2019 15th International Conference on Distributed Computing in Sensor Systems (DCOSS)

VICINITY: IoT Semantic Interoperability based on the Web of Things

Andrea Cimmino[2], Viktor Oravec[1], Fernando Serena[2], Peter Kostelnik[3], María Poveda-Villalón[2], Athanasios Tryferidis[4], Raúl García-Castro[2], Stefan Vanya[1], Dimitrios Tzovaras[4], and Christoph Grimm[5]
[1]*bAvenir, s.r.o., Bratislava, Slovakia.* {*viktor.oravec,stefan.vanya*}*@bavenir.eu*
[2]*Ontology Engineering Group, Universidad Politécnica de Madrid, Madrid, Spain.*
{*cimmino,fserena,mpoveda,rgarcia*}*@fi.upm.es*
[3]*InterSoft A.S., Košice, Slovakia. peter.kostelnik@intersoft.sk*
[4]*CERTH/ITI - Centre for Research and Technology Hellas/Information Technologies Institute, Thessaloniki, Greece.*
{*thanasic,dimitrios.tzovaras*}*@iti.gr*
[5]*Kaiserslautern University of Technology, Kaiserslautern, Germany. grimm@cs.uni-kl.de*

***Abstract*—Internet of Things ecosystems have been developed under different standards and semantics leading to sparse islands of information. The lack of consensus regarding both standards and semantics hinders the interoperability among Internet of Things ecosystems, preventing the exploitation of the huge potential expected by integrating such ecosystems. In this paper we present the H2020 project VICINITY, a decentralized bottom-up standards-based platform to integrate Internet of Things ecosystems avoiding the tedious task of adapting their semantics. VICINITY offers transparent interoperability among such environments as a service in the cloud. In addition, dynamic discovery of new ecosystems is included in VICINITY. We aim at implementing VICINITY in several real-world pilot scenarios in order to validate our approach.**

***Keywords*-Internet of Things, Web of Things, ecosystems interoperability**

I. Introduction

Nowadays, Internet of Things (IoT) ecosystems have become pervasive in the Web [1]. These infrastructures provide access to data from several physical or virtual devices [2]. IoT ecosystems bring suitable data for end-users, business models, and software agents. Unfortunately, these ecosystems are islands of information from different domains that are developed under different standards. In addition, there is no consensus among which standard should be used when providing access to a new IoT ecosystem.

The Web of Things approach aims at gathering these sparse islands to bring transparent interoperability among IoT ecosystems and the possibility of discovering new ecosystems as well. Web of Things approach handles IoT ecosystems as set of things that are described by means of a Thing Description to establish how it should be used (access) and how interact with such thing (semantics).

Unfortunately, the lack of consensus in IoT standards hinders the interoperability among IoT ecosystems and presents several challenges [3]:

- IoT ecosystems are built under different, often proprietary, non-common standards.
- IoT vendors or systems integrators may be reluctant to share the interface specifications due to intellectual property.
- Large-scale integration imposes rules that are disadvantageous for particular participants.
- Adapting existing ecosystems towards new standards requires significant change of management efforts regarding IoT users and operators.
- Information exchange entails serious privacy issues.

The H2020 project VICINITY implements an open virtual neighbourhood to interconnect IoT ecosystems and smart objects providing transparent interoperability based on the Web of Things. Up to authors' knowledge, VICINITY is the first approach that provides interoperability as a service rather than another standard to be adopted. The underlying idea consists in building a neighbourhood of nodes (IoT ecosystems) that are able to communicate among them through a secure P2P network transparently. VICINITY is based on the Web of Things approach; including new nodes in its cloud requires submitting a Thing Description and a privacy policy. VICINITY brings the opportunity to dynamically discover new ecosystems (always fulfilling defined privacy policies) offering an application layer to software agents, business models and end-users.

VICINITY is part of the IoT European Platforms Initiative[1], which also includes the following projects AGILE [4], BIG IoT [5], bIoTope [6], INTER-IoT [7], SymbIoTe [8], and TagItSmart [9]. As far as we know, all these platforms address IoT ecosystems interoperability differently from VICINITY.

Offering interoperability as a service based on Thing Descriptions allows VICINITY to integrate the existing IoT ecosystems, or new ones, defined under different standards. In addition, thanks to the privacy policies, reluctant vendors do not need to share their whole device specifications, only those willing to be accessed. Furthermore, the P2P network of VICINITY offers a decentralized secure environment for data exchange. The whole VICINITY environment provides fairness to particular participants to be integrated.

[1]http://iot-epi.eu/

2325-2944/19/$31.00 ©2019 IEEE
DOI 10.1109/DCOSS.2019.00061

VICINITY shall achieve its objectives by building and demonstrating a bottom-up, user-driven, decentralized, extensible ecosystem for interoperability in several real-world pilot scenarios in order to be validated. These scenarios include different domains such as health, smart buildings, and smart parking.

The paper is organized as follows: section II introduces the components in which VICINITY relies on to offer IoT ecosystem interoperability as a service; section III describes how VICINITY offers interoperability by means of discovery and data access services; section IV shows several real-world pilot scenarios in which VICINITY is being implemented in order to be validated; finally, section VI recaps our conclusions and forecast future work.

II. Components

VICINITY relies on several components that aim at creating a cloud in which a wide range of nodes may transparently exchange data through a secure channel. The nodes that conform the cloud integrate sparse IoT ecosystems developed under different standards. In addition, the nodes may define privacy policies narrowing down the devices that are allowed to exchange data with them, i.e., their neighbourhood, or the portion of data that is suitable to be shared with other nodes in the cloud.

A. Thing Description

VICINITY relies on the work of the W3C Web of Things (WoT) Interest Group[2], which is promoting the implementation of *Thing Descriptions* that aim at being a standard framework to semantically describe "Things" in the Web, turning them interoperable. *Thing Descriptions* are meant to cover the following aspects: 1) to contain semantic metadata that explicitly specifies the semantics used by a "Thing" in the Web, how to interact with it, its properties, actions, and events; 2) to specify the security requirements, mechanisms to access, and wrap the "Thing" interface in case only a portion is suitable to be shared; 3) to establish the communications, i.e., what kind of protocols and data exchange formats are supported by a "Thing", as well as the endpoints that the "Thing" exposes to access to its resources.

B. Thing Ecosystem Description

VICINITY relies on the concept of *Thing Ecosystem Description* in order to implement a dynamic discovery and a transparent data exchange. The *Thing Ecosystem Description* defines an ecosystem of *Thing Descriptions* that refer to common data, which can be accessed by a specific node. During a discovery task, the *Thing Ecosystem Description* is built up on the fly by aggregating already submitted *Thing Descriptions* in the VICINITY cloud. In addition, every *Thing Ecosystem Description* answers to the privacy policies related to each node in the cloud, narrowing down the access to those nodes that allow being discovered. As a result, the nodes within the cloud must understand the semantics of the *Thing Ecosystem Description* to be interoperable.

[2]https://www.w3.org/TR/wot-thing-description/

C. VICINITY Nodes

VICINITY integrates sparse IoT ecosystems by wrapping them into *VICINITY Nodes*. The *Nodes* consist of a *VICINITY Adapter* that connects an IoT ecosystem with a *Gateway API*. The later component is in charge of communications with the cloud through a secure P2P network. When a new *VICINITY Node* is registered in the cloud, thus a new IoT ecosystem is integrated, submitting a *Thing Description* establishing its setup and a privacy policy is mandatory. When a *VICINITY Node* is created, or updated, its *Thing Description* is stored in a VICINITY component known as *Semantic Agent Platform*.

D. P2P Network

VICINITY relies on a P2P network that ensures a secure communication between peers (VICINITY Nodes) based on configuration of neighbourhood: firstly, when two *VICINITY Nodes* exchange data by means of their *Gateway APIs*; secondly, during a discovery task when a *VICINITY Node* submits its discovery criteria to the *Gateway API Services*.

When a new *VICINITY Node* submits its *Thing Description*, it establishes several security criteria that are defined within; for instance the encryption to be used during a data exchange with other nodes. The P2P network follows this security criteria established in the *Thing Descriptions* thanks to the VICINITY *Neighbourhood Manager* component.

E. Neighbourhood Manager

In VICINITY, nodes exist in the context of a neighbourhood that establishes which other nodes they are allowed to exchange data with. Neighbourhoods are defined in the privacy policy submitted with a new *VICINITY Node*, which is used during a discovery process to fulfil such policies. The idea in which neighbourhoods rely is the concept of friendships in social networks; similarly, *VICINITY Nodes* have "friends" with which they are allowed to exchange their data.

The neighbourhoods are handled by the *Neighbourhood Manager* whose goal is twofold, as depicted in Figure 1. One the one hand, it connects the *Gateway Services API* with the *Semantic Agent Platform*, which is crucial for the discovery task. On the other hand, during a discovery task required by a *VICINITY Node* the *Neighbourhood Manager* filters those *Thing Descriptions* that are retrieved from the *Semantic Agent Platform* and which do not belong to the *VICINITY Node* neighbourhood.

F. Ontology Network

The data to be exchanged among the *VICINITY Nodes* and VICINITY components covers a wide range of domains of

interest; ranging from general domains like time and space to specific custom domains such as the *Thing Ecosystem Descriptions*. As a result, VICINITY relies on a highly modular ontology network that has been built following the LOT (Linked Open Terms) methodology[3].

The *VICINITY ontology network* is organized in a modular way in which each component might represent different levels of detail or granularity. Currently, the VICINITY ontology network consist of five modules, namely: VICINITY core module, Adapters module, Ontology model for Web of Things, VICINITY WoT mappings mode and Ontology model for datatypes. All modules are accessible from the VICINITY ontology network portal[4]. For example, the VICINITY WoT module is general enough to be reused in other projects as it model the Web of Thing entities independently of the VICINITY specific needs. This is also the case for the datatypes modules. However, the core module takes into account particular needs of VICINITY as for example the modelling of neighbourhoods. It is worth mentioning that this module also include cross-domains information as time, space. However, it can be considered a general model as it does not represent information about specific domains in IoT like health, housing sensors, power, etc. which can be seen as a more specific layer of knowledge. This specific layer is represented by the VICINITY adapters module, which models specific information for concrete devices in the domains at hand by extending general classes for devices and properties. Finally, the WoT mappings module is a specific model to define mappings between vocabularies and data sources to be processed by VICINITY components.

Besides the domain-specific ontological requirements, the ontologies developed in VICINITY are based on the following non-functional requirements:

- Modularity: the ontology network is designed as a network of ontologies in which modules, i.e., other ontologies, might be interconnected and refer to others.
- Reuse: ontologies within the ontology network are developed doing our best effort to reuse already existing standard ontologies or models, increasing in this way the interoperability with external systems which also rely on such ontologies. For instance SSN and its module SOSA[5], or existing ontologies for representing units of measure[6].
- Extensibility: the ontologies within the ontology network allow the development of extensions by third-parties.

III. Interoperability as Service

The VICINITY architecture addresses interoperability following an approach whose main goal is to provide a service to both discover and access distributed *VICINITY nodes*, i.e., integrated IoT ecosystems. The main challenges to address when both discover or access *VICINITY nodes* are: 1) Relevant metadata about all the nodes in the cloud must be known; 2) discovery and data access among *VICINITY nodes* must always fulfil privacy policies; 3) node-to-node communications shall be established among *VICINITY nodes* so they may exchange data.

A. Discovery

One of the main challenges of implementing interoperability among IoT ecosystems is to enable consumers to discover, in a distributed and dynamic scenario, those IoT ecosystems that are relevant to their needs but without having any prior knowledge about them.

In order to implement such functionality, VICINITY makes the following assumptions:

- There is a common information model to semantically describe Things, i.e., the *VICINITY ontology network*. Every *VICINITY Node* in the cloud must understand its semantics and refer to them in their *Thing Ecosystem Description*.
- The *Thing Description* must be the framework to be used for describing any *VICINITY Node*. The cloud stores the submitted *Thing Descriptions* in a semantic repository known as *Semantic Agent Platform*; such repository is available to any *VICINITY Node*.
- Only those *VICINITY Nodes* whose *Thing Description* is included in the *Semantic Agent Platform* are suitable to be discovered. Therefore, integrating a new IoT ecosystem as a *VICINITY Node*, or changing an existent one, entails submitting a new *Thing Description* to the *Semantic Agent Platform*.
- Consumers learn and leverage the common model and are aware of the *Semantic Agent Platform*. The *Gateway APIs* of the *VICINITY Node* are the semantic mediators among them, and the *Semantic Agent Platform*. Therefore, the *Gateway APIs* provide an interface for discovery requests.
- *VICINITY Nodes* specify their discovery needs as a search criteria that makes use of the semantic models within the *VICINITY ontology network*. Any *Gateway API* must be able to specify discovery needs as semantic-based search criteria, i.e., a SPARQL query.

The steps that the dynamic discovery entails within VICINITY are depicted by Figure 1, which consist in the following sequence of interactions:

1) An IoT ecosystem provides a search criteria to discover new integrated IoT ecosystems expressed in its own semantics and format.
2) The *Gateway API* receives the search criteria and translates it into a search criteria expressed in SPARQL. Then, the later search criteria is securely

[3]http://lot.linkeddata.es/
[4]http://vicinity.iot.linkeddata.es
[5]https://www.w3.org/TR/vocab-ssn/
[6]http://www.wurvoc.org/vocabularies/om-1.8/

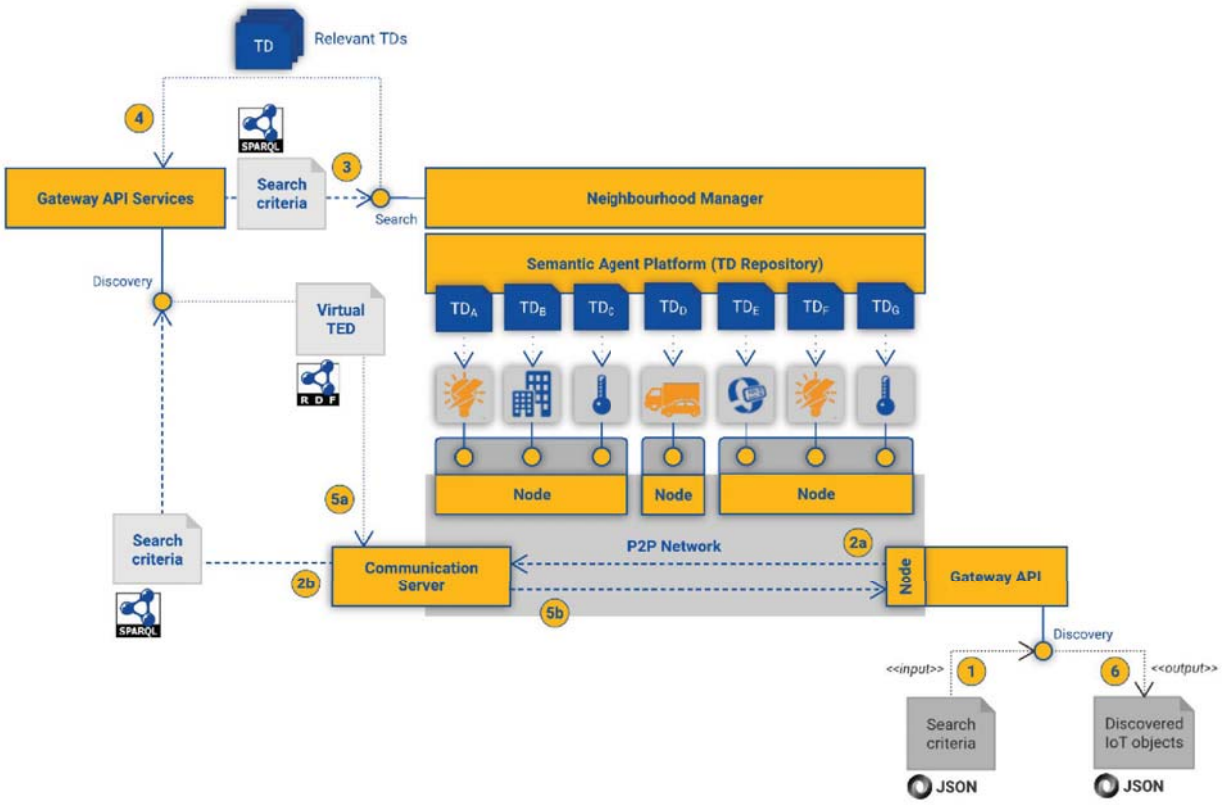

Figure 1. Discovery process in VICINITY.

transmitted trough the *P2P Network* (2a) to the *Gateway API Services* (2b).

3) The *Gateway API Services* receive the search criteria expressed in SPARQL and directly forwards it to the *Neighbourhood Manager*.
4) The *Neighbourhood Manager* retrieves and returns to *Gateway API Services* a set of *Thing Descriptions* that are within the neighbourhood of the *VICINITY Node* that sent the search criteria. In order to retrieve the *Thing Descriptions*, the *Neighbourhood Manager* has to query the *Semantic Agent Platform*. Therefore the *Neighbourhood Manager* is the component responsible of applying the neighbourhood-filtering among the retrieved *Thing Descriptions*.
5) The *Gateway API Services* receives the set of *Thing Descriptions* and builds up the Virtual TED, i.e., *Thing Ecosystem Description*, which is forwarded through the *P2P network* (5a) to the *Gateway API* (5b).
6) Finally, the Virtual TED is processed by the *Gateway API*. As a result, the set of discovered *VICINITY Nodes* that wrap IoT ecosystems, their identifiers, and relevant related metadata is translated from the common VICINITY format to the actual semantics of the specific IoT ecosystem.

B. Access

Another challenge of IoT ecosystems interoperability is the semantic interoperability among them, i.e., exchange data between them transparently. VICINITY ensures a transparent access to heterogeneous IoT ecosystems that rely on different semantics. On the one hand, the information provided and/or required among *VICINITY Nodes* is expressed in terms of its *Ontology Network*. On the other hand, the information received by a *VICINITY Node* can be always translated, and therefore understood, to the IoT-ecosystem-specific semantics.

The data access in VICINITY takes as ground truth the following assumptions in order to transparently exchange data among *VICINITY Nodes*:

- All data endpoints to access an IoT ecosystem referred in a *Thing Description* must be expressed in terms of the *VICINITY ontology network* common format and semantics, and shall be used to represent the data provided/required by such endpoints.

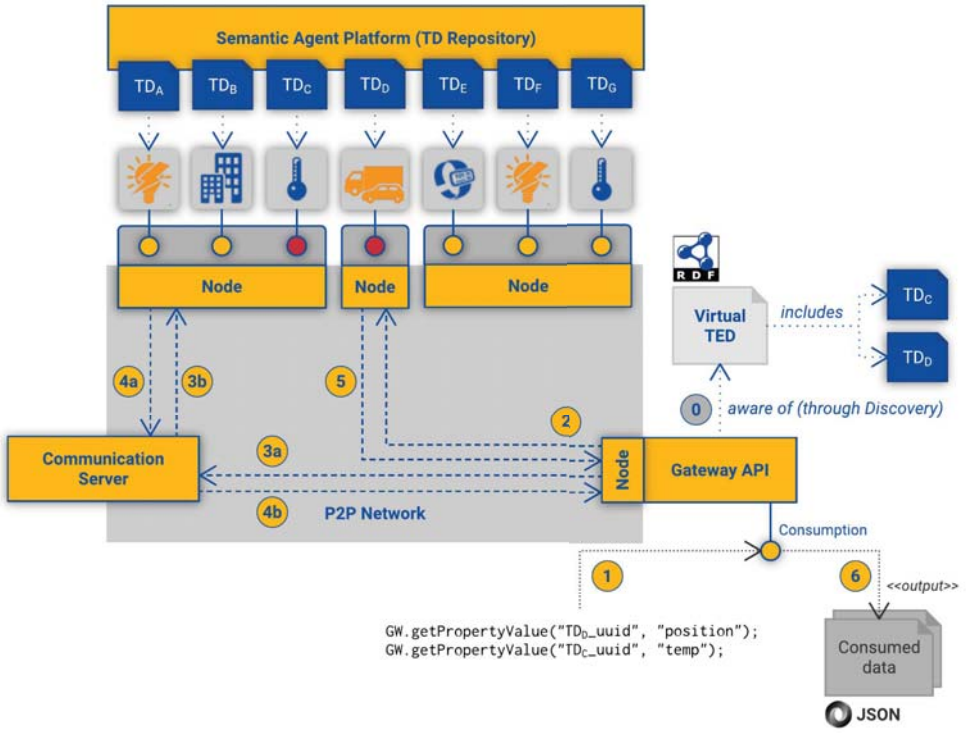

Figure 2. Accessing process in VICINITY.

- *Thing Descriptions* must contain semantic access mappings to explicitly specify the semantics of the data to be exchanged among the *VICINITY Nodes*. Applying such access mappings allows transforming schema-based data into semantic data. The *Thing Descriptions* of IoT ecosystems may include predefined access mappings for each endpoint within. If necessary, during discovery the *Gateway API Services* may attach query-relevant access mappings to the Virtual TED for each *Thing Description* retrieved by the *Neighbourhood Manager*.
- Enriching *Thing Descriptions* with access mappings may entail including concepts and predicates from other vocabularies. In VICINITY the *Ontology Network* comprises different ontologies to which these new concepts and predicates may refer when enriching the context of the *Thing Descriptions*.

Figure 2 depicts an example of how data is accessed in the VICINITY cloud. The example consists of the following steps:

1) The *Gateway API* offers an interface to exchange data with other *VICINITY Nodes*. We assume that the *Gateway API* has previously obtained a Virtual TED by means of a discovery task and thanks to that it knows other *VICINITY Nodes* by means of their *Thing Description* identifiers, i.e., TD_D_uuid and TD_C_uuid, and their properties, i.e., "position" and "temp". In this example the *Gateway API* interface is invoked with two "getPropertyValue" statements, each of which points at two properties that belong to different *VICINITY Nodes*.
2) The *Gateway API* digests the first "getPropertyValue" statement and sends a request to the *VICINITY Node* pointed by the first *Thing Description* identifier. In this case the *P2P Network* allows a directly channel between both *VICINITY Nodes*.
3) Then, the *Gateway API* digests the second "getPropertyValue" statement. In this case the recipient *VICINITY Node* is not directly reachable due to network constraints in-between. Therefore, the Communication Server takes care of the issue and makes sure that the request finally reaches its addressed *VICINITY Node*.
4) The last *VICINITY Node* receives the request forwarded by the Communication Server. Following, the *VICINITY Node* queries the *Thing Ecosystem Descriptor* that describes its own infrastructure so as to determine the specific endpoints that must be invoked. Having all raw data collected from its underlying

infrastructure, the recipient *Gateway API* composes a response message and sends it back to the requester through the Communication Server.

5) The same as described in step 4 with the exception that the response message is directly sent trough the *P2P Network* to the requester *Gateway API*.
6) As a result, the data received in the former *Gateway API* is finally translated into the custom semantics of the related IoT ecosystem.

IV. Pilot scenarios

To validate VICINITY we aim at implementing several pilot scenarios, for example:

- Intelligent Building Systems: Implemented at Oslo Science Park made of four semi-independent buildings and a parking basement garage; 55,000m^2 in total. Goals: 1) energy management, forecasting and consumption; 2) smart parking/booking/electric vehicle charging and optimizing.
- Smart energy systems: targets collaborative management of a community-scale energy ecosystem linking the SOLAR LAB, a demonstration platform, a weather station, and a cluster of municipal buildings. This infrastructure allows data exchange from Generation and Demand sides.
- eHealth and Sensitive Living home: shows how sensors, actuators and integrated communication devices installed at home can provide assisted living to elderly people and people with long term needs, allowing monitoring of health and providing them directly communication with a 24-hours call center.

V. Security and Privacy

Due to nature heterogeneity and distribution of VICINITY platform and connected pears the security and privacy became a complex issue. VICINITY addresses them on the different layers using existing security solutions to ensure and support standardization during integration and operation of the platform:

- Users controlled security and privacy: each IoT user operator defines in Neighbourhood manager the rules who (i.e. services) can access his/her devices (properties, actions and events). These rules are automatically enforced in peer-to-peer network on the level of VICINITY Gateway API. Any change of rules is logged and audited, for later verification. Based on these rules the user data are exchange in peer-to-peer directly between peers without storage of data in VICINITY Cloud infrastructure. Moreover, the IoT operator have access to ”safety button” where he can immediately revoke any access to the device in VICINITY.
- Peer-to-peer network: access to peer-to-peer network si controlled based on identity of device, services and gateway. These identities are defined by neighbourhood manager and enforced by P2P network itself. Confidentiality and integrity of P2P network is ensured by authenticated encryption and PKI infrastructure managed implemented protocol of P2P network. In this layer, confidentiality and integrity is completely transparent for IoT infrastructures and Value-added services. Here we are talking about Gateway API - to - Gateway API encryption. P2P loosing control over security beyond Gateway API. Addressing other security threads in P2P network, VICINITY relies on selected peer-to-peer network protocols and their implementation.
- end-to-end encryption: while exchanged payload between peers in P2P network is based on JSON standard, payload can be signed and encrypted based applying by Javascript Object Signing and Encryption (RFC 7515 and RFD 7516). With proper set-up media type of the messages (i.e. application/jose or application/jose+json) and correct PKI management (e.g. possible extension of the thing description with public key of thing) payload signing and encryption is transparent for any selected or future peer-to-peer engine. This can ensures confidentiality and integrity from source of information to its destination.

VI. Conclusions

The current Internet of Things landscape looks like sparse islands of data that describe physical or virtual devices information. For the Internet of Things to reach its full potential, transparent interoperability among ecosystems is required to generate valuable data that may be consumed by end-users, business models, or software agents. Unfortunately, the lack of consensus regarding which standards should be used to access their data hinders their interoperability.

In this paper we present VICINITY, a secured and decentralized network to interconnect and integrate IoT ecosystems. VICINITY provides transparent interoperability among these ecosystems, allowing the discovery of new ones and transparent access to their data with no prior knowledge of them. VICINITY does not aim at defining another standard for integration, instead it provides interoperability as a service; which as far as we know is the main novelty regarding other proposals. VICINITY is specially suitable because of: 1) does not require IoT ecosystems to be developed under a specific standard, instead only requires a Thing Description describing them to be submitted into the VICINITY cloud; 2) reluctant vendors have the opportunity to hide their system interfaces and still be integrated within VICINITY; and 3) VICINITY offers an on-the-fly discovery service of devices, taking their privacy policies into account.

Having a common data model is the cornerstone of semantic interoperability in the Internet of Things. Therefore, a semantic model for describing things and how to extract and understand relevant information from them is required. In VICINITY we align and contribute to the W3C

Web of Things initiative and use ontologies to represent things, their features and capabilities, and how to gather information from things through their own web interfaces, significantly extending the support to interoperability in the IoT ecosystem.

VICINITY also proposes an ontology-based approach to leverage things discovery and access that is transparent to the syntax, protocols and formats used in things' interfaces. This approach mainly builds on the ability of the involved actors to generate, publish, understand, and query thing descriptions. However, not all actors in VICINITY need to have all these abilities; it shall depend on the role each one plays in the ecosystem, e.g., consumers are not required to have the same abilities as publishers.

ACKNOWLEDGMENT

This research is partially founded by the VICINITY project (H2020-688467), funded by the European Commission Directorate-General for Research and Innovation, under the ICT-30 IoT action of its Horizon 2020 Research and Innovation Programme. The authors acknowledge help and contributions from all partners of the VICINITY project.

REFERENCES

[1] Gazis, V., Görtz, M., Huber, M., Leonardi, A., Mathioudakis, K., Wiesmaier, A., Zeiger, F., Vasilomanolakis, E.: A survey of technologies for the internet of things. In: Wireless Communications and Mobile Computing Conference (IWCMC), 2015 International, IEEE (2015)

[2] Desai, P., Sheth, A., Anantharam, P.: Semantic gateway as a service architecture for IoT interoperability. In: Mobile Services (MS), IEEE International Conference. (2015)

[3] Vermesan, O., Friess, P.: Internet of things-from research and innovation to market deployment. River Publishers Aalborg (2014)

[4] Felfernig, A., Erdeniz, S.P., Azzoni, P., Jeran, M., Akcay, A., Doukas, C.: Towards configuration technologies for IoT gateways. In: 18 th International Configuration Workshop. (2016)

[5] Bröring, A., Schmid, S., Schindhelm, C.K., Khelil, A., Käbisch, S., Kramer, D., Le Phuoc, D., Mitic, J., Anicic, D., Teniente, E.: Enabling IoT ecosystems through platform interoperability. IEEE software (2017)

[6] Werthmann, D., Hellbach, R.: Evaluation Report of the bIoTope Pilots. River Publishers Aalborg (2017)

[7] Ganzha, M., Paprzycki, M., Pawłowski, W., Szmeja, P., Wasielewska, K.: Semantic interoperability in the internet of things: An overview from the INTER-IoT perspective. Journal of Network and Computer Applications (2017)

[8] Soursos, S., Žarko, I.P., Zwickl, P., Gojmerac, I., Bianchi, G., Carrozzo, G.: Towards the cross-domain interoperability of IoT platforms. In: Networks and Communications (EuCNC), 2016 European Conference on, IEEE (2016)

[9] Georgoulas, S., Krco, S., van Kranenburg, R.: TagItSmart–SmartTags for unlocking business potential. Newsletter (2014)