

Internal distribution code:

- (A) [-] Publication in OJ
- (B) [-] To Chairmen and Members
- (C) [-] To Chairmen
- (D) [X] No distribution

**Datasheet for the decision
of 3 December 2020**

Case Number: T 1873/14 - 3.5.01

Application Number: 11010010.4

Publication Number: 2469457

IPC: G06Q10/06

Language of the proceedings: EN

Title of invention:

System and method for improved service oriented architecture

Applicant:

SAP SE

Headword:

System and method for improved service oriented architecture/
SAP SE

Relevant legal provisions:

EPC Art. 56

Keyword:

Inventive step - caching service-output values (no - common
general knowledge)



Beschwerdekammern
Boards of Appeal
Chambres de recours

Boards of Appeal of the
European Patent Office
Richard-Reitzner-Allee 8
85540 Haar
GERMANY
Tel. +49 (0)89 2399-0
Fax +49 (0)89 2399-4465

Case Number: T 1873/14 - 3.5.01

D E C I S I O N
of Technical Board of Appeal 3.5.01
of 3 December 2020

Appellant: SAP SE
(Applicant) Dietmar-Hopp-Allee 16
69190 Walldorf (DE)

Representative: Müller-Boré & Partner
Patentanwälte PartG mbB
Friedenheimer Brücke 21
80639 München (DE)

Decision under appeal: **Decision of the Examining Division of the
European Patent Office posted on 17 April 2014
refusing European patent application No.
11010010.4 pursuant to Article 97(2) EPC.**

Composition of the Board:

Chairman W. Chandler
Members: N. Glaser
C. Schmidt

Summary of Facts and Submissions

- I. This appeal is against the decision of the examining division to refuse the European patent application No. 11010010.4 pursuant to Article 97(2) EPC on the ground of lack of inventive step (Article 56 EPC).

The invention was considered to be an obvious automation of an administrative scheme using a general purpose networked computer system and the techniques, all known, for example, from M. Wieland et al.:

"Towards Reference Passing in Web Service and Workflow-Based Applications," IEEE International Enterprise Distributed Object Computing Conference, Auckland, 2009, pages 109-118 (D1).

- II. In the statement setting out the grounds of appeal, the appellant requested that the decision under appeal be set aside and that a patent be granted on the basis of the refused main or auxiliary request re-filed therewith.
- III. In the communication accompanying the summons to oral proceedings, the Board set out its preliminary opinion that the invention seemed not to involve an inventive step (Article 56 EPC).
- IV. In a reply, the appellant gave further arguments in favour of inventive step.
- V. Oral proceedings took place on 3 December 2020. At the end of the oral proceedings, the Chairman announced the decision.

VI. Independent claim 1 of the main request reads as follows:

"1. A computer-implemented method of executing a service-oriented task in a distributed network environment, comprising:

- providing a task-coordinating unit (202) comprising an enterprise service bus;*
 - saving (102), by the task-coordinating unit (202), the task as a sequence of services including a first service and a second service, wherein the saved task includes a first service-input location (202A) that indicates where first service-input values are stored for the first service, wherein the first service-input location (202A) is a file directory under the control of the task-coordinating unit (202);*
 - sending (104) a first task-pending notification from the task-coordinating unit (202) to a first service provider (204) corresponding to the first service, wherein the first task-pending notification includes the first service-input location (202A);*
 - receiving (106) at the task-coordinating unit (202) a first task-results notification from the first service provider (204), wherein the first task-results notification includes a first service-output location (204A) that indicates where the first service provider has stored corresponding first service-output values, wherein the first service-output location (204A) is a file directory under the control of the first service provider (204);*
- sending (108) a second task-pending notification from the task-coordinating unit (202) to a second service provider (206) corresponding to the second service, wherein the second task-pending notification includes the first service-output location (204A) identified as a second service-input location that indicates where*

second service-input values are stored for the second service;

- receiving (110), at the task-coordinating unit (202) a second task-results notification from the second service provider (206), wherein the second task-results notification includes a second service-output location (206A) that indicates where the second service provider (206) has stored corresponding second service-output values, wherein the second service-output location (206A) is a file directory under the control of the second service provider (206);
- providing the first service at the first service provider (204) by accessing the first service-input values at the first service-input location (202A), generating the first service-output values from the first service-input values, and storing the first service-output values at the first service-output location (204A);
- providing the second service at the second service provider (206) by accessing the second service-input values at the second service-input location, generating the second service-output values from the second service-input values and storing the second service-output values at the second service-output location (206A), wherein providing at least one of the services includes:
 - maintaining a cache of service-output values for the at least one service; and
 - generating at least some service-output values for the at least one service by accessing cache values stored prior to receiving a corresponding task-pending notification; and
- sending, by the task-coordinating unit (202) the second service-output location (206A) to a user."

VII. Auxiliary request 1 adds to claim 1 of the main request that "*the file directories (202A, 204A, 206A) of the task-coordinating unit (202) and the first (204) and second service provider (206) are accessible through a Uniform Resource Locator, URL*" at the end of the receiving step (110).

Reasons for the Decision

1. The invention

1.1 The invention relates to service-oriented computing in the form of Service Oriented Architectures (SOA), where components provide services, such as hotel reservations or map data, through a communication protocol over a network. SOA is characterised by platform-independence, loose coupling, dynamic search and binding, and location-independence. Service providers are decoupled from service consumers, but operate as passive building blocks that simply perform a task when queried and reply with the result, paragraph [0002] of the application.

1.2 The application points out, see paragraph [0002], that typical SOA implementations are based on synchronous web services which are orchestrated by an Enterprise Service Bus (ESB). The ESB may become a bottleneck, because all data requests and responses have to pass through it.

1.3 The invention aims to provide a more efficient approach in which the execution of service-oriented tasks is improved by coordinating service providers that access service-input values from other service providers and by generating service-output values that are accessible by other service providers (end of [0002] and [0003]).

1.4 In the example of Figure 2, a task comprises a sequence of service A (e.g. providing hotel reservation information) and service B (e.g. providing map information), and it includes a first service-input location 202A which indicates where input values (area where a hotel is desired) for service A are stored.

This service-input location 202A is a file directory under control of a task-coordinating unit 202. The task coordinating unit notifies the first service provider, providing service A, of the location.

The first service provider then accesses the input values, generates output values (hotel addresses, rates, on-line reservation access) and stores them as first-service output values at a first service-output location 204A, which is a file directory under the control of the first service provider.

The task coordinating unit, having received a notification of the result and of location 204A, notifies the second service provider, providing service B (map information), of location 204A where the input values for service B are stored.

The second service provider then accesses the input values (hotel addresses), generates output values (map with icons representing the hotels) and stores them as second-service output values at a second service-output location 206A, which is a file directory under control of the second service provider.

The task coordinating unit, having received a notification of the result and of location 206A, sends location 206A to the user.

2. Article 56 EPC - main request
- 2.1 The Board agrees with the appellant that the terms "service", "task", "service provider", "task coordination unit", "enterprise service bus" should be given their well-defined meaning in the technical field of service oriented architectures (SOA) and Enterprise Service Buses (ESB). These terms have technical character and should not be interpreted in a purely abstract non-technical way.
- 2.2 In the Board's view, the document D1 referred to by the examining division represents the closest prior art as it discloses a computer-implemented method of executing a service-oriented task in a distributed network environment having the same idea as the invention, which is how to transfer data more efficiently in service-oriented architectures (SOA), by passing pointers, i.e. service input/output data locations, so-called End Point References (EPR) instead of the data itself (cf. point 1.3, above and D1, page 117, right column, third paragraph).
- 2.3 D1 discloses "a task-coordinating unit" (the devices labelled "workflow" in Figures 3 and 5 process workflows which orchestrate the (web) services, corresponding to the first and second service providers). The appellant did not dispute that an enterprise service bus was implicit in D1.
- 2.4 Figures 2(b) and 5 of D1 show three services, namely a "Filtering" service, a "Mapping" service and a "Rendering" service which make a "sequence of services" as claimed (see paragraph bridging pages 114 and 115). The signals that pass between the workflow engine and the services can be considered to be the task-pending

and task-results notifications. They include pointers/locations to the corresponding input/output data (thin lines) rather than the data itself (thick lines). The processing thus corresponds to the saving, sending, receiving and providing service features of claim 1.

Each service has a Reference Resolution System (RRS) that writes values to a storage location, creates a reference/pointer for that location and returns the reference for later retrieval of the data (see also sections 1.2 and 3). Figure 3 shows the architecture of the RRS with five exemplary EPR adaptors for accessing data, among others, "FTP access" and "file access". The "file access" adaptor refers to a file system, page 113, left column, fourth bullet point, which equates to the claimed "file directory". Contrary to the appellant's view, the RRS can also write values (see page 112, right column, second paragraph, last sentence). D1 states that data of the pointers is stored in any kind of system optimised for the kind of data, page 111, right column, last paragraph.

- 2.5 The three services in Figure 5 could be mapped to the claimed first and second service providers in a number of ways. The distinction is important because the first task-pending notification of claim 1 contains a reference to the input data whereas in D1 it is the data itself (thick line to the Filtering service in D1).

If the "Filtering" service is mapped to the first service provider and the "Mapping" service or "Rendering" service to the second service provider, then there would be a difference between claim 1 and D1, because the "Filtering" service receives data and not a reference. However, in the Board's view it would be obvious to the skilled person to replace the (input)

data with a reference. D1 discloses at page 113, left column, last paragraph, that a RRS can be also connected to the workflow engine. This would enable the workflow engine to store the data and to use its RSS to transmit a reference to the "Filtering" service.

Alternatively, the "Mapping" service could be mapped to the first service provider and the "Rendering" service to the second service provider. Both services then receive references and there would be no difference. This mapping is possible, because the application, paragraph [0020] discloses that "first" and "second" are not intended to denote any specific spatial or temporal ordering.

2.6 Claim 1 thus differs by:

- (1) *"maintaining a cache of service-output values for the at least one service"* and
- (2) *"generating at least some service-output values for the at least one service by accessing cache values stored prior to receiving a corresponding task-pending notification"*.

2.7 In the Board's view, these features boil down to the fact that a service provider provides its service either by accessing service-input values indicated in a task pending notification which it received from the central task-coordinating unit, or from service-output values which it generated perviously and which were stored in a cache.

2.8 Therefore, the two features can be considered to have the technical effect of being an intermediate and efficient accessible storage of data, which is the function of a conventional cache memory used in

conventional computer systems, such as those on which the invention is said to be implemented, see paragraph [0021], or disclosed in D1.

- 2.9 Accordingly, claim 1 of the main request does not involve an inventive step (Article 56 EPC) over D1 in combination with common general knowledge, in particular about the "caching of data".
3. The appellant argued that the invention achieved more far reaching effects.
 - 3.1 Firstly, an *outsourcing* of most of the tasks from the task-coordinating unit which led to a decentralisation of scheduling. In the Board's view, however, *scheduling* is more than the exchange of references to data values. Scheduling would need to specify communication delays, a non or delayed response of a service provider, non-availability of a file system for storing values, and other events.
 - 3.2 Secondly, an *active role* of service providers compared to a passive role in traditional SOA architectures. The fact that the locations where the data was stored were under the control of the relevant service provider gave them an active role. In the Board's view, however, this feature does not imply any effect other than, for example, a read/write access control for a service provider as disclosed in D1.
 - 3.3 Thirdly, a more active role of service providers made them self-contained units and allowed them to reuse previous results more easily. For example, a user might request vacancies in nearby hotels that was handled first by a hotel address service, then by a hotel category service and finally by an availability service

based on the user's requirements for the dates of the stay. If this request did not produce any available hotels, the user might want to change the category of the hotels and/or the hotel requirements. This could be done by re-using the results of the first service with different requirements for the second and third service. In the Board's view this argument is connected with the previous arguments. The general idea of using a cache to re-use previous results is considered to be an obvious normal feature as set out above. Moreover, no scheduling or "active role" of the services that might achieve this is claimed.

3.4 Finally, a more efficient exchange of data (beyond that achieved by the use of references), an adaptation of *different data formats* and a more efficient routing of data exchanged between service providers. The Board accepts that the service providers may be self-contained units, as discussed above, but according to Figure 1 and paragraphs [0010] to [0016] of the application, it is still the task-coordinating unit 202 which, as the central unit, controls the processing of a task and the exchange of data between the different service providers. It is the task which not only defines the sequencing of services, but also the data which is exchanged between them.

4. Article 56 EPC - auxiliary request

4.1 Claim 1 of the first auxiliary request adds the feature "*wherein the file directories (202A, 204A, 206A) of the task-coordinating unit (202) and the first (204) and second service provider (206) are accessible through a Uniform Resource Locator, URL*".

- 4.2 The Board agrees with the examining division that URLs were an obvious choice for locating resources. Furthermore, this feature is disclosed in D1. The Listing 1 on page 113 illustrates the structure of a EPR reference which includes the parameter *anyURI* in line 3. Also paragraphs [0019] and [0020] refer to a URI as Uniform Resource Identifier, of which a Uniform Resource Locator (URL) is an obvious example.
- 4.3 The appellant's argument that the feature had the effect of improving flexibility in coordinating service-oriented tasks in a service-oriented architecture is not convincing, because the provision of a URL does not have any such coordinating effect on the execution of tasks. Rather, a URL is employed in the invention, see paragraph [0018], lines 7 to 8, as a conventional network access.
- 4.4 Accordingly, claim 1 of the first auxiliary request does not involve an inventive step (Article 56 EPC).

Order

For these reasons it is decided that:

The appeal is dismissed.

The Registrar:

The Chairman:



T. Buschek

W. Chandler

Decision electronically authenticated