Metadata Management in S-OGSA

Oscar Corcho, Pinar Alper, Paolo Missier,
Sean Bechhofer, Carole Goble, and Wei Xing

School of Computer Science, University of Manchester
Oxford Road, Manchester M13 9PL, United Kingdom
{ocorcho,penpecip,pmissier,seanb,carole,wxing}@cs.man.ac.uk

Abstract. Metadata-intensive applications pose strong requirements for metadata management infrastructures, which need to deal with a large amount of distributed and dynamic metadata. Among the most relevant requirements we can cite those related to access control and authorisation, lifecycle management and notification, and distribution transparency. This paper discusses such requirements and proposes a systematic approach to deal with them in the context of S-OGSA.

1 Introduction

Metadata can be attached to many different types of objects, available in different formats and locations, and can be expressed in a wide range of languages (e.g., natural language, lists of terms, formal languages) and using a wide range of vocabularies (e.g., keyword sets, concept taxonomies). Depending on these aspects, many authors use the terms semantic versus non-semantic, or rich versus non-rich metadata.

Many technologies are available to manage metadata, such as Jena, Sesame, Boca, Oracle-RDF, Annotea, Technorati, etc. They give a good level of support for current metadata-intensive applications. However, they fall short for some of the metadata management requirements that new applications are posing, such as metadata distribution, lifecycle management and authorisation.

This paper starts by providing a set of requirements for advanced metadata management (Section 2), which characterise metadata-intensive applications in domains like bioinformatics, social sciences, engineering, market analysis, etc. Existing approaches and technologies do not meet all of the previous requirements [1]. Hence we make a proposal for managing metadata as first-class resources in distributed systems, so that we can deal with the previous requirements (Section 3). Finally we provide some conclusions and future work.

2 Metadata Management Requirements

In this Section we describe some of the advanced requirements for metadata management that are needed in some metadata-intensive applications.

Metadata should be stored and accessible in a distributed manner. Many systems will need to integrate distributed metadata (and ontologies),

which may be supplied by multiple parties and with different technologies. Most of the current metadata management systems provide repositories designed for centralized use, with metadata consumers and producers acting as local-access clients using specialized APIs. Remote distributed access to metadata and ontologies has received little attention, causing each system to depend on one particular technology, thus reducing interoperability.

A service-oriented approach to metadata and ontology access could improve this situation. However, this should not be done by simply wrapping current metadata management systems with web services that replicate their local APIs, as described below.

Metadata should be accessible with technology-independent service-oriented protocols. Metadata can be available in multiple forms (in RDF, as social tags, in natural language, etc.). Hence we require systems that enable the integration and exploitation of this heterogeneous metadata.

Even in the case of a single form of metadata (e.g. RDF(S)) each technology uses a different abstraction and means to handle it. There is a clear need for a common abstraction for metadata, which would enable the systematic development of services that combine other services and applications with varying levels of semantic capabilities for dealing with and interpreting metadata.

Metadata should evolve together with the resources that it describes and the vocabularies that it uses. Metadata evolves due to different reasons. However, the dynamics of metadata is poorly supported by existing technologies, although some work has been done with respect to the semantics of change and its propagation to the related metadata [2,3].

Metadata should be maintained up-to-date in a cost-effective manner. This includes maximising the automation of different aspects of the knowledge life-cycle, managing the evolution and change of metadata and knowledge models in distributed contexts, and synchronising adequately the evolution of all these related entities by means of notification mechanisms.

Metadata access should be controlled. Metadata owners may want to define the conditions and rules under which others can access their metadata. However, existing technologies vary in their support and granularity for security, which normally means access-control only. For instance, Sesame provides built-in user/role-based access control per repository, while Jena has no built-in access-control support since it assumes that this can be provided by the underlying database technology. Moreover, all these access control mechanisms have their own specific APIs and conventions, which creates overhead at the client layer (need for multiple usernames and passwords, or multiple sign-ons).

Securing message exchanges, authentication and access-control are well-studied problems in distributed environments, with standards¹ and open-source reference implementations² for global user identification, single sign-on, communication encryption and representation and decision of resource-sharing policies.

¹ E.g., XACML: www.oasis-open.org/committees/xacml/, WS-Security: www.oasis-open.org/committees/wss/

² E.g., Globus GSI: <http://www.globus.org/toolkit/docs/4.0/security/>

3 Metadata Management in S-OGSA

In this Section we describe the approach for metadata management proposed in the context of the Semantic-OGSA (S-OGSA) architecture [4]. S-OGSA is an extension of the Open Grid Service Architecture [5] for the development of distributed applications that need to use explicit and distributed metadata.

The S-OGSA model is driven by the principle that *associations between resources and their metadata should be first-class resources that can be distributed and managed in a service-oriented manner*. This gives support to the first two requirements from Section 2, and is the basis for giving support to the rest.

The term Semantic Binding (SB) [6] denotes this new type of resource, which constitutes the core of S-OGSA. Semantic Bindings represent associations between any set of resources and any set of knowledge entities (e.g., ontologies, rule sets, controlled lists of terms). These associations contain metadata, which *can be encoded in different formats*: RDF, natural language, as social tags, etc. Besides, *different Semantic Bindings can describe the same set of entities*: different tools or persons may create different annotations for the same resource.

The following subsections describe the capabilities associated to Semantic Bindings and how they support our metadata management requirements.

3.1 Semantic Binding Capabilities

In [4] we describe the model used to represent Semantic Bindings, expressed as an ontology that extends the Grid ontology described in [7]. The main properties of a Semantic Binding are the set of resources to which it refers (that is, the resources for which it contains metadata), the set of knowledge entities that the metadata is based on, and the actual metadata that they store. Other properties (state, creation time, last modification time, etc.) are stored and used for managing lifetime, notification and authorisation mechanisms.

Figure 1 shows the basic operations provided by the Semantic Binding service suite (the *Semantic Binding Service*, its corresponding *Factory* and a *Metadata Service* that unifies the metadata stored by several Semantic Bindings).

- *Create*. It creates a Semantic Binding, given the described resources, the Knowledge Entities used for the description, and the metadata to store.
- *Update Resource and Knowledge Entity References*. They allow managing the references to Resources and Knowledge Entities of the Semantic Binding.
- *Update Semantic Binding Content*. It updates the metadata stored in the Semantic Binding, due to its reannotation or curation.
- *Destroy*. It destroys the Semantic Binding, together with its content, immediately or at a scheduled point in time.
- *Archive*. It archives the Semantic Binding content so that it is not active but can be retrieved in case that it is needed later (e.g., for provenance).
- *Query*. It executes a query over the metadata stored by the Semantic Binding. Queries will be sent in a query language that the Semantic Binding supports, and can take into account the knowledge entities to which the Semantic Binding refers or not.

Fig. 1. Functionality of the Semantic Binding Service

This suite is supported by the S-OGSA reference implementation³, which complies with WS-Resource Framework [8] (WSRF). It can be deployed on the Globus Toolkit 4 platform⁴ and in Apache Tomcat⁵.

Next we describe how we deal with the rest of requirements, namely lifetime management, metadata change notifications and controlled access to metadata.

3.2 Semantic Binding Lifetime

In some applications, metadata is a dynamic entity subject to frequent changes: resources and knowledge entities that they refer to can evolve, become suddenly unavailable, be destroyed, etc. Some of these changes may cause metadata to become invalid. Other changes may not have an influence on metadata validity (e.g., the removal of a concept in an ontology for which there are no instances in the stored metadata). Finally, metadata may become invalid after a given period of time, when a new annotation tool has been made available and the resource has to be reannotated, when a metadata curation process is in place, etc.

In order to deal with all these changes in a principled way, S-OGSA defines SBs as stateful resources with a defined lifetime and identifies the states and state transitions that a SB can go through throughout its lifetime. The corresponding state diagram is presented in the next subsection.

Figure 2 shows the state diagram associated to an SB, which includes fundamental states and state transitions, and the external events that cause the transitions. The SB lifetime specification extends WS-ResourceLifetime, a part of WSRF that standardizes the way that resources are destroyed, and defines resource properties for the inspection and monitoring of a resource lifetime. While

³ Available at the OntoGrid CVS.

⁴ <http://www.globus.org/toolkit/>

⁵ <http://tomcat.apache.org/>

WS-ResourceLifetime is focused exclusively on resource destruction, we extend it to include any life-changing event that may affect the validity and updates of an SB. Furthermore, the basic state machine presented here can be extended with sub-states if needed in a certain application.

The explanation of the state transition diagram is as follows. When it is first created, a Semantic Binding SB is in the *Valid* state. We denote with Res_{SB} and KE_{SB} , respectively, the set of Resources and Knowledge entities that are part of the association, and with $content_{SB}$ the metadata payload within SB .

Fig. 2. State transition diagram for a generic Semantic Binding

State transition events are of the following types:

- Changes in the described resources, denoted by $Res_{SB} \rightarrow Res'_{SB}$.
- Changes in the Knowledge entities, i.e., $KE_{SB} \rightarrow KE'_{SB}$
- Updates to the SB content: $content_{SB} \rightarrow content'_{SB}$.

Note the Resources and Knowledge entities can also be destroyed: $Res_{SB} \rightarrow \emptyset$, $KE_{SB} \rightarrow \emptyset$. In addition to these external events, a content expiration date can also be associated to an SB, so that it gets stale upon expiration.

For a *Valid* SB, these events cause its transition to either one of two possible *Validate* states, *Validate Res* and *Validate KE*. These are interim states in which the SB may be invalid, and is awaiting re-validation. A re-validation process, either manual or automated, is any procedure that updates any or all of Res_{SB} , KE_{SB} , or $content_{SB}$, and which results in a decision as to whether the updated entities represent a new valid combination⁶.

- For a *ValidateRes* SB, such procedure determines whether the existing meta-data can be associated to the new Resources, and provides an update to the references in SB to Res'_{SB} . For example, following a change in a workflow

⁶ In both cases, the SB goes back to the *Valid* state in case of successful validation, and to *Invalid* otherwise.

that is described with a piece of metadata, the procedure determines whether the same metadata can be associated to the new workflow.

- For a *ValidateKE* SB, the problem is to determine whether the new ontology can still be used to interpret the old metadata.

Finally, the *Archived* state indicates that a SB is still available for inspection.

3.3 Notification of Semantic Binding Changes

Requirement 3 suggests that the metadata-aware services that use SBs should be informed of state change for those SBs. For this purpose, S-OGSA defines a set of notification mechanisms based on WS-Notification. S-OGSA proposes using a set of pre-defined topics associated to the changes described above, which any consumers can subscribe to.

Services that receive any of these notifications will decide, as part of their business logic, how to react to the changes that they are notified about. The S-OGSA specification does not enforce any specific type of behaviour. However, as part of our future work, we are designing a service, called *SB housekeeping service*, which monitors SB lifetime by subscribing to all their topics. This service will be responsible for activating application-dependent re-validation procedures (validation of the SB content, triggering of re-annotation processed, etc.).

3.4 Security over Semantic Bindings

The last requirement suggests the need to control the access to metadata in more fine-grained ways than what it is currently done with current technology.

S-OGSA provides the framework needed to support security in metadata management. Metadata is a first-class resource; hence standard security mechanisms can be applied to metadata in the same way as it is done with other resources, including the possibility of specifying and enforcing access control policies over each SB. Besides, since the reference implementation of S-OGSA can be deployed on top of Globus Toolkit 4, its associated security mechanisms, such as Globus GSI, can be also applied, ensuring both message and transport level security, based on standard X.509 end entity and proxy certificates.

With this framework, it is possible to allow or deny access to the different annotations of an object, stored by different SBs, based on the users or groups that have created them and on their access control policies.

4 Conclusions and Future Work

In this paper we have described the requirements posed by some of the existing and the envisaged metadata-intensive distributed applications:

- Applications may require their explicit metadata to be encoded in multiple forms, supplied by multiple parties and coming from different contexts. Moreover, metadata may need to be in different physical locations.

- Applications may use heterogeneous metadata storage and query technologies, each of which has specific advantages over the others (e.g. Sesame is good at query performance, Jena has a rich API, etc.). To ease metadata sharing, common technology-independent means for metadata access (meta-models and protocols) become necessary.
- Metadata normally evolves, either because of new metadata generation means or because of the evolution of the resources or knowledge entities that it refers to. Adequate means to deal with this evolution have to be made available.
- Access to metadata may need to be secured with different levels of granularity and different access control policies.

We have shown how the approach followed in S-OGSA complies with the previous requirements without the need to create expensive and difficult-to-maintain ad-hoc solutions for each of them. Basically, S-OGSA proposes to treat metadata as a first-class resource (Semantic Binding) in the application, so that standard resource management and sharing solutions used in the context of distributed applications can be easily applied to it. This includes aspects like service orientation, lifecycle management, security, etc.

The S-OGSA approach is being used in the development of different types of applications, all of which are characterised by being metadata-intensive and by needing support for some of the previous requirements. The S-OGSA reference implementation is being or will be used in several system prototypes: a satellite quality image analysis system [9], an information service for the EGEE Grid [10], and a coral bleaching alert system (Semantic Reef) [11].

Our future work will be devoted to improve our S-OGSA reference implementation, which includes the Semantic Binding Service suite presented in Section 3. We will consider the additional requirements that the early adopters of S-OGSA are providing and create an implementation with industrial standards.

Among the aspects that will be improved, we can cite the following:

- *Security*. We will provide a set of pre-defined security configurations that cover the most common metadata management security aspects.
- *Naming*. Our current implementation uses WS-Addressing EndPoint References (EPRs) to identify Semantic Bindings and the Resources that they refer to, and URIs to identify Knowledge Entities. However, there are more ways to identify entities in a distributed environment (URIs, ARK Identifiers, LSIDs). We will implement a metadata identification model using WS-Naming, which builds on WS-Addressing and extends it with URIs.
- *Semantic Binding Housekeeping Service*. We will build a configurable SB Housekeeping Service that gives support to the most common behaviours for metadata evolution found in applications.
- *Support for more forms of metadata*. We will give support to several forms of annotation (besides RDF), such as social tags, natural language comments, other ontology languages, etc.

Finally, as part of our future work we will also evaluate our system with respect to other metadata management systems, in terms of memory consumption, query execution performance, etc., in distributed settings.

Acknowledgements

This work is supported by the EU FP6 OntoGrid project (STREP 511513) funded under the Grid-based Systems for solving complex problems, and by the Marie Curie fellowship RSSGRID (FP6-2002-Mobility-5-006668). We also thank the other members of the OntoGrid team at Manchester for their helpful discussions: Ian Dunlop and Ioannis Kotsiopoulos.

References

1. O. Corcho, P. Missier, P. Alper, S. Bechhofer, and C. Goble, "Principled metadata management for next generation metadata-intensive systems," in *4th European Semantic Web Conference (ESWC2007)*. Submitted, Innsbruck, Austria, 2007.
2. L. Stojanovic, "Methods and tools ontology evolution," Ph.D. dissertation, Univ Karlsruhe (TH), 2004.
3. M. Klein, "Change management for distributed ontologies." Ph.D. dissertation, Vrije Universiteit Amsterdam, 2004.
4. Ó. Corcho, P. Alper, I. Kotsiopoulos, P. Missier, S. Bechhofer, and C. A. Goble, "An Overview of S-OGSA: A Reference Semantic Grid Architecture," *Journal Web Semantic*, vol. 4, no. 2, pp. 102–115, 2006.
5. I. Foster, H. Kishimoto, A. Savva, D. Berry, A. Grimshaw, B. Horn, F. Maciel, F. Siebenlist, R. Subramaniam, J. Treadwell, and J. V. Reich, *The Open Grid Services Architecture, Version 1.5*, gfd-i.080 ed., GGF, July 2006, <http://forge.gridforum.org/projects/ogsa-wg>.
6. P. Missier, P. Alper, O. Corcho, I. Kotsiopoulos, I. Dunlop, W. Xing, S. Bechhofer, and C. Goble, "Managing semantic grid metadata in S-OGSA," in *Cracow Grid Workshop 2006*. Submitted, Cracow, Poland, 2006.
7. M. Parkin, S. van den Burghe, O. Corcho, D. Snelling, and J. Brooke, "The Knowledge of the Grid: A Grid Ontology," in *Proceedings of the 6th Cracow Grid Workshop*, Cracow, Poland, October 2006.
8. K. Czajkowski, D. Ferguson, I. Foster, J. Frey, S. Graham, I. Sedukhin, D. Snelling, S. Tuecke, and W. Vambenepe, "Web Services Resource Framework (WSRF)," Globus Alliance and IBM, Technical report,, March 2005.
9. M. Sánchez-Gestido, L. Blanco-Abruña, M. de los Santos Pérez-Hernández, R. González-Cabrero, A. Gómez-Pérez, and O. Corcho, "Complex data-intensive systems and semantic grid: Applications in satellite missions," in *Proceedings of the 2nd IEEE International Conference on e-Science and Grid Computing (e-Science 2006)*, Amsterdam, The Netherlands, December 2006.
10. W. Xing, O. Corcho, C. Goble, and M. Dikaiakos, "Active ontology: An information integration approach for highly dynamic information sources," in *Submitted to ESWC2007*, May 2007.
11. T. S. Myers, "The Semantic Reef: Managing complex knowledge to predict coral bleaching on the great barrier reef." in *In proceedings of AusGrid 2007*, 2007, <http://eprints.jcu.edu.au/1131>.