

**Internal distribution code:**

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

**Datasheet for the decision  
of 20 September 2024**

**Case Number:** T 1132/20 - 3.5.06

**Application Number:** 13732760.7

**Publication Number:** 2862074

**IPC:** G06F9/50

**Language of the proceedings:** EN

**Title of invention:**

Multi-tenant middleware cloud service technology

**Applicant:**

Microsoft Technology Licensing, LLC

**Headword:**

Multi-tenant cloud service/MICROSOFT

**Relevant legal provisions:**

EPC Art. 56

RPBA 2020 Art. 13(2)

**Keyword:**

Late-filed requests - main request and first and second auxiliary requests (admitted)

Inventive step - main request and first and second auxiliary requests (no)

Late-filed auxiliary requests - main requests A to E and auxiliary requests 1A to 1E and 2A to 2E (not admitted)

**Decisions cited:**

T 0911/98, T 1030/06, T 0953/16



**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

Case Number: T 1132/20 - 3.5.06

**D E C I S I O N**  
**of Technical Board of Appeal 3.5.06**  
**of 20 September 2024**

**Appellant:** Microsoft Technology Licensing, LLC  
(Applicant) One Microsoft Way  
Redmond, WA 98052-6399 (US)

**Representative:** Grünecker Patent- und Rechtsanwälte  
PartG mbB  
Leopoldstraße 4  
80802 München (DE)

**Decision under appeal:** **Decision of the Examining Division of the  
European Patent Office posted on 10 January 2020  
refusing European patent application  
No. 13732760.7 pursuant to Article 97(2) EPC**

**Composition of the Board:**

**Chairman** M. Müller  
**Members:** R. de Man  
K. Kerber-Zubrzycka

## **Summary of Facts and Submissions**

I. The applicant appealed against the decision of the examining division refusing European patent application No. 13732760.7.

II. The examining division decided that the subject-matter of claim 1 of the main request and of auxiliary requests 1 and 2 lacked an inventive step over the following document:

D3: US 2011/0292792 A1, 1 December 2011.

III. In its statement of grounds of appeal, the appellant maintained the requests considered in the contested decision.

IV. In a communication accompanying the summons to oral proceedings, the board expressed the preliminary opinion that the subject-matter of claim 1 of the main request and of auxiliary requests 1 and 2 lacked an inventive step. The board also raised objections under Articles 84 and 123(2) EPC.

The board introduced the following document:

D4: "Sandbox (computer security)", Wikipedia, 2 April 2012, retrieved from [https://en.wikipedia.org/w/index.php?title=Sandbox\\_\(computer\\_security\)&oldid=485242182](https://en.wikipedia.org/w/index.php?title=Sandbox_(computer_security)&oldid=485242182).

V. With its submissions in preparation for the oral proceedings, the appellant replaced its requests with a new main request, main requests A to E, a new auxiliary

request 1, auxiliary requests 1A to 1E, a new auxiliary request 2, and auxiliary requests 2A to 2E.

VI. Oral proceedings were held on 20 September 2024. At the end of the oral proceedings, the chairman announced the board's decision.

VII. The appellant's final requests were that the decision under appeal be set aside and that a patent be granted on the basis of the claims of the main request or, in the alternative, of one of the main requests A to E, or of one of the auxiliary requests 1, 1A to 1E, or auxiliary requests 2, 2A to 2E.

VIII. Claim 1 of the main request reads as follows:

"A system (210) comprising:

a plurality of hosts (211;300), each running a plurality of virtual machines (310), each virtual machine being assigned to a tenant;

a middleware enforcement mechanism (400) that is configured to apply a middleware function (401) for network traffic corresponding to at least some of the plurality of virtual machines running on at least a particular host of the plurality of hosts;

a plurality of services (212) including a middleware management service (500) that is configured to maintain per-tenant middleware policy (501) for each of multiple tenants, wherein, based on the tenant identity, the middleware policy is referred to in order to identify the middleware function(s) to apply to the network traffic;

and wherein the middleware management service causes the middleware policy to be applied to network traffic by directing network traffic to the middleware enforcement mechanism, the middleware policy depending on an identity of a tenant that is assigned to a virtual machine that corresponds to the network traffic; and

a service coordination system that communicates with the plurality of hosts and with the plurality of services."

IX. Claim 1 of auxiliary request 1 reads as follows:

"A system (210) comprising:

a plurality of hosts (211;300), each running a plurality of virtual machines (310), each virtual machine being assigned to a tenant;

a middleware enforcement mechanism (400) that is configured to apply a middleware function (401) for incoming network traffic that is destined to a target virtual machine, the target virtual machine being one of the plurality of virtual machines;

a plurality of services (212) including a middleware management service (500) that is configured to maintain per-tenant middleware policy (501) for each of multiple tenants, wherein, based on the tenant identity, the middleware policy is referred to in order to identify the middleware function(s) to apply to the incoming network traffic, wherein the middleware management service identifies the target virtual machine to which the incoming network traffic is destined, and wherein the middleware management service causes the middleware

policy to be applied to the incoming network traffic by directing the incoming network traffic to the middleware enforcement mechanism, the middleware policy depending on an identity of a tenant that is assigned to the target virtual machine; and

a service coordination system that communicates with the plurality of hosts and with the plurality of services."

- X. Claim 1 of auxiliary request 2 differs from claim 1 of auxiliary request 1 in that the following text has been added at the end of the claim:

"wherein the middleware enforcement mechanism is a specialized virtual machine running on the same host as the target virtual machine."

- XI. Claim 1 of main request B differs from claim 1 of the main request in that:

- the text ", wherein the middleware function is a conditional filtering operation or a conditional transformation operation" has been inserted after "on at least a particular host of the plurality of hosts";
- the text "to maintain per-tenant middleware policy (501) for each of multiple tenants" has been replaced with "to maintain a middleware policy (501), wherein, for the middleware function, the middleware policy includes a policy for each of multiple tenants";
- in the penultimate paragraph, two instances of "middleware policy" have been replaced with "policy".

XII. Claim 1 of main request D differs from claim 1 of main request B in that the text "wherein the middleware function is a conditional filtering operation or a conditional transformation operation" has been replaced with:

"wherein the middleware function includes one of 1) a firewall function configured to filter incoming traffic and to not allow certain traffic types to reach a virtual machine and/or configured to filter outgoing traffic and to not allow a virtual machine to dispatch certain traffic types, 2) an antivirus function configured to perform an antivirus check of certain types of outgoing or incoming network traffic, 3) a demilitarized zone function configured to execute certain types of network traffic in a sandboxed environment to identify what the execution causes to have happen, and to filter the network traffic depending on the effects, 4) an encryption and/or decryption function configured to encrypt outgoing network traffic and/or to decrypt incoming network traffic, and 5) a compression and/or decompression function configured to compress outgoing network traffic and/or to decompress incoming network traffic".

XIII. Claim 1 of main requests A, C and E differs from claim 1 of the main request and main requests B and D, respectively, in that the text ", each virtual machine being assigned to a tenant" and ", wherein, based on the tenant identity, the middleware policy is referred to in order to identify the middleware function(s) to apply to the network traffic" has been deleted.

XIV. Auxiliary requests 1A to 1E and 2A to 2E correspond to auxiliary requests 1 and 2 with the amendments made in main requests A to E relative to the main request.



## **Reasons for the Decision**

1. The application relates to multi-tenant cloud computing.

### *Main request and auxiliary requests 1 and 2*

2. *Admittance into the appeal proceedings*

- 2.1 The main request and auxiliary requests 1 and 2, which were filed in response to the board's communication, correspond to the main request and auxiliary requests 1 and 2 considered in the decision under appeal with two minor amendments made in response to objections under Article 84 and 123(2) EPC raised for the first time in the board's communication.

- 2.2 Since these amendments represent a reasonable reaction to the board's objections raised in its communication and do not raise new issues, the board is of the opinion that exceptional circumstances are present and admits these requests into the appeal proceedings (Article 13(2) RPBA).

### *Main request*

3. *Interpretation of claim 1*

- 3.1 Claim 1 of the main request is directed to a system comprising a plurality of hosts, each running a plurality of virtual machines, each virtual machine being assigned to a tenant.

The system further includes a "middleware enforcement mechanism", a plurality of "services" and a "service coordination system".

One of the services is a "middleware management service". The remaining services are not further specified. The service coordination system "communicates with the hosts and with the plurality of services.

3.2 Claim 1 defines the "middleware enforcement mechanism" and the "middleware management service" as follows:

- the middleware enforcement mechanism is configured to apply a middleware function for network traffic corresponding to [some of the virtual machines];
- the middleware management service is configured to maintain [per-tenant middleware policies], wherein, based on the tenant identity, the middleware policy is referred to in order to identify the middleware function(s) to apply to the network traffic;
- wherein the middleware management service causes the middleware policy to be applied to network traffic by directing network traffic to the middleware enforcement mechanism, the middleware policy depending on an identity of a tenant that is assigned to a virtual machine that corresponds to the network traffic.

3.3 The meaning of the term "middleware" depends on the context in which it is used. The present application uses the term without trying to define it. The appellant argued that, in the context of a multi-tenant cloud computing system, the term "middleware function" encompassed conditional filtering or transformation functions applied to network traffic, examples being a firewall function and an encryption/decryption function

(see also the amendments proposed in main requests B and D as set out in sections XI. and XII. above).

In the board's view, the term "middleware" limits the "enforcement mechanism", "management service" and "policy" of claim 1 only in the sense that they relate to such "middleware functions".

- 3.4 In Figure 5 of the application, a middleware policy 501 is depicted as a matrix of policies 501Xy, where X refers to the tenant and y refers to the "function" (see paragraph [0041] of the description).

Hence, "based on the tenant identity, the middleware policy is referred to in order to identify the middleware function(s) to apply to the network traffic" is intended to express that the set of functions y to be applied to the network traffic of a tenant X is determined by the middleware policy.

- 3.5 The claim further states that "the middleware policy" is applied to network traffic. At the oral proceedings, the appellant argued that this does not just mean that the middleware policy is used to identify the functions to be applied. Rather, the per-tenant middleware policy includes configurations of the relevant functions (corresponding to the policies 501Xy in Figure 5). For the purpose of assessing inventive step, the board will adopt this interpretation.

#### 4. *Inventive step*

- 4.1 Document D3 discloses a multi-tenant computing cloud comprising a plurality of hosts 255 and 265, each host running a plurality of virtual machines 230, 235, 270 and 275 (see paragraphs [0032] and [0033] and

Figure 2). Each virtual machine is assigned to a service application, which is owned by a tenant (paragraphs [0033] and [0034]).

The system enforces per-tenant quality-of-service (QoS) policies, which have been laid down in contracts (paragraph [0035]). This includes limiting the rate of outgoing network traffic generated by a virtual machine in accordance with the QoS policies by means of a scheduler 401 and token-bucket queues 231, 236, 271 and 276 (paragraphs [0043], [0049] and [0050] and Figure 4). A routing mechanism 410 directs network traffic to the token-bucket queues (paragraph [0046] and Figure 4).

- 4.2 The board considers that the rate-limiting functionality implemented by the token-bucket queues is a "middleware function" within the meaning of claim 1 (see point 3.3 above).

Hence, the scheduler 401 and token-bucket queues 231, 236, 271 and 276 form a "middleware enforcement mechanism" for enforcing the QoS policies. The routing mechanism 410 represents a "middleware management service" that directs network traffic to the middleware enforcement mechanism.

- 4.3 At the oral proceedings, the appellant argued that the claimed system differed from the system of document D3 in that it supported a plurality of middleware functions, whereas document D3 only supported one, and in that it required the middleware policies for these middleware functions to be maintained in a central location.

The board agrees with the appellant that document D3 does not disclose any other "middleware function".

However, it does not agree that claim 1 requires the middleware policies to be maintained "in a central location". Claim 1 defines the "middleware management service" and the other components of the system only at a conceptual level and does not rule out that each of these components is implemented in a distributed manner.

4.4 The subject-matter of claim 1 therefore differs from the disclosure of document D3 in that:

- (a) for each tenant, the middleware policy includes configurations for a tenant-specific set of middleware function(s);
- (b) there are further "services" and a "service coordination system" which communicates with the hosts and the services.

4.5 As for difference (b), since these features are not specified in any detail and do not interact with the other features of the claim to solve a technical problem, they cannot support an inventive step.

4.6 As for difference (a), although document D3 does not disclose any other "middleware function" such as firewall or encryption/decryption functionality, these types of functionality are well known in the art. In the context of document D3, a cloud-computing provider's decision to provide such functionality on an optional basis and configured in accordance with the tenants' individual preferences (as negotiated in per-tenant policy agreements or "contracts") is a non-technical business decision.

To implement this business decision, the skilled person would add suitable firewall and encryption/decryption

components to the cloud-computing system of document D3 and direct the relevant network traffic to those components. They would thereby arrive at feature (a) without the exercise of inventive skill.

4.7 Hence, the subject-matter of claim 1 lacks an inventive step (Article 56 EPC).

*Main requests A to E*

5. *Admittance into the appeal proceedings*

5.1 Main requests A to E, which were filed in response to the board's communication, are based on the main request and include various amendments intended to address objections under Articles 84 and 123(2) EPC raised for the first time in the board's communication.

5.2 The board has admitted the main request into the appeal proceedings and has dealt with it not under Article 84 or 123(2) EPC but under Article 56 EPC. As the appellant confirmed at the oral proceedings, the amendments made in main requests A to E were not intended to address the board's inventive-step objection. In these circumstances, the board's fresh objections under Articles 84 and 123(2) EPC do not represent exceptional circumstances which can justify the admittance of main requests A to E (see also decision T 953/16, Reasons 3.1).

5.3 Hence, the board does not admit main requests A to E into the appeal proceedings (Article 13(2) RPBA).

*Auxiliary request 1*

6. *Inventive step*

6.1 Claim 1 of auxiliary request 1 adds to claim 1 of the main request the following features:

(c) the middleware enforcement mechanism applies the middleware function to "incoming network traffic that is destined to a target virtual machine";

(d) the middleware management service identifies the target virtual machine to which the incoming network traffic is destined.

6.2 The board agrees with the appellant that document D3 discloses limiting the rate of outgoing but not of incoming network traffic.

6.3 However, the obvious addition of firewall functionality (see point 4.6 above) results in a middleware function being applied to "incoming network traffic that is destined to a target virtual machine" in accordance with feature (c) and in the middleware management service identifying "the target virtual machine to which the incoming network traffic is destined" in order to direct the incoming network traffic to the firewall enforcement mechanism in accordance with feature (d).

6.4 At the oral proceedings, the appellant argued, albeit primarily in connection with auxiliary request 2, that, in the context of the problem-and-solution approach, once the closest prior art had been chosen and the initial feature mapping had been made, it was no longer allowed to deviate from those choices. Since document D3 was the closest prior art, and since the only

middleware enforcement mechanism disclosed in document D3 were the token-bucket queues, an inventive-step reasoning which mapped the middleware enforcement mechanism to a different entity was in conflict with the problem-and-solution approach. The board notes that the appellant's argument also applies to the board's reasoning for auxiliary request 1.

6.4.1 However, the problem-and-solution approach as developed in the case law of the board of appeal is not a mechanistic tool for assessing inventive step but should rather be seen as a useful framework which can help answering in an objective manner the question whether, having regard to the state of the art, the skilled person could and would have arrived at the invention. The problem-and-solution approach must be applied having in mind the purpose of Article 56 EPC, which is to prevent a patent from being granted on routine or otherwise obvious modifications of what, at the effective filing date, had been available to the public.

6.4.2 In the problem-and-solution approach, the skilled person is not presented with a partial feature mapping as starting point but with technical subject-matter, typically disclosed in a document, and the question to be answered is whether they could and would have modified this starting point to obtain something that falls within the terms of the claim.

In the present case, the appellant has not disputed that the skilled person "could" have added firewall functionality to the system disclosed in document D3, i.e. there is no reason why such a modification was technically infeasible or outside the abilities of the



skilled person. And in point 4.6, the board has explained why the skilled person "would" have done so.

- 6.5 Hence, the amendments made in auxiliary request 1 do not overcome the objection of lack of inventive step over document D3 (Article 56 EPC).

*Auxiliary requests 1A to 1E*

7. *Admittance into the appeal proceedings*

- 7.1 Auxiliary requests 1A to 1E were filed in response to the board's communication and include the same amendments, intended to address objections under Articles 84 and 123(2) EPC, as made in main requests A to E, but now with respect to auxiliary request 1.

- 7.2 Since the board has admitted auxiliary request 1 into the appeal proceedings and has dealt with it not under Article 84 or 123(2) EPC but under Article 56 EPC, the board's reasons for not admitting main requests A to E (see point 5. above) also apply to auxiliary requests 1A to 1E.

- 7.3 Hence, the board does not admit auxiliary requests 1A to 1E into the appeal proceedings (Article 13(2) RPBA).

*Auxiliary request 2*

8. *Inventive step*

- 8.1 Claim 1 of auxiliary request 2 adds to claim 1 of auxiliary request 1 the following feature:

(e) the middleware enforcement mechanism is a specialized virtual machine running on the same host as the target virtual machine.

8.2 The board notes that "the middleware enforcement mechanism" of feature (e), read in combination with feature (c), refers to the part of the middleware enforcement layer of the application's multi-tenant cloud-computing system which applies a middleware function to "incoming traffic that is destined to" the target virtual machine, i.e. to a particular virtual machine.

8.3 Referring to paragraph [0057] of the description, the appellant argued that feature (e) achieved the technical effect of improving security and provided a safer way to allow third parties to add middleware functions while reducing risk of harm to the host. Starting from document D3, the skilled person would not have considered implementing the middleware enforcement mechanism as a specialised virtual machine, since the only middleware enforcement mechanism disclosed in document D3 were the token-bucket queues, which were not provided by third parties.

8.4 The board notes that running software components in virtual machines is well-known in the art, and in the system of document D3 each host already runs software components inside a plurality of virtual machines (see paragraph [0032] and [0033]).

Moreover, the alleged technical effect of improved security by running untrusted code inside a "sandbox" in the form of a virtual machine is also well-known in the art, as evidenced by document D4.

8.5 The appellant's argument that the skilled person would not have run the token-bucket queues in a virtual machine because the token-bucket queues of document D3 were not provided by untrusted third parties is not convincing. If security were not improved by placing the token-bucket queues in a virtual machine, this would just mean that the alleged effect is not achieved over the system of document D3 and that, lacking a technical effect, such a modification of the system of document D3 is an arbitrary and thus obvious modification.

However, since the token-bucket queues of document D3 process outgoing traffic and not incoming traffic, they in any event do not correspond to a middleware enforcement mechanism that is configured to apply a middleware function "for incoming traffic", as required by feature (e) read in combination with feature (c).

8.6 As noted in points 4.6 and 6.3 above, the board considers adding firewall functionality to the system of document D3 to be obvious. To carry out this modification, the skilled person inevitably has to make a number of implementation choices. Since it is well-known to place software components in a virtual machine, and since this well-known implementation choice in the context of document D3 at most achieves the known effect of improving security, the board considers that the skilled person would take this implementation decision without the exercise of inventive skill, thereby obtaining a system complying with feature (e).

8.7 The appellant's argument that, since the only middleware enforcement mechanism disclosed in document D3 were the token-bucket queues, it was illegitimate to

map the middleware enforcement mechanism of claim 1 of auxiliary request 2 to a different entity, is not convincing for the reasons given in point 6.4 above.

- 8.8 The appellant further argued that, under the problem-and-solution approach, the skilled person would not, in a first step, decide to add firewall functionality and then, in a second step, place the corresponding enforcement functionality in a virtual machine in accordance with feature (e).

However, as mentioned in point 8.6 above, the decision to add firewall functionality to the closest prior art inevitably leads to a number of implementation choices. The skilled person does not stop at the decision to add the functionality; they also carry out this decision by modifying the closest prior art accordingly (see also decision T 1030/06, Reasons 20).

Moreover, assessing the contributions of two distinguishing features A and B separately is not generally precluded even if A is a prerequisite for B, provided that, in the context of the invention, the two features do not combine to provide a technical effect going beyond the sum of their individual effects (see e.g. decision T 911/98, Reasons 8 and 10).

- 8.9 Hence, the subject-matter of claim 1 lacks an inventive step (Article 56 EPC).

#### *Auxiliary requests 2A to 2E*

#### 9. *Admittance into the appeal proceedings*

- 9.1 Auxiliary requests 2A to 2E were filed in response to the board's communication and include the same

amendments, intended to address objections under Articles 84 and 123(2) EPC, as made in main requests A to E, but now with respect to auxiliary request 2.

- 9.2 Since the board has admitted auxiliary request 2 into the appeal proceedings and has dealt with it not under Article 84 or 123(2) EPC but under Article 56 EPC, the board's reasons for not admitting main requests A to E (see point 5. above) also apply to auxiliary requests 2A to 2E.
- 9.3 Hence, the board does not admit auxiliary requests 2A to 2E into the appeal proceedings (Article 13(2) RPBA).
10. Since none of the requests admitted into the proceedings is allowable, the appeal is to be dismissed.

## Order

### **For these reasons it is decided that:**

The appeal is dismissed.

The Registrar:

The Chairman:



L. Stridde

Martin Müller

Decision electronically authenticated