



- (51) International Patent Classification:  
H04L 29/06 (2006.01)
- (21) International Application Number:  
PCT/GB2013/000564
- (22) International Filing Date:  
23 December 2013 (23.12.2013)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:  
12250186.9 21 December 2012 (21.12.2012) EP
- (71) Applicant: BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY [GB/GB]; 81, Newgate Street, London EC1A 7AJ (GB).
- (72) Inventor: KARALE, Sachin; PP: C5A, 81 Newgate Street, London, EC1A 7AJ (GB).
- (74) Agent: GEFFEN, Nigel, Paul; BT Legal, Intellectual Property Department, PP: C5A, BT Centre, 81 Newgate Street, London EC1A 7AJ (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, BN, 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, IR, IS, JP, KE, KG, 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, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, 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, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, 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, CL, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: NETWORK SECURITY MANAGEMENT

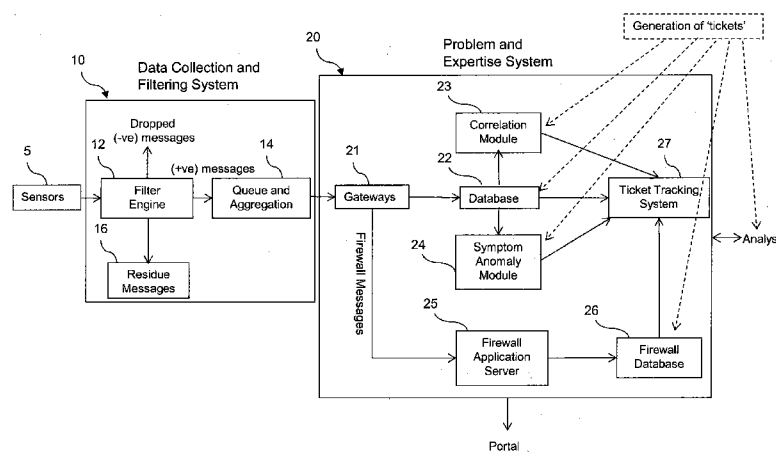


Figure 1 - Message Flow in a Network Security Monitoring Process

(57) Abstract: Methods and apparatus are disclosed for processing status messages for use in network security management in respect of a network of computing devices, the status messages comprising data relating to a plurality of attributes; the method comprising: filtering received status messages according to filtering rules in order to identify (i) status messages indicative of potential or actual network security events in respect of which a predetermined response is deemed applicable; (ii) status messages in respect of which a null response is deemed applicable; and (iii) residue messages not identified as (i) or (ii); processing messages identified as (i) such that a predetermined response may be initiated; performing analysis involving clustering in respect of messages identified as residue messages; and updating the filtering rules for use in subsequent filtering of received status messages in dependence on the result of the analysis.

WO 2014/096761 A1

## **NETWORK SECURITY MANAGEMENT**

### **Technical Field**

The present invention relates generally to network security management and, more specifically, to methods and systems for processing status messages for use in network  
5 security management in respect of a network of computing devices. Aspects of the invention relate also to dynamic network intrusion monitoring, detection and response.

### **Background to the Invention and Prior Art**

In general, computer and network security products focus on the prevention of various possible types of attack on the security of the network and the computing devices within it.  
10 Firewalls are generally used to prevent unauthorised traffic from entering the internal, local, or otherwise bounded network or networks of a company, a particular group of users, an individual user, or another such entity; authentication mechanisms may be used to prevent unauthorised persons from logging on to a company's (or other entity's) computers; and encryption may be used to prevent unauthorised persons from reading files on, or sent to or  
15 from those computers. Such products cannot be relied upon to work perfectly, however, and because security 'bugs' may exist in other software or hardware, complete network security also requires monitoring and detection, and appropriate response in the event of a breach.

An effective monitoring, detection and response system for a network may be provided internally (e.g. by a system administrator of a company network, or by a user of an individual  
20 network). It may include firewalls, secure servers and routers, dedicated intrusion detection systems, and other security products, all of which may provide audit and other information about their status, about possible security/related issues, and about other characteristics or events. These may be assessed by internal system administrators or in an automated or semi-automated manner. Whatever "internal" security provisions an entity uses, these may  
25 be augmented by making use of externally-provided network security services. While externally-provided network security services should not generally be used by a customer (e.g. a company, an individual, etc.) as a complete substitute for provisions used by themselves or their own system administrators, individuals and system administrators normally do not have the time or ability to read through large amounts of constantly updated  
30 audit information, looking for attacks on their systems. They also may not have the time to continuously monitor the activities of 'hackers' (or other such parties), looking out for new tactics, tools and trends in what they may be using. Also, they may not have the time to become experts on every kind of intrusion or attack and to maintain that expertise.

A monitoring, detection and response system that employs human intelligence, uses trained personnel, and takes advantage of network security intelligence and other knowledge databases can provide network users and system administrators with the advice and coaching they may need, when they need it, to help them repel or otherwise respond to attacks and maintain network integrity and 'uptime'. While completely automatic defences may be used, and may work against some attacks (particularly automated attacks), they may be at a disadvantage against an intelligent attack, against which a specialist intelligent monitoring, detection and response capability may be needed. Such a specialist capability may be provided by a dedicated external service provision entity (although it will be noted that a corresponding service could also be provided by a similarly dedicated service provision entity within what could still be regarded as a customer's network). In any case, for such a dedicated service to be provided in respect of a network, the network (generally by way of one or more computing devices within it) generally needs to provide information about possible network security incidents or 'events' (such as possible intrusion of unauthorised traffic, possible attempts by unauthorised executable software that has successfully entered the network to take action against the interests of the network or its users, possible attempts by unauthorised to breach or otherwise overcome authentication mechanisms, etc.).

In order to allow such a capability to be provided (whether externally or internally), computing devices in the network and/or software applications running thereon may be arranged to send "status messages" to an entity providing such a capability. These may filtered or otherwise controlled such that only those that may relate to potential network security issues are sent, in which case these may be regarded as "network security messages". Such messages can also be provided by dedicated sensors or probes set up to monitor devices, exchanges of data, and other such interactions between devices in the network. Such messages may be provide information about the devices affected (e.g. their identity, their location in the network, software applications that they have been using, etc.) and/or software applications affected in a form that enables them to be processed manually (i.e. by a human analyst), automatically (i.e. by a computer processor), or by a combination of the two. Such messages generally provide information in the form of data fields of several defined types, in order to allow for efficient processing. This may be necessary because the computing devices in a large network, such as individual computers, firewalls, servers, routers, specific intrusion detection systems, etc. can generate millions of lines of security-related or other status or audit information each day. This may contain information indicative of ongoing network attacks or intrusions, but this may not get noticed in amongst a large amount of other less-critical or non-security-related audit information.

In view of this, it is known that Managed Security Monitoring (MSM) services may be provided and used, which can assist a customer by receiving "network security messages" and/or more general status messages from the customer's network (and from the devices of that network), and filter and analyse the information therein effectively in order to detect such attacks or intrusions, and if required, suggest or provide an appropriate response.

Once a possible attack or intrusion (i.e. a network security "incident" or "event") is detected, its characteristics and particulars may then be examined and analysed by trained security analysts continuously monitoring the customer's network to further understand the incident and eliminate false positives. In analysing the incident, security analysts can draw upon information and knowledge from a variety of sources, including but not limited to security intelligence databases containing information about the characteristics of various hacker techniques and tools and known vulnerabilities in various operating systems and commercial software products and hardware devices. If necessary, security analysts can escalate the handling of the incident according to a variety of possible escalation procedures to stop the attack and shut down the vulnerability before the attacker does any damage. In effect, an MSM service can act as an additional defensive shield for a customer's network.

United States patent US7895641 and corresponding International application WO 01/71499 ("Schneier"/"Counterpane") relate to a known technique for dynamic network intrusion monitoring, detection and response. According to this, a probe attached to a customer's network collects status data and other audit information from monitored components of the network, looking for footprints or evidence of unauthorised intrusions or attacks. The probe filters and analyses the collected data to identify potentially security-related events occurring in the network. Identified events are reported to a human analyst for problem resolution. The analyst has access to a variety of databases to aid in problem resolution, and may follow an escalation procedure in the event he or she is unable to resolve the problem. Various customer personnel can be alerted in a variety of ways depending on the nature of the problem and the status of its resolution. Feedback from problem resolution efforts can be used to update the knowledge base available to analysts for future attacks, and may be used to update the filtering and analysis capabilities of the probe and other systems.

Aspects of an existing product, which will be referred to as the "Counterpane product", for providing managed security monitoring (MSM) services essentially in accordance with the technique described in US7895641 will briefly be described, and will be referred to again later in more detail. The Counterpane product uses sensors to monitor customers' data networks for security threats. Data from the sensors in the form of messages are received by a module referred to in the Counterpane product as the "Sentry". This contains a filtering

subsystem which classifies the messages as "positive messages" (messages that need to be monitored, as they appear to relate to security issues that may require action to be taken), "negative messages" (messages that may be discarded, because they do not relate to security issues), and "residue messages" (messages that are not possible to parse).

5 Messages classed as "negative" are discarded.

Messages classed as "positive" are passed to a module referred to as "SOCRATES", an acronym which stands for "Secure Operations Center Responsive Analyst Technical Expertise System". This collects and formats messages into "problem tickets" (each of which represents a discrete security-related event or incident of possible intrusive activity on a customer's network), associates with each such ticket information useful for problem investigation, resolution and response, presents such tickets to security analysts for handling, and generally serves as a repository of useful information and procedures.

10 The filtering subsystem may be unable to parse the residue messages because it does not have enough filters, or does not have enough information to parse the messages. This may also occur if there has been a change or an error in the format of the security messages.

With the Counterpane product (and as described in US7895641), these residue messages are passed to a human analyst for checking and analysis. As this is done manually (i.e. by human analysts) and takes time and human efforts, there may be a significant delay before the residue messages have been analysed. Further, human error may affect the analysis, and/or some residue messages may even be left un-analysed because of lack of manpower.

Residue messages that have been manually analysed may thus get manually classified (correctly or incorrectly) as "positive messages" or "negative messages". Further to this, messages that have been classified (manually or automatically) as "positive messages" or "negative messages" may get manually analysed in a separate operation, and it is possible that this may result in a human decision to update the filter engine.

25 Techniques such as the above thus depend on significant manual input from human analysts both in analysing residue messages and in reviewing messages classed as positive.

A known technique in data analysis is "clustering". Clustering algorithms generally aim to divide a set of objects into groups (clusters), where objects in each cluster are similar to each other (and as dissimilar as possible to objects from other clusters). Objects that do not fit well to any of the clusters detected by the algorithm may be considered as "outliers", or to form a special cluster of outliers.

30 A paper by Risto Vaarandi entitled: "A Data Clustering Algorithm for Mining Patterns From Event Logs" (Proceedings of the 2003 IEEE Workshop on IP Operations and Management,

ISBN: 0-7803-8199-8) relates to a problem whereby event logs such as those used in system and network management contain vast amounts of data that can easily overwhelm a human. Identifying that mining patterns from event logs is an important system management task, the paper presents a clustering algorithm for log file data sets which can help in the detection of  
5 frequent patterns from log files, in the building of log file profiles, and in the identification of anomalous log file lines. The technique outlined considers standard "sys log" formats that contain free-form text, and suggests using clustering on the full set of event logs in order to try to form clusters that would cover "outliers". It will be understood that clustering is thus used in what can be regarded as a single processing stage on full sets of event logs with the  
10 aim of avoiding any of these from being left outside of the clusters.

A paper by Feng Xuewei, Wang Dongxia, Zeng Jiemei, Ma Guoqing, and Li Jin, entitled: "Analyzing and Correlating Security Events Using State Machine" (2010 10th IEEE International Conference on Computer and Information Technology (CIT 2010)) relates to attack scenario reconstruction, clustering analysis, causal analysis, use of an attack scenario  
15 tree, and use of a correlating state machine. Recognising that it may be unfeasible for a security manager to analyse security events manually, the paper propose use of an attack scenario reconstruction technology based on a state machine. The processes of attackers can be replicated and more comprehensive attack scenario description information can be generated. The paper appears to suggest using clustering in conjunction with existing  
20 knowledge in order to alert an analyst about possible attacks, based on types of attack that have happened previously.

Referring to other (cited) prior art, a paper entitled "Mining Alarm Clusters to Improve Alarm Handling Efficiency" by K. Julisch et al (<http://www.acsac.org/2001/papers/115.pdf>) discusses a problem wherein intrusion detection systems overload their human operators by  
25 triggering thousands of alarms per day, and presents the results of some research indicating that alarms should be managed by identifying and resolving their root causes. It discusses alarm clustering as a method that supports the discovery of root causes.

A paper entitled "An efficient algorithm for clustering intrusion alert" by J. D. Adelina et al (<http://www.iatit.org/volumes/Vol37No2/11Vol37No2.pdf>) discusses intrusion detection  
30 systems, noting the problems of relying on alert supervision. It proposes using "Meta alerts", which are generated for appropriate clusters and which form generalisations of alerts, as a way of identifying origins of alerts. A hybrid clustering algorithm is proposed which is applied to the data set. Redundant data are filtered in order to reduce the rate of false positives.

US patent application US2004/015719 ("Lee et al") relates to network security protection,  
35 and to integrated security systems in which individual security agents are actively inter-

related. In particular, it proposes a security system comprising a firewall for interconnecting and controlling access between external and internal networks, a plurality of security agents for monitoring a data flow and system calls over the internal network, an "intelligent" security engine (ISE) for analysing alert messages, traffic information and event information from the security agents to decide if there is an attack and to generate a signature through a learning process, and a security policy manager for managing and applying a security policy to each of the security agents based on a decision of the ISE.

US patent application US2005060562 ("Bhattacharya et al") relates to techniques for displaying network security incidents.

10 US patent US8176527 ("Njemanze et al") relates to a "correlation engine" or "rules engine" with support for time-based rules. The rules engine receives security events generated by network devices, which are aggregated and provided to the rules engine at specific times associated with time-based rules. The security events are cross-correlated with the time-based rules; and one or more first stage meta-events are reported.

15 The present inventors have identified a need to provide further assistance in relation to the analysis of messages such as security messages that are sent for processing by a managed security monitoring system, in particular in relation to those that do not match with an existing knowledge base. In relation to parts of the processing that are currently performed manually, the fact that such parts are time consuming when performed by human analysts may mean that they are done incorrectly, done too slowly to be of use in real-time or near real-time processing, or not done at all. While automation is often seen as a solution in relation to tasks that are time-consuming when performed by human analysts, it is not generally possible to automate parts of a process that rely on decision-making that is not based on existing rules.

### **Summary of the Invention**

25 According to a first aspect of the present invention, there is provided a method of processing status messages for use in network security management in respect of a network of computing devices, the status messages each comprising information relating to the network, information relating to at least one of said computing devices therein, or information relating to an event that has been detected by or in respect of the network or detected by or in respect of at least one of said computing devices therein, the information of which a status message is comprised comprising data relating to a plurality of attributes of said network, said at least one computing device, or said event; the method comprising:

filtering received status messages according to filtering rules in order to identify, in dependence on one or more attributes of said status messages:

- (i) status messages indicative of potential or actual network security events in respect of which a predetermined response is deemed applicable;
  - (ii) status messages in respect of which a null response is deemed applicable;
- and

- 5** (iii) residue messages in respect of which the filtering rules neither indicate that a predetermined response is deemed applicable, nor indicate that a null response is deemed applicable;

processing received status messages identified as status messages indicative of potential or actual network security events in respect of which a predetermined response is  
**10** deemed applicable, such that a predetermined response may be initiated;

in respect of one or more received status messages identified as residue messages, subjecting the one or more residue messages to analysis involving clustering in dependence on one or more attributes thereof; and

updating the filtering rules for use in subsequent filtering of received status  
**15** messages in dependence on the analysis.

According to preferred embodiments, the filtering of a received status message is performed in dependence on a comparison of one or more attributes of said message with corresponding attributes of one or more of said filtering rules.

According to preferred embodiments, the clustering of residue messages comprises  
**20** subjecting the one or more received status messages identified as residue messages to analysis involving clustering in dependence on one or more attributes of one or more existing filtering rules as well as on one or more attributes of the one or more received status messages identified as residue messages. It will be understood however that the clustering  
**25** of residue messages need not be dependent on attributes of existing filtering rules, and may be dependent just on attributes of the residue messages. Such analysis may identify clusters of residue messages indicative of a need for the addition of one or more new filtering rules, rather than merely identifying clusters of residue messages indicative of a need for an existing filtering rule to be updated, for example.

According to preferred embodiments, the clustering of residue messages is performed in  
**30** dependence on a plurality of attributes of each message.

According to preferred embodiments, the clustering of residue messages is performed in dependence on one or more disparity measures each calculated according to a predetermined disparity function dependent on one or more attributes of a residue message and one or more corresponding attributes of one or more of said filtering rules.

According to preferred embodiments, the clustering of residue messages is performed in dependence on one or more disparity measures each calculated according to a predetermined disparity function dependent on one or more attributes of a residue message and one or more corresponding attributes of one or more of other residue messages.

- 5 According to preferred embodiments, the updating of the filtering rules may comprise updating one or more of the existing filtering rules. Alternatively or additionally, the updating of the filtering rules may comprise adding one or more new filtering rules.

According to preferred embodiments, the filtering step comprises automatically filtering said received status messages.

- 10 According to preferred embodiments, the step of subjecting residue messages to analysis involving clustering comprises subjecting the residue messages to an automated analysis procedure involving clustering. This step may further comprise subjecting at least some of the residue messages to an automated analysis procedure involving sub-space clustering.

According to preferred embodiments, the method may further comprise performing a pre-

- 15 filtering step prior to the filtering step, the pre-filtering step involving preventing status messages according to one or more predetermined categories from being included in the filtering step. Such a pre-filtering step may be used to prevent internal or confidential messages from being forwarded to a remote site or to an external organisation, for example.

According to preferred embodiments, the method may further comprise performing a pre-

- 20 filtering step prior to the filtering step, the pre-filtering step involving identifying received status messages either as being potentially related to network security or as not being potentially related to network security, and performing the filtering step in respect of status messages identified as being potentially related to network security. Such a pre-filtering step may be used to prevent messages known to be unrelated to network security from affecting  
25 the post-filtering analysis.

According to preferred embodiments, the step of processing received status messages identified as being indicative of potential or actual network security events in respect of which a predetermined response is deemed applicable may comprise aggregating a plurality of messages so-identified that have one or more corresponding or similar attributes, whereby to  
30 allow a predetermined response to be taken in relation to the aggregated messages. Such aggregation may be used to facilitate subsequent analysis by a human analyst of messages indicative of potential or actual network security events, for example.

According to a second aspect of the present invention, there is provided apparatus for processing status messages for use in network security management in respect of a network

of computing devices, the status messages each comprising information relating to the network, information relating to at least one of said computing devices therein, or information relating to an event that has been detected by or in respect of the network or detected by or in respect of at least one of said computing devices therein, the information of which a status message is comprised comprising data relating to a plurality of attributes of said network, said at least one computing device, or said event; the apparatus comprising:

a filtering module operable to filter received status messages according to filtering rules in order to identify, in dependence on one or more attributes of said status messages:

- (i) status messages indicative of potential or actual network security events in respect of which a predetermined response is deemed applicable;
- (ii) status messages in respect of which a null response is deemed applicable; and
- (iii) residue messages in respect of which the filtering rules neither indicate that a predetermined response is deemed applicable, nor indicate that a null response is deemed applicable;

a processing module operable to process received status messages identified as status messages indicative of potential or actual network security events in respect of which a predetermined response is deemed applicable, such that a predetermined response may be initiated;

an analysis module operable to subject one or more received status messages identified as residue messages to analysis involving clustering in dependence on one or more attributes thereof; and

a filter-rule updating module operable to update the filtering rules for use in subsequent filtering by said filtering module of received status messages in dependence on the analysis.

The various options and preferred embodiments referred to above in relation to the first aspect are also applicable in relation to the second aspect.

### **Brief Description of the Drawings**

A preferred embodiment of the present invention will now be described with reference to the appended drawings, in which:

**Figure 1** shows the possible message flow in a network security monitoring process;

**Figure 2** illustrates the logging of network traffic for analysis;

**Figure 3** illustrates the message flow in a data collection and filtering system;

**Figure 4** illustrates a manual residue message handling process;

**Figure 5** illustrates a residue message handling process using clustering;

**Figure 6** illustrates the concept of clustering for residue messages analysis;

**Figure 7** is a possible residue message process according to a preferred embodiment;

**Figure 8** is a flowchart illustrating a possible process for updating a filter engine following the  
5 processing of residue messages, which may be used in a method according to a preferred embodiment;

**Figure 9** shows three examples of messages, and illustrates possible data fields of which status messages or network security messages may be comprised;

**Figure 10** illustrates the concept of subspace clustering for residue message analysis; and

10 **Figure 11** illustrates various stages of residue message analysis by clustering.

### **Description of Preferred Embodiments of the Invention**

With reference to the accompanying figures, methods and apparatus according to preferred embodiments will now be described.

**Figure 1** shows the flow of messages within a network security monitoring process. As will  
15 become apparent, parts of this correspond partly to parts of a network security monitoring process such as may be performed by the Counterpane product referred to earlier.

Overall, messages from sensors 5 are received by a data collection and filtering system 10. The sensors may monitor and provide messages relating to the customer network in general, the devices therein, customer network traffic (internal communication and communication  
20 with external devices), sending status messages about these or other aspects of the network to the data collection and filtering system 10. This may thus receive messages relating to the customer network and to computing devices therein, which may comprise information relating to their status, and/or to events that have been detected by or in respect of the network or detected by or in respect of the computing devices therein.

25 Sensors generally sit within customer networks and log network traffic. As indicated by **Figure 2**, the traffic may be logged depending upon specific customer requirements. This be in dependence on the devices therein and/or applications used therein, or could be dependent on specific threats mentioned by a customer, for example. Data obtained by the sensors 5 may be pre-processed before being passed to the data collection and filtering  
30 system 10. This may be done in order to remove messages that clearly relate to internal matters specific to the customer, for example, or to remove messages that clearly relate to other, clearly non-security-related matters, for example, that the customer may not wish to

pass to an external service provider. This may be done simply in order to reduce the amount of messages that are passed to the external service provider, or to allow the external service provider to use their expertise specifically in relation to messages that the customer may class as "network security messages", on the grounds that they might include messages  
5 indicative of security problems but need expert analysis. Alternatively, the customer may wish to provide all data obtained by the sensors to the data collection and filtering system 10, in order to allow it to perform its filtering function on all such data.

The data collection and filtering system 10 may also sit (topologically) within the customer network and/or be (geographically) located on the customer's premises, or it may be located  
10 remotely from the customer's network and/or premises, on the network of and/or at the premises of the external service provider, for example.

The data collection and filtering system 10 contains a filter engine 12, which uses filter rules to classify the messages received as "positive messages", "negative messages", and "residue messages". The filter rules may (initially) be generic or may be customer-specific.  
15 Messages classed as "negative" may be discarded, because they do not relate to security issues, although they may be stored for use in further analysis and updating of the filter rules. Messages classed as "positive" are passed to a queue and aggregation module 14, which aggregates them before forwarding them to a problem and expertise system 20, the function of which will be discussed further later. These may also be used in further analysis and  
20 updating of the filter rules, however. Messages classed neither as positive nor as negative are passed to a "residue messages module" 16.

As indicated by **Figure 3**, there are generally two types of filters, generally distinguished from each other more by the use made thereof than by the general manner in which they function. For convenience, these two types will be referred to as "negative filters" 32 and "positive  
25 filters" 34, although it will be noted that this terminology is optional, and should not be taken to imply that one type simply blocks or allows messages with values above a particular threshold while the other blocks or allows messages with values below a particular threshold. Instead these terms are used in accordance with the idea that messages may be identified as "positive" or "negative" messages as a result of being passed to the filter engine 12 (or as  
30 explained earlier, may be identified as neither, in which case they may instead be identified as "residue messages").

Filters of either type (i.e. "positive" or "negative") implement filtering rules, which generally involve a comparison of one or more attributes of a status message with one or more corresponding attributes as defined by a filtering rule. Filtering rules may have attributes

defined in terms of thresholds (upper or lower), ranges (inclusive or exclusive), equalities or inequalities, or in any or a variety of other ways, allowing corresponding attributes of messages to be compared therewith to determine, on the basis of one or a combination of attributes, whether a message concerned meets the filtering rule of a particular (positive or negative)

5 filter or not.

Negative filters 32 are essentially used to identify and remove messages in respect of which a null response is deemed applicable. These may be messages deemed as irrelevant in that they do not need any further scrutiny, and in respect of which it is known that there is no need for action to be taken in relation to network security.

10 Positive filters 34 are essentially used to identify and select messages indicative of potential or actual network security events in respect of which a predetermined response is deemed applicable. These may be messages in respect of which it is known or decided that a particular countermeasure is appropriate. Alternatively, the predetermined response may be that further analysis is to be performed, either by a human analyst or by a machine.

15 Messages in respect of which the filtering rules of the filter engine neither indicate that a predetermined response is deemed applicable, nor indicate that a null response is deemed applicable, are not "caught" by either type of filter, and are referred to as residue messages.

As explained earlier, if using a system such as the existing Counterpane product, considerable manual effort is often needed (or at least should be performed, but may not

20 always be performed) following the corresponding filtering stage of the process, in order to analyse and "understand" the residue messages, and possibly classify them. As indicated by **Figure 4**, analysts may analyse residue messages using SQL queries, depending upon the analyst's understanding of the domain, or other factors. They try to understand the messages, whether they were classified as residue messages due to a change in message  
25 format and/or content or due to a new type of message that does not match the existing knowledge-base. As indicated earlier, residue messages are sometimes left un-analysed because of lack of time, allowing important messages to be ignored.

As indicated by **Figure 5**, techniques according to preferred embodiments, while not necessarily involving complete automation of the process, may be used in place of human  
30 analysts, or to assist analysts in processing the residue messages faster, better, and more efficiently. Such techniques may also help to prioritise the processing of high-severity messages that fall into the residue.

Referring back to **Figure 1**, the problem and expertise system 20 categorises, prioritises, investigates and responds to customer incidents or problems. Further detailed explanation of

aspects of a possible version of a problem and expertise system 20 (i.e. the "SOCRATES" system), including its components and its functionality, are described in US7895641, the contents of which is incorporated herein by reference, but a brief summary of the overall functionality of such a system will follow in order to illustrate the context in which the

5 presently-described embodiment may be used.

One or more gateway systems 21 may serve a "gatekeeper" function in linking the data collection and filtering system 10 to a Secure Operations Centre (SOC) in which the problem and expertise system 20 is located. Gateways 21 may reformat, add information to, or delete information from incoming messages from the data collection and filtering system 10 to

10 ensure maximum utility of the output, or may allow mutual "pulse monitoring" by the respective systems of their counterparts to ensure that each is continuously operational.

After being received and possibly processed by gateway systems 21, "gateway messages" may be passed for analysis in conjunction with information in a database 22, which may include security intelligence, customer information, problem resolution information, etc. This

15 may pass messages for analysis in a correlation module 23 and/or a symptom anomaly module 24.

The correlation module 23 may be an engine or a series of engines that receive gateway messages as inputs, and provide an output in the form of new gateway messages. It may be used to detect when a combination of events has occurred at various locations or devices

20 and/or times within a customer's network. From this, patterns may emerge that indicate that "correlated" security alerts based on prior non-correlated security alerts should be created.

The symptom anomaly module 24 may receive gateway messages and analyse them in the light of existing event information in order to explain what event(s) may have occurred and caused the messages, and to identify how to verify this and what should be done about it

25 Some gateway messages may be the result of messages from the data collection and filtering system 10 that are marked as relating to firewall issues, or which are otherwise recognisable as "firewall messages". These may be passed for analysis to a firewall application server 25 and a firewall database 26. The firewall application server 25 may receive firewall messages (possibly in compressed or "zipped" form), decompress/unzip

30 them if necessary and forward individual firewall messages to an assigned firewall database. They may be process similarly to other messages, but using a simpler and more parallel workflow than that of non-firewall messages.

In any of the above manners, and possibly others, the problem and expertise system 20 system thus collects and formats gateway messages into "problem tickets", each of which

represents a discrete security-related event or incident of possible intrusive activity happening on a customer's network. The correlation module 23, symptom anomaly module 24, firewall database 26, and analysts themselves may all generate tickets.

5 A ticket tracking system 27 keeps track of the tickets, and may also keep track of any follow-ups and updates that result therefrom. The information from the problem and expertise system 20 may be shared with customers through a customer portal. Tickets and/or advice based thereon may be passed to customers, who may act on them to maintain the security of their networks. Customers may also communicate directly with the analysts if they need to, if the information available to them via their customer portal is insufficient, for example.

10 Now that the overall message flow within a network security monitoring process has been outlined, there now follows a more detailed explanation of the processing of residue messages, and an explanation of the concept of "clustering" as applied thereto.

As shown in **Figure 6**, sensors provide messages which then get filtered using an existing knowledgebase represented by the filters of the filter engine into positive, negative and  
15 residue messages. The residue is then sent to a 'clustering engine' where analysis of the residue messages (which is explained below) is carried out using a 'clustering' technique. The results of this analysis may be provided to human analysts via a user interface 60 for further analysis, or be used directly and/or automatically in relation to the existing knowledgebase, in order to update the knowledgebase. As will be explained later, this may  
20 involve updating (or possibly removing or merging) existing filters from the filter engine, by altering (or possibly cancelling) the rules thereof, or it may involve adding one or more new filters to the filter engine by including new rules therein.

In this manner, the results of analysis of the residue messages using clustering are fed back to the knowledgebase, which may thus be updated, thereby enabling classification of  
25 subsequent residue messages to be carried out using the updated knowledgebase. It will be noted that such updating may therefore be done essentially in real-time.

**Figure 7** is a flowchart symbolising the processing of the residue messages. Residue messages are received (step s70) by the data collection and filtering system 10 and passed to the filter engine 12. In a filtering step (s71), the received messages are classified as  
30 negative messages or positive messages, or as residue messages. Messages classified as "negative" may be dropped (s72), and messages classified as "positive" may be passed on for further analysis or action (s73). These processes have been discussed elsewhere, and will therefore not be discussed further here.

The residue messages are sent to a 'clustering engine', which may be a separate processing module within the data collection and filtering system 10, or may be provided by other means (remote processing, for example). As will be explained, in the clustering step (s74), processing is performed on messages in dependence on their attributes, and generally in

5 dependence on attributes of existing filters as well, in order to determine whether the residue messages are in fact "near" to existing filters (despite not having been "caught" by them in the earlier filtering process), and also to determine in dependence on their attributes and in dependence on the attributes of other residue messages whether they are similar to other

10 messages may be dependent just on attributes of the residue messages, and not on those of existing filters. Such analysis may identify clusters of residue messages indicative of a need for the addition of one or more new filtering rules, rather than merely identifying clusters of residue messages indicative of a need for updating of an existing filtering rule, for example.

Such classification of the residue messages (i.e. into clusters and as clusters near existing

15 positive or negative filters or as clusters not near existing filters) happens in step s75, allowing filter updates to be determined. These are then used to update the filter engine as appropriate (s76), such that subsequent residue messages received may be subjected to filtering (s71) by the filter engine 14 according to the updated and/or new filtering rules.

As well as being used to determine updates to the filtering rules, residue messages that

20 appear at this stage to be near to "high-severity" filters may also trigger alarms at this stage, allowing action to be taken even though these messages may not have been caught by positive filters in the earlier filtering process (due to them not quite matching the rules of those filters), or may be passed to an analyst who may take such messages into consideration when performing further analysis.

25 In determining whether residue messages should form part of a cluster, a disparity measure referred to as a "distance" may be used. This is a measure, in a (generally) multi-dimensional "cluster-space", of the overall disparity between a particular message and a particular cluster, with reference to the respective attributes by which each is categorised.

The distance "d" of a residue message "m" from respective clusters "C<sub>i</sub>" may be calculated

30 using a function in the following form, for example:

$$d_{m,i} = \sqrt{(A_1 - A_{Ci,1})^2 + (A_2 - A_{Ci,2})^2 + \dots + (A_n - A_{Ci,n})^2}$$

where  $A_x$  is the value of the  $x^{\text{th}}$  attribute of message "m",  $A_{Ci,y}$  is the value of the  $y^{\text{th}}$  attribute of cluster "C<sub>i</sub>", and "n" is the number of attributes being taken into consideration. Other

functions may of course be used that take account of some or all of the respective attributes by which the messages and the clusters are categorised, and indicate in dependence thereon a measure of disparity between a particular message and the respective clusters.

Referring back to **Figure 6**, this attempts to indicate, using the two dimensions available in a drawing, how residue messages, or newly-formed clusters of residue messages (R1, R2, R3) may be regarded as being "near" existing positive filters (R1 is near "+ve"), "near" existing negative filters (R2 is near "-ve"), or not near any existing filters (R3).

If the residue messages appear near normal positive filters and can be classified as positive, they can then join the normal queue of positive messages and be processed accordingly. If the messages appear near normal negative filters and can be classified as negative, they can be dropped. The knowledge gained by this process is fed back to the knowledge system by way of filter engine updates. If the new filters need to be formed (due to a determination that messages such as R3 should in fact be considered as negative, and that subsequent messages similar to them should be caught by a negative filter) then the knowledgebase is upgraded to incorporate a new filter having appropriate filtering rules to catch subsequent residue messages having attributes such as those in R3.

Referring to the flowchart shown as **Figure 8**, this illustrates two ways in which change in the knowledge-base from residue message analysis can be incorporated into the filter engine. One involves modifying one or more of the existing filters, e.g. if the characteristics of the residue messages are found to be similar (using a disparity measure such as that discussed above, or otherwise) to those of one or more existing filters. Another is to add new negative and/or positive filters (again depending upon the characteristics of the residue message).

As explained earlier, the residue is provided to the clustering engine, with existing positive and negative filters being used as seeds in the clustering process (s74 of Figure 7). The largest radius of the clusters may be denoted by ' $r_c$ ', or each cluster ' $C_i$ ' may have a radius ' $r_{ci}$ '. Residue messages are clustered as appropriate. Some messages and/or clusters of messages may be found to be near existing filters, indicating that the rules for those filters should be updated in order that similar such messages subsequently received should be caught by those filters. Any message that is sufficiently distant from any of the filters (i.e. outside the appropriate radius for each) may be considered as an outlier, and may be clustered with other such messages or serve as a basis for forming a new cluster.

Looking into this in more detail, at step s80, the distance of each cluster from the existing filters is determined. For those that are sufficiently near an existing filter, an update to the

rules of that filter may be determined (s75 of Figure 7) and implemented in the filter engine (s76 of Figure 7).

At step s81, messages found to have a distance from any existing filter greater than the relevant threshold (i.e. ' $r_c$ ' ' $r_{ci}$ ', etc.) are labelled as outlier messages.

- 5 If it is determined at step s82 that any messages are in the vicinity of any of the existing filters, the closest negative or positive filters are identified (s83). For messages near a particular filter, the cluster characteristics are then studied using subspace clustering (s84) and compared with the filter characteristics. This may indicate that the rules of that filter are failing to catch messages due to the value or identity of a particular attribute, but would
- 10 otherwise (correctly) have been caught by that filter, for example. An analyst may then use this information to identify an update to the rules of that filter (s85) that would result in those messages being caught (if deemed appropriate), possibly by changing that filter's requirement in relation to the attribute highlighted in the subspace clustering step. An appropriate update may then be implemented (s76 of Figure 7).
- 15 At step s86, the remaining outlier messages (i.e. those that are not in the vicinity of any existing filters) are considered. At step s87, clustering is carried out in respect of the remaining outliers. At step s88, subspace clustering is carried out to understand these clusters and their relationships if any. Similarly, this may be in order to identify particular attributes that are resulting in a significant number of these remaining outlier messages not
- 20 being near let alone caught by any existing filter. At step s89, new filters may be created based on the knowledge obtained in step s88 if appropriate.

A sample of message data for the analysis is shown in **Figure 9**. In this, the fields shown may represent the following attributes (although it should be noted that these fields and their explanations are given by way of example only):

- 25 - '*src*' represents the source IP address
- '*src\_port*' represents the source port number
- '*Dst*' represents the destination IP address, i.e. the IP address of the device where the last target destination of the message is.
- '*dst\_port*' represents the destination port address
- 30 - '*dst\_country*' represents the country where the destination device is located.
- '*service*' represents the service currently being used, and/or to which the message relates
- '*REDefID*' refers to the table of contents for the syslog matching expression sets
- '*date*' and '*time*' are self-explanatory
- '*device\_id*' represents the sensor ID, or may represent the device IP address

- '*duration*' indicates the length of time for which the application was accessed (or may indicate how long a transaction took, for example)
  - '*app*' represents the application currently being used, and/or to which the message relates
  - '*app\_cat*' represents the application category
- 5 - '*user*' represents the user of the device, or the user currently registered as using the device

By way of example, we will take the simple case of a positive filter that covers the following:

*src\_ip: All, src\_port: 3108, dst\_port: 80, app\_cat: file-transfer, user: All*

On filtering, the first message from Figure 9 would pass through the positive filter and would therefore be sent for further analysis. The second message does not pass through the filter,  
10 so (assuming it is not caught by any other filter, positive or negative), it lands in the residue. On clustering, it may be found to be near to the above filter but not caught by it (due to the value in the '*src\_port*' field). Continuing further with subspace clustering may indicate that '*src\_port*' and '*app\_cat*' are the attributes that are keeping the message apart from the filter cluster in question.

15 From the above attributes, an analyst can decide to update the existing filter to accommodate subsequent messages the same as or similar to the message in question, or include a new filter, perhaps creating a new negative filter to ensure that such messages are automatically regarded as negative and dropped in the future.

Residue messages form clusters around positive and negative filters as well as independent  
20 clusters (depending on the messages). Clusters of residue messages around filters may be distant from actual filters for reasons that may need to be analysed. The details of each cluster may be analysed separately using the technique of subspace clustering, which is illustrated by **Figure 10**. Features selected for subspace clustering may provide more information, indicating, for example, the attribute or attributes by virtue of which messages or  
25 clusters of messages are distant from one or more filters. For example, as explained in relation to the example messages above, it may be identified that '*src\_port*' and '*app\_cat*' are the attributes by virtue of which the messages are being kept apart from a particular filter cluster. Thus, some or all of the residue messages may automatically be compared with  
30 characteristic of the clusters and the filters, thereby saving time for analysts, who can then make use of this information to modify and/or to update the existing filters.

Using the correct attributes may help in better understanding the data by producing subspace clusters. This concept can be used to learn more about the data.

As indicated, each of the clusters produced may then be analysed using subspace clustering. This may lead to clusters being found in different subspaces within a dataset. Thus each cluster around a filter may get further divided into groups that may help to understand the reasons why the messages were not filtered, and assist in the process of

5 updating or modifying the filter.

Referring to **Figure 11**, in a first stage outlier clustering is used to separate the messages that are not close to any of the positive or negative messages and appear to need completely new filters. This happens because the knowledgebase of filters is not rich enough to deal with these messages or the knowledgebase is still under development.

10 The aim here may be not to remove the data only because it is "noise" but also to separate the data that could help in the forming of new filters. In the case of outliers, important information may be provided by data that on initial analysis may appear to be part of the background noise, relating to messages that do not confirm to the rules of any existing filters. While the values of attributes in messages may be well beyond the attribute value range of

15 an existing filter, it may not be appropriate to dismiss this as noise in the data. Instead, it may indicate that a particular attribute of an existing filter needs to be changed, or that a new filter should be created. In some cases, it may turn out that this is simply noise, but analysis using clustering (and possibly subspace clustering) may allow this to be determined.

For outliers, i.e. messages or clusters thereof that are not within the vicinity of any positive or

20 negative clusters, clustering may be carried out using ' $r_c$ ' (or ' $r_{ci}$ ', as appropriate) as a minimum distance for the cluster centres. Analysts may use these new clusters to create new negative and/or positive filters.

Looking again at the sample messages shown in Figure 9, the third message has a few attributes by virtue of which it differs from the filter (defined earlier) and the other two sample

25 messages. This third message may therefore be considered as an "outlier". On carrying out clustering on the outliers, an analyst can see characteristics of the clustered outliers and can form a new filter using the cluster information. An analyst can use subspace clustering on the clusters formed from outliers to understand the information in more detail, and form more crisp filters.

30 An analyst may have a graphic user interface 60 as shown in **Figure 6**, in which clusters and messages may be represented visually so that the analyst may take action to update or to create filters using procedures such as those described above.

**Claims**

- 1) A method of processing status messages for use in network security management in respect of a network of computing devices, the status messages each comprising information relating to the network, information relating to at least one of said computing devices therein,  
5 or information relating to an event that has been detected by or in respect of the network or detected by or in respect of at least one of said computing devices therein, the information of which a status message is comprised comprising data relating to a plurality of attributes of said network, said at least one computing device, or said event; the method comprising one or more processors performing steps of:
- 10 filtering received status messages according to filtering rules in order to identify, in dependence on one or more attributes of said status messages:
- (i) status messages indicative of potential or actual network security events in respect of which a predetermined response is deemed applicable;
  - (ii) status messages in respect of which a null response is deemed applicable;
  - 15 and
  - (iii) residue messages in respect of which the filtering rules neither indicate that a predetermined response is deemed applicable, nor indicate that a null response is deemed applicable;
- processing received status messages identified as status messages indicative of  
20 potential or actual network security events in respect of which a predetermined response is deemed applicable, such that a predetermined response may be initiated;
- in respect of one or more received status messages identified as residue messages, subjecting the one or more residue messages to analysis involving clustering in dependence on one or more attributes thereof; and
- 25 updating the filtering rules for use in subsequent filtering of received status messages using an automated updating procedure dependent on one or more results of the analysis.
- 2) A method according to Claim 1 wherein the filtering of a received status message is  
30 performed in dependence on a comparison of one or more attributes of said message with corresponding attributes of one or more of said filtering rules.
- 3) A method according to Claim 1 or 2 wherein the clustering of residue messages is performed in dependence on a plurality of attributes of each message.

- 4) A method according to any of the preceding claims wherein the clustering of residue messages is performed in dependence on one or more disparity measures each calculated according to a predetermined disparity function dependent on one or more attributes of a residue message and one or more corresponding attributes of one or more of said filtering
- 5 rules.
- 5) A method according to any of the preceding claims wherein the clustering of residue messages is performed in dependence on one or more disparity measures each calculated according to a predetermined disparity function dependent on one or more attributes of a
- 10 residue message and one or more corresponding attributes of one or more of other residue messages.
- 6) A method according to any of the preceding claims wherein updating the filtering rules comprises updating one or more of the existing filtering rules.
- 15
- 7) A method according to any of the preceding claims wherein updating the filtering rules comprises adding one or more new filtering rules.
- 8) A method according to any of the preceding claims wherein the filtering step
- 20 comprises automatically filtering said received status messages.
- 9) A method according to any of the preceding claims wherein the step of subjecting residue messages to analysis involving clustering comprises subjecting the residue messages to an automated analysis procedure involving clustering.
- 25
- 10) A method according to Claim 9 wherein the step of subjecting residue messages to analysis involving clustering further comprises subjecting at least some of the residue messages to an automated analysis procedure involving sub-space clustering.
- 30 11) A method according to any of the preceding claims, further comprising performing a pre-filtering step prior to the filtering step, the pre-filtering step involving preventing status messages according to one or more predetermined categories from being included in the filtering step.
- 35 12) A method according to any of the preceding claims, further comprising performing a pre-filtering step prior to the filtering step, the pre-filtering step involving identifying received

status messages either as being potentially related to network security or as not being potentially related to network security, and performing the filtering step in respect of status messages identified as being potentially related to network security.

13) A method according to any of the preceding claims wherein the step of processing  
5 received status messages identified as status messages indicative of potential or actual network security events in respect of which a predetermined response is deemed applicable comprises aggregating a plurality of messages so-identified that have one or more corresponding or similar attributes, whereby to allow a predetermined response to be taken in relation to the aggregated messages.

10

14) A method according to any of the preceding claims wherein the clustering of residue  
messages comprises subjecting the one or more received status messages identified as  
residue messages to analysis involving clustering in dependence on one or more attributes  
of one or more existing filtering rules and on one or more attributes of the one or more  
15 received status messages identified as residue messages.

15) Apparatus for processing status messages for use in network security management  
in respect of a network of computing devices, the status messages each comprising  
information relating to the network, information relating to at least one of said computing  
20 devices therein, or information relating to an event that has been detected by or in respect of  
the network or detected by or in respect of at least one of said computing devices therein, the  
information of which a status message is comprised comprising data relating to a plurality of  
attributes of said network, said at least one computing device, or said event; the apparatus  
comprising:

25 a filtering module operable to filter received status messages according to filtering  
rules in order to identify, in dependence on one or more attributes of said status messages:

- (i) status messages indicative of potential or actual network security events in  
respect of which a predetermined response is deemed applicable;
- (ii) status messages in respect of which a null response is deemed applicable;
- 30 and
- (iii) residue messages in respect of which the filtering rules neither indicate  
that a predetermined response is deemed applicable, nor indicate that a null  
response is deemed applicable;

a processing module operable to process received status messages identified as  
35 status messages indicative of potential or actual network security events in respect of which

a predetermined response is deemed applicable, such that a predetermined response may be initiated;

an analysis module operable to subject one or more received status messages identified as residue messages to analysis involving clustering in dependence on one or  
5 more attributes thereof; and

a filter-rule updating module operable to update the filtering rules for use in subsequent filtering by said filtering module of received status messages using an automated updating procedure dependent on one or more results of the analysis.




Figure 1 - Message Flow in a Network Security Monitoring Process

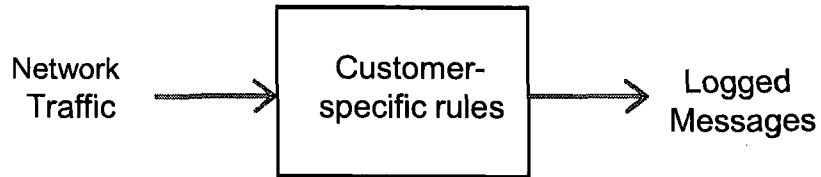


Figure 2 – Logging of network traffic for analysis




Figure 3 - Message flow in Data Collection and Filtering System

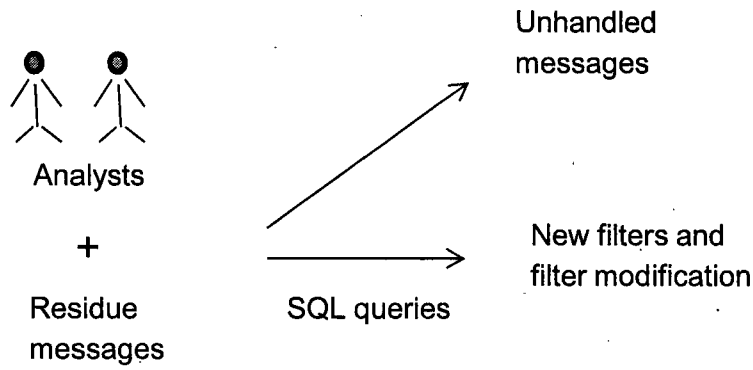


Figure 4 - Manual residue message handling process

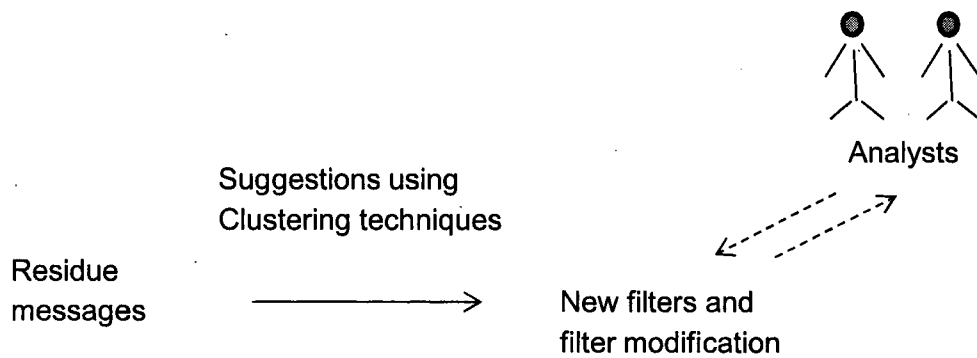


Figure 5 - Residue message handling process using clustering




Figure 6 - Concept of clustering for residue messages analysis

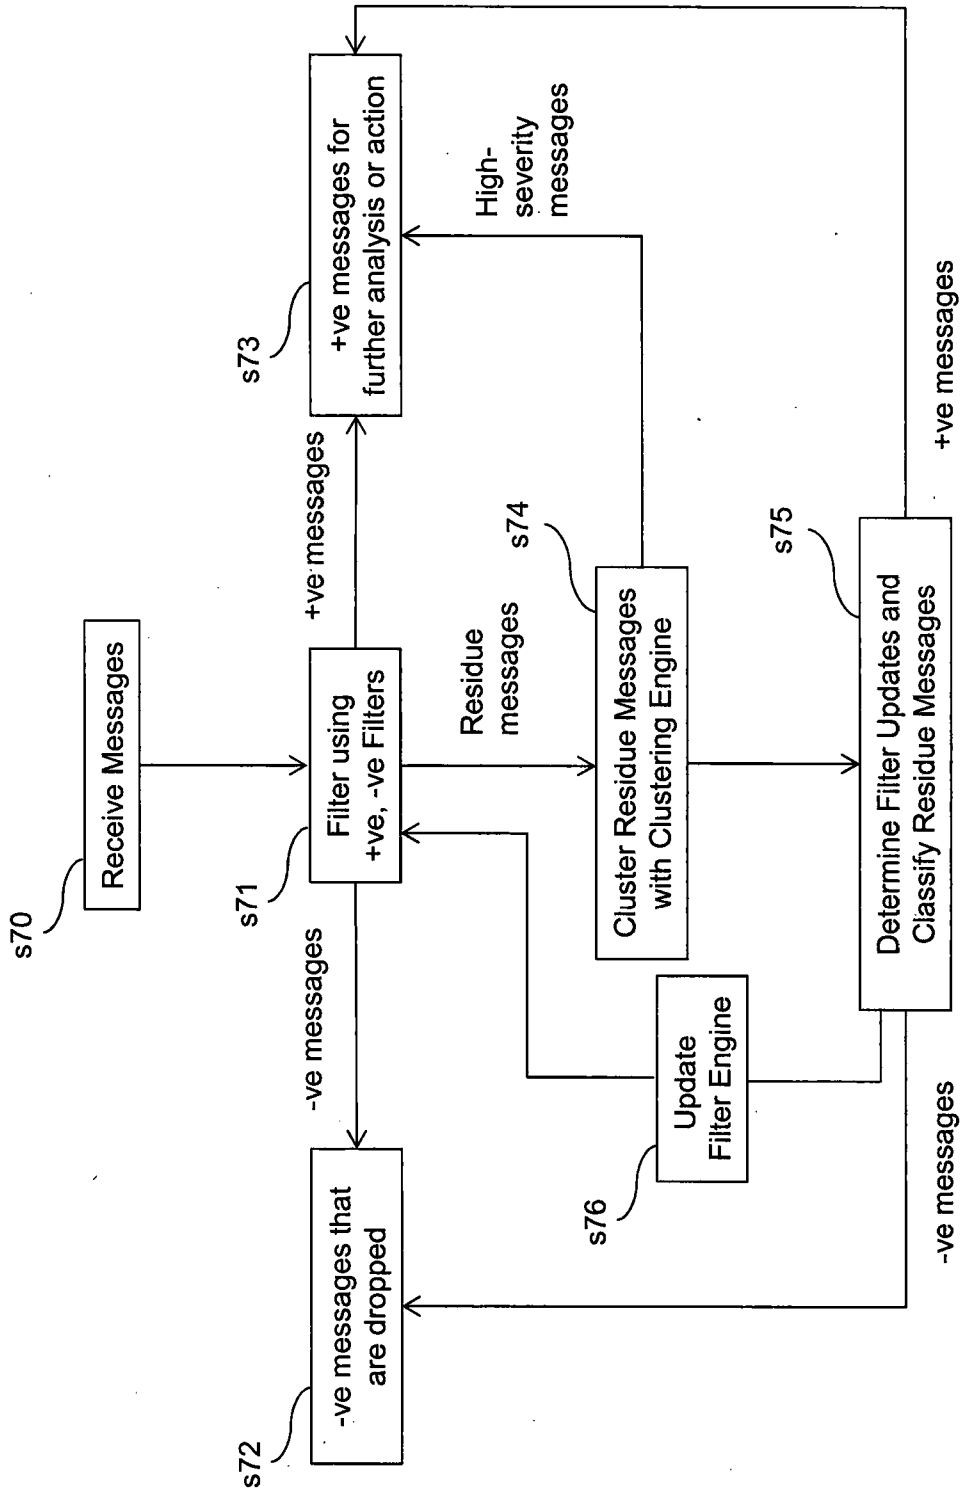


Figure 7 - Residue Message Process




Figure 8 - Flowchart for filter engine updating process

src	src_port	Dst	dst_port	dst_country	service	REDefID	date	time	device_id	duration	app	app_cat	user
172.30.12.42	3108	74.125.230.110	80	US	HTTP	FortiGate	13/04/2011	12:43:17	FGT3KB3408600 584	105	HTTP.Image	file-transfer	SALAS.A
172.30.12.42	3101	74.125.230.114	80	US	HTTP	FortiGate	13/04/2011	12:43:17	FGT3KB3408600 584	106	Google.Maps	web	SALAS.A
172.18.186.143	2685	172.16.254.20	88	N/A	88/udp	FortiGate	13/04/2011	12:43:17	FG300B3908607 012	182	N/A	N/A	N/A

Figure 9 - Examples of message data




Figure 10 - Subspace clustering for residue message analysis




Figure 11 - Various stages of residue message analysis by clustering

INTERNATIONAL SEARCH REPORT

International application No  
PCT/GB2013/000564

A. CLASSIFICATION OF SUBJECT MATTER  
INV. H04L29/06  
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)  
H04L

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 practicable, search terms used)  
EPO-Internal, WPI Data, INSPEC

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	WO 01/71499 A1 (COUNTERPANE INTERNET SECURITY [US]) 27 September 2001 (2001-09-27) cited in the application page 12, line 16 - page 13, line 2 figures 1,2	1-15
Y	K. Julisch: "Mining alarm clusters to improve alarm handling efficiency", 2001, XP002697689, Retrieved from the Internet: URL:http://www.acsac.org/2001/papers/115.pdf [retrieved on 2013-05-13] cited in the application the whole document	1-15
	----- -/--	

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 application or patent 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  14 February 2014	Date of mailing of the international search report  24/02/2014
---	--

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  Olaechea, Javier
--	--

## INTERNATIONAL SEARCH REPORT

International application No  
PCT/GB2013/000564

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	Adelina J D; Anushiadevi R; Lakshminarayanan T R: "An efficient algorithm for clustering intrusion alert",  31 March 2012 (2012-03-31), XP002697690, Retrieved from the Internet: URL:http://www.jatit.org/volumes/Vol137No2/ 11Vol137No2.pdf [retrieved on 2013-05-13] cited in the application the whole document	1-15
A	----- US 2004/015719 A1 (LEE DAE-HYUNG [KR] ET AL) 22 January 2004 (2004-01-22) cited in the application paragraph [0068] - paragraph [0086]	1-15
A	----- US 2005/060562 A1 (BHATTACHARYA PARTHA [US] ET AL) 17 March 2005 (2005-03-17) cited in the application paragraph [0031] - paragraph [0042]	1-15
A	----- US 8 176 527 B1 (NJEMANZE HUGH S [US] ET AL) 8 May 2012 (2012-05-08) cited in the application page 3, line 66 - page 5, line 45 page 9, line 1 - page 14, line 37	1-15
A	----- WO 2010/114363 A1 (UNIV TWENTE [NL]; BOLZONI DAMIANO [NL]; ETALLE SANDRO [NL]) 7 October 2010 (2010-10-07) page 2, line 18 - page 4, line 27 page 5, line 19 - page 8, line 13	1-15
A	----- US 7 506 373 B2 (MORIN BENJAMIN [FR] ET AL) 17 March 2009 (2009-03-17) column 1, line 14 - column 4, line 16 column 11, line 1 - column 12, line 3	1-15

# INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/GB2013/000564

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 0171499	A1	27-09-2001	AU 4735101 A 03-10-2001
			US 2002087882 A1 04-07-2002
			US 2007162973 A1 12-07-2007
			WO 0171499 A1 27-09-2001
-----			
US 2004015719	A1	22-01-2004	NONE
-----			
US 2005060562	A1	17-03-2005	CN 1829953 A 06-09-2006
			EP 1665011 A2 07-06-2006
			US 2005060562 A1 17-03-2005
			US 2010058165 A1 04-03-2010
			WO 2005026900 A2 24-03-2005
-----			
US 8176527	B1	08-05-2012	US 8176527 B1 08-05-2012
			US 2012260306 A1 11-10-2012
-----			
WO 2010114363	A1	07-10-2010	EP 2415229 A1 08-02-2012
			US 2012036577 A1 09-02-2012
			WO 2010114363 A1 07-10-2010
-----			
US 7506373	B2	17-03-2009	AT 358373 T 15-04-2007
			DE 602004005616 T2 24-01-2008
			EP 1695485 A2 30-08-2006
			FR 2864392 A1 24-06-2005
			US 2007118905 A1 24-05-2007
			WO 2005060160 A2 30-06-2005
-----			