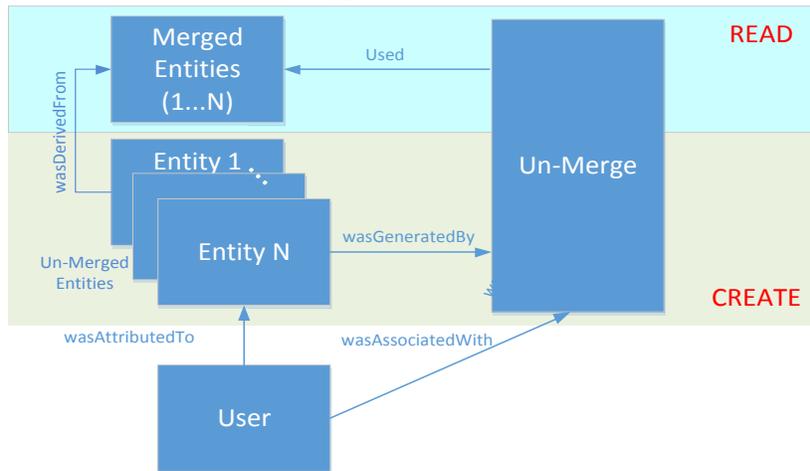


Figure 51: Un-Merge Provenance Event

8.10.27 Verify

This event captures the provenance when an Agent evaluates the compliance of the information content of an Entity data objects with regulations, requirements, specifications, or other internally imposed conditions based on organizational policy.

Verify

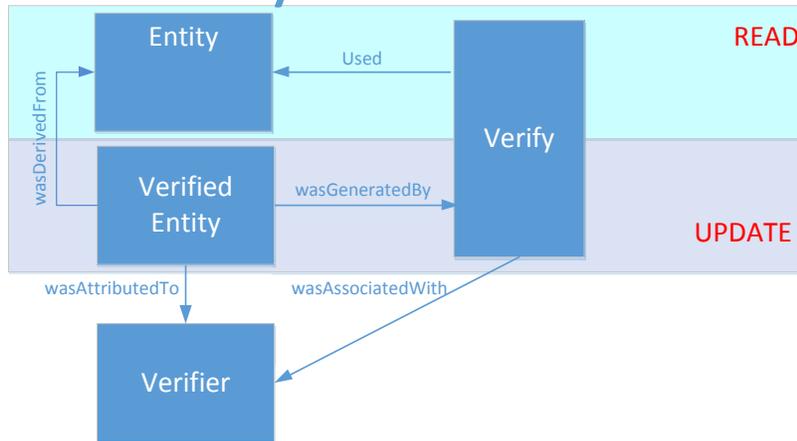


Figure 52: Verify Provenance Event

9 Requirements

Table 16 summarizes the requirements for a federated provenance service in the healthcare domain. This requirements are partly based on the W3C Requirements for Provenance on the Web [W3C Prov Req]. This table contains both functional requirements, which are related to the technical functionality of the system, and non-functional requirements that capture the general requirements applicable to the system in its entirety. Non-functional requirements are organized into the the following groups:

- *Content*: Requirements prescribing the type and form of information to be represented in a provenance record.
- *Management*: Requirements applicable to collecting provenance information and making it available and accessible.
- *Administration*: Requirements pertaining to the administration of the federation and its members.

Table 16: Healthcare Provenance Requirements (FPA: Federated Provenance Authority; PS: Provenance Store)

ID	Requirement Title/Text	Assigned To	Further Source of Guidance
PS-1	Functional Requirements Functional requirements are those		
1	Provenance Store SHALL accept direct push of provenance information by a client from a participating domain acting as an Agent.	PS	
2	Provenance Store SHALL accept redirect push of provenance information by a client from a participating domain acting as an Agent and deliver the provenance information to a client from a participating domain acting as the Recipient.	PS	
3	Provenance Store SHALL accept queries from Recipients requesting specific provenance information and deliver the matching provenance information to the requesting Recipient.	PS	
4	Provenance Store SHALL accept queries from Recipients requesting collections of provenance information specified by patterns or parameters and deliver the matching provenance information to the requesting Recipient.	PS	
5	Provenance Store SHALL accept subscription requests from Recipients requesting recurring delivery of provenance information and deliver the matching provenance information to the requesting Recipient.	PS	
6	A client application acting as a Recipient SHALL accept provenance information sent by an Agent or by the Provenance Store.	PS, Recipient	
7	Provenance Store SHALL accept requests for reports from Recipients via a Provenance Analysis Service interface and deliver the results to the requesting Recipient.	PS, Analysis Service	
8	Provenance Store SHALL accept requests for subscription to receive recurring reports from Recipients via a Provenance Analysis Service interface and deliver the results to the requesting Recipient.	PS, Analysis Service	

ID	Requirement Title/Text	Assigned To	Further Source of Guidance
9	A client application acting as a Recipient SHALL accept analysis reports sent by the Analysis Service of the Provenance Store.	Recipient	
10	Provenance Store SHALL accept requests for subscription to receive recurring notifications from Recipients via a Provenance Analysis Service interface and deliver the notifications to the requesting Recipient.	PS, Notification Service	
11	A client application acting as a Recipient SHALL accept notifications sent by the Analysis Service of the Provenance Store when the Recipient has previously requested for the notifications.	Recipient	
PS-2 Non-Functional Requirements			
Content			
1	Provenance information SHALL be recorded conformant with the data model presented in this document.	PS, Agent, Recipient	W3C Provenance Specifications
2	Provenance information SHALL be recorded in a machine-readable form.	PS, Agent, Recipient	FPA policy
3	Provenance information SHALL be serializable into standard formats with minimum information loss.	PS, Agent, Recipient	W3C Provenance Specifications, HL7 FHIR, FPA policy
4	Provenance data model classes and attributes SHOULD be linked, when relevant, to HL7 healthcare semantics, data models and formats.	PS, Agent, Recipient	HL7 FHIR, HL7 CDA, FPA policy
5	Provenance information SHOULD be captured and associated with each update to an Entity based on policy.	PS, Agent	FPA policy
6	Provenance information SHOULD make it possible to derive the chronological sequence of Activities.	PS	FPA policy
7	Entities, Activities and Agents SHALL be uniquely identifiable within a domain and should have persistent identifiers.	PS, Agent, Recipient	
8	Released Entities SHOULD include a value for the Contact attribute.	PS, Agent	
9	Activities and Entities SHOULD have a value for the Description attribute representing a short description or link to a description.	PS, Agent, Recipient	
Management			
10	Query: Provenance Store SHALL define query formulation and search parameters enabling search queries to retrieve provenance information.	PS	HL7 FHIR
11	Access Control: Provenance information SHALL be subject to dissemination control, access, use and licensing controls as established by Federation policy. Provenance Agents MAY withhold provenance information as required to meet privacy protection requirements per applicable law.	PS	FPA policy, HL7 FHIR

ID	Requirement Title/Text	Assigned To	Further Source of Guidance
12	<p>Granularity: Federation policy SHALL stipulate the minimum level of granularity for recording provenance information by considering trade-offs of scale and the requirements of the use-cases.</p> <p>The scale of provenance information is a major concern, as the size of the provenance records may by far exceed the scale of the artifacts themselves. Despite the presence of large amounts of provenance, efficient access to provenance records must be possible.</p>	PS	FPA policy, HL7 FHIR
Administration			
13	<p>Level of Trust: Provenance Federation Authority SHALL verify the trust level of its member.</p> <p>Trust level can be established based on past reliability ratings, presented credentials, or third-party attestations.</p>	FPA	FPA policy, NIST SP-800-63-r3
14	<p>Accountability: Provenance Federation Authority SHALL address policies to assure member accountability.</p>	FPA	FPA policy

Appendix A – Life-Cycle Events Definitions

This Appendix provides the definitions of the life-cycle events. A summary of all the LCEs and their source is given in **Table 17**.

Table 17: Complete List of Healthcare Lifecycle Events (LCEs)

LCE	LCE Description
Originate/Retain	<p>(a) Initiate capture of potential record content, and (b) incorporate that content into the storage considered permanent part of the health record. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent performs two Activities: the Agent initiates the entry of data as potential content for an EHR record (originate) and enters that data into storage considered permanent (retain). [HL7 EHR LCE]</p>
Amend (Update)	<p>Agent makes any change to record entry content currently residing in storage considered permanent (persistent). [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent makes any changes to the content of data currently residing in storage considered permanent. For the purposes of Amend (Update) Lifecycle Event, amend and update are considered synonymous. [HL7 EHR LCE]</p>
Transform/Translate	<p>Agent causes system to change the form, language or code system used to represent record entry content. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent makes any changes to the form (transform), language, or coding system (translate) used to represent data currently residing in “permanent” storage. [HL7 EHR LCE]</p>
Attest	<p>Agent causes system to capture the Agent’s digital signature (or equivalent indication) during formal validation of record entry content. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent performs a formal validation on the contents of data objects. [HL7 EHR LCE]</p>
Access/View	<p>Agent causes system to obtain and open a record entry for inspection or review. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent is obtaining data from one or more record entries. [HL7 EHR LCE]</p>
Report (Output)	<p>Agent causes system to produce and deliver record entry content in a particular form and manner. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent produces and delivers the content of a record in the form expected by the recipient. Note: For the purposes of the Output (Report) Life Cycle Event, report and output are considered synonymous. [HL7 EHR LCE]</p>
Disclose	<p>Agent causes system to release, transfer, provision access to, or otherwise divulge record entry content. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent releases, transfers, provisions access to, or divulges in any other manner, information to third parties within or outside the healthcare provider organization from an individual’s health record, with or without the consent of the individual to whom the record pertains. [HL7 EHR LCE]</p>
Transmit	<p>Agent causes system to send record entry content from one (EHR/PHR/other) system to another. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent sends EHR content from one system (EHR/PHR/other) to another. [HL7 EHR LCE]</p>

Receive/Retain	<p>Agent causes system to: a) initiate capture of data content from elsewhere, and b) incorporate that content into the storage considered a permanent part of the health record. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent both acquires data that exist elsewhere as potential content for an EHR record (receive) and enters that data into storage considered permanent (retain). (See “1. Originate/Retain Lifecycle Event” for the definition of “Retain (v).”) [HL7 EHR LCE]</p>
De-Identify (Anonymize)	<p>Agent causes system to scrub record entry content to reduce the association between a set of identifying data and the data subject in a way that may or may not be reversible. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent performs the process of reducing the association between a set of identifying data and the data subject in a way that is not reversible. Note: For the purposes of the De-identify (Anonymize) Life Cycle Event, de-identify and anonymize are considered synonymous. [HL7 EHR LCE]</p>
Pseudonymize	<p>Agent causes system to remove record entry content to reduce the association between a set of identifying data and the data subject in a way that may be reversible. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent performs de-identification which may be reversible. [HL7 EHR LCE]</p>
Re-Identify	<p>Agent causes system to restore information to data that allows identification of information source and/or information subject. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent restores individual identity in Record Entry content, usually from a previously pseudonymized record, that allows the identification of the source of the information or the information subject. [HL7 EHR LCE]</p>
Extract	<p>Agent causes system to selectively pull out a subset of record entry content, based on explicit criteria. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent pulls out a set of health data or record content from a larger volume of data using explicit criteria. [HL7 EHR LCE]</p>
Archive	<p>Agent causes system to create and move archive artifacts containing record entry content, typically to long-term offline storage. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent moves the contents of a data object to long-term storage. [HL7 EHR LCE]</p>
Restore (from Archives)	<p>Agent causes system to recreate record entries and their content from a previous created archive artifact. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent recreates Record Entries and their content from a previously created archive artifact. [HL7 EHR LCE]</p>
Destroy/Delete	<p>Agent causes system to permanently erase record entry content from the system. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent either permanently erases data from the system (destroy) or just makes the data inaccessible to the application by removing the information about an object from memory or storage (delete). [HL7 EHR LCE]</p>
Deprecate	<p>Agent causes system to tag record entry(ies) as obsolete, erroneous or untrustworthy, to warn against its future use. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent designates data or record content as obsolete, erroneous or untrustworthy in order to warn against its use in the future. [HL7 EHR LCE]</p>

Re-activate	<p>Agent causes system to recreate or restore full status to record entries previously deleted or deprecated. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent recreates previously deleted or deprecated record entries and restores them to full active status. [HL7 EHR LCE]</p>
Merge	<p>Agent causes system to combine or join content from two or more record entries, resulting in a single logical record entry. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent performs an Activity which combines the content of two or more Record Entries, resulting in a single record entry. [HL7 EHR LCE]</p>
Unmerge	<p>Agent causes system to reverse a previous record entry merge operation, rendering them separate again. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent performs an Activity which reverses a previously executed merge operation. [HL7 EHR LCE]</p>
Link	<p>Agent causes system to connect related record entries. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent performs an Activity which connects two or more separate record entries so that access or use of one record entry means equal access to and ability to use all of the connected record entries. [HL7 EHR LCE]</p>
Unlink	<p>Agent causes system to disconnect two or more record entries previously connected, rendering them separate (disconnected) again. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent performs an Activity which undoes any linked record entries, rendering them separate again. [HL7 EHR LCE]</p>
Add Legal Hold	<p>Agent causes system to tag or otherwise indicate special access management and suspension of record entry deletion/destruction, if deemed relevant to a lawsuit or which are reasonably anticipated to be relevant or to fulfill organizational policy under the legal doctrine of “duty to preserve” [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent places a tag or otherwise indicates special access management and suspension of destruction for record entries deemed relevant to a law suit, are reasonably anticipated to be relevant, or are consistent with organization policy under the legal doctrine of “duty to preserve.” [HL7 EHR LCE]</p>
Remove Legal Hold	<p>Agent causes system to remove a tag or other cues for special access management had required to fulfill organizational policy under the legal doctrine of “duty to preserve”. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent removes a tag or other cues for special access management and suspension of destruction for record entries deemed relevant to a law suit, are reasonably anticipated to be relevant, or are consistent with organization policy under the legal doctrine of “duty to preserve.” [HL7 EHR LCE]</p>
Verify	<p>Agent causes system to confirm compliance of data or data objects with regulations, requirements, specifications, or other imposed conditions based on organizational policy. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent evaluates the compliance of data objects with regulations, requirements, specifications, or other internally imposed conditions based on organizational policy. [HL7 EHR LCE]</p>

Encrypt	<p>Agent causes system to encode record entry content in a cipher. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent performs an Activity which renders information unreadable by algorithmically transforming plain text into ciphertext. [HL7 EHR LCE]</p>
Decrypt	<p>Agent causes system to decode record entry content from a cipher. [ISO 21089]</p> <p>As part of trusted record management, this is the record lifecycle event describing when an Agent performs an Activity which renders information readable by algorithmically transforming ciphertext into plaintext. (ENCRYPT concept in HL7 ActCode code system, HL7 v3 ObligationPolicy value set, modified) [HL7 EHR LCE]</p>

Appendix B – Provenance Model Listing

Table 18: Complete List of Provenance Model Classes

Provenance Item	Designation	Provenance Event Description
Activity	Class	An Activity is something that occurs over a period of time and acts upon or with Entities; it may include consuming, processing, transforming, modifying, relocating, using, or generating Entities.
ActivityInfluence	Class	ActivityInfluence is the capacity of an Activity to have an effect on the character, development, or behavior of another by means of generation, invalidation, communication, or other.
Agent	Class	An Agent is something that bears some form of responsibility for an Activity taking place, for the existence of an Entity, or for another Agent's Activity.
AgentInfluence	Class	AgentInfluence is the capacity of an Agent to have an effect on the character, development, or behavior of another by means of attribution, association, delegation, or other.
Association	Class	An Activity association is an assignment of responsibility to an Agent for an Activity, indicating that the Agent had a role in the Activity. It further allows a plan to be specified, which is the plan intended by the Agent to achieve some goals in the context of this Activity.
Attribution	Class	Attribution is the ascribing of an Entity to an Agent. When an Entity e is attributed to Agent ag, Entity e was generated by some unspecified Activity that in turn was associated to Agent ag. Thus, this relation is useful when the Activity is not known, or irrelevant.
Bundle	Class	A bundle is a named set of provenance descriptions, and is itself an Entity, so allowing provenance of provenance to be expressed.
Collection	Class	A collection is an Entity that provides a structure to some constituents, which are themselves Entities. These constituents are said to be member of the collections.
Communication	Class	Communication is the exchange of an Entity by two Activities, one Activity using the Entity generated by the other.
Delegation	Class	Delegation is the assignment of authority and responsibility to an Agent (by itself or by another Agent) to carry out a specific Activity as a delegate or representative, while the Agent it acts on behalf of retains some responsibility for the outcome of the delegated work. For example, a student acted on behalf of his supervisor, who acted on behalf of the department chair, who acted on behalf of the university; all those Agents are responsible in some way for the Activity that took place, but we do not say explicitly who bears responsibility and to what degree.
Derivation	Class	A derivation is a transformation of an Entity into another, an update of an Entity resulting in a new one, or the construction of a new Entity based on a pre-existing Entity.
EmptyCollection	Class	An empty collection is a collection without members.
End	Class	End is when an Activity is deemed to have been ended by an Entity, known as trigger. The Activity no longer exists after its end. Any usage, generation, or invalidation involving an Activity precedes the Activity's end. An end may refer to a trigger Entity that terminated the Activity, or to an Activity, known as ender that generated the trigger.
Entity	Class	An Entity is a physical, digital, conceptual, or other kind of thing with some fixed aspects; Entities may be real or imaginary.
EntityInfluence	Class	EntityInfluence is the capacity of an Entity to have an effect on the character, development, or behavior of another by means of usage, start, end, derivation, or other.

Generation	Class	Generation is the completion of production of a new Entity by an Activity. This Entity did not exist before generation and becomes available for usage after this generation.
Influence	Class	Influence is the capacity of an Entity, Activity, or Agent to have an effect on the character, development, or behavior of another by means of usage, start, end, generation, invalidation, communication, derivation, attribution, association, or delegation.
InstantaneousEvent	Class	The PROV data model is implicitly based on a notion of instantaneous events (or just events), that mark transitions in the world. Events include generation, usage, or invalidation of Entities, as well as starting or ending of Activities. This notion of event is not first-class in the data model, but it is useful for explaining its other concepts and its semantics.
Invalidation	Class	Invalidation is the start of the destruction, cessation, or expiry of an existing Entity by an Activity. The Entity is no longer available for use (or further invalidation) after invalidation. Any generation or usage of an Entity precedes its invalidation.
Location	Class	A location can be an identifiable geographic place (ISO 19112), but it can also be a non-geographic place such as a directory, row, or column. As such, there are numerous ways in which location can be expressed, such as by a coordinate, address, landmark, and so forth.
Organization	Class	An organization is a social or legal institution such as a company, society, etc.
Person	Class	Person Agents are people.
Plan	Class	A plan is an Entity that represents a set of actions or steps intended by one or more Agents to achieve some goals.
PrimarySource	Class	A primary source for a topic refers to something produced by some Agent with direct experience and knowledge about the topic, at the time of the topic's study, without benefit from hindsight. Because of the directness of primary sources, they 'speak for themselves' in ways that cannot be captured through the filter of secondary sources. As such, it is important for secondary sources to reference those primary sources from which they were derived, so that their reliability can be investigated. A primary source relation is a particular case of derivation of secondary materials from their primary sources. It is recognized that the determination of primary sources can be up to interpretation and should be done according to conventions accepted within the application's domain.
Quotation	Class	A quotation is the repeat of (some or all of) an Entity, such as text or image, by someone who may or may not be its original author. Quotation is a particular case of derivation.
Revision	Class	A revision is a derivation for which the resulting Entity is a revised version of some original. The implication here is that the resulting Entity contains substantial content from the original. Revision is a particular case of derivation.
Role	Class	A role is the function of an Entity or Agent with respect to an Activity, in the context of a usage, generation, invalidation, association, start, and end.
SoftwareAgent	Class	A software Agent is running software.
Start	Class	Start is when an Activity is deemed to have been started by an Entity, known as trigger. The Activity did not exist before its start. Any usage, generation, or invalidation involving an Activity follows the Activity's start. A start may refer to a trigger Entity that set off the Activity, or to an Activity, known as starter, that generated the trigger.
Usage	Class	Usage is the beginning of utilizing an Entity by an Activity. Before usage, the Activity had not begun to utilize this Entity and could not have been affected by the Entity.

Table 19: Complete List of Provenance Model Properties

Provenance Item	Designation	Provenance Event Description
actedOnBehalfOf	Property	Delegation is the assignment of authority and responsibility to an Agent (by itself or by another Agent) to carry out a specific Activity as a delegate or representative, while the Agent it acts on behalf of retains some responsibility for the outcome of the delegated work. For example, a student acted on behalf of his supervisor, who acted on behalf of the department chair, who acted on behalf of the university; all those Agents are responsible in some way for the Activity that took place, but we do not say explicitly who bears responsibility and to what degree.
AlternateOf	Property	Two alternate Entities present aspects of the same thing. These aspects may be the same or different, and the alternate Entities may or may not overlap in time.
atLocation	Property	A location can be an identifiable geographic place (ISO 19112), but it can also be a non-geographic place such as a directory, row, or column. As such, there are numerous ways in which location can be expressed, such as by a coordinate, address, landmark, and so forth.
atTime	Property	The PROV data model is implicitly based on a notion of instantaneous events (or just events), that mark transitions in the world. Events include generation, usage, or invalidation of Entities, as well as starting or ending of Activities. This notion of event is not first-class in the data model, but it is useful for explaining its other concepts and its semantics.
endedAtTime	Property	End is when an Activity is deemed to have been ended by an Entity, known as trigger. The Activity no longer exists after its end. Any usage, generation, or invalidation involving an Activity precedes the Activity's end. An end may refer to a trigger Entity that terminated the Activity, or to an Activity, known as ender that generated the trigger.
Entity	Property	The prov:Entity property references a prov:Entity which influenced a resource. This property applies to a prov:EntityInfluence, which is given by a subproperty of prov:qualifiedInfluence from the influenced prov:Entity, prov:Activity or prov:Agent.
generated	Property	Generation is the completion of production of a new Entity by an Activity. This Entity did not exist before generation and becomes available for usage after this generation.
generatedAtTime	Property	Generation is the completion of production of a new Entity by an Activity. This Entity did not exist before generation and becomes available for usage after this generation.
hadActivity	Property	An Activity is something that occurs over a period of time and acts upon or with Entities; it may include consuming, processing, transforming, modifying, relocating, using, or generating Entities.
hadGeneration	Property	Generation is the completion of production of a new Entity by an Activity. This Entity did not exist before generation and becomes available for usage after this generation.
hadMember	Property	A collection is an Entity that provides a structure to some constituents, which are themselves Entities. These constituents are said to be member of the collections.
hadPlan	Property	A plan is an Entity that represents a set of actions or steps intended by one or more Agents to achieve some goals.

hadPrimarySource	Property	A primary source for a topic refers to something produced by some Agent with direct experience and knowledge about the topic, at the time of the topic's study, without benefit from hindsight. Because of the directness of primary sources, they 'speak for themselves' in ways that cannot be captured through the filter of secondary sources. As such, it is important for secondary sources to reference those primary sources from which they were derived, so that their reliability can be investigated. A primary source relation is a particular case of derivation of secondary materials from their primary sources. It is recognized that the determination of primary sources can be up to interpretation and should be done according to conventions accepted within the application's domain.
hadRole	Property	A role is the function of an Entity or Agent with respect to an Activity, in the context of a usage, generation, invalidation, association, start, and end.
hadUsage	Property	Usage is the beginning of utilizing an Entity by an Activity. Before usage, the Activity had not begun to utilize this Entity and could not have been affected by the Entity.
influenced	Property	Influence is the capacity of an Entity, Activity, or Agent to have an effect on the character, development, or behavior of another by means of usage, start, end, generation, invalidation, communication, derivation, attribution, association, or delegation.
influencer	Property	This property is used as part of the qualified influence pattern. Subclasses of prov:Influence use these subproperties to reference the resource (Entity, Agent, or Activity) whose influence is being qualified.
invalidated	Property	Invalidation is the start of the destruction, cessation, or expiry of an existing Entity by an Activity. The Entity is no longer available for use (or further invalidation) after invalidation. Any generation or usage of an Entity precedes its invalidation.
invalidatedAtTime	Property	Invalidation is the start of the destruction, cessation, or expiry of an existing Entity by an Activity. The Entity is no longer available for use (or further invalidation) after invalidation. Any generation or usage of an Entity precedes its invalidation.
qualifiedAssociation	Property	An Activity association is an assignment of responsibility to an Agent for an Activity, indicating that the Agent had a role in the Activity. It further allows a plan to be specified, which is the plan intended by the Agent to achieve some goals in the context of this Activity.
qualifiedAttribution	Property	Attribution is the ascribing of an Entity to an Agent. When an Entity e is attributed to Agent ag, Entity e was generated by some unspecified Activity that in turn was associated to Agent ag. Thus, this relation is useful when the Activity is not known, or irrelevant.
qualifiedCommunication	Property	Communication is the exchange of an Entity by two Activities, one Activity using the Entity generated by the other.
qualifiedDelegation	Property	Delegation is the assignment of authority and responsibility to an Agent (by itself or by another Agent) to carry out a specific Activity as a delegate or representative, while the Agent it acts on behalf of retains some responsibility for the outcome of the delegated work. For example, a student acted on behalf of his supervisor, who acted on behalf of the department chair, who acted on behalf of the university; all those Agents are responsible in some way for the Activity that took place, but we do not say explicitly who bears responsibility and to what degree.
qualifiedDerivation	Property	A derivation is a transformation of an Entity into another, an update of an Entity resulting in a new one, or the construction of a new Entity based on a pre-existing Entity.

qualifiedEnd	Property	End is when an Activity is deemed to have been ended by an Entity, known as trigger. The Activity no longer exists after its end. Any usage, generation, or invalidation involving an Activity precedes the Activity's end. An end may refer to a trigger Entity that terminated the Activity, or to an Activity, known as ender that generated the trigger.
qualifiedGeneration	Property	Generation is the completion of production of a new Entity by an Activity. This Entity did not exist before generation and becomes available for usage after this generation.
qualifiedInfluence	Property	Influence is the capacity of an Entity, Activity, or Agent to have an effect on the character, development, or behavior of another by means of usage, start, end, generation, invalidation, communication, derivation, attribution, association, or delegation.
qualifiedInvalidation	Property	Invalidation is the start of the destruction, cessation, or expiry of an existing Entity by an Activity. The Entity is no longer available for use (or further invalidation) after invalidation. Any generation or usage of an Entity precedes its invalidation.
qualifiedPrimarySource	Property	A primary source for a topic refers to something produced by some Agent with direct experience and knowledge about the topic, at the time of the topic's study, without benefit from hindsight. Because of the directness of primary sources, they 'speak for themselves' in ways that cannot be captured through the filter of secondary sources. As such, it is important for secondary sources to reference those primary sources from which they were derived, so that their reliability can be investigated. A primary source relation is a particular case of derivation of secondary materials from their primary sources. It is recognized that the determination of primary sources can be up to interpretation and should be done according to conventions accepted within the application's domain.
qualifiedQuotation	Property	A quotation is the repeat of (some or all of) an Entity, such as text or image, by someone who may or may not be its original author. Quotation is a particular case of derivation.
qualifiedRevision	Property	A revision is a derivation for which the resulting Entity is a revised version of some original. The implication here is that the resulting Entity contains substantial content from the original. Revision is a particular case of derivation.
qualifiedStart	Property	Start is when an Activity is deemed to have been started by an Entity, known as trigger. The Activity did not exist before its start. Any usage, generation, or invalidation involving an Activity follows the Activity's start. A start may refer to a trigger Entity that set off the Activity, or to an Activity, known as starter, that generated the trigger.
qualifiedUsage	Property	Usage is the beginning of utilizing an Entity by an Activity. Before usage, the Activity had not begun to utilize this Entity and could not have been affected by the Entity.
specializationOf	Property	An Entity that is a specialization of another shares all aspects of the latter, and additionally presents more specific aspects of the same thing as the latter. In particular, the lifetime of the Entity being specialized contains that of any specialization. Examples of aspects include a time period, an abstraction, and a context associated with the Entity.
startedAtTime	Property	Start is when an Activity is deemed to have been started by an Entity, known as trigger. The Activity did not exist before its start. Any usage, generation, or invalidation involving an Activity follows the Activity's start. A start may refer to a trigger Entity that set off the Activity, or to an Activity, known as starter, that generated the trigger.
used	Property	Usage is the beginning of utilizing an Entity by an Activity. Before usage, the Activity had not begun to utilize this Entity and could not have been affected by the Entity.

value	Property	Provides a value that is a direct representation of an Entity.
wasAssociatedWith	Property	An Activity association is an assignment of responsibility to an Agent for an Activity, indicating that the Agent had a role in the Activity. It further allows for a plan to be specified, which is the plan intended by the Agent to achieve some goals in the context of this Activity.
wasAttributedTo	Property	Attribution is the ascribing of an Entity to an Agent.
wasDerivedFrom	Property	A derivation is a transformation of an Entity into another, an update of an Entity resulting in a new one, or the construction of a new Entity based on a pre-existing Entity.
wasEndedBy	Property	End is when an Activity is deemed to have been ended by an Entity, known as trigger. The Activity no longer exists after its end. Any usage, generation, or invalidation involving an Activity precedes the Activity's end. An end may refer to a trigger Entity that terminated the Activity, or to an Activity, known as ender that generated the trigger.
wasGeneratedBy	Property	Generation is the completion of production of a new Entity by an Activity. This Entity did not exist before generation and becomes available for usage after this generation.
wasInfluencedBy	Property	Influence is the capacity of an Entity, Activity, or Agent to have an effect on the character, development, or behavior of another by means of usage, start, end, generation, invalidation, communication, derivation, attribution, association, or delegation.
wasInformedBy	Property	Communication is the exchange of an Entity by two Activities, one Activity using the Entity generated by the other.
wasInvalidatedBy	Property	Invalidation is the start of the destruction, cessation, or expiry of an existing Entity by an Activity. The Entity is no longer available for use (or further invalidation) after invalidation. Any generation or usage of an Entity precedes its invalidation.
wasQuotedFrom	Property	A quotation is the repeat of (some or all of) an Entity, such as text or image, by someone who may or may not be its original author. Quotation is a particular case of derivation.
wasRevisionOf	Property	A revision is a derivation for which the resulting Entity is a revised version of some original. The implication here is that the resulting Entity contains substantial content from the original. Revision is a particular case of derivation.
wasStartedBy	Property	Start is when an Activity is deemed to have been started by an Entity, known as trigger. The Activity did not exist before its start. Any usage, generation, or invalidation involving an Activity follows the Activity's start. A start may refer to a trigger Entity that set off the Activity, or to an Activity, known as starter, that generated the trigger.

Appendix C – Inverse Provenance Names

Table 20: Inverse Provenance Names

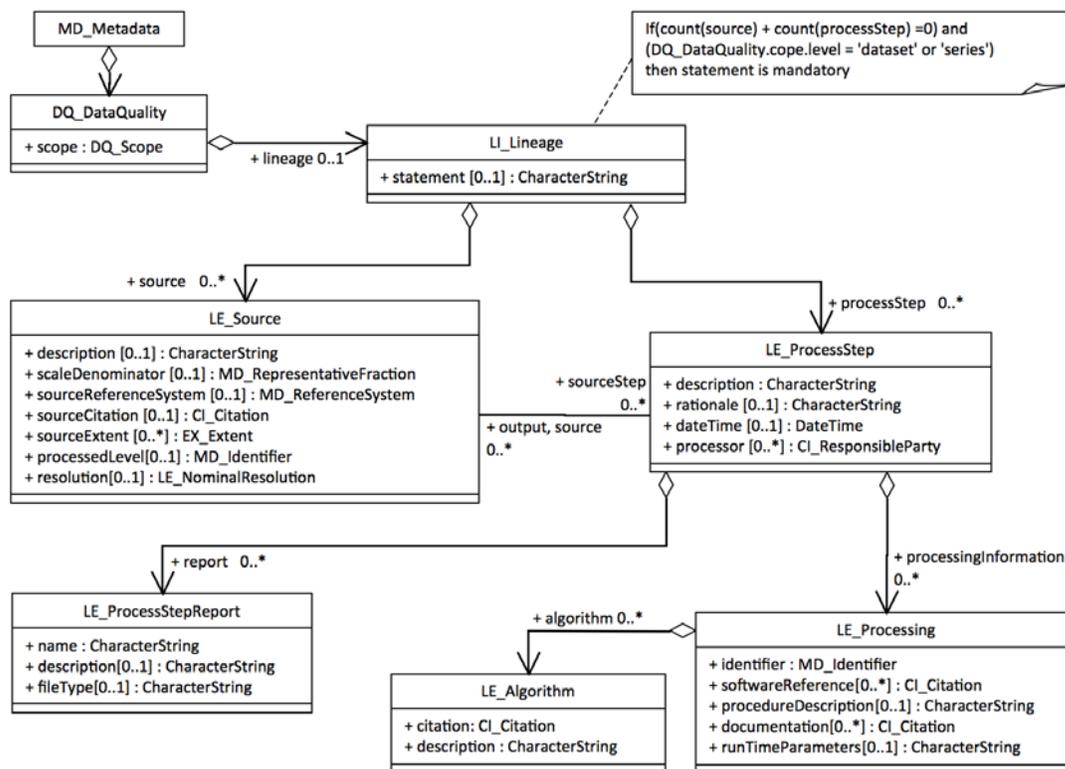
Provenance Event	Recommended inverse name
<u>actedOnBehalfOf</u>	hadDelegate
<u>Activity</u>	ActivityOfInfluence
<u>Agent</u>	AgentOfInfluence
<u>alternateOf</u>	alternateOf
<u>atLocation</u>	locationOf
<u>Entity</u>	EntityOfInfluence
<u>generated</u>	wasGeneratedBy
<u>hadActivity</u>	wasActivityOfInfluence
<u>hadGeneration</u>	generatedAsDerivation
<u>hadMember</u>	wasMemberOf
<u>hadPlan</u>	wasPlanOf
<u>hadPrimarySource</u>	wasPrimarySourceOf
<u>hadRole</u>	wasRoleIn
<u>hadUsage</u>	wasUsedInDerivation
<u>influenced</u>	wasInfluencedBy
<u>influencer</u>	hadInfluence
<u>invalidated</u>	wasInvalidatedBy
<u>qualifiedAssociation</u>	qualifiedAssociationOf
<u>qualifiedAttribution</u>	qualifiedAttributionOf
<u>qualifiedCommunication</u>	qualifiedCommunicationOf
<u>qualifiedDelegation</u>	qualifiedDelegationOf
<u>qualifiedDerivation</u>	qualifiedDerivationOf
<u>qualifiedEnd</u>	qualifiedEndOf
<u>qualifiedGeneration</u>	qualifiedGenerationOf
<u>qualifiedInfluence</u>	qualifiedInfluenceOf
<u>qualifiedInvalidation</u>	qualifiedInvalidationOf
<u>qualifiedPrimarySource</u>	qualifiedSourceOf
<u>qualifiedQuotation</u>	qualifiedQuotationOf
<u>qualifiedRevision</u>	revisedEntity
<u>qualifiedStart</u>	qualifiedStartOf
<u>qualifiedUsage</u>	qualifiedUsingActivity
<u>specializationOf</u>	generalizationOf

<u>used</u>	wasUsedBy
<u>wasAssociatedWith</u>	wasAssociateFor
<u>wasAttributedTo</u>	contributed
<u>wasDerivedFrom</u>	hadDerivation
<u>wasEndedBy</u>	ended
<u>wasGeneratedBy</u>	generated
<u>wasInfluencedBy</u>	influenced
<u>wasInformedBy</u>	informed
<u>wasInvalidatedBy</u>	invalidated
<u>wasQuotedFrom</u>	quotedAs
<u>wasRevisionOf</u>	hadRevision
<u>wasStartedBy</u>	started

Appendix D – Provenance Data Lineage

This Appendix is informative content only. It is not part of the PSAF Volume 3 Provenance normative material. Comments on this section will be considered as informative comments in ballot reconciliation.

Figure 53 shows more detail in the UML model used by the ISO Standard to describe lineage. In some cases, a simple descriptive statement can describe the lineage effectively. In more complex cases, multiple sources and process steps might be required. The definitions of sources and processSteps are also shown in the UML. The capability to specify the spatial and temporal extent of the source and to describe the rationale for a process step are new in the ISO Standard. Note that each source can have any number of associated sourceSteps and that each processStep can have any number of sources (and outputs in ISO 19115-2).



DQ_Lineage (19115-2)

Figure 53: ISO Data Lineage Information Model

Appendix E – IBM Watson Research Model

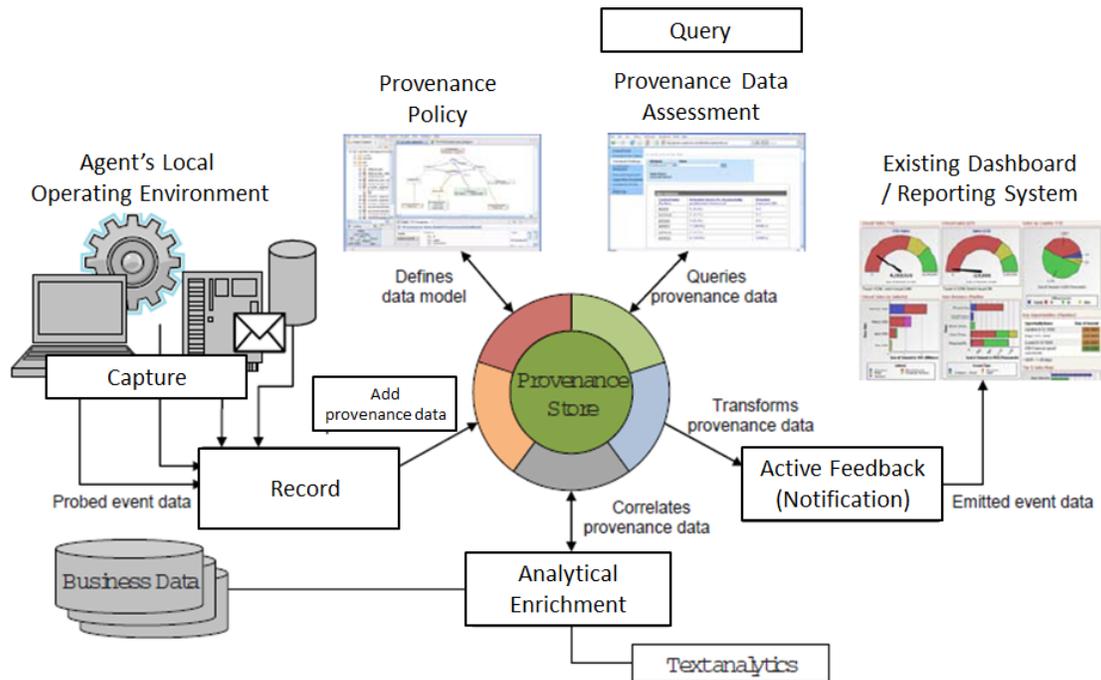


Figure 54: IBM Watson Research Model

Appendix F – Provenance Intra-Domain Model

Figure 55 illustrates an intra-domain Provenance model for the creation of provenance information shared under the requirements of a federated provenance authority.

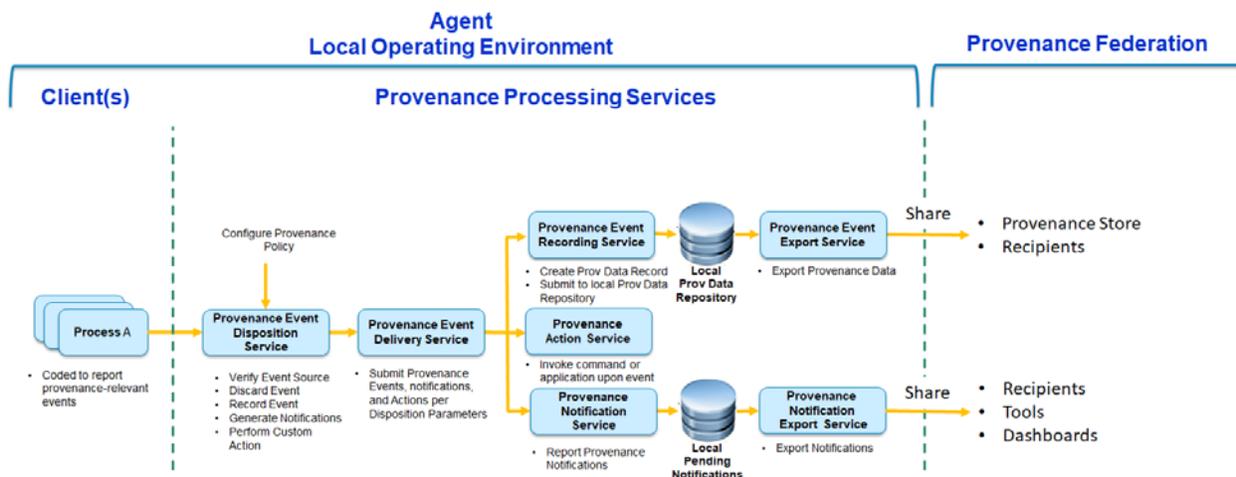


Figure 55: Intra-Domain Provenance Model

Capture/Recording Service

This service allows for the detection of provenance events and the creation of corresponding provenance records to memorialize them. The service performs any processing as may be configured including but not limited to formatting provenance data and adding additional security or other context. Capture and creation occur at the local participant’s level. Recording moves local provenance data to the Provenance Store shared and accessible by all participants. Other information may be captured and recorded to support enhanced provenance analytics.

Query/Analytics Service

This service allows searching and analysis of the Provenance Store which stores the aggregated set of provenance data records associated with resources. A search can be parameter-based to provide granular, flexible searching, for example, retrieve for a specified date range in which all provenance data records meeting the specified search criteria, regardless of contributing source, are returned.

Data analytics services enable enhanced understanding of the aggregated provenance information. This deeper insight into the aggregated provenance data enable a more informed decisions about “fitness for purpose” and facilitate identifying circumstances that require notification.

Analytics services enable deeper insight into provenance data singularly and in the aggregate and may include correlation between provenance records and the content and context of the data including business data that was captured and recorded with the provenance event.

Feedback/Notification Service

This service proactively initiates notifications based on configured rules and triggers. Notifications are sent to destinations specified by the configurations and per Provenance Policy.

Notifications allow external entities (machine or human) to receive updates about certain provenance-related events without continuously polling and querying the provenance store. Example use-cases include alert systems, dashboard systems, and reporting systems.

Notifications can be set based on ordinary provenance events or based on analytics; for example, a suspicious pattern of updates to a resource may require a notification.

Notifications can be delivered using different protocols and communication mechanisms such as web calls (e.g. to a REST endpoint), email, or text messages.

Appendix G – Acronyms and Abbreviations

CDA	Clinical Document Architecture
CRUDE	Create, Read, Update, Delete, and Execute
EHR	Electronic Health Record
FHIR	Fast Healthcare Interoperability Resources
HIPAA	Health Insurance Portability and Accountability Act
HL7	Health Level Seven
ISO	International Organization for Standardization
LCE	Lifecycle Event
MDA	Model Driven Architecture
ODP	Open Distributed Processing
PII	Personally identifiable information
PFA	Provenance Federation Authority
PHR	Personal Health Record
PROV	Provenance Data Model
PSAF	Privacy and Security Architecture Framework
RBAC	Role Based Access Control
REST	Representational State Transfer
RIM	Reference Information Model
RM-ODP	Reference Model of Open Distributed Processing
SAEAF	Services Aware Enterprise Architecture Framework
TF4FA	Trust Framework for Federated Authorization
UML	Unified Modeling Language
W3C	World Wide Web Consortium

Appendix H – Glossary

Activity	An Activity is something that occurs over a period of time and acts upon or with Entities; it may include consuming, processing, transforming, modifying, relocating, using, or generating Entities [W3C Prov Ontology].
Agent	An Agent is something that bears some form of responsibility for an Activity taking place, for the existence of an Entity, or for another Agent's Activity [W3C Prov Ontology].
Author	An Agent (person, system, or an organization) which is the initial creator of an Entity.
Broker	A subsystem or service that facilitates distribution of provenance information between an Agent and the Provenance Store, or between the Provenance Store and a Recipient. Brokers can simplify technical integration between Provenance Store and the Recipient by providing a common interface for integration.
Broker Use Case Scenario	A use case scenario where intermediary acts to facilitate communications between the Provenance Store and Recipient(s).
Collection	A collection is an Entity that provides a structure to some constituents, which are themselves Entities. These constituents are said to be member of the collections [W3C Prov Ontology].
Conceptual Model	A high-level model of a system based on identifying its components, their relationships, and interactions.
CRUD Security Operations	<p>Create is a fundamental operation in an information system that results in the act of bringing an object into existence.</p> <p>Read is fundamental operation in an information system that results only in the flow of information about an object to a subject.</p> <p>Update is a fundamental operation in an information system that results only in the revision or alteration of an object.</p> <p>Delete is a fundamental operation in an information system that results only in the removal of information about an object from memory or storage.</p>
Digital Signature	<p>The result of a cryptographic transformation of data that, when properly implemented, provides a mechanism for verifying origin authentication, data integrity and signatory non-repudiation [NIST/FIPS 186-4].</p> <p>A digital signature is a cryptographic technique used to validate the authenticity and integrity of a message, software or digital document.</p> <p>https://searchsecurity.techtarget.com/definition/digital-signature</p>
Direct Use Case Scenario	Scenario in which Agent provides provenance data directly to Recipient.
Domain	<p>A group of systems or organizations under a single authority and subject to the same policies.</p> <p>A domain is characterized by a domain identifier, domain name, domain authority, and domain qualifier (ISO/TS 22600-2:2006).</p>
Domain Analysis Model	Scope of a problem domain, and to introduce its information content, the parties involved in creating and managing the information, and the relevant behaviors of those parties. [HL7 DAM Guidance doc]
Entity	A physical, digital, conceptual, or other kind of thing with some fixed aspects; Entities may be real or imaginary [W3C Prov Ontology].
Event	A change in an (external or internal) input to a system that triggers a change of state, for example, creation of a new Entity, or the start of an Activity by an Agent.

Facilitating Software	System software responsible for implementing predecessor-successor state changes.
Federated Provenance	Provenance information exchanged among members of a federation.
Federated Provenance Service	A federation of participating systems focused on capturing, collecting, and sharing provenance records.
Federated Store (Repository)	The shared repository(ies) where federation members have agreed to provide specific provenance information as required by federation by-laws.
Federation	A group of information system collaborating and sharing information while maintaining autonomy [Heimbigner-McLeod 1985].
Federation Authority	An organization that governs the federation in accordance with policies and oversees and facilitates member systems or organizations joining or leaving the federation in a dynamic fashion.
Federation Member	A system or organization that participates in a federation.
Fitness for Purpose	Suitable for the reason, objective or goal. The extent to which the information resource is of appropriate quality for the situation in which it is to be used [Klobas 1995].
Lifecycle Event	Any event consequential to the state or content of an Entity. A healthcare event as defined by ISO 21089.
Lineage	A representation of source(s) and production process(es) used in producing an Entity (adapted from [ISO 19115-1:2014]). A description of the Entities and processes that generate, transmit, or influence data. The W3C defines it as “a resource that describes the Entities and processes that generate, transmit, or influence other resources. Lineage information is the basis for resource credibility assessment, giving trust, and allowing regeneration”.
Lineage (ISO 19115)	This is a OWL 2 DL extension of PROV-O that models part of the ISO 19115 UML metadata standard; in particular the concepts relating to lineage. The modelling covers the standard classes prefixed by LI_ ("lineage") and LE_ ("lineage extended") and provides placeholders for the other classes referenced by them. The intention of this ontology is to enable ISO lineage records (typically presented in XML) to be represented according to this ontology and therefore supporting interoperability with other PROV-O provenance records. The design has treated PROV-O as an upper ontology extended with the ISO 19115 concepts, faithfully carrying through the names and structure of the ISO 19115. https://www.w3.org/2001/sw/wiki/PROV
Metadata	Data that provides information about other data. Metadata is machine understandable information for the web.
Performer	This is the actor(s) who performed a provenance relevant state transition. (e.g.the facilitator of an event that changed a predecessor to a successor Entity) HL7 CDA® R2 Implementation Guide: Data Provenance, Release 1 – US Realm Draft Standard for Trial Use (Errata Release) September 2016
Policy	A set of legal, political, organizational, functional and technical obligations for communication and cooperation [ISO/TS 22600-1:2014]. 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) The policy represents the rules and criteria that constrain Activities of the objects to make the domain secure. (OMG Security Services Specification)

Predecessor	<p>An Activity (or Entity) that precedes another Activity (or Entity) based on the <i>wasInformedBy</i> (or <i>wasDerivedFrom</i>) relation.</p> <p>A predecessor is an Activity (or Entity) that precedes another Activity (or Entity) – not in the chronological sense but according to their dependency to each other. A predecessor can have several direct successors. https://www.inloox.com/project-management-glossary/predecessor/</p>
Provenance	<p>Information about Entities, Activities, and people involved in producing a piece of data or thing, which can be used to form assessments about its quality, reliability or trustworthiness [W3C Prov Overview].</p> <p>Provenance of a resource is a record that describes Entities and processes involved in producing and delivering or otherwise influencing that resource. Provenance provides a critical foundation for assessing authenticity, enabling trust, and allowing reproducibility. Provenance assertions are a form of contextual metadata and can themselves become important records with their own provenance.</p> <p>https://www.w3.org/2005/Incubator/prov/wiki/What_Is_Provenance</p>
PROV	The suite of W3C Provenance specification and notes including the Provenance Data Model.
Provenance Activity Chain	A chain of Activities each of which is linked to the predecessor based on an instance of the <i>wasInformedBy</i> relation [W3C Prov Ontology].
Provenance Analysis	<p>Structured provenance assessments based on rules, including trust and compliance assessments,</p> <p>The compilation of a collection of Provenance Records, which may include individual records as well as aggregations, summaries, and other analyses such as correlations.</p>
Provenance Chaining	Metadata that enables linking predecessor and successor provenance information (adapted from [HL7 DPROV CDA IG]).
Provenance Data	See <i>Provenance Record</i> .
Provenance Entity Chain	Provenance chains comprising only Entities formed using the <i>wasDerivedFrom</i> relation [W3C Prov Ontology].
Provenance Event	Any event in the system which is consequential from a provenance perspective and leads to capturing and recording of a Provenance Record, such as creation of an Entity, or start and end of an Activity.
Provenance Metadata	<p>A class that enables one to link predecessor and successor entries, both within and external to the CDA instance by specifying the Provenance Event that changed the former to the latter as well as the performer of the change, the author of the Provenance Metadata, any facilitating software, whether there was a signature on the predecessor prior to incorporation, the applicable provenance policies for recording this information, and Provenance Security Labels that enable recipient systems to evaluate confidence in the successors authenticity, integrity, and reliability without having to register more detail than necessary for access control and integration processing than necessary [HL7 DPROV CDA IG].</p> <p>A class that enables linking predecessor and successor entries and the attributes of the chaining (adapted from [HL7 DPROV CDA IG]).</p>
Provenance Notification	An event sent to federation members using a push-back mechanism. Events could be pre-defined based on policies (e.g, events indicting security, integrity or other events requiring timely notification) or could be custom based on the definition provided by a Recipient at time of requesting a notification subscription.

Provenance Policy	The rules governing federation member participation in a provenance federation. The policy that determines the granularity and details for capturing, collecting and recording provenance information in a system, or a federated system.
Provenance Store	That service that receives, retains, and shares provenance records in a Federated Provenance system.
Provenance Record	An instance of provenance information recorded based on the model presertend in the Class Model of this DAM.
Provenance Report	The compilation of Provenance Records for information disclosed. See <i>Provenance Analysis</i> .
Recipient	Federation member that can request and receive provenance records, provenance reports, or provenance notifications.
Redirect	Direct (something) to a new or different place or purpose.
Relation	A semantic connection between two (or generally more) logical Entities, for example, the <i>used</i> relation between an Activity and an Entity.
Resource	A generalized logical or material object that participates in, or in some way supports federated provenance. See <i>Entity</i> .
Security Authority	An Agent that must be identifiable and responsible for defining the policies to be applied to the domain, but may delegate that responsibility to a number of subauthorities, forming subdomains where the subordinate authorities' policies are applied. (OMG Security Services Specification)
Security Domain	A set of subjects, their information objects, and a common security policy (NIST Special Publication 800-33). <ul style="list-style-type: none"> • 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.
Security Policy	The complex of legal, ethical, social, organizational, psychological, functional, and technical rules for ensuring trustworthiness of health information systems. (ISO 22600-2)
Security Policy Domain	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)
Security Label	The means used to associate a set of security attributes with a specific information object as part of the data structure for that object [ISO/IEC 10181-3:1996]. A security label, sometimes referred to as a confidentiality label, is a structured representation of the sensitivity of a piece of information. Security labels classify constraints on the WHO, HOW, WHEN, WHERE, and WHY – or in other words, on the actor (which is not necessary a person) and the context – for accessing and using a Resource. [HL7 Healthcare Privacy and Security Classification System (HCS)] A <i>security label</i> is a concept attached to a resource or bundle that provides specific <i>security</i> metadata about the information it is fixed to.FHIR V4.0.0

Successor	An Activity (or Entity) that proceeds another Activity (or Entity) based on the <i>wasInformedBy</i> (or <i>wasDerivedFrom</i>) relation.
Successor Event	A successor is an Activity that follows another Activity – not in the chronological sense but according to their dependency to each other. A successor Activity can have several direct predecessor Activities. https://www.inloox.com/project-management-glossary/successor/
Term	A provision stipulated in as an obligation in a contract to which parties to the contract must comply.
Trust	Trust is a term with many definitions and uses, but in many cases establishing trust in an object or an Entity involves analyzing its origins and authenticity. Trust is often equated with provenance, and it is indeed related but it is not the same. Trust is derived from provenance information, and typically is a subjective judgment that depends on context and use. https://www.w3.org/2005/Incubator/prov/wiki/What_Is_Provenance#Provenance_and_Trust
Trustworthy Interoperability	See Trust
Used	Usage is the beginning of utilizing an Entity by an Activity. Before usage, the Activity had not begun to utilize this Entity and could not have been affected by the Entity [W3C Prov Ontology].
wasAssociatedWith	An Activity association is an assignment of responsibility to an Agent for an Activity, indicating that the Agent had a role in the Activity. It further allows for a plan to be specified, which is the plan intended by the Agent to achieve some goals in the context of this Activity [W3C Prov Ontology].
wasAttributedTo	Attribution is the ascribing of an Entity to an Agent [W3C Prov Ontology].
wasDerivedFrom	A derivation is a transformation of an Entity into another, an update of an Entity resulting in a new one, or the construction of a new Entity based on a pre-existing Entity [W3C Prov Ontology].
wasGeneratedBy	Generation is the completion of production of a new Entity by an Activity. This Entity did not exist before generation and becomes available for usage after this generation [W3C Prov Ontology].

Appendix I – References

[Heimbigner-McLeod 1985]	Heimbigner D. and McLeod D., A Federated Architecture for Information Management, ACM Transactions on Information Systems (TOIS), Vol. 3, No. 3, Jul. 1985, pp. 253-278. https://dl.acm.org/citation.cfm?doid=4229.4233
[HL7 DPROV CDA IG]	HL7, <i>Implementation Guide for CDA Release 2 Data Provenance, Release 1 Draft Standard for Trial Use 2015</i>
[HL7 Da Vinci PDex]	HL7 Da Vinci Project, Da Vinci Payer Data Exchange Implementation Guide Release 0.1.0, Jun. 2019. http://hl7.org/fhir/us/davinci%2Dpdex/2019Jun/
[HL7 EHR LCE]	HL7, <i>EHR Lifecycle Events Vocabulary</i>
[HL7 FHIR LCE IG]	HL7, <i>FHIR Record Lifecycle Events Implementation Guide</i> See also: FHIR Object Lifecycle Events Valueset: https://www.hl7.org/fhir/valueset-object-lifecycle-events.html
[HL7 PSAF Vol. 4 Audit]	HL7, Privacy and Security Architecture Framework Volume 3 Audit, Release 1 (planned for May 2019 Ballot)
[HL7 PSAF TF4FA Vol.1]	HL7, <i>Privacy and Security Architecture Trust Framework for Federated Authorization (TF4FA), Conceptual Model Volume 1</i>
[HL7 PSAF TF4FA Vol.2]	HL7, <i>Privacy and Security Architecture Trust Framework for Federated Authorization (TF4FA), Behavioral Model Volume 2</i>
[HL7 PSAF TF4FA Guide]	HL7, <i>Privacy and Security Architecture Trust Framework for Federated Authorization (TF4FA), Guide</i>
[IVOA Prov DM]	International Virtual Observatory Alliance (IVOA), IVOA Provenance Data Model, Version 1.0, IVOA Proposed Recommendation, 15 Oct. 2018. http://www.ivoa.net/documents/ProvenanceDM/20181015/PR-ProvenanceDM-1.0-20181015.pdf
[ISO 19115-1:2014]	ISO 19115-1:2014, Geographic Information -- Metadata -- Part 1: Fundamentals, Apr. 2014. https://www.iso.org/standard/53798.html
[ISO/TS 22600-1:2014]	ISO/TS 22600-1:2014, Health informatics -- Privilege management and access control -- Part 1: Overview and Policy Management, Oct. 2014. https://www.iso.org/standard/62653.html
[ISO/TS 21089]	ISO/TS 21089:2018, Health informatics --Trusted End-to-End Information Flows, Apr. 2018. https://www.iso.org/standard/66936.html
[ISO/IEC 10181-3:1996]	ISO/IEC 10181-3:1996, Information Technology -- Open Systems Interconnection -- Security Frameworks for Open Systems: Access Control Framework, Sep. 1996. https://www.iso.org/standard/18199.html
[ISO/IEC 10746-3]	ISO/IEC 10746-3:2009 <i>Open Distributed Processing – Reference Model: Architecture – Part 3</i>
[ISO/HL7 10781]	ISO/HL7 10781:2015 Health Informatics, HL7 Electronic Health Records-System Functional Model, Release 2 (EHR FM), August 2015. https://www.iso.org/standard/57757.html
[Klobas 1995]	J. E. Klobas, Beyond Information Quality: Fitness for Purpose and Electronic Information Resource Use, Journal of Information Science, Vol. 21, No. 2, April 1995. https://doi.org/10.1177/016555159502100204

[NIST/FIPS 186-4]	National Institute of Standards and Technology (NIST) publication, Digital Signature Standard (DSS), July 19, 2013. https://doi.org/10.6028/NIST.FIPS.186-4
[NIST SP 800-53]	National Institute of Standards and Technology (NIST), <i>Security and Privacy Controls for Federal Information Systems and Organizations</i>
[NIST SP 800-63-r3]	National Institute of Standards and Technology (NIST), Digital Identity Guidelines, Federation and Assertions, NIST Special Publication 800-63, Revision 3, June 2017. https://pages.nist.gov/800-63-3/sp800-63-3.html
[ONC HIT S&I PI]	<i>ONC Privacy & Security Update: Data Provenance (DPROV) S&I Initiative</i>
[W3C PROV AQ]	W3C, PROV-AQ: Provenance Access and Query, 30 April 2013. https://www.w3.org/TR/prov-aq
[W3C Prov Constraints]	W3C, Constraints of the PROV Data Model, 30 April 2013, https://www.w3.org/TR/prov-constraints
[W3C Prov DM]	W3C, Semantics of the PROV Data Model, W3C Recommendation, 30 April 2013, https://www.w3.org/TR/prov-dm
[W3C Prov Ontology]	W3C, PROV-O: The PROV Ontology, W3C Recommendation, 30 April 2013. https://www.w3.org/TR/prov-o
[W3C Prov Overview]	W3C, PROV-Overview, An Overview of the PROV Family of Documents, W3C Working Group Note, 30 April 2013. https://www.w3.org/TR/prov-overview
[W3C Prov Primer]	W3C, PROV Model Primer, 30 April 2013 https://www.w3.org/TR/prov-primer/
[W3C Prov Req]	W3C Provenance Incubator Group, Requirements for Provenance on the Web, 9 April 2010. https://www.w3.org/2005/Incubator/prov/wiki/User_Requirements
[W3C Prov XG FR]	W3C Provenance Incubator Group, Provenance XG Final Report, 8 December 2010. http://www.w3.org/2005/Incubator/prov/XGR-prov