Internal distribution code:

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

Datasheet for the decision of 14 May 2025

Case Number: T 0157/23 - 3.5.06

Application Number: 17852247.0

Publication Number: 3509253

IPC: G06F9/455, H04L12/46

Language of the proceedings: EN

Title of invention:
INTER-CLOUD COMMUNICATION METHOD AND RELATED DEVICE, INTER-CLOUD COMMUNICATION CONFIGURATION METHOD AND RELATED DEVICE

Applicant:
Huawei Cloud Computing Technologies Co., Ltd.

Headword:
Inter-cloud communication/HUAWEI

Relevant legal provisions:
EPC Art. 56
RPBA 2020 Art. 13

Keyword:

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

Case Number: T 0157/23 - 3.5.06

D E C I S I O N
of Technical Board of Appeal 3.5.06
of 14 May 2025

Appellant: (Applicant)	Huawei Cloud Computing Technologies Co., Ltd. Huawei Cloud Data Center Jiaoxinggong Road Qianzhong Avenue Gui'an New District Guizhou 550025 (CN)
Representative:	Gill Jennings & Every LLP The Broadgate Tower 20 Primrose Street London EC2A 2ES (GB)
Decision under appeal:	**Decision of the Examining Division of the European Patent Office posted on 14 July 2022 refusing European patent application No. 17852247.0 pursuant to Article 97(2) EPC.**

Composition of the Board:

Chairman	M. Müller
Members:	T. Alecu
	A. Jimenez

Summary of Facts and Submissions

I. The appeal lies from the decision of the Examining Division to refuse the application for lack of inventive step.

II. The decision cited inter alia documents
D1: US 2015/063102 A1, 5 March 2015, and
D2: VIRTUAL EXTENSIBLE LOCAL AREA NETWORK (VXLAN): A FRAMEWORK FOR OVERLAYING VIRTUALIZED LAYER 2 NETWORKS OVER LAYER 3 NETWORKS; RFC7348.TXT, INTERNET ENGINEERING TASK FORCE, IETF; STANDARD, INTERNET SOCIETY (ISOC) 4, 27 August 2014

III. The Appellant requested with the statement of grounds of appeal that the decision of the Examining Division be set aside and that a patent be granted on the basis of a main request or one of two auxiliary requests, all refiled with the statement of grounds of appeal, and corresponding to those underlying the decision (filed on 17 February 2022).

IV. The Board informed the Appellant of its preliminary opinion that the decision under appeal was correct. In response, the Appellant filed a further (third) auxiliary request.

V. The decision was taken and announced during oral proceedings.

VI. Claim 1 of the main request defines:

An inter-cloud communication method, used to send a data packet of a first virtual machine to a second virtual machine, wherein the first virtual machine and

the second virtual machine are located in a first cloud and a second cloud respectively, the first virtual machine and a first virtual switch are deployed on a first computing node, and the second virtual machine and a second virtual switch are deployed on a second computing node; and the method comprises:

determining (S42), by the first virtual switch after receiving the data packet of the first virtual machine, that a first-hop node of the data packet is a gateway node in the second cloud;

sending (S43), by the first virtual switch, the data packet to the gateway node through a tunnel between the first computing node and the gateway node, wherein the tunnel between the first computing node and the gateway node is established in a Layer 2 virtual network;

determining (S44), by the gateway node, after receiving the data packet, that a second-hop node of the data packet is the second computing node;

sending (S45), by the gateway node, the data packet to the second computing node through a tunnel between the gateway node and the second computing node, wherein the tunnel between the gateway node and the second computing node is established in a Layer 2 virtual network; and

sending (S46), by the second virtual switch deployed on the second computing node, the data packet to the second virtual machine.

VII. Claim 1 of the first auxiliary request differs form that of the main request by adding the following at the end of the claim:

wherein the determining (S42), by the first virtual switch after receiving the data packet of the first virtual machine, that a first-hop node of

the data packet is a gateway node in the second cloud comprises:

after receiving the data packet of the first virtual machine, determining, by the first virtual switch based on a destination address of the data packet and a pre-generated first forwarding relationship table, the gateway node corresponding to the destination address as the first-hop node, wherein the first forwarding relationship table is used to indicate a correspondence between the second virtual machine and the gateway node in the second cloud.

VIII. Claim 1 of the second auxiliary request differs form that of the first auxiliary request by adding the following at the end of the claim:

wherein the determining (S44), by the gateway node after receiving the data packet, that a second-hop node of the data packet is the second computing node comprises:

after receiving the data packet, determining, by the gateway node based on a destination address of the data packet and a pre-generated second forwarding relationship table, the second computing node corresponding to the destination address as the second-hop node, wherein the second forwarding relationship table is used to indicate a correspondence between the second virtual machine and the computing node on which the second virtual machine is located.

IX. Claim 1 of the third auxiliary request defines:

A method of inter-cloud communication between a plurality of clouds, used to send a data packet of a first virtual machine to a second virtual machine,

wherein the first virtual machine and the second virtual machine are located in a transmit end cloud and a receive end cloud respectively, the transmit end cloud and the receive end cloud are different clouds within the plurality of clouds, the first virtual machine and a first virtual switch are deployed on a first computing node within the transmit end cloud, the second virtual machine and a second virtual switch are deployed on a second computing node within the receive end cloud, and each cloud within the plurality of clouds comprises a respective gateway node; and the method comprises a sequence of steps comprising the following steps performed at the transmit end cloud:

determining (S42), by the first virtual switch after receiving the data packet of the first virtual machine, that a first-hop node of the data packet is the gateway node in the receive end cloud; and

sending (S43), by the first virtual switch, the data packet to the gateway node through a tunnel between the first computing node and the gateway node, wherein the tunnel between the first computing node and the gateway node is established in a Layer 2 virtual network; and the sequence of steps comprises the following steps performed at the receive end cloud:

determining (S44), by the gateway node, after receiving the data packet, that a second-hop node of the data packet is the second computing node;

sending (S45), by the gateway node, the data packet to the second computing node through a tunnel between the gateway node and the second computing node, wherein the tunnel between the gateway node and the second computing node is established in a Layer 2 virtual network; and

sending (S46), by the second virtual switch deployed on the second computing node, the data packet to the second virtual machine,

wherein the sequence of steps is repeated at least once; and

the transmit end cloud used in the repeated sequence of steps is the receive end cloud used in the sequence of steps and/or the receive end cloud used in the repeated sequence of steps is the transmit end cloud used in the sequence of steps.

Reasons for the Decision

The application

1. The application relates to inter-cloud communication technologies. A cloud service provider provides to a tenant an integrated cloud platform (paragraph 3). For security and stability reasons, the platform may be deployed on a plurality of clouds; inter-cloud communication between virtual machines (VMs) of the same tenant is therefore required (paragraphs 4 and 5).

1.1 According to the application there are two prior art solutions for inter-cloud communication.

1.1.1 One is communication over the internet, which is stated to have two disadvantages: on the one hand low availability, because it is a layer 3 interworking solution, and on the other hand high communication costs (paragraphs 33 and 34).

1.1.2 The other (paragraphs 35 and 36), called "tunnelling", works by establishing a tunnel between the computing nodes where the virtual machines are located. Tunnelling is an "overlay" encapsulation technology that constructs a layer 2 "logic" network on a layer 3 or layer 4 underlay (paragraph 53). This solution is

stated to have the disadvantages that many tunnels are required (one for each pair of communicating computing nodes) and that a new tunnel must be established whenever a VM changes location.

2. The application proposes a solution according to which each computing node comprises a virtual switch communicating with the virtual machines on that node. To establish communication between a VM on a first cloud and a second VM on a second cloud, a tunnel is established between the virtual switch on the first cloud and a gateway node on the second (paragraph 50). VXLAN (virtual extensible local area network) tunnelling may be used for this purpose (paragraph 53).

2.1 A forwarding table on the switch determines whether the MAC destination address of a packet is in a different, second cloud, and, if so, the packet is redirected through the tunnel to the gateway of that second cloud (paragraphs 47 to 50). When receiving the packet, the gateway determines the destination MAC address and sends the packet to the destination virtual machine (paragraph 59) through another tunnel between the gateway node and the virtual switch on the corresponding computing node.

Main request

The decision

3. The Examining Division found claim 1 of the main request to lack inventive step over D1 (decision, reasons 2). This document teaches inter alia a communication scheme for a multi-tenant network where VMs for one tenant may be hosted on different computing hosts.

3.1 In the specific configuration of figure 5C, referred to by the Examining Division, two VMs on two different computing hosts communicate through a router using VXLAN tunnels between the nodes and the router (see paragraphs 29, 41 and 42).

3.2 The Examining Division equated the router with the gateway of the claimed invention, and also considered that virtual switches were implied on each of the hosts (decision, reasons 2.1). It therefore saw the only difference as being the application context, i.e. inter-cloud instead of inter-host communication.

3.3 It considered this difference to "*[relate] to an administrative aspect of networking, or to a particular perception/view of network configuration*" without any technical effect (decision, reasons 2.2). Furthermore, the person skilled in the art would be able to use the method of D1 for inter-cloud communication without any change, so, as the Board understands the argument, the difference was obvious anyway (decision, reasons 2.3).

The Appellant's arguments

4. The Appellant did not challenge this difference per se nor that it was obvious to the skilled person. The Appellant argued, however, that there were other differences between the claimed invention and D1, and that these involved an inventive step over D1.

5. In the Appellant's view, D1 taught a centralized architecture with a central router and a central controller which pre-determined the flows for the deployed services from one host to another (D1 paragraphs 39, 42, 63, 64). So the router itself was not on one of the hosts (or clouds) and did not act as a gateway to a

cloud. It did not determine itself, as the claim required, the next "hop", but only routed packets along pre-determined routes. So the claimed step 44 was not disclosed in D1.

5.1 Also, D1 did not explicitly mention a switch. There was no need for one, because the packets were to be sent through the VXLAN tunnels as pre-determined by the controller. A relay was sufficient for this purpose.

5.2 Whichever the device routing packets from the first host to the router, this device did not *determine* the router as the next destination; rather, the router as a destination was pre-determined by the central controller.

5.3 The *determining* in steps 42 and 44 as claimed had to be interpreted as something more than automatic routing of packets arriving at an input port to a predetermined output port according to a pre-determined table. The *determining* took place at the device, after receiving a packet.

5.4 In contrast to D1, the application taught a decentralised approach, according to which small forwarding tables were held in the switch and in the gateway, the devices being used to determine only the next hop (see e.g. application, paragraph 59). This reduced the size of the tables vis-à-vis those of D1, which pre-determined and stored flow tables for the whole path.

5.5 D1 used a centralised approach in order to achieve a number of advantages explained in paragraph 69. Renouncing this centralisation would go against the teachings

of D1, so that the person skilled in the art would not do it.

6. The Board agrees with the Appellant that D1 teaches a centralised approach where flow tables are pre-determined in advance. However, it does not see that this implies all the differences that the Appellant identifies.

6.1 First, both in D1 and in the application, the routing in the different devices takes place based on forwarding tables (see paragraph 46 and 55 in the application describing steps 42 and 44). It is true that in D1 these are pre-determined. But neither the claims nor the application exclude this. The claims simply define that the next destination is determined.

6.2 Second, the Appellant argues that switches are neither needed nor explicitly mentioned in D1. The Board remarks that D1 generically teaches the usage of flow programming to "*establish forwarding flows through switches, routers, service nodes, etc., in a mapped fashion. In essence, the flow programming function maps service flows onto the graph to obtain desired or requested packet services between network endpoints*" - (D1 paragraph 19). So the fact that a flow is pre-determined does not imply that only simple devices such as relays are to be used. The person skilled in the art will use the devices it needs depending on the network at hand.

6.3 In view of figure 5C in D1, the Board considers as a realistic scenario one in which a virtual machine on the first host not only sends packets to the second host via the VXLAN tunnel and the router, but also to other VMs on the same host. In this situation, it would

be obvious to use a switch (virtual or not) to direct traffic towards the tunnel or towards other machines on the same host.

6.4 The Board adds that switches are commonly used in VXLAN tunnelling (see D2 at point 6, page 17, first paragraph). Switching functionality may even be considered implicit, as the Examining Division stated, but this may be left open.

7. So the only other difference (in addition to the undisputed distinction between inter-host and inter-cloud communication; see points 3.2, 3.3 and 4 above) that the Board can see is that the router is placed on the second cloud, i.e. the "receiving" one.

8. If the person skilled in the art is interested in applying the solution of D1 to communication between two different clouds (cf. decision, reasons 2.3), which assertion the Appellant did not contest, the Board is of the view that, if only for practicality reasons, an obvious option is to place the router (and the controller) on one of the two clouds (and not on a third computing entity).

8.1 This does not, per se, change the centralised approach of D1. Rather, it is merely an obvious choice, given that the central components must be hosted somewhere.

8.2 The Board further points out that the communication in D1 is bidirectional (D1, paragraph 42), so that the claimed subject-matter is obtained irrespective of the cloud on which the router is placed: If the router is placed on the first host, one needs to consider a communication initiated at the other cloud. This

central router inherently acts as a gateway to the receiving cloud.

9. The Board concludes that claim 1 of the main request lacks inventive step over D1 in view of common general knowledge as evidenced by D2.

First and second auxiliary requests

10. The Board agrees with the Examining Division that the existence of forwarding tables as claimed in these two requests is implied by the functionality of the switches and routers discussed above when using VXLAN tunnelling (see e.g. D2 at 4.1, and at page 17, second paragraph).

10.1 The Appellant did not bring any further arguments.

10.2 The above conclusion that there is a lack of inventive step thus remains valid for both auxiliary requests.

Third auxiliary request

11. The Appellant argued that the new request should be admitted because the amendments (which were based on the main request) were not complex and solved all the issues regarding novelty and inventive step. They clarified the decentralized approach of the invention, according to which each cloud had its own gateway to receive packets and sent packets through a switch.

12. The Board does not see any exceptional circumstances (Article 13(2) RPBA): the Board merely confirmed the view of the Examining Division, that the claims do not define an inventive difference over D1.

12.1 Also, this would be the first time that the claim clearly defined a configuration with one gateway per cloud. If admitted, the Board would have to remit, as the matter might require a further search, which would be against procedural economy.

12.2 Further, the technical significance of having another pair of clouds repeating the transmitting and receiving steps is not clear. The Appellant appears to wish to define a system and its use through method steps.

12.3 For these reasons, the Board does not admit this request (Article 13 RPBA).

Order

For these reasons it is decided that:

The appeal is dismissed.

The Registrar: The Chairman:

L. Stridde Martin Müller

Decision electronically authenticated