FIDDLE: Federated Infrastructure Discovery and Description Language

Alexander Willner and Robyn Loughnane
Architekturen der Vermittlungsknoten (AV)
Technische Universität Berlin, Germany

Thomas Magedanz
Next Generation Network Infrastructures (NGNI)
Fraunhofer FOKUS, Berlin, Germany

***Abstract*—Considerable efforts have been spent on designing architectures to manage heterogeneous resources across multiple administrative domains. Specific fields of application are federated cloud computing (Intercloud) approaches and distributed testbeds, among others. An important interoperability challenge that arises in this context is the exchange of information about the provided resources and their dependencies. Existing work usually rests upon schematic data models, which impede the discovery and management of heterogeneous resources between autonomous sites. One way of addressing this issue is to exchange semantic information models. In this paper, we exploit such approaches to formally define federations, including their infrastructures and the life-cycle of the offered resources and services. The requirements of this work have been derived from several research projects and the results are in process of being standardized by an international body. The main contribution of this work is a higher level (upper) ontology and initial integration concepts for it. These contributions form a basis for further work in the general context of distributed semantic resource management.**

IEEE ***Index Terms*—Intercloud, Infrastructure Federation, Testbeds, e-infrastructures, Semantic Web, SDN, NFV, IoT**

I. Introduction

moti-vation

The development of the Internet has paved the way to design concepts to distribute computing tasks between geographically distinct areas. Specific fields of application include Intercloud computing approaches and the spread of experiments over geographically distributed test environments in the context of Future Internet (FI) research. This also includes federated Virtualized Network Function (VNF) provisioning over Software Defined Networkings (SDNs) for e.g. mobile network operators, big-data analysis and the Internet of Things (IoT). In any case, information about the provided services and resources of the distributed infrastructures needs to be exchanged one way or another. This leads to a wide range of frameworks, protocols, and architectures that have mainly been designed to exchange arbitrary schematic data models.

impor-tance

In the context of interconnected, federated, autonomous administrative domains, this approach paradoxically often aggravates interoperability issues. For example, XML and JSON based data models define the structure of the data in each single domain as a tree. However, the union of two tree structures is not a tree and the structure quickly becomes unwieldy [1]. More importantly, the actual meaning of the transferred data is encoded in functional code running in the particular domain. This aggravates mutual understanding and interoperation, which has been the major objective of the above mentioned approaches. With n domains participating in a federation this leads to a combinatorial problem of n^2 required transformations that have to be implemented.

related work

Instead of relying on data model compatibility, one approach to address these issues is to handle interoperability on the information model level. A specific field of application is the Semantic Web[2], where information can be formally defined and abstracted from existing data models. The basic concept is to model information as labeled directed graphs instead of trees using the Resource Description Framework (RDF). These foundations have already bee used to manage e.g. networks by defining formal languages such as the Network Mark-Up Language[3] (NML) or clouds using the Open-Source API and Platform for Multiple Clouds[4] (mOSAIC). Further, first approaches to map this to the management of more general infrastructures have been developed, like the Infrastructure and Network Description Language[5] (INDL). However, while the existing work provides highly interesting insights and describes valuable concepts, it has been developed and used in rather narrow contexts.

as-sump-tions

In this paper, we assume that (i) federation mechanisms are either already in place or under active development in the area of Intercloud computing or FI experimentation; (ii) the concept can be abstracted to a generic resource life-cycle management over administrative domains; (iii) it will be adopted by more fields of application in the future, while the underlying architectures and Application Programmers Interfaces (APIs) evolve and change; and (iv) managing resources on an abstract information level allows re-usability of the developed concepts and interoperability between independent domains.

contri-bu-tions

We present a formal information model to describe federations, their infrastructures and the contained resources and services in order to allow software integration concepts to semantically manage distributed objects. This upper ontology can be used and extended to exchange information between federated infrastructures in order to discover and use available resources and services independently from specific APIs or architectures.

valida-tion

We have initially validated our approach by intense discussions within the Future Internet Research and Experimentation[1] (FIRE) and adjoining communities, the analysis of their

[1] http://ict-fire.eu

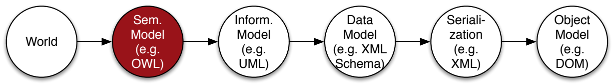

Fig. 1. Relationship between models and the syntax (based on [6], [7])

requirements and the inspection of related work; also taking work from the Grid Computing context into account. This discussion has then been extended to the fields of Interclouds, Internet of Things, and Network Function Virtualization. Further, we have developed reference implementations for two contexts to demonstrate the applicability of the model. Finally, the underlying concept is currently in the process of being refined and standardized by an international consortium that is independent of specific research projects or products.

structure The remainder of the paper is structured as follows. We give a brief overview of related work in the context of semantic resource management in Sec. II. In Sec. III, we describe our approach in more detail and present the upper ontology. In the Sec. IV-A example assertions based on the presented ontology are given. The results obtained from an initial implementation are discussed in Sec. IV. Finally, we close giving some conclusions and considerations and describe future work in Sec. V.

II. Related Work

In order to put the contribution of this paper into context, it is important to elaborate on the difference between an Object Model (OM), a Syntax or Serialization, a Data Model (DM), an Information Model (IM) and a Semantic Model (SM). Further, some specific fields of application for semantic resource descriptions are given to highlight the applicability and universality of the underlying principles.

A. Information Modeling

Foundations In this paper we're following the argumentation from RFC 3444[6]. The corresponding relationships are highlighted in Fig. 1. Typically, developments of middleware APIs in the Information and Communication Technology (ICT) sector concentrate on the exchange of semistructured tree data models. Often serialized as XML or JSON, these are then transformed to programming language dependent OMs.

Issue As already noted, the composition of such data in a federated environment becomes rather complex and involves significant amount of functional code. Even if two trees refer to the same resource, it is likely that related information will be found in distinct locations in the tree. As a result, additional choices must be made to obtain a new well-formed XML document (see Listing 1).

Listing 1. Same semantics but different syntax

```
<resource name="PC−133">
  <administrator>Alexander Willner</administrator>
</resource>
<administrator name="Alexander␣Willner">
  <resource>PC−133</resource>
</administrator>
```

Ontologies One approach to address these issues is to move the interoperability from the structure to the semantic level[8] by introducing semantic IMs and defining ontologies. Ontologies are formal, explicit specifications of shared conceptualizations[9]. In other words, they are shared formalized vocabularies about individuals (instances), classes (concepts), attributes, and relations. They allow us to link, combine, validate, reason and query information using only Description Logics (DL), to automatically infer knowledge about the relationships of different resources.

Semantic Web In particular the growth of the Semantic Web[2] and the adoption of the accompanied graph-based data model in RDF has seen an expansion of accompanying languages, protocols and standards. These include the knowledge representation languages Resource Description Framework Schema (RDFS) and Web Ontology Language (OWL); and the query language SPARQL Protocol And RDF Query Language (SPARQL). For example, one of the larger and better known data sets is DBpedia[2][10].

B. Fields of Application

Applications This transition from syntactic to semantic interoperability can be observed in different ICT related areas, such as Grid Computing, Web Services, Cloud Computing, the Internet of Things and FI testbeds.

Grid Computing For example in the Grid Computing context, the main purpose of the Grid Laboratory for a Uniform Environment[11] (GLUE) schema was to allow interoperability between grid projects in the US and the EU by defining a schematic vocabulary with serializations to XML, LDAP and SQL[12]. The lack of formalism, and, as a consequence, the missing means to reason over the information, motivated the transition to a Semantic Open Grid Service Architecture[13] (S-OGSA).

Internet of Things Likewise, in the IoT domain, the use of ontologies is needed to allow for interoperability between different realms and approaches[14]. In general, two main standardization activities can be identified in this context: the European Telecommunications Standards Institute (ETSI) M2M[3] Working Group has identified the need for a semantic resource description; and the successor OneM2M has already established a dedicated group on Management, Abstraction and Semantics[4].

Cloud Computing In Cloud Computing and all its application sectors, it is inevitable that interoperability on a semantic level will be required sooner or later. There is currently a focus on Intercloud approaches and 20 such proposals are presented in [15]. Only one uses a semantic Intercloud resource exchange, namely the IEEE Intercloud architecture. Based on the IEEE Standard for Intercloud Interoperability and Federation[16][5] (P2302) Working Group, it builds on the existing mOSAIC ontology.

Testbeds Finally, another field of application with a body of existing work is the research conducted on federated networks and FI testbeds, e.g., in the FIRE and Global Environment for Network

[2]http://dbpedia.org
[3]http://www.etsi.org/technologies-clusters/technologies/m2m
[4]http://onem2m.org/mashome.cfm
[5]http://standards.ieee.org/develop/project/2302.html

Innovations[6] (GENI) projects. In order to address interoperability issues within the management of heterogeneous networks, Semantic Web approaches have again been introduced[17]. One notable thesis that continued this work was the definition of the Network Description Language[18] (NDL), a formalization of the ITU-T Generic Functional Architecture of Transport Networks (G.805) using RDFS. NDL formed the basis for further work, such as the Network Description Language based on the Web Ontology Language[19][7] (NDL-OWL) and the Open Grid Forum (OGF) standardized version NML. These ontologies mainly focus on the underlying network and consequent work analyzed challenges for semantically describing federated virtualized infrastructures[20]. As a result, the Networking innovations Over Virtualized Infrastructures[21] (NOVI) model and its successor INDL have been created to take these additional aspects into account. An overview of current work in this area, such as Testbed as a Service Ontology Repository[22] (TaaSOR), is given in [23] and [24].

C. Summary

Synopsis These approaches serve their dedicated purposes in different domains of application and represent valuable foundations. However, multiple SMs have defined the same high level concepts and properties in similar but not equal ways. Therefore, the present work is intended to act as a seeding document for an upper ontology for the management of federated ICT infrastructures. This can be compared to the Suggested Upper Merged Ontology[25] (SUMO), which served as a starter document for the IEEE Standard Upper Ontology (P1600.1).

III. FIDDLE Upper Ontology

overview The goal of the Federated Infrastructure Description and Discovery Language (FIDDLE) is threefold: first, to build on existing work by importing and referencing related ontologies; second, to allow its use for resource life-cycle management in the field of FI testbeds (and the related e-infrastructure area); and, finally, to act as a seeding document for further standardization in these contexts. These objectives are reflected in the overview Fig. 2, which shows the proposed concepts covered by the ontology at four main levels of semantic abstraction: Federation, Infrastructure, Object and Life Cycle.

A. Federation

Fig. 3 In order to describe the federation itself, we introduce the concepts of fiddle:Federation and fiddle:FederationMember. These are both subclasses of the Friend of a Friend[8] (FOAF) foaf:Organization class. The federation level of abstraction is illustrated in Fig. 3, along with the relevant object properties.

argumentation This approach was chosen because these concepts are not purely technical and may have properties like an address, an administrative structure, a funding body and so forth. If these were grouped with the other entities in the ontology, then it would not be appropriate to attribute such properties to them, from either an abstract semantic or technical perspective.

infrastructure A fiddle:FederationMember may also administer one or more fiddle:Infrastructures, denoted by the object property fiddle:administers (and its inverse, fiddle:isAdministeredBy).

B. Infrastructure

Fig. 4 Whereas administrative and other aspects of an infrastructure are described by fiddle:FederationMember, a fiddle:Infrastructure focuses on the technological aspects. An infrastructure may exist whether or not it is part of a fiddle:Federation. The related ontology is illustrated in Fig. 4, along with the relevant object properties.

argumentation The term Infrastructure is introduced here with a similar sense to novi:Platform, but with broader scope. In particular, it is not restricted to testbeds, but expands to general ICT infrastructures that can be federated, such as Cloud computing facilities. Further, the class fiddle:Infrastructure is a subclass of fiddle:Group and can therefore be seen as a collection of fiddle:ICTObjects, expressed by the relations fiddle:hasResource and fiddle:hasService (analog to the properties nml:hasService and nml:hasNode).

C. Generic Concepts

Fig. 5 A number of the generic common concepts have been previously identified in the NML, NOVI and INDL vocabularies. These are linked to here with owl:equivalentClass or rdfs:seeAlso annotations. The generic FIDDLE concepts are shown in Fig. 5, together with all subclass relations and selected object properties. Note that most of the object properties have inverse relations that are not shown in the figure. Equivalent classes and see also references are also not shown.

argumentation The class fiddle:ICTObject is comparable to nml:NetworkObject and novi:Resource and acts as the superclass of the basic concepts fiddle:Resource, fiddle:Group and fiddle:Service. These are again based on the NOVI classes of the same name and comparable with the NML concepts nml:Node, nml:Group and nml:Service. As per NOVI and INDL, a fiddle:Resource has the subclasses fiddle:VirtualResource and fiddle:Component. These can be related to a fiddle:Resource by the object properties fiddle:implements and fiddle:hasComponent respectively.

attribute Each of these basic concepts can also take on a fiddle:Attribute via the fiddle:hasAttribute object property, to allow extension for, for example, read-only monitoring attributes or read-write resource control attributes.

outlook Further work is planned to design the fiddle:Dependency concept to specify dependencies between individuals based on output or input values for service chaining; and the fiddle:Policy concept to specify authorization related information.

D. Specific Concepts

out of scope In order to describe the heterogeneous services and resources in a federated environment, domain-specific ontologies have to be defined that usually re-uses and links existing work. While the definition of these ontologies is purposely out of scope, to

[6] http://geni.net
[7] http://geni-orca.renci.org/owl
[8] http://xmlns.com/foaf/0.1/

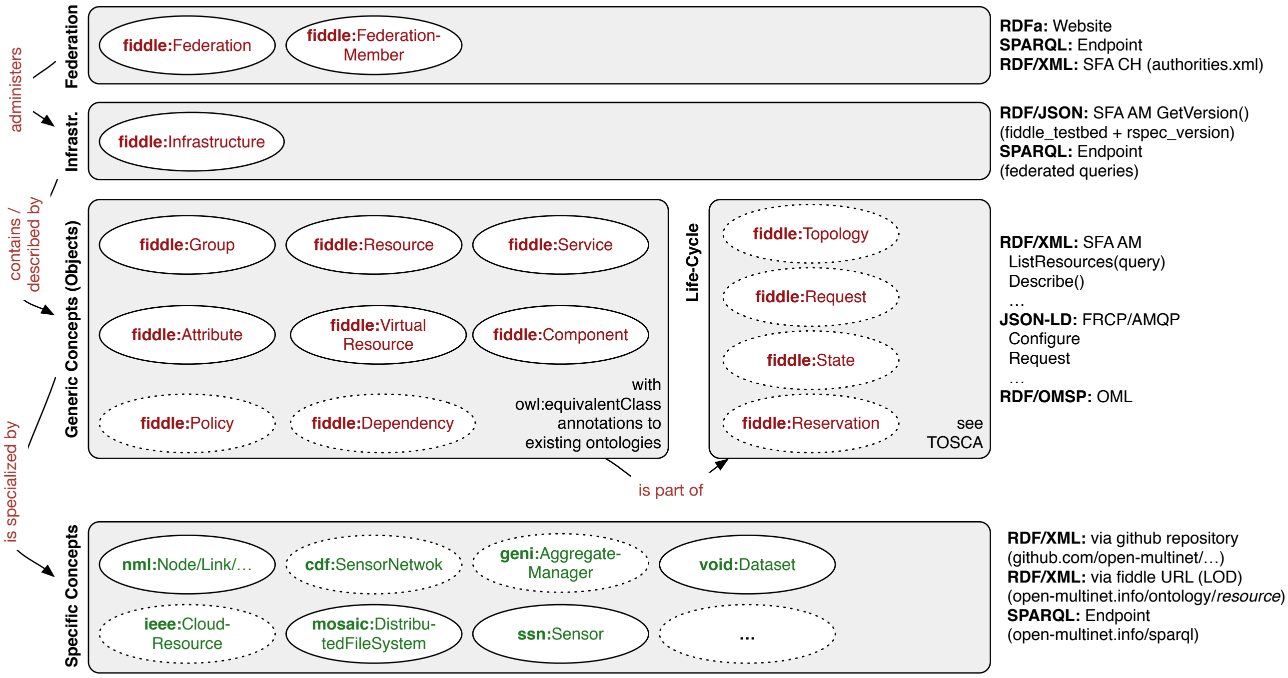

Fig. 2. FIDDLE: Overall Ontology and Integration Concepts (dotted lines: work in progress)

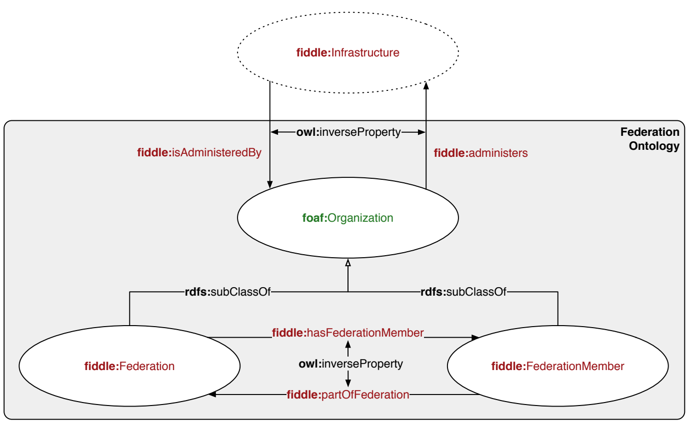

Fig. 3. FIDDLE: Federation Level

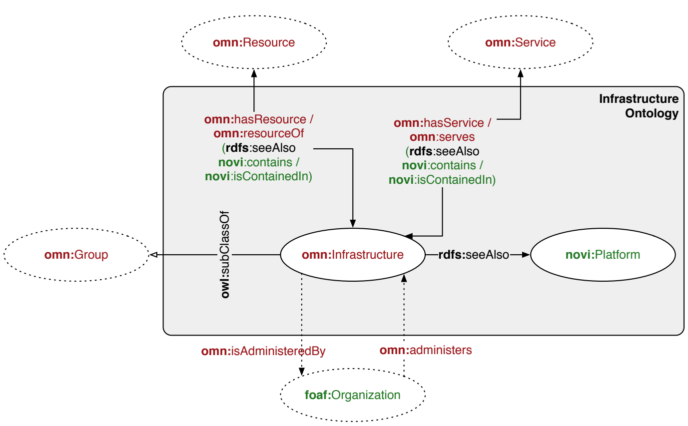

Fig. 4. FIDDLE: Infrastructure Level

support their discovery and management they should be linked to the aforementioned generic concepts.

E. Life-Cycle

One purpose of FIDDLE is to support the life-cycle of heterogeneous resources in federated environments. Therefore, the following concepts are needed and are currently under development, while evaluating existing work in parallel, such as the Topology and Orchestration Specification for Cloud Applications[26] (TOSCA). *moti-vation*

The key definition in this area is the concept of a fiddle:Topology. Analog to the nml:Topology and novi:Topology it is a subclass of fiddle:Group and acts as a container for a dynamic topology/testbed/infrastructure/slice. *topol-ogy*

The fiddle:Topology is then provisioned, described, modified or deleted within a fiddle:Request for a given set of fiddle:Infrastructures. This includes a bound or unbound request for a specific set of resources and services. Further, a provisioning request can contain time constraints specified in the fiddle:Reservation concept. *re-quest*

The fiddle:State concept then indicates the current status of an fiddle:ICTObject, e.g., whether a requested topology or resource is starting or stopped. *state*

F. Outlook

The definition of an upper ontology for the given broad application context of federated infrastructures requires the involvement of many stakeholders, such as tool developers, facility owners and federation operators. As a result, the Open-Multinet[9] (OMN) Forum has been established by a group of *Open-Multinet Fig. 6*

[9] http://open-multinet.info

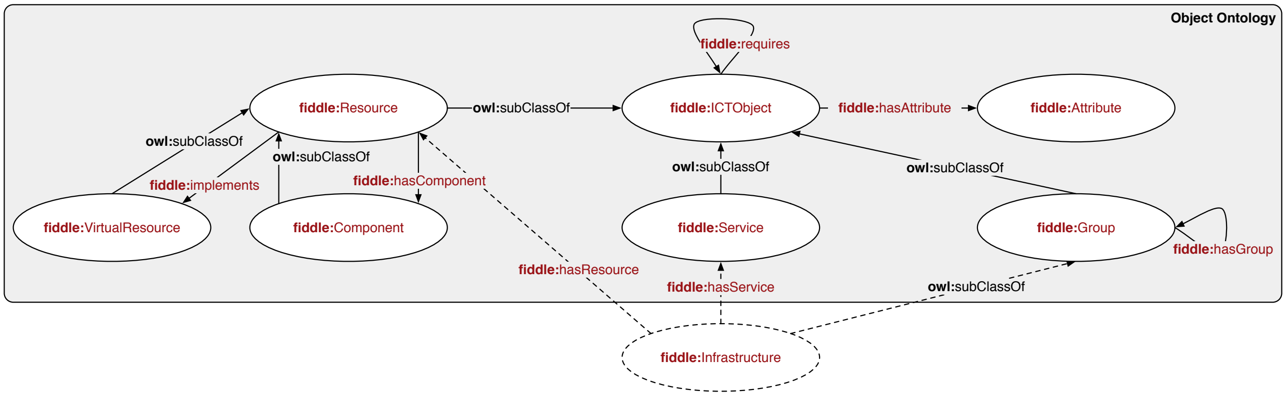

Fig. 5. FIDDLE: Object Level

international experts to define an according upper ontology. The presented work served as a starting point and in Fig. 6 a high level overview of the modularized approach has been given. Further, the proposed ontology has already started to be discussed within the IEEE P2302 Working Group.

IV. Evaluation

overview

As mentioned in Sec. III, one requirement for FIDDLE is support to manage federated infrastructures in general and FI testbeds in particular. Therefore, in order to validate the proposed work, three approaches have been followed. First, the general requirements, concepts and drafts have been discussed in the research project Federation for FIRE (Fed4FIRE), the IEEE standardization effort P2302 and with authors of the related work. Second, a federation that is already in place has been modeled using the ontology. Third, the usage of the information model has been integrated into existing management systems.

A. Exemplary Descriptions

federation

As an example of the applicability of the proposed ontology, the following Assertion-Box (A-Box) in Listing 2 describes the Fed4FIRE federation from a high-level point of view.[10] The subject http://fed4fire.eu/about#me is an instance of the fiddle:Federation concept and re-uses well-known vocabularies to describe common meta-information. The federation consists of multiple partners of which the member http://fokus.fraunhofer.de/about#me is described in more detail. While the Fraunhofer FOKUS might also provide further details at the given URL, in this example we decided to include basic information directly in the A-Box and to link to the corresponding DBpedia entry.

Listing 2. Fed4FIRE Federation A-Box Excerpt

```
<http://fed4fire.eu/about> rdf:type owl:Ontology .
:me rdf:type fiddle:Federation , owl:NamedIndividual ;
        rdfs:label "Fed4FIRE" ;
        foaf:homepage <http://fed4fire.eu> ;
        foaf:depiction <http://fed4fire.eu/.../logo.jpg> ;
        foaf:mbox <mailto:contact@fed4fire.eu> ;
        fiddle:hasFederationMember
                <http://iminds.be/about#me>, ...,
                <http://fokus.fraunhofer.de/about#me> .
<http://fokus.fraunhofer.de/about#me> rdf:type fiddle:FederationMember ,
                owl:NamedIndividual ;
        rdfs:label "Fraunhofer␣FOKUS" ;
    foaf:based_near [ a wgs84:Point; wgs84:lat 52.526; wgs84:long 13.314] ;
        rdfs:isDefinedBy <http://de.dbpedia.org/page/Fraunhofer−
            Institut_für_Offene_Kommunikationssysteme> ;
    fiddle:administers <http://fuseco.fokus.fraunhofer.de/about#me> .
```

[10]The nonsemantic version is available at http://fed4fire.eu/home/partners.html

infrastructure

Following in Listing 3 is a description of the FUSECO Playground infrastructure, administered by the Fraunhofer FOKUS. Of particular interest is the possibility to define endpoints in order to allow infrastructures to offer their services via different APIs, which might be needed for different federation contexts. In this example Slice-based Federation Architecture[27] (SFA) Aggregate Manager (AM) and Slice Authority (SA) APIs are being offered. A SPARQL endpoint is also defined using the Vocabulary of Interlinked Datasets[28].

Listing 3. FUSECO Playground A-Box Excerpt

```
<http://fuseco.fokus.fraunhofer.de/about> rdf:type owl:Ontology .
:me rdf:type fiddle:Infrastructure, owl:NamedIndividual ;
            rdfs:label "FUSECO␣Playground" ;
                foaf:homepage <http://fuseco−playground.org> ;
            foaf:based_near :location ;
            fiddle:isAdministeredBy <http://fokus.fraun...> ;
            fiddle:hasService :am, :sa ;
            void:sparqlEndpoint <http://fuseco−playground.org/sparql>;
:am rdf:type geni:AggregateManager ;
    fiddle:hasURI <https://fuseco.fokus.fraunhofer.de/api/sfa/am/v3> .
:sa rdf:type geni:SliceAuthority ;
    fiddle:hasURI <https://fuseco.fokus.fraunhofer.de/api/sfa/sa/v1> .
```

B. Software Integration Concepts

overview

To demonstrate first integration concepts, the ontology has been used to exchange information within SFA and P2302 federation contexts. For this purpose, the management framework FITeagle[11], a reference implementation of the Federated Infrastructure Resource Management Architecture[29] (FIRMA), had been extended and the SFA integration will shortly be described. To the right of Fig. 2, related implementation opportunities for FIRE protocols and Semantic Web approaches are sketched.

[11]http://fiteagle.org

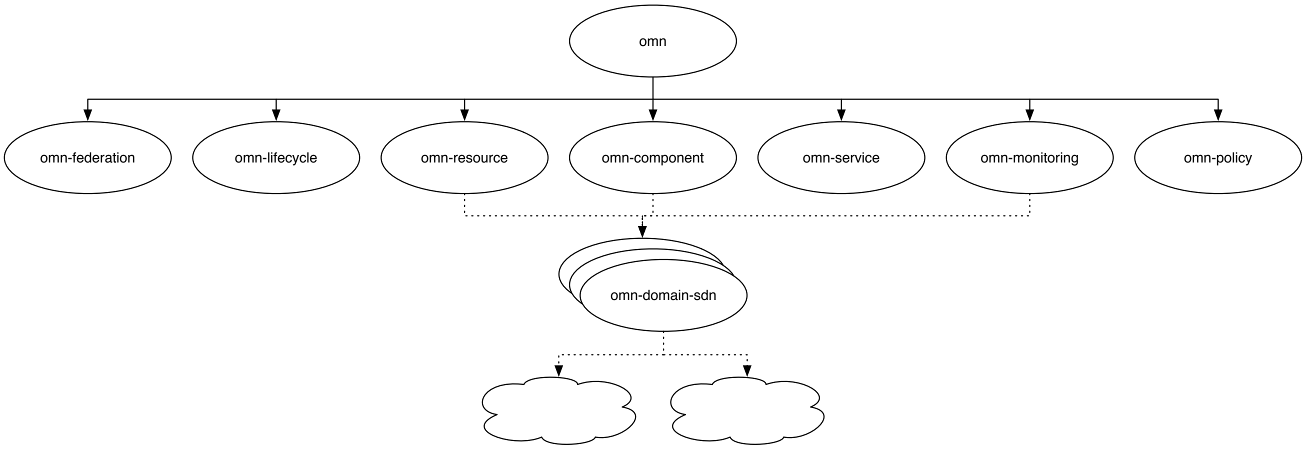

Fig. 6. Preliminary Open-Multinet Upper-Ontologies (dotted lines: link to domain-specific ontologies)

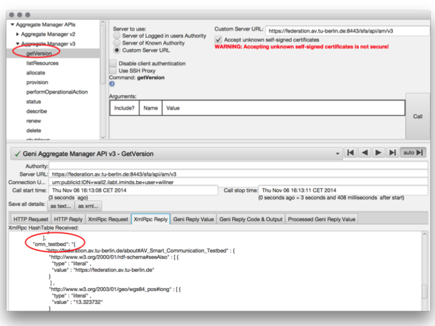

Fig. 7. Integration into the SFA AM GetVersion() Call

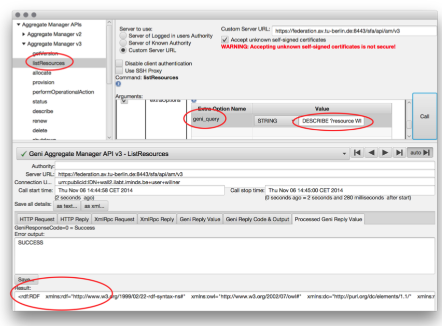

Fig. 8. Integration into the SFA AM ListResources() Call

meta data

In order to gather details about an infrastructure and the offered APIs, the AM method call GetVersion() is used in SFA. Fig. 7 shows the result of such a call using the client library jFed[12]. As highlighted, an RDF/JSON serialized description of the testbed has been embedded in the JSON based response data structure by adding a new node in the tree.

discov-ery

To discover available resources in SFA, the AM method call ListResources() is executed. In Fig. 8 the integration of FIDDLE in this method call is shown. First, instead of the usual GENI Resource Specification (RSpec) XML tree, an RDF/XML serialized graph is being returned. Second, an additional parameter has been introduced in order to allow server side filtering of resources by passing a SPARQL query, which was not supported before.

Out-look

While only two basic examples are presented in the context of this paper, work on a deeper integration into the life-cycle management of dynamic topologies has already been started.

V. Conclusion and Future Work

results

We have shown which concepts and relationships are generally needed to semantically describe federated infrastructures and how this information can be used to manage resources within such environments. The crucial result is a generic upper ontology that forms a foundation for further standardization efforts in different areas of application.

bene-fits

This work can be used by infrastructure owners to describe their resources and services once, in order to publish this information using different APIs. Users who want to build or extend a requested topology have means to build complex queries and can precisely define their requirements. Developers of user or infrastructure tools do not need to write functional code to extract the meaning out of different data structures and can focus on the provided information instead. To the extent to which it may be applicable, federation owners have the possibility to harmonize different specifications using description logic.

out-look

Further work includes the refinement and publication of the proposed information model in the short term. The midterm goal is to refactor and standardize the work on an international level; this has already been started within the OMN and P2302 consortia. The long-term objective is to extend the standardization activities further to cover the fields Software-Defined Network Exchanges (SDX) for multidomain SDN control, and Network Function Virtualization (NFV) orchestration for mobile network operators, as well as IoT ontologies for Smart Cities and Smart Factories.

[12]http://jfed.iminds.be

VI. Acknowledgments

Research for this paper was partially financed by the European Fed4FIRE (Grant No. 318389) project. We thank our project partners for their contributions and their collaboration on this research work.

References

[1] P. Hitzler, M. Krotzsch, and S. Rudolph, *Foundations of semantic web technologies*. CRC Press, 2011.

[2] T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific American*, vol. 284, no. 5, pp. 34–43, May 2001.

[3] J. van der Ham, F. Dijkstra, R. Lapacz, *et al.*, *GFD.206: Network markup language base schema version 1*, 2013.

[4] F. Moscato, R. Aversa, B. Di Martino, *et al.*, "An analysis of mOSAIC ontology for Cloud resources annotation," in *Federated Conference on Computer Science and Information Systems (FedCSIS)*, IEEE, 2011, pp. 973–980.

[5] M. Ghijsen, J. van der Ham, P. Grosso, *et al.*, "Towards an Infrastructure Description Language for Modeling Computing Infrastructures," in *IEEE 10th International Symposium on Parallel and Distributed Processing with Applications*, IEEE, Jul. 2012, pp. 207–214.

[6] A. Pras and J. Schoenwaelder, *On the Difference between Information Models and Data Models*, RFC 3444 (Informational), 2003.

[7] A. Pras, J. Schonwalder, M. Burgess, *et al.*, "Key research challenges in network management," *Communications Magazine, IEEE*, vol. 45, no. 10, pp. 104–110, 2007.

[8] A. P. Sheth, "Changing Focus on Interoperability in Information Systems: From System, Syntax, Structure to Semantics," in *Interoperating Geographic Information Systems*, Springer, 1999, pp. 5–29.

[9] T. Gruber, "A translation approach to portable ontology specifications," *Knowledge acquisition*, vol. 5, no. 2, pp. 199–220, 1993.

[10] S. Auer, C. Bizer, G. Kobilarov, *et al.*, *Dbpedia: A nucleus for a web of open data*. Springer, 2007.

[11] S. Andreozzi, S. Burke, L. Field, *et al.*, *Glue Schema Specification*, 2008.

[12] M. Drozdowicz, M. Ganzha, M. Paprzycki, *et al.*, "Ontologies, agents and the grid: An overview," in *Parallel, Distributed and Grid Computing for Engineering*, Stirlingshire: Saxe-Coburg Publications, 2009, pp. 117–140.

[13] O. Corcho, P. Alper, I. Kotsiopoulos, *et al.*, "An overview of S-OGSA: A reference semantic grid architecture," *Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 4, no. 2, pp. 102–115, 2006.

[14] S. Gaglio and G. Re, *Advances Onto the Internet of Things: How Ontologies Make the Internet of Things Meaningful*. Springer Publishing Company, Incorporated, 2014.

[15] N. Grozev and R. Buyya, "Inter-Cloud Architectures and Application Brokering: Taxonomy and Survey," *Software: Practice and Experience*, vol. 44, no. 3, pp. 369–390, 2012.

[16] D. Bernstein and V. Deepak, "Draft Standard for Intercloud Interoperability and Federation (SIIF)," *IEEE P2303*, 2014.

[17] A. Wong, P. Ray, N. Parameswaran, *et al.*, "Ontology mapping for the interoperability problem in network management," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 10, pp. 2058–2068, Oct. 2005.

[18] J. Van Der Ham, "A Semantic Model for Complex Computer Networks: The Network Description Language," PhD thesis, Universiteit van Amsterdam, 2010, p. 164.

[19] Y. Xin, C. Hill, I. Baldine, *et al.*, "Semantic Plane : Life Cycle of Resource Representation and Reservations in a Network Operating System," RENCI, Tech. Rep., 2013.

[20] J. van der Ham, C. Papagianni, J. Steger, *et al.*, "Challenges of an information model for federating virtualized infrastructures," in *5th International DMTF Academic Alliance Workshop on Systems and Virtualization Management: Standards and the Cloud (SVM)*, IEEE, Oct. 2011, pp. 1–6.

[21] J. van der Ham, J. Stéger, S. Laki, *et al.*, "The {NOVI} information models," *Future Generation Computer Systems*, vol. 42, pp. 64–73, 2015.

[22] M. Tosic, I. Seskar, and F. Jelenkovic, "TaaSOR: Testbed as a Service Ontology Repository," *Testbeds and Research Infrastructure Development of Networks and Communities*, vol. 44, pp. 419–420, 2012.

[23] M. Tosic and I. Seskar, "Resource specification and intelligent user interaction for federated testbeds using Semantic Web technologies," *Computer Networks*, vol. 63, pp. 84–100, Jan. 2014.

[24] Y. Xin, I. Baldin, J. Chase, *et al.*, "Leveraging Semantic Web Technologies for Managing Resources in a Multi-domain Infrastructure-as-a-Service environment," Renci, Tech. Rep., 2014, pp. 1–20. arXiv: arXiv:1403.0949v1.

[25] I. Niles and A. Pease, "Towards a standard upper ontology," *Proceedings of the international conference on Formal Ontology in Information Systems (ACM)*, vol. 2001, pp. 2–9, 2001.

[26] D. Palma and T. Spatzier, *Topology and Orchestration Specification for Cloud Applications (TOSCA) Version 1*, Nov. 2013.

[27] L. Peterson, S. Sevinc, J. Lepreau, *et al.*, "Slice-based Federation architecture," GENI, Tech. Rep., 2009.

[28] K. Alexander, R. Cyganiak, M. Hausenblas, *et al.*, "Describing linked datasets with the void vocabulary," *W3C Interest Group Note*, 2011.

[29] A. Willner and T. Magedanz, "FIRMA: A Future Internet Resource Management Architecture," in *Workshop on Federated Future Internet and Distributed Cloud Testbeds (FIDC)*, Karlskrona, Sweden: IEEE Xplore Digital Library, Sep. 2014.