

Internal distribution code:

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

**Datasheet for the decision
of 7 June 2016**

Case Number: T 1731/10 - 3.5.06

Application Number: 06023720.3

Publication Number: 1788481

IPC: G06F9/50, G06F9/48

Language of the proceedings: EN

Title of invention:

Hierarchical, multi-tiered mapping and monitoring architecture
for service-to-device re-mapping smart items

Applicant:

SAP SE

Headword:

Hierarchical mapping / SAP

Relevant legal provisions:

EPC 1973 Art. 84, 56

Keyword:

Inventive step - (no)

Decisions cited:

Catchword:



Beschwerdekammern
Boards of Appeal
Chambres de recours

European Patent Office
D-80298 MUNICH
GERMANY
Tel. +49 (0) 89 2399-0
Fax +49 (0) 89 2399-4465

Case Number: T 1731/10 - 3.5.06

D E C I S I O N
of Technical Board of Appeal 3.5.06
of 7 June 2016

Appellant:
(Applicant)

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

Representative:

Raible Deissler Lehmann Patentanwälte
Bolzstraße 3
70173 Stuttgart (DE)

Decision under appeal:

**Decision of the Examining Division of the
European Patent Office posted on 8 March 2010
refusing European patent application No.
06023720.3 pursuant to Article 97(2) EPC.**

Composition of the Board:

Chairman W. Sekretaruk
Members: G. Zucka
M. Müller

Summary of Facts and Submissions

I. The appeal is against the decision by the examining division, with reasons dispatched on 8 March 2010, to refuse European patent application 06023720.3, on the basis that none of the requests met the requirements of Article 56 EPC 1973. The following documents were referred to in the appealed decision:

D1 = I. Foster et al.: "The Open Grid Services Architecture, Version 1.0"; Global Grid Forum, 29 January 2005; downloaded on 16 March 2007 from <http://www.gridforum.org/documents/GWD-I-E/GFD-I.030.pdf>; XP002425854

D3 = L. Ferreira et al.: "Introduction to Grid Computing with Globus"; September 2003; downloaded on 21 December 2006 from www.redbooks.ibm.com/redbooks/pdfs/sg246895.pdf; XP002412649

II. A notice of appeal was received on 5 May 2010, the appeal fee being paid on the same day. A statement of the grounds of the appeal was received on 7 July 2010.

III. The appellant requested that the decision under appeal be set aside and a patent granted on the basis of the claims of either the main or the auxiliary request filed with the grounds of appeal. The appellant made a conditional request for oral proceedings.

IV. The board issued a summons to oral proceedings. In an annex to the summons, the board set out its preliminary, negative opinion on the appeal.

- V. The board understands that the appellant is seeking the grant of a patent based on the following documents.

Description:

Pages 1 to 82, as originally filed.

Claims:

Main request: 1 to 18, as received on 7 July 2010.

Auxiliary request: 1 to 16, as received on 7 July 2010.

Figures:

Figures 1, 2a, 2b, 3, 4a, 4b, 5 to 17, as originally filed.

- VI. Claim 1 of the main request reads as follows:

A hierarchical, multi-tiered mapping and monitoring system for use with device networks, the system comprising a plurality of heterogenous computing smart item devices with the same or different computing capabilities, the smart item devices having respective processors and memories, the memories containing instructions that when loaded onto the processors cause the processors to implement:

a global service mapper (120, 120a, 132, 902) building the hierarchy's root of the system and serving as a first addressee of a service mapping request, the global service mapper including a global metadata table (136) that stores and provides global information associated with a plurality of local networks (102, 108-114, 306-310, 324, 326) that are associated with the global service mapper, the information including, for each local network, device metadata, and information about services that are provided in the respective

local network wherein the local networks may be dispersed across a large geographical region;

a local service mapper (120, 120b, 132, 904) that is specific to at least one local network (102) of the plurality of local networks (102, 108-114, 306-310, 324, 30 326) and the devices contained therein, and that includes a local metadata table (138) that stores and provides for the at least one local network (102) device information along with service information associated with each device, quality of service information and more specific location information for each device, and that is configured to update the global metadata table (136) based thereon; and

a group leader mapper (120, 120c, 132, 132a, 906) that is contained in a dedicated device of the at least one local network that serves as a representative of a corresponding group of devices into which the at least one local network is ultimately clustered, used to provide information about respective group members as well as about services provided by the group members and configured to query the group of devices of the at least one local network (102, 108-114, 306-310, 324, 326) and aggregate group-level device metadata for transmission to the local service mapper (120, 120b, 132, 904) and updating of the local metadata table (138), wherein the group of devices is constructed in association with a group ID that is unique to the group of devices so that a service deployment proceeds at least in part with respect to the group ID and the device metadata are collected in a timely and scaleable manner, the group creation within the at least one local network being based on a number of specific parameters/requirements,

wherein at least one of the global service mapper (120, 120a, 132, 902), the local service mapper (120, 120b, 132, 904), and the group leader service mapper (120, 120c, 132, 132a, 906) are configured to determine a motivation for re-deployment of an executing service (706) on an originating device (110) by determining device metadata associated with the originating device and indicating that the originating device currently has insufficient device characteristics to continue adequately implementing the executing service based on service metadata associated with the executing service, map the executing service to a selected device (108) from among a plurality of devices that includes the originating device (110) and the selected device (108), and re-deploy the executing service on the selected device (108).

VII. Claim 1 of the auxiliary request differs from that of the main request in that some features of the group leader mapper have been deleted and an additional feature has been introduced, *viz.* that re-deploying the executing service on the selected device comprises re-instating, on the selected device, a state of the executing service on the originating device prior to re-deploying.

VIII. At the end of the oral proceedings, the chairman announced the board's decision.

Reasons for the Decision

1. *The appellant's non-attendance at the oral proceedings*

As announced in advance, the duly summoned appellant did not attend the oral proceedings. In accordance with Article 15(3) RPBA, the board relied for its decision only on the appellant's written submissions. The board was in a position to decide at the conclusion of the oral proceedings, since the case was ready for decision (Article 15(6) RPBA), and the voluntary absence of the appellant was not a reason for delaying a decision (Article 15(3) RPBA).

2. *Overview of the invention*

2.1 The application relates to a method of service migration in a wide area network (106 in figure 1) connecting local area networks (102 in figure 1) including devices such as sensors, embedded systems and RFID readers (108, 110, 112 and 114 in figure 1; paragraphs [0018] and [0019]; all references are to the application documents as originally filed). The method is executed by a "service mapper" (120 in figure 1), a three-tier hierarchy ([0064], first sentence) of service mappers called respectively "global service mapper" (120a in figure 1; [0041], "GSM"), "local service mapper" (120b in figure 1; [0041], "LSM") and "group leader service mapper" (120c in figure 1; [0041], "GLSM"). The GSM component includes a GSM metadata table (136 in figure 1; figure 4A) comprising local network device descriptions ([0045]) which may not need to include descriptions about singular devices ([0046], first sentence), whereas the LSM component includes an LSM metadata table (138 in figure 1; figure

4) storing more specific device, quality of service and device location information ([0046]). Collection of device metadata and monitoring of devices occurs at the lowest level in the hierarchy at group level ([0047], [0049], [0050]).

2.2 Figure 2A and figure 2B illustrate respectively flowcharts for a method of mapping a service request to a specific device and deploying the service thereon and a method of migrating an already deployed service from one device to another.

2.3 According to the method of service migration depicted in figure 2B, the system monitor determines the "motivation" for migrating the service (212 in figure 2B) which might be the detection of the availability of a more powerful or better-suited device (page 30, lines 8 to 12) or the detection of insufficient resources at the originating device such as low power or limited memory (page 30, lines 19 to 21). The service is then removed from the originating devices and migrated to a more suitable device thereby maintaining its state on the originating device ([0060] to [0062]).

3. *Clarity, Article 84 EPC 1973*

3.1 The expression in claim 1 of the main request (page 2, lines 21 to 22) stating that "the device metadata are collected in a timely and scal[.]able manner" is not clear: under which conditions could the collection be called timely and scalable?

3.2 It is not clear in claim 1 of the main request how "service deployment" can "proceed ... with respect to the group ID" of a group of devices (page 2, lines 20 to 21). A group ID is an identifier given to a group

and an identifier by itself cannot serve any purpose other than identifying what it is supposed to identify. It is unclear to the board how a service can be deployed "with respect to" the group ID.

4. *Inventive step, Article 56 EPC 1973*

4.1 In the decision under appeal D1 was considered to be the closest prior art and claim 1 of the main and the first auxiliary requests, which were then the first and the second auxiliary requests respectively, were found not to establish an inventive step over the disclosure of D1 (see points 2.2 and 2.3 of the appealed decision).

4.2 D1 discloses the specification of the Open Grid Services Architecture (OGSA). Grid systems integrate resources within a distributed and heterogenous environment comprising a variety of operating systems and devices such as computers, sensors, networks and storage devices (page 4, first paragraph and page 6, last paragraph). Resource sharing in a grid environment requires *inter alia* all entities within the grid to be easily addressable in a global name space (page 7, §2.2, second paragraph, first bullet point) and entity metadata to be accessed, aggregated and managed for finding, invoking and tracking entities (*ibid.*, next bullet point). In order to be scalable to thousands of resources of a widely varied and dynamic nature, the management of a grid system needs to be done in a hierarchical or peer-to-peer fashion (page 6, last sentence; page 10, penultimate paragraph). Services on a grid environment must fulfill agreed quality of service (QoS) requirements and it should be possible to migrate executing services in the grid due to performance or availability (page 8, §2.4, in particular last

bullet point). If the execution of a service is migrated, this might involve "checkpointing" the state of the service at the originating device to resume at the destination device (page 18, third bullet point).

4.3 In the decision under appeal (see page 8) the examining division considered claim 1 of the main request to differ from the disclosure of D1 in that

(1) the system includes a group leader mapper that is contained in a dedicated device of the at least one local network that serves as a representative of a corresponding group of devices into which the at least one local network is ultimately clustered, used to provide information about respective group members as well as about services provided by the group members and is configured to query a group of devices of the at least one network and aggregate group-level metadata and is configured to determine a motivation for re-deployment of said executing service, and

(2) the group of devices is constructed in association with a group ID that is unique to the group of devices so that a service deployment proceeds at least in part with respect to the group ID and the device metadata are collected in a timely and scalable manner, the group creation within the at least one local network being based on a number of specific parameters/requirements.

4.4 According to the board, however, these features are disclosed by D1, as far as they can be understood (see 3. above), and the only distinguishing feature of claim 1 of the main request is that the architecture of the system of D1 has exactly three tiers:

In particular, D1 discloses that for the architecture of the system to scale to thousands of resources, the management of a grid system needs to be done in a hierarchical or peer to peer (federated/collaborative) fashion (page 10, penultimate paragraph) and the metadata of the entities in the grid is "aggregated" and "managed" across administrative domains (page 7, §2.2, second bullet point). In view of the board this disclosure is sufficient for the skilled person to read in D1 that the disclosed architecture involves a hierarchy of levels of devices which collect information from devices in their group and aggregate them in a timely and scalable manner.

As D1 discloses the migration of executing services "for performance or availability" (page 8, §2.4, last bullet point), the board considers D1 to disclose a "motivation" for re-deploying executing services, too, *i.e.* either a decrease in the performance of the current device executing the service or a more suitable device becoming available in the grid.

As all entities in the grid have a name (page 7, §2.2, first bullet point), the board also considers the groups of devices, which would be read in D1 by the skilled person as indicated above, to have group IDs and, as one cannot allocate a task to an entity without being able to address the entity, the service deployment in D1 has to involve the group ID as well.

4.5 Thus the only distinguishing feature of claim 1 of the main request is that the claimed hierarchical architecture has exactly three tiers, whereby the particular names given to different levels of the hierarchy by the invention, *i.e.* "global service mapper", "local service mapper" and "group leader

mapper", are immaterial. The board considers this not to establish an inventive step over the disclosure of D1, as three-tier architectures are quite common in the field of computing and the board is not aware of any reason why the skilled person would not consider using a three-tier architecture in the present case.

It is observed that the present application itself also does not consider the actual number of layers in the hierarchy to be important; see page 21, lines 12 to 16 of the description.

- 4.6 Although the board considers D1 to sufficiently disclose for the skilled person a distributed grid management system, for the sake of completeness, it also mentions a passage in D3 which was cited in the appealed decision "for illustration only" (see the decision, page 9, penultimate paragraph). The board regards the cited passage to explain more explicitly what is already disclosed in D1. D3, which is an introductory document to grid computing, discloses on page 24 in the paragraph titled "Distributed grid management" that the grid is organised, in order to increase the scalability, in a hierarchy consisting of "clusters of clusters" of devices. The collection of data as well as service deployment are distributed to match the topology of the grid. Lower level clusters receive information from individual machines, aggregate it and send it to higher level nodes.
- 4.7 The board therefore concludes that claim 1 of the main request lacks inventive step over each of D1 and D3, Article 56 EPC 1973.
- 4.8 Claim 1 of the first auxiliary request differs from claim 1 of the main request, apart from a deletion of a

number of features of the group leader mapper, in that re-deploying the executing service on the selected device comprises re-instating, on the selected device, a state of the executing service on the originating device prior to re-deploying. The board considers this feature to be disclosed in D1, page 18, third bullet point, where it is stated that the state of a service should be checkpointed to ensure its restartability at another location. Thus the reasons for the lack of an inventive step in claim 1 of the main request apply *mutatis mutandis* to claim 1 of the first auxiliary request, Article 56 EPC 1973.

- 4.9 In the statement of grounds of appeal (page 1, last paragraph to page 2, third line) the appellant argues that grids provide a much more controlled environment than the claimed invention in the sense that the available nodes are under some kind of a central control and management and nothing in D1 indicates that the system described in D1 is not based on such a central control and management. The board cannot follow this argument, as it does not believe that the skilled person would consider grid systems to implicitly disclose central control and management, as alleged but not further substantiated by the appellant. Indeed both D1 and D3 provide indications to the contrary: D1 explicitly discloses that, in order to provide for scalability, the management of grid systems "needs to be done in a hierarchical or peer-to-peer (federated/ collaborative) fashion" (page 10, penultimate paragraph). D3, likewise, discloses distributed management for grids (page 24, second full paragraph).
- 4.10 The appellant also argues in the statement of grounds of appeal (page 2, lines 3 to 18) that the system of the invention comprises heterogenous computing devices

including different kinds of programmable sensor nodes at different layers of a hierarchical network. In this regard the board cannot see any difference over the disclosure of D1, which enumerates several heterogeneous device types comprising sensors on page 6, last paragraph.

4.11 The appellant further argues that the physical location of a computing node does not matter in a typical grid environment, thus location constraints are not taken into account in a grid unlike in the system of claim 1 of the main request. The appellant gives the example of a temperature sensing service for a cooling container which has to be deployed on one of the programmable sensor devices physically contained in said container (page 2, lines 18 to 28). The board considers these arguments to be the statement of obvious facts. A grid environment comprises a variety of shared devices for a variety of services. For most grid services, such as data storage or computational tasks, the physical location at which the service is executed is indeed immaterial, except for the communication overhead it might involve. A sensor service, however, such as measuring the temperature of a particular container, can only be executed by a sensor physically located in said container. The board considers it to be obvious for the skilled person to take into account all constraints relevant for deployment of a particular grid service.

4.12 The appellant also argues that smart item devices like sensor nodes are mobile and have typically only a limited battery life span, and might therefore suddenly appear or disappear. The appellant alleges that the adding or removal of nodes in a grid system is a much more controlled and manual process (page 2, line 28 to

last line). The board cannot follow this argument, as it does not believe that the skilled person would understand grid systems to be more strictly and manually controlled, as alleged but not further substantiated by the appellant. Indeed, to the contrary, D1 explicitly discloses on page 6, last sentence that grid environments are typically dynamic and evolve in ways not anticipated.

Order

For these reasons it is decided that:

The appeal is dismissed.

The Registrar:

The Chairman:



B. Atienza Vivancos

W. Sekretaruk

Decision electronically authenticated