Internal distribution code:

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

Datasheet for the decision of 6 December 2017

Case Number:	T 0821/15 - 3.5.06
Application Number:	12164876.0
Publication Number:	2597569
IPC:	G06F9/50
Language of the proceedings:	EN

Title of invention:
System and method for distributing processing of computer security tasks

Applicant:
Kaspersky Lab, ZAO

Headword:
Distributing computer security tasks/KASPERSKY

Relevant legal provisions:
EPC Art. 56, 123(2)

Keyword:
Amendments - added subject-matter (yes)
Inventive step (no)

Decisions cited:

This datasheet is not part of the Decision.
It can be changed at any time and without notice.

Catchword:

This datasheet is not part of the Decision.
It can be changed at any time and without notice.

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 0821/15 - 3.5.06

DECISION
of Technical Board of Appeal 3.5.06
of 6 December 2017

Appellant: Kaspersky Lab, ZAO
(Applicant) 39A/3 Leningradskoe Shosse
Moscow 125212 (RU)

Representative: Sloboshanin, Sergej
V. Füner, Ebbinghaus, Finck, Hano
Mariahilfplatz 3
81541 München (DE)

Decision under appeal: **Decision of the Examining Division of the European Patent Office posted on 17 October 2014 refusing European patent application No. 12164876.0 pursuant to Article 97(2) EPC.**

Composition of the Board:

Chairman W. Sekretaruk
Members: M. Müller
S. Krischer

Summary of Facts and Submissions

I. The appeal lies against the decision of the examining division, with reasons dispatched on 17 October 2014, to refuse European patent application No. 12 164 876.0. The following documents were cited:

D2: de Camargo R Y *et al.*, "Grid: An Architectural Pattern", Proceedings of the 11th Conference on Pattern Languages of Programs, 2004,

D4: Ferreira L *et al.*, "Introduction to Grid Computing with Globus", IBM Redbooks, 2003, and

D5: Oberheide J *et al.*, "CloudAV: N-Version Antivirus in the Network Cloud", Proceedings of the 17th Usenix Security Conference, 2008,

and it was found that claim 1 lacked inventive step over D5 in combination with common general knowledge of the person skilled in the art as known from D2 or D4.

II. Notice of appeal was filed on 17 December 2014, the appeal fee being paid on the same day. A statement of grounds of appeal was filed on 6 February 2015. The appellant requested that the decision be set aside and that a patent be granted on the basis of claims 1-15 of the main request, or on the basis of the first or second auxiliary requests, all as filed with the grounds of appeal, in combination with figures 1, 2, 3A-3E, 4, 5A, 5B, and 6-9 and description pages 2, 6, 7, and 9-27 as originally filed, and description pages 1, 3, 3b, 4, 5, and 28 as filed on 9 January 2013 and 3a as filed on 20 September 2013.

III. By way of an annex to the summons to oral proceedings, the board informed the appellant of its preliminary opinion that the pending claims were deficient under Articles 56, 83, 84, and 123(2) EPC.

IV. In response to the summons, by way of letter dated 3 November 2017 the appellant filed amended claims 1-13 as its sole request.

V. Independent claim 1 reads as follows:

"1 A computer system for operation in a distributed computation system in which security-related tasks are delegated, the computer system comprising:

computing hardware including a processor (904), a memory device (906), a user interface (908, 914), and a communications interface (912);

a distributed computing service module (110) adapted to:

receive a request for a distribution of a security-related task for analyzing an unknown file for malware;

divide the received security-related task into a plurality of distinct task parts for performing non-overlapping types of antivirus analysis using respective malware databases to achieve a common objective of analyzing the unknown file for malware, and delegate each of the plurality of distinct task parts to a different one of multiple remote agent computers (100) for execution in response to a suitability determination as to whether each of the multiple remote agent computers (100) is suitable to perform the execution of the respective distinct task part, and

determine computing capacity requirements to perform each of the distinct task parts that are to be delegated to the respective remote agent computers (100) for execution, wherein the distributed computing service module (110) comprises:

a result analysis module (598) adapted to:

obtain results of each of the plurality of distinct task parts processed by each of the multiple remote agent computers (100), respectively,

determine whether the security-related task has been completed based on the obtained results from each remote agent computer (100) and task parameters, which specify requirements for results of task processing, to determine whether each of the plurality of distinct task parts has been completed such that the common objective of analyzing the unknown file has been achieved, and

determine whether the unknown file includes malware upon determining that the security-related task is complete and based on the obtained results,

a task acceptance module (540) for each agent computer (100) adapted to compute a determination of suitability of the agent computer (100) to accept a delegation of the at least one distinct task part via the distributed computing service module (110), the determination including obtaining the computing capacity requirements determined by the distributed computing service module (110) for performing a respective distinct task part, determining computing capability of the agent computer (100) based on available resources that includes types of anti-virus software capable of analyzing the unknown file and the respective malware database, and rendering a decision of whether the computing capability of the respective agent computer (100) is sufficient to meet the computational requirements; and

a task execution module (550) for each agent computer (100) coupled with the at least one task acceptance module (540) and adapted to obtain a respective distinct task part from the distributed computing service module (110) in response to the determination of suitability of the respective agent computer (100), and to execute the delegated distinct task part."

The claims also comprise an independent method claim 8, which corresponds closely to an independent system claim 1, and a computer program claim 13, which refers to the preceding method claims 8-13.

VI. Oral proceedings were held on 6 December 2017, at the end of which the chairman announced the decision of the board.

Reasons for the Decision

The invention

1. The application relates to the distributed execution of "security related" operations (such as antivirus scanning) in a (grid or peer-to-peer (P2P)) network (see page 1, lines 11-13 and page 7, lines 13-15, all references to the application being to its version as originally filed). Although it was known from the prior art for end-users to carry out malware analysis locally for the benefit of the entire network, a more effective solution was desirable (see page 2, line 27, to page 3, line 5).

1.1 As a solution, the application describes a distributed computer system in which tasks are delegated "on behalf" or "for the benefit" of a "beneficiary computer" (see page 3, lines 9-13) to "agent computers" based on a determination of their "suitability" in view of their computing capacity or current availability (page 3, lines 13-19). The distribution itself is aided by a "distributed computing service" (page 3, lines 9-10).

1.2 Figure 5A depicts an agent computer (100) in communication with a distribution server (110). Each agent computer has a "task acceptance module" which determines whether it is "available" to accept the execution of tasks, e.g. in view of "user activity" and computing capacity (see page 19, lines 9-15 and 27-32). The decision of which task should be delegated to which agent may be taken by the distributed computing service based on information from each agent, or the agents may decide themselves whether or not they are able to execute a proposed task (see page 13, lines 17-24 and figure 7, in particular step 720).

1.3 The application also discloses that a task can be broken into parts that are to be processed in parallel on separate computers (page 13, lines 7-11, and page 21, line 3). For example, "different agent computers" may "apply different non-overlapping portions of their antivirus databases" or perform a security analysis on different parts of a network (see page 21, lines 4-8). It is also disclosed that different agents may apply different methods for the detection of malware to an unknown file (for instance signature analysis and sandboxing) or to use different versions of an anti-virus database (see page 21, lines 26-32, and page 22, lines 4-5). The results of the different engines need to be combined into one overall result. Depending on the circumstances, this may mean combining the results in a "report" or determining the first or best result (see page 21, lines 22-25, or page 22, lines 1-10).

The prior art

2. D5 provides antivirus scanning as an "in-cloud network service". Each computer runs a lightweight "host agent" which detects suspicious files and forwards them for antivirus analysis to the "cloud" (see page 1, right column, last paragraph, and figure 3). In the network, several "heterogeneous" analysis engines with "complementary detection capabilities", operating in parallel, scan the files and return their analysis reports (*loc. cit.*; see also page 2, left column, paragraph 2 and the six lines just below it; page 5, left column, section 3.3; and section 4, paragraph 1). The results from the individual detection engines are then combined to determine whether the file in question is safe (page 7, left column, section 4.2.2, and page 8, right column, paragraph 4). The primary example in D5 is that the different engines use different analysis techniques so as to increase the "detection coverage" of the overall system (see section 4.2, in particular section 4.2.1, paragraph 1, and page 10, right column, paragraph 2). This approach is referred to as "N-version protection" (see section 3.3), suggesting that it provides N versions of the same kind of protection.

3. D2 discusses grid computing as a "pattern" of software architecture. It discloses the idea of splitting a job into parts and parallelising them across a distributed computing system. For illustration purposes, it is disclosed that a geographical area, for which a weather forecast is to be computed, may be split into smaller areas to be processed separately, as far as it is

possible to do so (see page 1 and figure 1). The remainder of D2 is, however, generic and independent of that particular example. D2 discloses the use of a "resource monitoring service" which monitors the availability of resources at the individual computing nodes. Based on that information, it searches for a set of nodes that provide the resources required to execute a query (see page 10, penultimate paragraph, and page 11, paragraph 3 from the bottom).

Added subject-matter, Article 123(2) EPC

4. Claims 1 and 8 refer to "divid[ing] the [...] task into [parts] for performing non-overlapping types of antivirus analysis using respective malware databases".

4.1 The board notes that the term "type" of antivirus analysis is not literally defined in the description. In the board's view, however, the skilled person would understand it to refer to the "different methods" of antivirus analysis as illustrated on page 21, last paragraph. For instance, signature analysis and sandboxing would be understood as two different "types" of analysis.

4.2 On page 21, lines 2-8, it is disclosed that a database may be split into "non-overlapping portions" for processing by different remote agents. Analogously, the entirety of a network to be analysed may be split into "parts", which the skilled person would understand to be "non-overlapping" as well.

4.3 However, if signature analysis was carried out on a given file vis-à-vis different "portions" of a signature database, the skilled person would not, in the board's view, talk about different "types" of analysis, because the signature analysis would be the same (and thus of the same "type") for each portion. Likewise, the board takes the view that qualifying two "types" of analysis as "non-overlapping" has no established meaning in the art and no clear meaning beyond marking them as "different". What, for instance, would it mean for sandboxing and virus scanning to be not only "different" but also "non-overlapping"?

4.4 The board therefore concludes that the term "non-overlapping" and the term "type" in the recited phrase refer to different, incompatible embodiments. In specifying them in combination, claims 1 and 8 thus go beyond the disclosure of the application as originally filed, in conflict with Article 123(2) EPC.

5. In response to this objection, the appellant indicated its willingness to limit the claimed invention to the embodiment disclosed on page 21, paragraph 2, according to which a signature database was split into several "non-overlapping portions" so that signature analysis vis-à-vis each of these portions could be distributed to a different remote agent computer. The corresponding amendment would imply, in particular, the deletion of any reference to "types of anti-virus analysis" or "software" from the claims (see, in particular claim 1, page 1, line 9, and page 2, line 14).

5.1 An amended set of claims was not actually filed. The board has no doubt, however, that the formulation of such claims would have been straightforward.

5.2 The appellant stressed its interest in obtaining the board's view as to the extent to which the claims limited to this embodiment involve an inventive step, rather than only in relation to the grounds perceived to be merely "formal" such as Article 84 and 123(2) EPC. The board thus continued the discussion of inventive step, interpreting the claims in the light of this embodiment.

Inventive step, Article 56 EPC

6. The board agrees with the decision and the appellant that D5 constitutes a suitable starting point for the assessment of inventive step.

6.1 D5 discloses (see in particular figure 3) a distributed computer system in which the "security-related task" of analysing a "suspicious" file for malware is delegated by a "distributed computing service module" (see in particular the "network service" discussed in section 5.2 and depicted in figure 3) to various "remote agent computers" (see the "analysis engines" of figure 3 and, equivalently, the "detection engines" in section 4.2.1).

6.2 The results from the detection engines are eventually combined to come to a final decision as to whether the file is considered safe or not (see section 4.2.2). This is performed by a "result analysis module" called an "aggregator", in view of the "common objective of analysing the unknown file", such as the security policy in place (*loc. cit.*). The aggregator is located in the network service (see section 4.2 and figure 3) and can thus be considered a component of the "distributed computing service module" as claimed.

7. Accordingly, the subject-matter of claim 1 differs from D5 in how the task is divided into portions and in that

(and how) the "suitability" of the remote agent computer is assessed before distribution.

7.1 More specifically, D5 does not disclose that

(a) a signature database is divided into non-overlapping parts, each of which is delegated to a different remote agent computer,
(b) each agent computer determines its "suitability" based on the "computing capacity requirements [...] for performing a [...] task part" and its own "computing capability".

In essence, these differences correspond to the distinguishing features I and II identified in the decision under appeal (see point 4.2 of the reasons).

7.2 D5 discloses that at least some of the deciding agents may carry out an antivirus analysis based on a signature database (see e.g. section 3.2, paragraph 3, and section 6.2, paragraph 2). D5 also discloses that the result of the detection engines may reach the aggregator at different times (see section 4.2.2, paragraph 2). Beyond that, the board deems it to be obvious that different antivirus "products" (see figure 2) will have different run-time behaviour and that, therefore, one may produce its result considerably faster than another.

7.3 Feature a) has the effect of speeding up the processing of any antivirus product in D5 which happens to be based on processing a signature database. This has the potential of speeding up the entire system of D5, for instance if this product is computationally more intensive (i.e. takes longer to complete) than the other ones used.

7.4 The board considers that the problem of speeding up an antivirus analyser using a signature database is one that can reasonably be assumed to arise in the context of D5.

7.5 Moreover, it would have been obvious for the skilled person to consider parallelising the execution of just that analyser as a solution to the given problem. The board takes the view that this would be the case in general, but even more so because D5 already discloses a parallel computation scenario.

7.6 Parallel computing rests on the idea that large tasks can be split into smaller parts for simultaneous execution by several computing agents. In other words, any parallelisation of a given task requires that the task be split into parts.

7.7 Furthermore, given that an virus scanner has to (independently) process a large number of virus signatures in the same way, the board considers that an obvious way to split this task into parts would to be to split the signature database into "non-overlapping parts" and to carry out the same analysis on each of these parts.

7.8 Feature b) addresses the problem that not every computing agent may be capable of processing a given task or sub-task at any point in time, for instance when it happens to be busy processing something else. This is an aspect of task scheduling, which, in the board's view, is also fundamental in parallel computing, and it involves, by necessity, a comparison of what has to be done (i.e. the requirements) with what a computing agent can do (i.e. its capacity). Again, D2 discloses a "resource monitoring service" determining "which nodes have available resources to execute the application", based on up-to-date information from the "resource providers", and providing this information to the scheduler (see page 3, section "Solution", paragraph 3, and page 4, last paragraph, to page 5, paragraph).

7.9 In the board's judgment, where exactly the comparison is carried out, i.e at the computing agent as claimed or at a separate scheduling component (as known, for instance, from D2), is a marginal issue and, more specifically, moving that comparison into an "acceptance module" of the computing agent is an obvious way of taking load from the scheduler.

7.10 The board notes *obiter* that the disclosure of D2 is not in conflict with that of D5 as the appellant argued (see the grounds of appeal, page 6, last paragraph). Notably, the relevant combination of D5 and D2 does not replace the "N-version protection" as a whole by a grid architecture like that of D2, but uses the "grid" of D2 to speed up the implementation of an individual decision engine within the N-version architecture of D5.

8. In summary, the board comes to the conclusion that the claimed invention is an obvious solution of speeding up the system of D5 in view of commonly known fundamental principles of parallel programming - as they are known, for instance, from D2 - and thus lacks inventive step, Article 56 EPC.

Order

For these reasons it is decided that:

The appeal is dismissed.

The Registrar: The Chairman:

Beschwerdekammern des Europäischen Patentamts
Boards of Appeal of the European Patent Office
Chambres de recours de l'Office européen des brevets

B. Atienza Vivancos W. Sekretaruk

Decision electronically authenticated