

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
5 April 2012 (05.04.2012)


(10) International Publication Number  
**WO 2012/041596 A1**

- (51) **International Patent Classification:**  
*G06F 17/30* (2006.01)
- (21) **International Application Number:**  
PCT/EP2011/064187
- (22) **International Filing Date:**  
17 August 2011 (17.08.2011)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**  
10182332.6 29 September 2010 (29.09.2010) EP
- (71) **Applicant (for all designated States except US):** INTERNATIONAL BUSINESS MACHINES CORPORATION [US/US]; New Orchard Road, Armonk, New York 10504 (US).
- (72) **Inventor; and**
- (75) **Inventor/Applicant (for US only):** DUNN, Tim [GB/GB]; IBM United Kingdom Limited, M/P 211, Hursley Park, Hursley, Winchester Hampshire SO21 2JN (GB).
- (74) **Agent:** LITHERLAND, David, Peter; IBM United Kingdom Limited, Intellectual Property Law, Hursley Park, Winchester Hampshire SO21 2JN (GB).
- (81) **Designated States (unless otherwise indicated, for every kind of national protection available):** AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, QA, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States (unless otherwise indicated, for every kind of regional protection available):** ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG,

[Continued on next page]

(54) **Title:** ADAPTIVE CONTENT-BASED PUBLISH/SUBSCRIBE MESSAGING

FIG. 3



(57) **Abstract:** Method, system, and computer program product publish/subscribe messaging are provided. One aspect of the method comprises: initiating a subscription comprising designating a subscription queue to which published messages relating to the subscription are to be put; defining multiple content-based filtering expressions for the subscription, each of the filtering expressions being defined for a given utilisation of the subscription queue. Another aspect of the method comprises: receiving a published message relating to a subscription; checking the utilisation of a subscription queue designated in the subscription; determining a content-based filtering expression defined for the given utilisation of the subscription queue; and applying the content-based filtering expression.

ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG). **Published:** — *with international search report (Art. 21(3))*

## ADAPTIVE CONTENT-BASED PUBLISH/SUBSCRIBE MESSAGING

### Technical Field

5 This invention relates to the field of publish/subscribe messaging. In particular, the invention relates to adaptive content-based publish/subscribe messaging.

### Background Information

10 The publish subscribe processing model is one in which producers of information (publishers) and consumers of that same information (subscribers) are not directly connected to each other. The passage of information between the two parties is handled by a publication engine or broker. The publication engine or broker has the responsibility of matching interested parties with the information being published. As not all subscribers  
15 necessarily want to receive all published messages, the techniques of topic and content-based routing are used to control what information a subscriber receives.

Topic based routing uses a simple classification mechanism which is implemented as a hierarchal tree. A subscriber can subscribe to receive messages at a particular point in that  
20 tree. It could be the root and so all messages are received, or it could be a leaf node in which case only messages published to the leaf node will be received.

Content-based routing provides a mechanism whereby a subscriber can provide a filter expression that will be applied to all messages that match the named point in the topic tree.  
25 This filter is typically a Structured Query Language (SQL) like expression that will reference information in the header or body of the message. If the filter condition evaluates to true the published message is written to the subscribers queue. If it evaluates to false it is not written to the subscribers queue.

30 Content-based filtering is useful in that it provides an additional filtering mechanism. In current implementations only one filter expression is supplied although this could be a complex expression. Although the application designer or business user can think carefully

about the construction of the filter statement for the subscriber, this one filter is the only mechanism available to control the amount of information received since they cannot influence the rate at which published messages are made.

5 When a subscriber's nominated queue becomes full of published messages, no more messages can be received until the queue size is increased, the subscribing application consumes messages off the queue, or an operator or automated command clears some or all messages off the queue.

10 Therefore, there is a need in the art to address the aforementioned problems.

### SUMMARY

15 According to a first aspect of the present invention there is provided a method for publish/subscribe messaging, comprising: initiating a subscription including designating a subscription queue to which published messages relating to the subscription are to be put; defining multiple content-based filtering expressions for the subscription, each of the filtering expressions being defined for a given utilisation of the subscription queue.

20 According to a second aspect of the present invention there is provided a method for publish/subscribe messaging, comprising: receiving a published message relating to a subscription; checking the utilisation of a subscription queue designated in the subscription; determining a content-based filtering expression defined for the given utilisation of the subscription queue; and applying the content-based filtering expression.

25 According to a third aspect of the present invention there is provided a computer software product for publish/subscribe messaging, the product comprising a computer-readable storage medium having computer readable program code embodied therewith, the computer readable program code configured to: initiate a subscription including designating a  
30 subscription queue to which published messages relating to the subscription are to be put; and define multiple content-based filtering expressions for the subscription, each of the filtering expressions is defined for a given utilisation of the subscription queue.

According to a fourth aspect of the present invention there is provided a computer software product for publish/subscribe messaging, the product comprising a computer-readable storage medium having computer readable program code embodied therewith, the computer readable program code configured to: receive a published message relating to a subscription;  
5 check the utilisation of a subscription queue designated in the subscription; determine a content-based filtering expression defined for the given utilisation of the subscription queue; and apply the content-based filtering expression.

According to a fifth aspect of the present invention there is provided a system for  
10 publish/subscribe messaging, comprising: a publication engine including a subscription mechanism, the subscription mechanism including: a subscription queue designation component for designating a subscription queue to which published messages relating to the subscription are to be put; a content-based filter settings component for defining multiple content-based filtering expressions for a subscription, each of the filtering expressions being  
15 defined for a given utilisation of the designated subscription queue.

According to a sixth aspect of the present invention there is provided a system for  
publish/subscribe messaging, comprising: a publication engine including: a published  
20 message receiving component for receiving published messages relating to a subscription; a queue utilisation checking component for checking the utilisation of a subscription queue designated in the subscription; a filter determining component for determining a content-based filtering expression defined for the given utilisation of the subscription queue; and a filter applying component for applying the content-based filtering expression.

25 Viewed from a further aspect, the present invention provides a computer program stored on a computer readable medium and loadable into the internal memory of a digital computer, comprising software code portions, when said program is run on a computer, for performing the steps of the invention.

30 The subject matter regarded as the invention is particularly pointed out and distinctly claimed in the concluding portion of the specification. The invention, both as to organization and method of operation, together with objects, features, and advantages

thereof, may best be understood by reference to the following detailed description when read with the accompanying drawings in which:

#### Brief Description of the Drawings

5

The present invention will now be described, by way of example only, with reference to preferred embodiments, as illustrated in the following figures:

Figure 1 is a block diagram of a system in accordance with the present invention;

10

Figure 2 is a block diagram of a computer system in which the present invention may be implemented;

15

Figure 3 is a flow diagram of a subscriber registration in accordance with an aspect of the present invention; and

Figure 4 is a flow diagram of a message filtering in accordance with an aspect of the present invention.

20

#### Detailed Description of the Invention

25

It will be appreciated that for simplicity and clarity of illustration, elements shown in the figures have not necessarily been drawn to scale. For example, the dimensions of some of the elements may be exaggerated relative to other elements for clarity. Further, where considered appropriate, reference numbers may be repeated among the figures to indicate corresponding or analogous features.

30

In the following detailed description, numerous specific details are set forth in order to provide a thorough understanding of the invention. However, it will be understood by those skilled in the art that the present invention may be practiced without these specific details.

In other instances, well-known methods, procedures, and components have not been described in detail so as not to obscure the present invention.

Conventional topic and content-based publish/subscribe processing does not take account of the amount of data already on a queue and perhaps more importantly how close to the capacity of the queue the current queue depth is. When the subscriber's queue fills, typically what will happen is that the publication engine will rollback the published message if the message is persistent and transacted, or the message will be lost if it is non-persistent. Therefore, published messages may be received and processed on an unpredictable basis which is not desirable. Conventionally, the subscribing application has no ability to change the filter expression as the queue limit begins to be approached.

A user, for example, a business user or application designer, who is the subscriber may be happy for the content filter to become more restrictive if they know the queue is reaching capacity, rather than continuing to lose all published messages from the point at which the queue is full.

It would be valuable for the user to have more influence over what is kept and what is discarded as the subscription queue fills, rather than have a condition imposed by the transactional behaviour of the publication engine.

The described system and method provide a mechanism to allow the user to supply a series of filters for content-based publish/subscribe to take account of higher queue utilisations. The execution of a particular filter is dependent on the utilisation of the individual subscription queue. This utilisation is determined by the publication engine and the appropriate filter executed at the time of published message. Subscription queue utilization may include: remaining queue capacity, queue rate of use, amount of data currently on queue, or any other method of measuring the queue usage.

By being more restrictive in the filter expression, it is possible for a user to specify that fewer or even no new messages be accepted once a certain queue utilisation had been reached. This gives the power of whether to accept more published messages to the user rather than having it imposed on them because the subscription queue has filled.

The advantage is that the subscribing application will in effect still continue to see new published messages for much longer as the queue full condition will have been averted for much longer by reducing the rate at which new published messages are accepted onto the subscription queue.

5

In the described method and system, when providing a filter for content-based publish/subscribe, the subscriber specifies a set of filter expressions and queue utilisations (or other queue threshold/characteristics) at which to apply each filter expression.

10

The subscriber provides the filters but does not implement them. The publication engine is responsible for applying the different filters. Prior to invoking a filter for the content-based published message, the publication engine checks the queue utilisation of the subscription queue and then applies the appropriate filter.

15

Only the subscriber user understands the business context of the information being consumed and so is the only one capable of making a judgement about how restrictive he is prepared to be in order to try and avoid the queue full condition whilst seeing as many new published messages as possible.

20

Referring to Figure 1, a block diagram shows an example system 100 for adaptive content-based publish/subscribe messaging.

25

A publication engine 110 is provided as an intermediary between publishers 101 and subscribers 102. A subscriber 102 registers with a publication engine 110 which provides routing and filtering between publishers 101 and subscribers 102. In some cases, there may be a publication engine that performs publish/subscribe processing only. In another case, a publish/subscribe function may be provided as one of a number of functions in a multi-function broker. The term publication engine is intended to include any publish/subscribe processing mechanism including a message broker.

30

A publication engine 110 includes a subscriber mechanism 111 at which a subscriber 102 registers a subscription. The subscriber mechanism 111 includes subscriber settings 120



including the name 121 of the subscription queue 103 to receive published messages and the subscriber criteria 122. The subscriber criteria 122 includes the topic or topics of the published messages the subscriber is interested in.

5 The subscription queue 103 is the destination to which the publication engine 110 sends messages matching a particular subscription. When a subscriber 102 registers to a topic, it indicates to the publication engine 110 to which queue it wants published messages to be forwarded. One queue may be used by several registrations. This queue is specified in one of the parameters of the register subscriber command.

10

In addition, in the described system 100 the subscriber settings 120 of the subscriber mechanism 111 includes content-based filter settings 123 including multiple filter expressions 124, 125 to be applied at different subscription queue utilisations 126, 127.

15 The publication engine 110 includes a published message receiving mechanism 130 for receiving published messages from publishers 101 and determining subscriptions to the published messages. The published message receiving mechanism 130 includes a publication queue 131 for receiving published messages at the publication engine 110.

20 The subscriber mechanism 111 includes a queue component 150 for holding messages to be put to the subscription queue 103. In an example implementation, an area of disk is set aside to hold messages for a particular queue. Messages are forwarded to the subscription queue 103 when a communication program and link are available. The queue component 150 tracks the current queue depth for a subscription queue 103 as it sees all write requests and  
25 all retrieve requests against that queue including the number of bytes in the messages. A queue utilisation checking component 153 calculates the subscription queue 103 utilisation before putting a message to the subscription queue 103.

30 In one embodiment, the publication engine 110, publisher 101 and subscriber 102 may optionally include queue managers for managing their local queues.

A content-based filter mechanism 140 is provided including a filter expression determining component 141 for determining which of the filter expressions 124, 125 to apply based on the subscription queue 103 utilisation determined by the queue utilisation checking component 153 and a filter applying component 142.

5

The filter mechanism 140 receives a published message from a publisher 101, identifies a subscription in the subscriber mechanism 111 for which the message is valid. The filter mechanism 140 then obtains the name of the subscription queue 103 and the queue utilisation checking component 153 checks the utilisation of the queue before the filter applying component 142 applies the appropriate content-based filter setting for the queue utilisation as specified in the content-based filter settings 123.

10

Referring to Figure 2, an exemplary system for implementing aspects of the invention may include a data processing system 200 suitable for storing and/or executing program code including at least one processor 201 coupled directly or indirectly to memory elements through a bus system 203. The memory elements can include local memory employed during actual execution of the program code, bulk storage, and cache memories which provide temporary storage of at least some program code in order to reduce the number of times code must be retrieved from bulk storage during execution.

15

20

The memory elements may include system memory 202 in the form of read only memory (ROM) 204 and random access memory (RAM) 205. A basic input/output system (BIOS) 206 may be stored in ROM 204. System software 207 may be stored in RAM 205 including operating system software 208. Software applications 210 may also be stored in RAM 205.

25

The system 200 may also include a primary storage means 211 such as a magnetic hard disk drive and secondary storage means 212 such as a magnetic disc drive and an optical disc drive. The drives and their associated computer-readable media provide non-volatile storage of computer-executable instructions, data structures, program modules and other data for the system 200. Software applications may be stored on the primary and secondary storage means 211, 212 as well as the system memory 202.

30

The computing system 200 may operate in a networked environment using logical connections to one or more remote computers via a network adapter 216.

5 Input/output devices 213 can be coupled to the system either directly or through intervening I/O controllers. A user may enter commands and information into the system 200 through input devices such as a keyboard, pointing device, or other input devices (for example, microphone, joy stick, game pad, satellite dish, scanner, or the like). Output devices may include speakers, printers, etc. A display device 214 is also connected to system bus 203 via an interface, such as video adapter 215.

10 Implementations of a data processing system may not have all the above components. There are small or micro brokers running on limited capacity devices such as phones or pipeline monitoring devices which could use the described system and which may have a limited set of the described functions.

15 Referring to Figure 3, a flow diagram 300 shows an embodiment of the registration of a subscriber in the described method and system. A new subscription is initiated 301. The subscription specifies 302 a name of a subscription queue for receiving the messages. The subscription specifies 303 a topic. A content filter is applied 304 on top of the topic based  
20 subscription and the content filter specifies 305 multiple filter expressions with associated queue utilisations.

25 Referring to Figure 4, a flow diagram 400 shows an embodiment of message filtering at a publication mechanism in the described method and system. A published message is received 401 and a subscription is identified 402. The subscription may be for specified topics.

30 The subscription queue designated in the subscription is checked 403 for its utilisation or capacity. Given a certain queue utilisation, an appropriate content filter is determined 404 from predefined filter expressions.

It is determined 405 from the content filter if the message is to be delivered to the subscriber.

If the message is not to be delivered, the process loops 406 to await the next message. If the message is to be delivered, the message is placed 407 on the subscription queue and the process loops 406 to await the next message.

5 Allowable languages for expressing the logic needed by the publication engine for content-based filtering vary depending on the implementation of the publication engine. In one example, message selectors may allow the use of embedded SQL (ESQL) expressions to filter on an entire message.

10 Content filtering requires a method of pointing to portions of the header and body of a message that is to be published. In one implementation, field references are used in filters for content-based filtering. To reference a field in a filter, a path must be specified. Each element of the path consists of a, possibly indexed, field name.

15 The following are example field name and correlation name identifiers as an example implementation. These identifiers represent all messages as a hierarchical syntax element tree. Each path identifies a route through that tree, which leads to a particular syntax element, starting from one of the predefined correlation names that refer to fixed points that every message has. The following correlation names shown in Table 1 are supported for  
 20 content-based filtering.

Table 1

Root	Identifies the root of a published message.
Properties	Identifies the portion of the message in which the standard properties of a message lie.
Body	Identifies the last child of the root of the message, which is usually, but not always, the application data that follows any headers.

Some examples of field references, together with their meanings, are shown in Table 2.

Table 2

Body.Person.Address	Refers to the first Address field in the Person entity in the body of the message.
Properties.Topic	Refers to the "Topic" field in the standard properties of a message.
Root.MQMD.UserIdentifier	Refers to the UserIdentifier field in the MQMD of the message.

Some examples of content-based filtering are shown in Table 3.

Table 3

Body.Person.Salary>10000	Filtering against an integer literal
"Body.Person.Address"[1]NOT LIKE'Blen%'AND"Body.Person.Salary">15000	A more complex filter. Note that field identifiers can optionally be surrounded by double quotation marks.
Body.Date1='2000-02-14'	Filtering against a date. The date is matched as a string and care must be taken with its layout.
Body.Person.ApprovalFlag	Filtering against a Boolean field.
Body.Person.Salary+Person.Bonus>Body.Person.Limit	An arithmetic filter.
Properties.Topic='employees/marketing'	Filtering on a message property.
Root.MQMD.UserIdentifier='Jones'	Filtering on a message attribute.
Body.Person.HourlyRate = 10.24	Filtering against a float literal
Body.Planet.DistanceFromSun = 0.93E8	Filtering against a float literal in exponential format

In the described method and system, the content-based filter includes multiple filter expressions which are applied depending on the utilisation of the subscription queue. The utilisation of the subscription queue may be defined, for example, in terms of percentage capacity used.

5

For example, if a filter integer range is provided, the range may become more restricted as the queue utilisation increases. In this way, more selective filtering is carried out as the subscription queue utilisation increases and has reduced capacity.

10

### Examples

In a first example, a system looks at the spending habits of shoppers. The subscriber wants to capture details of all shoppers spending more than \$100 at a time for further analysis of their spending habits. As the subscriber queue depth increases as utilisation grows, the subscription filter adapts to only capture the very high spending shoppers. A business decision was made when the filter was constructed to loose the lower spending shoppers under such circumstances. However, the subscriber certainly does not want to risk loosing the very high spending shopper and this is reflected in the construction of the filter.

15

20

If Message.Body.ReceiptTotal > \$100 as the content filter default expression;  
If Message.Body.ReceiptTotal > "2 \* \$100" if the queue utilisation hits 70%;  
If Message.Body.ReceiptTotal > "3 \* \$100" if the queue utilisation hits 90%;  
If Message.Body.ReceiptTotal > "30 \* \$100" if the queue utilisation hits 95%.

25

A second example is a system which wants to capture stocks whose price varies significantly one from one day to the next so that analysis of these stocks can be made, possibly with a view to trading them in the future.

30

Initially the subscriber is happy to capture those with a variation on +/- 5%. However, as the queue depth decreases due to utilisation increases, the subscriber is prepared to only capture only the most volatile. Again a business decision was made when constructing the filter that meant it was acceptable to discard more data in order to capture only the most volatile.

If Message.Body.PriceSwingToday > 5% as the content filter default expression;  
If Message.Body.PriceSwingToday > "2 \* 5%" if the queue utilisation hits 70%;  
If Message.Body.PriceSwingToday > "3 \* 5%" if the queue utilisation hits 90%;  
If Message.Body.PriceSwingToday > "10 \* 5%" if the queue utilisation hits 95%.

5

In these examples as the queue utilisation has grown the filter has only accepted larger and larger values. This is not always going to be the case. In other situations it may be appropriate to accept only smaller and smaller values as the queue utilisation grows.

10

In these examples, a threshold value is set of a parameter (in the first example, the spending value, and in the second example, the variation value) and the filter expressions are varied according to a multiple of the threshold value. This is one example of how the filter expressions may be varied in relation to queue utilisation. Other method may also be used.

15

An adaptive content-based publish/subscribe system may be provided as a service to a customer over a network.

20

The invention can take the form of an entirely hardware embodiment, an entirely software embodiment or an embodiment containing both hardware and software elements. In a preferred embodiment, the invention is implemented in software, which includes but is not limited to firmware, resident software, microcode, etc.

25

The invention can take the form of a computer program product accessible from a computer-usable or computer-readable medium providing program code for use by or in connection with a computer or any instruction execution system. For the purposes of this description, a computer usable or computer readable medium can be any apparatus that can contain, store, communicate, propagate, or transport the program for use by or in connection with the instruction execution system, apparatus or device.

30

The medium can be an electronic, magnetic, optical, electromagnetic, infrared, or semiconductor system (or apparatus or device) or a propagation medium. Examples of a computer-readable medium include a semiconductor or solid state memory, magnetic tape, a

removable computer diskette, a random access memory (RAM), a read only memory (ROM), a rigid magnetic disk and an optical disk. Current examples of optical disks include compact disk read only memory (CD-ROM), compact disk read/write (CD-R/W), and DVD.

- 5 Improvements and modifications can be made to the foregoing without departing from the scope of the present invention.



**CLAIMS**

1. A method for publish/subscribe messaging, comprising the steps:  
initiating (301) a subscription comprising designating a subscription queue to which  
5 published messages relating to the subscription are to be put; and  
defining (305) multiple content-based filtering expressions for the subscription, each  
of the filtering expressions being defined for a given utilisation of the subscription queue.
2. The method as claimed in claim 1, wherein the multiple filtering expressions are  
10 increasingly restrictive filters defined for increasing utilisation of the subscription queue.
3. The method as claimed in either of claims 1 or 2, further comprising a step of:  
defining subscription queue utilisation parameters for which a filtering expression is  
15 defined.
4. The method as claimed in claim 3, wherein the subscription queue utilisation  
parameters comprise one or more of the group of: remaining queue capacity, queue rate of  
use, amount of data currently on queue.
- 20 5. A method for publish/subscribe messaging, comprising the steps of:  
receiving (401) a published message relating to a subscription;  
checking (403) the utilisation of a subscription queue designated in the subscription;  
determining (404) a content-based filtering expression defined for the given  
utilisation of the subscription queue; and  
25 applying (405) the content-based filtering expression.
6. A computer program product for publish/subscribe messaging, the computer program  
product comprising a computer readable storage medium readable by a processing circuit  
and storing instructions for execution by the processing circuit for performing a method  
30 according to any one of claims 1 to 5.
7. A system for publish/subscribe messaging, comprising:

a publication engine (110) comprising a subscription mechanism (111), the subscription mechanism (111) comprising:

a subscription queue designation component (121) for designating a subscription queue (103) to which published messages relating to the subscription are to be put;

5 a content-based filter settings component (123) for defining multiple content-based filtering expressions (124, 125) for a subscription, each of the filtering expressions (124, 125) being defined for a given utilisation (126, 127) of the designated subscription queue (103).

10 8. The system as claimed in claim 7, wherein the multiple filtering expressions (124, 125) are increasingly restrictive filters defined for increasing utilisation of the subscription queue (103).

15 9. The system as claimed in either of claims 7 or 8, wherein the content-based filter settings component (123) is further operable for defining subscription queue utilisation parameters for which a filtering expression (124, 125) is defined.

20 10. The system as claimed in claim 9, wherein the subscription queue utilisation parameters comprise one or more of the group of: remaining queue capacity, queue rate of use, amount of data currently on queue.

11. A system for publish/subscribe messaging, comprising:

a publication engine (110) comprising:

25 a published message receiving component (130) for receiving published messages relating to a subscription;

a queue utilisation checking component (153) for checking the utilisation of a subscription queue (103) designated in the subscription;

30 a filter determining component (141) for determining a content-based filtering expression (124, 125) defined for the given utilisation (126, 127) of the subscription queue (103); and

a filter applying component (142) for applying the content-based filtering expression (124, 125).

12. The system as claimed in claim 11, wherein the queue utilisation checking component (153) is operable for tracking the usage of the subscription queue (103) by seeing all write requests and all retrieve requests against the subscription queue (103).

5 13. An apparatus for publish/subscribe messaging, comprising:  
means for initiating (301) a subscription comprising means for designating a  
subscription queue to which published messages relating to the subscription are to be put;  
and  
means for defining (305) multiple content-based filtering expressions for the  
10 subscription, each of the filtering expressions being definable for a given utilisation of the  
subscription queue.

14. The apparatus as claimed in claim 15, wherein the multiple filtering expressions are increasingly restrictive filters defined for increasing utilisation of the subscription queue.

15

15. The apparatus as claimed in either of claims 13 or 14, further comprising:  
means for defining subscription queue utilisation parameters for which a filtering  
expression is definable.

20

16. The apparatus as claimed in claim 15, wherein the subscription queue utilisation parameters comprise one or more of the group of: remaining queue capacity, queue rate of use, amount of data currently on queue.

25

17. An apparatus for publish/subscribe messaging, comprising:  
means for receiving (401) a published message relating to a subscription;  
means for checking (403) the utilisation of a subscription queue designated in the  
subscription;  
means for determining (404) a content-based filtering expression definable for the  
given utilisation of the subscription queue; and

30

means for applying (405) the content-based filtering expression.

18. A computer program stored on a computer readable medium and loadable into the internal memory of a digital computer, comprising software code portions, when said program is run on a computer, for performing the method of any of claims 1 to 5.

FIG. 1




FIG. 2




FIG. 3





FIG. 4





**INTERNATIONAL SEARCH REPORT**

International application No  
PCT/EP2011/064187

**A. CLASSIFICATION OF SUBJECT MATTER**  
INV. G06F17/30  
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

**B. FIELDS SEARCHED**

Minimum documentation searched (classification system followed by classification symbols)  
G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)  
EPO-Internal, WPI Data

**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2008/065604 A1 (TIU WILLIAM K [US] ET AL) 13 March 2008 (2008-03-13) paragraph [0034] paragraph [0033] paragraph [0035]	1-18
A	US 2010/198920 A1 (WONG LIK [US] ET AL) 5 August 2010 (2010-08-05) paragraph [0034] paragraph [0042] paragraph [0066]	1-18

Further documents are listed in the continuation of Box C.

See patent family annex.

\* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
- "E" earlier document but published on or after the international filing date
- "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- "O" document referring to an oral disclosure, use, exhibition or other means
- "P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- "&" document member of the same patent family

Date of the actual completion of the international search  28 September 2011	Date of mailing of the international search report  06/10/2011
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016	Authorized officer  Haffner, Ronald

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2011/064187

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2008065604 A1	13-03-2008	US 2008065701 A1	13-03-2008
US 2010198920 A1	05-08-2010	NONE	