Internal distribution code:

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

Datasheet for the decision of 30 October 2023

Case Number: T 2027/20 - 3.5.06

Application Number: 12745327.2

Publication Number: 2673706

IPC: G06F9/455, G06F11/30

Language of the proceedings: EN

Title of invention:
VIRTUAL SWITCH INTERCEPTOR

Applicant:
Microsoft Technology Licensing, LLC

Headword:
Virtual Machine Migration/MICROSOFT

Relevant legal provisions:
EPC Art. 56, 84
RPBA 2020 Art. 13(2)

Keyword:
Claims - clarity (no)
Inventive step - (no)
Amendment after summons - exceptional circumstances (no)

Decisions cited:

Catchword:

Europäisches Patentamt
European Patent Office
Office européen des brevets

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 2027/20 - 3.5.06

D E C I S I O N
of Technical Board of Appeal 3.5.06
of 30 October 2023

Appellant: (Applicant)	Microsoft Technology Licensing, LLC 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 31 July 2020 refusing European patent application No. 12745327.2 pursuant to Article 97(2) EPC.**

Composition of the Board:

Chairman	M. Müller
Members:	G. Zucka
	A. Jimenez

Summary of Facts and Submissions

I. The appeal is against the decision by the examining division, dispatched with reasons on 31 July 2020, to refuse European patent application 12745327.2, on the basis that the subject-matter of the independent claims 1 and 5 of the main and the auxiliary request was not inventive (Article 56 EPC), and claim 1 of the auxiliary request was not clear (Article 84 EPC). The following document cited during the first instance proceedings is referred to in the present decision:

D5: Jason Sonnek *et al.*: "Starling: Minimizing Communication Overhead in Virtualized Computing Platforms Using Decentralized Affinity-Aware Migration", Parallel Processing (ICPP), 2010 39th International Conference on, IEEE, Piscataway, NJ, USA, 13 September 2010 (2010-09-13), pages 228-237, XP031773708, ISBN: 978-1-4244-7913-9 (PDF version from the conference proceedings).

II. A notice of appeal was received on 25 September 2020, the appeal fee being paid on the same day. A statement of grounds of appeal was received on 9 December 2020.

III. The appellant requested that the decision under appeal be set aside and a patent granted on the basis of claims 1 to 9 of the main or the auxiliary request which were the object of the appealed decision, both re-filed with the statement of 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 opinion on the appeal.

V. On 7 September 2023, the appellant filed three additional auxiliary requests, labelled "Auxiliary Request 0a" to "Auxiliary Request 0c".

VI. During the oral proceedings, the appellant filed another auxiliary request, labelled "Auxiliary Request 0".

VII. The appellant requests that the decision under appeal be set aside and a patent be granted on the basis of claims 1 to 9 of

- the main request underlying the decision under appeal and re-filed with the statement of grounds of appeal;

- auxiliary request 0 filed during the oral proceedings;

- auxiliary requests 0a-0c filed with letter of 7 September 2023; or

- the auxiliary request underlying the decision under appeal and re-filed with the statement of grounds of appeal.

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

"A method of facilitating application management, comprising:

employing at least one processor configured to execute computer-executable instructions stored in memory to perform the following acts:

observing virtual-application messages, external to application-hosting virtual machines running on different host servers; and

performing an action as a function of the messages; wherein the method further comprises:

gathering information regarding which virtual machine is communicating with which other virtual machines, wherein the gathering information comprises identifying message pairs of requests and replies for a particular technology and stack layer and associating the messages with a specific application; and wherein performing an action comprises:

determining placement of one or more of the virtual machines as a function of the gathered information so that communicating virtual machines are placed on the same host.

IX. Claim 1 of auxiliary request 0 differs from that of the main request in that "as a function of the gathered information" is replaced by "as a function of the identified message pairs".

X. Claim 1 of auxiliary request 0a differs from that of the main request in that it specifies that the method comprises observing and analyzing virtual-application messages, that the action is performed based on the analyzed messages, and that the step of gathering the information is part of the analyzing step.

XI. Claim 1 of auxiliary request 0b differs from that of the main request in that it specifies that the performing an action comprises placing the one or more

of the virtual machines so that communicating virtual machines are placed on the same host.

XII. Claim 1 of auxiliary request 0c corresponds to a combination of that of auxiliary requests 0a and 0b.

XIII. Claim 1 of the auxiliary request filed with the statement of grounds of appeal contains the following features in addition to the features of claim 1 of the main request:

(a) The application management takes place in a two-tier hierarchical management structure comprising a first tier and a second tier;
(b) the host servers are part of the first tier (and the management server is part of the second tier; see feature (e));
(c) each of the virtual machines includes a network protocol stack that enables cross-virtual-machine communication;
(d) the performing of an action further comprises moving a virtual machine during execution from host to host, including performing virtual-machine live migration within a Transport Control Protocol, TCP, timeout;
(e) moving the virtual machine comprises running a copy of the network protocol stack of the virtual machine being moved on a management server, which is part of the second tier.

XIV. Claim 5 of each request relates to a system having apparatus features corresponding to the method features of claim 1 of that request.

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

Reasons for the Decision

1. *The invention*

The application relates to virtual applications in virtual machines running on host servers (see claim 1).

According to the statement of grounds of appeal (section II.1 "The invention"), the aim is to provide for the situation where a single server can host a large number of virtual machines.

To this end, messages between applications are observed externally, to determine which applications communicate with one another, and communicating virtual machines are placed on the same host (*ibid.*).

2. *Main request - inventive step; Article 56 EPC*

2.1 The board agrees with the examining division that D5 is a suitable starting point for an inventive step analysis.

2.2 D5 discloses a method of facilitating application management, comprising:

employing at least one processor configured to execute computer-executable instructions stored in memory (page 232, right-hand column, last paragraph, first sentence) to perform the following acts:

observing virtual-application messages, external to application-hosting virtual machines running on different host servers (page 229, left-hand column, third paragraph, first sentence); and

performing an action as a function of the messages (page 229, left-hand column, first paragraph, penultimate sentence).

The method further comprises:

gathering information regarding which virtual machine is communicating with which other virtual machines, wherein the gathering information comprises associating the messages with a specific application (page 229, left-hand column, third paragraph, first sentence; page 229, second paragraph, second sentence; and page 230, left-hand column, last paragraph, first and third sentence); and
wherein performing an action comprises:
determining placement of one or more of the virtual machines as a function of the gathered information so that communicating virtual machines are placed on the same host (page 230, right-hand column, last full paragraph, last sentence; page 229, left-hand column, first paragraph, penultimate sentence).

2.3 As indicated in the appealed decision (point 3.2), the difference between the subject-matter of claim 1 and the disclosure of D5 is therefore that the step of gathering information comprises identifying message pairs of requests and replies for a particular technology and stack layer.

2.4 According to the appellant (statement of grounds of appeal, page 6, first paragraph), the distinguishing feature achieves the effect of enabling a more fine-tuned approach to migrating VMs that communicate with one another.

2.5 The appellant has not convinced the board that such improved fine-tuning brings an advantage at a technical level.

The appellant gives the example of "cases where the volume of traffic between two applications is comparatively small, but it is desirable to have low latency in the communication between these applications" (*ibid.*, page 4, last paragraph).

In the board's view, however, such reduced latency between particular applications would not be a direct consequence of the identification of message pairs of requests and replies for a particular technology and stack layer. Instead, it would depend on how the placement of the VMs as a function of the gathered information is determined, which the claim does not limit to message pairs, as they are only "comprised" in it.

Moreover, upon examination of the identified message pairs, one would first need to establish that certain VMs communicate with each other and possibly also establish the relative volume of the traffic between them. One would then possibly decide to prioritise certain pairs of VMs, and ultimately decide that all or a selection of the VMs which communicate with each other should be placed on the same host. The claim, however, does not contain the corresponding method steps. The method as it is claimed therefore does not achieve the alleged effect.

2.6 During the oral proceedings, the appellant stressed that whereas the method disclosed in D5 merely measures the volume data traffic between VMs, the method of claim 1 establishes whether two-way communication in

the form of requests and replies takes place, as opposed to mere one-way (e.g. broadcast) communication.

This firstly means that fewer messages need to be observed before a decision about VM migration can be made, as it is only necessary to establish that some (two-way) communication takes place, whereas in D5 the volume of traffic can only be measured after a sufficient quantity of messages has been exchanged.

Secondly, the method of claim 1 ensures that only VMs which are communicating with each other are placed on the same host, and not VMs where only a one-way traffic exists from one VM to the other.

The appellant sees the technical effect in a reduction of the latency of (two-way) communication between VMs, in contrast to D5, where the aim is to reduce overall traffic.

2.7 The board observes that the claim is not so limited as it is understood in this argument. The gathered information may indeed contain much more information than identification of message pairs of requests and replies, and it is not apparent which advantage this identification brings over the observation of message traffic which takes place in the method of D5.

2.8 The board therefore holds that the subject-matter of claim 1 of the main request is not inventive (Article 56 EPC).

3. *Auxiliary request 0*

3.1 The appellant argued during the oral proceedings that the amendment in claim 1 of auxiliary request 0 intends

to address the arguments which the board formulated during those proceedings against the presence of an inventive step, specifically the argument that VM migration according to claim 1 of the main request does not only depend on observed message pairs. The amendment in auxiliary request 0 aims to make this explicit.

The board however points out that not only was the reasoning that a detection of message pairs has no apparent technical effect already present in its summons (point 5.5), but even the appealed decision already argued (point 3.3) that such detection has no technical effect. For that reason, auxiliary request 0 could and should already have been filed together with the statement of grounds of appeal.

3.2 Regarding Article 123(2) EPC, the appellant argued during the oral proceedings that the amendment finds a basis in the description as originally filed:

- Par. [0029] discloses that pairs of requests and replies are identified and provided to the message processor for further processing;
- Par. [0030] discloses that any number of operations can be performed with respect to the message data, which necessarily means the message data of the identified message pairs of requests and replies, given that these are the ones that are passed on to the message processor; and
- Par. [0042] discloses that message data can be analyzed etc., and the results can be used among other things to place, migrate and update virtual machines.

According to the appellant, the skilled person would deduce from those passages that the action is performed as a function of the identified message pairs, not just generally as a function of the messages, or even more broadly as a function of the gathered information.

The board however *a priori* fails to see an unambiguous indication in the cited passages, or in the application as a whole, that placement of virtual machines could take place as a function of the identified message pairs. In particular, it fails to see support for the appellant's argument that the message data in par. [0030] must be construed as referring to the pairs of requests and replies, or that the message data in par. [0042] is necessarily the same as the message data in par. [0030].

3.3 Regarding Article 56 EPC, the board observes that *a priori* nothing in the application supports the argument, starting from the teaching of D5, that a placement of virtual machines as a function of identified message pairs (for a particular technology and stack layer), instead of observed traffic, would provide any specific advantage. In particular, even if one could say that the effect of improved latency is implicit from the application when compared with D5 in certain *specific* situations, it is not apparent why those specific situations would be technically relevant.

3.4 For those reasons, auxiliary request 0 is not admitted under Article 13(2) RPBA.

4. *Auxiliary requests 0a to 0c*

Auxiliary requests 0a to 0c *prima facie* do not overcome the specific objection regarding lack of inventive step over the disclosure of D5 raised in the summons, since it was in any case already assumed in the summons that some kind of "analysis" of the messages takes place and that communicating VMs are placed on the same host. Auxiliary requests 0a to 0c are therefore also not admitted under Article 13(2) RPBA.

5. *Auxiliary request filed with the statement of grounds of appeal - Clarity; Article 84 EPC*

5.1 According to the appealed decision (point 5.1), feature (d) above is unclear because it is not clear how a virtual machine of an arbitrary size can be live migrated in a consistent manner during an arbitrary TCP timeout over a network with an arbitrary transmission capacity.

In this respect, the board firstly notes that a TCP timeout can be adjusted and that it has a priori no minimum value, except that it is normally expressed in seconds. Secondly, the transmission capacity of a network depends on a great variety of factors, such as the network topology, the available communication channels, and the stability of the network. All those factors together can lead to substantial delays in VM migration.

5.2 The appellant submits (statement of grounds of appeal, second half of page 7) that features (d) and (e) have to be read together. The skilled person would understand that the migration is started but not necessarily finished within the TCP timeout.

Feature (e) ensures that messages that are sent after the TCP timeout is finished do not get lost.

5.3 The board observes that feature (d) is clearly worded so that the VM live migration is *performed*, i.e. also *finished* and not only started, within a TCP timeout. Also step (e) refers to a copy of the network protocol stack of the virtual machine *being moved*, i.e. during migration and not after. The passage in the description indicated as basis for the amendment (page 9, lines 5 to 8, of the description; see also the minutes of the oral proceedings before the examining division, point 3.5) indicates nothing else. In this respect, the board agrees with the examining division's objection regarding lack of clarity.

5.4 During the oral proceedings, the appellant firstly reiterated its position that for a skilled person the wording "performing VM live migration within a TCP timeout" does not mean that the migration needs to be completed within such timeout, and secondly that a TCP timeout can be relatively long; the default value is for instance 300 seconds in the Microsoft Windows operating system. Thirdly, the appellant defended the point of view that, rather than rendering the claim unclear, the fact that some VMs could not be completely migrated over some networks within a TCP timeout, if that is how the claim should be interpreted, only meant that such VMs and networks are not covered by the definition of claim 1, not that the claim is unclear.

The board holds that the claim's wording is unambiguous: The live migration needs to be completed within a TCP timeout. The alternative interpretation provided by the appellant is meaningless, because "starting" refers to a point in time; it has zero

duration. The claim would fail to indicate what exactly, if anything, must be guaranteed to be completed "within" a TCP timeout. The effect of the feature on the claim's scope, if so interpreted, would be unclear.

Further, the appellant's point of view according to which VMs which cannot be migrated within such timeout are not covered by the claim is firstly contradicted by the preceding wording in the claim that "communicating virtual machines *are* placed on the same host", i.e. it does not depend on the time which such process would take.

Secondly, there is no apparent reason why the skilled person would take the fact that migration within a TCP timeout is not always possible as limiting the claim to certain (sizes of) VMs or to certain networks. The claim itself does not suggest such limitation, which would in any case be vague. Such a limitation would in fact constitute an essential feature of the invention, and one cannot assume the "synthetic propensity" of a skilled person to go as far as reading missing essential features into a claim.

5.5 Claim 1 of the auxiliary request is therefore not clear (Article 84 EPC).

Order

For these reasons it is decided that:

The appeal is dismissed.

The Registrar: The Chairman:

L. Stridde Martin Müller

Decision electronically authenticated