

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
20 September 2007 (20.09.2007)

PCT

(10) International Publication Number
WO 2007/104988 A1

- (51) International Patent Classification:
H04L 29/06 (2006.01) *G06F 21/00* (2006.01)
- (21) International Application Number:
PCT/GB2007/000901
- (22) International Filing Date: 14 March 2007 (14.03.2007)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
0605117.1 14 March 2006 (14.03.2006) GB
- (71) Applicant (for all designated States except US):
STREAMSHIELD NETWORKS LIMITED [GB/GB];
St. Mary's Court, The Broadway, Amersham, Bucking-
hamshire HP7 0UT (GB).


- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, 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, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, 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, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

- (72) Inventor; and
- (75) Inventor/Applicant (for US only): **CURNYN, Jon** [GB/GB]; 6 Bakers Orchard, Wooburn Green, Bucking-
hamshire HP10 0LS (GB).
- (74) Agent: **GILL JENNINGS & EVERY LLP**; Broadgate House, 7 Eldon Street, London EC2M 7LH (GB).

Published:
— with international search report

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: A METHOD AND APPARATUS FOR PROVIDING NETWORK SECURITY



(57) Abstract: The invention relates to the provision of virus scanning capabilities in a network environment. The volume of traffic sent to one or more virus scanners is reduced through preliminary analysis of content carried by the network. This preliminary analysis is effective to identify whether or not a virus scan on the content is necessary. This analysis involves calculating a digest for the content in dependence upon a source address of the content. The content is identified as safe if the digest matches a stored digest. The content that is not safe is passed to a virus scanner.

WO 2007/104988 A1

A METHOD AND APPARATUS FOR PROVIDING NETWORK SECURITY

Field of the Invention

The present invention relates to network security. In particular, the present
5 invention relates to an apparatus and method of providing high-throughput anti-virus
(AV) services to a large number of subscribers.

Background to the Invention

There are many proven AV scanners in use today, and these scanners have
10 gained considerable market acceptance for use in desktop, file server and gateway
applications. Customers are able to rely on independent information and advice to
select a scanner vendor, and then trust that vendor's product to reliably detect malware.

However, while the performance of these scanners is acceptable for desktop,
server and gateway usage, it is not sufficient for use in high speed network
15 infrastructures such as the core of the internet. The production of a new, high
performance scanner presents not only technical difficulties but also issues of market
acceptance (users are understandably unwilling to rely on untried products for their
security). As such, it is advantageous to develop a solution incorporating existing
scanners in such a manner that the overall performance of the solution is sufficient for
20 deployment in these high speed network infrastructures.

It is known to use existing third party scanners within network applications. For
example, organisations known as Managed Security Service Providers (MSSPs) offer
services such as scanning all e-mail that passes through a subscriber's internet
connection for viruses. Typically, this is done by diverting customer traffic through the
25 MSSP's site. The traffic is then scanned by conventional software running on
conventional personal computers (PCs). However, to scale the scanner performance to
the required levels of both high throughput and low latency, it is often necessary to
deploy of a large number of PCs operating scanners. Where this number of PCs grows
large, the amount of external infrastructure such as switches and load balancers
30 required to coordinate the system also increases. This results in both expense and
unreliability.

Typically, in such an installation the large number of PCs all operate the same
set of tasks. These tasks include:

- receiving and transmitting data into and out of the PC;
- 35 - decoding and operating the protocols that carry this data;
- copying this decoded data to the computer's main memory or disk;
- invoking one or more AV scanners;

- sending the data to one or more AV scanners;
- undertaking the scanning tasks such as decompression, content decode, signature matching, heuristics analysis;
- processing the results from the scanners;
- 5 - transmitting the data (if not infected), or an alternative to it (if infected), onto the intended destination; and,
- finally collecting and storing any statistics or other logging information on the tasks undertaken.

As such, the scanner on each PC receives data regardless of the type or level of
10 threat from the content. However, the threat level depends on the application being used (e.g. web browsing, e-mail, peer to peer (P2P)) and the program being used to operate the application (for example, the Internet Explorer web browser). These factors are discussed further below:

- the application for which the content is intended: there are numerous types of
15 malware in existence today ranging from mass-mailers to Trojans. However some of these threats are specific to certain applications, such that they can only be propagated and become active through a single application but no other; for example a mass mailing virus cannot be picked up and propagated through web browsing;
- 20 - the program by which the content is used: in addition to traditional forms of files based malware such as viruses, Trojans, worms etc., there exist a number of vulnerabilities in the programs (such as web browsers) that operate applications, and these vulnerabilities may be exploited by specially crafted pieces of content. These vulnerabilities are specific to each program. As such, a vulnerability in
25 one program used as a web browser will not exist in a second program used as an e-mail client.

In addition to the above, the type of content being supplied will have a bearing on the threat level. In this context, content will broadly fall into two categories, executable and non-executable. Executable content poses a significantly higher threat. Executable
30 content is able, once executed, to gain control of a computer and subsequently can then execute any payload it chooses (for example, it could delete the contents of a hard drive). Moreover, executable content can come in many forms and can use complex techniques to disguise itself (such as encryption and metamorphism). In contrast, non-executable content can only pose a threat by exploiting vulnerabilities in the programs
35 which use the content. As a result, the content cannot take variable forms since it exploits static vulnerabilities; consequently threats due to non-executable content are often easier to detect than those due to executable content.

Summary of the Invention

According to a first aspect of the present invention, there is provided a network security apparatus, comprising: one or more network traffic processors adapted to receive network traffic and to extract content from the network traffic to yield a data stream; content processing means adapted to calculate a digest for the content in dependence upon a source address of the content and identify the content as safe if the calculated digest matches a stored digest; scanning means adapted to scan content for viruses; and, a stream manager adapted to pass content that is not safe to the scanning means.

According to a second aspect of the present invention, there is provided a method for reducing the load on scanning means adapted to scan content for viruses, the method comprising the steps of: extracting content from the network traffic to yield a data stream; calculating a digest for the content in dependence upon a source address of the content; identifying the content as safe if the calculated digest matches a stored digest; and, passing content that is not safe to the scanning means.

The present invention allows more effective use of the available resources by ensuring that only necessary scans are performed. In this respect, content is considered to be safe if a virus scan is not required. It should be understood that the term safe does not necessarily imply that the content does not contain malware. Indeed, it may indicate that the malware has already been identified and for this reason it is not necessary to (and thus safe not to) scan the data. Accordingly, a number of measures may be employed to ensure that the AV scanners only perform required tasks.

To form the data stream, content is extracted from the payload of the data packets passed over the network. The headers and other network information are stripped from the traffic so that the data stream is able to provide entire pieces of content (that may stretch across a number of packets) in uninterrupted form. The data stream may then require further processing to be carried out in order to yield the content in a form suitable for analysis. Typically such processing will include a number of sequential transformation steps, each transformation generating a derived stream that is operated on by the next transformation, ultimately generating the content in a suitable form. Examples of such transformations are MIME decoding and decompression.

Digests are used to prevent the same piece of content from the same source being scanned twice. According to this technique, a digest is calculated for each piece of content that passes through the system. The results are then stored in a cache, along with the results of any scan. This allows the content to be recognised if it is seen

again. Consequently, there is no need to repeat a scan on content that has previously been seen. The digest may be calculated for an entire piece of content, or just a segment thereof. It depends upon the source of the content (for example, the IP address, domain or URL). This means that content will not be passed as safe unless it has been found to be trustworthy from its particular source in the past. This protects against the possibility of malware developers attempting to fool a digest system that depends only upon the form of the content, since they are unable to disguise its source. Note, a piece of content may refer solely to an item such as a file, or may also apply to an e-mail which in turn carries a file as an attachment.

The source may be defined in a number of ways. For example, a given domain may represent a number of IP addresses. As such, a single piece of content from a single source may have a number of different digests, each digest depending upon the definition of source used to calculate it.

In the case where content is compressed, digests are preferably calculated on the content in both compressed and decompressed form. This is advantageous when the same piece of content may be compressed using a variable seed or other random element such that the compressed version is different in each time the original content is compressed.

Download Managers are increasingly used to enhance download speeds. Typically, a Download manager will do this by requesting different parts (or segments) of the same piece of content from various sources. However, conventional scanners would operate on the different segments separately, and a full scan of the content would thus never be completed. In this context, where Download Managers are used instead of, or in conjunction with, conventional download programs (such as web browsers), the calculation of digests for segments of a piece of content is particularly advantageous, allowing potentially dangerous content to be identified in a situation where conventional scanning techniques would fail.

In a preferred embodiment, the present invention is capable of recognising when a segment of content is being downloaded (for example, when a Download Manager is in use) and only passing such segments to the scanners if they meet certain criteria. When Download Managers are in use, these programs request segments of a piece of content as opposed to the entire piece of content. These segments are requested over web protocols (HTTP) through the use of the Byte Range fields in the HTTP protocol Methods, where these Byte Ranges specify the segment of the file being requested. The present invention may detect this use of byte ranges in the HTTP Methods to determine that the content being requested is only a part of a file, and will then use the parameters detected in the Byte Range to determine whether to send to this piece of

content to the scanners. As stated previously, scanners are typically unable to scan content segments for viruses effectively, and accordingly the present invention does not pass these segments to the scanners. However, it may be the case that certain segments (such as the start or end of a file) can be usefully scanned and the present invention may therefore be programmed to send content segments that fulfill predetermined criteria to the scanner. For example, the criteria may stipulate that the first and last segments of a piece of content (or segments specified by a certain byte range) are passed to the scanners.

The present invention may also reduce the load on the AV scanners through the use of override lists. These lists contained allowed and denied sites. When data is received the source and destination are checked against the override lists and data is not sent to a scanner where it is found either to come from a trusted source (the data is automatically allowed) or from a source known to be a source of infected content (the data is blocked). Though the override lists may be generated offline, they are preferably generated dynamically as the system operates. In this way when a piece of content is determined to be infected, the source of that content may be added to the override list. Various configurable thresholds or settings may be adopted to control this automatic updating of the override list. For example, the addition or subtraction of a source from the override list may depend on the number of infected pieces of content seen from that source or on the time period that elapsed since content from that source was last found to be infected.

The majority of content flowing through the Internet is usually free from malware, and digests are thus able to reduce the traffic load on the scanners significantly. However, in certain circumstances, such as the outbreak of a new mass-mailing virus, the amount of infected content passing through the network will be considerably higher than is usual, and such malware may use polymorphism and other techniques to change the form of the infected content each time it replicates, thereby making the use of digests redundant. For this reason, the present invention is preferably also capable of using variable pattern matching, variable digest, and traffic anomaly detection techniques to detect known pieces of malware. If these techniques detect malware then content is marked as infected and not passed to the scanners, thereby reducing further the load on the scanners.

Preferably, a subscriber is able to specify that certain content types (such as executables) should never be transferred to or from their network. Data is not passed to the scanners if the content type is one that is automatically blocked by such subscriber preferences.

In a preferred embodiment of the present invention, other content security services (such as anti-spam and anti-phishing) are performed on the data in parallel to the AV service. If any one of these services returns a result that indicates that the data is to be blocked, then the present invention ceases the AV processing (thus reducing
5 the load on the AV scanners).

Brief Description of the Drawings

An example of the present invention will now be described in detail with reference to the accompanying drawings, in which:

10 Figure 1 is a simplified schematic representation of a Content Security Gateway (CSG);

Figure 2 is a flow diagram illustrating processing steps taken in accordance with one embodiment of the present invention;

15 Figure 3 illustrates function offloading and digest calculation according to a preferred embodiment of the present invention; and,

Figure 4 illustrates the components upon which processing functions are executed in a preferred embodiment of the present invention.

Detailed Description

20 As will be clear to one skilled in the art, the present invention may be implemented on a number of platforms (including a conventional PC). However, the preferred embodiment of the present invention exploits the capabilities of a dedicated hardware analysis device such as the Content Security Gateway (CSG) devices described in the Applicant's co-pending British patent application nos. 0523739.1 and
25 0522862.2. The CSG is capable of simultaneous performance of a number of content processing services on data sent and received by a large number of subscribers. These services include Anti-Virus (AV) capability and a variety of other content processing options (such as Anti-Spam and Anti-Phishing). Each service may be customised for each subscriber (for example, a subscriber may not have signed up for
30 anti-spam or may specifically request that web pages are not checked for phishing).

Figure 1 shows a broad schematic outline of the composition of an example of a CSG. Network Ports 100 receive data packets from any type of network. Network Traffic Processor 110 then identifies the transport protocol (such as TCP) used by the data, and extracts the payload from each data packet and combines it with others in the
35 same communication to yield a data stream. By extracting the payload in this way, a continuous flow of content (the data stream) is provided to the rest of the CSG, allowing content level processing of the traffic. In this way, a full piece of content (which may

have been spread across a number of network data packets) may be analysed by the CSG.

5 The CSG terminates TCP connections locally within itself. This means that instead of a TCP connection forming end-to-end between the subscriber machine and a destination machine, one connection forms between the subscriber and the CSG, and a second forms between the CSG and the destination machine. When a new flow using TCP is detected, and the CSG determines it belongs to a subscriber, at this point the two connections are set-up. Note, the session layer protocol (e.g. HTTP) is still end-to-end, although the CSG may manipulate information passed over this session. The CSG
10 may operate the TCP termination in the manner of a conventional network proxy (e.g. each connection utilises distinct network and link layer addresses), or in a transparent manner such that these link layer and network layer addresses are identical on the pair of TCP connections.

The same "transparent" approach is used for UDP and other protocols.

15 The termination of these TCP connections permits the CSG to modify content as it passes between end-points, ensuring that any changes to the content made by the CSG do not cause communication problems. If the TCP connections were still end to end, as the CSG modifies the content, the acknowledgement functionality of TCP would cause problems, as the information sent by one party would be different to that received
20 by the other (as the CSG has modified it), causing the session to fail and be aborted by the devices.

It should be recognised that the data stream, while continuous, will contain discrete pieces of content to be processed. For example, each file within the stream may be treated as a separate piece of content.

25 The data stream is then passed to a Streams Manager 130. Further information may also be passed to the Streams Manager 130, such as: a stream ID, a subscriber ID, network layer source and destination address, a policy for the stream including which services are to be operated (for example, AV service enabled), and the configuration of the or each selected service (for example, instructions to scan all traffic or block certain types of applications). The Content Processor Controller (CPC) 120 is
30 also illustrated in Figure 1. The CPC 120 collates results from the services performed by the CSG and effects the ultimate decision as to whether to block or allow the subscriber's access to the data.

The CSG contains a number of content engines. The content engines may
35 broadly be categorised either as hardware content engines 150 or software content engines 140. In this particular embodiment, the hardware content engines are Generic Content Engines (GCEs) 150 optimised to perform various content processing tasks.

An example of a GCE 150 is described in Applicant's co-pending British patent application no. 0522862.2. The GCEs are extremely fast at performing the tasks for which they are designed. The software content engines 140, referred to hereinafter as CESofts, may comprise conventional computer platforms capable of operating conventional software (such as a conventional AV scanner). It should be recognised that each CESoft 140 provides a flexible option and that the tasks undertaken by the GCEs 150 in the following description could also be undertaken by one or more CESofts 140.

When the Streams Manager 130 receives a new stream it passes it to the appropriate GCE (or GCEs) 150 to identify the data protocol used (for example, HTTP, SMTP, P2P) and to decode the identified protocol (step 202 in Figure 2). This identifies the application for which the network traffic is intended (such as web browsing or e-mail). During this decode the program used by the application will be identified if such an identifier exists in the stream. For example, HTTP streams usually contain a 'User Agent Field' that indicates which program generated the stream (such as a specific web browser or an update utility such as WindowsUpdate Manager).

In the case of SMTP (for example), the protocol decode will also yield the IP addresses of the source of the information (step 204 in Figure 2). This source information, along with the source IP addresses extracted by the NTP are then sent to the CPC 120 by the GCEs 150 (via the Streams Manager) and used in a check against a number of Realtime Blacklists (RBLs) (step 206) to determine if the stream originates from a source deemed to issue malware or inappropriate content (step 208). If the source is suspected of issuing such content then the stream is blocked (step 210), and no further work is undertaken on this stream (thereby eliminating an unnecessary load on the AV scanners). Additionally, the CPC 120 contains some defined override lists that can be set to ensure the stream is always propagated, or always blocked, again ensuring no unnecessary load is placed on the AV scanners. These override lists can be specified down to a per subscriber level.

If the stream is not deemed to originate from a malware or inappropriate content source then it is operated on further by the GCEs 150 in order to identify what type of content, if any, is being carried in the stream (step 212). The GCEs 150 then perform a lookup if this content type against the service settings. The service settings indicate a service mode for the content type, the result of which is that traffic is allowed, blocked or scanned (step 214). If the result is:

- Allow: the stream is then released onto its destination (step 216);
- Block: the CPC is informed and the CPC then blocks the stream (step 210);

- Scan: the content must be sent to one or more appropriate scanners based on the content type, application (as indicated by the protocol), and program.

When a stream is blocked from reaching a subscriber various other actions may occur, and these may be in dependence on the subscriber's preferences. For example, a block page may be transmitted to a subscriber when HTTP data is blocked.

If the service mode is scan, then the GCEs continue to process the stream in order to capture the piece of content (for example a file or web page). Whilst capturing the content, hereinafter referred to as the derived stream, the GCEs also calculate a digest of the content. Moreover, if the content type is compressed (step 218), and the compression format is one the GCE can decompress, the GCE then decompresses the content (step 220) yielding a new stream bearing the decompressed form of the content.

At this juncture in processing the following information about the stream is now available to the present invention:

- protocol type (indicating the application for which the content is intended);
- content type (in particular whether the content is executable); and,
- program type (such as a specific web browser or utility).

This information allows a threat profile for the content to be established. The present invention makes optimal use of the resources available to it by using the threat profile to send the stream and/or content to the most appropriate scanner available. In the preferred embodiment, the scanners to be used are determined by means of a simple algorithm (step 222). This algorithm is not fixed and may vary over time as the number and type of scanners vary, but an example is shown here below:

- 1) Look up the content type in a table. If the file is of the type 'image' then send to an Image Scanner to be scanned.
- 2) If the content type is ASCII or HTML and does not contain active content (such as scripts or specific HTML tags), and is carried over SMTP protocol, then send to an anti-spam service for checking.
- 3) If the content is executable and is carried over HTTP then send to a Web Threat Scanner.
- 4) If the content is script based and is destined for a known web application carried over HTTP then send to a Web Script Scanner.
- 5) If the program used is known then send to the scanner which checks for exploits (vulnerabilities) of this program.
- 6) If the content has no active parts, and the application is web browsing, then send to an anti-phishing service.

This scanner selection algorithm is usually implemented by a simple lookup in a database using tuples of protocol, content and program where each can be wildcarded. The result of the algorithm indicates which scanners are to be used, and whether they operate on the stream, the content (derived stream) or decompressed content (derived stream), and these streams and scanner instructions are then sent to the relevant AV scanners.

The GCE now sends the stream and scanner instructions to the relevant scanners (step 224). The scanners are implemented both in hardware on the GCEs and in software on a general use platform using standard PC components (such as a CESoft 150) that accepts industry standard software. A piece of software, known hereinafter as the Scanner Controller (SC), allows a plurality of software scanners to appear as one. If the stream is sent both to the SC and to hardware GCE scanners then the CPC is instructed to await results from both sets of scanners before data is finally blocked or released to reach the subscriber (step 226).

Similarly, if stream is also scheduled to be processed by other services (as well as the Anti-Virus Service), such as the Anti-Spam Service or Anti-Phishing Service then the CPC is informed of this activity so that a release decision is not made before the results of all the separate scheduled processes can be combined. However, note that early block decisions can be made if a single result requires a block. In such a situation incomplete tasks may be terminated immediately.

The presence of multiple GCEs in the invention allows the tasks of decompression and digest calculation to be performed in parallel. Similarly the multiple GCEs permit pipelining such that multiple streams can be processed in parallel.

The SC and GCE scanners then return to the CPC the results from the scanners used on the content, and the CPC then releases or blocks the content accordingly.

Scanners that may be used in accordance with the present invention include:

- Image Scanner: images are non executable content which can only form malware if they contain an exploit crafted for a specific vulnerability in an application. The number of these vulnerabilities is small (measured in tens and perhaps hundreds) and the exploit must be of a fixed (i.e not polymorphic or metamorphic) nature. As a result, the image scanner is typically implemented in accelerated hardware optimised for pattern matching, or in a targeted software pattern matching scanner. In the case of software the size of the pattern database is minimised in order to increase speed;
- Web Script Scanner: where a specific program is running, such as Internet Explorer, it can execute content within the defined limits of its architecture of this program (e.g. Java Security Architecture). This scanner is designed solely to

detect this executable content type that is relevant to this controlled security environment.

- 5 - Web Threat Scanner: this is a conventional third party industry scanner, where this scanner is only configured to deal with threats residing on the web such as adware, spyware Trojans etc. This scanner has less work to do and therefore operates faster than conventional scanners in their normal mode.
- 10 - Web Browsing Program Scanner: a program such as Internet Explorer has a number of vulnerabilities that are unique to that program, and this scanner is designed solely to protect against such threats. Again these threats are typically static and simple hence this scanner is typically a fast pattern matcher
- 15 - Spoofing Scanner: a piece of malware may attempt to masquerade as another legitimate program in order to avoid detection. A spoofing scanner dedicated to validating that the stream data of an identified program is in fact being generated by that program may therefore be included in the present invention.
- 20 - E-mail Application Scanner: an application such as e-mail has a number of vulnerabilities that are unique to that application, and this scanner is designed solely to protect against such threats. Again these threats are typically static and simple hence this scanner is typically a fast pattern matcher. In this case it is not always possible to identify the particular e-mail program in use (e.g. Microsoft Exchange Server) so the threats for all e-mail program may be combined together in a single scanner.
- 25 - Instant Messaging Application Scanner: an application such as IM has a number of vulnerabilities that are unique to that application, and this scanner is designed solely to protect against such threats. Again these threats are typically static and simple hence this scanner is typically a fast pattern matcher.
- 30 - Conventional industry scanner: this scanner is used in exceptional conditions where a piece of content, stream or application is unknown, unusual or suspicious. These scanners may be obtained from well known third party organisations such as Symantec, Kaspersky, and FRISK.

35 As stated previously, the scanners implemented in software are resident on a platform using standard PC components such that they accept industry standard software. Though this has the benefit of allowing known, established, and trusted AV scanners to be incorporated in the present invention, performance is consequently limited by the platform itself (no pipelining or parallelism in the hardware) and the software (which is not designed for high throughput). These scanners are grouped together to provide a single interface to the system, and made to appear as a single scanner by a software module known as the Scanner Controller (SC). Figure 3

illustrates the SC 340, which coordinates a number of scanners 342 and incorporates a result content store 344 to combine the results of the various scanners 342.

Other scanners may be introduced as and when needed for the mode of use of the invention. For example, if the invention is deployed in an environment where FTP traffic is prevalent then a scanner specifically designed for FTP may be included. Similarly, if the content being passed through the CSG features a large degree of content of a particular type then a specific scanner for that type of content is introduced (for example, if music downloads are common then a scanner which scans this type of content for known exploits may be introduced).

As would be clear to one skilled in the art, specialised scanners of this type may be implemented in a number of ways. For example, they could take the form of conventional third party scanners with limited configuration pattern matching databases. Alternatively, it is possible that the scanners will be developed specifically for use in the context of the present invention.

In order to reduce the load on the platform running software AV scanners, the other services (such as anti-spam and anti-phishing) are operated on separate platforms.

The scanners are selected so that their performance and characteristics complement each other. For example, conventional scanners are relatively good (i.e. fast) at scanning large pieces of content, and relatively poor (i.e. slow) at scanning small pieces of content (due to the overhead of opening a file). However, web browsing includes many very small image files. To counter this, a specialised image scanner (as described above) may be incorporated into the invention. As images provide a low threat profile such a scanner is relatively easy to implement. In addition, conventional AV scanners can also be slow at scanning text files, and for this reason a specialised web script scanner may be incorporated (along with the additional anti-spam and anti-phishing services).

The present invention is also capable of improving performance by offloading tasks typically done by conventional AV scanners to dedicated hardware units. For example, as mentioned above, the GCEs may decompress the data before it is sent to a scanner.

Figure 3 conceptually illustrates the flow of a data stream through the CSG and in particular shows the way in which functions are offloaded from the conventional scanners. Figure 4 shows which components of the CSG host the various tasks illustrated in Figure 3:

As detailed previously, once a stream is received from the NTP, the content, program, and protocol type of the stream is identified, and the protocol is decoded 300.

A preliminary check 302 of the IP address against block and override lists is made to ensure that further processing is required. The content stream is then decompressed (if required) and a digest is calculated 304.

In addition to decompression, a number of further compute intensive functions
5 may be performed before the derived stream is passed to the scanners. The functions available are typically implemented as dedicated hardware blocks in a GCE, where these functions can be programmed in for each available combination of protocol (i.e. application), content and program. Preferably, the scanners are aware that these functions have been offloaded so as to ensure that the scanners do not unnecessarily
10 repeat these tasks. Since the offloaded functions are performed on high performance hardware and software building blocks and the scanner is no longer required to perform these tasks, overall performance is significantly improved.

A non-exhaustive list of possible function offloads 320 includes:

- 15 - Pattern Matcher (PM) 324: the PM is programmed with a set of patterns which are searched for across the stream/derived stream/content, and a set of results indicating the following are returned: number of matches, offsets in stream where found. The patterns are defined as per conventional Regular Expression matcher found in the PERL language and are of the same format, or similar industry standard pattern matching languages.
- 20 - Attribute Checker 326: this function checks each content stream for a series of attributes against a set of defined thresholds. For example, the size and format of the file header may be checked. The function may also check for a number of attributes across all streams, again checking against defined thresholds.
- 25 - Instruction Decoder 328: where the content identification check performed indicates the file is for a specific hardware platform (e.g. Windows executable) this function then performs a count of each instruction found within the data and code segments of the file, and checks the densities of these values across the file. The decoder will then report any unusual results, which may include the most commonly used op-codes or byte values (particularly if their density
30 exceeds any threshold values) and also any sudden change in density of such features.

The above is not an exhaustive list, and a number of other functions 329 could be offloaded in accordance with the present invention. For example, a Statistics and Logging function offload 327 may be included (storing such details as the most
35 frequently occurring viruses, the source of most viruses, and the type of viruses being scanned). Similarly, a Sample Capture function offload 325 may also be utilised,

allowing samples of viruses to be collected (without imposing a processing burden on the partial scanners) for subsequent further analysis.

The function controller 322 coordinates the actions of the various function offloads. In particular, note that a particular result of one function may cause the function controller 322 to call another function. The function controller is programmed to forward the results of the various functions, and the streams on which they operate, to one or more partial scanners 332, coordinated by a Partial Scanner Controller (PSC) 330. The manner in which the functions are utilised is configured through a programmable interface (API) 338. The API 338 is used to configure parameters for each function, the outputs each function generates, and how the function controller 322 should process these outputs. For example, the API 338 may be used to specify that if the Pattern Matcher 324 detects a match of a certain type, then a certain portion of the streamed content is sent to the Instruction Decode 328 function, or that a portion of the streamed content is sent to a defined partial scanner 332. This configuration information is stored on a user-defined function controller configuration 336. The partial scanners may also have access to the API. For example, after acting on streams and results received from the function offloads, the Partial Scanner 332 may then request, through the API, that further functions are executed. Moreover, the Partial Scanner 332 may be able to control the manner in which they are executed by passing parameters to the Function Controller so that the function is operated and returns results in a defined manner. The PSC 330 operates in a manner analogous to the mode of operation of the SC 340, including partial scanners 332 in the place of scanners 342, as well as a result content store 334. The partial scanners 332 are adapted to interpret the results of the offloaded functions. Examples of such partial scanners include:

- 25 - PM scanners: such scanners use pattern matching regularly and the results of the searches are presented to the scanners; the scanner simply uses these results as an indication of infection, and if no infection is present the scanner then moves on to undertake its remaining checks;
- 30 - Heuristics based scanners: such scanners use the presence of certain attributes to determine whether content is malicious. In this case the Attribute Checker function has checked for the presence of these attributes in advance and returned a result summary to the heuristics engine simply to interpret these results;
- 35 - Instruction Distribution Scanners: such scanners utilise these checks to look for anomalies in code which may indicate presence of 'foreign' code (i.e. a virus) in a file.

The partial scanners may be implemented by configuration of conventional scanners to operate with the function offloads. For example, a conventional scanner may be compiled with a pattern database containing only a single entry. In this case, the pattern matching function offload performs the pattern search and the results are passed to the PSC. The partial scanner then undertakes a further (redundant) pattern search, but this runs quickly due to the small size of the pattern database. The partial scanner then performs the scanning functions for which no function offload is available. Alternatively, the partial scanner may comprise a conventional scanner adapted not to use its pattern search engine.

In one example, a partial scanner is compiled without various modules such as the scanning of image files. Accordingly, image files are not sent to this partial scanner and as such the partial scanner need not undertake a full range of functions.

In another example a partial scanner is designed to operate specifically with a Function Offload such as the pattern matcher 324. The pattern matcher 324 will generate a set of results indicating that it has detected a number of patterns at specific locations within the streamed content, therefore allowing the partial scanner 332 to analyse the parts of the content identified by the pattern matcher. In this manner the partial scanner and pattern matcher provide parallelism, thereby increasing the throughput of the overall scan operation.

In combination, the partial scanners and the various offloaded functions essentially provide different elements of a single overall AV scanner providing a defence against all types of malware. Each element performs one or more of the AV techniques required to offer this comprehensive service. The combination of a distributed set of partial scanners each with a specific purpose where each has its compute intense function offloaded that provides the overall high throughput of the scanner.

As shown in Figure 3, the results of the partial scanners are collated by the PSC. The results are then combined 360 with those of any other services 350 and the CPC takes action (for example, blocking or allowing data) accordingly. Figure 4 shows that results analysis 370 also occurs at the CPC.

Figure 3 also illustrates the creation one or more digests of the content stream before the stream is passed to any service (including AV) for action. The digest acts as a unique identifier, or fingerprint, for the content. It may be used to identify content that has previously been scanned (for example, as part of a separate transmission), and consequently to prevent unnecessary repetition of a task that has already been performed.

Figure 3 illustrates the use of digest in combination with an SC 340. One skilled in the art will readily understand that the principles of this use may equally be applied to

a PSC or, indeed, to other services (such as Anti-Spam or Anti-Phishing). As illustrated in Figure 3, the digest is first calculated and then transferred to the SC 340, which contains means 346 to receive the content. The SSC 340 operates a cache 348 of scanned pieces of content, storing the result of each scan within the cache 348. The cache 348 is indexed by the digest of the content. Note the cache 348 is flushed or cleared each time the scanner signatures or definitions are updated.

When a piece of content arrives at the SC, the SC first looks up the content digest in the cache. If the entry is not present then the SC 'connects' this stream to the appropriate scanner(s), and returns the result(s) to the CPC. The digest entry is then added to the cache with this scan result.

If, on the other hand, the digest has been previously stored then the SC takes the cache results and returns these result(s) to the CPC without undertaking a scan.

It is important to bear in mind that multiple digests may be created for a given piece of content. That is, digests may be calculated for one or more segments of a larger piece of content. In particular, digests may be updated as additional data is received. The choice whether to use each of these multiple digests may be static (i.e. always or never) or selected on the basis of application type. For example, the application WindowsUpdate transmits large pieces of invariant content to millions of users, and it may therefore be beneficial to recognise the content at the earliest available stage, thereby reducing unnecessary load on the resources of the CSG. As such, a digest calculated on the basis of an initial content segment may be deemed appropriate for this application. Note, when using partial digests there are multiple results supplied to the SC at each juncture when a digest is available, and the stream available up to that point is also sent to the SC.

The advantages of calculating digests not only on an entire piece of content but also on segments of the content are also apparent when the content does contain malware. Consider the case where a user attempts to download a large file and a virus is only discovered in that file once the majority of the file has been transferred to the user. At this point, the CSG will prevent transfer of the remainder of the file. However, if the user were to attempt to re-start the download then only the last section of the file would be requested. This can only be recognised if a digest had been calculated on that segment of the content.

Analysis of content segments also proves valuable in, for example, the context of download managers or peer-to-peer file sharing. In these cases, a single large piece of content is downloaded in segments from a variety of sources. In this case, each data stream will only contain segments of the content. Typically, a virus scanner cannot perform without access to the entire piece of content, and there is therefore no reason

to scan the segments individually and the present invention will therefore not pass the content segments to the virus scanner. There may be exceptions to this rule, where certain segments can be scanned (for instance, the start and the end of the file may betray the presence of a virus) and the present invention may therefore be adapted to
5 identify certain segments from a piece of content and pass these on to the virus scanners. In the case of HTTP, the segments may be identified through use of the HTTP protocol Methods.

The digest is dependent upon the source of the content. The source may be defined as, for example, the IP address, the domain or the URL and digests may be
10 calculated for each definition of source that is adopted. This provides a number of advantages. For example, it is theoretically possible to introduce malware to a piece of content in such a way that a digest calculated for it is not affected. A hacker may try to exploit this by altering content that has been previously scanned in this way since content having a known digest is not scanned. However, if the digest also depends
15 upon the source of the content then this evasion technique will not prevent the content being scanned, as the content will now originate from a different source (i.e. the hacker's website rather than the original source).

There is a small probability that the digests of two unrelated pieces of content will be identical, due to the manner in which digests are calculated. In order to
20 overcome this problem the present invention may calculate more than one digest for any given piece of content (or content segment), with each of these digests being calculated using a different digest calculation algorithm. For example, digests may be calculated using both MD5 and SHA-1.

A source-dependent digest also finds particular utility in the context of content
25 segments. As mentioned above, different segments of a single piece of content are often downloaded from a variety of sources. A scan on each individual segment may not be enough to identify malware, so a scan on the entire piece of content is preferably performed. If the content in its entirety is found not to contain a virus this does not necessarily indicate that each segment is virus-free, as some segments may have
30 originated from an infected version of the content while others did not. It is therefore necessary to scan the entire piece of content from a single source to establish that each content segment from that source is not infected. Digests that are dependent on the source are able to indicate whether or not content segments originate from a source for which the entire piece of content has been found to be virus free. Once the entire
35 content has been found to be virus free from a number of sources it is possible to download any segment from any of those sources, without the requirement for a scan.

As such, a piece of content may still be obtained in segments originating at a number of different sources, thereby maintaining the advantage of Download Managers.

The digests calculated for a given source may only be valid for a limited period so that if content from that source is adapted to contain a virus then this is recognised.
5 Moreover, if any content from a specific source is found to contain malware then the present invention may be adapted to invalidate all digests calculated for content from that source.

Certain content may only be allowed if its digest indicates that it originated from a trusted source. In this way, content can be identified and trusted both on the basis of
10 its origin and on the basis of an earlier scan. In this way, certain pieces of content may be allowed only from certain sites (for example, a Microsoft update may only be allowed from an official Microsoft site).

For simple types of malware that replicate in such a manner that each copy of the malware is identical to all others (typically worms or Trojans) then the use of digests
15 is an effective method to reduce traffic sent to the scanners; the invention processes many pieces of the malware which are identical and after performing a single scan the digest computed is then used to detect all further instances of this invariant piece of malware, which are not sent to the scanners. However, more complex forms of
20 malware tend to vary each time they replicate. For example, mass-mailers spread by infecting a machine then reading the address book of the user logged onto the infected machine. New copies of the malware are then sent to recipients found in the address book (consequently each mail will be different as each address book is different). In addition to this, mass mailers will typically also change other fields inside the e-mail
25 such as the subject line or phrases inside the e-mail body and the e-mails sent therefore differ each time the malware spreads. Moreover, any file sent in an e-mail carrying malware may vary in each replication through the use of polymorphic or metamorphic replication techniques used by malware writers. Therefore in such circumstances use of digests computed on the entirety of the e-mail or any attachment are not effective in
30 reducing traffic sent to the scanners.

To counter the threat of variable malware, a number of detection techniques
30 may be adopted to identify and prevent the spread of such content.

For example, variable pattern matching techniques are known in the art. According to such techniques, a number of samples of particular malware may be collected as it spreads. A comparison of these samples will typically show some
35 commonality between the different instances of the malware (for example, common words or phrases). As such, a pattern may be identified that indicates an instance of the malware. This pattern may be a simple word or phrase, or a combination of words

or phrases (for example, word A, followed by a variable number of spaces, followed by word B, followed by a variable number of spaces followed by word C), and can be detected by conventional complex variable pattern matchers. An example of such a pattern matcher is the GCE described in Applicant's co-pending British patent application no. 0522862.2 which loads the patterns into a high speed hardware engine for high throughput detection of patterns. Other forms of high speed complex pattern matching are pieces of software running on general microprocessors, an example of which is the open source AV scanner CLAMAV which simply looks for the patterns it is loaded with; this CLAMAV pattern matcher running on a general purpose microprocessor is not as fast as the GCE hardware implementation but by limiting the pattern database used it still provides a performance benefit over a conventional scanner. In this case the pattern database is tailored to the highest traffic loads at any period of time. For example, only patterns for malware currently propagating the internet (known as active in the wild) are loaded into the database, as opposed to all malware patterns that have ever been known.

The present invention may also use traffic anomaly detection to identify outbreaks of malware. For example, a worm may propagate over the TCP protocol, and have an exceptionally high replication rate. Consequently, the levels of TCP traffic on certain TCP ports will increase dramatically in comparison with the usual amount of traffic on those TCP ports. As such, an effective method of identifying such malware is to compare the usual level of traffic on a defined TCP port against the level of traffic over a defined period of time. For example, if the average transfer rate for TCP over a port is 100 files per second over a 60 minute period, the content may be determined to be malware (and thus not sent to the scanners) if the detected rate is greater than N times this average rate (where N may be specified for each particular instance of malware). Similarly, packet rates can be compared and if deemed to be malware the content carried over those packets is not sent to the scanners. There are a number of further metrics that may be used in addition to simple traffic levels. For example, the number of instances of files sent of a certain size, or the rate at which TCP connections are opened and closed. With each metric the typical (or usual) traffic level is compared to the current traffic level.

In the alternative, the present invention may make use of traffic anomaly detection to identify large amounts of legitimate content that need not be scanned. For example, a large supermarket chain may send out a mass-email to its customers which is tailored to their shopping preferences. It is clearly disadvantageous to scan every one of these e-mails. For this reason, the present invention may analyse the traffic flow for the source of the data, together with other attributes (for example, certain

expressions in the content, the size of the content, and the use of certain TCP or UDP ports). Through a comparison of these details with the typical behaviour of each source, a judgment may be made as to whether it is necessary to scan the content or not.

5 Many pieces of content passing over the Internet, or other public networks, are in fact subtly different forms of the same piece of source content; examples are spam messages, which account for over 75% of all e-mail traffic, where the source of the spam messages wishes to send the same piece of content to as many recipients as possible, but changes each incarnation of the message being sent so as to subvert anti-
10 spam filers operating in the network and at the recipients. The present invention may be arranged to combat such variable content by first operating a number of techniques which distil the content down into the 'core' content message (i.e. the characteristics that are invariant between each piece of content) that is being communicated, and then calculating a digest (referred to hereinafter as a 'variable digest') on this piece of core
15 content. Accordingly, differing content may have the same variable digest as long as the selected core parts are invariant

A number of techniques may be adopted to identify the invariant, core content upon which variable digests are calculated. In the case of a spam e-mail, these techniques may include the generation of MIME-decoded streams, HTML to ASCII
20 conversion, and textual parsing (this step being performed with knowledge of how spam e-mails are constructed). For example, the open source anti-spam detection system Distributed Checksum Clearing (DCC) identifies parts of an e-mail thought to be invariant (by removing variable parts such as the intended recipients, the white space in content, and the non-renderable content) and these may be used to calculate a variable
25 digest. Clearly, while the recipient address of a mass mailed spam e-mail is variable, certain other parts will be invariant (such as the purpose of the spam).

There are also image manipulation techniques that may be used to identify core content, and consequently to calculate variable digests. These include colour space techniques effective to remove colour and image re-sizing algorithms.

30 Variable digests may be calculated both on entire pieces of content and on content segments as required.

These variable digests will be used in combination with the fixed digests described earlier. The same approach can be used on parts of files (or attachments) that are thought to be invariant. Analysis of the latest malware trends is used to identify
35 which parts of the content are likely to be invariant. The invariant parts could be, for example, the file header or the last 4 kbytes of the file. According to this technique, the digest calculation algorithm will vary over time. The variable digests will be able to

detect malware and thus prevent content being unnecessarily passed to the scanners, thereby reducing the load on the scanners.

As with fixed digests, a number of variable digests may be calculated for each piece of content. Similarly, variable digests may be calculated for both the compressed
5 and decompressed forms of the content, and may depend on the source address.

The use of the override lists and the blocking of certain application types also reduces scanner load. A further reduction is available by allowing subscribers to implement a policy defining types of content to be blocked for a given application. For example, the subscriber may specify that all executable files are to be blocked when
10 using e-mail.

It may be that text-based content with no embedded active content (such as HTML with no active tags, or an ASCII text file) is not considered to be a virus threat. In this case, the stream is not acted upon by the Anti-Virus service at all but is instead passed to services that deal with, for example, social engineering attacks such as
15 hoaxes or phishing. Since these other services are not performed on the same platform as the AV scanners, the workload on these scanners is reduced.

CLAIMS:

1. A network security apparatus, comprising:
one or more network traffic processors adapted to receive network traffic and to
5 extract content from the network traffic to yield a data stream;
content processing means adapted to calculate a digest for the content in
dependence upon a source address of the content and identify the content as safe if the
calculated digest matches a stored digest;
scanning means adapted to scan content for viruses; and,
10 a stream manager adapted to pass content that is not safe to the scanning
means.
2. An apparatus according to claim 1, wherein the content comprises one or more
content items, each content item formed of a plurality of content segments, and the
15 content processing means is adapted to calculate a digest for each content segment.
3. An apparatus according to claim 1 or 2, wherein the content processing means
is adapted to decompress the content prior to calculating a digest.
- 20 4. An apparatus according to any preceding claim, wherein the content processing
means is adapted to recognise when a segment of a content item is received and to
identify such a segment as safe unless it fulfils predetermined criteria.
5. An apparatus according to any preceding claim, wherein the content processing
25 means is adapted to check the source and destination of the data stream and to identify
content as safe if the source or destination is contained in an override list.
6. An apparatus according to any preceding claim, wherein the content processing
30 means is adapted to identify a content type and to identify content as safe in
dependence on the content type.
7. An apparatus according to any preceding claim, wherein the content processing
means is adapted to search for a pattern in the content and to identify content as safe if
the pattern corresponds to a known piece of malware.

35

8. An apparatus according to any preceding claim, wherein the content processing means is adapted to detect distribution anomalies in the network traffic and to identify content as safe if such an anomaly is detected.
- 5 9. An apparatus according to any preceding claim, wherein the content processing means is adapted to calculate a variable digest for the content and to identify the content as safe if the calculated variable digest matches a previously calculated variable digest.
- 10 10. An apparatus according to any preceding claim, wherein the content processing means is adapted to perform one or more additional services on the content and to identify the content as safe in dependence on one or more results of the additional services.
- 15 11. An apparatus according to claim 10, wherein the additional services comprise one or more of anti-phishing and anti-spam.
12. A method for reducing the load on scanning means adapted to scan content for viruses, the method comprising the steps of:
- 20 extracting content from the network traffic to yield a data stream;
calculating a digest for the content in dependence upon a source address of the content;
identifying the content as safe if the calculated digest matches a stored digest;
and,
- 25 passing content that is not safe to the scanning means.
13. A method according to claim 12, wherein the content comprises one or more content items, each content item formed of a plurality of content segments, and a digest is calculated for each content segment.
- 30 14. A method according to claim 12 or 13, further comprising the step of decompressing the content prior to the step of calculating a digest.
15. A method according to any of claims 12 to 14, further comprising the step of
35 recognising when a segment of a content item is received and identifying such a segment as safe unless it fulfils predetermined criteria.

16. A method according to any of claims 12 to 15, further comprising the step of checking the source and destination of the content and wherein the content is identified as safe if the source or destination is included in an override list
- 5 17. A method according to any of claims 12 to 16, further comprising the step of identifying a content type and wherein content is identified as safe in dependence on the content type.
- 10 18. A method according to any of claims 12 to 17, further comprising the step of searching for a pattern in the content and identifying content as safe if the pattern corresponds to a known piece of malware.
- 15 19. A method according to any of claims 12 to 18, further comprising the step of detecting distribution anomalies in the network traffic and identifying content as safe if such an anomaly is detected.
- 20 20. A method according to any of claims 12 to 19, further comprising the step of comprises calculating a variable digest for the content and identifying the content as safe if the calculated variable digest matches a stored variable digest.
21. A method according to any of claims 12 to 20, further comprising the step of performing one or more additional services on the content and wherein content is identified as safe in dependence on results of the additional services.
- 25 22. A method according to claim 21, wherein the additional services comprise one or more of anti-phishing and anti-spam.


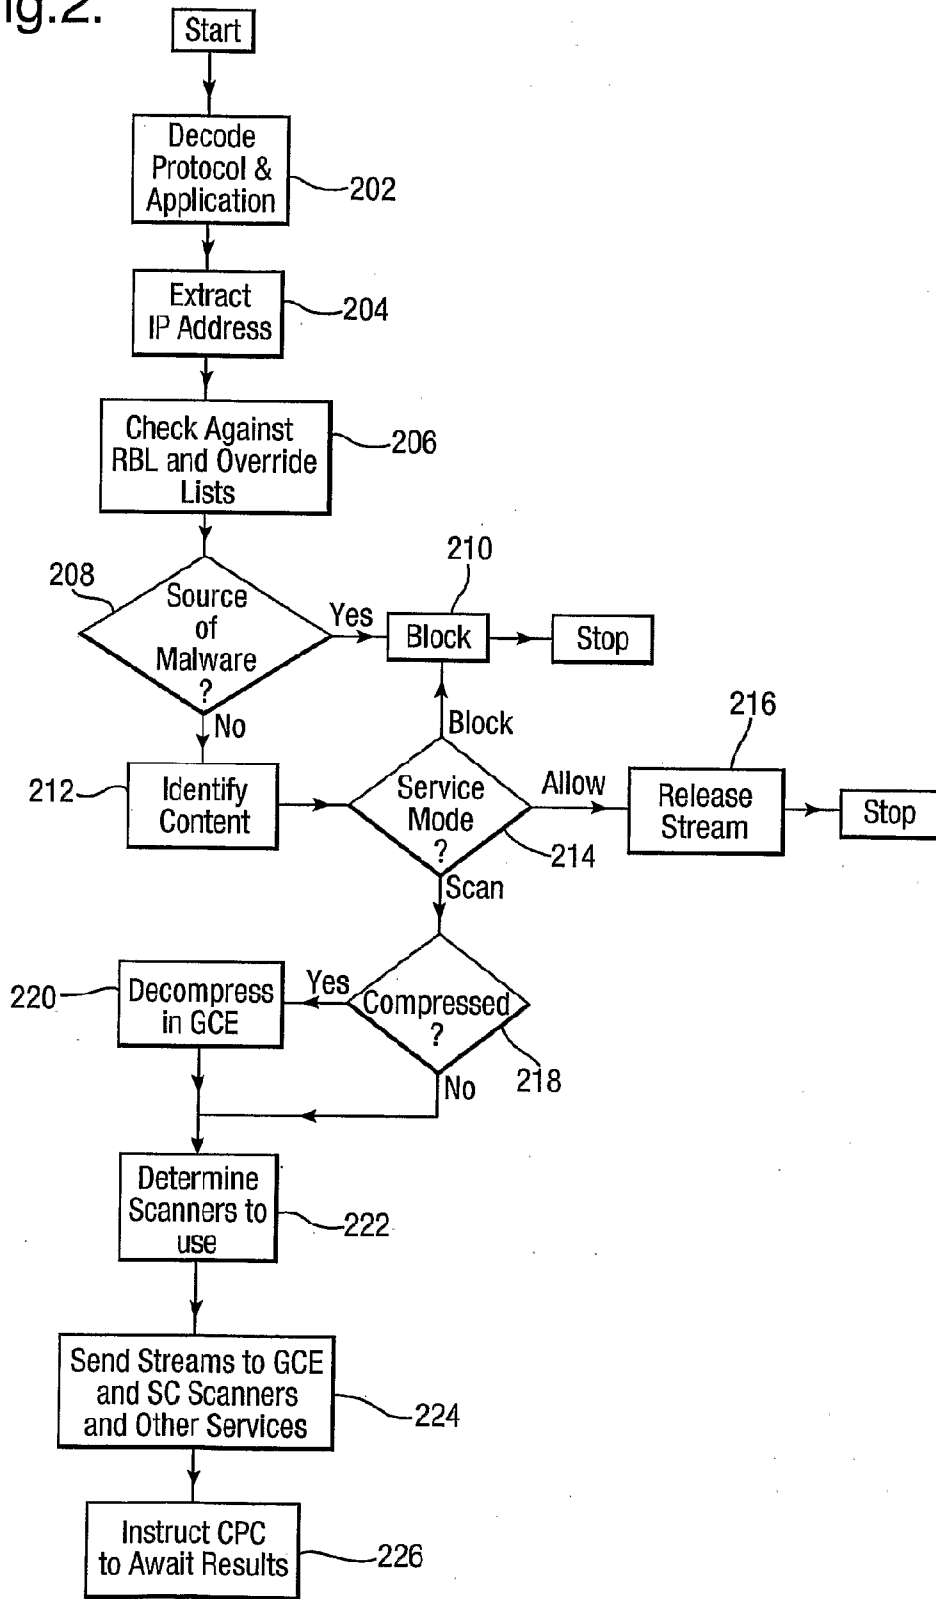


Fig.1.

Fig.2.






Fig. 3.




Fig. 4.

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2007/000901

A. CLASSIFICATION OF SUBJECT MATTER INV. H04L29/06 G06F21/00		
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 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.
X	US 6 993 660 B1 (LIBENZI DAVIDE [US] ET AL) 31 January 2006 (2006-01-31) column 5, line 64 - column 9, line 18; figures 7,8	1-22
X	US 2004/181687 A1 (NACHENBERG CAREY S [US] ET AL) 16 September 2004 (2004-09-16) paragraphs [0019], [0020] paragraph [0027] paragraphs [0030] - [0032] paragraph [0036] paragraphs [0039] - [0043]	1-22
A	WO 2004/053621 A (NOKIA INC [US]) 24 June 2004 (2004-06-24) page 16, line 25 - page 17, line 8 page 18, line 8 - page 21, line 4	1-22
-/--		
<input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> 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	Date of mailing of the international search report	
19 June 2007	28/06/2007	
Name and mailing address of the ISA/ European Patent Office, P.B. 5618 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Tx. 31 651 epo nl, Fax: (+31-70) 340-3016	Authorized officer RUIZ SANCHEZ, J	

INTERNATIONAL SEARCH REPORT

International application No
PCT/GB2007/000901

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	GB 2 417 655 A (STREAMSHIELD NETWORKS LTD [GB]) 1 March 2006 (2006-03-01) cited in the application page 10, line 6 - page 12, line 14 page 13, line 28 - page 15, line 35 page 17, line 33 - page 19, line 14 page 26, line 29 - page 27, line 9 page 3, line 4 - page 37, line 20	1-22

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/GB2007/000901

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
US 6993660	B1	31-01-2006	NONE	
US 2004181687	A1	16-09-2004	NONE	
WO 2004053621	A	24-06-2004	AU 2003302790 A1	30-06-2004
GB 2417655	A	01-03-2006	NONE	