

Article

Using Semantic Web Technologies to Query and Manage Information within Federated Cyber-Infrastructures

Alexander Willner^{1,2,†,*}, Mary Giatili^{3,‡}, Paola Grosso^{5,‡}, Chrysa Papagianni^{3,4,‡}, Mohamed Morsey^{5,‡}, Ilya Baldin^{6,‡}

¹ Fraunhofer FOKUS, Berlin, Germany; alexander.willner@fokus.fraunhofer.de

² Technische Universität Berlin, Germany; alexander.willner@tu-berlin.de

³ NETwork Management and Optimal DEsign Laboratory, National Technical University of Athens, Greece; mgiatili@netmode.ntua.gr, chrisap@noc.ntua.gr

⁴ The Institute for Systems Research, University of Maryland, USA; chrisap@isr.umd.edu

⁵ Informatics Institute, University of Amsterdam, Netherlands; {p.grosso | m.morsey}@uva.nl

⁶ RENCI/UNC, Chapel Hill, NC, USA; ibaldin@renci.org

* Correspondence: alexander.willner@fokus.fraunhofer.de; Tel.: +49-30-3463-7116

† Current address: Kaiserin-Augusta-Allee 31, 10589 Berlin, Germany

‡ These authors contributed equally to this work.

Academic Editor:

Version June 21, 2017 submitted to Information

Abstract: A standardized descriptive ontology supports efficient querying and manipulation of data from heterogeneous sources across boundaries of distributed infrastructures, particularly in federated environments. In this article, we present the Open-Multinet (OMN) set of ontologies, which were designed specifically for this purpose as well as to support management of life-cycles of infrastructure resources. We present their initial application in Future Internet testbeds, their use for representing and requesting available resources, and our experimental performance evaluation of the ontologies in terms of querying and translation times. Our results highlight the value and applicability of Semantic Web technologies in managing resources of federated cyber-infrastructures.

Keywords: Semantic Web, Federated Infrastructures, RDF, OWL, SPARQL

1. Introduction

Cloud computing supports many distributed applications that are vital to the economy and to the advancement of science. The rising popularity of cloud computing and the diversity of available resources create an urgent need to match those resources optimally to the requests of end-users.

The desired level of self-serve operation within the cloud obviates the need for intervention by IT departments, allowing end-users direct and independent access to the computational infrastructure. As a result, end-users can deploy the necessary infrastructure and software solutions rapidly. Toward this end, accurate modeling of the infrastructure must support abstract representation of various resources, simplify interactions with them, and expose the right levels of information. The next frontier in cloud computing lies in supporting widely distributed clouds and the various aspects of the architectures needed to manage resources across multiple administrative domains. These problems are also closely related to future Internet research in academia, as well as to emerging commercial technologies like the Internet of Things (IoT) [1].

Modeling cloud infrastructures in a manner that supports effective matching of users' requests with available resources is a challenging task. The issue becomes even more complex in the context

25 of distributed cloud systems with multiple infrastructure owners. In the academic research the
26 same problem is encountered when trying to describe computational resources, scientific instruments
27 and testbeds, which belong to different institutions and must be used by inter-disciplinary and
28 inter-institutional collaborative teams. In such environments each infrastructure owner may model
29 resources using their particular information and data modeling approach to set up the system quickly
30 and attract users. The end-result, however, is user lock-in and inability to easily leverage available
31 resources if they belong to different owners. Thus, resource matching and recommendation *based on*
32 *common models* becomes of great importance.

33 This paper describes Open-Multinet (OMN) [2], a set of ontologies that rely on Semantic Web [3]
34 technologies. It was designed by an international team of academic researchers who are intimately
35 familiar with the related problems. The OMN researchers are also involved in multiple efforts to
36 design a federation of Future Internet and cloud testbeds spanning the US and the EU, to be used for
37 at-scale experimentation with novel concepts in networking and distributed systems. While we briefly
38 introduced the ontology set in [2] and presented a preliminary description of its application in the
39 context of a federated cloud environments in [4], in this paper we complement our previous work by
40 an extended description of the OMN ontology set and we further added new evaluation results of the
41 overall OMN framework.

42 Motivation for our work comes largely from our experience with the growth of academic
43 networking, including the proliferation of cloud testbeds. Their *ad hoc* attempts to federate with
44 each other, *i.e.*, to make their resources available to wider communities of users through common
45 interfaces, suffer from a lack of common models to describe available resources. Testbed owners use
46 such models chiefly to provide their users with information about available resources, *e.g.*, compute
47 nodes, storage, network switches, and wireless access points. Each user, in turn, employs similar
48 models to request resources from the testbeds, describing in some detail the exact configuration of
49 available resources needed from the testbed.

50 Most testbeds are small when they first launch. Their designers often spend little time thinking
51 through the information model that they wish to use to present resource information to users. Testbeds
52 frequently rely on home-brewed solutions utilizing syntactic schema specifications serialized using
53 XML or JSON, sometimes referred to as *RSpecs*, although RSpec is also a name of a specific XML dialect
54 used by a subset of testbeds in the US. Documents expressed in those languages are passed between
55 the users and the testbed management software in order to describe the available resources and to
56 request specific configurations of resources for the experiments. While the built-in mechanisms in those
57 languages allow for straightforward verification of document syntax, few mechanisms are available
58 for validation of semantic correctness. These solutions typically rely on *structure-implied semantics* to
59 validate correctness by associating semantic meaning rigidly with the position of information elements
60 within the document.

61 These approaches tend to work in early phases of the design. As the diversity of resources
62 grows, however, and as the sophistication of users increases, the need arises for extension mechanisms.
63 Demand emerges for more powerful resource descriptions. The extension mechanisms then inevitably
64 relax the structure-implied semantics, thus making validation of documents progressively more
65 difficult. We observed this development first-hand in the case of US Global Environment for
66 Network Innovations (GENI) [5] and EU Future Internet Research and Experimentation (FIRE) [6]
67 testbed-federation efforts. XML schema extensions were introduced to allow different federation
68 members to describe the unique attributes of their cloud testbeds. The extensions, we found, made it
69 possible to create syntactically valid but semantically invalid documents requesting resources from a
70 testbed, *e.g.*, by requesting that a particular operating-system image be attached to a network interface
71 instead of to a compute node.

72 Informed by these experiences, we decided to adopt Semantic Web technologies, which provided
73 us with a number of advantages:

- 74 • A common standardized model is used to describe cloud and testbed infrastructures. The
75 extensibility of this model is built into it from the start in the form of additional ontologies that
76 describe new types of resources. The machinery to deal with extensions is built into standard
77 semantic web toolkits, leaving the designers free to think about the information model while
78 affixing the data model.
- 79 • Different resources and descriptions easily can be related and connected semantically. Semantic
80 web mechanisms intuitively represent computer network graph structures. Network topologies
81 are embedded into the RDF graph using graph homeomorphisms and then are annotated with
82 additional information, addressing *structural and semantic constraints* in a single structure.
- 83 • Model errors can be detected early, before testbed resources are provisioned, by using many
84 standard inference tools.
- 85 • Rules can in particular be used to complement queries. Rules for harmonizing relationships
86 should to be defined and applied on the federation level. This is where specialties and
87 commonalities of the involved testbeds are known and this approach lifts the burden from
88 users to formulate complex queries.
- 89 • The annotation process, *i.e.*, the conversion from XML-based RSpecs to RDF-based graphs, is
90 automatic and configurable to take testbed specific extensions and federation-wide agreements
91 into account.
- 92 • Using standard Semantic Web tools, complex queries can be formulated to discover resources. A
93 common way for testbeds to operate is by ingesting JSON/XML or other encoding of the user
94 request or resource advertisement and then converting it into a non-portable native form on
95 which queries and embeddings are performed. Semantic web tools allow us to store testbed-state
96 information natively in RDF and to operate on that information using a multitude of native
97 inference and query tools, thus simplifying and abstracting many parts of testbed operations.
- 98 • Once cloud resources are described semantically, they can be interlinked to other Linked Open
99 Data (LOD) [7] cloud data sets. These linkages provide additional information about resource
100 availability or constraints and help to link resources, *e.g.*, to policies governing their allocation.
- 101 • Semantic resource descriptions support convergence from multiple syntactic-schema based
102 representations of testbed resources to a single semantically enriched representation that
103 combines information from multiple sources. Such sources include various RSpecs describing
104 testbed resources, out-of-band knowledge that may be encoded in resource names or contained
105 in human-readable online Web pages, an approach consistent with Ontology-based Data Access
106 (OBDA). Encoding this information in a structured way into a single representation prepares
107 it for direct analysis, without need of an intermediate representation. Answers are derived by
108 matching resources required by the user to those available at one or more different testbeds,
109 federating the testbeds automatically, with minimal human intervention.

110 We believe that our approach represents an interesting application of OBDA to a novel area of use
111 that combines information search and retrieval with active infrastructure-resource management.

112 The OMN development effort consisted of several phases, starting with the upper ontology
113 design, followed by the design of several critical subordinate ontologies for, *e.g.*, resource monitoring.
114 We relied heavily on previous work in this area, directly incorporating, for instance, the Infrastructure
115 and Network Description Language (INDL) [8,9] ontology for describing computer networks and
116 the Networking innovations Over Virtualized Infrastructures (NOVI) [10] ontology for federated
117 infrastructures. We then started developing tools that utilize the ontology, including converters from
118 the various testbed resource description formats to OMN, inference rules for knowledge extension to
119 complement the conversion process, and rule sets for semantic validation of the documents. We also
120 developed standard queries that assist the testbed resource-matching algorithms in extracting needed
121 information from the testbed resource descriptions.

122 The remainder of the paper is structured as follows. We give a brief overview of related work
123 in the context of (federated) heterogeneous computing infrastructures in Section 2. In Section 3, we
124 present the OMN ontology set. Section 4 shows how we extract information from RSpecs and annotate

125 it using additional knowledge extraction from out-of-band information. Querying and validation using
126 traditional semantic web stools are then performed by the tools built on this framework. Section 5
127 shows the performance and applicability of our tools. Finally, we close in Section 6 with conclusions,
128 considerations, and a description of future work.

129 2. Related Work

130 Many application disciplines shifted the focus from tree-based data models (*e.g.*, XML-based
131 syntactic schemas) to semantic models. This change is reflected in development of ontologies to
132 support, for example, the operation of Grids, Clouds, and now the Internet of Things. These efforts
133 have informed our own OMN development. In the coming section, we provide an overview of these
134 efforts.

135 2.1. Semantic models for Grids, Clouds and IoT

136 In the context of Grid Computing, the Grid Laboratory for a Uniform Environment (GLUE) [11]
137 schema was started 15 years ago to support interoperability between Grid projects by defining a
138 schematic vocabulary with serializations to XML, LDAP, and SQL [12]. A lack of formalism and a
139 consequent inability to perform logical operations on data motivated the transition to Semantic Open
140 Grid Service Architecture (S-OGSA) [13].

141 Semantic Web service discovery [14,15] addresses the automated discovery of Web services that
142 satisfy given requirements. The discovery process uses a matchmaking algorithm to find potential
143 Web services that might solve the problem at hand. Such methods, however, are inadequate to handle
144 the complex interconnected computing infrastructures addressed by our work. Research on matching
145 concentrates mainly on Web services [16], specifically, on semantic similarities between input and
146 output parameters of various services. Our resource-matching involves more than matching available
147 resources to the requirements of the end-user. We also need to identify homeomorphic embeddings
148 of requested topologies within available resource topologies. The combination of such semantic
149 and structural constraints leads to a substantially greater challenge. Pedrinaci *et al.* introduced
150 Linked USDL (Linked USDL) [17], a vocabulary that applies research conducted on Semantic Web
151 services to USDL [18,19]. Linked USDL provides comprehensive means to describe services in support
152 of automated processing. It focuses only on services, and is unsuited to the description of cloud
153 infrastructures. Ontologies such as the Semantic Markup for Web Services (OWL-S) [20] or Good
154 Relations (GR) [21], however, are of interest to our work, and are referenced in part in our ontology.

155 In the domain of Cloud Computing, researchers are working to ensure interoperability on a
156 semantic level. Since 2008, work has progressed in the development of ontologies and of tools
157 for semantic cloud computing [22–24]. Haase *et al.* [25], for example, introduced an approach to
158 administration of enterprise cloud environments, using semantic Web technologies. They proposed
159 a Semantic Web-based product called eCloudManager, which incorporates an ontology to model its
160 cloud data. However, the system and its ontology only focus on the management aspect of cloud
161 systems, and the data are not open for usage. In another example, Haak *et al.* [26] proposed an
162 ontology-based optimization methodology that enables cloud providers to detect the best resource set
163 to satisfy a user's request. Their framework handles only a single administrative domain, whereas we
164 seek to cover a distributed set of provider domains.

165 A paradigm shift is in progress in favor of Intercloud Computing. For instance, 20 approaches
166 to this new challenge are presented in [27]. Within this context, Manno *et al.* proposed the use of
167 the semantic Federated Cloud Framework Architecture (FCFA) [28] to manage resource life cycles
168 based on formal models. In contrast, the Intercloud architecture developed within the Institute of
169 Electrical and Electronics Engineers (IEEE) Standard for Intercloud Interoperability and Federation
170 (P2302) [29,30] Working Group uses graphs to describe and to discover cloud resources based on the
171 existing Open-Source API and Platform for Multiple Clouds (mOSAIC) [31] ontology. Both approaches
172 are being considered as domain-specific extensions to our work. In addition, Santana-Pérez *et al.* [32]

173 proposed a scheduling algorithm that was suitable for federated hybrid cloud systems. The algorithm
174 applies semantic techniques to scheduling and to matching tasks with the most suitable resources. The
175 information model is based on the Unified Cloud Interface (UCI) project ontologies, which cover a
176 wide range of details but which cannot handle Intercloud systems. Le and Kanagasabai [33,34] also
177 proposed ontology-based methodologies to discover and to broker cloud services. They use Semantic
178 Web technologies for user requirements and for cloud provider advertisements, and then apply an
179 algorithm to match each requirement list to advertised resource units. Multiple levels of matching
180 are defined, ranging from an exact match to no match. These methodologies concentrate only on
181 Infrastructure as a Service (IaaS) provisioning. Moreover, they typically neither export their data nor
182 provide a SPARQL Protocol And RDF Query Language (SPARQL) [35] endpoint, thereby hindering
183 reuse of and access to data.

184 Interest has soared recently in the uses and challenges of the Internet of Things in which many
185 heterogeneous devices from different administrative domains communicate with each other. Semantic
186 models are needed for the IoT. The European Research Cluster on the Internet of Things (IERC) has
187 established Activity Chain 4 – Service Openness and Interoperability Issues/Semantic Interoperability
188 (AC4) [36], and semantic models such as the Semantic Sensor Network (SSN) [37] ontology have
189 been developed. Support for semantics in Machine-To-Machine Communication (M2M) [38] has
190 received further attention [39]. The primary applicable standardization activity from the European
191 Telecommunications Standards Institute (ETSI) M2M Working Group has identified the need for
192 a semantic resource descriptions in [40]. The successor, oneM2M¹ [41], already has established
193 the OneM2M Working Group 5 Management, Abstraction and Semantics (MAS). With the recent
194 establishment of the World Wide Web Consortium (W3C) Web of Things (WoT) [42] Working Group,
195 semantic vocabularies will be developed to describe data and interaction models.

196 2.2. OMN background

197 Development of our approach, the OMN ontology set, started within the Federation for FIRE
198 (Fed4FIRE) [43] project. The aim was to extend and to harmonize related work for the FIRE initiative,
199 which has been developed within the context of federated networks and e-infrastructures. Our main
200 motivation was the state of the art in the Future Internet (FI) experimentation context, which considers
201 only simple schema-based models. The Slice-based Federation Architecture (SFA) [44] is the de-facto
202 standard Application Programming Interface (API) for testbed federation. It uses XML-based RSpecs
203 to describe, to discover, and to manage resources in a simple declarative way. However, it cannot
204 support complex queries combining structural and semantic constraints or knowledge analysis. OMN
205 ontology design reuses concepts previously defined in RSpecs, but also leverages significant prior
206 efforts to define ontologies targeting cyber-infrastructure management.

207 The Open Grid Forum (OGF) Network Mark-Up Language (NML) [45] is a well established
208 ontology for modeling computer networks. It provides a framework for definition and description
209 of topologies ranging from simple networks comprising a few nodes and connections to massive,
210 complex networks with hundreds or thousands of nodes and links. The model underwent a thorough
211 review and definition process, finally becoming an OGF standard. While NML lacks concepts and
212 properties required to describe federated infrastructures, OMN adopts NML in order to model the
213 networking aspects of the infrastructure.

214 In comparison with NML, the INDL addresses virtualization of resources and services. It supports
215 description, discovery, modeling, composition, and monitoring of those resources and services. The
216 INDL actually imports NML to describe attached computing infrastructures in a manner that is
217 independent of technology and vendor. It offers the capacity to extend coverage to emerging network

¹ <http://onem2m.org>

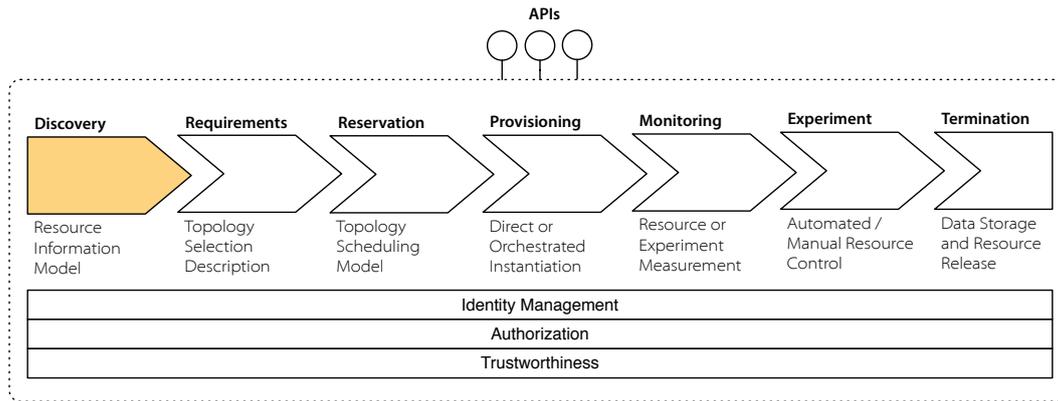


Fig. 1. The experiment life-cycle phases and protocols

218 architectures. The INDL, however, does not support infrastructure federation, in which several
 219 different testbeds are interconnected experimentally.

220 Semantic models developed within the European NOVI and GEYSERS [46] projects have been
 221 used to describe federated infrastructures, user requests, policies, and monitoring information. They
 222 also support virtualization concepts for computing and networking devices. They have been adopted
 223 by OMN where their incorporation is appropriate. In the first project, the proposed Information
 224 Modeling Framework (IMF) [47] represents resources from the same or from different infrastructure
 225 providers.

226 In parallel to this, within the GENI initiative, the Network Description Language based on the
 227 Web Ontology Language (NDL-OWL) [48–51] model specifies capabilities to control and to manage
 228 complex networked testbed infrastructures. Lessons learned from live deployments of NDL-OWL in
 229 GENI proved informative in OMN modeling discussions.

230 Efforts related to describing APIs via OWL-S or DAML-S are not applicable directly to our problem,
 231 since they focus on the description of Web-service APIs. The testbeds in our research converged on
 232 a set of simple API calls like *createSlice* (requesting that a desired network topology be created) or
 233 *listResources* (requesting information about available infrastructure resources). The complexity lies
 234 in the parameters passed in those calls, not in the diversity of types of parameters serving as inputs
 235 or outputs to the APIs. Those parameters often are represented by XML documents describing
 236 requested or available testbed resource topologies. Our goal is to replace those syntactic-schema-based
 237 representations with semantic-web based views. We want them to include enough information to
 238 support native querying based on both structural and semantic constraints, either by the users or by
 239 testbed management algorithms.

240 3. Open-Multinet Ontology Set

241 Following Noy and McGuinness[52], the first step for defining a formal information model is to
 242 determine the specific domain and scope of the proposed ontology. As stated in Sections 1 and 2, the
 243 initial objective was to support resource management in federated infrastructures for experimentation.
 244 The related phases to this management effort are depicted in Figure 1. Each step embodies a wide
 245 range of requirements and challenges. We particularly highlight the first phase in this paper; however,
 246 our approach was to provide a hierarchical set of ontologies to cover the whole resource life-cycle.

247 3.1. Design

248 After identifying the scope of the reusable work (cf. Section 2), we have defined the significant
 249 concepts and properties. Consequently, our ontology bundle consists of nine ontologies, specifically
 250 the *omn* upper ontology and eight descendant ontologies (cf. Figure 2): *omn-federation*; *omn-lifecycle*;

251 *omn-resource*; *omn-component*; *omn-service*; *omn-monitoring*; *omn-policy*; and domain-specific extensions
 252 called *omn-domain-xxx*.

253 These ontologies can be used to describe formally a federation of e-infrastructures, including
 254 types and attributes of resources as well as services available within the federation. Ontologies also
 255 describe the life-cycle phases of usage. The various ontologies extend the upper OMN ontology (solid
 256 lines), which contains common concepts used within the other models. To describe concrete resources
 257 within a particular infrastructure, domain-specific ontologies might need to be defined that use and
 258 extend selected subsets of the OMN ontologies (dotted lines, see Section 3.1.9).

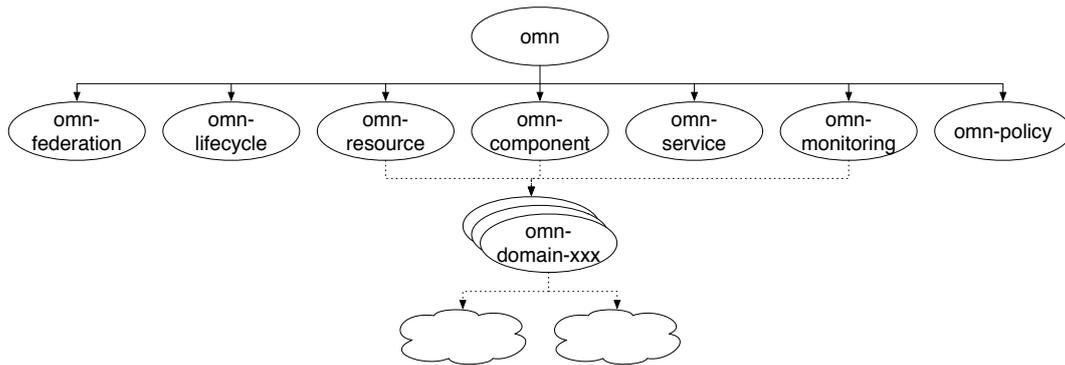


Fig. 2. Open-Multinet ontology hierarchy [4]

259 3.1.1. OMN Upper Ontology

260 The *omn* upper ontology defines the abstract terms required to describe federated infrastructures
 261 in general.

262 It includes a set of classes representing concepts providing general terms to model federated
 263 infrastructures, along with their respective components and services. These concepts are as follows:

- 264 • *Resource*: a stand-alone component of the infrastructure that can be provisioned, *i.e.*, granted to a
 265 user such as a network node.
- 266 • *Service*: is a manageable entity that can be controlled and/or used via either APIs or capabilities
 267 that it supports, such as a SSH login.
- 268 • *Component*: constitutes a part of a *Resource* or a *Service*, such as a port of a network node.
- 269 • *Attribute*: helps to describe the characteristics and properties of a specific *Resource*, *Group*, *Resource*,
 270 or *Component*, such as Quality of Service (QoS).
- 271 • *Group*: is a collection of resources and services, for instance, a testbed or a requested network
 272 topology logically grouped together to perform a particular function.
- 273 • *Dependency*: describes a unidirectional relationship between two elements such as *Resource*,
 274 *Service*, *Component*, or *Group*. It may define, for example, an order in which particular resources
 275 need to be instantiated: first, a network link, and then, the compute nodes attached to it. This
 276 class opens up the possibility of adding more properties to a dependency via annotation.
- 277 • *Layer*: describes a place within a hierarchy to which a specific *Group*, *Resource*, *Service*, or
 278 *Component* can adapt. Infrastructure resources naturally fall into layers, with resources at higher
 279 layers requiring presence of resources at lower layers in order to function.
- 280 • *Environment*: the conditions under which a *Resource*, *Group*, or *Service* is operating, as in, *e.g.*,
 281 concurrent virtual machines.
- 282 • *Reservation*: a specification of a guarantee for a certain duration. Hence, it is a subclass of the
 283 "Interval" class of the W3C Time ontology [53].

284 The OMN upper ontology has 23 properties, of which the following are the most significant:

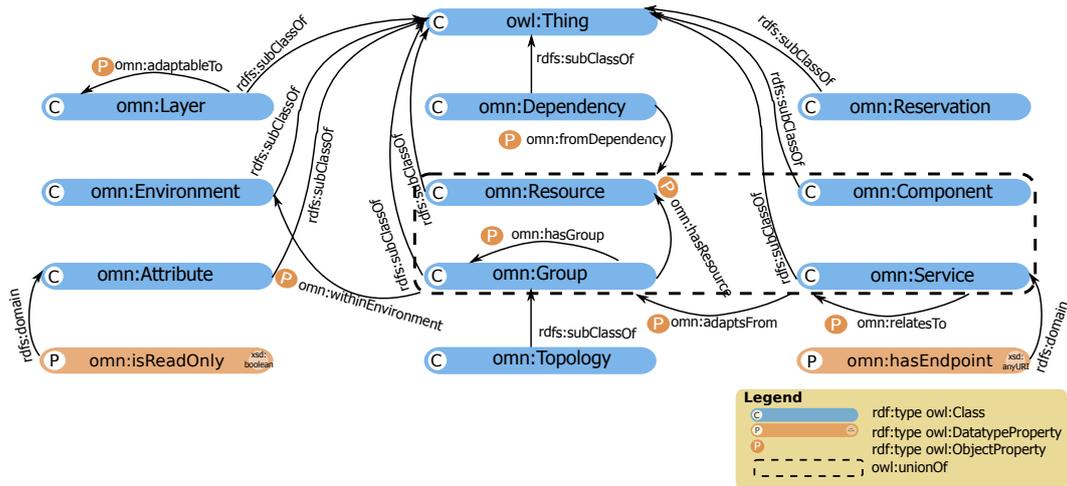


Fig. 3. The key concepts and properties of the omn upper ontology

- 285 • *hasAttribute*: the Attribute associated with a *Component*, *Resource*, *Service*, or *Group*; e.g., CPU
- 286 speed, or uptime.
- 287 • *hasComponent*: links a *Component*, *Resource*, or *Service* to its subcomponent.
- 288 • *hasGroup*: connects a *Group* to its subgroup; it is the inverse of *isGroupOf*.
- 289 • *hasReservation*: relates *Group*, *Resource* or *Service* to its *Reservation*.
- 290 • *hasResource*: declares that a specific *Group* has a *Resource*.
- 291 • *hasService*: declares that a *Group*, *Resource* or *Service* provides a *Service*.
- 292 • *withinEnvironment*: defines the “Environment” in which a *Group*, *Resource*, *Service*, or *Component*
- 293 operates. An example of environment is the operating system under which a resource works.

294 To support rich querying and inferences, inverse counterparts have been declared for most
 295 properties. Figure 3 illustrates the key concepts and properties of the OMN ontology.

296 3.1.2. OMN Federation

297 A crucial part of the developed ontology set is the formal description of the relationship between
 298 the involved e-infrastructure (see Figure 4). This allows to describe how resources relate to each
 299 other from the highest organizational level and depicts the starting point to discover capabilities and
 300 offered services. Federated providers maintain their autonomy in making resource-allocation decisions;
 301 however, they inter-operate with some federation-level functions, such as identity management or
 302 resource advertisement. To model these aspects, the *omn-federation* ontology introduces the concepts
 303 of a *Federation*, *FederationMember*, and *Infrastructure*, along with properties *hasFederationMember* and
 304 *isAdministeredBy*. The first two are subclasses of the schema:Organization class, which allows them to
 305 be described by properties of the Schema vocabulary. The latter concept relates infrastructures to a
 306 federation or to its members, and finally subclasses the *Group* concept which allows infrastructures to
 307 expose services with endpoints, such as an SFA Aggregate Manager (AM).

308 3.1.3. OMN Life Cycle

309 Another important ontology is the *omn-life cycle*, which addresses life-cycle management of a
 310 collection of resources and services (e.g., a requested network topology) that are grouped together
 311 to perform a particular function (e.g., to conduct an experiment or to deploy a service architecture).
 312 The life-cycle of the resources is described by a set of allocation and operational state changes such as
 313 *Allocated*, *Provisioned*, *Unallocated*, *Pending*, *Ready*, *Started*, and *Stopped*. The life-cycle of the collection of
 314 resources reflects the first four phases of Figure 1:

- 315 1. the infrastructure provider advertises an *Offering* describing the available resources;

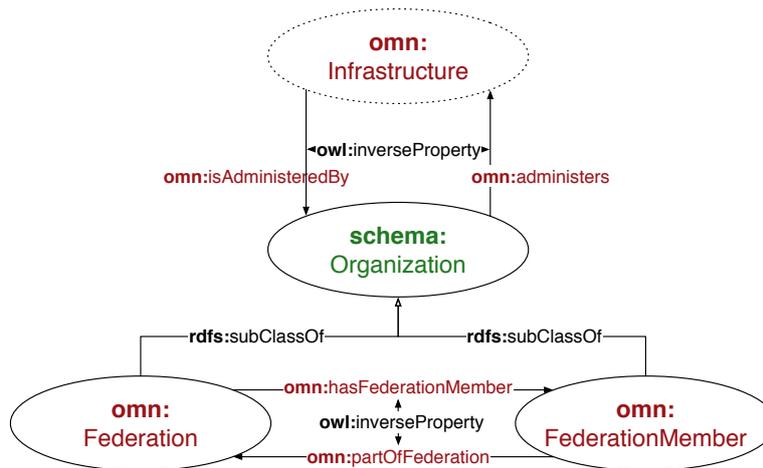


Fig. 4. OMN Federation Ontology

- 316 2. the user forms a *Request* defining the required collection of resources to the infrastructure
 317 provider;
 318 3. the *Confirmation* contains an agreement by the provider, termed bound (tied to a specific set of
 319 physical resources) or unbound, to provide the requested resources;
 320 4. and, finally, a *Manifest* describes the provisioned resources and their properties.

321 Each of these stages is represented as a subclass of the *Group* concept.

322 3.1.4. OMN Monitoring

323 It describes the main concepts and relations to support monitoring services within federated
 324 infrastructures. It includes, therefore, multiple classes and properties that define measurement *Metrics*,
 325 their *Data*, *Unit*, *Tool*, and further *Generic-Concepts*. The monitoring ontology therefore comprises an
 326 upper-level ontology. It describes the common, basic concepts and properties, which then are reused
 327 and specialized in the subjacent ontologies.

328 3.1.5. OMN Resource

329 The OMN Resource ontology deals with the networking aspect of the infrastructure. It supports
 330 creation of complex networks of interconnected resources. It include concepts and properties, *e.g.*,
 331 *Node*, *Link*, and *IPAddress*, which are required for defining complex networks. It also supports defining
 332 single or bi-directional links, which can be utilized for defining the direction of packet flow across the
 333 link(s).

334 3.1.6. OMN Component

335 This ontology describes any entity that is part of a *Resource* or a *Service*; however, in itself, it is
 336 not a *Resource* or a *Service*. The OMN Component ontology describes concepts that are subclasses of
 337 the *Component* class defined in the OMN Upper ontology. It covers several classes to describe a set of
 338 basic entities in any Information and Communication Technology (ICT) infrastructure such as *CPU*,
 339 and *Memory*. Any class or instance of these can be the range of the property *hasComponent* that has a
 340 *Resource*, *Service*, or even another *Component* as a domain.

341 3.1.7. OMN Service

342 This ontology deals with ICT services. Any entity that delivers a value for its user is considered
 343 by the OMN Service ontology as a service. Examples can be services that offer APIs or login
 344 capabilities such as SSH. This ontology includes a set of classes to describe those services being

345 used in ICT infrastructures. The current version covers a set of services being used and implemented
 346 by OMN ontology within the context of the application area addressed in this paper, namely the FI
 347 experimentation.

348 3.1.8. OMN Policy

349 This ontology will cover policy-related concepts and relations. We consider the NOVI policy
 350 ontology as a starting point for its design, as it supports [10]:

- 351 • Authorization policies that specify authorization rights of users within the federation.
- 352 • Event-condition-action policies that enforce control and management actions upon certain events
 353 within the managed environment.
- 354 • Role-based-access control policies that assign users to roles, with different permissions/usage
 355 priorities on resources.
- 356 • Mission policies that define inter-platform obligations in a federation.

357 3.1.9. OMN Domain Specific

358 OMN provides a way to define domain-specific ontologies, which customize the definition of
 359 concepts and relations for a particular ICT application. This allows a set of concepts and relations
 360 that are specific to a particular domain to be grouped along with some concepts and relations from
 361 other OMN ontologies to form a domain-specific ontology. Examples of these ontologies include, for
 362 instance, OMN Wireless ontology and OMN Cloud ontology, used to define the behavior of wireless
 363 networks and of cloud infrastructures, respectively. Another example includes the specification of an
 364 OS version within a disk image, using *omn-domain-pc:hasDiskimageVersion*.

365 3.2. Use of Existing Ontologies

366 As described in Section 2 the OMN ontology set is inspired by and based on a number of
 367 existing formal information models. As an indicator, in Listing 1 a list of referenced vocabularies
 368 are shown that are used within the upper OMN ontology. An OMN Service, for example, has
 369 relationships to *novi:Service*, *dctype:Service*, *gr:ProductOrService*, *service:Service*, *schema:Service*, *nml:Service*,
 370 and *owl-s:Service*.

Listing 1: Used ontologies within *omn.ttl*

```

@prefix : <http://open-multinet.info/ontology/omn#> .
@prefix cc: <http://creativecommons.org/ns#> .
@prefix dc: <http://purl.org/dc/elements/1.1/> .
@prefix gr: <http://purl.org/goodrelations/v1#> .
@prefix nml: <http://schemas.ogf.org/nml/2013/05/base#> .
@prefix owl: <http://www.w3.org/2002/07/owl#> .
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix xml: <http://www.w3.org/XML/1998/namespace#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .
@prefix foaf: <http://xmlns.com/foaf/0.1/> .
@prefix indl: <http://www.science.uva.nl/research/sne/indl#> .
@prefix move: <http://www.ontologydesignpatterns.org/cp/owl/move.owl#> .
@prefix novi: <http://fp7-novi.eu/im.owl#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix time: <http://www.w3.org/2006/time#> .
@prefix vann: <http://purl.org/vocab/vann/> .
@prefix voaf: <http://purl.org/vocommons/voaf#> .
@prefix color: <http://geni-orca.renci.org/owl/app-color.owl#> .
@prefix owl-s: <http://www.daml.org/services/owl-s/1.0DL/Service.owl#> .
@prefix dctype: <http://purl.org/dc/dcmitype/> .
@prefix schema: <http://schema.org/> .
@prefix dcterms: <http://purl.org/dc/terms/> .
@prefix service: <http://purl.org/ontology/service#> .
@prefix collections: <http://geni-orca.renci.org/owl/collections.owl#> .
  
```

371 3.3. Implementation

372 We selected OWL2 to encode the OMN ontology suite due to its expressiveness, wide acceptance,
373 and available tools. To ensure quality, changes to the ontologies are automatically checked using
374 Apache Jena Eyeball inspectors; other validators such as the Ontology Pitfall Scanner (OOPS) [54] are
375 executed manually.

376 As part of the design process we are taking steps to ensure the broadest possible dissemination of
377 the ontologies. As a result, we are using Dublin Core (DC), Vocabulary for Annotating Vocabulary
378 Descriptions (VANN), and Vocabulary of a Friend (VOAF) vocabularies to describe the associated
379 meta information. We are publishing the files by following best practices². The URL [http://open-](http://open-multinet.info/ontology/omn)
380 [multinet.info/ontology/omn](http://open-multinet.info/ontology/omn) provides both a human-readable documentation and machine-readable
381 serializations. We have registered the permanent identifier <https://w3id.org/omn> and published the
382 root ontology to the Linked Open Vocabulary (LOV) repository³. Additionally, we have registered the
383 *omn* name space⁴. The source code of, and an issue tracker for, the ontologies are publicly available⁵.

384 In order to make the work recognizable to the international community, we established the
385 Open-Multinet Forum, which is named after the ontology, and created the W3C OMN Community
386 Group⁶.

387 4. Information Querying and Validation

388 4.1. DBcloud Application For Federated Experimental Infrastructures

389 Most of the requirements for the development of OMN are rooted in research issues within the
390 life-cycle management of resources across federated experimental infrastructures. In such a distributed
391 environment, resource discovery is highly constrained, as it is based on (multi-) attribute matching.
392 It requires an increased level of coordination between users and infrastructure providers as well
393 as among infrastructure providers in the federation. For this purpose, we essentially propose a
394 federation-wide knowledge layer over the federated infrastructures to support semantic representation
395 of such information and to facilitate semantic-based resource discovery.

396 A large amount of semistructured information is available describing the GENI and FIRE testbed
397 federations, including details about the testbeds involved and about the heterogeneous resources
398 offered, reservation information, and monitoring data. This information is encoded mainly as
399 human-readable text on websites as well as in the forms of JSON and XML trees via secured API
400 calls. To extract this information and to make it semantically accessible on the Web, we previously
401 introduced the OMN extraction framework [4].

402 In essence, the OMN extraction framework (Figure 5) follows the design of the DBpedia extraction
403 framework [55]. Information is retrieved from the infrastructures, calling periodically according
404 to methods of the SFA AM API⁷. The downloaded documents are translated into a semantically
405 annotated Resource Description Framework (RDF) [56] graph using the *OMN translator* and the OMN
406 ontology suite. To extend the knowledge encoded in this graph, the Apache Jena inference engine is
407 used within this process by applying infrastructure-specific rules (Section 4.2).

408 Finally, the resulting knowledge graph is written in an in-memory triplet database (Sesame v.2.8.6)
409 and in a Turtle (TTL) [57] serialized file (DBcloud Dump). A SPARQL endpoint on top of the triplet
410 data store implements a **federation-wide lookup service** that enables resource discovery by end-users.
411 The result is currently available at <http://lod.fed4fire.eu> using, among others, the Vocabulary of

2 <http://www.w3.org/TR/swbp-vocab-pub/>

3 <http://lov.okfn.org/dataset/lov/vocabs/omn>

4 <http://prefix.cc/omn>

5 <https://github.com/w3c/omn>

6 <https://www.w3.org/community/omn>

7 http://groups.geni.net/geni/wiki/GAPI_AM_API_V3_DETAILS

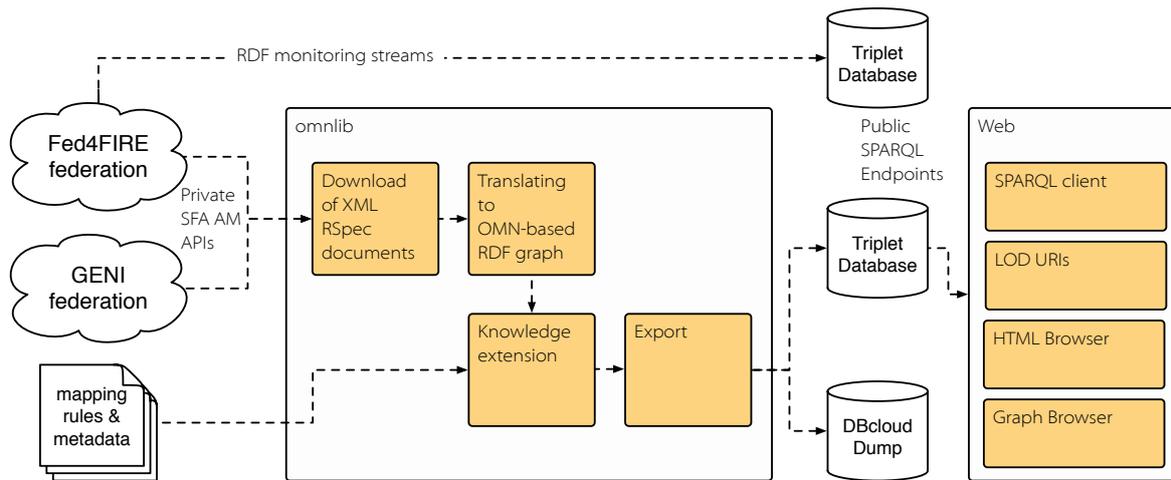


Fig. 5. OMN extraction framework and lookup service (based on [4,55])

412 Interlinked Datasets. The knowledge base currently describes approximately 100 aggregates, 3,000
 413 nodes, 30,000 links, and about 25,000 interfaces. This consists of 4.1 million statements, with the
 414 potential to grow substantially as new testbeds join the federation.

415 The *OMN translator* is a Java-based extensible translation mechanism introduced in [2], allowing
 416 the automated transformation of semi-structured data into an OMN based knowledge graph. It
 417 translates statelessly between GENI, Resource Specifications (RSpecs), and OMN; applies inferencing
 418 rules for validation and knowledge injection; and has been extended to support Topology and
 419 Orchestration Specification for Cloud Applications (TOSCA) [58] and Yet Another Next Generation
 420 (YANG) [59] data models as well.

421 The implementation of the translation tool follows a Test Driven Development (TDD) approach, is
 422 included in a Continuous Integration (CI) environment with test coverage analytics, and is offered as a
 423 Java-based open-source library ("*omnlib*") in a public maven repository. It uses the Java Architecture
 424 for XML Binding (JAXB) and Apache Jena to map between XML, RDF, and Java objects. It supports a
 425 number of APIs: (i) a native API to be included in other Java projects; (ii) a CLI to be used within other
 426 applications; and (iii) a REST-based API to run as a Web service.

Listing 2: RSpec Advertisement (excerpt)

```

<rspec xmlns="http://www.geni.net/resources/rspec/3" type="advertisement">
<node component_manager_id="urn:publicid:IDN+ple+authority+cm"
component_id="urn:publicid:IDN+ple:netmodeple+node+stella.planetlab.ntua.gr"
exclusive="false" component_name="stella.planetlab.ntua.gr"> <----- Node Name
<hardware_type name="http://open-multinet.info/ontology/resources/pc#PC"/> <----- Hardware Type
<sliver_type name="http://open-multinet.info/ontology/resources/plab-vserver#PLAB-VSERVER"/> < Provisions Sliver Type
<available now="true"/>
</node>
</rspec>

```

1
2
3
4
5
6
7
8
9

427
 428 The *OMN translator* parses the XML tree and converts the tags and attributes to their
 429 corresponding classes or properties. To give a better understanding of this translation process, we
 430 provide an illustrative example for the conversion of a GENI Advertisement RSpec used to publish
 431 available resources within a federation of experimental infrastructures.

432 The example in Listing 2 shows a single node of type *PC* that can provision the sliver type
 433 *PLAB-VSERVER* (virtual server for PlanetLab). Traditionally, hardware type and sliver type used
 434 to be simple strings, but unique Uniform Resource Identifiers (URIs) are used here to provide
 435 machine-interpretable information.

Listing 3 shows the converted graph, serialized in Turtle. The overall approach is to define an *omn:Topology* (the subclass *omn-lifecycle:Offering* is used in this case) that contains pointers to the offered resources. Each resource is an individual of a specific type that *can implement* (i.e., can provision) one or more specific sliver types.

Listing 3: OMN Offering

```

<urn:uuid:7eb7b551-7566-4d3c-ac5f-f41a63236baa> <----- Offering
a <http://open-multinet.info/ontology/omn-lifecycle#Offering>;
<http://www.w3.org/2000/01/rdf-schema#label> "Offering";
<http://open-multinet.info/ontology/omn#hasResource> <urn:publicid:IDN+ple:netmodeple+node+stella.planetlab.ntua.gr> .
1
2
3
4
5
<urn:publicid:IDN+ple:netmodeple+node+stella.planetlab.ntua.gr>
a <http://open-multinet.info/ontology/omn-resource#Node> ,
  <http://open-multinet.info/ontology/resources/plab-vserver#PLAB-VSERVER> ,
  <http://open-multinet.info/ontology/resources/pc#PC>;
  <http://www.w3.org/2000/01/rdf-schema#label>
  "stella.planetlab.ntua.gr"^^<http://www.w3.org/2001/XMLSchema#string>;
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
<http://open-multinet.info/ontology/omn-lifecycle#canImplement> <----- Implements Sliver Type
<http://open-multinet.info/ontology/resources/plab-vserver#PLAB-VSERVER>;
<http://open-multinet.info/ontology/omn-lifecycle#hasComponentName>
  "stella.planetlab.ntua.gr";
<http://open-multinet.info/ontology/omn-lifecycle#managedBy>
  <urn:publicid:IDN+ple+authority+cm>;
<http://open-multinet.info/ontology/omn-resource#hasHardwareType> <----- Node Hardware Type
  <http://open-multinet.info/ontology/resources/pc#PC>;
  <http://open-multinet.info/ontology/omn-resource#hasSliverType>
  <http://open-multinet.info/ontology/resources/plab-vserver#PLAB-VSERVER>;
  <http://open-multinet.info/ontology/omn-resource#isAvailable>
  true;
  <http://open-multinet.info/ontology/omn-resource#isExclusive>
  false .
<http://open-multinet.info/ontology/resources/plab-vserver#PLAB-VSERVER>
a <http://open-multinet.info/ontology/omn-resource#SliverType>;
<http://open-multinet.info/ontology/omn-lifecycle#hasSliverName>
  "http://open-multinet.info/ontology/resources/plab-vserver#PLAB-VSERVER" .
<http://open-multinet.info/ontology/resources/pc#PC>
a <http://open-multinet.info/ontology/omn-resource#HardwareType>;
<http://www.w3.org/2000/01/rdf-schema#label>
  "http://open-multinet.info/ontology/resources/pc#PC" .

```

440

441 4.2. Knowledge Extension and Information Querying

442 Having described the framework for semantic-based resource discovery in the context of federated
 443 experimental infrastructures, we will now focus on the specifics of the discovery process. Given a user
 444 request (query) and the aforementioned knowledge base, the resource-discovery problem amounts
 445 to automatically finding the resources from the triplet data store that match the query requirements
 446 along with policies set by infrastructure providers, since a request can be expressed at different levels
 447 of abstraction (*Resource Matching*). The adoption of the OMN ontology suite provides the necessary
 448 flexibility of expression as well as tools for querying and inference that simplify the typical problems
 449 encountered in the process of resource matching. Rules can capture domain background knowledge or
 450 infer resource requirements from the request model; specifically regarding the latter these are added as
 451 additional information to the initial request model. In addition, they can be used to check the request
 452 model's validity [49]. These benefits are highlighted in the following text.

453 4.2.1. Knowledge Extension

454 Background knowledge captures additional knowledge about the domain. This information can
 455 be used in matching a request with available resources. Knowledge is expressed in terms of rules that

456 use the vocabulary of the ontology to add axioms. The knowledge graph can be extended by applying
457 such rules.

458 For example, infrastructure providers in the federation do not advertise explicitly the hardware
459 configurations of their resources in the RSpec XML documents provided. Such data are not translated
460 into RDF. Instead, the information is encoded in each resource's *hardware type*, arbitrarily set by the
461 infrastructure provider as highlighted in the advertisement excerpt provided in Listing 2, (i.e., *hardware*
462 *type: PC*).

463 In Table 1, we provide sample hardware specifications for a subset of the federated experimental
464 infrastructures as they are described by the corresponding infrastructure providers, namely,
465 NETMODE⁸ and Virtual Wall 2⁹ testbeds.

Tab. 1. Resource Specifications.

HW Type	Description
alix3d2	500 MHz AMD Geode LX800, 256 MB DDR DRAM, 1GB flash card storage
pcgen3	2x Hexacore Intel E5645 (2.4GHz) CPU, 24GB RAM, 250GB harddisk

466 Figure 6 depicts a rudimentary *offering* (advertisement) excerpt from the NETMODE infrastructure
467 provide. For the sake of readability, only a single advertised resource is depicted (*omf.netmode.node1*).
468 Moreover, the diagram does not show all the details of the resource description, although it
469 identifies the distinct OMN ontologies used for this purpose, in the upper part the figure. In the
470 excerpt provided the offered resource *omf.netmode.node1* is *managedBy* the infrastructure provider
471 *omf.netmode* (*AMService*) and is part of (*isResourceOf*) the offering (advertisement) identified by
472 *urn:uuid:c9c34c9c-08d6-4dc6-91e2-2e5fac9dd418*. The resource is related via the object property
473 *hasHardwareType* to the *HardwareType* individual with the label *alix3d2*. It is associated (*hasSliverType*) to
474 the *SliverType* individual, with the label *miniPC*, attributed with specific *Disk Image* properties (e.g., OS
475 Voyage). As noted in this example, infrastructures advertise node capacities by their hardware type
476 name (*alix3d2* in this case).

Listing 4: Infrastructure knowledge 1 (excerpt)

```

[rule1:
(?node omnres:hasHardwareType ?hwtype)
(?hwtype rdfs:label ?label)
regex (?label, "pcgen0?3.*") ← ----- For every compute node with a hardware
makeTemp(?cpuComp) ← ----- type that has a label matching "pcgen0?3.*"
->
(?cpuComp rdf:type owl:NamedIndividual)
(?cpuComp rdf:type omncomp:CPU)
(?cpuComp rdfs:label "Intel E5645 CPU")
(?cpuComp omn:hasModelType "Intel E5645") ← ----- Insert standard information about this node
(?cpuComp rdfs:label "Hexa Core Processor") ← ----- type: CPU Type, Core Count, CPU Frequency
(?cpuComp dbp:fastest "2.4"^^xsd:double)
(?cpuComp dbp:fastUnit <http://dbpedia.org/resource/GHZ>)
(?cpuComp omncomp:hasCores 6)
(?cpuComp dbp:arch <http://dbpedia.org/resource/X86-64>)
(?node omn:hasComponent ?cpuComp) ← ----- Link new information to the compute node
]

```

477

8 <http://www.netmode.ntua.gr/testbed>

9 <http://doc.ilabt.iminds.be/ilabt-documentation/virtualwallfacility.html>

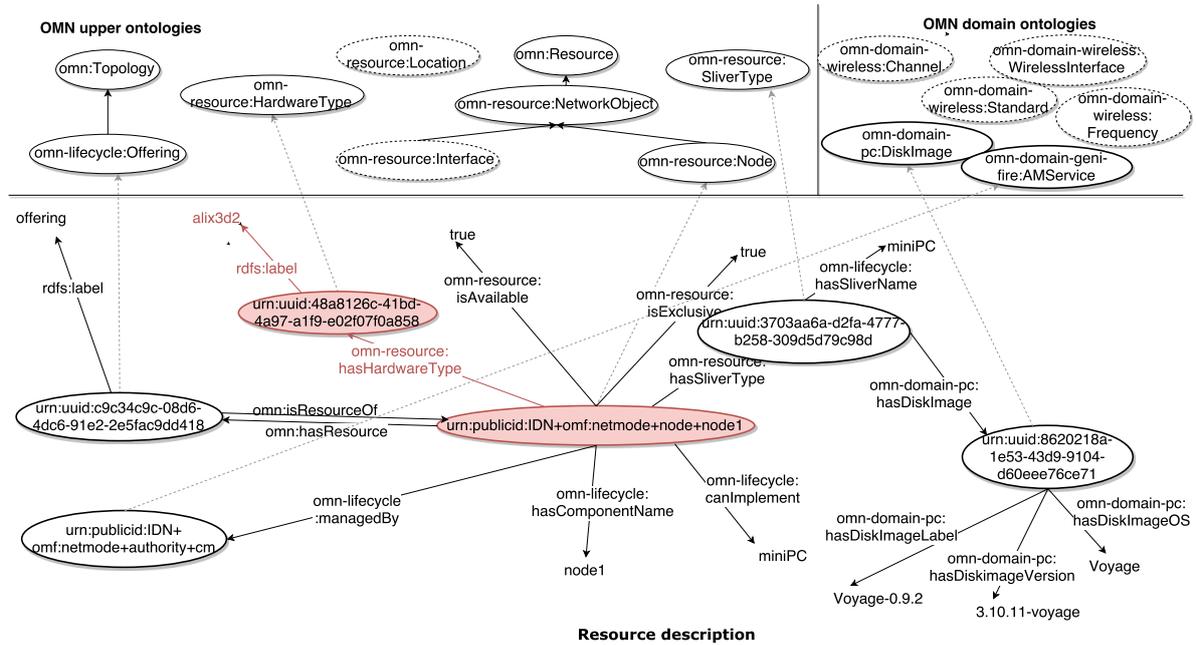


Fig. 6. Partial NETMODE offering

478 A simple example of background knowledge on the context of the "hardware type" is provided
 479 in Listing 4. The listing represents a subset of the rules used to expand the knowledge base with
 480 CPU-related information regarding *pcgen3* nodes listed in Table 1. Such information can be used in
 481 the resource matchmaking process. In the specific application, it is the responsibility of the federator,
 482 which maintains/provides the extraction framework, to apply such rules.

483 In our second example, shown in Listing 5, rules 1 to 3 mandate that each node identified by
 484 hardware type *alix3d2* have the hardware capacity described in Table 1 in terms of CPU, memory, and
 485 storage. Rules 4 to 6 link this information to the node.

Listing 5: Infrastructure knowledge 2 (excerpt)

```

[rule1: uriConcat(omncomp:,"alix3d2_mem", ?memComp) noValue(?memComp rdf:type owl:NamedIndividual)->
(?memComp rdf:type owl:NamedIndividual)
(?memComp rdf:type omncomp:MemoryComponent)(?memComp omnmonunit:hasValue "256000000"^^xsd:integer) ]
[rule2: uriConcat(omncomp:,"alix3d2_cpu", ?cpuComp) noValue(?cpuComp rdf:type owl:NamedIndividual)->
(?cpuComp rdf:type owl:NamedIndividual)
(?cpuComp rdf:type omncomp:CPU) (?cpuComp omnmonunit:hasValue "500000000"^^xsd:integer) ]
[rule3: uriConcat(omncomp:,"alix3d2_sto", ?stoComp) noValue(?stoComp rdf:type owl:NamedIndividual)->
(?stoComp rdf:type owl:NamedIndividual)
(?stoComp rdf:type omncomp:StorageComponent)(?stoComp omnmonunit:hasValue "1000000000"^^xsd:integer) ]
[rule4: (?node omnres:hasHardwareType ?hwtype) (?hwtype rdfs:label "alix3d2"^^xsd:string)
uriConcat(omncomp:,"alix3d2_mem", ?memComp) -> (?node omncomp:hasComponent ?memComp) ]
[rule5: (?node omnres:hasHardwareType ?hwtype) (?hwtype rdfs:label "alix3d2"^^xsd:string)
uriConcat(omncomp:,"alix3d2_cpu", ?cpuComp) -> (?node omncomp:hasComponent ?cpuComp) ]
[rule6: (?node omnres:hasHardwareType ?hwtype) (?hwtype rdfs:label "alix3d2"^^xsd:string)
uriConcat(omncomp:,"alix3d2_sto", ?stoComp) -> (?node omncomp:hasComponent ?stoComp) ]
    
```

4.2.2. Information Querying

486 Having applied the rules in Listing 4, a user may make a request for cloud resources with, for
 487 example, specific CPU requirements. In the sample SPARQL query provided in Listing 6, the user
 488 submits a request for two virtual machines with a specific number of CPU cores and OS type, e.g.,
 489 Fedora:6cores. The results are shown in Listing 7.

490 In a more complex example, a user may submit a request for two nodes running a Linux distribution,
 491 with specific hardware requirements; e.g., 256MB of RAM and storage capacity greater than 500MB. The
 492 query is described in Listing 8. The resource-matching process is not straightforward, as it was in the
 493

494 previous case, even if we apply the rules in Listing 5. In most cases, Infrastructure Providers advertise
 495 the exact Linux Distribution (e.g., Voyage in Figure 6). Thus, the condition for *Linux OS variant* needs
 496 to be either incorporated into the request requirements or advertised explicitly by the testbeds. We
 497 follow the first approach in this case; additional rules are added to infer automatically the resource
 498 characteristics, e.g., acceptable Linux distribution, without explicit statements needed from the user, as
 499 proposed in [60]. The rule set is an appropriately defined set of axioms from which additional implicit
 500 information can be derived. A sample rule used is provided in Listing 9 stating Linux compatibility
 501 (Voyage is a Linux-variant OS).

Listing 6: SPARQL query 1

```

SELECT ?resource1 ?resource2 WHERE {
?resource1 rdf:type omnres:Node . <----- Find me two hosts, resource1 and resource2
?resource2 rdf:type omnres:Node .
?resource1 omnres:hasSliverType/omndpc:hasDiskImage/omndpc:hasDiskimageOS ?os1.
?resource2 omnres:hasSliverType/omndpc:hasDiskImage/omndpc:hasDiskimageOS ?os2.
?resource1 omn:hasComponent ?cpuComp1. <----- Both with 6 cores
?cpuComp1 rdf:type omncomp:CPU.
?cpuComp1 omncomp:hasCores ?cpuvalue1.FILTER (?cpuvalue1 = "6"^^xsd:integer).
?resource2 omn:hasComponent ?cpuComp2.
?cpuComp2 rdf:type omncomp:CPU.
?cpuComp2 omncomp:hasCores ?cpuvalue2.FILTER (?cpuvalue2 = "6"^^xsd:integer).
FILTER (xsd:string(?os1) = "Fedora"^^xsd:string). <----- Both running Fedora
FILTER (xsd:string(?os2) = "Fedora"^^xsd:string).
FILTER (?resource1 = ?resource2)limit 1 ! !<----- Just one answer, please
  
```

502

Listing 7: Query results

```

RESULTS
urn:publicid:IDN+wall2.ilabt.iminds.be+node+n095-05a
urn:publicid:IDN+wall2.ilabt.iminds.be+node+n096-02
TIME EXECUTION: 0.016sec
  
```

Listing 8: Initial SPARQL query 2

```

SELECT ?node1 ?node2 WHERE {
?node1 rdf:type omn_resource:Node. <----- Find me two hosts, node1 and node2
?node2 rdf:type omn_resource:Node.
?node1 omn:hasComponent ?memComp1.
?node2 omn:hasComponent ?memComp2.
?memComp1 rdf:type omn_component:MemoryComponent.
?memComp2 rdf:type omn_component:MemoryComponent.
?memComp1 omn_monitoring_unit:hasValue ?mvalue.
FILTER (?mvalue >= "256000000"^^xsd:integer) <----- Both with RAM greater than 256MB
?memComp2 omn_monitoring_unit:hasValue ?mvalue.
FILTER (?mvalue >= "256000000"^^xsd:integer)
?node1 omn:hasComponent ?stoComp1.
?node2 omn:hasComponent ?stoComp2.
?stoComp1 rdf:type omn_component:StorageComponent.
?stoComp2 rdf:type omn_component:StorageComponent.
?stoComp1 omn_monitoring_unit:hasValue ?svalue1.
FILTER (?svalue1 >= "500000000"^^xsd:integer) <----- Both with disk storage greater than 500MB
?stoComp2 omn_monitoring_unit:hasValue ?svalue2.
FILTER (?svalue2 >= "500000000"^^xsd:integer)
?node1 omn_resource:hasSliverType/omn_domain_pc:hasDiskImage/omn_domain_pc:hasDiskimageOS ?os1. < Both running Linux
?node2 omn_resource:hasSliverType/omn_domain_pc:hasDiskImage/omn_domain_pc:hasDiskimageOS ?os2.
FILTER (xsd:string(?os1) = "Linux"^^xsd:string)
FILTER (xsd:string(?os2) = "Linux"^^xsd:string)
FILTER (?node1 = ?node2)LIMIT 1
  
```

503

Listing 9: IT background knowledge (excerpt)

```

[rule7:(?node rdf:type omn_resource:Node)
(?node omn_resource:hasSliverType ?stype)
(?stype omn_domain_pc:hasDiskImage ?dimage)
(?dimage omn_domain_pc:hasDiskimageOS "Voyage"^^xsd:string) ->
(?dimage omn_domain_pc:hasDiskimageOS "Linux"^^xsd:string)]
  
```

504 Once the rules are applied, OR-AND clauses are built and added to the initial request [61]. Given
 505 the additional information injected into the graph, Listing 10 shows the new, expanded SPARQL query,
 506 with OR-AND clauses included in Lines 22-25. The results are restricted to one feasible matching
 507 solution, which is shown in Listing 11.

Listing 10: SPARQL query 2

```

SELECT ?node1 ?node2 WHERE {
?node1 rdf:type omn_resource:Node. <----- Find me two hosts, node1 and node2
?node2 rdf:type omn_resource:Node.
?node1 omn:hasComponent ?memComp1.
?node2 omn:hasComponent ?memComp2.
?memComp1 rdf:type omn_component:MemoryComponent.
?memComp2 rdf:type omn_component:MemoryComponent.
?memComp1 omn_monitoring_unit:hasValue ?mvalue.
FILTER (?mvalue >= "256000000"^^xsd:integer) <----- Both with RAM greater than 256MB
?memComp2 omn_monitoring_unit:hasValue ?mvalue.
FILTER (?mvalue >= "256000000"^^xsd:integer)
?node1 omn:hasComponent ?stoComp1.
?node2 omn:hasComponent ?stoComp2.
?stoComp1 rdf:type omn_component:StorageComponent.
?stoComp2 rdf:type omn_component:StorageComponent.
?stoComp1 omn_monitoring_unit:hasValue ?svalue1.
FILTER (?svalue1 >= "500000000"^^xsd:integer) <----- Both with disk storage greater than 500MB
?stoComp2 omn_monitoring_unit:hasValue ?svalue2.
FILTER (?svalue2 >= "500000000"^^xsd:integer)
?node1 omn_resource:hasSliverType/omn_domain_pc:hasDiskImage/omn_domain_pc:hasDiskimageOS ?os1. < Running Linux Variant
?node2 omn_resource:hasSliverType/omn_domain_pc:hasDiskImage/omn_domain_pc:hasDiskimageOS ?os2.
FILTER (xsd:string(?os1) = "Voyage"^^xsd:string || xsd:string(?os1) = "Fedora"^^xsd:string || xsd:string(?os1) = "Ubuntu"
^^xsd:string || xsd:string(?os1) = "Linux"^^xsd:string)
FILTER (xsd:string(?os2) = "Voyage"^^xsd:string || xsd:string(?os2) = "Fedora"^^xsd:string || xsd:string(?os2) = "Ubuntu"
^^xsd:string || xsd:string(?os2) = "Linux"^^xsd:string)
FILTER (?node1 = ?node2)LIMIT 1

```

Listing 11: Query results

```

RESULTS
:node1=> <urn:publicid:IDN+omf:netmode+node+node18>,
:node2=> <urn:publicid:IDN+omf:netmode+node+node14>
TIME EXECUTION: 0.299sec

```

508 4.3. Validation

509 Documents created using OMN vocabularies can be validated semantically in part by using
 510 traditional OWL entailments, which verify that the domains and ranges of properties used in a
 511 particular model match those defined in the vocabulary. We found, however, that the expressivity of
 512 those mechanisms was not always sufficient to validate the user requests being sent to the testbed.
 513 Procedural verification is not portable. It is hard to ensure correctness and consistency across
 514 implementations. To supplant traditional OWL mechanisms, we developed Datalog rule-sets that
 515 trigger inference errors when processing a document that either lacks specific information or is
 516 semantically ambiguous. In this section, we explore several examples of such rules.

517 For instance, if a user is attempting to request a network connection that loops to the same node
 518 on which it started, a request may be represented by a valid OMN model; however, semantically, it
 519 doesn't make sense to the resource-matching algorithm that is attempting to reproduce the topology.
 520 To guard against cases like this, we validate the user's request using the following Datalog rule in
 521 Listing 12.

Listing 12: Validating self-looping links in requests

```

(?Z rb:violation error('Connection Validation', 'Connection cannot loop on itself', ?Y))
<- (?X rdf:type pc:PC), (?X nml:hasOutboundPort ?P1), (?X nml:hasInboundPort ?P2),
(?Y rdf:type nml:Link), (?P1 nml:isSink ?Y), (?P2 nml:isSource ?Y) ]

```

522 In some requests by end-users, every Virtual Machine node must specify an OS image to be
 523 booted. At the same time, a VM-server node does not need an image, since it operates using only
 524 a pre-determined image. The *pc : hasDiskImage* property is defined for all *PC* types, including VM
 525 Servers and VMs, so a cardinality restriction cannot be used in this case. This request validation rule is
 526 expressed as follows in Listing 13.

Listing 13: Validating presence of OS image in VM requests

```
(?Z rb:violation error("Validating that VM nodes have OS images", ?R)) <- (?R rdf:type pc:VM),
noValue(?R, pc:hasDiskImage, ?I)
```

1
2

527 It is important to emphasize that the set of the rules that we use continues to evolve with the
 528 schema and with the resource-matching algorithms used to allocate CI resources for the users. For
 529 example, as the algorithms become more sophisticated, they are able to function without some of the
 530 guards protecting them from poorly formed requests, reducing the need for some rules. Nonetheless,
 531 the designing of resource-matching and of embedding algorithms in testbeds is an active field of study.
 532 The availability of declarative rule-based semantic validation significantly simplifies the continuing
 533 evolution of these algorithms by clearly associating a particular algorithm with its own set of validation
 534 rules that prevent errant executions and simplify the algorithm code.

535 5. Performance Evaluation

536 By adopting formal information models and semantically annotated graphs, our approach allows
 537 operations to link, relate, enhance, query, and conduct logical manipulations of heterogeneous data, all
 538 of which would be impossible otherwise.

539 One of the most important measure for the applicability of our work is the amount of time required
 540 to translate and to query resources using our ontology. This time needs to range in a practicable span
 541 for the given context. Our initial work [4] looked at the sizes of the advertisements for testbeds in the
 542 FIRE and GENI projects and evaluated the performance of the translation to RDF of the respective
 543 XML files. The novel work we present in this paper show a more comprehensive comparison of the
 544 queries performance; namely, we look at the time needed to translate resource information to the one
 545 needed to list resources, as well as the performance of queries of different complexity.

546 We have analyzed the result of the *ListResources* method call of the 99 SFA AMs that are
 547 monitored¹⁰ within the Fed4FIRE project. This list contains 82 valid XML based GENI RSpec replies
 548 with 762.634 XML elements in total, of which 3.043 are *Nodes*, 31.155 are *Links*, and 25.493 *Interfaces*.
 549 Figure 7 shows the side of the RSpec advertisements in the testbeds we considered.

550 To estimate the time needed to translate the advertisements, the actual RSpecs from these testbeds
 551 has been downloaded. The XML files were then translated to TTL serialized RDF graphs using
 552 the OMN translator. Of great importance to the potential scalability of our approach is the time
 553 taken for such translations, particularly with regard to the number of XML elements involved. 100
 554 Advertisement RSpecs had been extracted, of which six contained errors, *e.g.*, not adhering to the
 555 RSpec XML Schema Definition (XSD) file, and could not be translated without manual changes. Tests
 556 were run on a MacBookPro with OS X Yosemite, a 2.8GHz Intel Core i7 processor, and 8GB of RAM.
 557 Running a translation over all correct RSpecs produced median values of 24 milliseconds from XML to
 558 Java Architecture for XML Binding (JAXB) and 20 milliseconds from JAXB to RDF, yielding a total
 559 median translation time of 44 milliseconds from XML to RDF. As shown in Figure 8, translation times
 560 appear to be roughly linearly correlated with the number of XML elements translated, with a median
 561 of 180 elements and a maximum of 159,372 translated. This linear correlation indicates upwards scaling
 562 should be possible, although more data are required to confirm this point. At this stage, no major

¹⁰ <https://flsmonitor.fed4fire.eu/>

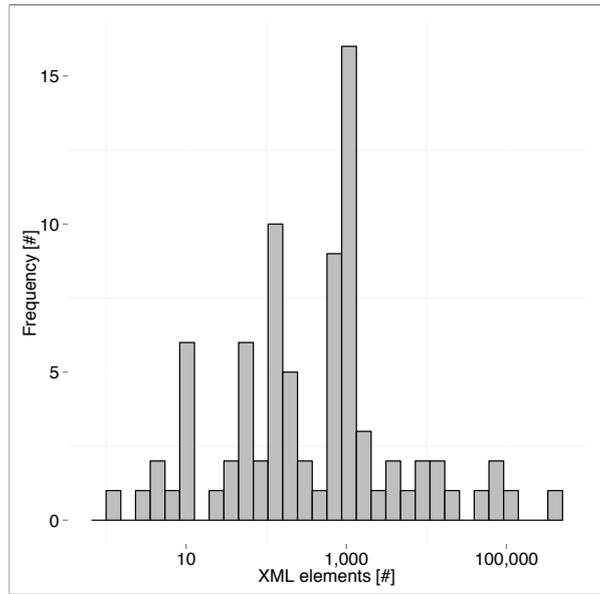


Fig. 7. Size distribution of RSpec Advertisements (logarithmic)

563 limiting factors have been identified, and, given appropriate processing power, translation should be
 564 possible in most foreseeable use cases.

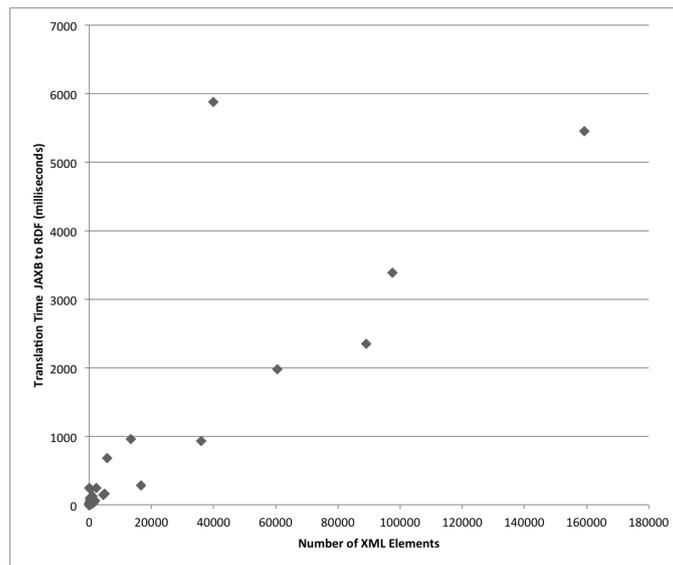


Fig. 8. JAXB to RDF translation times versus number of XML elements [4]

565 To put the duration needed for the translation of an RSpec *Advertisement* into relation with the
 566 duration of the underlying function call needed in the FI experimentation context, we quantified the
 567 query and translation time for a single testbed. As indicated in Figure 7, about 95% of the testbeds
 568 expose fewer than 20,000 XML elements; therefore, we have used the CloudLab Wisconsin testbed¹¹,
 569 which exposes 19,371, for our measurements. The results in Figure 9 show that the average translation
 570 time of 583 ms +/- 9 ms (95% CI) would add about 10% to the average response time of 5453 ms +/-

¹¹ <https://www.cloudlab.us>

571 131 ms (95% CI). This effect, however, could be mitigated by translating in advance or by distributing
572 the work load. The delay of over 5 seconds for listing resources using a single API call, is influenced by
573 mainly two factors. First, the available bandwidth to transmit the resulting XML document from the
574 testbed to the caller. Second, the testbed internal communication architecture to gather the required
575 information, as CloudLab is a distributed infrastructure itself that is composed by three different sites.

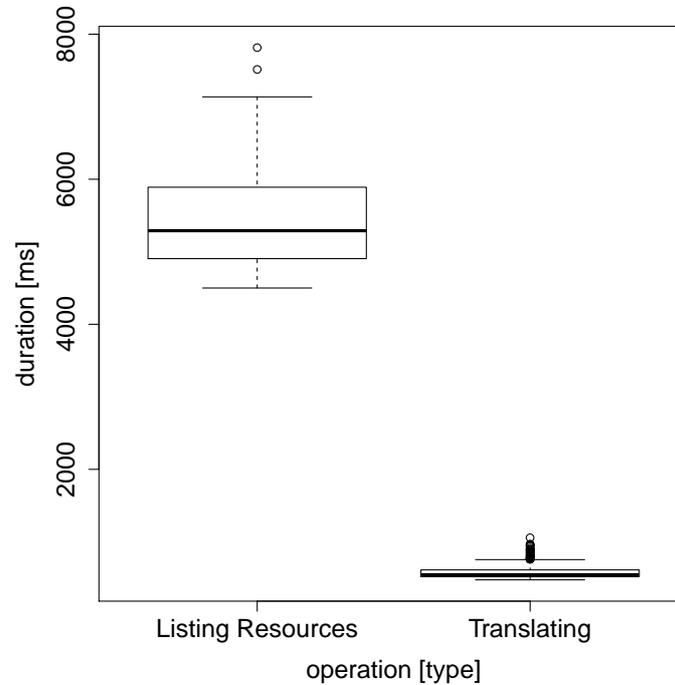


Fig. 9. Listing/translating resource information

576 Assuming that a testbed accepts the potentially enhanced response time, in favor of the added
577 value of merging its information into a global linked data set, its resources can be found by applying
578 the aforementioned resource-matching queries. The translation of all available tree data structures into
579 an RDF-based graph, using our OMN vocabulary and rules, resulted in a set of 2.911.372 statements. It
580 builds the basis for our conclusion that adding further rules, infrastructures, and other data sources
581 will increase the potential for significant growth.

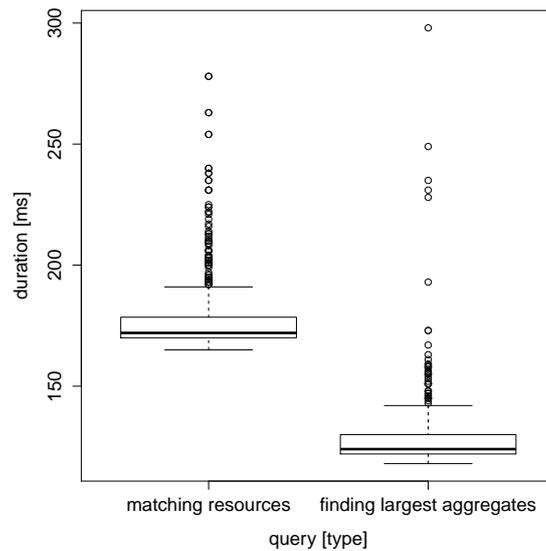


Fig. 10. Performance comparison of queries

582 In Figure 10, the duration of Listing 10 against this graph is shown. To assess the performance
 583 impact of the complexity of the query, it has been compared with a simpler one, which is shown in
 584 Listing 14 together with its result in Listing 15. While finding the three largest aggregates took on
 585 average 129 ms +/- 3 ms (95% CI), the matching query took on average 168 ms +/- 1 ms (95% CI) and
 586 therefore took about 30% longer, yet much less time than a single *ListResources* call in a single testbed.
 587 Finally, we have summarized our findings in Table 2.

Listing 14: Finding the largest aggregate via query

```
SELECT (COUNT(?am) as ?fre) ?am WHERE {
  ?node omn--lifecycle:managedBy ?am .
} GROUP BY (?am) ORDER BY DESC (?fre) LIMIT 3
```

1
2
3

Listing 15: Largest aggregates

```
?fre ?am
719 <urn:publicid:IDN+emulab.net+authority+cm>
326 <urn:publicid:IDN+utah.cloudlab.us+authority+cm>
255 <urn:publicid:IDN+ple+authority+cm>
```

1
2
3
4

Tab. 2. Results of the performance evaluation.

Median duration [ms]	Phase
24	Translation from XML to JAXB (on average)
20	Translation from JAXB to RDF (on average)
583	Translation of 19.371 XML elements (CloudLab)
5453	Listing resources (CloudLab)
129	Querying three largest aggregates (Listing 10))
168	Matching resources (Listing 14)

588 6. Conclusion and Future Work

589 The OMN set of ontologies that we presented in this article has been developed to support
 590 resource management in federated and distributed computing infrastructures. OMN provides a
 591 federation-wide knowledge layer that eases the process of resource selection and matching.

592 In this article, we described the OMN framework, which allows the extraction of underlying
593 information from tree-based data structures. It exposes this information in the form of OMN triples
594 to interested parties via the Web. DBcloud is an application developed in support of the federation
595 of experimental cyber-infrastructures which relies on OMN translators that automatically transform
596 semi-structured data into OMN graphs. An important aspect that we have assessed is the performance
597 of such translations, as this is crucial to OMN usability and adoption. We have shown that the
598 translation and query times require additional time (on the order of 10% in our experiment), which,
599 however, we expect to be acceptable to all resource providers given the added value of merging
600 information.

601 We have also shown how users can query OMN information that represents the resources available
602 in the underlying infrastructures and match them with their own computational requirements. In such
603 case, we evaluated the time needed to find matching resources. We have shown that more complex
604 queries complete within times that are acceptable to end-users.

605 In the long run, we expect that our contributions will outlive the specific use case of the cloud
606 testbed resource management. We believe that it will be accepted by the broader community of
607 academic and commercial cloud providers. It will help to create an ecosystem of flexible, extensible
608 tools and mechanisms that will see the use of cloud platforms become even more pervasive. We expect
609 it to open up the marketplace to competing cloud providers, large and small, catering to specific
610 market niches. We are also promoting adoption of OMN in new domains such as the Internet of Things
611 (IoT). As a specific Industrial Internet of Things (IIoT) [62] example, things, services and data can be
612 connected between federated manufacturing facilities. As an analog to the federation of testbeds the
613 involved facilities, the digital factories, their available APIs and services, have to be described formally
614 to allow for matchmaking capabilities required for the envisioned autonomous production within the
615 fourth industrial revolution. Our ontology set could act as a basis. Following discussions within the
616 German initiative Plattform Industrie 4.0 (PI4.0), linking information based on the LOD paradigm and
617 using interfaces such as the W3C WoT could build a technological base for implementing this vision.
618 Another focus area for the OMN ontologies is the integration with ontologies defining data-access
619 policies among cooperating entities that make use of the cloud infrastructures. The support provided
620 by OMN for the definition of complex usage of heterogeneous resources will be the backbone for novel
621 kinds of open data services, both in industrial and commercial settings, as well as in the scientific
622 community.

623 **Acknowledgments:** Research for this paper was financed in part by the Fed4FIRE (#318389) and Fed4FIREplus
624 (#732638) projects, by the US NSF award CNS-1526964 and GENI funding, and by the Dutch national program
625 COMMIT. We thank our project partners for their contributions and for their collaboration in this work.

626 **Author Contributions:** A.W. initiated development of the presented ontology set and of the DBcloud extraction
627 framework; contributed to related work as well as to the sections about the ontology and validation; and evaluated
628 the performance of the system. M.G. contributed to the development of the upper and domain-specific ontologies,
629 of knowledge extension, and of information querying. P.G. contributed to the development of the ontologies. C.P.
630 contributed to the design and development of the presented ontology set as well as to the ontology, application,
631 knowledge extension, and information-querying sections. M.M. contributed to the development of the ontologies
632 and wrote parts of the ontology section. I.B. contributed to the upper ontology design and provided feedback
633 based on using NDL-OWL ontology in GENI, particularly w.r.t. RSpec translation and request validation.

634 **Conflicts of Interest:** The authors declare no conflict of interest. The funding sponsors had no role in the collection,
635 analyses, or interpretation of the data, in the writing of the manuscript, nor in the decision to publish the results.

636 Abbreviations

637 The following abbreviations are used in this manuscript:

638

639 **AC4** Activity Chain 4 – Service Openness and Interoperability Issues/Semantic Interoperability

640 **AM** Aggregate Manager

641 **API** Application Programming Interface
642 **DC** Dublin Core
643 **ETSI** European Telecommunications Standards Institute
644 **FCFA** Federated Cloud Framework Architecture
645 **Fed4FIRE** Federation for FIRE
646 **FI** Future Internet
647 **FIRE** Future Internet Research and Experimentation
648 **GENI** Global Environment for Network Innovations
649 **GLUE** Grid Laboratory for a Uniform Environment
650 **GR** Good Relations
651 **IaaS** Infrastructure as a Service
652 **ICT** Information and Communication Technology
653 **IEEE** Institute of Electrical and Electronics Engineers
654 **IERC** European Research Cluster on the Internet of Things
655 **IIoT** Industrial Internet of Things
656 **IMF** Information Modeling Framework
657 **INDL** Infrastructure and Network Description Language
658 **IoT** Internet of Things
659 **JAXB** Java Architecture for XML Binding
660 **JSON** JavaScript Object Notation
661 **LDAP** Lightweight Directory Access Protocol
662 **LOD** Linked Open Data
663 **M2M** Machine-To-Machine Communication
664 **MAS** OneM2M Working Group 5 Management, Abstraction and Semantics
665 **mOSAIC** Open-Source API and Platform for Multiple Clouds
666 **NDL-OWL** Network Description Language based on the Web Ontology Language
667 **NML** Network Mark-Up Language
668 **NOVI** Networking innovations Over Virtualized Infrastructures
669 **OGF** Open Grid Forum
670 **OMN** Open-Multinet
671 **OOPS** Ontology Pitfall Scanner
672 **OWL-S** Semantic Markup for Web Services
673 **P2302** Standard for Intercloud Interoperability and Federation
674 **PI4.0** Plattform Industrie 4.0
675 **QoS** Quality of Service
676 **RDF** Resource Description Framework
677 **RSpec** Resource Specification
678 **S-OGSA** Semantic Open Grid Service Architecture
679 **SFA** Slice-based Federation Architecture
680 **SPARQL** SPARQL Protocol And RDF Query Language
681 **SQL** Structured Query Language
682 **SSH** Secure Shell
683 **SSN** Semantic Sensor Network
684 **TOSCA** Topology and Orchestration Specification for Cloud Applications
685 **TTL** Turtle
686 **UCI** Unified Cloud Interface
687 **URL** Uniform Resource Locator
688 **VANN** Vocabulary for Annotating Vocabulary Descriptions
689 **VOAF** Vocabulary of a Friend
690 **VoID** Vocabulary of Interlinked Datasets
691 **W3C** World Wide Web Consortium
692 **WoT** Web of Things
693 **XML** Extensible Markup Language
694 **XSD** XML Schema Definition
695 **YANG** Yet Another Next Generation

696 **References**

- 697 1. Ashton, K. That 'Internet of Things' Thing. *RFiD Journal* **2009**, p. 4986.
- 698 2. Willner, A.; Papagianni, C.; Giatili, M.; Grosso, P.; Morsey, M.; Al-Hazmi, Y.; Baldin, I. The Open-Multinet
699 Upper Ontology Towards the Semantic-based Management of Federated Infrastructures. 10th EAI
700 International Conference on Testbeds and Research Infrastructures for the Development of Networks &
701 Communities (d); ACM: Vancouver, Canada, 2015; p. 10.
- 702 3. Berners-Lee, T.; Hendler, J.; Lassila, O. The Semantic Web. *Scientific American* **2001**, *284*, 34–43.
- 703 4. Morsey, M.; Willner, A.; Loughnane, R.; Giatili, M.; Papagianni, C.; Baldin, I.; Grosso, P.; Al-Hazmi, Y.
704 DBcloud: Semantic Dataset for the cloud. 2016 IEEE Conference on Computer Communications Workshops
705 (INFOCOM WKSHPs); IEEE: San Francisco, CA, 2016; pp. 207–212.
- 706 5. Berman, M.; Chase, J.S.; Landweber, L.; Nakao, A.; Ott, M.; Raychaudhuri, D.; Ricci, R.; Seskar, I. GENI: A
707 federated testbed for innovative network experiments. *Computer Networks* **2014**, *61*, 5–23.
- 708 6. Gavras, A.; Karila, A.; Fdida, S.; May, M.; Potts, M. Future internet research and experimentation. *ACM*
709 *SIGCOMM Computer Communication Review* **2007**, *37*, 89.
- 710 7. Bauer, F.; Kaltenböck, M. Linked open data: The essentials. *Edition mono/monochrom, Vienna* **2011**.
- 711 8. Ghijsen, M.; van der Ham, J.; Grosso, P.; de Laat, C. Towards an Infrastructure Description Language
712 for Modeling Computing Infrastructures. 10th International Symposium on Parallel and Distributed
713 Processing with Applications. IEEE, 2012, pp. 207–214.
- 714 9. Ghijsen, M.; van der Ham, J.; Grosso, P.; Dumitru, C.; Zhu, H.; Zhao, Z.; de Laat, C. A Semantic-Web
715 Approach for Modeling Computing Infrastructures. *Computers and Electrical Engineering* **2013**,
716 *39*, 2553–2565.
- 717 10. van der Ham, J.; Stéger, J.; Laki, S.; Kryftis, Y.; Maglaris, V.; de Laat, C. The NOVI information models.
718 *Future Generation Computer Systems* **2015**, *42*, 64–73.
- 719 11. Andreozzi, S.; Burke, S.; Field, L.; Fisher, S.; Konya, B.; Mambelli, M.; Schopf, J.M.; Viljoen, M.; Wilson, A.
720 GFD 147: Glue Schema Specification. Gfd, Open Grid Forum (OGF), 2007.
- 721 12. Drozdowicz, M.; Ganzha, M.; Paprzycki, M.; Olejnik, R.; Lirkov, I.; Telegin, P.; Senobari, M.; Others.
722 Ontologies, agents and the grid: An overview. In *Parallel, Distributed and Grid Computing for Engineering*;
723 Saxe-Coburg Publications: Stirlingshire, 2009; pp. 117–140.
- 724 13. Corcho, O.; Alper, P.; Kotsiopoulos, I.; Missier, P.; Bechhofer, S.; Goble, C. An overview of S-OGSA: A
725 reference semantic grid architecture. *Web Semantics: Science, Services and Agents on the World Wide Web* **2006**,
726 *4*, 102–115.
- 727 14. Junghans, M.; Agarwal, S.; Studer, R. Towards Practical Semantic Web Service Discovery. European
728 Semantic Web Conference (ESWC). Springer, 2010, Vol. 6089, *Lecture Notes in Computer Science*, pp. 15–29.
- 729 15. Stollberg, M.; Keller, U.; Lausen, H.; Heymans, S. Two-Phase {Web} Service Discovery Based on Rich
730 Functional Descriptions. Proceedings of the 4th European Semantic Web Conference (ESWC); Franconi, E.;
731 Kifer, M.; May, W., Eds. Springer, 2007, Vol. 4519, *Lecture Notes in Computer Science*, pp. 99–113.
- 732 16. Paolucci, M.; Kawamura, T.; Payne, T.R.; Sycara, K. Semantic matching of web services capabilities.
733 International Semantic Web Conference. Springer, 2002, pp. 333–347.
- 734 17. Pedrinaci, C.; Cardoso, J.; Leidig, T. Linked USDL: A Vocabulary for Web-Scale Service Trading. In *The*
735 *Semantic Web: Trends and Challenges*; Springer, 2014; pp. 68–82.
- 736 18. Cardoso, J.; Barros, A.P.; May, N.; Kylau, U. Towards a Unified Service Description Language for the
737 {Internet} of Services: {Requirements} and First Developments. IEEE SCC. IEEE Computer Society, 2010,
738 pp. 602–609.
- 739 19. Oberle, D.; Barros, A.P.; Kylau, U.; Heinzl, S. A unified description language for human to automated
740 services. *Inf. Syst.* **2013**, *38*, 155–181.
- 741 20. Martin, D.; Burstein, M.; Hobbs, J.; Lassila, O.; McDermott, D.; McIlraith, S.; Narayanan, S.; Paolucci, M.;
742 Parsia, B.; Payne, T.; Others. OWL-S: Semantic Markup for Web Services. Member submission, World Wide
743 Web Consortium (W3C), 2004.
- 744 21. Hepp, M. GoodRelations: An Ontology for Describing Products and Services Offers on the Web. In
745 *Knowledge Engineering: Practice and Patterns*; Springer Berlin Heidelberg: Berlin, Heidelberg, 2008; pp.
746 329–346.

- 747 22. Youseff, L.; Butrico, M.; Da Silva, D. Toward a Unified Ontology of Cloud Computing. *Grid Computing*
748 *Environments Workshop*. IEEE, 2008, pp. 1–10.
- 749 23. Han, T.; Sim, K.M. An ontology-enhanced cloud service discovery system. *International MultiConference*
750 *of Engineers and Computer Scientists*. Springer, 2010, Vol. 1, pp. 17–19.
- 751 24. Ma, Y.B.; Jang, S.H.; Lee, J.S. Ontology-Based Resource Management for Cloud Computing. In *Intelligent*
752 *Information and Database Systems*; Springer, 2011; pp. 343–352.
- 753 25. Haase, P.; Mathäß, T.; Schmidt, M.; Eberhart, A.; Walther, U. Semantic Technologies for Enterprise Cloud
754 Management. *International Semantic Web Conference*, 2010.
- 755 26. Haak, S.; Grimm, S. Towards Custom Cloud Services – {Using} Semantic Technology to Optimize Resource
756 Configuration. *Proceedings of the 8th Extended Semantic Web Conference, ESWC 2011*, 2011.
- 757 27. Grozev, N.; Buyya, R. Inter-Cloud architectures and application brokering: taxonomy and survey. *Software:*
758 *Practice and Experience* **2014**, *44*, 369–390.
- 759 28. Manno, G.; Smari, W.W.; Spalazzi, L. FCFA: A semantic-based federated cloud framework architecture.
760 *International Conference on High Performance Computing & Simulation (HPCS)*. IEEE, 2012, pp. 42–52.
- 761 29. Bernstein, D.; Deepak, V.; Chang, R. Draft Standard for Intercloud Interoperability and Federation (SIIF).
762 Technical report, IEEE P2303, 2015.
- 763 30. Martino, B.D.; Cretella, G.; Esposito, A.; Willner, A.; Alloush, A.; Bernstein, D.; Vij, D.; Weinman, J. Towards
764 an Ontology-Based Intercloud Resource Catalogue – The IEEE P2302 Intercloud Approach for a Semantic
765 Resource Exchange. *International Conference on Cloud Engineering*; IEEE: Tempe, Arizona, 2015; pp.
766 458–464.
- 767 31. Moscato, F.; Aversa, R.; Di Martino, B.; Fortis, T.; Munteanu, V. An analysis of mOSAIC ontology for Cloud
768 resources annotation. *Federated Conference on Computer Science and Information Systems (FedCSIS)*.
769 IEEE, 2011, pp. 973–980.
- 770 32. Santana-Pérez, I.; Perez-Hernández, M.S. A semantic scheduler architecture for federated hybrid clouds.
771 *Cloud Computing (CLOUD)*, 2012 IEEE 5th International Conference on. IEEE, 2012.
- 772 33. Dastjerdi, A.V.; Tabatabaei, S.G.H.; Buyya, R. An Effective Architecture for Automated Appliance
773 Management System Applying Ontology-Based Cloud Discovery. *10th IEEE/ACM International*
774 *Conference on Cluster, Cloud and Grid Computing (CCGrid)*, 2010. IEEE, 2010.
- 775 34. Ngan, L.D.; Kanagasabai, R. {OWL-S} Based Semantic Cloud Service Broker. *2012 {IEEE} 19th International*
776 *Conference on Web Services, Honolulu, HI, USA, June 24-29, 2012*, 2012.
- 777 35. Aranda, C.B.; Corby, O.; Das, S.; Feigenbaum, L.; Gearon, P.; Glimm, B.; Harris, S.; Hawke, S.; Herman, I.;
778 Humfrey, N.; Michaelis, N.; Ogbuji, C.; Perry, M.; Passant, A.; Polleres, A.; Prud'hommeaux, E.; Seaborne,
779 A.; Williams, G.T. SPARQL 1.1 Overview. *Recommendation, World Wide Web Consortium (W3C)*, 2013.
- 780 36. Serrano, M.; Barnaghi, P.; Cousin, P. IoT Semantic Interoperability: Research Challenges, Best Practices,
781 Solutions and Next Steps. Technical report, IERC, 2013.
- 782 37. Compton, M.; Barnaghi, P.; Bermudez, L.; García-Castro, R.; Corcho, O.; Cox, S.; Graybeal, J.; Hauswirth,
783 M.; Henson, C.; Herzog, A.; Huang, V.; Janowicz, K.; Kelsey, W.D.; Le Phuoc, D.; Lefort, L.; Leggieri,
784 M.; Neuhaus, H.; Nikolov, A.; Page, K.; Passant, A.; Sheth, A.; Taylor, K. The SSN ontology of the W3C
785 semantic sensor network incubator group. *Web Semantics: Science, Services and Agents on the World Wide*
786 *Web* **2012**, *17*, 25–32.
- 787 38. Geng Wu.; Talwar, S.; Johnsson, K.; Himayat, N.; Johnson, K.D. M2M: From mobile to embedded internet.
788 *IEEE Communications Magazine* **2011**, *49*, 36–43.
- 789 39. Čačković, V.; Popović, Ž. Abstraction and Semantics support in M2M communications. *36th International*
790 *Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO)*,
791 2013, pp. 404–408.
- 792 40. European Telecommunications Standards Institute. *Machine-to-Machine Communications (M2M);*
793 *Functional Architecture*. Technical Report TS 102 690 V2.1.1, ETSI, 2013.
- 794 41. Swetina, J.; Lu, G.; Jacobs, P.; Ennesser, F.; Song, J. Toward a standardized common M2M service layer
795 platform: Introduction to oneM2M. *IEEE Wireless Communications* **2014**, *21*, 20–26.
- 796 42. Raggett, D. *Web of Things Framework*. Draft charter, World Wide Web Consortium (W3C), 2015.
- 797 43. Vandenberghe, W.; Vermeulen, B.; Demeester, P.; Willner, A.; Papavassiliou, S.; Gavras, A.; Quereilhac, A.;
798 Al-Hazmi, Y.; Lobillo, F.; Velayos, C.; Vico-oton, A.; Androulidakis, G. Architecture for the Heterogeneous

- 799 Federation of Future Internet Experimentation Facilities. Future Network and Mobile Summit (FNMS);
800 IEEE: Lisboa, Portugal, 2013; pp. 1 – 11.
- 801 44. Peterson, L.; Sevinc, S.; Lepreau, J.; Ricci, R. Slice-based Federation architecture. Draft version, GENI, 2009.
- 802 45. van der Ham, J.; Dijkstra, F.; Lapacz, R.; Zurawski, J. GFD.206: Network Markup Language Base Schema.
803 Technical report, Open Grid Forum (OGF), 2013.
- 804 46. Escalona, E.; Peng, S.; Nejabati, R.; Simeonidou, D.; Garcia-Espin, J.A.; Riera, J.F.; Figuerola, S.; de Laat, C.
805 GEYSERS: A Novel Architecture for Virtualization and Co-Provisioning of Dynamic Optical Networks and
806 IT Services. Future Network and Mobile Summit. IEEE, 2011, pp. 1—8.
- 807 47. Garcia-Espin, J.A.; Riera, J.F.; Figuerola, S.; Ghijsen, M.; Demchenko, Y.; Buysse, J.; De Leenheer, M.;
808 Develder, C.; Anhalt, F.; Soudan, S. Logical infrastructure composition layer, the GEYSERS holistic
809 approach for infrastructure virtualisation. Terena Networking Conference (TNC); Open Access TCN:
810 Reykjavík, 2012; pp. 1—16.
- 811 48. Baldine, I.; Xin, Y.; Mandal, A.; Renci, C.H.; Chase, U.C.J.; Marupadi, V.; Yumerefendi, A.; Irwin, D.
812 Networked cloud orchestration: A GENI perspective. Globecom Workshops. IEEE, 2010, pp. 573–578.
- 813 49. Yufeng, X.; Baldine, I.; Chase, J.; Anyanwu, K. TR-13-02: Using Semantic Web Description Techniques for
814 Managing Resources in a Multi-Domain Infrastructure-as-a-Service Environment. Technical Report April,
815 RENCI Technical Report Series, 2013.
- 816 50. Xin, Y.; Hill, C.; Baldine, I.; Mandal, A.; Heermann, C.; Chase, J. Semantic Plane : Life Cycle of Resource
817 Representation and Reservations in a Network Operating System. Technical report, RENCI, 2013.
- 818 51. Xin, Y.; Baldin, I.; Chase, J.; Ogan, K.; Anyanwu, K. Leveraging Semantic Web Technologies for Managing
819 Resources in a Multi-Domain Infrastructure-as-a-Service Environment. Technical report, Renci, 2014,
820 [1403.0949].
- 821 52. Noy, N.F.; Mcguinness, D.L. Ontology Development 101: A Guide to Creating Your First Ontology.
822 Technical report, Stanford, 2001.
- 823 53. Hobbs, J.R.; Pan, F. Time Ontology in OWL. Working draft, World Wide Web Consortium (W3C), 2006.
- 824 54. Poveda-Villalón, M.; Suárez-Figueroa, M.C.; Gómez-Pérez, A. Validating Ontologies with OOPS! In
825 *Knowledge Engineering and Knowledge Management*; Springer, 2012; pp. 267–281.
- 826 55. Lehmann, J.; Isele, R.; Jakob, M.; Jentzsch, A.; Kontokostas, D.; Mendes, P.N.; Hellmann, S.; Morsey, M.; van
827 Kleef, P.; Auer, S.; Others. DBpedia-a large-scale, multilingual knowledge base extracted from wikipedia.
828 *Semantic Web Journal (SWJ)* **2014**, *5*, 1–29.
- 829 56. Cyganiak, R.; Wood, D.; Lanthaler, M. Resource Description Framework (RDF) 1.1 Concepts and Abstract
830 Syntax. Recommendation, World Wide Web Consortium (W3C), 2014.
- 831 57. Beckett, D.; Berners-Lee, T.; Prud'hommeaux, E. Turtle-terse RDF triple language. Team submission,
832 World Wide Web Consortium (W3C), 2008.
- 833 58. Palma, D.; Spatzier, T. Topology and Orchestration Specification for Cloud Applications (TOSCA) Version
834 1. Standard November, OASIS, 2013.
- 835 59. Bjorklund, M. RFC 6020: YANG - A Data Modeling Language for the Network Configuration Protocol
836 (NETCONF). Rfc 6020 (proposed standard), Internet Engineering Task Force (IETF), 2010.
- 837 60. Urbani, J.; van Harmelen, F.; Schlobach, S.; Bal, H. QueryPIE: Backward Reasoning for OWL Horst over
838 Very Large Knowledge Bases. 10th International Semantic Web Conference (ISWC). Springer, 2011, pp.
839 730–745.
- 840 61. Urbani, J.; Harmelen, F.V.; Schlobach, S.; Bal, H. QueryPIE: Backward reasoning for OWL Horst over very
841 large knowledge bases. In Proc. 10th Int. Semantic Web Conf. (ISWC'11). Springer, 2011.
- 842 62. Various. *Industrial Internet of Things*; Springer Series in Wireless Technology, Springer International
843 Publishing: Cham, 2017.