

IWBase: A Justification Metadata Infrastructure

Deborah L. McGuinness Paulo Pinheiro da Silva
Cynthia Chang
Stanford University, USA,
{dlm,pp,csc}@ksl.stanford.edu

Abstract

If users are expected to depend on answers from applications and services on the web, they need to have methods for asking for information that will increase their understanding and their trust in the answers. One way of increasing a user’s understanding and trust is to provide a summarized or detailed description of how the answer was obtained. This process would typically depend on some kind of proof structure defending the answer but this alone (no matter how understandable and/or thorough the proof is) is not enough. Any presentation of why a user should depend on an answer also critically depends on metadata concerning elements used in the proof such as underlying data sources, their authoritative-ness, reasoning methods used, any uncertainties, etc. In this paper, we introduce IWBase, a comprehensive OWL-based distributed and scalable infrastructure for effectively creating, storing, evolving, and using proof-related metadata for explanations. IWBase metadata is used to annotate answer justifications encoded in the Proof Markup Language (PML). IWBase is currently in use by a number of DARPA and ARDA projects.

1 Introduction

Our research aims to improve user acceptance of answers by making them “actionable” and defensible. Answers will be actionable if users have enough information about the answer that allows them to believe the answer enough to take action using the information. Answers are defensible if the user requesting the answer can understand and explain why the question answering systems believes the answer is valid. Our work makes answers actionable and defensible by providing explanatory information about the question answering process that exposes how the answer was generated. Explanations depend on provenance metadata and provenance metadata infrastructure to make answers actionable and defensible. Without metadata, users may be unable to understand explanations because, for example, they may be unable to gather critical information about sources accessed, question answering systems, and inference methods. Without metadata, tools may be unable to derive alternative (potentially more understandable) explanations because the tools may not know which

rules (rewriting rules) to use for abstracting explanations. Without metadata, users may not trust answers because they may not have access to tools that check the explanations and compute trust values for answers. Without metadata infrastructure, it may become difficult for users (humans and agents) to include, maintain, evolve, and query provenance metadata.

In this paper, we describe the metadata that is required in addition to answer justifications¹ to generate explanations that improve understanding and trust. Answer justifications are encoded in the Proof Markup Language (PML) described in [12]. These justifications may include information about the knowledge sources used as well as the processes involved for accessing and integrating knowledge. If answers are obtained simply by looking up stored information in a data source, then justifications need to include provenance information about the data source such as its recency, authoritativeness, etc. If answers are retrieved from multiple sources and combined in some manner, then information about all sources used as well as combination techniques needs to be accessible to interested users. If answers are obtained by information manipulation processes, then justifications need to be informed by descriptions of the manipulation processes such as their authors, versions, inference strategies, and assumptions. This kind of meta information is required in addition to the actual proof trace of the manipulations that were performed.

It may be worth noting that our original explanation focus was on the information manipulation area since our early target question answering applications used reliable and familiar data sources and used complicated (and potentially distributed) reasoning processes to reach conclusions. In those situations, users were familiar with and trusted the source data but needed help in understanding the deduction paths leading to conclusions. As our application focus expanded to more typical web applications with diverse data sources, some of which were likely to be unknown to users obtaining the final answers, it became clear that providing the option of accessing provenance information was required, even if just in a summary mode so that users could quickly check source usage. For example, some question answering applications rely on knowledge bases that were built and maintained using information extraction techniques, such as those in UIMA [2]. The original data may be in free text sources where some of the free text may be from known places and other text sources may be unfamiliar (and of unknown quality and recency). Users obtaining answers based on these knowledge bases need to have the option of finding out which raw text sources were used, when they were updated, how they were transformed into structured knowledge bases, and then later how the information was manipulated.

This paper introduces our research thrust on provenance metadata and its associated infrastructure that, along with information manipulation traces, are required to present comprehensible and useful explanations of answers. The paper presents an example of an answer justification annotated with provenance meta data. We introduce the terminology used in our provenance infrastruc-

¹The terms *justification*, *information manipulation trace* and *proof* are used interchangeably in this paper s.

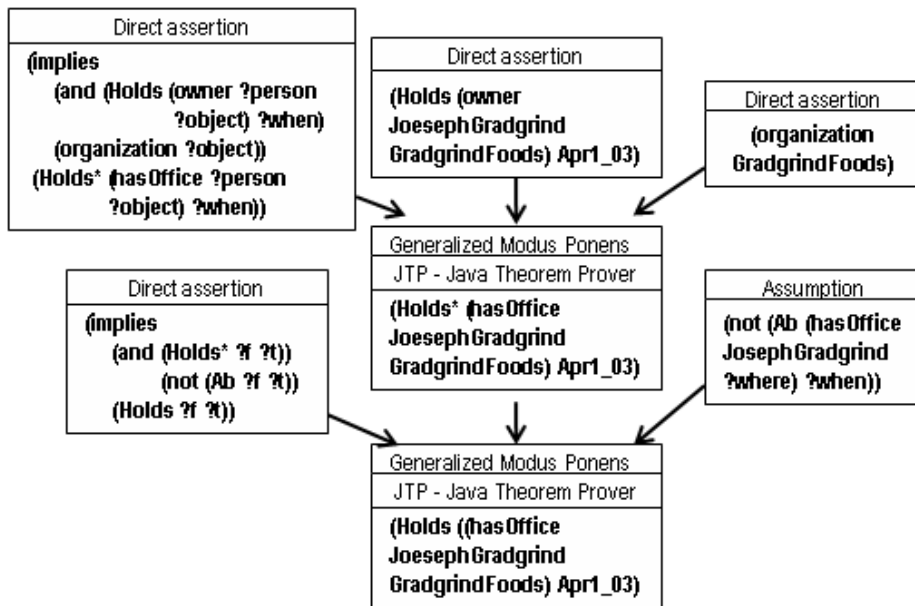


Figure 1: A typical proof

ture. We present IWBase as a distributed repository of proof-related metadata that provides a solution to our provenance requirements. We introduce the major components of our infrastructure (including registry and support services) and describe how our solution is maintainable, scalable, distributable, while supporting privacy and usability requirements. We will conclude with a discussion of the relative merits of the approach, current and planned usage of the implementation, and related work.

2 Metadata in Action: An Example

The need for provenance elements is compelling when any kind of meta-cognition activity is performed that relies on proofs that describe the answer derivation (or answer composition, answer extraction, etc.). For example, raw information manipulation traces are often used by reasoner developers to debug their engines. More sophisticated uses of proofs involve the exchange of proofs among reasoners in a hybrid reasoning environment [13] or the use of proofs as an input for generating human understandable explanations [8]. In these examples and in many others, providing information to users that helps them understand the information manipulations typically requires access to information from the question answering systems about what they did and this is often unavailable. The rendering of a proof in Figure 1 has several examples of uses

of the IWBase elements. There, each statement is annotated with the name of the engine responsible for adding the statement to the proof, i.e., **JTP - Java Theorem Prover** [3]. Further, if a statement is derived, the proof is annotated with the software responsible for the statement derivation in addition to the source responsible for creating of the node set. For derived statements, the proof is annotated with the inference rule applied by the engine to the statement’s antecedents allowing the derivation of the consequent statement. For example, in Figure 1, the **generalized modus ponens** rule was used to conclude² that Joeseph Gradgrind had an office at GradgrindFood on April 1, 2003, which is represented in the proof as“(Holds (hasOffice JoesephGradgrind GradgrindFoods) Apr1.03)”.

The proof annotations mentioned above are just few examples of IWBase elements that are associated with the proof in Figure 1. As presented in the proof, provenance meta data is used to increase the understanding of proofs for reasoner developers and even for other users not involved on the answer generation process, e.g., a user trying to understand how the answer was produced. In Figure 1, each PML element in the proof is rendered as a hot link leading to further information about proof elements.

3 Terminology

The IWBase architecture described in Section 4 has two main concepts: *node elements* and the IWBase *node*. By defining node elements, we describe how meta-data is classified and represented. By defining the IWBase node, we describe the evolving infrastructure for storing, maintaining and using proof-related meta-data on the Web.

3.1 Node Elements

Node elements³ are also PML elements. They are instances of PML classes, which are OWL [9] classes. They are used to build OWL documents representing both proofs and proof provenance information. PML concepts can be considered to be either *proof level concepts* representing proof elements directly or *provenance level concepts* representing information about terms used in a proof.

Proof elements are PML elements at the proof level. They are used for building interchangeable, distributed proofs and explanations. *NodeSet* and *InferenceStep* are the main types of proof elements. A *NodeSet* represents a step in a proof whose conclusion is justified by any of a set of inference steps associated with the *NodeSet*. Each node set has one URI. PML adopts the term “node set” because each instance of *NodeSet* can be viewed as a set of nodes gathered from one or more proof trees having the same conclusion. An *InferenceStep*

²This used three antecedents representing a provisional conclusion about Joeseph Gradgrind’s office and the lack of atypicality with respect to Joeseph Gradgrind’s office.

³Terms defined in this Section are typed in **bold**.

represents a justification for the conclusion of a node set. Inference steps are anonymous OWL classes defined within node sets. We assume that applications handling PML proofs are able to identify the node set of an inference step and thus inference steps have no URIs.

Provenance elements are PML elements describing the origin of proof elements. *ProvenanceElement* is a superclass of the PML concepts at the provenance level (such as source, language, and inference rule). A taxonomy of provenance elements is shown in Figure 2. The *ProvenanceElement* attributes include:

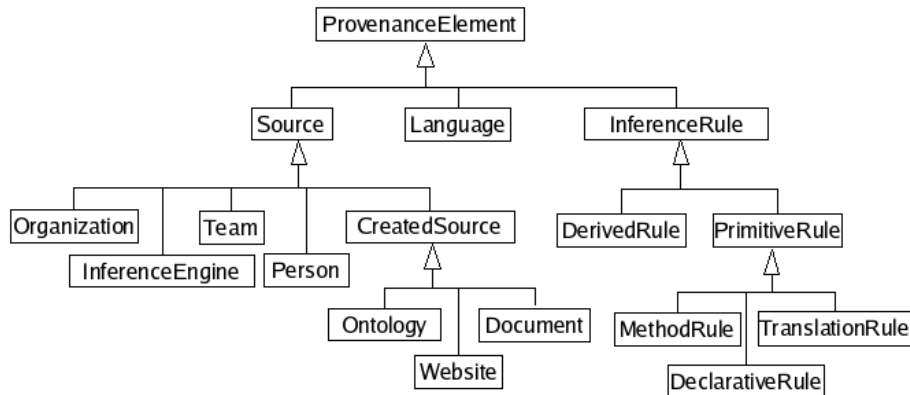


Figure 2: Upper PML taxonomy of provenance elements

- The **URI** of a provenance element is the unique identifier of the provenance element. Every provenance element has one well-formed URI, which is an instance of the primitive type *anyURI*.
- The **URL** of a provenance element describes an URL used to browse an element's web document. For instance, if the provenance element is an organization named the New York Times, then the <http://www.nytimes.com> URL can be used to access a web document about the organization. In this case the URL points to the organization's web site. A provenance element can have zero or one URLs.
- The **Name** of a provenance element describes a short name (or "nickname") for the element within the IWBBase. Every provenance element has one name, and the name is an instance of the primitive type *string*.
- The **Submitter** of a provenance element represents the team of people⁴ responsible for the registration of the provenance element in IWBBase. Every provenance element has one submitter, and the submitter is an instance of the type *Team*, which is a subclass of *Source* (see Figure 2).

⁴Of course a team can consist of any number of people including zero known members to date or a single member.

- The `DateTimeInitialSubmission` of a provenance element is the date when the provenance element was first registered in IWBBase. Every provenance element has one `DateTimeInitialSubmission`, and that is an instance of the primitive type *dateTime*.
- The `DateTimeLastSubmission` of a provenance element is the last date when the provenance element was registered in IWBBase. Every provenance element has one `DateTimeLastSubmission`, and that is an instance of the primitive type *dateTime*.
- The `EnglishDescription` of a provenance element is a description in English of the provenance element. A provenance element can have zero or one descriptions in English, and `EnglishDescription` is an instance of the primitive type *string*. The description in English is intended to be used by tools to present provenance elements to human agents. For example, the description in English of the *Modus Ponens* inference rule may be a better presentation and more informative for most users browsing a PML document than the presentation of the formal logical specification of the rule.

Provenance elements are formally specified in OWL⁵ and are briefly described below. Further descriptions of provenance elements including the description of some relationships among provenance elements are available in [12].

A *Source* represents an entity which is the source of the original data. A source can be either an *Organization*, an *InferenceEngine*, a *Team*, a *Person* or a *CreatedSource*. *CreatedSources* can be *Ontologies*, *Websites* or, generally, *Documents*. *InferenceEngine* is a source representing an engine that is able to produce a justification for a given conclusion. Note that the use of the term “inference engine” in this paper is not limited to engines with reasoning capabilities. For example, a search engine retrieving information may serve as an inference engine and it may provide a justification of its answer by a direct assertion inference step. Similarly extraction modules may be viewed as inference engines and, in fact, we have registered a number of extraction modules from the UIMA extraction toolkit[2] so that extracted answers may be explained in Inference Web using PML[11].

The *Language* provenance element is used to encode a language used to write conclusions of node sets. Any language can be registered in IWBBase including formal logical languages, such as KIF, and natural languages, such as English, and representation languages such as DAML+OIL and OWL.

An *InferenceRule* is a specialization of *ProvenanceElement* used for describing rules applied to premises deriving node set conclusions. An *InferenceRule* can be either a *PrimitiveRule* or a *DerivedRule*. A *PrimitiveRule* is a type of an inference rule that is implemented by one or more inference engines. A given rule R_1 may only be called primitive when it becomes associated with one or more inference engines. Thus, assuming that R_1 is implemented by an inference

⁵The PML ontology version 0.9.2 is available at <http://iw.stanford.edu/2004/07/iw.owl>

engine E_1 , the inference engine may declare R_1 to be a primitive rule. The notion that R_1 is a primitive rule for one specific engine is relevant since R_1 may also be derived from R_2 , which is a primitive rule for another engine E_2 . In this case, R_1 may be registered once as a primitive rule and zero or more times as a derived rule, depending on how many combinations of rules are used to derive R_1 . A *DerivedRule* is an *InferenceRule* specified from a PML node set schema with the restriction that each node set schema must have one and only one inference step.

3.2 IWBase Node

IWBase Node critically depends on the IWBase registry, IWBase registrar and IWBase configuration file so we will define these first.

An **IWBase registry** is a repository of provenance elements. In order to support interoperability when sharing provenance metadata among Inference Web tools and between Inference Web tools and other Semantic Web tools in general, elements in the registry are stored as PML files. Thus, with the use of Semantic Web tools, one can retrieve, parse, and use proof-related metadata. For querying and maintaining large quantities of metadata, the parsing of PML files has shown to be too expensive. Therefore, to increase scalability, provenance elements are also stored in a database.

An **IWBase registrar** is a collection of applications used for maintaining a registry. From a human user point of view, the registrar is an interactive application where the user can add, update, and browse the registry contents. From a software agent point of view, the registrar is a collection of services for querying and updating the registry. The registrar is responsible for keeping the synchronization between the registry repository of OWL files and the registry database.

An **IWBase configuration file** contains configuration and deployment information for IWBase node applications. Information is defined as parameters in name/value pairs and is loaded at application startup time. Global parameters are shared among all applications deployed on the node. Each IWBase application deployed on the node has its own collection of parameters, which can override global parameters.

An **IWBase node** is a named collection of computing resources available on a machine. One machine may have multiple nodes. In concrete terms, a node is characterized by having the following resources:

- a registry;
- a registrar;
- access to read and write in a filesystem in order to store the registry, the registrar and any optional IWBase applications;

- access to a web server platform⁶ supporting servlets/JSP and SOAP services, user authentication and authorization;
- permission to change the configuration of the web server;
- access to a database in a relational database management system (DBMS)⁷ including the following: permission to update the schema of the database, permission to add and update tuples in the database;
- a collection of optional applications based on IWBBase elements;
- a configuration file used to set parameters for the registry, some parameters for the web server and DBMS, and parameters for IWBBase applications installed on the node.

One machine is allowed to have multiple nodes. Nodes on a common machine may share filesystems, web servers and DBMSs. However, to accommodate its registry, each node needs to have a non-shared directory in a shared filesystem for the OWL files and a non-shared database.

3.3 IWBBase Applications

An **IWBBase application** is an application installed under an IWBBase node web server, using the IWBBase configuration file, and providing some functionality that depends on PML elements of the node. The registrar is an IWBBase application.

Engines can use any strategy for creating PML documents including the use of **PML generation services (PGSs)**, the use of the PML API, or even the implementation of their own code to generate PML documents. IWBBase node's PGSs provide support for querying the registry through a combination of queries to database(s) and OWL files. In addition to the use of indexes for querying the registry, PGSs provide a mechanism for selecting and associating provenance elements with proof elements. Even if none of those services are critical for applications, users may still choose to use PGSs since they provide an uniform way of generating PML documents thus supporting consistent usage if the PML specification evolves.

We envision the addition of many IWBBase applications to IWBBase nodes. For instance, at the moment we are developing an **AnswerTrustComputation** service[14] to compute trust values for answers. These services rely on a combination of trust relations, some of which are already implemented in IWBBase. These services are supported by a trust network of users and node elements.

⁶IWBBase currently has support for Tomcat and Axis.

⁷IWBBase currently has support for MySQL.

4 Architecture

We use the term **IWBase** to refer to an hypothetical “entire” collection of IWBase nodes, whether nodes are interconnected or not and whether nodes are available for use or not. In this section we discuss the different types of nodes and node inter-communications in IWBase. Before discussing node configurations, we may need to briefly discuss the roles of IWBase users.

4.1 Users

IWBase users can play one or more of the following roles when interacting with an IWBase node.

A **node visitor** is some “web user”, human or agent, using the node either to inspect the registry or to interact with the registrar for the purpose of browsing node elements. Node visitors do not need to be registered with the node and do not have ownership of node elements. In public nodes, node visitors may browse the entire content of the registry.

A **node member** is some user, human or agent, using the registrar for maintaining node elements. Node members are registered with the node administrator so that they may have access to the registrar for adding and updating node elements. Node members are also registered as members of PML Teams (a PML provenance element).

A **node administrator** is a human user responsible for the following tasks: maintaining and tuning the IWBase node web server and DBMS, creating node members, adding and removing IWBase applications, and setting IWBase application parameters in the node configuration file.

4.2 Nodes and Node Associations

An IWBase node is either a **core node** or a **domain node**. Some PML elements such as *InferenceEngine*, *InferenceRule*, and *Language* (see Figure 2) are used in all Inference Web applications and they are collected in a core node. The core node is publicly available to all other IWBase nodes. The current demonstration registrar for the core node is available at: <http://inferenceweb.stanford.edu/iwregistrar/> and is one example core node. The core node architecture is convenient when there is one set of entries that describe the main meta information about the common engines, their rules, and representation and reasoning languages. Empirically, we have found our current uses of Inference Web benefit from such a core node. Domain ontologies and their related meta information can vary widely from project to project and thus projects benefit from having project-specific domain nodes. Just as the notion of “upper ontologies” is both popular and contentious, we anticipate some upper level ontologies to emerge that are popular enough and reused enough that Inference Web users will find that it may be beneficial to have these included in the core node registry. Alternatively, a user can have nodes owning entries for relevant

ontologies associated with domain nodes where the user is a node member as described below in this section. We expect these decisions to evolve with usage.

Private and public nodes. IWBase relies on typical internet and web security apparatus, i.e., firewalls, for protecting access to node metadata. A public node implements no restrictions for web users or agents to access the node's contents. If any kind of restriction is enforced, the node is private. Other than the core node that is always public, nodes can be either public or private.

The IWBase architecture also specifies some services supporting the collaboration between the nodes described as follows:

Domain node–domain node associations. Inter-domain communications are accomplished by configuring **association links** between nodes. An association link from node *B* to node *A* means that services provided by node *B* are based on meta-information coming from both *B* and *A*. To illustrate the usefulness of association links, assume a link from a *vehicles* node to *cars* node. Thus, automated services for browsing the registry of *vehicles* would list elements registered both in *vehicles* and *cars*. In this case, *vehicles* could leverage the relevant list of sources about cars available in *cars*. Moreover, *vehicles* could set association links to other domain nodes, such as *bicycles*, increasing the amount of meta-information available to proofs. Association links increase elements' re-usability, eliminate data duplication among domain nodes, and help prevent inconsistency.

Mirrored core registry. Information in the core node registry is required for running services for a domain node, such as proof generation services. In certain situations, however, node services may be unable to access the core node, i.e., when a domain node has no internet connectivity. In other situations, although a domain node may have internet connectivity, the cost of getting information through the web can be too expensive or unreliable, and having a local copy of the core registry reduces the internet traffic. If any of these situations occurs, domain node administrators have the option of copying the content of the core registry to a local filesystem and setting the domain node configuration file to point to the local copy of the core registry, referred to as the *mirrored core registry*. The currency of the mirrored core registry depends on how often the domain node administrator updates the copy. Core node elements are only registered and maintained by the core node registrar and elements in the mirrored core registry are not supposed to be updated.

Core node–domain node associations. Every domain node must have an association with either the core node or a mirrored core registry. The association is required since some classes of provenance elements supporting services in the domain node, such as *Language* and *Inference Engine*, are available only in the core registry. Using the configuration file, a domain node administrator can choose which core registry to use.

5 Discussion and Status

In this section, we will return to our goals and summarize our contributions with respect to those goals. Our approach is distributed in that it supports a number of IWBase nodes. The IWBase core node contains information common to applications using proof-related metadata. The registry for the core node could be either local or remote. IWBase may have multiple domain nodes that contain information specific to any particular application. There is no limit to the number of domain nodes and they can be physically located on any machine. IWBase administration is also distributed since node administrators may make the decisions about who can make updates by registering other members and administrators.

The scalability of our approach is supported from a few perspectives. We have utilized a database back-end for IWBase so that we may leverage the scalability and efficiency of DBMSs for querying and updating databases. We also do not require registration of all meta data in our repositories; as long as meta data is accessible, that is all that is required. For example, in our joint effort with IBM to explain text analytics, we found the sheer volume of information to be enough to drive us to use a special purpose repository that is tuned to the needs of text extraction. This database is accessed when PML proofs are generated so that provenance information is available for explanations on demand but is not pro-actively registered in IWBase for all possible content. Additionally, when we designed the Proof Markup Language, we chose to represent it in OWL and stayed within the description logic (DL) portion of the language. Thus, when we rely on reasoning with the OWL representations, we can claim the same efficiency of reasoning that OWL-DL can claim. OWL-DL was designed specifically to be a language for which efficient reasoners may be implemented.

The IWBase infrastructure includes tools to help minimize maintenance effort. This is intentional since we understand that the quality of explanations will be related to the quality and recency of the provenance metadata. We include a registration service that supports automatic registration of sources. We also have included many reasoners and inferences in the core node in the hopes that new registrations of components will be more likely to require identification of registered items that may be reused rather than requiring original specifications of new inferences. We also provide checking services to identify if PML dumps are consistent in their usage with the inference rule registrations that they claim to contain⁸. The IWBase solution addresses privacy using internet and web security apparatus as described in Section 4.2. Nodes can be private or public and access to update the IWBase requires registration. While this solution has been adequate to date, we do anticipate augmenting our privacy solution in order to address more of the concerns of the intelligence community.

The primary contributions of our work with respect to provenance metadata infrastructure for explanations revolves around specifications, infrastructure, and tools. We have gathered requirements from a broad diversity of users from

⁸See <http://iw.stanford.edu/infmldescription.html> for more information.

many perspectives - backgrounds, usage, sophistication, etc. We have collected the requirements in one place and have provided a description of a solution architecture that meets these requirements. The solution is not only described in papers but is also implemented in a freely available, open-source implementation distributed through the Inference Web website⁹ and soon through SemWebCentral¹⁰. This implementation can serve as infrastructure for other's explanation and trust efforts, or if required, it could serve as an operational specification and starting point for someone to generate their own customized solution. Our approach is also modular and stresses interoperability so one could, for example, use our IWBase core node, some existing domain nodes, some new internal domain nodes, our proof generation services, and an internal interface/browser interface.

IWBase, as part of the Inference Web, is gathering users and is currently integrated with a large text analytics tool suite - UIMA. It has been integrated with the JTP hybrid reasoning engine and is being used by a number of projects that require explanations of hybrid reasoners - in particular in the CALO project of the DARPA PAL program. As a result, it is explaining results from applications such as ISI's Ariadne system [7] where it must explain where information came from after web scraping is done and how queries are broken up into smaller queries that can answer a larger question. It is also being integrated with SRI's SPARK system [10] where it needs to explain task-oriented processing. It is also being used by the KANI project within ARDA's NIMD program which means it needs to range from explaining text analytics to temporal reasoning to context reasoning with hypotheticals.

Work is proceeding to both broaden and deepen IWBase capabilities. Since explanations are improved when metadata information is more current and more complete, we continue to improve the registration services. Currently IWBase has services providing automatic support for maintaining node elements. There are still open issues concerning automated inclusion of attributes to sources. For instance, an agent needs to select and add authors to a registered document but it may have difficulty since it is required to select the appropriate metadata corresponding to the authors. We are also adding an automatic configuration service that helps users register their question answering components.

We are also working to support more flexible security strategies for accessing domain nodes behind firewalls. For instance, some organizations may require their domain nodes to be behind a firewall but they still want selected external users to access parts of their registries.

6 Related Work

When we first specified our metadata registry requirements, we did a literature search and did not find any existing solution that we believed could meet our requirements. Our initial approach was to extend Stanford's WebBase [6]

⁹<http://iw.stanford.edu>

¹⁰<http://www.semwebcentral.org/>

to register proof-related metadata. We learned from our own work and that of Heery [5] that accessing metadata can be quite complex. IWBase requires elements to be grouped and indexed by their types. This led us away from an approach such as WebBase that indexes through a uniform attribute. We decided to design an solution for registering metadata that was based on a combination of a web repository and of a database system.

An agent's ability to share metadata with other agents is a major reason for IWBase to store proof-related metadata in the registry as OWL files. In fact, the lack of a standard format for metadata was a concern in previous efforts to build metadata registries [1].

The Dublin Core Metadata Initiative (DCMI)¹¹ is probably the best-known collection of bibliographic metadata. In the PML provenance element taxonomy, *Document* is the element representing Dublin Core metadata. *Document* is the only non-extensible PML element and it is expected to conform with DCMI. Furthermore, DCMI does have a registry¹² for storing and maintaining DCMI metadata and the registry shares a number of functional requirements as described in [4] with IWBase. For instance, the IWBase and DCMI Registry are expected to store reliable trusted information, to provide authoritative definitions, and to manage the evolution of their elements. One difference between IWBase and DCMI Registry is that IWBase includes instances of metadata types.

There are many registries of domain-specific metadata available on the web. For example, the DoD Metadata Registry and Clearinghouse¹³ is a domain-specific registry of software related metadata for use by the U.S. Department of Defense. Other government agencies throughout the world have their own registries. In the future, we expect applications that use government registries such as this and/or maintainers of this registry to use information from IWBase.

7 Summary

This paper presents IWBase as a distributed, scalable, maintainable provenance metadata infrastructure with a privacy strategy. Provenance metadata in IWBase may help users to understand and trust answers from applications and services on the web. We gathered a set of requirements for provenance metadata for web applications and designed and implemented a solution that addressed these requirements. Specific contributions from this effort include:

- A distributed registry of provenance metadata with a taxonomy of provenance elements specified in an interoperable proof interlingua (PML) based on the W3C Recommendation for a Web Ontology Language(OWL).
- A registrar that provides support for the registration and maintenance of provenance metadata. For humans, the registrar provides its functional-

¹¹<http://dublincore.org/>

¹²<http://dublincore.org/dcregistry/>

¹³<http://diides.ncr.disa.mil/mdregHomePage/mdregHome.portal>

ities through an interactive interface. For computer agents, it provides automated services.

- Services for generating proofs in PML, checking those proofs against inference specifications, abstracting those proofs into understandable explanations, and computing trust values for answers supported by those proofs and explanations. We are able to provide these services because of the infrastructure for provenance metadata.

IWBase is being used to support explanations of question answering applications that range from text analytics and aggregation to hybrid theorem proving in academic, industrial, and government settings.

Acknowledgments

The authors would like to thank the following programs that funded portions of this work: DARPA's DAML and PAL programs and ARDA's NIMD program. Additionally, we gratefully acknowledge valuable comments from Jose-Luis Ambite, Priyendra Deshwal, J William Murdock, and Ilya Zaihrayeu.

References

- [1] Christophe Blanchi and Jason Petrone. Distributed Interoperable Metadata Registry. *D-Lib Magazine*, 7(12), December 2001.
- [2] D. Ferrucci and A. Lally. UIMA: An Architectural Approach to Unstructured Information Processing in the Corporate Research Environment. *Journal of Natural Language Engineering*, 10(3/4):327–348, June 2004.
- [3] Richard Fikes, Jessica Jenkins, and Gleb Frank. JTP: A System Architecture and Component Library for Hybrid Reasoning. Technical Report KSL-03-01, Knowledge Systems Laboratory, Stanford University, Stanford, CA, USA, 2003.
- [4] Rachel Heery. Draft DCMI Open Metadata Registry Functional Requirements, October 2001. http://dublincore.org/groups/registry/fun_req_ph1-20011031.shtml as accessed on 6 Nov 2004.
- [5] Rachel Heery and Harry Wagner. A Metadata Registry for the Semantic Web. *D-Lib Magazine*, 8(5), May 2002.
- [6] Jun Hirai, Sriram Raghavan, Hector Garcia-Molina, and Andreas Paepcke. WebBase : A repository of Web pages. *Computer Networks*, 33(1-6):277–293, May 2000.
- [7] Craig A. Knoblock, Steven Minton, Jose-Luis Ambite, Naveen Ashish, Ion Muslea, Andrew G. Philpot, and Sheila Tejada. The ariadne approach to web-based information integration. *International Journal of Cooperative Information Systems (IJCIS), Special Issue on Intelligent Information Agents: Theory and Applications*, 10(1/2):145–169, 2001.

- [8] Deborah L. McGuinness and Paulo Pinheiro da Silva. Trusting Answers on the Web. In Mark T. Maybury, editor, *New Directions in Question Answering*, chapter 21. AAAI/MIT Press, October 2004.
- [9] Deborah L. McGuinness and Frank van Harmelen. OWL Web Ontology Language Overview. Technical report, World Wide Web Consortium (W3C), December 9 2003. Proposed Recommendation.
- [10] David Morley and Karen Myers. The SPARK Agent Framework. In *Proc. of the Third Int. Joint Conf. on Autonomous Agents and Multi Agent Systems (AAMAS-04)*, 2004.
- [11] J. William Murdock, Paulo Pinheiro da Silva, David Ferrucci, Christopher Welty, and Deborah L. McGuinness. Encoding Extraction as Inferences. In *Proc. of AAAI Spring Symposium on Metacognition on Computation*, pages 92–97, Stanford University, USA, 2005. AAAI Press.
- [12] Paulo Pinheiro da Silva, Deborah L. McGuinness, and Richard Fikes. A Proof Markup Language for Semantic Web Services. *Information Systems*, 2004. (to appear).
- [13] Jörg H. Siekmann, Christoph Benzmüller, Vladimir Brezhnev, Lassaad Cheikhrouhou, Armin Fiedler, Andreas Franke, Helmut Horacek, Michael Kohlhase, Andreas Meier, Erica Melis, Markus Moschner, Immanuel Normann, Martin Pollet, Volker Sorge, Carsten Ullrich, Claus-Peter Wirth, and Jürgen Zimmer. Proof Development with OMEGA. In *Proceedings of 18th International Conference on Automated Deduction*, volume 2392 of *LNCS*, pages 144–149, Copenhagen, Denmark, July 2002. Springer.
- [14] Ilya Zaihrayeu, Paulo Pinheiro da Silva, and Deborah L. McGuinness. IWTrust: Improving User Trust in Answers from the Web. In *Proc. of iTrust2005*, Rocquencourt, France, 2005. Springer. (to appear).