Towards a Reference Architecture for Service-Oriented Cross Domain Security Infrastructures

Wen Zhu[1], Lowell Vizenor[2], and Avinash Srinivasan[3]

[1] Alion Science and Technology, Vienna, VA, USA
wzhu@alionscience.com
[2] National Center for Ontological Research, Reston, VA, USA
[3] George Mason University, Fairfax, VA, USA
asriniv5@gmu.edu

Abstract. Today's Cross Domain Communication (CDC) infrastructure largely consists of guards built to vendor specifications. Such an infrastructure often fails to provide adequate protections for CDC workflows involving Service Oriented Architectures. Focusing on the transport layer and oblivious to the context of the information exchanges, the guards often rely on rudimentary filtering techniques that require frequent human intervention to adjudicate messages. In this paper, we present a set of key requirements and design principles for a Service Oriented Cross Domain Security Infrastructure in form of a CDC Reference Architecture, featuring domain-associated guards as active workflow participants. This reference architecture will provide the foundation for the development of protocols and ontologies enabling runtime coordination among CDC elements, leading to more secure, effective, and interoperable CDC solutions.

Keywords: Cross Domain Communications, Security Guard, Workflow, Service Oriented Architecture, Reference Architecture, Ontology, Protocol.

1 Introduction

A common network security practice is to separate computer systems into secure domains or enclaves based on the classification and sensitivity of data stored and processed by these systems. Within each domain, a certain level of trust among systems is assumed. The domains are protected by Cross Domain Communication (CDC) infrastructures, which largely consist of security guards placed at the network links between two domains. These guards are responsible for enforcing security policies by inspecting and filtering information that flows between domains. However, information needed to support a mission often cut across two or more security domains. Currently, CDC flows are impeded by time-consuming release procedures that require frequent human intervention. While there have been research efforts in this area, most of them address particular aspects of CDC [1][2][3]. The lack of a comprehensive CDC framework to provide a systematic examination of CDC issues, a necessary step toward standardization, contributes to issues mentioned above.

G. Fortino et al. (Eds.): IDCS 2014, LNCS 8729, pp. 275–284, 2014.

The wide adaption of Service Oriented Architectures (SOA), and web service technologies in particular, has presented both new challenges and new opportunities in the area of CDC. It is now possible to accomplish complex workflows, carried out at the application layer across organizational boundaries with security implications. At the same time, service description metadata [4] could be used to understand the context and semantics of service interactions, and automate the enforcement of policies.

The alignment of CDC infrastructures with SOA will help extend SOA across the boundaries of security domains. In this paper, we propose a reference architecture to delineate responsibilities among various CDC participants, and describe how they interact with one another. It addresses the many facets of CDC:

1. From a workflow perspective, what role does the CDC infrastructure play in an application workflow and business process management (BPM) in particular?
2. From an information perspective, how does the CDC infrastructure interpret and act upon the information carried in CDC messages and web services in particular?
3. From a network perspective, how does the CDC infrastructure fit into the transport protocols' stack and the web services technology stack in particular?

The rest of this paper is organized as follows. We will start with a survey of current CDC solutions and highlight the key issues we seek to resolve. Then in sections 3 and 4, we introduce the reference architecture and discuss how a security ontology based on such an architecture enables coordination among CDC participants and between security domains. Section 5 identifies opportunities for standardizing CDC interactions in the forms of protocol candidates. Finally, we conclude our discussion and evaluation of an implementation approach.

2 Background and Issues

2.1 Cross Domain Solutions Today

A survey of current cross-domain security solutions revealed a number of critical issues related to application design and network infrastructure. From the perspective of mission applications design [5], current CDC solutions require mission application programs to design and implement their own individual solutions around particular guard designs, resulting in vendor lock-ins. Even then, the solutions are limited to simple cases without full-duplex architectures. Current CDC offerings include email integration, file transfer, chat, and browse down capabilities. Yet, an application workflow is likely to involve more complex interactions and different transport mechanisms, which are difficult to implement with current, inflexible CDC solutions.

From the perspective of enterprise security infrastructure, existing CDC solutions use a Transport-Oriented Guard (TOG), associated with the links between domains. TOGs monitor traffic on the links rather than at nodes. While most guards today understand XML formats and HTTP protocol, they make security decisions solely based on the bytes over the wire, without the benefit of application context. TOG creates several issues. First, the guard is required to have the highest security privilege to inspect the information flow and prevent confidentiality breach. Consequently, the guard may have to be (pre-)loaded with cryptographic keys and other sensitive infor-

mation. Contrary to the best security practices, such models of CDC make the guard a target for attack. Furthermore, this design implies that the same security terminology is required at both domains connected by the guard. For example, the same set of security labels have to be used. Such an assumption is not always true, especially when the information exchange is across organizational boundaries. Second, associating the guards with the links often fails to scale as the network grows in size. The number of guards required grows exponentially as the number of interconnected security domains increases, commonly known as the "n-squared" problem. Finally, CDC guards often have limited configurability and vendor-specific API, resulting in locked-in stacks, increased development cost and the lack of flexibility to support mission requirements.

2.2 CDC in the Context of SOA

The need for aligning CDC with services and their supporting infrastructure has become increasingly evident as organizations adopt SOA. For example, the US Intelligence Community has developed a set of service specifications, including XML schemas and REST [6] API, for content discovery and retrieval across multiple repositories and potential different security domains[7]. While these specifications identify required services, CDC concerns are notably absent. Leaving the alignment of CDC infrastructure to service infrastructure to the interpretation of their perspective developers could result in inconsistent and potentially non-interoperable implementations.

We should note that the industry has developed innovative CDC solutions for specific service implementations despite of the lack of a CDC reference architecture. For example, XDDS [3] is a cross domain service discovery solution designed to work with existing guards. In this solution, a Local Discovery Agent (LDA) is deployed in a domain to intercept service discovery requests using Universal Description Discovery and Integration protocol (UDDI) [8]. It coordinates with LDA instances in other domains to locate appropriate services for the request. When the security policy requires the service provider's identity be masked, the solution uses a Global Discovery Service (GDS) for anonymization. Since the GDS is aware of all LDAs across all domains, the GDS will need the highest security privilege. There are several interesting obversations. First, this solution introduces a new infrastructure component (LDA) to address a particular CDC need (service discovery). Add-hoc components like LDA will be needed for other CDC usecases without an overall framework to address CDC concerns. Second, even though the security guards are considered transparent to the mission applications, the LDAs are not. The mission applications need to address UDDI requests to the LDA for local domain. This inconsistency calls for a convention on CDC infrastructure's role in application workflows. Finally, UDDI is an application layer protocol. There is clearly a need for the guards to understand the context of application-level interactions in order to make more intelligent security decisions.

2.3 Relevant Security Ontology Work

Current CDC solutions also require excessive amounts of human intervention, mainly due to the lack of a standard and flexible framework for describing information exchanged. Many guards use a dirty word list or some rudimentary rules expressed in

eXtensible Stylesheet Language Transformations (XSLT) to filter information passing through them. These techniques often fail to take into account the context of messages and the meaning of the words, leading to high error rates, i.e., false positives and negatives. Human review is often required to adjudicate ambiguities.

In order to build applications that can more precisely analyze information flows across domains**,** we argue that the security community should adopt a standard security ontology. A standard security ontology would provide the community with a common set of concepts around which they could form a shared understanding to advance the theory and practice of security, privacy, and trust of Web-based applications. A number of security ontologies have been developed and are currently in use. Some notable examples are the DAML Services Security and Privacy ontology[1], the Navy Research Lab (NRL) Security ontology[9] and SecOnt[2,3], which is based on the security relationship model described in the National Institute of Standards and Technology (NIST) Special Publication 800-12. These security ontologies focus on the areas of assets, threats, vulnerabilities, and countermeasures. Examples of how security ontologies are being applied include: the use of formal representations of policies in ontology and algorithms in order to support machine-aided reason about the policies [10] and the use of ontologies to annotate generic resources from simple documents to interactive services with security-related metadata and not just Web services.

Our review identified a number of relevant security ontology work but they are not specific to CDC [11]. We believe more ontology work needs to be done in this area.

3 Cross Domain Security Reference Architecture (CDC-RA)

3.1 Overview

CDC-RA is a key to CDC standardization. It provides: 1) a common framework and vocabulary to describe CDC mechanisms; 2) the abstract interaction patterns among CDC participants (basis for standardization through protocols); and 3) CDC infrastructure design patterns in the form of protocol constraints and assumptions.

We are primarily concerned with standardized interfaces for products whose primary responsibility is to facilitate secure cross-domain communications. A security guard is one such product, and perhaps the most important. However, CDC is a shared responsibility between the CDC infrastructure and the mission application taking advantage of such an infrastructure. We should not lose sight of the critical role that mission applications play. As illustrated in Fig. 1, we separate the CDC concerns into two categories: Application Aspects and Infrastructure Aspects.

The Application Aspects of CDC address the following:

- Architecture Concerns: We believe it is reasonable to assume that mission applications are aware of the fact that they interact with systems in other security domains, and therefore presence of the security guards. Application is responsible for properly handling the cases where a message is rejected by the guard.

[1] http://www.daml.org/services/owl-s/security.html

[2] http://www.securityontology.com

[3] http://www.ida.liu.se/~iislab/projects/secont/

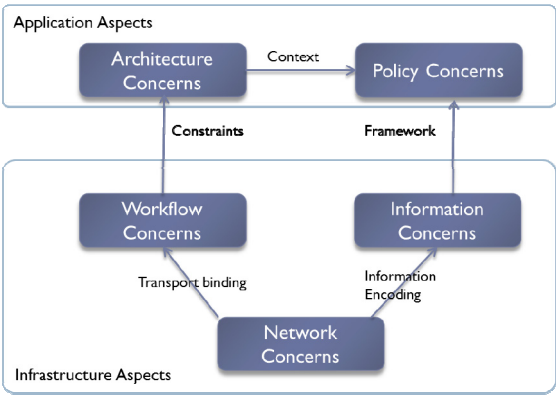

Fig. 1. Cross Domain Security Concerns

- Policy Concerns: Security attributes for application messages need to be defined so that proper security policies can be enforced by the infrastructure. We consider this an application specific concern since the security attributes associated with individual data elements are processed by that application.

The Infrastructure Aspects of CDC address the following:

- Network Concerns: We would like to address the concerns of how a guard interacts with the network, including how the CDC-specific communications are carried in the network protocols, for instance, in a SOAP message for web services-based communications or in the HTTP header for web traffic. Doing so may require extensions to existing protocols such that security-specific information could be added to the messages. A particular aspect of integrating guard with network protocol is a mechanism to handle end-to-end encryption and authentication. For example, if a mission application encrypts payload using WS-Security [12], the guard would not be able to inspect the message content unless the message is addressed to the guard (instead of the target system) and encrypted using the guard's key(s).
- Information Concerns: We would like to address how guards interact with the information flowing through them as part of the information concerns. There needs to be a convention for determining how application-specific messages are interpreted and acted upon by the guards, to enable automation and interoperability.
- Workflow Concerns: We would like to address how guards interact with other participants of the workflow, i.e. mission application and other guards. Compared to others, there hasn't been an extensive research on this aspect of CDC. Much less discussion exists about how existing standards such as Business Process Modeling Notation (BPMN) [13] and WS-BPEL [14] languages can be leveraged for CDC. One important workflow consideration is whether or not the guard is an active participant in the workflow, and if so, how the guard acts as a service intermediary.

Standards like BPMN and WS-BEPL allow complex workflows to be defined, but their use in CDC has been limited because guards have been largely absent in workflow definitions. As an active workflow participant, the guards will be able to enforce policies based on models expressed in BPMN and WS-BEPL. Guard vendors may even include BPMS functionality in their products to manage cross-domain workflows.

3.2 CDC Participants

With the CDC concerns cataloged, we need to discuss how CDC participants address these concerns collaboratively in order to determine their roles and responsibilities. As illustrated in Fig. 2, cross-domain solutions include the following four key participants:

i. Security Domain: We assume that, for a single security domain, there is a consistent security vocabulary for all actors, activities, and information. Examples include a single information classification hierarchy (Top Secret, Secret, FOUO). Furthermore, a security domain may have one or more Security Guards to enforce policies, described using the domain's security vocabulary.

ii. Mission Application: For the purpose of CDC, mission applications associate mission-specific concepts with the security vocabulary.

iii. Security Monitor (Optional): A domain may utilize a centralized security management and monitoring system. With the Security Monitor, the domain security administrator can define consistent security policies for communication with other domains using the domain's security vocabulary. A Security Monitor may communicate with the Security Guard at runtime.

iv. Security Guard: A security guard enforces security policy defined by the mission application, and may act as a policy enforcement point for the domain. A guard may coordinate with other guards, in the same or different domains, to enforce the security policies.

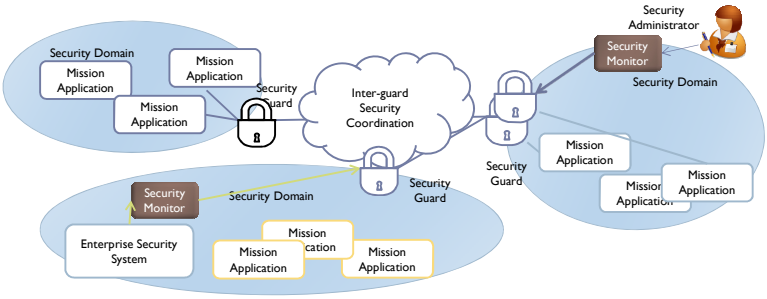

Fig. 2. CDC Participants

Deviating from the existing TOG approach of CDC security, we advocate associating security guards with domains instead of the links. Domain-Oriented Guards (DOG) would operate at the same security level as the associated domain, without unnecessary privilege. In addition, the same security monitor will be able to manage both the domain and the guard, avoiding policy conflicts and duplication. The number of guards required for securing inter-connected domains increases linearly with the number of domains unlike TOG where the increase is exponential in nature. The implication is that, for any communication path, there will be at least two guards, hence the need for inter-guard coordination. Because the guards need to trust each other without necessarily revealing mission information, the issues of identity and trust must be addressed in the context of inter-guard coordination/ synchronization. We envision a CDC protocol could require mutual authentication between guards through a mechanism such as Public Key Infrastructure (PKI) while assuming mutual trust to be established out of band, e.g. through a white list distributed among the guards.

3.3 Guards as Active Workflow Participants

In many CDC solutions, guards are considered transparent to the mission applications. In reality, a well-designed application must be aware of the fact that certain communications involve systems in other domains in order to handle the unique nature of CDC. For example, CDC communications may be blocked by the guards or delayed due to pending human review. Without the knowledge of the guards' existence along with a feedback mechanism, an application could siliently fail in the background.

Our reference architecture assumes mission applications are aware of the guards. As such, we see security guards as active paticipants in CDC workflows. By having the guard as an active participant in the workflow, it becomes possible to define a notification mechanism. The notification mechanism enables the mission application to be informed in case the guard blocks an application message, for security reasons, thereby allowing the mission application to take appropriate remedial actions. Having a guard as an active participant in the workflow also solves another otherwise difficult issue: end-to-end encryption. Without a message being explicitly addressed to the guard, encryption will prevent the guard from inspecting the message due to lack of cryptographic keys at the guard. It is now possible for the guards to encrypt the message on behalf of the application after taking appropriate security actions (redaction for example) and forward the message on to its ultimate destination.

For separation of business and security concerns, it may be possible to design an application workflow in a CDC-independent way using BPMN. When the process is deployed in a CDC environment, security guards are injected into a business process through such approaches as Object Management Group's (OMG) Model Driven Architecture® (MDA) [15], as shown in Fig. 3.

A guard-aware workflow opens up other possibilities as well. For example, the guard could also act as brokers for cross domain service discovery. It could also proxy the service provider and consumer to avoid unnecessary disclosure of system identity, while eliminating the needs for ad-hoc CDC infrastructure components as described in [3]. An implication of this approach is that the guard may have to expose the same interface as the invocation target, an issue that CDC protocol design needs to take into account. For example, to proxy web service, the guard may have to implement the same service interface as defined in Web Service Description Language (WSDL) [16].

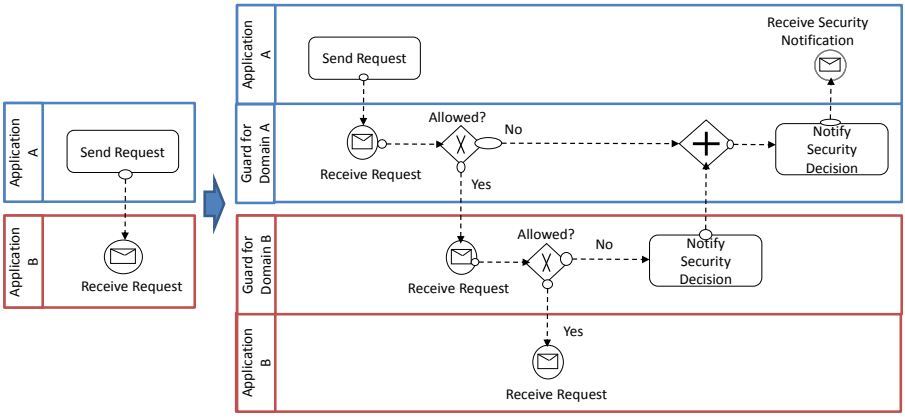

Fig. 3. Transformation of CDC Workflow

4 Security Ontology

The security community has begun to recognize the need for controlled vocabularies, taxonomies, and ontologies to make progress toward a science of (cyber) security [17]. In 2010, DOD sponsored a study to examine the theory and practice of security, and evaluate whether it is possible to adopt a more scientific approach [18]. In the context of CDC, the use of a security ontology would reduce the need for humans to adjudicate CDC messages, a problem exacerbated by the use of user-defined labels and keywords without precise meanings. These labels and keywords make it difficult for computer agents to analyze information flows without human intervention.

We envision a number of distinct, yet related ontologies to mediate the vocabularies and policies among different security domains. This will reduce the need to standardize on a single set of security policies across all domains. In addition, these ontologies will enable machine-to-machine communication, minimizing the need for human review. A community supported security ontology can be used to support the semantic annotation of generic resources such as documents, enterprise architectures, business process models and web service with security-related metadata. Semantic annotation is the act of associating an ontology term with a resource or some part of resource. More precisely, this means embedding a Uniform Resource Identifier (URI) within an information resource. Semantic annotation differs from user-generated tags in that the meaning of the metadata is defined in an external ontology that can be used disambiguate the metadata and supports the automatic discovery of resources.

The security ontology would also provide the capability to use reasoning to match mission requirements with service capabilities. Different domains may use different security vocabularies, making the discovery of services across domains difficult. The security ontology would help to harmonize the differences in security terminology and make it possible for agents to use a domain specific vocabulary and discover resources described using a different vocabulary. In [19], we described a semantic mediation infrastructure that uses ontologies to mediate the data model difference between SOA services. Similarly, semantic-aware guards could translate security vocabularies between domains based on these ontologies and enforce security policies accordingly.

We recommend that the security community should start from existing efforts to create security ontologies and, in order to ensure that others can share and reuse the ontology. We recognized the fact that it is both unlikely and undesirable that there be a single security ontology. Instead, we recommend that a suite of modular ontologies be developed, with a single core security ontology at the center. This makes it possible for different groups to extend the core security ontology to address their respective needs.

5 CDC Protocol Candidates

We see opportunities to define protocols that specify the interaction among CDC participants. These protocols will ensure interoperability among CDC participants and guards, and allow organizations to tailor security configurations based on mission needs. CDC protocol candidates include:

- CDC Application Interface: Enables interactions between mission applications and the security guard within a security domain. This interface can be specified in two levels: 1) an abstract protocol for communication with the guards that is

transport-independent and 2) protocol bindings for realizing the abstract protocol with a particular transport mechanism. For web services, the mission application could potentially use WS-Addressing [20] to indicate to the security guard the ultimate destination of the message. Business processes carried out collectively by mission applications in different domains can be best understood by analyzing the interactions between applications. As such, we recommend that the security domains be configured to allow only specific application level protocols and accordingly, the guards be implemented at the application layer.

- Inter-Guard Coordination Protocol: Enables interactions among the security guards. Leveraging the security ontology, a mechanism can be defined for guards to authenticate among themselves, correlate security attributes of the source and the target applications, the activities, and the information carried in the payload. Using annotations, we can associate the application-specific metadata with the concepts in the ontology, and further associate ontology concepts with the security attributes that will be used by the guard to make runtime decisions.
- Security Monitor Interface: Defines an interface to manage CDC infrastructure, perhaps by extending the Simple Network Management Protocol (SNMP) [21].

With Inter-Guard Coordination Protocol, we are envisioning a more peer-to-peer coordination among the guards, avoiding the need a single global security system to manage multiple domains. Such system is often not practical in an inter-organizational environment and it introduces the potential of a single point of failure.

6 Road to Implementation

The CDC reference architecture can be used to guide the development of solutions that compliments and enhances existing CDC guard products to secure SOA interactions. As the first step, we see software components developed to compliment existing TOG. This component, Cross Domain Service Proxy, is associated with a particular domain and works with the domain's SOA infrastructure such as an Enterprise Service Bus (ESB) and Identify and Access Management (IdAM). Because the Proxy is an active participant of CDC workflows, applications or services always address cross domain messages to the Proxy, which can then filter the messages based on policies and inspect the content using ontologies. Through the Guard Application Interface described earlier, the Proxy can inform mission applications security decisions it takes so that the application can respond properly. Only compliant messages are then sent out via communication link protected by the traditional guard.

This solution does introduce redundant message inspections – at the Proxy and again at the traditional cross domain guard, mainly to alleviate concerns with the new architecture. We expect that, as organizations gain confidence in the architecture, the Proxy can be either integrated into the guard products or replace the current guards.

7 Conclusion

Standardization of interactions among these CDC participants is a pre-requisite for achieving the interoperability and flexibility required by business. Due to the complex and multi-faceted nature of CDC security, standardization is only possible within a framework where interactions can be abstracted and discussed in a structured manner.

This paper represents our attempt at establishing such a framework, and we hope it will encourage further discussions within the community, resulting in more interoperable, flexible and efficient CDC solutions to serve the needs of the business.

References

[1] Swamy, N., Hicks, M.: Verified Enforcement of Security Policies for Cross-Domain Information Flows, http://www.cs.umd.edu/~mwh/papers/selinks-cpa.pdf

[2] Irvine, C.E., et al.: MYSEA: the Monterey security architecture. In: Proc. of the Workshop on Scalable Trusted Computing (ACM STC), Conference on Computer and Communications Security (CCS), pp. 39–48. Association for Computing Machinery (ACM), Chicago (2009)

[3] Atighetchi, M., et al.: XDDS: A Salable Guard-Agnostic Cross Domain Discovery Service, http://www.dtic.mil/cgi-bin/GetTRDoc?AD=ADA532504

[4] W3C, Web Services Architecture, W3C Working Group Note (February 11, 2004)

[5] Shader, M.: Cross-Domain Application Architecture: The Need for an End-to-End Approach (2012), http://yellowhouseassociates.net/download/YHA_CDAA_WP.pdf

[6] Fielding, R.: Architectural styles and the design of network-based software architectures. Diss. University of California, Irvine (2000)

[7] Intelligence Community and Department of Defense Content Discovery and Retrieval Integrated Project Team. IC/DoD Content Discovery and Retrieval Reference Architecture (February 2011)

[8] OASIS, Universal Description, Discovery and Integration v3.0.2, OASIS Standard (February 2005)

[9] Kim, A., Luo, J., Kang, M.: Security ontology for annotating resources. In: Meersman, R. (ed.) OTM 2005. LNCS, vol. 3761, pp. 1483–1499. Springer, Heidelberg (2005)

[10] Denker, G., Kagal, L., Finin, T.: Security in the Semantic Web using OWL. Information Security Technical Report 10(1), 51–58 (2005)

[11] Blanco, C., et al.: A Systematic Review and Comparison of Security Ontologies, ares. In: 2008 Third International Conference on Availability, Reliability and Security, pp. 813–820 (2008)

[12] OASIS, Web Services Security: SOAP Message Security 1.1, OASIS Standard (February 2006)

[13] Object Management Group (OMG), Business Process Model and Notation (BPMN) Version 2.0, OMG Standard (January 2011)

[14] OASIS, Web Services Business Process Execution Language 2.0, OASIS Standard (April 2007)

[15] Object Management Group (OMG), Model Driven Architecture ®, http://www.omg.org/mda/

[16] W3C, Web Services Description Language (WSDL) 1.1, W3C Note (March 15, 2001)

[17] Mundie, D.A., McIntire, D.M.: The MAL: A Malware Analysis Lexicon. CERT® Program - Carnegie Mellon University. Technical (2013)

[18] The MITRE Corporation, Science of Cyber-Security, The MITRE Corporation. Technical (2010)

[19] Zhu, W.: Semantic Mediation Bus: An Ontology-based Runtime Infrastructure for Service Interoperability. In: 2012 IEEE 16th International Enterprise Distributed Object Computing Conference Workshops (EDOCW), September 10-14, pp. 140–145 (2012)

[20] W3C, Web Services Addressing 1.0 – Core, W3C Recommendation (May 9, 2006)

[21] Harrington, D., Presuhn, R., Wijnen, B.: An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks. Internet Engineering Task Force RFC (December 2002)