

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
24 October 2002 (24.10.2002)

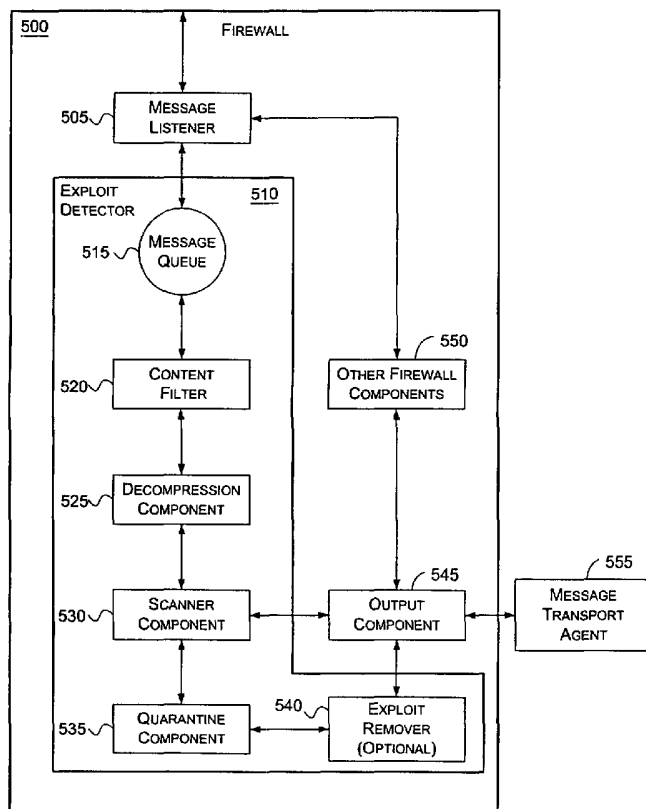
PCT

(10) International Publication Number
WO 02/084495 A1

- (51) International Patent Classification⁷: G06F 12/14, H04L 9/00, G06F 13/00
 - (21) International Application Number: PCT/US02/11685
 - (22) International Filing Date: 12 April 2002 (12.04.2002)
 - (25) Filing Language: English
 - (26) Publication Language: English
 - (30) Priority Data: 60/283,757 13 April 2001 (13.04.2001) US
 - (71) Applicant: NOKIA, INC. [US/US]; 6000 Connection Drive, Irving, TX 75039 (US).
 - (72) Inventor: SMITH, Gregory, J.; 3576 Sunnydays Lane, Santa Clara, CA 95051 (US).
 - (74) Agent: BRANCH, John, W.; Merchant & Gould P.C., P.O. Box 2903, Minneapolis, MN 55402-0903 (US).
 - (81) Designated States (national): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, OM, PH, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZM, ZW.
 - (84) Designated States (regional): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:** — with international search report

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR PROVIDING EXPLOIT PROTECTION FOR NETWORKS



(57) Abstract: A method and system for providing protection from exploits to devices connected to a network (500). The system and method include a component (520) for determining whether an encapsulation has been applied to an attachment and unencapsulating such encapsulated attachments, a component (525) that performs at least one decompression of the attachment when the attachment is compressed, a component that determines whether a header, body, and/or attachment of a message includes an exploit, and a component that holds and optionally cleans messages that include exploits. A device that receives messages that are directed to the network employs the components above to provide exploit protection for at least one of the messages.

WO 02/084495 A1



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.

**SYSTEM AND METHOD FOR PROVIDING EXPLOIT PROTECTION FOR
NETWORKS**

This application is being filed as a PCT International Patent Application in the name of Nokia Inc., a U.S. national corporation and resident, on 12 April 2002, designating all countries except the US, and claiming priority to U.S. Serial No. 60/283,757 filed 13 April 2001.

Field of the Invention

The present invention relates to computer network security, and in particular to exploit protection for networks.

10 **Background**

The Internet connects millions of nodes located around the world. With the click of a button, a user in one part of the world can access a file on another computer thousands of miles away. In addition, the Internet has facilitated the exchange of information in the form of electronic messages known as email. Although, first used to transfer short text messages, email can now be used to send digital pictures, sound files, documents, spreadsheets, executable programs, and other electronic files. Sending such files is as easy as attaching them to an email message and clicking a send button.

The ease of transmitting information, however, has also been exploited by some for other purposes. One of the first well-publicized cases of exploitation involved using email servers to propagate a program. Once an email server became "infected" with the program, it would begin sending email messages containing the program to other email servers that it was aware of. Like a virus, the program spread from email server to email server with amazing speed. Although the program did not erase files or harm data on the email servers, because of the volume of email messages sent by the infected email servers, the program caused retrieval of email messages from email servers to slow to an excruciating pace.

Now, the news reports virus-like programs (hereinafter "exploits") on an almost daily basis. Some of these exploits are relatively benign; others destroy data or capture sensitive information. Unless properly protected against, these

exploits can bring a company's network or email system to its knees or steal sensitive information, even if only a few computers are infected.

The most prevalent method for dealing with these exploits is to install virus protection software on every computer. As soon as a version of virus protection software is released, those seeking to circumvent the exploit protection software create new exploits that are not recognized by the virus protection software. This prompts virus protection software developers to create updates for their virus protection software to detect these new exploits. Previously, updating the virus protection software on each computer required obtaining the updates on disk and going from computer to computer to install the updates. Now, virus protection software allows computer users to download updates using the Internet. Since installing updates requires effort on the part of each user (or a computer support team), quite often the virus protection software is not brought up-to-date on every computer. In addition, there is often a significant delay between the introduction of a new virus and the creation and distribution of an update aimed to protect against the new virus. Additionally, virus protection software typically does not protect against the more general class of virus-like programs known as exploits.

Summary

In accordance with the present invention, there is provided a method and system for providing protection from exploits to devices connected to a network. The system and method include a component for determining whether an encapsulation has been applied to an attachment and unencapsulating such encapsulated attachments, a component that performs at least one decompression of the attachment when the attachment is compressed, a component that determines whether a header, body, and/or attachment of a message includes an exploit, and a component that holds and optionally cleans messages that include exploits. A device that receives messages that are directed to the network employs the components above to provide exploit protection for at least one of the messages.

In one aspect of the invention, the message includes a header, a body, and an attachment. The header includes at least one field having a defined size. When the size of data in the field is other than the defined size, an exploit is found. Exploits may also be found in the body and/or attachment of a message.

In another aspect of the invention, a client executing on the system determines when an update is available by polling servers associated with vendors of exploit protection software. When the client determines that an update is available, it automatically retrieves the update.

5 In another aspect of the invention, the system employs exploit protection software from at least two vendors to determine whether an attachment includes an exploit.

In another aspect of the invention, components of the system are implemented in software and may be included on at least one of a firewall, router, 10 switch, and traffic manager.

In another aspect of the invention an encapsulation includes at least one of Multipurpose Internet Mail Extensions (MIME), Base 64, and uuencode.

15 These and various other features as well as advantages, which characterize the present invention, will be apparent from a reading of the following detailed description and a review of the associated drawings.

Brief Description of the Drawings

FIGURES 1-3 show components of an exemplary environment in which the invention may be practiced;

20 FIGURE 4 illustrates an exemplary environment in which a system for providing exploit protection for a network operates;

FIGURE 5 illustrates components of a firewall operable to provide exploit protection; and

FIGURE 6 illustrates a flow chart for detecting exploits in accordance with the invention.

Detailed Description

25 In the following detailed description of exemplary embodiments of the invention, reference is made to the accompanied drawings, which form a part hereof, and which are shown by way of illustration, specific exemplary embodiments of which the invention may be practiced. These embodiments are 30 described in sufficient detail to enable those skilled in the art to practice the invention, and it is to be understood that other embodiments may be utilized, and

other changes may be made, without departing from the spirit or scope of the present invention. The following detailed description is, therefore, not to be taken in a limiting sense, and the scope of the present invention is defined by the appended claims.

5 In the following description, first definitions of some terms that are used throughout this document are given. Then, illustrative components of an illustrative operating environment in which the invention may be practiced is disclosed. Next, an illustrative operating environment in which the invention may be practiced is disclosed. Finally, a method of detecting and removing exploits is
10 provided.

Definitions

The definitions in this section apply to this document, unless the context clearly indicates otherwise. The phrase “this document” means the
15 specification, claims, and abstract of this application.

“Including” means including but not limited to. Thus, a list including A is not precluded from including B.

A “packet” refers to an arbitrary or selectable amount of data which may be represented by a sequence of one or more bits. A packet may correspond to
20 a data unit found in any layer of the Open Systems Interconnect (OSI) model, such as a segment, message, packet, datagram, frame, symbol stream, or stream, a combination of data units found in the OSI model, or a non OSI data unit.

“Client” refers to a process or set of processes that execute on one or more electronic devices, such as computing device 300 of FIGURE 3. A client is
25 not constrained to run on a workstation; it may also run on a server such as a WWW server, file server, or other server, other computing device, or be distributed over a group of such devices. Where appropriate, the term “client” should be construed, in addition or in lieu of the definition above, to be a device or devices upon which one or more client processes execute, for example, a computing device, such as
30 computing device 300, configured to function as a World Wide Web (WWW) server, a computing device configured as a router, gateway, workstation, etc.

Similarly, “server” refers to a process or set of processes that execute on one or more electronic devices, such as computing device 300 configured as a

WWW server. Like a client, a server is not limited to running on a computing device that is configured to predominantly provide services to other computing devices. Rather, it may also execute on what would typically be considered a client computer, such as computing device 300 configured as a user's workstation, or be
5 distributed among various electronic devices, wherein each device might include one or more processes that together constitute a server application. Where appropriate, the term "server" should be construed, in addition or in lieu of the definition above, to be a device or devices upon which one or more server processes execute, for example, a computing device configured to operate as a WWW server, router,
10 gateway, workstation, etc.

An exploit is any procedure and/or software that may be used to improperly access a computer through email. Exploits include what are commonly known as computer viruses but may also include other methods for inappropriately gaining access to a computer. For example, computer viruses are typically included
15 in an attachment to an email message. Some exploits, however, are contained in the header or body of an email message. For example, some exploits attempt to overflow buffers allotted for portions or all of a header or body of an email message. In bytes of data contained in the overflow, these exploits often contain executable code. This executable code is arranged in such a fashion as to be executed by the
20 host computer. The executable code may then improperly access data and/or execute unauthorized programs on the host computer.

Referring to the drawings, like numbers indicate like parts throughout the figures and this document.

Definitions of terms are also found throughout this document. These
25 definitions need not be introduced by using "means" or "refers" to language and may be introduced by example and/or function performed. Such definitions will also apply to this document, unless the context clearly indicates otherwise.

Illustrative Operating Environment

30 FIGURES 1-3 show components of an exemplary environment in which the invention may be practiced. Not all the components may be required to practice the invention, and variations in the arrangement and type of the components may be made without departing from the spirit or scope of the invention.

FIGURE 1 shows wireless networks 105 and 110, telephone phone networks 115 and 120, interconnected through gateways 130A-130D, respectively, to wide area network/local area network 200. Gateways 130A-130D each optionally include a firewall component, such as firewalls 140A-140D, respectively. The letters FW in each of gateways 130A-130D stand for firewall.

Wireless networks 105 and 110 transports information and voice communications to and from devices capable of wireless communication, such as such as cell phones, smart phones, pagers, walkie talkies, radio frequency (RF) devices, infrared (IR) devices, CBs, integrated devices combining one or more of the preceding devices, and the like. Wireless networks 105 and 110 may also transport information to other devices that have interfaces to connect to wireless networks, such as a PDA, POCKET PC, wearable computer, personal computers, multiprocessor systems, microprocessor-based or programmable consumer electronics, network PCs, and other properly-equipped devices. Wireless networks 105 and 110 may include both wireless and wired components. For example, wireless network 110 may include a cellular tower (not shown) that is linked to a wired telephone network, such as telephone network 115. Typically, the cellular tower carries communication to and from cell phones, pagers, and other wireless devices, and the wired telephone network carries communication to regular phones, long-distance communication links, and the like.

Similarly phone networks 115 and 120 transport information and voice communications to and from devices capable of wired communications, such as regular phones and devices that include modems or some other interface to communicate with a phone network. A phone network, such as phone network 120, may also include both wireless and wired components. For example, a phone network may include microwave links, satellite links, radio links, and other wireless links to interconnect wired networks.

Gateways 130A-130D interconnect wireless networks 105 and 110 and telephone networks 115 and 120 to WAN/LAN 200. A gateway, such as gateway 130A, transmits data between networks, such as wireless network 105 and WAN/LAN 200. In transmitting data, the gateway may translate the data to a format appropriate for the receiving network. For example, a user using a wireless device may begin browsing the Internet by calling a certain number, tuning to a particular

frequency, or selecting a browsing feature of the device. Upon receipt of information appropriately addressed or formatted, wireless network 105 may be configured to send data between the wireless device and gateway 130A. Gateway 130A may translate requests for web pages from the wireless device to hypertext transfer protocol (HTTP) messages which may then be sent to WAN/LAN 200. Gateway 130A may then translate responses to such messages into a form compatible with the wireless device. Gateway 130A may also transform other messages sent from wireless devices into message suitable for WAN/LAN 200, such as email, voice communication, contact databases, calendars, appointments, and other messages.

Before or after translating the data in either direction, the gateway may pass the data through a firewall, such as firewall 140A, for security, filtering, or other reasons. A firewall, such as firewall 140A, may include or send messages to an exploit detector. Firewalls and their operation in the context of embodiments of the invention are described in more detail in conjunction with FIGURES 4-6. Briefly, a gateway may pass data through a firewall to determine whether it should forward the data to a receiving network. The firewall may pass some data, such as email messages, through an exploit detector which may detect and remove exploits from the data. If data contains an exploit, the firewall may stop the data from passing through the gateway.

In other embodiments of the invention, exploit detectors are located on components separate from gateways and/or firewalls. For example, in some embodiments of the invention, an exploit detector may be included within a router inside a wireless network, such as wireless network 105, that receives messages directed to and coming from the wireless network, such as wireless network 105. This may negate or make redundant an exploit detector on a gateway between networks, such as gateway 130A. Ideally, exploit detectors are placed at ingress locations to a network so that all devices within the network are protected from exploits. Exploit detectors may, however, be located at other locations within a network, integrated with other devices such as switches, hubs, servers, routers, traffic managers, etc., or separate from such devices.

In another embodiment of the invention, an exploit detector is accessible from a device that seeks to provide exploit protection, such as a gateway.

Accessible, in this context, may mean that exploit protector is physically located on the server or computing device implementing the gateway or that the exploit detector is on another server or computing device accessible from the gateway. In this embodiment, a gateway, may access the exploit detector through an application programming interface (API). Ideally, a device seeking exploit protection directs all messages through an associated exploit detector so that exploit detector is “logically” between the networks that the device interconnects. In some instances, a device may not send all messages through an exploit detector. For example, an exploit detector may be disabled or certain messages may be explicitly or implicitly designated to avoid the exploit detector.

Typically, WAN/LAN 200 transmits information between computing devices as described in more detail in conjunction with FIGURE 2. One example of a WAN is the Internet which connects millions of computers over a host of gateways, routers, switches, hubs, and the like. An example of a LAN is a network used to connect computers in a single office. A WAN may be used to connect multiple LANs.

It will be recognized that the distinctions between WANs/LANs, phone networks, and wireless networks are blurring. That is, each of these types of networks may include one or more portions that would logically belong to one or more other types of networks. For example, WAN/LAN 200 may include some analog or digital phone lines to transmit information between computing devices. Phone network 120 may include wireless components and packet-based components, such as voice over IP. Wireless network 105 may include wired components and/or packet-based components. Network means a WAN/LAN, phone network, wireless network, or any combination thereof.

FIGURE 2 shows a plurality of local area networks ("LANs") 220 and wide area network ("WAN") 230 interconnected by routers 210. Routers 210 are intermediary devices on a communications network that expedite packet delivery. On a single network linking many computers through a mesh of possible connections, a router receives transmitted packets and forwards them to their correct destinations over available routes. On an interconnected set of LANs--including those based on differing architectures and protocols--, a router acts as a link between LANs, enabling packets to be sent from one to another. A router may be

implemented using special purpose hardware, a computing device executing appropriate software, such as computing device 300 as described in conjunction with FIGURE 3, or through any combination of the above.

Communication links within LANs typically include twisted pair,
5 fiber optics, or coaxial cable, while communication links between networks may utilize analog telephone lines, full or fractional dedicated digital lines including T1, T2, T3, and T4, Integrated Services Digital Networks (ISDNs), Digital Subscriber Lines (DSLs), wireless links, or other communications links known to those skilled in the art. Furthermore, computers, such as remote computer 240, and other related
10 electronic devices can be remotely connected to either LANs 220 or WAN 230 via a modem and temporary telephone link. The number of WANs, LANs, and routers in FIGURE 2 may be increased or decreased arbitrarily without departing from the spirit or scope of this invention.

As such, it will be appreciated that the Internet itself may be formed
15 from a vast number of such interconnected networks, computers, and routers. Generally, the term "Internet" refers to the worldwide collection of networks, gateways, routers, and computers that use the Transmission Control Protocol/Internet Protocol ("TCP/IP") suite of protocols to communicate with one another. At the heart of the Internet is a backbone of high-speed data
20 communication lines between major nodes or host computers, including thousands of commercial, government, educational, and other computer systems, that route data and packets. An embodiment of the invention may be practiced over the Internet without departing from the spirit or scope of the invention.

The media used to transmit information in communication links as
25 described above illustrates one type of computer-readable media, namely communication media. Generally, computer-readable media includes any media that can be accessed by a computing device. Computer-readable media may include computer storage media, communication media, or any combination thereof.

Communication media typically embodies computer-readable
30 instructions, data structures, program modules, or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term "modulated data signal" means a signal that has one or more of its characteristics set or changed in such a manner as to encode

information in the signal. By way of example, communication media includes wired media such as twisted pair, coaxial cable, fiber optics, wave guides, and other wired media and wireless media such as acoustic, RF, infrared, and other wireless media.

The Internet has recently seen explosive growth by virtue of its
5 ability to link computers located throughout the world. As the Internet has grown, so has the WWW. Generally, the WWW is the total set of interlinked hypertext documents residing on HTTP (hypertext transport protocol) servers around the world. Documents on the WWW, called pages or Web pages, are typically written in HTML (Hypertext Markup Language) or some other markup language, identified
10 by URLs (Uniform Resource Locators) that specify the particular machine and pathname by which a file can be accessed, and transmitted from server to end user using HTTP. Codes, called tags, embedded in an HTML document associate particular words and images in the document with URLs so that a user can access another file, which may literally be halfway around the world, at the press of a key
15 or the click of a mouse. These files may contain text (in a variety of fonts and styles), graphics images, movie files, media clips, and sounds as well as Java applets, ActiveX controls, or other embedded software programs that execute when the user activates them. A user visiting a Web page also may be able to download files from an FTP site and send packets to other users via email by using links on the
20 Web page.

A computing device that may provide a WWW site is described in more detail in conjunction with FIGURE 3. When used to provide a WWW site, such a computing device is typically referred to as a WWW server. A WWW server is a computing device connected to the Internet having storage facilities for storing
25 hypertext documents for a WWW site and running administrative software for handling requests for the stored hypertext documents. A hypertext document normally includes a number of hyperlinks, i.e., highlighted portions of text which link the document to another hypertext document possibly stored at a WWW site elsewhere on the Internet. Each hyperlink is associated with a URL that provides
30 the location of the linked document on a server connected to the Internet and describes the document. Thus, whenever a hypertext document is retrieved from any WWW server, the document is considered to be retrieved from the WWW. As is known to those skilled in the art, a WWW server may also include facilities for

storing and transmitting application programs, such as application programs written in the JAVA programming language from Sun Microsystems, for execution on a remote computer. Likewise, a WWW server may also include facilities for executing scripts and other application programs on the WWW server itself.

5 A user may retrieve hypertext documents from the WWW via a WWW browser application program located on a wired or wireless device. A WWW browser, such as Netscape's NAVIGATOR[®] or Microsoft's INTERNET EXPLORER[®], is a software application program for providing a graphical user interface to the WWW. Upon request from the user via the WWW browser, the
10 WWW browser accesses and retrieves the desired hypertext document from the appropriate WWW server using the URL for the document and HTTP. HTTP is a higher-level protocol than TCP/IP and is designed specifically for the requirements of the WWW. HTTP is used to carry requests from a browser to a Web server and to transport pages from Web servers back to the requesting browser or client. The
15 WWW browser may also retrieve application programs from the WWW server, such as JAVA applets, for execution on a client computer.

FIGURE 3 shows a computing device. Such a device may be used, for example, as a server, workstation, network appliance, router, bridge, firewall, exploit detector, gateway, and/or as a traffic management device. When used to
20 provide a WWW site, computing device 300 transmits WWW pages to the WWW browser application program executing on requesting devices to carry out this process. For instance, computing device 300 may transmit pages and forms for receiving information about a user, such as address, telephone number, billing information, credit card number, etc. Moreover, computing device 300 may transmit
25 WWW pages to a requesting device that allows a consumer to participate in a WWW site. The transactions may take place over the Internet, WAN/LAN 100, or some other communications network known to those skilled in the art.

It will be appreciated that computing device 300 may include many more components than those shown in FIGURE 3. However, the components
30 shown are sufficient to disclose an illustrative environment for practicing the present invention. As shown in FIGURE 3, computing device 300 may be connected to WAN/LAN 200, or other communications network, via network interface unit 310. Network interface unit 310 includes the necessary circuitry for connecting

computing device 300 to WAN/LAN 200, and is constructed for use with various communication protocols including the TCP/IP protocol. Typically, network interface unit 310 is a card contained within computing device 300.

5 Computing device 300 also includes processing unit 312, video display adapter 314, and a mass memory, all connected via bus 322. The mass memory generally includes random access memory ("RAM") 316, read-only memory ("ROM") 332, and one or more permanent mass storage devices, such as hard disk drive 328, a tape drive (not shown), optical drive 326, such as a CD-ROM/DVD-ROM drive, and/or a floppy disk drive (not shown). The mass
10 memory stores operating system 320 for controlling the operation of computing device 300. It will be appreciated that this component may comprise a general purpose operating system including, for example, UNIX, LINUX™, or one produced by Microsoft Corporation of Redmond, Washington. Basic input/output system ("BIOS") 318 is also provided for controlling the low-level operation of
15 computing device 300.

The mass memory as described above illustrates another type of computer-readable media, namely computer storage media. Computer storage media may include volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information, such as
20 computer readable instructions, data structures, program modules or other data. Examples of computer storage media include RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM, digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the
25 desired information and which can be accessed by a computing device.

The mass memory may also store program code and data for providing a WWW site. More specifically, the mass memory may store applications including special purpose software 330, and other programs 334. Special purpose software 330 may include a WWW server application program that includes
30 computer executable instructions which, when executed by computing device 300, generate WWW browser displays, including performing the logic described above. Computing device 300 may include a JAVA virtual machine, an SMTP handler application for transmitting and receiving email, an HTTP handler application for

receiving and handing HTTP requests, JAVA applets for transmission to a WWW browser executing on a client computer, and an HTTPS handler application for handling secure connections. The HTTPS handler application may be used for communication with an external security application to send and receive sensitive information, such as credit card information, in a secure fashion.

Computing device 300 may also comprise input/output interface 324 for communicating with external devices, such as a mouse, keyboard, scanner, or other input devices not shown in FIGURE 3. In some embodiments of the invention, computing device does not include user input/output components. For example, computing device 300 may or may not be connected to a monitor. In addition, computing device 300 may or may not have video display adapter 314 or input/output interface 324. For example, computing device 300 may implement a network appliance, such as a router, gateway, traffic management device, etc., that is connected to a network and that does not need to be directly connected to user input/output devices. Such a device may be accessible, for example, over a network.

Computing device 300 may further comprise additional mass storage facilities such as optical drive 326 and hard disk drive 328. Hard disk drive 328 is utilized by computing device 300 to store, among other things, application programs, databases, and program data used by a WWW server application executing on computing device 300. A WWW server application may be stored as special purpose software 330 and/or other programs 334. In addition, customer databases, product databases, image databases, and relational databases may also be stored in mass memory or in RAM 316.

As will be recognized from the discussion below, aspects of the invention may be embodied on routers 210, on computing device 300, on a gateway, on a firewall, on other devices, or on some combination of the above. For example, programming steps protecting against exploits may be contained in special purpose software 330 and/or other programs 334.

30 Exemplary Configuration of System to Protect from Exploits

FIGURE 4 illustrates an exemplary environment in which a system for providing exploit protection for a network operates, according to one embodiment of the invention. The system includes outside network 405, firewall

500, network appliance 415, workstation 420, file server 425, mail server 430, mobile device 435 application server 440, telephony device 445, and network 450. Network 450 couples firewall 500 to network appliance 415, workstation 420, file server 425, mail server 430, mobile device 435, application server 440, and telephony device 445. Firewall 500 couples network 450 to outside network 405.

Network appliance 415, workstation 420, file server 425, mail server 430, mobile device 435, application server 440, and telephony device 445 are devices capable of connecting with network 450. The set of such devices may include devices that typically connect using a wired communications medium such as personal computers, multiprocessor systems, microprocessor-based or programmable consumer electronics, network PCs, and the like. The set of such devices may also include devices that typically connect using a wireless communications medium such as cell phones, smart phones, pagers, walkie talkies, radio frequency (RF) devices, infrared (IR) devices, CBs, integrated devices combining one or more of the preceding devices, and the like. Some devices may be capable of connecting to network 450 using a wired or wireless communication medium such as a PDA, POCKET PC, wearable computer, or other device mentioned above that is equipped to use a wired and/or wireless communications medium. An exemplary device that may implement any of the devices above is computing device 300 of FIGURE 3 configured with the appropriate hardware and/or software.

Network appliance 415 may be, for example, a router, switch, or some other network device. Workstation 420 may be a computer used by a user to access other computers and resource reachable through network 450, including outside network 405. File server 425 may, for example, provide access to mass storage devices. Mail server 430 may store and provide access to email messages. Mobile device 435 may be a cell phone, PDA, portable computer, or some other device used by a user to access resources reachable through network 450. Application server 440 may store and provide access to applications, such as database applications, accounting applications, etc. Telephony device 445 may provide means for transmitting voice, fax, and other messages over network 450. Each of these devices may represent many other devices capable of connecting with network 450 without departing from the spirit or scope of the invention.

Outside network 405 and Network 450 are networks as previously defined in this document. Outside network may be, for example, the Internet or some other WAN/LAN.

5 Firewall 500 provides a pathway for messages from outside network 405 to reach network 450. Firewall 500 may or may not provide the only pathway for such messages. Furthermore, there may be other computing devices (not shown) in the pathway between outside network 405 and network 450 without departing from the spirit or scope of the invention. Firewall may be included on a gateway, router, switch, or other computing device or simply accessible to such devices.

10 Firewall 500 may provides exploit protection for devices coupled to network 450 by including and/or accessing an exploit detector (not shown) as described in more detail in conjunction with FIGURE 5. Firewall 500 may be configured to send certain types of messages through an exploit detector. For example, firewall 500 may be configured to perform normal processing on non-
15 email data while passing all email messages through an exploit detector.

Exemplary Exploit Detector

FIGURE 5 illustrates components of a firewall operable to provide exploit protection, according to one embodiment of the invention. The components
20 of the firewall 500 include message listener 505, exploit detector 510, output component 545, and other firewall components 550. Exploit detector 510 includes message queue 515, content filter 520, decompression component 525, scanner component 530, quarantine component 535, and exploit remover 540. Also shown is message transport agent 555.

25 Firewall 500 may receive many types of messages sent between devices coupled to network 450 and outside network 405 of FIGURE 4. Some messages may relate to WWW traffic or data transferred between two computers engaged in a communication while other messages may relate to email. Message listener 505 listens for a message and, upon receipt of an appropriate message, such
30 as an email or file, sends the message to exploit detector 510 to scan for exploits. Some messages may be inappropriate for exploit detection. Such messages are passed by message listener 505 to other firewall components 550.

When processing email messages, exploit detector 510 provides exploit protection, in part, by scanning and verifying the fields of an email message. An email message typically includes a header (which may include certain fields), a body (which typically contains the text of an email), and one or more optional
5 attachments. As described earlier, some exploits are crafted to overflow buffers in a header or body. Exploit detector 510 may examine the lengths of the fields of an email message to determine whether they are longer than they should be. Being “longer than they should be” may be defined by standards, mail server specifications, or selected by a firewall administrator. If an email message includes
10 any fields that are longer than they should be, the message may be sent to quarantine component 535 as described in more detail below.

Exploit detector 510 may utilize exploit protection software from many vendors. For example, a client may execute on exploit detector 510 that connects to a virus protection update server. Periodically, the client may poll a
15 server associated with each vendor and look for a flag to see if an exploit protection update is available. If there is an update available, the client may automatically retrieve the update and check it for authenticity. For example, the update may include a digital signature that incorporates a hash of the files sent. The digital signature may be verified to make sure that the files came from a trusted sender, and
20 the hash may be used to make sure than none of the files have been modified in transit. Another process may unpack the update, stop the execution of exploit detector 510, install the update, and restart exploit detector 510.

Exploit detector 510 may be configured to poll for customized exploit protection updates created by, for example, an information technology team. This
25 process may execute in a manner similar to the polling for vendor updates described above.

In addition to, or in lieu of polling, updates may be pushed to exploit detector 510. That is, a client may execute on exploit detector 510 that listens for updates from exploit protection update servers. To update the exploit protection
30 executing on firewall 410, such servers may open a connection with the client and send exploit protection updates. A server sending an update may be required to authenticate itself. Furthermore, the client may check the update sent to make sure that files have not changed in transit by using a hash as described above.

The components of exploit detector 510 will now be explained.

Upon receipt of a message to scan for exploits, exploit detector 510 stores the message in message queue 515. Content filter 520 processes messages from message queue 515 to determine encapsulation methods that have been applied to
5 the message prior to its entry into the system. For example, a message may be encapsulated using Multipurpose Internet Mail Extensions (MIME), Base 64, and uuencode. Content filter 520 may also strip out attachments from email in order to examine them more closely. A message or attachment (hereinafter each referred to as a “message”) that is output from content filter 520 is then processed by
10 decompression component 525.

Decompression component 525 determines whether a message is compressed. If the message is not compressed, the bits that make up the message are sent serially to scanner component 530. If the message is compressed, decompression component 525 may decompress the message one or more times
15 before sending it to scanner component 530. Decompressions may be done in a nested fashion if a message has been compressed multiple times. For example, a set of files included in a message may first be zipped and then tarred using the UNIX “tar” command. After untarring a file, decompression component 525 may determine that the untarred file was previously compressed by zipping software such
20 as WinZip. To obtain the unzipped file(s), decompression component 525 may then unzip the untarred file. There may be more than two levels of compression that decompression component 525 decompresses to obtain decompressed file(s).

Scanner component 530 receives decompressed messages and messages that were not compressed from decompression component 525. Scanner
25 component 530 includes software that scans the message for exploits. Scanner component 530 may scan messages using exploit protection software from many vendors. For example, scanner component 530 may pass a message through software from virus protection software vendors such as Norton, MacAfee, Network Associates, Inc., Kaspersky Lab, Sophos, etc. In addition, scanner component 530
30 may apply proprietary or user-defined algorithms to the message to scan for exploits. For example, a user-defined algorithm testing for buffer overflows may be used to detect exploits.

Scanner component 530 may also include an internal mechanism that creates digital signatures for messages and content that an administrator wants to prevent from being distributed outside a network. For example, referring to FIGURE 4, a user on one of the computing devices may create a message or try to forward a message that is confidential to outside network 405. Scanner component 530 may examine each message it receives (including outbound messages) for such digital signatures. When a digital signature is found that indicates that the message should not be forwarded, scanner component 530 may forward the message to quarantine component together with information as to who sent the message, the time the message was sent, and other data related to the message.

When a message is determined to have an exploit, the message is sent to quarantine component 535. Quarantine component 535 may store messages that contain exploits for further examination by, for example, a network administrator. In addition, quarantine component 535 may send an infected message to exploit remover 540 to remove an exploit.

When scanner component 530 does not find an exploit in a message, the message may be forwarded to output component 545. Output component 545 forwards a message towards its recipient. Output component 545 may be hardware and/or software operative to forward messages over a network. For example, output component 545 may include a network interface such as network interface unit 310.

Exploit remover 540 may remove exploits from a message. Some exploits may be removed from a message after detection yielding a cleaned message. The cleaned message, now free from exploits, may then be forwarded to its intended recipient. After cleaning a message, exploit remover may forward the message to output component 545. If exploit remover cannot remove an exploit, it may send the message back to quarantine component 535.

A firewall may perform other tasks besides passing messages to an exploit detector. For example, a firewall may block messages to or from certain addresses. Such other tasks may be accomplished by other firewall components 550. When other firewall components 550 determines that a message should be passed through firewall 500, other firewall components 550 forwards the message to output component 545.

Message transport agent 555 is a computing device that receives email. Email receiving devices include mail servers. Examples of mail servers include Microsoft Exchange, Q Mail, Lotus Notes, etc. Referring to FIGURE 4, firewall 500 may forward a message to mail server 430.

5

Illustrative Method of Scanning for Exploits

FIGURE 6 illustrates a flow chart for detecting exploits, according to one embodiment of the invention. The process begins at block 605 when a listener, such as message listener 505 of FIGURE 5, is ready to receive a message.

10

At block 610, the message is received by a listener. The listener determines whether the message should be scanned for exploits. If the message is to be scanned for exploits, processing continues at block 615; otherwise other processing (not shown) may be performed on the message. For example, referring to FIGURE 5, a message including an email message is received by message listener 505. Message listener 505 determines that the message should be scanned for exploits and sends the message to message queue 515.

15

At block 615 the message is unencapsulated, if necessary. A message may be encapsulated in many ways, including MIME, Base 64, and uuencode. To retrieve the message, the message may be unencapsulated. For example, referring to FIGURE 5, the email message may include an attachment that is encoded using MIME. Content filter 520 may unencapsulate the attachment. After block 615, processing continues at block 620.

20

At block 620, the message and/or its attachment, if any, may be decompressed one or more times. For example, referring to FIGURE 5, an email message may include an attachment that has been compressed by WinZip. Decompression component 525 may determine the compression algorithm used and then decompress the attachment. After block 620, processing continues at block 625.

25

At block 625, a message is scanned for exploits. The message may be scanned using conventional exploit detection software and/or proprietary or user-defined exploit detection software. For example, referring to FIGURE 5, the header, body, and attachment fields of an email message may be scanned to determine if they are less than or equal to the maximum length of such fields. In addition, the

30

attachments of an email, if any, may be passed through virus detection software from various vendors to determine if they include any exploits. After block 625, processing continues at block 630.

At block 630, a determination is made as to whether the scan detected
5 any exploits. If exploits are found, processing continues at block 635; otherwise processing continues at block 640.

At block 635, a message is quarantined and optionally one or more exploits are removed. Quarantined may mean that the message is stored together with other information regarding the message, such as who sent the message, to
10 whom it was addressed, and when the message arrived. This may be done for further examination or analysis. Alternatively, quarantined may mean that the message is discarded. When exploits are removed from a message processing, may continue at block 640; otherwise, processing finishes for a particular message and another message may be scanned for exploits. For example, referring to FIGURE 5,
15 quarantine component receives an email including exploits and stores the email for further examination.

At block 640, a message is forwarded towards its recipient. The message may be an original message received by an exploit detector or may be a message from which exploits have been removed. For example, referring to
20 FIGURE 5, output component 545 forwards a message to message transport agent 555.

At block 645, processing ends. At this point a message has been scanned for exploits. If any part of the message has been encapsulated, the message has been unencapsulated. If the message was compressed one or more time, the
25 message has been decompressed one or more time. A scan for exploits has occurred on the message. If exploits were found they have been quarantined and/or optionally been removed from the message. The message or cleaned message has then been forwarded towards the recipient. The process outlined above may be repeated for each message received.

30 The various embodiments of the invention may be implemented as a sequence of computer implemented steps or program modules running on a computing system and/or as interconnected machine logic circuits or circuit modules within the computing system. The implementation is a matter of choice dependent

on the performance requirements of the computing system implementing the invention. In light of this disclosure, it will be recognized by one skilled in the art that the functions and operation of the various embodiments disclosed may be implemented in software, in firmware, in special purpose digital logic, or any
5 combination thereof without deviating from the spirit or scope of the present invention.

The above specification, examples and data provide a complete description of the manufacture and use of the composition of the invention. Since many embodiments of the invention can be made without departing from the spirit
10 and scope of the invention, the invention resides in the claims hereinafter appended.

WHAT IS CLAIMED IS:

1. A system for providing protection from exploits to devices connected to a network, comprising:

(a) a content filter that receives a message that is directed to at least one of the devices and that includes a header, a body, and an attachment, wherein the content filter determines an encapsulation that has been applied to the attachment prior to the system receiving the message and unencapsulates the attachment;

(b) a decompression component that is coupled to the content filter and that performs at least one decompression of the attachment when the attachment is compressed;

(c) a scanner component that is coupled to the decompression component and that determines whether at least one of the header and the body includes an exploit;

(d) a quarantine component that is coupled to the scanner component and that holds the message when the message includes an exploit; and

(e) a device that receives messages that are directed to the network and that employs at least the scanner component to provide exploit protection for at least one of the messages.

2. The system of claim 1, wherein the header includes a field having a defined size and wherein the scanner determines that the header includes the exploit when a size of data in the field is other than the defined size.

3. The system of claim 1, wherein the scanner component further determines whether the attachment includes an exploit.

4. The system of claim 3, further comprising a client that automatically applies an update to at least one of the content filter, the decompression component, the scanner component, and the quarantine component to enable detection of at least one exploit.

5. The system of claim 4, wherein the client determines when the update is available by polling servers associated with vendors of exploit protection software.

6. The system of claim 5, wherein the client automatically retrieves the available update.

7. The system of claim 4, wherein the scanner component employs at least two separate exploit protection applications to determine whether the attachment includes an exploit.

8. The system of claim 1, wherein the content filter, the decompression component, the scanner component, and the quarantine component are each implemented in software.

9. The system of claim 1, wherein the content filter, the decompression component, the scanner component, and the quarantine component are all included on at least one of a firewall, router, switch, and traffic manager.

10. The system of claim 1, wherein the encapsulation includes at least one of Multipurpose Internet Mail Extensions (MIME), Base 64, and uuencode.

11. The system of claim 1, wherein the quarantine component removes the exploit from the message and forwards the message towards a recipient of the message.

12. A method for providing protection from exploits to devices connected to a network, comprising:

(a) receiving a message at a node that receives messages that are directed to any of the devices and that causes the message to be scanned for an exploit before forwarding the message toward at least one of the devices, wherein the message includes a header and at least one of a body and an attachment;

(b) determining whether at least one of the header and the body includes the exploit; and

(c) when at least one of the header and the body of the message includes the exploit, quarantining the message.

13. The method of claim 12, further comprising:

(a) unencapsulating the attachment when the attachment is encapsulated; and

(b) decompressing the attachment at least one time when the attachment is compressed.

14. The method of claim 13, further comprising removing the exploit and forwarding the message towards the at least one of the devices.

15. The method of claim 13, wherein the header includes a field having a defined size and wherein the header includes the exploit when a size of data in the field is other than the defined size.

16. The method of claim 13, further comprising determining whether the attachment includes an exploit.

17. The method of claim 16, wherein exploit protection software from at least two vendors is employed to determine whether the attachment includes an exploit.

18. The method of claim 13, wherein the attachment is encapsulated using at least one of Multipurpose Internet Mail Extensions (MIME), Base 64, and uuencode.

19. A system for providing protection from exploits to devices connected to a network, comprising:

(a) means for receiving a message that includes a header and at least one of a body and an attachment;

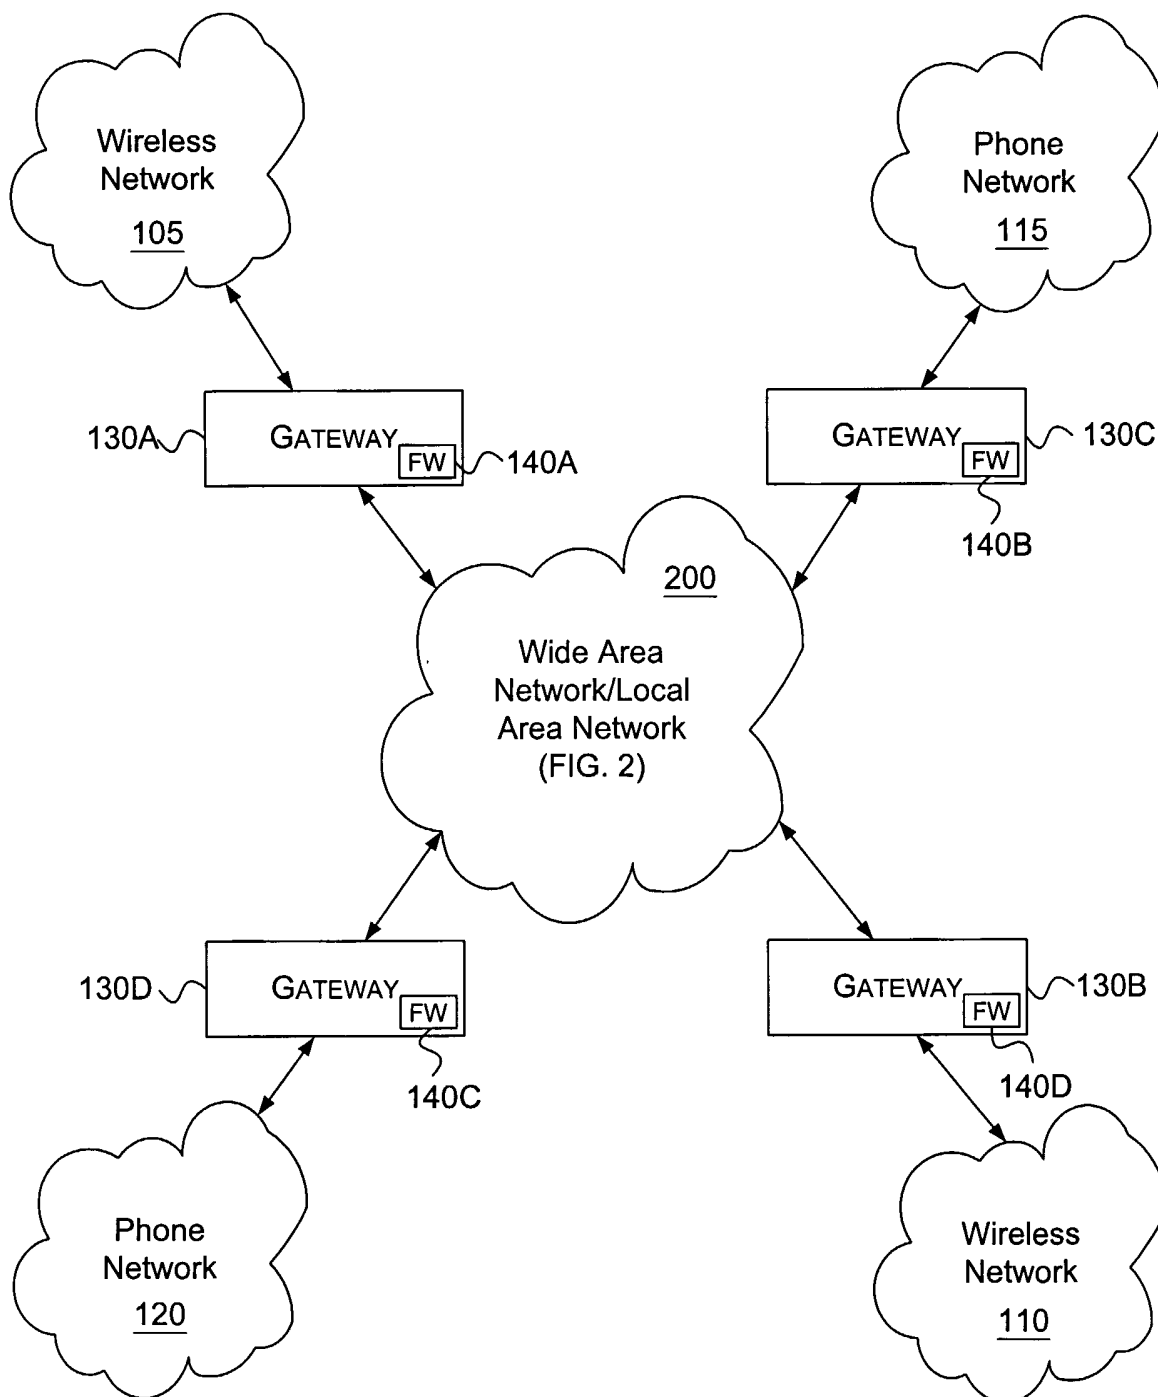
(b) means for determining whether the attachment is encapsulated and for unencapsulating the attachment when the attachment is encapsulated;

(c) means for decompressing the attachment at least one time when the attachment is compressed;

(d) means for determining whether at least one of the header and the body includes an exploit; and

(e) means for quarantining the message when the message includes the exploit.

FIG. 1






FIG. 2




FIG. 3

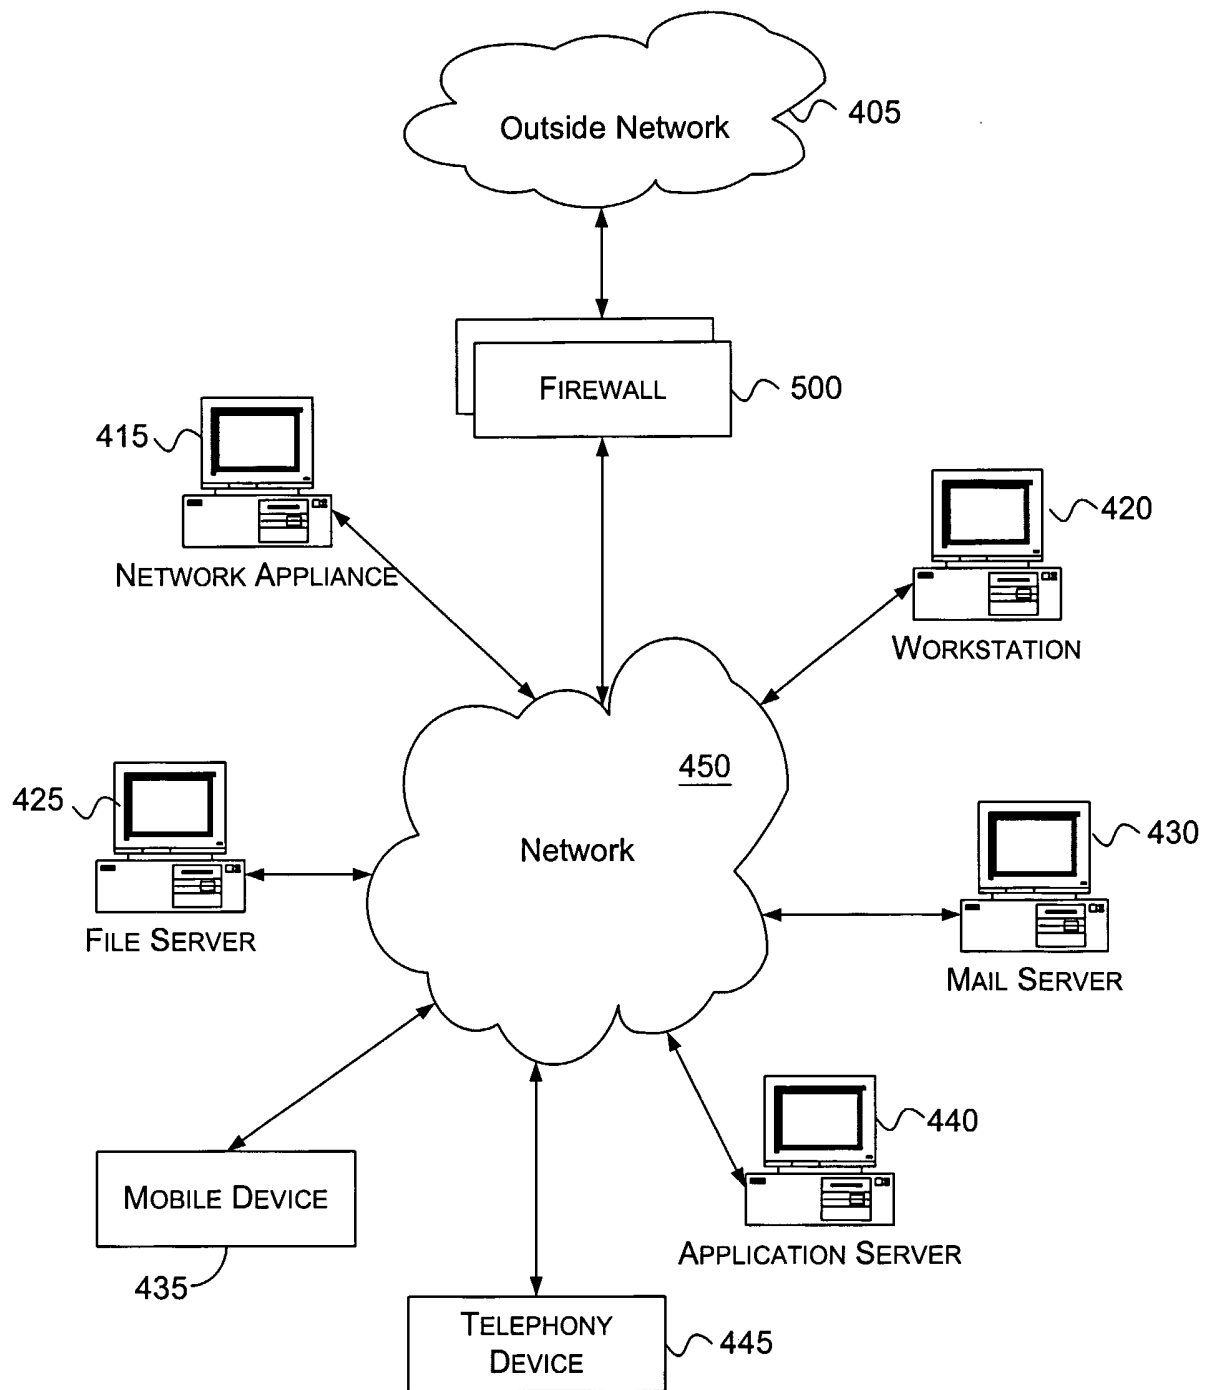


FIG. 4

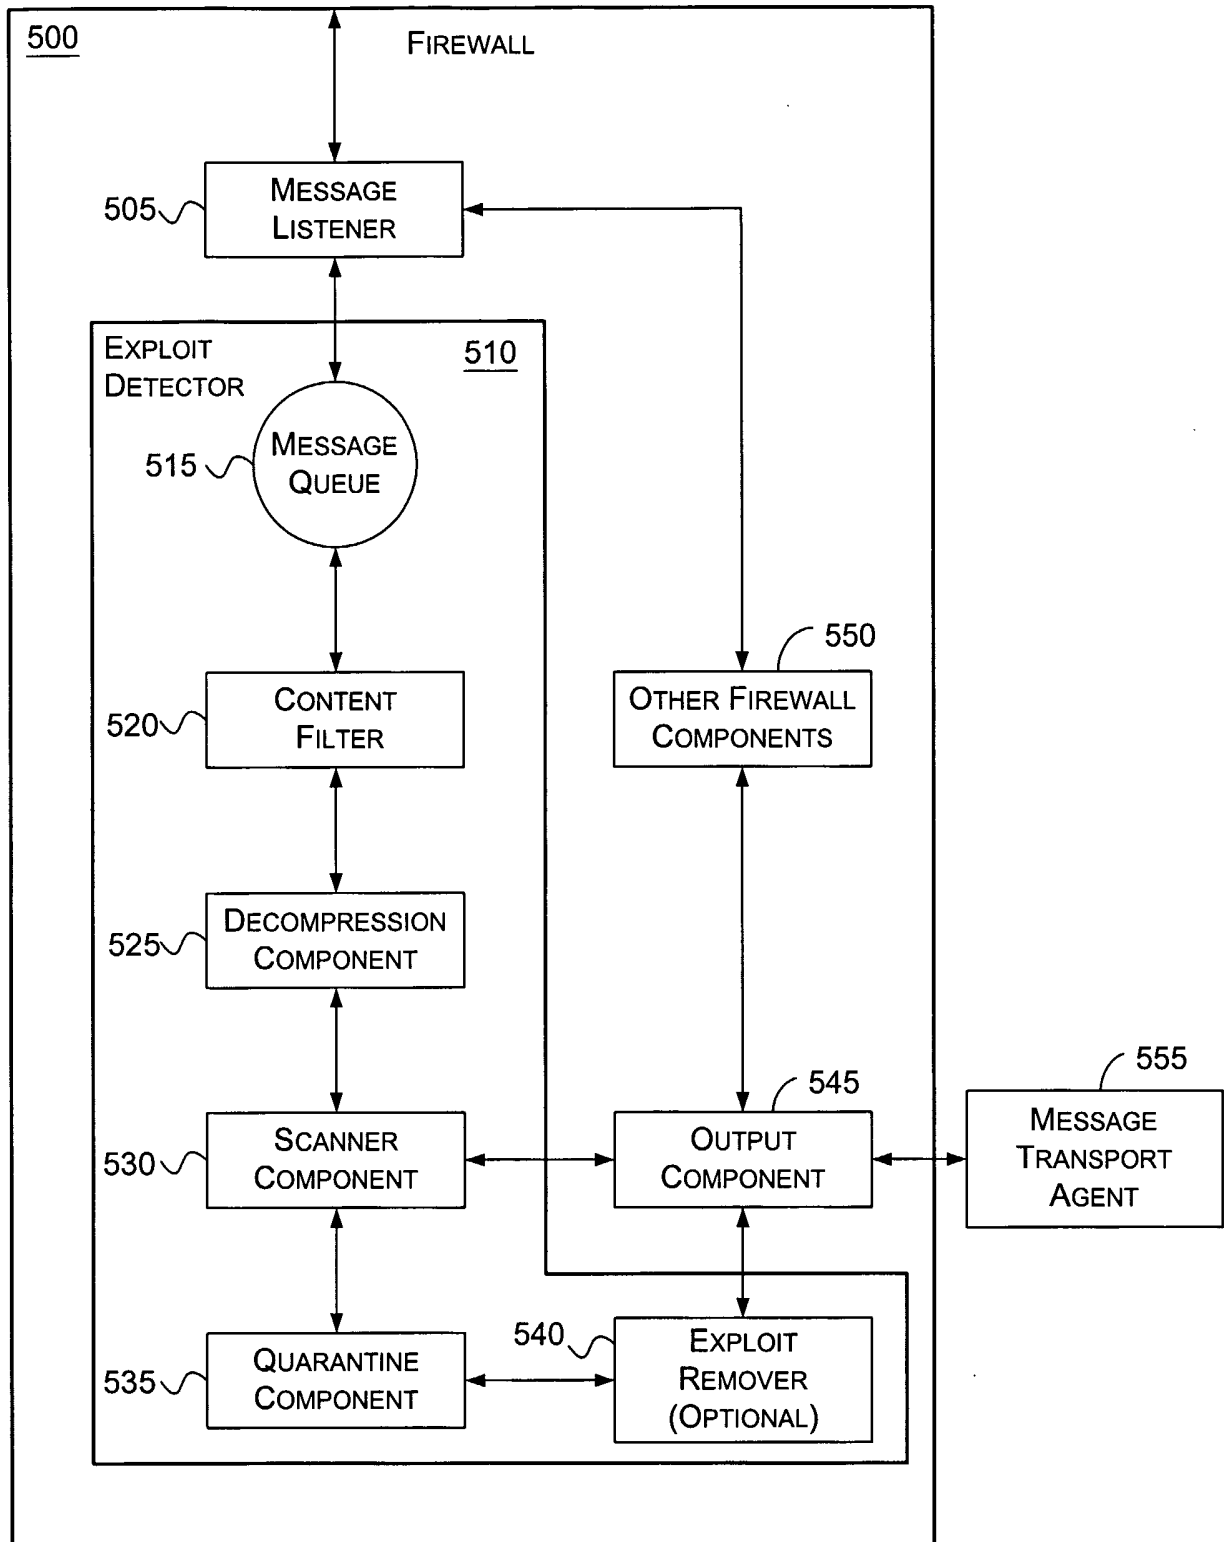
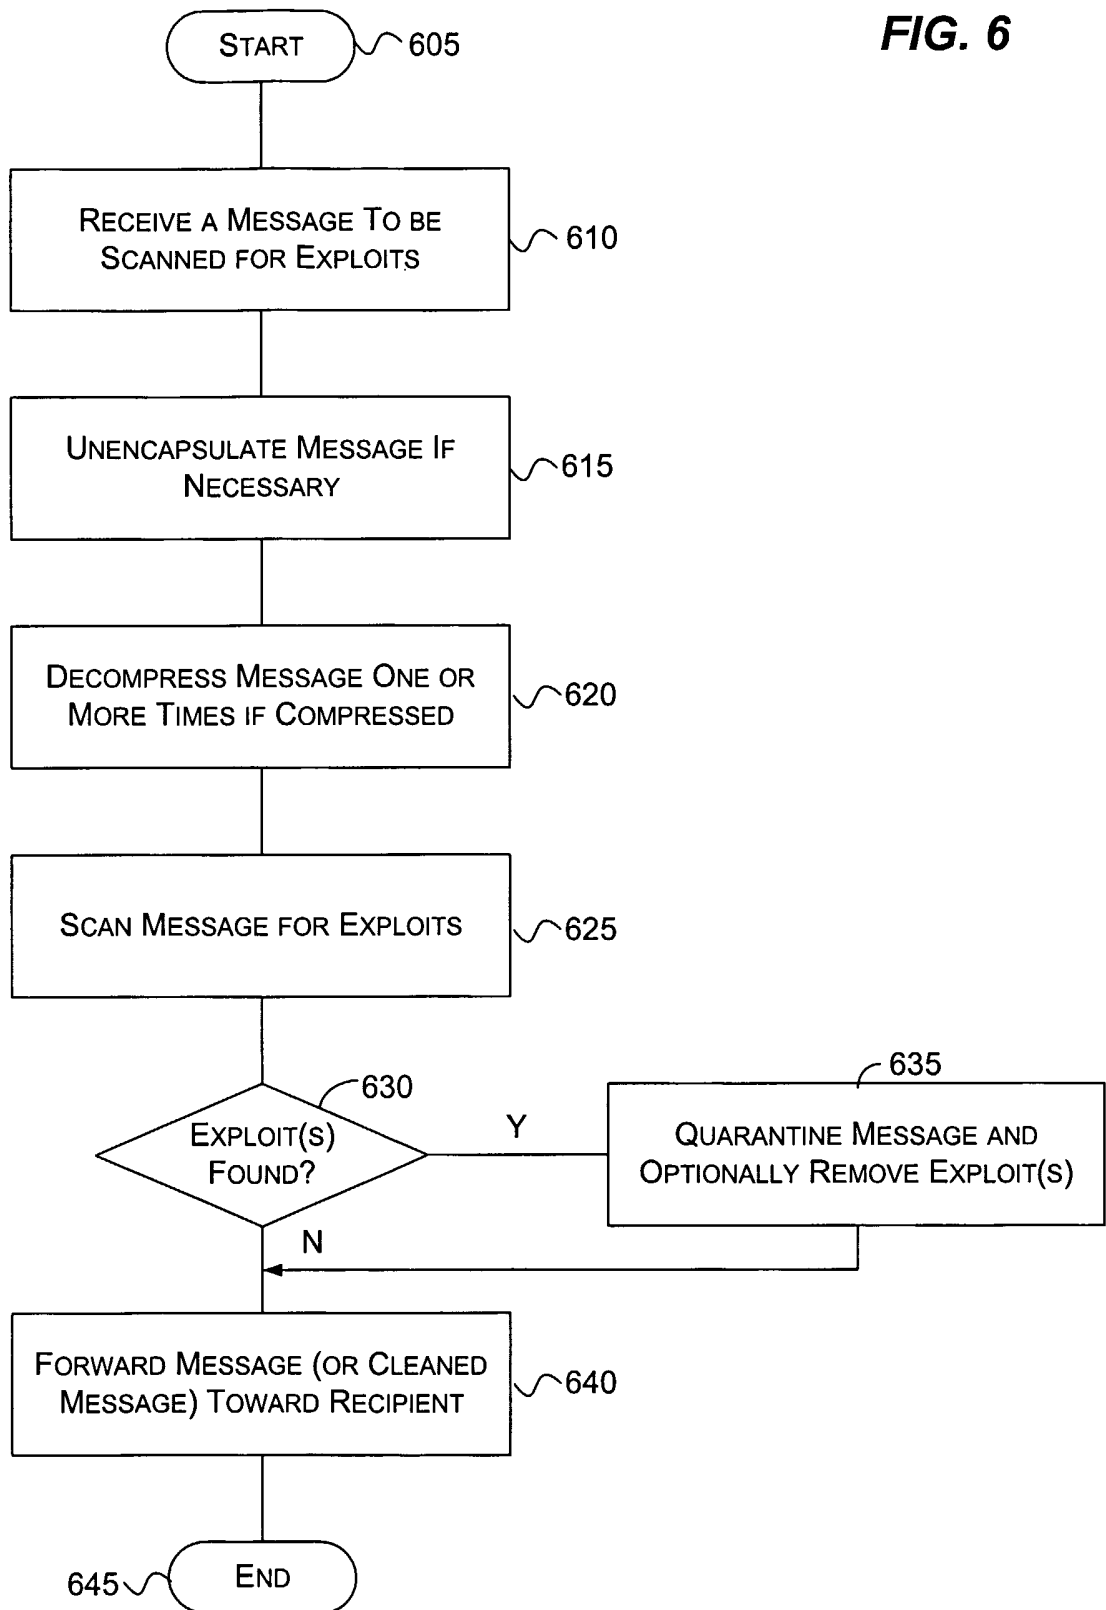


FIG. 5

FIG. 6



INTERNATIONAL SEARCH REPORT

International application No.

PCT/US02/11685

A. CLASSIFICATION OF SUBJECT MATTER				
IPC(7) : G06F 12/14; H04L 9/00; G06F 13/00				
US CL : 713/189; 713/201; 713/200; 712/225				
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) U.S. : 713/189; 713/201; 713/200; 712/225				
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched None				
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) East				
C. DOCUMENTS CONSIDERED TO BE RELEVANT				
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.		
Y	US 5,832,208 A (CHEN et al) 03 November 1998 (03.11.1998), column 4, lines 12-55.	1-19		
---		-----		
Y	US 6,205,551 B1 (GROSSE) 20 March 2001 (20.03.2001), column 3, lines 7-34.	1-19		
Y	US 5,815,571 A (FINLEY) 29 September 1998 (29.09.1998), column 3, lines 1-67.	1-19		
---		-----		
A	US 5,987,610 A (FRANCZEK et al) 16 November 1999 (16.11.1999), column 3, lines 1-58.	1-19		
A	US 5,832,228 A (HOLDEN et al) 03 November 1998 (03.11.1998), column 2, lines 21-48.	1-19		
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input type="checkbox"/> See patent family annex.				
<table style="width:100%; border:none;"> <tr> <td style="width:50%; border:none;"> * 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 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 </td> <td style="width:50%; border:none;"> "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 </td> </tr> </table>			* 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 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
* 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 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 03 July 2002 (03.07.2002)		Date of mailing of the international search report 02 AUG 2002		
Name and mailing address of the ISA/US Commissioner of Patents and Trademarks Box PCT Washington, D.C. 20231 Facsimile No. (703)305-3230		Authorized officer Zarni Maung <i>Peggy Hancock</i> Telephone No. (703)305-9669		