International Civil Aviation Organization

The Fourth Meeting of System Wide Information Management Task Force (SWIM TF/4)

Web-conference, 3 – 6 November 2020

Agenda Item 5: Updates on the assigned tasks by task leads/contributors including progress Report and issues

d) Governance
- Registry - implementation guidance for Interoperable Registry Model
  - Task 5 (Contains previous Task 1-4, Task 1-5, Task 2-1-2, Task 2-1-4)

SWIM DISCOVERY SERVICE (SDS): INTRODUCTION

(Presented by FAA SWIM, USA, and Korea Airports Corporation (KAC), Republic of Korea)

SUMMARY

This Working Paper describes the enabling technologies and practices for instituting a *SWIM Discovery Service*. The described approach will support federated service discovery among independently developed and autonomously managed APAC SWIM initiatives.

The paper also introduces the *SWIM Discovery Service (SDS) Implementation Specification* and discusses the guidelines and general technical principles described in this document.

1. INTRODUCTION

1.1 Service discovery has been regarded as a critical aspect in the development of service-centric systems. If a potential consumer is unable to locate (i.e., discover) a service, the service -- and SWIM in general -- fails to achieve its objectives. [Disc-IMP]

1.2 Every service-centric environment such as SWIM supports the ability of services to be discovered through some form of retrieval mechanism (e.g., a service registry). The growing number of SWIM implementations has introduced new challenges to SWIM stakeholders to discover services across geographical and organizational boundaries.

1.3 To address these concerns, APAC SWIM TF developers and architects examined various scenarios for integrating service information stored across multiple and diversified SWIM registries [WP16-ROK]. The analysis demonstrated that the integration of SWIM service registries is difficult to attain due to the fact that all SWIM registries are developed and managed under different organizational and technological constraints.

1.4 Subsequent analysis deduced that establishing a virtual network of interoperable services interacting on a well-defined composition pattern [WP16-ROK] [SIRF-CONOPS] is the most feasible approach for cross-domain service discovery.

1.5 To this end, the USA FAA SWIM program and Korean Airports Corporation (KAC) formed a collaborative effort to develop discovery services capable of enabling federated service discovery. The FAA SWIM team developed the SWIM Discovery Service (SDS)

Implementation Specification version 1.0.0 [SDS-Spec] that formalized an interface, interaction patterns, schemas, and message exchange formats.

1.6 Both teams are currently developing their own discovery services in accordance with SDS 1.0.0 and have produced demonstrable results. Bilateral testing is scheduled to be conducted in the near future.

1.7 This paper examines the SWIM Discovery Service concept and discusses the notion of federated service discovery and its objectives and advantages. It also introduces the SWIM Discovery Service (SDS) Implementation Specification and discusses the guidelines and general technical principles described in this document.

2. DISCUSSION

2.1 Discovery as a Service

2.1.1. Service-Oriented Architecture, or SOA, is an *architectural style* that supports a way of thinking in terms of services and service-based development and services outcomes. According to SWIM architectural principles, services may provide information (those are known as *information services*) or support interaction among services (referred to as *core services*). A discovery service, as discussed in this paper, is a core service whose purpose is enabling discovery of available information services.

2.1.2. Within the context of a single SWIM service inventory, the most common discovery tool is a service registry. A discovery service is not as familiar, and it is important to explain the distinction between a service registry and a discovery service.

2.1.3. A *service registry* is an authoritative, centrally controlled store of information. It is commonly realized as a catalog that allows a user to store and manage service-relevant information [SWIM-CV]. The information, generally known as a *service description*, is almost always entered by a human user. The human user may retrieve this information through some query or browsing interface.

2.1.4. A *discovery service* is first and foremost a *service* as defined and understood in SOA [SWIM-CV]. It is loosely-coupled, discoverable, reusable, and composable. It exchanges data with other components via messages using a formal language (XML, JSON). It has an interface that can be described in a formal language.

2.1.5. However, it should be noted that a discovery service is not expected to replace a service registry; both components may coexist and even complement each other in a single SWIM environment. Figure 1 demonstrates how two SWIM instances (in this case, FAA SWIM and ROK KAC) both deploy registries and discovery services. Both service registries support the entry and management of service metadata by human users, after which the metadata can be further disseminated by a discovery service. (It could be said that a discovery service uses a registry as a data source.) The discovery service in turn can augment the information stored in its associated registry with the information made available by other discovery services (*peers*). A service registry may also serve as a client for accessing a discovery service.

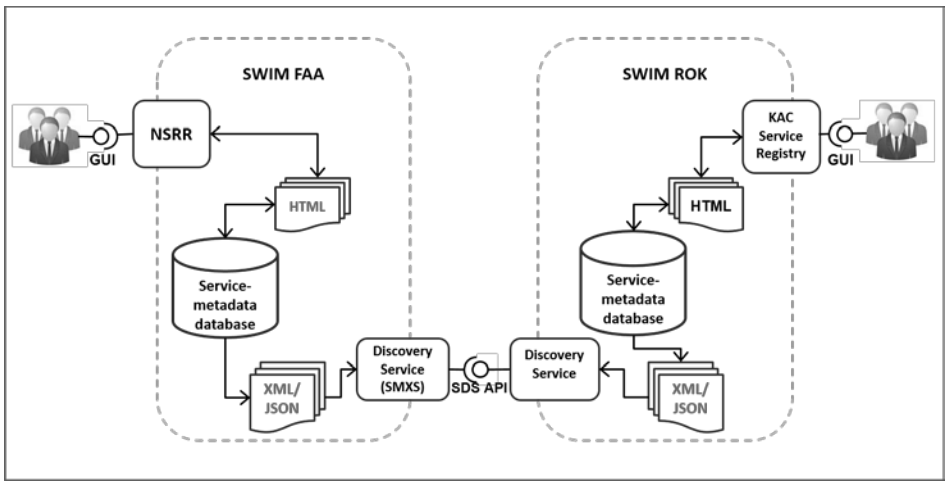

Figure 1 Two SWIM environments with service registries and interacting discovery services

2.2 Architectural Vision

2.2.1 The following factors motivated the development of the architectural vision of SDS:

2.2.1.1 To enable service discovery among independently developed and autonomously managed SWIM domains;

2.2.1.2 To allow a user to obtain service metadata from multiple sources simultaneously in a single, consolidated, semantically consistent result;

2.2.1.3 To avoid reliance on a single centralized discovery mechanism by allowing a discovery service to describe itself and other discovery services.

2.2.2 To realize these objectives, the following design principles were asserted:

2.2.2.1 To conform to SOA architectural principles;

2.2.2.2 To be based on World Wide Web architecture and standards;

2.2.2.3 To use a formal standard language for information exchange;

2.2.2.4 To be self-describing and self-advertising;

2.2.2.5 To be composable but not coupled.

2.2.3 From a technological perspective, the SDS architecture draws upon multiple works in information technology. However, two architectural models should be especially noted for understanding the SDS's architectural vision:

2.2.3.1 Representational State Transfer (REST) [REST]. REST defines an approach for implementing services (RESTful Web services) that provide and access a representation of information resources by using a uniform and predefined set of stateless operations. In the SDS context, the interface defined in REST is used to define the behavior of an individual discovery service.

2.2.3.2 Peer-to-Peer (P2P) Discovery [WS-Arch] [SIRF-CONOPS]. A P2P distributed application architecture defines a network where every node ("peer") is an equally privileged, equipotent participant with the same capabilities and responsibilities. Following a P2P Discovery architectural vision allows establishing a pattern for interactions among discovery services.

2.2.4 Together these two architectural paradigms address the design objectives defined in 2.2.2. Because every discovery service is a Web service, it conforms to SOA

architectural principles (2.2.2.1), it is based on WWW architecture and standards (2.2.2.2), and it uses a formal standard language for information exchange (2.2.2.3). Discovery services do not need a centralized discovery mechanism; they are self-advertising and also advertise each other (2.2.2.4), and removing or adding a service from or to a network does not preclude other services from continuing to function (2.2.2.5).

2.2.5 According to SDS architecture, a discovery service may request information from another discovery service about its identity and the operations it offers, including the operations' access constraints (self-advertising). Upon receiving this information, the service may request a filtered or unfiltered list of references to all service descriptions found by the other discovery service. Subsequently, the service may request one or more service descriptions using the previously obtained list of references. The service may also request information about other discovery services known to the invoked service (advertising). Figure 2 illustrates the described behavior pattern.

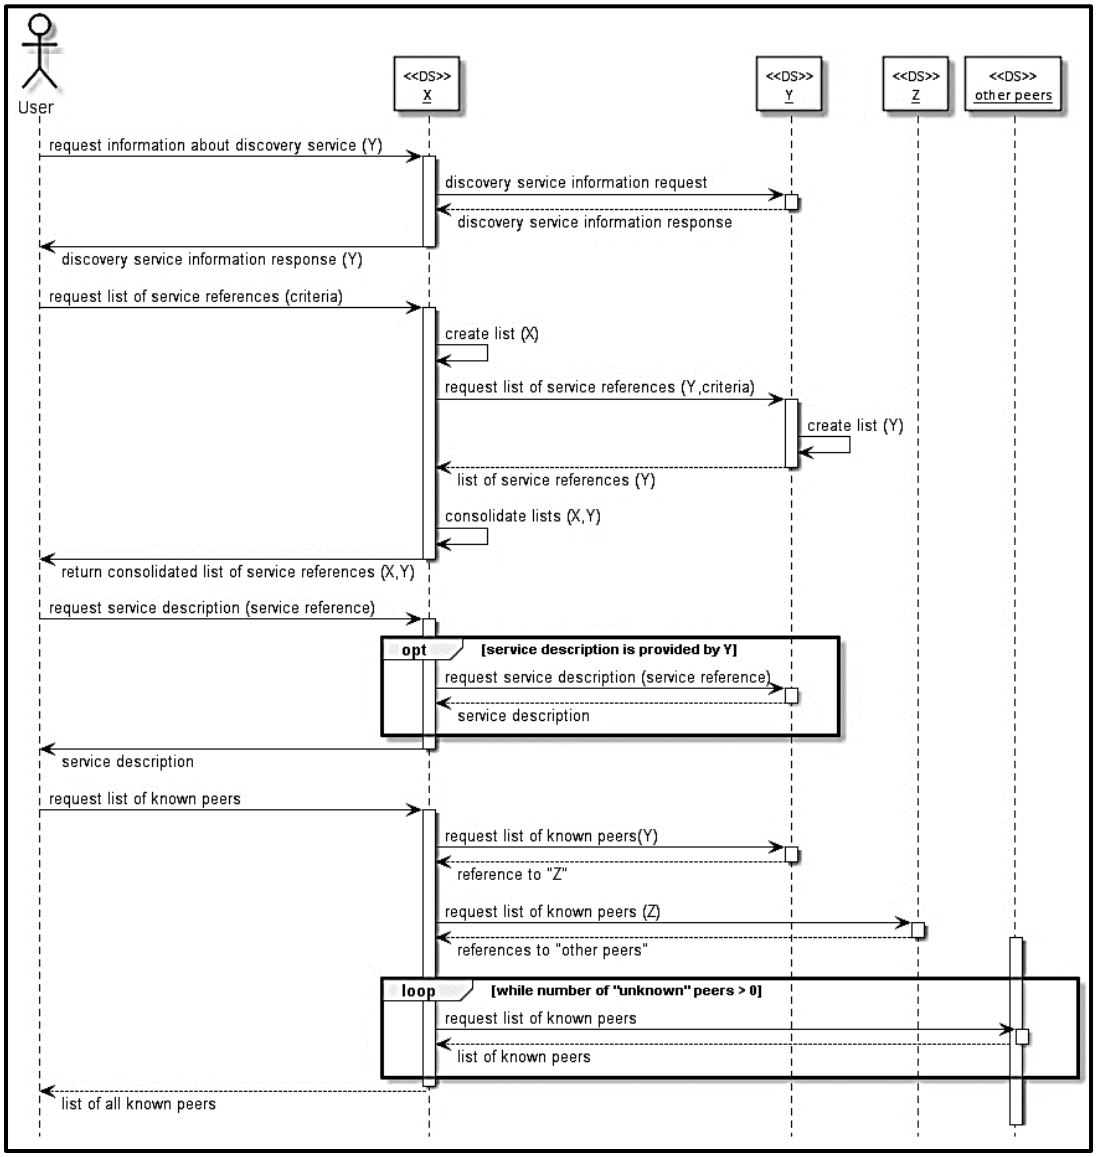

Figure 2 Discovery services interactions pattern

2.3 SDS Implementation Specification

2.3.1 To formalize the architectural vision of SDS, the FAA SWIM team developed "SWIM Discovery Service (SDS) Implementation Specification Version 1.0.0" [SDS-Spec] in July of 2020.

2.3.2 The purpose of the specification is to establish guidelines and general technical principles for the development of SDS. It presents the enabling technologies and practices that support federated service discovery among independently developed and autonomously managed SDS implementations. [SDS Spec]

2.3.3 The specification consists of five sections: Behavior Model, Information Model, Resource Model, Interface Requirements, and Security Requirements. Each section is described as follows:

2.3.3.1 *Behavior Model* presents a collection of use cases, which together describe the discovery service behavior, i.e., how the service interacts with a user and other services.

2.3.3.2 *Information Model* defines common structures for information exchanged among discovery services. All information elements are also rendered in a formal exchange language (UML, JSON). Information Model defines two kinds of information exchanged in the SDS context:

2.3.3.2.1 Information that supports discovery services' interactions, and which may include an identification of a discovery service, functionalities provided by the service, access policies, and references to other discovery services.

2.3.3.2.2 Information provided by a discovery service to support service discovery, such as lists of services or detailed descriptions of these services. This kind of information is derived from the international Service Description Conceptual Model (SDCM) version 2.0 [SDCM-2.0].

2.3.3.3 *Resource Model* defines a collection of interlinked resources. It closely follows the principles defined in REST architecture as well as structural content expressed in OpenAPI Specification, version 3.0 [OPEN-API]. (SDS OpenAPI document is included in the specification).

2.3.3.4 *Interface Requirements* prescribes the operations and associated messages supported by a discovery service, with each operation representing a simple interaction between the service and a user or another service. All operations defined in the context of the SDS API comply with requirements laid out by RFC 7231, Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content, IETF, 2014 [RFC-7231].

2.3.3.5 *Security Requirements* prescribes the security measures to be implemented by a discovery service.

2.4 Conclusion

2.3.4 The ability of SWIM stakeholders to find (discover) services across geographical and organizational boundaries is a precursor for achieving global information exchange.

2.3.5 SWIM Discovery Service enables global discoverability, that is, SWIM stakeholders' ability to find services, regardless of which SWIM registry originates and maintains the service metadata.

2.3.6 It promotes a technological means for presenting all aspects of a service's metadata in a manner suitable for both human-readable and machine-processable representations.

2.3.7 It should serve as a case study for further advancing the notion of service composition and the development of commonly-shared artifacts and practices in the international aviation community.

3. ACTIONS BY THE MEETING

3.1. The meeting is invited to:

3.1.1. note the contents of this working paper;

3.1.2. provide feedback on the proposed approach;

3.1.3. advise on the way forward, and in particular set priorities for future developments in the area of service discovery in the context of APAC SWIM.

4. REFERENCES

[Disc-IMP] Discovery Services in the Context of Global SWIM, ICAO IMP WG-G, Mark Kaplun (FAA), August 28, 2017

[WP16-ROK] SWIM TF/2 – WP/16, APAC SWIM Registry Approach, APAC SWIM TF3, Task 1-5, 07-10/05/19
https://www.icao.int/APAC/Meetings/2019SWIMTF3/WP16_ROK%20AI3%20-%20SWIM%20Registry%20Approach%20Revised.pdf

[REST] Architectural Styles and the Design of Network-based Software Architectures, Dissertation; Roy Thomas Fielding, 2000
https://www.ics.uci.edu/~fielding/pubs/dissertation/top.htm

[SDS-Spec] SWIM Discovery Service (SDS) Implementation Specification, Version 1.0.0; USA FAA SWIM; July 2020
https://discovery.swim.aero/sds/1.0.0/SDS%20Specification%20v.1.0.0.pdf

[SWIM-CV] SWIM Controlled Vocabulary V. 1.0.0; FAA SWIM; 2019-03-25
https://semantics.aero/pages/swim-vocabulary.html

[OPEN-API] OpenAPI Specification, Version 3.0.3; Swagger;
https://swagger.io/specification/

[SIRF-CONOPS] Concept of Operations for the SWIM Inter-Registry Framework (SIRF), Version 1.0.0; U.S. FAA SWIM, October 11, 2018
https://www.faa.gov/air_traffic/technology/swim/governance/international_collaboration/media/conops-sirf-1.0.0.pdf

[SDCM-2.0] Service Description Conceptual Model (SDCM) 2.0, FAA/SESAR, June 3, 2016
http://swim.aero/sdcm/2.0.0/sdcm-2.0.0.html

[WS-Arch] Web Services Architecture; W3C Working Group Note 11 February 2004
https://www.w3.org/TR/ws-arch/

– – – – – – – – – – – – – – –