Extending the FHIR standard to handle provenance

John Moehrke
johnmoehrke@gmail.com

Arnon Rosenthal
The MITRE Corporation
arnie@mitre.org

Adriane Chapman
The MITRE Corporation
achapman@mitre.org

Abstract

In this work, we look at the problems a domain-specific standard committee faces when trying to "involve" provenance to meet domain requirements, without committing to the major expansion of adding a general provenance capability[1]. We also begin the discussion of what the provenance community can do to assist the domain-specific creation committee with the easy inclusion and usage of well-specified, and provenance-community approved guidance.

Categories and Subject Descriptors • Information systems ~Data provenance • Software engineering ~Entity relationship modelling

General Terms Standards

Keywords *provenance, HL7®, FHIR®©*

1. Introduction

Interoperability Standards are crucial to the exchange of data among systems and organizations. The provenance community has proposed some standards to address provenance in a generally applicable way. A different question is examined here: *How is provenance information to be exchanged in situations under the control of other communities?* In particular, we discuss the situation as seen by mainstream standards groups in healthcare.

Within the provenance community, provenance is the *raison d'être*, the paramount concept to model and enable. The health data standards community is well aware of the importance of provenance for certain information, but provenance is just one among many features they desire, not a top goal. We use the experience of the Fast Healthcare Interoperability Resources (FHIR) standard creation to gain insight into how a domain-specific group approaches provenance. We also begin the discussion of what the provenance community can do to assist the domain-specific creation committee with the easy inclusion and usage of well-specified, and provenance-community approved guidance.

The HL7® FHIR® standard is currently undergoing rapid and active development. The HL7 Security workgroup is considering creating a model for provenance, intended to be usable (but not mandated) throughout FHIR. This gives an opportunity to begin a discussion with the provenance community; we are not aware of prior discussions of this sort with a domain community.

While HL7 Security workgroup will go its own route (for scheduling and skillset reasons), we suggest that provenance community participants may wish to examine what it would take to apply the provenance community's general standard [11] in the context of a mainstream standard such as HL7® FHIR® [3]. It provides us with an interesting opportunity to observe and reflect upon the multi-varied needs of the domain-specific standards body, and how well they are served by how the provenance community defines provenance.

Our contributions to the provenance community include:

1. A description of the FHIR standard, and its current provenance needs
2. Discussion of how other domain-specific standards, such as HL7 CDA, IRM and ISO 19115 use provenance
3. A description of the problems a domain-specific standard creation committee faces when attempting to consider/implement provenance
4. A discussion of what the provenance community can undertake to help domain-specific standards committees implement high-quality, useful provenance constructs in their standards.

In Section 2, we provide background on provenance standards and on the FHIR data standard effort. Section 3 illustrates how treatments of provenance in domain-specific standards are not pristine, while Section 4 outlines the problems met by the domain-specific standards committee when trying to incorporate provenance. A discussion of possible mitigations the provenance community can undertake to facilitate inclusion of useful provenance concepts in a domain-specific standard is in Section 5.

2. Background and Related Work

2.1 Provenance Standards

We distinguish two types of standards of interest to the provenance community: native provenance standards and domain-specific standards that include provenance. W3C's PROV [11, 12] and

[1] There is an old joke that in a bacon and eggs breakfast, the chicken is involved but the pig is committed

The author's affiliation with The MITRE Corporation is provided for identification purposes only, and is not intended to convey or imply MITRE's concurrence with, or support for, the positions, opinions or viewpoints expressed by the author.

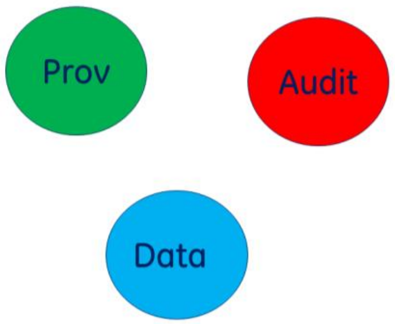

Figure 1: View of a data standard if showing "tracking" components as first class citizens.

the Open Provenance Model [8] are native provenance standards. Provenance is treated as a freestanding capability, and all of the needs, requirements, and uses of provenance are analysed and taken into account for their creation. However, provenance can be found in many domain-specific standards, including ISO19115 [1], HL7 CDA [2] and IRM [9]. Within these standards, provenance is "involved", but not the primary focus of the standard itself.

Provenance tracks *data flows* and *computations* that lead to a value – what operators with what inputs. A major challenge is how to insert the capture into various kinds of software systems, as they are built or after the fact. It records functions invoked, and the arguments and (ideally) the hidden variables that affect those functions. In essence, it tries to capture the computation graph for every significant output. The provenance community considers provenance a freestanding and essential component of any system, as in Figure 1. In contrast, Healthcare is concerned mainly with two functions: attribution or chain of custody (the data flow case), and audit, making provenance usage look more like Figure 2.

2.2 FHIR and HL7

HL7 (originally named Health Level 7 as an homage to the OSI stack) is an influential standards organization that covers the primary healthcare treatment scenarios. It includes standards for managing patient identity, provider identity, orders for procedures, observations of medical facts, dispensing of drugs, and documentation. HL7 focuses primarily on Interoperability Standards, the standards used to enable systems to communicate and act upon that communication. The DICOM standards organization is a companion standards organization that covers medical imaging.

FHIR® is a new effort to take the previous standards models and re-envision them leveraging current web standards. FHIR® Focuses on http RESTful interaction, but can also be exchanged in a message pattern, or as a document. The focus is on modular components called "Resources", expressed using simple XML and JSON. For example, there is a "Resource" defined for recording a clinical observation. A previous effort, (HL7 v3) was too complex for many to use and thus failed. FHIR thus aims for basic functionality and simplicity, rather than perfection.

3. Examples of imperfect provenance in other standards

Because provenance is "involved" in other standards, but not the primary focus, the domain specific standards often use provenance in a manner that provenance community purists find unsatisfying. Below, we list a few of the mutated forms of provenance:

1. 1-hop provenance: In IRM [9], there is room for provenance information, but the only information that is captured is 1-hop provenance, not a full chain.
2. Buried within a system. The classic proprietary Electronic Medical Record is modelled in FHIR as a database plus logs. Provenance might be represented in the data or might need to be extracted from the logs.
3. Buried within the data: In the HL7 CDA [2], the provenance information has a chain of custody model, however it is in the "data" portion of the standard, not a provenance section, or even a metadata section.
4. "Provenance" = "Audit": The argument is often made that provenance isn't needed, because there are audit logs that can be used. For instance, an alternative classic proprietary Electronic Medical Record model that has a database and logs, where the logs are used for both provenance and audit purposes.
5. "Provenance" is completely separate from "Audit" (the exact opposite of #4 above), as in IRM: There are records that are required for system auditing, such as the need to log when a system starts-up, is shut-down, temperature readings, virus-detection, port-use, user-authentication-failure, etc. These records are used for system auditing, not for data object derivation understanding. Thus, provenance tracking needs are considered distinct from audit needs, instead of an acknowledgement that provenance may often need some information from the audit logs.

The above was a *rough* categorization of differences between the provenance-community's conceptions and what is done by domain-specific bodies. However, the categories are not mutually exclusive and a single standard can use multiple tactics. For instance, FHIR is a hybrid. Some clinical resources are tied to the author and their provenance elements are built into the clinical resources. However, there are also a Provenance resource available for use against any resource, and an AuditEvent resource for recording all auditable events, possibly duplicating the Provenance record. FHIR's model constrains the Provenance to recording of Create, Update, and Import actions.

4. Problems modelling provenance in a domain-specific standard

Section 3 described usages of provenance that the provenance community finds awkward. In this section, we describe problems encountered by creators of a domain-specific standard.

4.1 Priority

To the provenance community, provenance is best done as a powerful, general solution. The provenance data stream and powerful graph operators can be applied to any number of uses from intrusion detection [5] to data trustworthiness [6] to scientific re-computations [7, 10]. However, a domain-specific community may choose to define a handful of data elements, and leave it to applications or humans to exploit it as they can. In FHIR, the provenance standard is restricted to: record of the authorship, for create, update, and import. It does not track access, use, and disclosure – key operations for traditional medical records, to enable a clinician to judge sources' trustworthiness, and to enable tracing responsibility when a mistake is found (e.g. for medical fraud, and malpractice). Implementing a full provenance facility would be costly and disruptive for the EHR and other health application systems. Therefore, the domain-specific standards often limit the support for provenance.

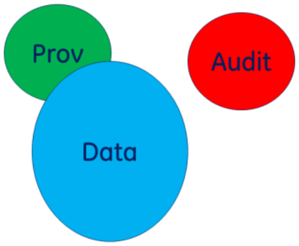

Figure 2: The inclusion of provenance within the data specification. [4]

4.2 Overlap with mainstream needs

When provenance is important to mainstream business processes or to meet legal obligations, the relevant standards include data elements for provenance. For instance, when prescribing drugs, it is important to know who prescribed them, and who filled them. It is embedded with the data as another attribute, not in a separate provenance section.

A domain committee needs to ask about the value added from new constructs, examining three cases:

- Where the domain-specific specification has already implemented provenance in special elements.
- Where the elements of provenance are fundamental to the Resource.
- How difficult inclusion will be for systems that have not yet implemented it.

4.3 One of Many

The healthcare area overlaps many technologies and capability needs that need provenance. For FHIR®, these included: care provision, medication, immunizations, diagnostics, personal management, organization/location management, encounter management, schedule management, order management, medical device use management, billing, payment, contracts, consents, accountability, and other. The creators of the domain-specific standard understand that systems that utilize that standard cannot invest in a separate infrastructure for each such area; infrastructure for a special need like provenance must leverage other general capabilities. Consequently, a variety of engineering compromises are made.

4.4 Provenance sub-committee

Most domain-specific standards committees are staffed by intelligent and passionate individuals who are very knowledgeable about their specific domain. They understand innately what the primary data in their standard is meant to be used for. However, they do not have a provenance committee. The overhead of creating one would be high. Because healthcare is enormously complex, standards-making is distributed across many committees, and the provenance-component retains the flavour and concerns of the sub-committee under which it fell. FHIR® has IT, standards, and medical records experts, but no provenance experts. The provenance component was done by the security committee, with oversight by medical records experts.

4.5 Overlap with other needs

Provenance is captured to facilitate interpretation and tracking of the mainstream data describing medical events. On the other hand, some provenance capabilities overlap with other “2^nd^ class citizens”, such as:

- Audit logging
- Trustworthiness of identity determinations – relevant to access privileges and also to merging data from multiple sources. (Since this is not tracked in FHIR, it will not be discussed further.)

From FHIR's perspective, provenance (P) and audit (A) are distinct; both were considered by the HL/7 Security WG. There were debates about whether one subsumed the other. The situation clarified when one compared along several explicit axes:

- *Are the event types known by Audit a superset of those known to Provenance?* Yes, since Audit examines Reads and System behaviours (intrusions, failures); in contrast; Provenance considers Reads only if they are part of creating another entity that has provenance.
- *Are the event instances tracked by Audit a superset of those tracked by Provenance?* Very likely not. For cost effectiveness, both may ignore items that seem not of interest track, and track others at coarse granularity. However, their selection criteria are likely to be quite different, e.g., for legal versus scientific justifications.
- *Are the analysis operators defined for Audit a sub- or superset of those for Provenance?* No. For example, audit systems may have exception reporting, while Provenance’s emphasis is to track reasoning.

To illustrate the divergences, consider how a Clinician might create an Order for some procedure. She would have reviewed many parts of the record, described observations and the Order as part of the encounter, and perhaps created ancillary resources and associated them with the Order (e.g., Specimen). To this end, two sets of event codes exist for provenance and audit, and are incompletely overlapping on both sides.

There might be one provenance record on the DiagnosticOrder, pointing at those parts of the record that the Clinician considered important, i.e., what she felt justified her conclusion. In contrast, there might be 50-60 Audit events recording all the pieces of evidence that the doctor reviewed, especially sensitive ones, and multiple specific audit events recorded to create the Order, and Specimen resource.

5. Discussion

While the problems we highlight in this work are common to all standards (difficult use, misuse, etc.), what can the provenance community do to help domain-specific standards committees? Essentially, the provenance community must find a way to maximize consistency among the treatment of provenance and related information. Additionally, the provenance community must find a way to give the domain specific communities the freedom to use the bits of provenance that they need. Moreover, because the domain-specific committees don’t know provenance inside-out, they need something to quickly “grab and insert”.

For all of this to happen, we must consider how standards can share and overlap. How can we define a provenance standard from pieces that can be pulled into a domain-standard for specific uses (e.g. auditing, security, etc.)? Or, should we gracefully combine standards (such as provenance standards with Audit standards)? How can provenance-like data elements be able to stand alone and also to coexist with an explicit provenance capability? Can we support the spectrum, between adding one of several provenance data elements, to different levels of capability?

Domain-specific standard creators could use advice on when to make the provenance record a standalone resource, vs a common part of all resources. If it is within resources, it will be in a form suited to that resource, not duplicated, and not lost. However, it is more difficult to provide general provenance functionality. How are the risks of separation of the data from provenance to be managed? Is it possible to support both logical viewpoints, abstracting away from the physical structure? On a narrower point, what should be the technical *and political* relationships between audit and provenance.

Additionally, the provenance community needs to help in creating provenance components that are efficient, simple and effective – both as components of standards and of software systems. Simplicity may be more valuable than great expressiveness. To this end, a web friendly encoding based on simple XML and JSON, using http RESTful interaction model would be viewed as helpful.

Finally, it bears repetition that there is a learning curve to provenance. When faced with many other tasks and priorities in the domain-specific standard, the domain-specific standard creators are going to do "just enough" investigation to get the few provenance-components that enable their current work, and may not understand deeper implications, e.g., for extensibility to other cases. The provenance community needs to find a way to help educate and share.

6. Conclusions and Future Work

This work introduces the provenance community to the HL7 FHIR® standard. FHIR comes from an interoperability standards organization with decades of experience that is taking a fresh look and designing a new standard. While the committee is addressing provenance, it is doing so in a manner that would not please a provenance purist.

We described the FHIR® standard and its current provenance needs. We also describe examples of how provenance has been used by other standards in the past that do not reflect the provenance community's current standards [11]. We describe many of the challenges facing a domain-specific standards committee, and present many questions to the provenance community in how we can facilitate and encourage the use of "good" provenance by other standards.

References

[1] "North American Profile of ISO19115:2003 - Geographic Information - Metadata." NAP Metadata Working Group 2005.

[2] "Clinical Document Architecture." http://www.iso.org/iso/iso_catalogue/catalogue_tc/catalogue_detail.htm?csnumber=44429: HL7® 2009.

[3] "FHIR® standard." https://www.hl7.org/fhir/: HL7® 2016.

[4] "Provenance vs Audit - it is not a competition," in *Healthcare Security/Privacy*, J. Moehrke, Ed.: http://healthcaresecprivacy.blogspot.com/2016/03/provenance-vs-audit-it-is-not.html, 2016.

[5] M. D. Allen, A. Chapman, L. Seligman, and B. Blaustein, "Provenance for Collaboration: Detecting Suspicious Behaviors and Assessing Trust in Information," *CollabCom*, 2011.

[6] D. L. Chenyun Dai, Murat Kantarcioglu, Elisa Bertino, Ebru Celikel, Bhavani Thuraisingham, "Query Processing Techniques for Compliance with Data Confidence Policies," in *SDM*, 2009, pp. 49-67.

[7] T. McPhillips, T. Song, T. Kolisnik, S. Aulenbach, K. Belhajjame, K. Bocinsky, Y. Cao, F. Chirigati, S. Dey, J. Freire, D. Huntzinger, C. Jones, D. Koop, P. Missier, M. Schildhauer, C. Schwalm, Y. Wei, J. Cheney, M. Bieda, and B. Ludaescher, "YesWorkflow: A User-Oriented, Language-Independent Tool for Recovering Workflow Information from Scripts," *International Journal of Digital Curation*, 2015.

[8] L. Moreau, B. Clifford, J. Freire, J. Futrelle, Y. Gil, P. Groth, N. Kwasnikowska, S. Miles, P. Missier, J. Myers, B. Plale, Y. Simmhan, E. Stephan, and J. Van den Bussche, "The Open Provenance Model core specification (v1.1)," *Future Generation Computer Systems*, 2010.

[9] Office of the Director of National Intelligence, "Information Resource Metadata," http://www.dni.gov/index.php/about/organization/chief-information-officer/information-resource-metadata, 2012.

[10] C. E. Scheidegger, H. T. Vo, D. Koop, J. Freire, and C. Silva, "Querying and Re-Using Workflows with VisTrails," *SIGMOD*, 2008.

[11] W3C, "Provenance Data Model http://www.w3.org/TR/prov-dm/," 2013.

[12] W3CProvenance, "http://www.w3.org/2011/prov/wiki/Main_Page," 2012.