Symbiosis of smart objects across IoT environments

688156 - symbIoTe - H2020-ICT-2015

2nd Open Call – Technical Details

The symbIoTe Consortium

Intracom SA Telecom Solutions, ICOM, Greece
Sveučiliste u Zagrebu Fakultet elektrotehnike i računarstva, UNIZG-FER, Croatia
AIT Austrian Institute of Technology GmbH, AIT, Austria
Nextworks Srl, NXW, Italy
Consorzio Nazionale Interuniversitario per le Telecomunicazioni, CNIT, Italy
ATOS Spain SA, ATOS, Spain
University of Vienna, Faculty of Computer Science, UNIVIE, Austria
Unidata S.p.A., UNIDATA, Italy
Sensing & Control System S.L., S&C, Spain
Fraunhofer IOSB, IOSB, Germany
Ubiwhere, Lda, UW, Portugal
VIPnet, d.o.o, VIP, Croatia
Instytut Chemii Bioorganicznej Polskiej Akademii Nauk, PSNC, Poland
NA.VI.GO. SCARL, NAVIGO, Italy

For more information on this document or the symbIoTe project, please contact:
Sergios Soursos, INTRACOM TELECOM, souse@intracom-telecom.com

The symbIoTe Approach

symbIoTe addresses a challenging objective to create an interoperable Internet of Things (IoT) ecosystem that will allow for the collaboration of vertical IoT platforms towards the creation of cross-domain applications. Thus, it designs an *interoperable mediation framework* to enable the discovery and sharing of connected devices across existing and future IoT platforms for rapid development of cross-platform IoT applications. symbIoTe allows for *flexible interoperability mechanisms* which can be achieved by introducing an incremental deployment of symbIoTe functionality across the platform's space, which will in effect influence the level of platform collaboration and cooperation with other platforms within a symbIoTe-enabled IoT ecosystem. Syntactic and semantic interoperability represent the essential mechanisms in the future symbIoTe-enabled ecosystem, while organizational interoperability has different flavors within symbIoTe (platform federations, dynamic Smart Spaces and roaming IoT devices) to enable platform providers to choose an adequate interoperability model for their business needs.

The symbIoTe approach defines **four interoperability levels**. We also refer to them as **compliance levels**, when considering them from the perspective of an IoT platform wanting to become interoperable. In all four levels, interoperability is achieved by offering a unified and secure way to advertise, discover and consume IoT resources, but in each level a different interoperability scenario is enabled offering various degrees of details about the involved resources.

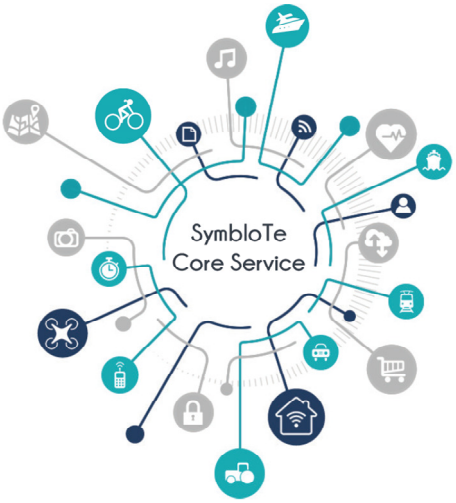

Figure 1: L1 compliance: application-to-resource

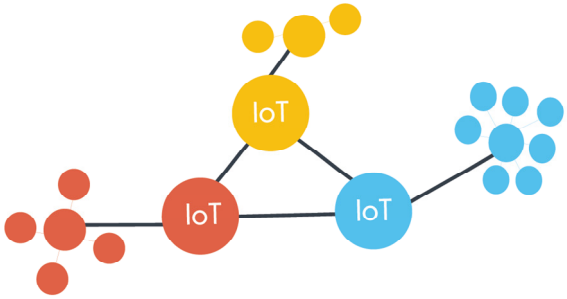

Figure 2: L2 compliance: platform-to-platform

Level 1 (L1) enables interactions between IoT applications and virtualized IoT resources, i.e., third party applications or systems can find and use resources across platforms through uniform interfaces. This is accomplished by the symbIoTe Core Services which allow IoT platforms to register and advertise their offered resources. At the same time, platforms can integrate symbIoTe components to offer uniform and secure access to their virtualized resources. Level 1 can be considered as a search engine for IoT resources.

Level 2 (L2) allows IoT platforms to closely collaborate by forming federations. Federations can be considered as a closed and distributed version of the Core services, i.e., the platforms can advertise only to the members of the federation the resources they want to make available and such resources can be discovered and used only by members

of the federations. The operation of IoT federations is governed by employed trust and resource bartering mechanisms.

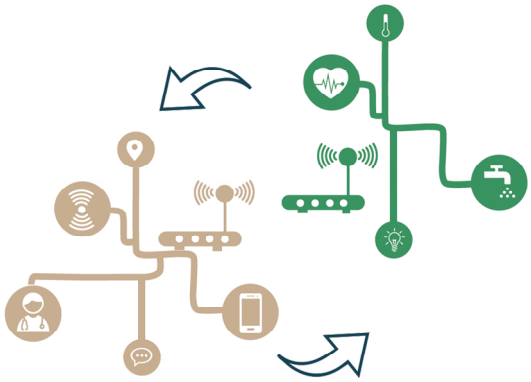

Figure 3: L3 compliance: IoT device-to-IoT device

Figure 4: L4 compliance: roaming IoT device-to-smart space

Level 3 (L3) enables dynamic smart spaces, i.e., the adoption of new resources at the gateway level and the direct interactions between symbIoTe-enabled IoT devices (e.g. mobile devices and Arduino boards) which are collocated in smart spaces, even if they are connected to different gateways and managed by different IoT platforms. This enables resource migration between collocated and in proximity IoT gateways and prevents vendor lock-in.

Level 4 (L4) offers support for roaming of IoT devices (e.g. smartphones and wearables) to interact with smart objects within a visited smart space managed by an IoT platform. A roaming device maintains its unique identity while on the move, and can interact with a smart space only in case the involved IoT platforms are collaborating.

symbIoTe: Technical Details

The symbIoTe Architecture

The symbIoTe architecture is built around a layered IoT stack connecting various devices (sensors, actuators and IoT gateways) within Smart Spaces with the Cloud. Smart Spaces share the available local resources (connectivity, computing and storage), while platform services running in the Cloud will enable IoT Platform Federations (associations between two platforms) and open up the Interworking Interface[1] to third parties. The architecture comprises four layered domains, 1) Application Domain, 2) Cloud Domain, 3) Smart Space Domain and 4) Device Domain, as depicted in Figure 5. Hereafter we list the main functional objectives for each of these domains:

Application Domain (APP): enables platforms to register IoT devices which they want to advertise and make accessible via symbIoTe to third parties, while symbIoTe provides the means for discovery of IoT devices across platforms by its Core Services. Domain-specific back-end services (called 'Domain Enablers') are envisioned to be placed in APP: they utilize the infrastructure provided by the underlying platforms to offer value-added services, e.g. data analytics on top of sensor data acquired from different platforms, which can ease the process of cross-platform and domain-specific application development (specifically for mobile and web applications).

Cloud Domain (CLD): provides a uniform and secure access to virtualized IoT devices exposed by platforms to third parties through an open API (Interworking Interface). In addition, it builds services for IoT Platform Federations enabling direct platform collaboration, in accordance with platform-specific business rules.

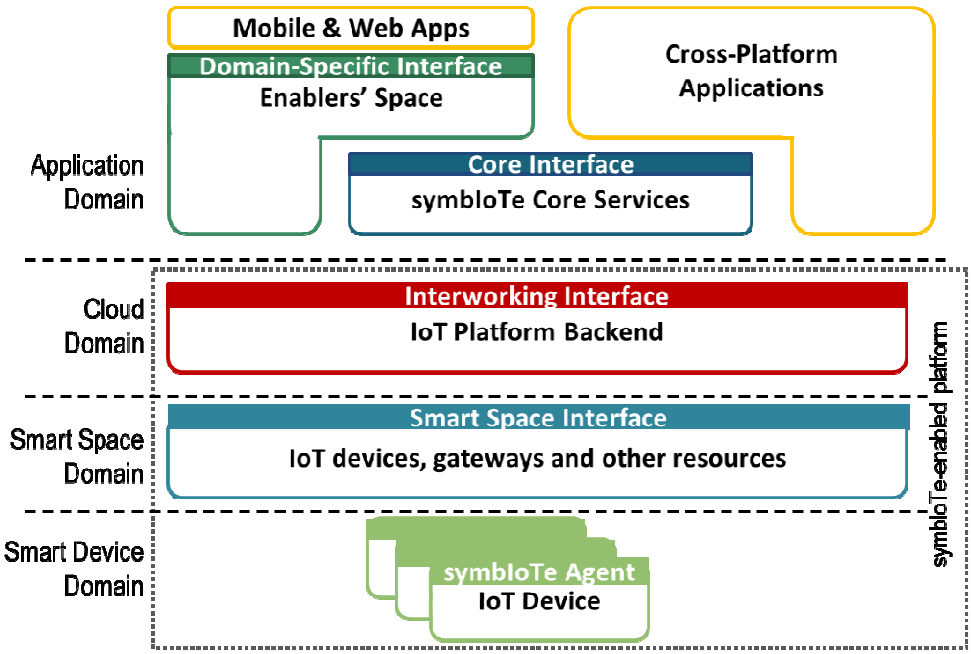

Figure 5: The symbIoTe high-level architecture

[1] Interworking Interface is a symbIoTe defined interface which opens up platform resources as IoT Services in the Cloud Domain.

Smart Space Domain (SSP): provides services for discovery and registration of new IoT devices in dynamic local smart spaces, dynamic configuration of devices in accordance with predefined policies in those environments, and well-documented interfaces for devices available in smart spaces.

Smart Device Domain (SDEV): relates to smart devices and their roaming capabilities. We assume that devices have the capabilities to blend with a surrounding smart space while they are on the move. In other words, smart devices can interact with devices in a visited smart space, which are managed by a visited platform, in accordance with predefined access policies.

Interoperability Aspects

symbIoTe allows for flexible interoperability mechanisms which can be achieved by introducing an incremental deployment of symbIoTe functionality across the listed domains (APP, CLD, SSP and SDEV). This approach will enable platform providers to choose an appropriate level of integration of symbIoTe-specific services within their platforms, which will, in effect, influence the level of platform collaboration and cooperation with other platforms within a symbIoTe-enabled ecosystem. For example, a platform may only choose to expose its Interworking Interface and selected IoT services to third parties in order to advertise them by using the symbIoTe Core Services, or it may opt for a closer collaboration with another platform by forming a platform federation. Platform federations require additional symbIoTe components to be included and integrated within a platform space in CLD.

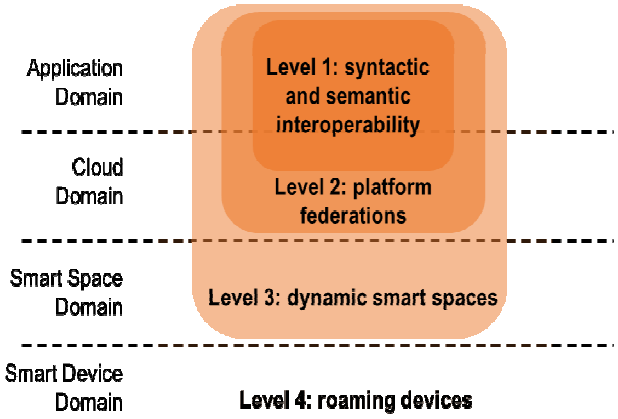

Figure 6: symbIoTe Compliance Levels

We define four different Compliance Levels (CLs) for IoT platforms, as depicted in Figure 6. They reflect different interoperability modes, which an IoT platform can support. Different interoperability modes affect the functionality which needs to be supported by symbIoTe-enabled platforms, and require specific symbIoTe components to be integrated within different domains.

Level 1 (L1) Compliant Platform: This is a "lightweight" symbIoTe CL since a platform opens up only its Interworking Interface to third parties to advertise and offer its virtualized IoT devices through the symbIoTe Core Services. It enables the syntactic and semantic interoperability of IoT platforms in a symbIoTe ecosystem, and affects only APP and CLD.

Level 2 (L2) Compliant Platform: This level assumes that platforms federate, which requires additional functionality to be included in CLD, for example for sharing/bartering of devices, as well as managing federated credentials. The functionality provided at this level enables the so-called organizational interoperability.

Level 3 (L3) Compliant Platform: This CL assumes that platforms integrate symbIoTe components within their smart spaces to simplify the integration and dynamic reconfiguration of IoT devices within local spaces.

Level 4 (L4) Compliant Platform: This level offers support for device roaming and can enable the interaction of smart devices which maintain their unique identity (L4 devices) with a visited smart space. A prerequisite for this interaction is that the smart space is already Level 3 compliant, so that smart spaces can discover visiting L4 devices and integrate them (e.g., grant access to certain local resources). Note that an L4 device may be registered either with symbIoTe Core Services of within a platform federation, and thus its new location and visited smart space should be updated (the idea is conceptually similar to Mobile IP).

OC2: Detailed description of topics

The topics of the symbIoTe's 2nd Open Call (OC2) are summarized as follows:

Topic Identifier	Targeted Applicants	Funding
symbIoTe-OC2-L1	IoT platform owners/operators	≤ €40,000
symbIoTe-OC2-L2	IoT platform owners/operators	≤ €50,000
symbIoTe-OC2-L3/4	IoT gateway manufacturers, Smart device manufacturers, integrators	≤ €40,000
symbIoTe-OC2-Apps	Mobile application companies	≤ €20,000
symbIoTe-OC2-Trials	NGOs, Municipalities, SMEs engaging user communities	≤ €15,000

The first three topics focus on making IoT platforms symbIoTe-complaint. The fourth topic searches for application developers to build innovative mobile apps on top of symbIoTe. The last topic looks for end users to support symbIoTe's planned trials. In the pages that follow we provide a detailed description for each topic.

symbIoTe-OC2-L1

Topic Summary

Purpose: To make 3rd party IoT platforms L1-compliant, so that they expose IoT resources to the symbIoTe Core Services. Applicants should offer platforms active in the Smart City domain, involving (but not limited to) city-wide IoT platforms for environmental monitoring, traffic/parking monitoring, mobility aspects, etc.

Commitment: Applicants should make their platforms and their resources available for a demo during the symbIoTe trials (mid/end 2018).

Type of Applicants: IoT Platform owners / operators

Funding: up to €40,000 per Extension, approx. two (2) Extensions to be funded

IoT platforms that become Level 1 (L1) compliant need to integrate the symbIoTe Interworking Interface with existing CLD components (it is assumed that IoT platforms have their backend hosted in the Cloud). This enables semantic interoperability and uniform access to IoT resources which a platform chooses to register and make discoverable via the symbIoTe Core Services. The access to those resources stays under the control of a platform provider.

APP is designed to offer a unified view on different platforms to a new generation of cross-platform IoT applications. This is achieved by the symbIoTe Core Services that can search for registered IoT resources across platforms. Note that the Core Services store and manage only IoT resource descriptions (i.e. resource metadata), while the access to those resources (e.g., sensor data and actuation) is provided by the underlying platforms. Thus, the symbIoTe Core Services are in close interaction and collaboration with the services

provided within the Cloud Domain which offer the actual access to virtualized IoT resources. In addition to the search functionality, the Core Services implement symbIoTe specific authentication and authorization mechanisms providing the means for secure access to underlying platform-specific resources.

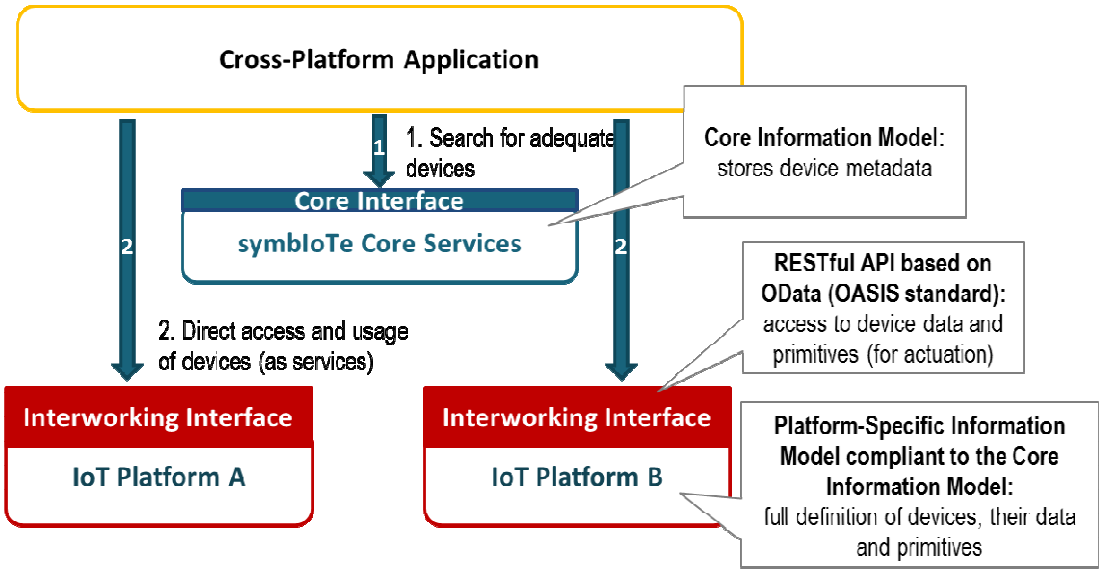

Figure 7: Illustrating Level 1 Compliance

Error! Reference source not found. shows the benefits of L1 compliance by an example depicting two platforms A and B using the symbIoTe Core Services. When an application searches for devices and identifies adequate ones, the application accesses the devices offered by the two platforms through the Interworking Interface. In other words, cross-platform applications i) use the symbIoTe Core Services to find adequate devices across platforms and ii) access, integrate and use those devices through an uniform and open interface. We are supporting a part of the OData standard for pull-based access to platform devices and provide a push-style mechanism for continuous delivery of sensor data to applications over WebSockets.

symbIoTe Information Model. To reach the goal of semantic interoperability, i.e., “the ability of computer systems to exchange data with unambiguous, shared meaning”, symbIoTe defines three types of Information Models related to the description of resources which platforms want to expose through symbIoTe:

- The **Core Information Model** (**CIM**) defines all information that symbIoTe needs to understand on an abstract level, e.g. a class *Sensor* that is related to the class *Location* via a relation/property *hasLocation*.
- **Platform-Specific Information Models** (**PIM**s) are platform-specific and contain all classes and their relations which a platform wants to expose through symbIoTe, e.g., a custom property of a *Sensor* called *hasColor*. A PIM is an extension of the CIM and must comply to it, e.g., by creating new sub-classes that are based on definitions from CIM.
- The **Best Practice Information Model** (**BIM**) is a special form of a PIM that is predefined by and shipped with symbIoTe. Its purpose is to provide a default and simplified way to use symbIoTe whenever a platform does not require having a custom PIM.

When exposing IoT devices to symbIoTe Core Services, devices need to be semantically described. This can be done using one of the following two approaches:

- Using the BIM: As the BIM is a complete information model and known to symbIoTe you will be able to register devices via a simple REST-based call with JSON payload.
- Using a custom PIM: If the BIM does not fit your needs, you can describe the exposed resources of your platform using a custom PIM which must be compliant to the CIM. This must be expressed as an ontology using the Resource Description Format (RDF)[2]. The symbIoTe team can provide support for this task.

symbIoTe components for L1 compliance. symbIoTe Core Services are composed of APP components, which enable the interaction between third-party applications and platforms. The main services of the Core are the following:

1. *Administration* – provides a web-based GUI as well as an API for platform administrators. The provided functionality includes: creation and management of the platform and its properties, generation of platform credentials for authentication.
2. *Registry* – repository for storing all platform-related metadata. It provides an API for registering and updating of platforms and their resources. During the registration process, unique resources' IDs are generated.
3. *Search Engine* – enables applications to find relevant registered devices/services within the Registry. Stores the required information in RDF store.
4. *Semantic Manager* – utility component for validating RDF descriptions used to describe PIM and resources. Also provides translation mechanism between RDF and Java objects.
5. *Resource Monitor* – tracks the availability of registered devices in order to ensure their availability.
6. *Resource Access Monitor* – provides access links to the resources and tracks information about resource popularity.
7. *Core Authentication and Authorization Manager* – authenticates third-party users and applications (i.e., users and applications that are not associated with any IoT platform) and provides credentials required to access symbIoTe Services. It also supports trust relationships between platforms, as it acts as the root certification authority.
8. *Anomaly Detection Module* – is responsible for detecting 0-day attacks and other types of security violations (malicious users, DoS attacks) by using a signature-less machine-learning approach.

The Core Services are designed and implemented based on the microservices architecture (using Spring Cloud), having in mind the scalability and distributed characteristics of the architecture. The listed services will be offered and managed by the symbIoTe consortium. The ones which are of particular interest to platforms are the Administration Service, Registry and Resource Monitor, since they are accepting requests from symbIoTe-enabled platforms.

[2] https://www.w3.org/RDF/

The following CLD components will facilitate the integration of a new platform with symbIoTe. All interfaces mentioned here are part of the Interworking Interface, which exposes platform devices as IoT services. The symbIoTe consortium will provide interface definitions for all components (REST and messaging system based on Rabbit AMQP) together with a supporting Java library implementing non-platform-specific parts of those components.

1. *Registration Handler* – handles platform-side registration of devices with the Registry by using the symbIoTe Information Model. An interface description will be provided along symbIoTe implementation in Java, which needs configuration according to platform specifics.
2. *Resource Access Proxy* – enables secure access to the IoT resources offered by the IoT platform and registered within the Registry. A plug-in template and interface description will be provided which needs to be implemented by the platform owner so as to forward requests to platform-specific actions and return results (resource data). Implementation of the non-platform specific parts will be provided by the symbIoTe consortium.
3. *Platform Authentication and Authorization Manager* – offers authentication and authorization mechanisms on the platform side. Authentication uses standardized Public Key Infrastructure approach whereas authorization is based on the Attribute-Based Access Control mechanism. The symbIoTe consortium provides a Java implementation for this component.
4. *Monitoring component* – monitors the status of registered devices/services and reports the status periodically to the Resource Monitor within the Core Services.

The aforementioned components will need to be integrated by the applicants and platform-specific parts will need to be implemented to connect symbIoTe Core Services with real platform resources. The symbIoTe team will provide interface definitions and component implementations in Java that require further extensions to be integrated with platform-specific components and security schemes. We acknowledge the fact that all extensions written in other languages require more resources.

To make a platform Level 1 compliant, applicants need to do the following:

- Analyze the symbIoTe solution for Level 1 compliance
- Expose the description IoT devices accessible through symbIoTe either using the symbIoTe BIM or a custom PIM (in the form of an ontology).
- Integrate the Interworking Interface required for Level 1 compliance within the CLD domain with their platforms
- Provide feedback and comments which can improve and simplify the process of creating L1 platforms

What will symbIoTe offer to applicants to make their platforms L1 compliant?

- A documented design and prototype implementation of the required symbIoTe Core Services
- A documented Interworking Interface and corresponding components in Java for CLD

- An example procedure for mapping of an existing information model to the symbIoTe information model
- A documented example explaining the process for creating a L1 platform based on the example of the open-source OpenIoT platform

Related Use Cases, Trials and IoT platforms. L1 supports the following symbIoTe use cases and trials:

i) *Smart Mobility and Ecological Urban Routing*, planned in the cities of Zagreb, Vienna and Porto, with 3 IoT platforms involved up to now, offering air quality, noise level and mobility/traffic/parking related data,

ii) *Smart Residence*, planned in Pisa, Vienna and Barcelona, with 3 IoT platforms involved up to now, offering data related to smart assisted living, indoor air quality as well as domotics, and

iii) *Smart Stadium*, planned in Barcelona, with 2 IoT platforms involved up to now, offering beacons platform for indoor location and promowalls.

Additionally, from the 1st Open Call, the selected Third Parties contribute to the following use cases: Smart Building, Smart Marinas and Smart Cities.

symbIoTe-OC2-L2

Topic Summary

Purpose: To make 3rd party IoT platforms L2-compliant, so that they join platform federations. Applicants should explicitly mention with which symbIoTe platforms they would like to federate with (see list of available platforms at the end of this topic) and what could be the business value out of this federation. A desired feature would be that apart from making their IoT platform L2-compliant, they also extend their native apps to display/access/control federated resources.

Commitment: Applicants should make their platforms and their resources available during the symbIoTe demos and trials (mid/end 2018).

Type of Applicants: IoT Platform owners / operators

Funding: up to €50,000 per Extension, approx. four (4) Extensions to be funded

symbIoTe Level 2 (L2) compliance enables organizational interoperability focusing on inter-platform communication and collaboration. In contrast to L1 compliance, L2 compliance targets the efficient resource sharing and access between compliant IoT platforms by forming federations. Thus, within a federation, platforms can securely interoperate, collaborate and share resources, which should be described based on the Information Model defined in L1 allowing the introduction of smart semantic approaches in order to expose IoT devices, and according to the accepted Service Level Agreement (SLA) defined on federation level.

By targeting interoperability between platforms on CLD, connected native applications are able to leverage the benefits and increase their business value by consuming and processing additional resources made available by other platforms and shared within the federation.

Based on the implemented bartering functionalities provided by the symbIoTe reference implementation, additional resource and data distribution channels on platform side are feasible. Through bartering, a level of fairness between platforms can be achieved, guaranteeing that the amount of shared resources between platforms is balanced. Moreover, the introduction of trust calculation at different levels (resources and platform) ensures the accurate reliability and reputation information for the members of the federation, influencing on the Bartering functionalities.

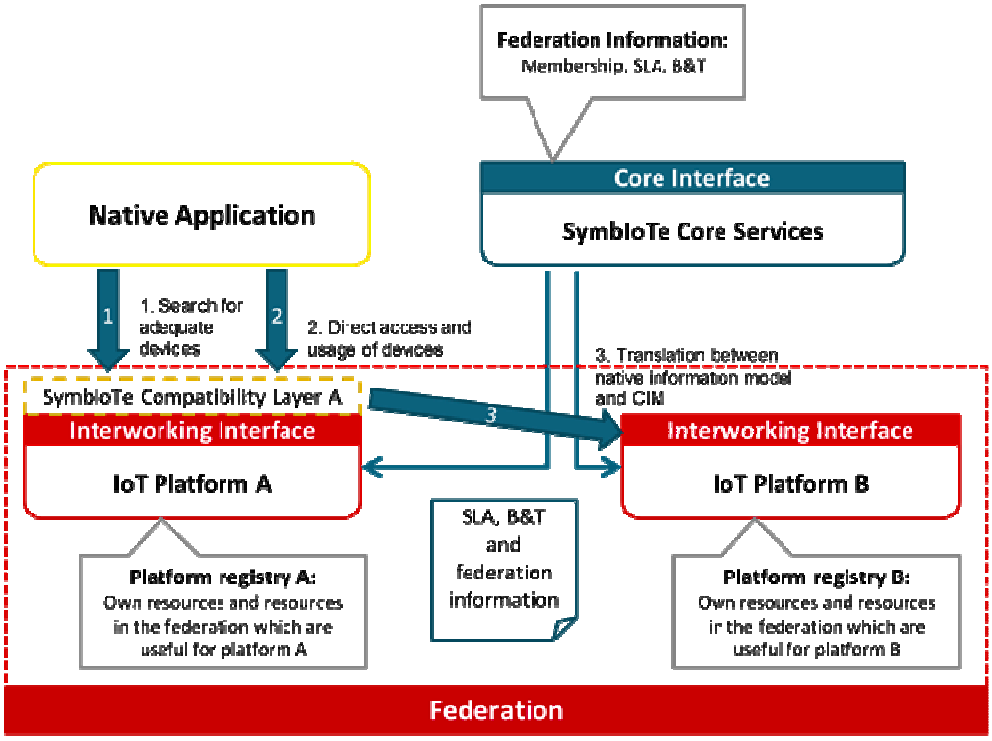

Figure 8: Level 2 compliance

Figure 8 shows an example of resource access for a native application to its home platform inside a federation. In it, the platforms have their own registry of shared resources, which contains their own devices and some devices of other platforms in the federation in which they might be interested. In this scenario, the symbIoTe core holds information about SLA agreements signed between the different platforms of the federation, some Bartering and Trading information as well as membership information and status. This information is propagated to the different participants in the federation as well.

When a native application of platform A searches for resources, it will receive resources from both, platform A and B since both are in the same federation and in this case, Platform A is interested in some resources of Platform B, which are valuable to the aforementioned application. As such, those resources are in its local registry so, when the application executes the search request, it will receive a list of available devices from both platforms. To make life easier to application developers, maintain to a minimum the work needed to adapt a native application and be able to work with symbIoTe on a compliance Level 2, the IoT Platform Owner might provide a compatibility layer that will translate from the Common Information Model, spoken by the Interworking API, to the platform specific information model before returning the data to the application. This layer can also talk to any other platform in the federation easily as the transformation of data will be the same for platform A and B since both of them provide it in the CIM. That way, the application will need just minimal modifications to adapt to the symbIoTe security paradigm. It even might remain unchanged if the compatibility layer also takes care of the security paradigm transformation.

In this section, we highlight the components that an IoT platform needs to integrate with existing components in CLD in order to become L2 compliant, as well as the Core Service components mainly required for managing federations and reaching an agreement for a

platform to join a federation. Since authenticated and authorized access to offered services is vital for an IoT ecosystem, we also include security-related components as outlined in Figure 9.

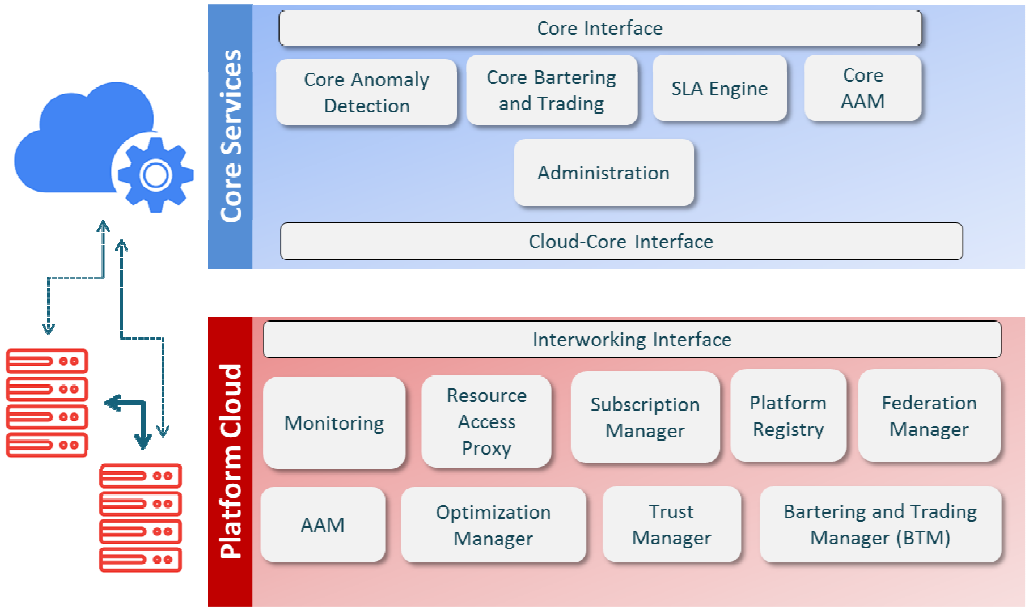

Figure 9: Core and Cloud Components applicable for Level 2 compliance

The Core Services are deployed on a central server and managed by the symbIoTe project team. Therefore, the platforms just have to integrate and interact with the well documented interfaces. On the other side, to make an IoT platform fully L2 compliant, the listed CLD components have to be installed, deployed and run on each platform interacting with the native platform functionalities.

symbIoTe Information Model. The Information models for the description of the exposed resources are the same that have been defined at L1 and the platforms should decide how to semantically describe their devices in the federation based on the two approaches introduced at L1. Therefore, the Information Model defined at L1 is valid for L2, however the main difference is that now the information is distributed among all the platforms and it is not centralized such as in L1.

symbIoTe components for L2 compliance. The following Core Services components interact to provide the L2 functionality:

1. *Core Bartering* – comprises all bartering functionalities that need to be centralized and coordinated by the symbIoTe Core Services. Especially the bartering features between federated platforms will be highlighted. These include validating issued vouchers and reporting to the Trust Manager when a platform within a federation is not cooperating.
2. *Core Authentication and Authorization Manager* – may also participate in federations. Acting as the root certification authority it also has the power to revoke misbehaving platforms certificates and invalidate all credentials originating from them.
3. *Administration* – provides a web-based GUI for the management of federations. The provided functionality includes: creation and management of federations and

their properties, management of security attributes used to access federated resources.

4. *SLA Engine* – manages the whole lifecycle of service level agreements (SLAs) for the federation (definition, negotiation, monitoring, etc).

The following CLD components will facilitate a new IoT platform becoming member of a federation. All interfaces mentioned here are part of the Interworking Interface, which exposes platform devices as IoT services. The symbIoTe consortium will provide interface definitions for all components (REST and messaging system based on Rabbit AMQP) together with a supporting Java library implementing non-platform-specific parts of those components.

1. *Platform Registry* – maintains the actual state of shared (own and foreign) resources within the federation. It also enables a local search of relevant resources, which are offered by other platforms, in contrast to L1.
2. *Subscription Manager* – in contrast to L1, the subscription and notification of any changed resource status or properties are propagated in a distributed fashion between the federated platforms.
3. *Federation Manager* – is responsible for managing all required federation information needed on platform level, like federation and SLA updates, updates on access policies and trust levels.
4. *Optimization Manager* – supports the suggestion of equivalent resources offered within the federation to ensure optimized resource usage by taking power consumption and availability aspects into account.
5. *Trust Manager* – introduces a multi-layer trust calculation on resource and platform level, which may impact the usage and acceptance of platforms and their offered resources.
6. *Resource Access Proxy* **–** enables secure access to the IoT resources offered to the federation. A plug-in template and interface description will be provided which needs to be implemented by the platform owner so as to forward requests to platform-specific actions and return results (resource data). Implementation of the non-platform specific parts will be provided by the symbIoTe consortium.
7. *Platform Authentication and Authorization Manager* – in addition to L1, supports the efficient and secure issuing of federated credentials used to access resources offered within the federation. Those credentials are then, in the federated platforms, validated against the existing federations definitions. The definitions (members and access policies) are provisioned into this module from the Federation Manager.
8. *Bartering and Trading Manager* – acts as the counterpart of the Core Services component, initializes and manages the voucher creation and assignment for each platform. It also interacts and communicates with the Core Bartering component for voucher consumption and validation.
9. *Monitoring* – collects and monitors the status and load of offered resources within the federation and also supports the transmission of aggregated metadata about the resource health to the SLA engine.

The aforementioned components will need to be integrated by the applicants and platform-specific parts will need to be implemented to connect symbIoTe Core Services

with real platform resources. The symbIoTe team will provide interface definitions and component implementations in Java that require further extensions to be integrated with platform-specific components and security schemes. We acknowledge the fact that all extensions written in other languages require more resources.

To make a platform Level 2 compliant, applicants need to do the following:

- Analyse the existing symbIoTe solution and reference implementation for Level 2 compliance
- Integrate, install, deploy and run the CLD components
- Integrate the Interworking interface required for Level 2 compliance within the CLD domain with the IoT platform
- Integrate the platform with the core services
- Create or join a federation and share their resources in the federation

To be able to improve our system, components and documentation continuously, we request all participants to provide feedback so we can enhance and simplify the process of creating L2 platforms.

What will symbIoTe offer to applicants to make their platforms L2 compliant?

- A documented design and reference implementation of the required symbIoTe Core Services
- The documented Interworking interface and corresponding components in Java for CLD
- An example procedure and documented example for setting up the federation L2 platform
- Core service components up and running for L2, in order to be used in the federation management

Adaptation of the native applications (optional).

Applicants could also adapt their existing native applications in order to use the resources of the federated platforms through the symbIoTe mechanisms.

Therefore, the applicants who want to integrate and evolve their native applications need to do the following:

- Make their platforms L2 compliant.
- Integrate and evolve their native applications to be L2 compliant – symbIoTe offers reference Java client libraries as well as standardized REST APIs.
- Participate in a federation and verify access to resources from other members through the symbIoTe APIs in the symbIoTe extended native applications.

To be able to improve the usability of L2 compliance, we request all participants to provide feedback so we can enhance and simplify the usage of the federation features.

What will symbIoTe offer to applicants to integrate their native applications at L2 compliant?

- A documented design and reference examples to integrate it with L2

- An example procedure and documented example for using the federation L2 platform.

Related Use Cases, Trials and IoT platforms. L2 supports the following symbIoTe use cases and trials:

i) *EduCampus*, planned in the city of Karlsruhe, involving indoor positioning and room reservation systems

ii) *Smart Mobility and Ecological Urban Routing*, planned in the cities of Zagreb and Vienna, with 2 IoT platforms involved up to now, offering air quality and noise level data.

iii) *Smart Residence*, planned in Pisa and Barcelona, with 2 IoT platforms involved up to now, offering data related to indoor air quality as well as domotics.

Applicants need to decide which of the existing IoT platforms they want to federate with so as to enhance or extend the aforementioned use cases.

symbIoTe-OC2-L3/4

Topic Summary

Purpose: To make 3rd party IoT gateways and/or IP native families of devices L3/4-compliant. The proposed Extensions must include at least 10 devices of at least 3 different types. Such devices can be: i) gateway-controlled devices, ii) IP-native smart devices or iii) a mix of the previous.

Commitment: Applicants must join a challenge event (end 2018) to demonstrate their solution and to pass a set of symbIoTe-defined interoperability tests.

Type of Applicants: IoT Gateway and/or IP native device manufacturers, system integrators

Funding: up to €40,000 per Extension, approx. four (4) Extensions to be funded

In L3/4 compliance, interoperability is considered at the gateway/smart device level. Devices and gateways can directly expose their resources to interested consumers (other devices, gateways or applications) in proximity. *Smart Space (SSP)* Domain provides services for discovery and registration of new IoT devices in dynamic local smart spaces, dynamic configuration of devices in accordance with predefined policies in those environments, and well-documented interfaces for devices available in smart spaces. *Smart Device (SDEV)* Domain relates to smart devices and their roaming capabilities. We assume that devices have the capabilities to blend with a surrounding smart space while they are on the move. In other words, smart devices can interact with devices in a visited smart space, which are managed by a visited platform, in accordance with predefined access policies. To achieve these points, the symbiote projects defines a SSP/SDEV infrastructure. An high level overview of a SSP is presented in the following figure.

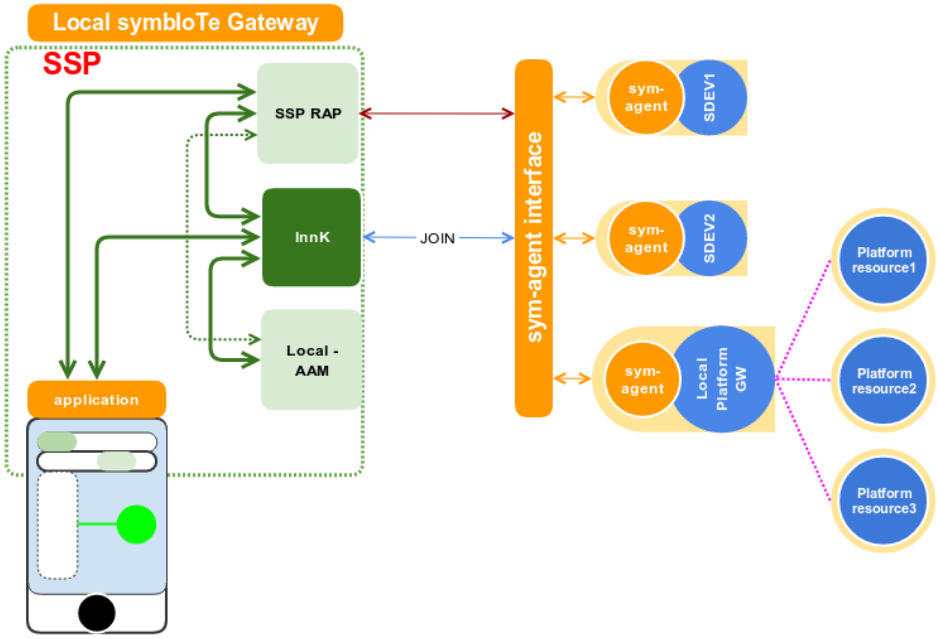

Figure 10 Level 2/3 Compliance

The idea is the following: an SSP is defined from a set of software modules that reside in a local gateway. The symbIoTe-enabled devices (sensors and actuators) could exist as a standalone device, communicating directly with the local gateway, or could be part of a local platform; in both cases, they will have to deploy an agent in order to communicate with the SSP. Moreover, any application running on a smartphone that wants to use the local resources should talk with the SSP using the interfaces provided. In order for a device to become symbIoTe-enabled it needs to develop a layer - the *symbIoTe agent* - that interacts with the SSP based on a well-defined communication protocol. This is all that a device/local platform manufacturer should develop to interact with the symbiote ecosystem at SSP level.

symbIoTe Information Model. In the smart space architecture, gateways and devices have their own symbIoTe agent component that lets them interact with the rest of the symbIoTe Smart Space ecosystem. The representation of the exchanged information is based on CIM introduced by symbIoTe (see L1 for more details).

symbIoTe components for L3/4 compliance. The following components interact to provide the L3/4 functionality:

1. *Innkeeper* – represents the point of entry for gateways and devices to the SmartSpace. This component is also tasked to communicate with the Resource Access Point component to keep the resources URIs consistent.
2. *Resource Access Proxy* – acts as connection between the symbIoTe ecosystem and devices/gateway, receiving requests from the symbIoTe's upper layers and adapting them to be processed by the devices/gateways symbIoTe agent.
3. *Local Authentication and Authorization Manager* – provides local security management system. Its main purpose is to interface directly the symbIoTe's upper layer security infrastructure, and provide al local interface to it. This mechanism allows the smart space to operate in some restricted fashion if local connectivity is not available, and therefore core security mechanisms are also unavailable.
4. *symbIoTe agent* – is the interface between devices and symbIoTe ecosystem. This agent can run on platform gateways or directly in devices; depending on where the agent runs, it could provide different API.

These components provide a set of services to interact with the local SSP: leaving out the registration/update process, an application or a symbIoTe agent can request the list of the available resources and get/send data to them.

To make a platform L3/4 compliant, applicants need to do the following:

- Analyze the existing symbIoTe solution and reference implementation for Level 3/4 compliance.
- Integrate, install, deploy and run the SSP components.
- Develop the symbIoTe agent for their gateways/devices.
- Participate in the challenge event where their gateways/devices will pass some test to validate their L3/4 compliance.

What will symbIoTe offer to applicants to make their platforms L3/4 compliant?

- symbIoTe provides documentation about the implemented interfaces between the agent and the modules inside the SSP, including also the semantic description that should adopt the resources;
- symbIoTe provides the software modules to create a local instance of SSP gateway with the symbIoTe middleware. This gateway can be hosted on common x86-based machine.

symbIoTe-OC2-Apps

Topic Summary

Purpose: To build a smart phone application (Android or iOS) that combines resources offered by L1-compliant platforms and/or available Domain Enablers and demonstrates cross-domain features. Proposed applications should focus in the domains of Smart City, Smart Residence and Smart Stadium/Buildings, as well as combination of the above.

Commitment: Applicants should make their applications available for a demo during the symbIoTe trials (mid/end 2018).

Type of Applicants: Mobile app companies

Funding: up to €20,000 per Extension, approx. three (3) Extensions to be funded

One of the key goals of symbIoTe is to demonstrate the easiness of developing IoT applications over the defined uniform and secure interface of L1-compliant platforms, as well as utilizing the advance intelligence and analytics capabilities of available Domain Enablers. In this topic, we look for Extensions that will build new IoT applications or will extend existing IoT applications so as to utilize the symbIoTe Core Services and be able to access the resources offered by the L1-compliant platforms. For more info on L1 compliance, please refer to the *symbIote-OC2-L1* topic. In addition, certain Domain Enablers that collect, aggregate and process data from various L1-compliant platforms can be utilized.

Available L1-compliant platforms. Below a list of the currently available L1-complaint platforms is provided. The list is expected to grow and include the beneficiaries from the 1st Open Call.

- Open IoT Platform: OpenIoT is an open-source Cloud platform for the Internet of Things. It manages the registration, data acquisition and deployment of different sensors using the Semantic Web technologies and the SSN ontology, thus enabling the semantic unification of diverse data and IoT applications in the Cloud. In symbIoTe, OpenIoT realizes a Mobile Crowdsensing Air Quality Platform offering support for discovering and collecting data in a crowdsensing fashion from wearable sensors through a Cloud-based Publish/Subscribe Middleware. Such a platform can be used for applications related to air quality monitoring and services, for which the users' wearables will provide air quality measurements. The crowd sensing air quality service operates on the data gathered through OpenIoT.
    - Offered data: Mobile sensors that measure air quality information (CO, NO2), temperature, humidity, pressure. Mobile phones provide noise and luminosity data.
- openUWEDAT: The cities of Zagreb and Vienna operate Air Quality Measurement Networks, whose data is typically collected from fixed stations using the platform called "Stationary Air Quality Platform" (openUWEDAT). This way, mobile data collected via OpenIoT platform can be complemented by data provided by openUWEDAT for improving the accuracy by minimizing measurement errors related to data from mobile sensors.

  - Offered data: Mainly urban area air quality data originating from stationary, high-accuracy measuring stations.

- MoBaaS Platform: The Mobility Backend as a Service (MoBaaS) offers a set of services, in the form of APIs, which intend to eliminate the friction created by having services from multiple vendors. This includes the integration of data from many sources, focusing on the mobility aspect of the city. In symbIoTe, MoBaaS is foreseen for the applications including the routing service based on preferences such as points of interest, air quality, distances, and amount of traffic and parking availability.
  - Offered data: Mainly mobility data such as POIs, spatial air quality, distances, traffic conditions, parking availability.
- Symphony Platform: Symphony is an IoT platform used for the integration of home/building control functionalities, devices and heterogeneous subsystems. Symphony can monitor, supervise and control many different building systems, devices, controllers and networks available from third-party suppliers. By intelligently correlating cross-system information, a flexible and highly efficient platform is delivered to the stakeholders. The system is a service-oriented middleware integrating several functional subsystems into a unified IP-based platform. As hardware/software compound, Symphony encompasses media archival and distribution, voice/video communications, home/building automation and management, and energy management.
  - Offered data: Sensors and actuators such as lights, dimmers, RGB lights, curtains.
- KIOLA Platform: KIOLA is an IoT telehealth platform. It is able to store health and person-related data, and register sensor information of personal health devices (e.g. body weight, blood pressure). Additionally, users can be identified using Bluetooth beacons (BLE). An additional plugin serves as a wrapper for commercially available fitness trackers. Using the plugin fitness trackers such as Fitbit or Nokia Health can also be registered with the symbIoTe core.
  - Offered data: Mainly health and person-related data, including activity tracker, body weights, blood pressure, as well as information on Bluetooth beacons and fitness trackers.
- nAssist Platform: The IoT platform is a software platform designed and conceived to allow agile, continuous management of data in the energy efficiency, security and automation fields. It is built following a Service Oriented Architecture paradigm and has been designed to be easily adapted to different areas of application that use, or implicitly need, data collection and data processing from logical or physical devices (sensors and actuators). One application of nAssist is that of monitoring and controlling a number of direct parameters related to indoor air quality, such as CO2 levels, humidity and temperature. In addition, this platform monitors and controls other factors that are important for indoor environmental quality considerations such as light and noise, as they also affect occupants.
  - Offered data: Mainly indoor air quality, humidity and temperature, and luminosity levels data.

- Beacons Platform / Visitor Platform / PromoWall Platform / Retailer Platform: These platforms are part of the smart stadium use case and represent the underlying IoT environments (beacons, visitor information, promo-wall control, and retailer platform). Visitors in the stadium, using the visitor app in their smart phone and the indoor location beacons, find and access nearby services provided by Retailers (shops and moving carts) through the retailer app in their tablets and the indoor location services. Retailers find Promowall devices and send info & promotions to them.
    - Offered data: Apps in each device (mobile, tablet, PromoWall, etc.) act as sensors (location) and actuators for the associated actions (show promotion, place order, etc.). The data interchanged is the beacon map, the devices (visitor, retailer) availability & location, the retailer offerings, the purchasing orders and the info & promotions for the Promowalls.
- Navigo Digitale Platform: The Navigo Digitale IoT Platform is a platform created to manage digital assets pertaining to harbors used for boating and yachting. Its scope embraces both physical entities (objects) and immaterial entities (documents and workflows). It consists of a distributed platform, with instances associated to different ports across Europe and running part in the cloud, and part on premise. The ultimate purpose of ND is to provide services to the harbor's activities (B2B) and to its end-users (B2C). An interface to Navigo's PortNet application provides access to the harbors control application.
    - Data: Mainly data on mooring process and parameters (location, distance, time, speed, water tanks levels, arrival time, latest routes, fuel situation, emissions, environmental sensor data, etc.).

Available Domain Enablers. Domain Enablers are virtual IoT platforms which combine data from different IoT platforms and provide this data either as is (raw data), or in an aggregated or processed manner. They can be accessed using L1 interfaces. Below, we describe the currently available Domain Enablers by symbIoTe.

- Green Route Enabler obtains air quality data from the platforms and interpolates the data with the street segments of the map being used in order to obtain the air quality of a given street. These data are provided to the routing services (either the ones residing within a platform or external services), which, combining with other data, such as traffic or parking, are able to compute green routes. Additionally, the data provided by the platforms can also be used to obtain POIs of interest to the users.

To make a symbIoTe-powered smart phone application, applicants need to do the following:

- Analyse the symbIoTe solution for Level 1 compliance
- Build or extend a smart phone application that uses resources offered by L1-compliant platforms and/or available Domain Enablers
- Provide feedback and comments which can improve and simplify the process of creating L1-compliant applications

What will symbIoTe offer to applicants to built symbIoTe-powered applications?

- A documented design and API of the symbIoTe Core Services and Domain Enablers
- A deployed instance of Core Services, with the metadata of the available L1-compliant platforms
- A deployed instance of Domain enablers, with the processed domain-specific data
- A documentation of the used information model
- Description of available IoT platforms that are symbIoTe enabled, and description of planned trials

symbIoTe-OC2-Trials

Topic Summary

Purpose: To involve end users and citizens that will use our applications and hardware (portable sensors, etc) to support our "Smart Mobility and Ecological Urban Routing" trial.

Commitment: Support the trials planned in Zagreb, Vienna and Porto (mid/end 2018).

Type of Applicants: NGOs, municipalities, organizations, companies with end users in Zagreb, Vienna and Porto.

Funding: up to €15,000 per Extension, approx. three (3) Extensions to be funded

The symbIoTe project has planned a number of trials in various locations. The purpose of the Extension is to strengthen the planned Smart City related trials in Vienna, Zagreb and Porto. More specifically, the **Smart Mobility and Ecological Urban Routing** use case involves the collection of air quality information from not only stationary sensing stations but also mobile wearable sensors, which is processed to obtain the air quality index for the various streets of the city. This information can later be used, through applications, to calculate the ecological routes to the users' destinations or to support the search for points of interest (POI) within the city.

As such, the expectation from the trials of this use case is for an end-user community to use the wearables during their daily commute, providing air quality data of the locations they visit. Additionally, users are also expected to use the applications developed within the Use Case to obtain the ecological routes (which are influenced by the data they provided) and to search for POIs. These trials are directed at NGOs (green/cyclists organizations) and municipalities of Vienna, Zagreb and Porto.

The Extensions must include the number of participants, the location of the trials as well as the necessary budget for trial implementation (which has to take into account the cost of wearables for air quality monitoring). Note that a smaller number of participants is need to perform air quality monitoring using wearables, while a larger community of users can use applications developed by the symbIoTe consortium for ecological routing and POI search. Thus, the number of trial participants is by no means limited by the number of available wearables. In addition, the participating organization needs to consider costs for the dissemination and promotion of their specific activity, in order to attract and incentivize an increased user community to participate in the organized trials.

The consortium will support the Extension to organize adequate air quality monitoring campaigns and lend whenever and if possible the wearables available among consortium partners for particular trials in order to optimize the usage of available wearables. The planning of the trials should be performed within the first month of the Extension.

Technical details for the wearables: The wearables for air quality monitoring should be reasonably small and portable, with autonomy of operation for at least one day. They should include a Bluetooth module for communication with a smartphone and use an open protocol for transmitting sensor readings directly to the smartphone. They should include at least one of the following sensors for gases: carbon monoxide, nitrogen dioxide,

particulate matter or volatile organic compounds (VOCs). The price range of such devices is €200 to €450 per unit.