

Position Paper – W3C Workshop on RDF Next Steps: OMG Ontology PSIG

Elisa Kendall¹, Roy Bell², Roger Burkhart³, Manfred Koethe⁴, Hugues Vincent⁵, and Evan Wallace⁶
Object Management Group (OMG) Ontology Platform Special Interest Group (PSIG)

¹Sandpiper Software, Los Altos, California, USA, ekendall@sandsoft.com

²Raytheon Company, Network Centric Systems Division, Fort Wayne, Indiana, USA, Roy_M_Bell@raytheon.com

³Deere & Company, Moline, Illinois, USA, BurkhartRogerM@JohnDeere.com

⁴88solutions Corporation, Newton, Massachusetts, USA, koethe@88solutions.com

⁵Thales, Neuilly-sur-Seine Cedex, France, hugues.vincent@thalesgroup.com

⁶National Institute of Standards and Technology (NIST), Manufacturing Systems Integration Division, Gaithersburg, Maryland, USA, ewallace@cme.nist.gov

Background

The Object Management Group (OMG) is an international, open membership, not-for-profit computer industry consortium, with liaison relationships with the World Wide Web Consortium (W3C) and International Standards Organization (ISO), among others. Any organization may join OMG and participate in the standards-setting process. From the OMG bylaws, organizational purposes include:

- Promoting frameworks for compatible and independent development of applications
- Enabling coordination among applications across heterogeneous networked systems in a multinational, multilingual environment
- Adopting a core of commercially available specifications of these frameworks, and
- Promoting international market acceptance and use.

In other words, one of the primary organizational goals of OMG is to develop standards for application programming interfaces for a wide variety of systems, networks, platforms, industries, and domains.

The OMG's Ontology Definition Metamodel (ODM) standard was adopted in 2006, finalized in 2008, and is planned for transition to an ISO standard in the coming year [1]. It includes a family of metamodels (models of the abstract syntax) for the Resource Description Framework [2], the Web Ontology Language [3], ISO Common Logic [4], ISO Topic Maps [5], and several UML (Unified Modeling Language [6]) profiles to enable use of UML tools for RDF vocabulary, OWL ontology, and topic map development. It also contains mappings (some partial and some complete) to support model transformations from one representation paradigm to another. It bridges standards and best practices from several communities and provides a foundation for defining, developing, and managing information models as independent but equal components of larger systems.

Given that the standard has matured considerably, the OMG Ontology Platform Special Interest Group (PSIG) is beginning to focus on next steps – interfaces with other OMG standards,

extensions to the ODM to support OWL 2, new standards that complement the work completed in the ODM, and so forth. Our intent with this position paper is to share some lessons learned in development of metamodels for RDF and RDF Schema as well as emerging requirements identified by members of the PSIG for further work that echo or augment some of the concerns identified in the Call for Participation.

Challenges

As OMG members and other interested parties have developed tools and services based on the ODM, and as our knowledge of requirements for further development grows, several critical issues with the current Semantic Web languages and language architecture have become apparent. These include:

1. There is no independent specification of the common elements of RDF vocabularies and OWL ontologies that connect them to the web. These are implicit, of course, and references to relevant specifications are embedded in the documents that comprise the current RDF and OWL 2 specifications, but much is left to the reader. Some of these elements, including documents, local names, namespaces, namespace definitions, and IRIs, could be collected in a common specification that both languages reference, as has been done in the RDFWeb package of the ODM RDF Metamodel. Currently some are re-defined independently in OWL 2, making it difficult to develop tools and APIs with a cohesive, common architecture. Another opportunity would be a common specification for literals and built-in datatypes (and their facets), rather than embedding them in the OWL 2 syntax specification.
2. Namespace organization in RDF, which may have developed based on historical usage, is tangled and made it impossible for separation of an RDF metamodel from an RDF Schema metamodel in the ODM. Examples include the definition of `rdfs:Resource` and `rdfs:Literal` in the RDF Schema namespace, but `rdf:Property` in the RDF namespace. Containers and collections also exhibit namespace entanglement. This latter entanglement is particularly irksome in that a circular relationship between independent metamodels for RDF and RDF schema would be required in order to maintain namespace separation, which is not permitted in UML.
3. Vocabulary and ontology alignment and mapping is a very high priority for the OMG community, which is interested in mapping the semantics of UML and domain specific language models, for example, mapping process models to service models, and so forth. To that end, named graphs and the related capabilities defined in *Named Graphs, Provenance and Trust* [7] were included in the ODM RDF metamodel, and we would like such facilities to be considered seriously for incorporation in any extension to RDF.

Importance of a Cohesive API Architecture

As tool vendors and users, systems integrators, and others building applications that leverage Semantic Web technologies, we are at a tremendous disadvantage without standardized APIs for

application development and knowledge base accessibility. The APIs available to date represent a significant effort on the part of a number of very thoughtful researchers, as well as contribution from broader communities of interest over the course of several years, but typically support either RDF or OWL, and in the case of OWL, OWL DL but not OWL Full. In one way or another, they lack features, error management mechanisms, change management support, or other capabilities that commercial vendors and users expect. None are well integrated with one another, and architectures in particular are not necessarily modularized in ways that would facilitate enterprise-level software development. While we anticipate that the recent launch of Epimorphics may lead to better commercialization of the Jena API and tools, a similar effort for the OWL API has yet to emerge, and the two are not yet tied together.

An architecture that supports interfaces for multiple languages and tools, that is designed in a cohesive and modular way, with an eye towards keeping typical integration problems at bay is essential to move Semantic Web software development out of the research realm and towards broader commercial adoption. Without this, a number of us, and colleagues in the Protégé group at Stanford, will continue to face issues in attempting to build tools that support multiple communities with varying language expressivity requirements. One example that comes quickly to mind is that the various APIs use different java libraries that have evolved independently. (In fact, at one point during an early integration effort, Pellet libraries clashed with OWL API libraries, which clashed with Eclipse libraries, which in turn all clashed with UML tool libraries – which we clobbered in order to get “something” to work.) This lack of synchronization, which is natural in any evolving effort, often leads to integration conflicts. A cohesive, coherent approach to integration and interoperability must be supported by the architecture, and in this case, something like an OSGi¹ approach as a platform-specific requirement on the Java API architecture.

While these are not necessarily requirements for next generation RDF language development, we believe that having a standard, integrated set of APIs for RDF and OWL is crucial to language and downstream technology adoption. And while APIs are not traditionally developed by the W3C, a coordinated effort by a combination of API developers and language authors, at W3C and OMG, may be the right next step.

Current Activities at OMG

OMG has a number of efforts related to RDF currently underway.

MOF 2 RDF

OMG has issued an RFP for a specification of structural mapping from the Meta Object Facility to RDF entitled “MOF to RDF Structural Mapping in support of Linked Open Data”². The intent

¹ <http://www.osgi.org/About/Technology>

² <http://www.omg.org/cgi-bin/doc?ad/09-12-09>

of this RFP is to provide a simple standard mapping from OMG's metamodeling language, the Meta Object Facility, to RDF such that any MOF compliant metamodel, model based on such a metamodel, or instance data based on such a model can be easily published as linked data. Since UML and other languages standardized by OMG are based on MOF, this capability would enable nearly any model related artifact compliant with OMG standards to be published as web resources in a standard form. This would include UML models, BPMN process models, models based on UML profiles, and models based on other metamodel standardized by OMG. The goal of this effort is to make the many information resources that have been created based on OMG standards more accessible and sharable.

Initial proposals for this specification are due to the OMG by 24 May 2010 and the specification should be voted on in final form in December 2010.

API 4 KBs

Thales Group has been leading an effort at OMG to draft a set of requirements (a Request for Proposal or RFP in OMG parlance) for a future OMG specification to standardize an API for knowledge bases with reasoning support using OWL and/or RDF. This RFP is currently scheduled to be issued on June 25 with initial submissions due in February 2011. The resulting API would be required to support:

- Retrieving content such as ontologies, classes, properties or instance data,
- Retrieving content related to other content such as all subclasses of a given class,
- Modifying or creating new content,
- Deleting instances,
- Querying for expressive power and capabilities of the content and reasoning system,
- Requesting execution of reasoning tasks,
- Performing queries based on class descriptions and related reasoning (and other similar queries involving reasoning), and
- Loading, deleting, and exporting ontologies.

The RFP is seeking support for the following platforms

- JAVA,
- WSDL such that a KBI can be accessed from a WS-I Basic Profile 1.1 compliant client,
- REST styled http client access.

Suggestions for additional requirements and platforms are very welcome.

Request for Community Involvement

The RDF Next Steps workshop specifically enumerates development of APIs as one of the topics of interest to this community. While the OMG is an appropriate place for standardization of APIs, such a specification would only be relevant if key stakeholders from the RDF and OWL communities were on board with the effort and would be willing to provide support for the interface in semantic web tools and toolkits. We would appreciate feedback from the RDF community as to their interest in standard API functionality described above for the API4KB effort, in particular. If there is interest, then we would also like to know:

- Are the requirements appropriate and practical?
- Who would be willing to participate in developing the specification?
- Is OMG an acceptable venue for such work (and if not, what would be)?
- Would semantic web tool builders build support for this API into their tools, and if so what schedule would they want to see for development of an API specification?
- Any other feedback concerning API standardization or other OMG efforts mentioned in this paper.

References

1. Ontology Definition Metamodel (ODM), Version 1.0 Specification, Object Management Group, Inc., Needham, MA, May 2009. Available at: <http://www.omg.org/spec/ODM/1.0/>.
2. Dan Brickley and R. V. Guha. RDF Vocabulary Description Language 1.0: RDF Schema, W3C Recommendation, World Wide Web Consortium, Amsterdam, The Netherlands, 10 February 2004. Latest version is available at <http://www.w3.org/TR/rdf-schema/>.
3. Boris Motik, Peter F. Patel-Schneider, and Bijan Parsia, eds., OWL 2 Web Ontology Language Structural Specification and Functional-Style Syntax, W3C Recommendation, World Wide Web Consortium, Amsterdam, The Netherlands, 27 October 2009. Latest version is available at <http://www.w3.org/TR/owl-syntax/>.
4. Harry Delugach, editor. ISO/IEC 24707 Information technology – Common Logic (CL) – A Framework for a Family of Logic-Based Languages. The formally adopted ISO specification is available at [http://standards.iso.org/ittf/PubliclyAvailableStandards/c039175_ISO_IEC_24707_2007\(E\).zip](http://standards.iso.org/ittf/PubliclyAvailableStandards/c039175_ISO_IEC_24707_2007(E).zip).
5. ISO 13250 Topic Maps – Data Model and XML Serialization.
6. Unified Modeling Language™ (UML®) Infrastructure and Superstructure Specifications, Version 2.1.2, Object Management Group, Inc., Needham, MA, November 2007. Latest versions of this and related UML specification components is available at <http://www.omg.org/spec/UML/2.1.2/>.
7. <http://www2005.org/cdrom/docs/p613.pdf>